Proceedings of the 2001 Winter Simulation Conference
B. A. Peters, J. S. Smith, D. J. Medeiros, and M. W. Rohrer, eds.

STOCHASTIC MODELING OF AIRLIFT OPERATIONS

Julien Granger
Ananth Krishnamurthy
Stephen M. Robinson

Department of Industrial Engineering
University of Wisconsin-Madison
1513 University Avenue
Madison, WI 53706-1572, U.S.A.

ABSTRACT

Large-scale military deployments require transporting equipment and personnel over long distances in a short time. Planning an efficient airlift system is complicated and several models exist in the literature. Particularly, a study conducted on a deterministic optimization model developed by the Naval Postgraduate School and the RAND Corporation has shown that incorporating stochastic events leads to a degradation of performance. In this paper we investigate the applicability of network approximation methods to take into account randomness in an airlift network. Specifically, we show that approximation methods can model key performance features with sufficient accuracy to permit their use for network improvement, while requiring only a small fraction of the computational work that would have been needed had simulation been used for all of the performance evaluations. Also, we predict that combining simulation and approximation may work substantially better than either one of these alone.

1 INTRODUCTION

In supporting military deployments and other operations, airlift practitioners must move equipment and personnel from multiple origins through a transportation network to destinations. Aircraft available for the purpose may be of diverse types. Generally the time allowed for movement is limited, and so are the resources (aircraft and airfield space) available. The task of planning an efficient airlift system is complicated, and modelers have developed both simulation models and optimization models to assist with such planning.

The Airlift Flow Model (AFM) is a stochastic simulation model of global airlift operations. As reported in (Rousseau, 1998) it is used by the Air Force's Air Mobility Command to assess movement of both cargo and passengers, examine the feasibility of war plans, and assess the impact of new resources or of modifications to policies or infrastructure. An example of a deterministic optimization model is the NPS/RAND Mobility Optimizer (NRMO) developed by the Naval Postgraduate School and the RAND Corporation (Baker et al., 1999).

It is well known that good operating policies may be very different for stochastic systems than for deterministic systems. For example, in a deterministic system an intermediate processor in the system, such as an airfield at which planes must land for fueling or maintenance enroute to the final destination, can support maximum utilization with no loss of efficiency. However, if elements of the system such as the service time and/or the interarrival times are stochastic, then loading the processor to maximum utilization is disastrous. In such situations one must plan for substantially less than full utilization in order to avoid extremely long average times in queue. Accordingly, the whole approach to planning the operation of the system has to be changed.

A predecessor of the present study (Niemi, 2000) investigated the effect on optimization of incorporating stochastic events into an airlift network, using the NRMO model as a testbed. It found a degradation of performance due to stochasticity. It also found that modeling the stochastic features of the airlift system could significantly improve effectiveness of new resources allocated to it. These findings underline the importance of taking into account the stochastic elements of airlift operations when trying to optimize such a system.

Current methods for modeling and optimizing stochastic systems frequently rely on simulation. However, for even a moderately complex network the amount of computational work required to simulate performance, let alone to optimize it with respect to parameters, can be very large. For effective optimization, something must be done to decrease the computational load involved in repetitively simulating network performance. However, simulation is an essential component of the process because it is the only known tool that enables us to assess system performance accurately.

Approximation methods developed in queuing theory have been extensively applied in manufacturing systems (Bitran and Morabito, 1996), and their use has led to significant improvements (Brown, 1988; Suri, 1998). One of the main advantages of these approximation methods is their rapidity of execution compared to the cost of simulation models. Therefore we decided to make a preliminary experimental investigation to see whether we could apply such approximation methods to airlift operations in order to reduce the computational burden of modeling and optimization while preserving enough accuracy so that the end results would be useful. We used small and simple networks and off-the-shelf commercial software, with the aim of determining whether a more extensive (and more costly) developmental effort would be warranted. This paper reports the results of that investigation.

Our findings, for the simple networks that we studied, are:

- Approximation methods model key performance features with sufficient accuracy to make their use in optimization packages feasible, while requiring only a small fraction of the computational work that would have been needed had simulation been used for all of the performance evaluations.
- Accurate and credible performance evaluation requires simulation, but simulation is costly. Therefore the combination of simulation and approximation works substantially better than either one of these alone. This combination can significantly improve network performance.

The rest of the paper provides details and data to support these conclusions. Our overall conclusions from this pilot study are that a program of developing methods for optimizing stochastic networks of this type by combining simulation with approximation has a good prospect of success, that if successful it would provide a capability that is unavailable with current tools, and that it might well have applications in areas other than transportation.

2 MODELING A SIMPLIFIED AIRLIFT NETWORK

2.1 Description of the Simplified Airlift Network

We consider a simplified version of the airlift network described in (Baker et al., 1999). It has five airfields, namely Aerial Port of Embarkation (APOE), Aerial Port of Debarkation (APOD) and three intermediate airfields: England, Spain and Germany.

Cargo is transmitted from APOE to APOD via these intermediate airfields using three types of fleet, namely C-5, C-17 and C-141. Each fleet consists of a specific number of aircraft. On intermediate airfields, aircraft can experience ground delays, for example due to repair. We assume that ramp space capacity at each airfield is finite. Figure 1 gives a representation of this network.

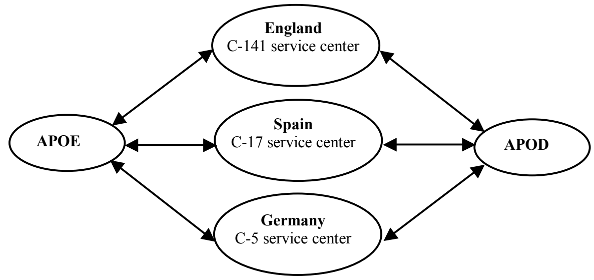

Figure 1: Simplified Airlift Network

Efficiency of the airlift operations is measured by outcomes such as completion time, average flight time for delivery and recovery aircraft, utilization of ramp capacity at airfields, and average number of aircraft in various parts of the network at any given time. For modeling convenience we make certain additional assumptions:

1. Each intermediate airfield is dedicated to a particular fleet type: all C-141s fly between APOE and APOD via England, C-17s fly via Spain and C-5s via Germany. (Baker et al., 1999) used this method to deal with the fact that servicing each type of aircraft requires specific skills and facilities.
2. No distinction is made between types of cargo (passengers (pax), outsized, oversized or bulk cargo), or between aircraft cargo capacities. Each aircraft carries one unit of cargo. The cargo delivery requirement from each fleet is known a priori.
3. Each type of aircraft (C-5, C-17, C-141) requires the same amount of ramp space at any airfield.
4. At APOE, aircraft never wait for cargo and cargo is loaded on the first available aircraft.
5. Ground times and inter-airfield flight times can be stochastic. Inter-airfield routes have infinite capacity.
6. The airlift operation is completed when all the cargo has been delivered at APOD. All aircraft return to APOE after the airlift has been completed.

Table 1 lists the initial choice of input parameters, namely cargo requirement, ramp capacities at airfields, number of aircraft in different fleets (fleet size), inter-airfield flight times and average ground times.

Table 1: Initial Data

Cargo Requirement per Fleet (units)		Fleet Size (units)	
C-141	500	C-141	60
C-5	500	C-5	60
C-17	500	C-17	60
Total Cargo	1,500	Total aircraft	180
Ramp Capacity (units)		**Average Run Times (hours)**	
APOE	200	APOE, APOD	10
APOD	30	Inter-airfield flying time	7
England, Germany & Spain	10	Average ground time at en-route airfield	5

2.2 Description of the Simulation Model

We built a simulation model of the simplified airlift network in ProModel® Version 4.2, a commercially available discrete event simulation package used for simulating manufacturing systems (see information on web page at `<http://www.promodel.com>`). The airlift network is therefore modeled as a manufacturing network, where airfields and aircraft are considered as locations and entities respectively. Figure 2 depicts the schematic of the simulation model. As can be seen in the figure, servers or locations sAPOE, sEng, sGer, sSpain and sAPOD represent the corresponding airfields of the airlift network.

Two locations represent each airfield: one location has finite capacity corresponding to ramp capacity available at airfield, and the other is an infinite-capacity queue corresponding to the competition between aircraft for ramp space. One could think of this second location as a ground holding area. For example, the sEng location models the airfield at England with finite ramp capacity, while qAPOEEng (qAPODEng) is the corresponding infinite capacity queue, where delivery (recovery) aircraft to (from) APOD wait for ramp capacity. After completing their delivery mission, aircraft return to a home base distinct from other airfields (sUSA in Figure 2). An empty aircraft, after arriving from the entry point of the network qArrivalAircraft, proceeds to APOE from sUSA only if there is cargo to be delivered. The distinction between sUSA and sAPOE serves modeling convenience, but does not necessarily have any physical reality. The amount of cargo present at APOE is modeled using a variable. Consequently, the airlift network is modeled as a closed network with a fixed number of entities equal to the fleet size of aircraft. Inter-airfield legs are modeled as locations with unlimited capacity where aircraft wait for times corresponding to flight durations. For example, sAPOEEng (sEngAPOE) represents the inter-airfield leg of a delivery (recovery) aircraft between England and APOE. Stochastic flying times and ground times are modeled using a random variable with gamma distribution having the appropriate mean. We chose a gamma distribution because it allows us to model a wide range of variability in flying and ground times.

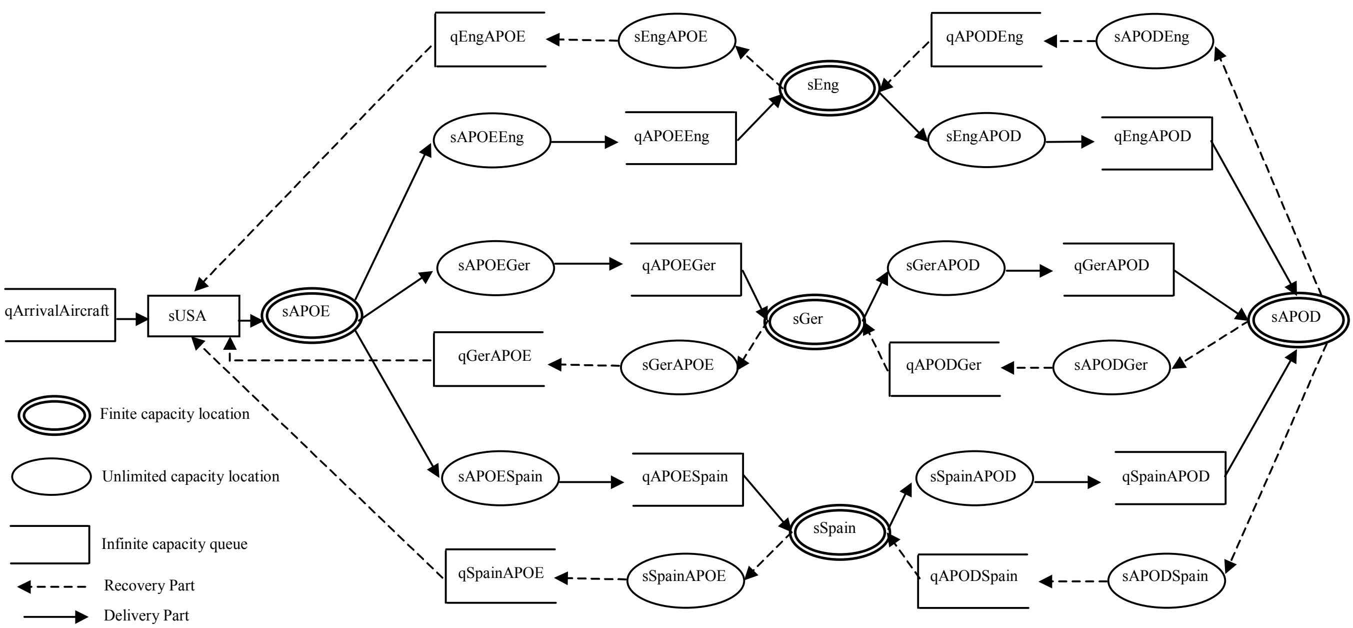

Figure 2: Schematic of the Airlift Network Simulated in ProModel®

2.3 Description of the Network Approximation Model

We built a network approximation model of the simplified airlift network using MPX® Version 3.3 (Network Dynamics, Inc., 1999), a commercially available package based on recent advances in approximation methods for queuing theory and reliability modeling. (Suri et al., 1993) and the references therein describe some of these advances. Case studies in (Suri, 1998) describe how network approximations have been successfully used in manufacturing.

Figure 3 depicts the schematic of the network approximation model. As in the simulation model, airfields and aircraft are modeled as equipment groups and part types. For example, equipment groups gAPOE, gEng, gGer, gSpain and gAPOD represent the corresponding airfields of the airlift network. The number of machines in each equipment group corresponds to ramp capacity. MPX models manufacturing networks as open networks of queues, while the strategic airlift network is a closed network with a fixed number of entities equal to the total number of aircraft. To model recovery aircraft returning to APOE, we define two part types corresponding to delivery aircraft and recovery aircraft. Recovery aircraft have a routing that is the exact reverse of delivery aircraft. Dock and Stock are the default starting and ending operations of any routing in MPX. Therefore, a delivery aircraft proceeds from Dock to gAPOE and then to gAPOD and Stock via an enroute airfield, while a recovery aircraft proceeds from Dock to gAPOD and then to gAPOE and Stock, also via an enroute airfield. The end demand for the delivery aircraft is set equal to the end demand for recovery aircraft. To approximately capture the effect of a fixed number of entities in the network, we set the demand equal to half the fleet size and the time horizon of the network approximation model to be equal to the average of the sum of delivery and recovery flow times. Therefore:

Demand for delivery aircraft in MPX = Demand for recovery aircraft in MPX = (0.5)(Fleet size of aircraft).

Time horizon for MPX model = (0.5){(Time horizon of simulation model)(Fleet size) / Cargo delivered}.

Cargo is modeled as a component of the bill of material of a loaded aircraft. Stochastic run times are modeled using run time at the equipment corresponding to the operation with appropriate mean. MPX output allows us to split the work in progress (WIP) at each equipment group into number of parts in process and in queue. The average number of parts in process at the equipment group indicates the number of aircraft occupying ramp space, while the average number of parts in queue indicates the number of aircraft waiting for ramp space.The performance metrics used to compare the simulation and network approximation models are:

- Completion time of airlift operations.
- Flow times of delivery and recovery aircraft.
- Utilization of each airfield.
- Average time spent by an aircraft at each airfield.
- Average number of aircraft at each airfield.

The average delivery (recovery) flow time is defined as the time taken by an aircraft to complete a delivery (recovery) trip. The utilization and the average time at each airfield indicate the ramp space usage and the time spent by an aircraft at each airfield.

3 DESIGN OF EXPERIMENTS

To compare the performance of the two models we conducted three sets of experiments. Experiment (a) evaluated the impact of stochastic ground times and flying times aircraft experience in the network, Experiment (b) evaluated the impact of fleet size on network performance, and Experiment (c) combined network approximations and simulation to improve the performance of the airlift network. For experiments (a) and (b) we ran the simulation model until all cargo had been shipped. After ensuring that the system was in steady state, we

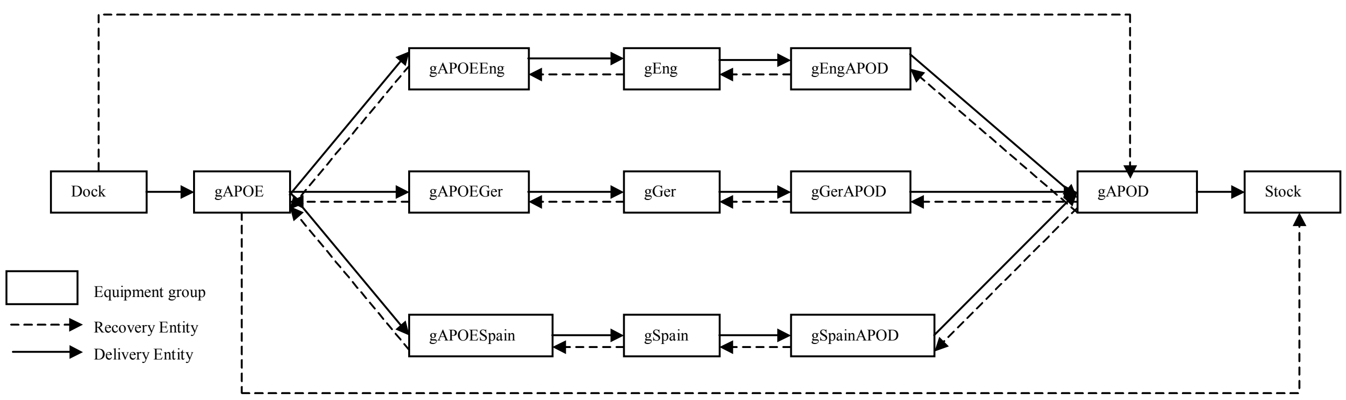

Figure 3: Schematic of the Airlift Network Modeled in MPX®

performed five replications and observed the mean value and the 95% confidence interval of the performance metrics. Then, we ran the corresponding network approximation model and compared its output with that obtained from simulation. We ran the models on a Pentium 230 MHz machine under WINNT 4.0. For each experiment, run times for the simulation model (warm-up plus five replications) were approximately five minutes. In contrast, the network approximation model needed to be run only once, and run times were approximately five seconds. Also, building the simulation model took more time than did the approximation model, because of the higher level of detail. This is evident even by comparing Figures 2 and 3.

3.1 Impact of Stochastic Flying Times and Ground Times

To study the effect of variability in flying times and ground times, we conducted experiments with the same average but varying squared coefficient of variation (SCV). We compared airfield utilization, flow time and number of aircraft in process. Table 2 contains SCV and output data for four scenarios, each using the numerical data shown in Table 1. Because the network is symmetric, we aggregated results for the intermediate airfields.

Table 2: Experiment (a) – Impact of Stochasticity on Network Performance

Scenario	a.1	a.2	a.3	a.4
SCV	0.0	0.5	2.0	5.0

System Performance Measures (hours)

Scenario			a.1		a.2		a.3		a.4	
			Simul.[1]	Appr.[2]	Simul.	Appr.	Simul.	Appr.	Simul.	Appr.
Completion Time			500	NA[4]	529	NA	545	NA	556	NA
Flow Time	Delivery	Average	30.6	31.3	32.4	32.5	33.6	38.9	34.3	41.2
		95% C.I.[3]	-		32.3 – 32.5		33.3 – 33.9		33.9 – 34.7	
	Recovery	Average	29.3	31.3	30.91	32.5	31.64	38.9	32.0	41.2
		95% C.I.	-		30.7 – 31.0		31.4 – 31.8		31.7 – 32.4	

[1]. Simulation [2]. Approximation [3]. 95% Confidence Interval [4]. Not Applicable

Utilization of Location (in %)

Scenario		a.1		a.2		a.3		a.4	
		Simul.	Appr.	Simul.	Appr.	Simul.	Appr.	Simul.	Appr.
APOE	Average	12.9	13.3	13.1	13.3	13.6	13.3	13.2	13.3
	95% C.I.	-		13.1 – 13.1		13.2 – 13.8		13.1 – 13.3	
APOD	Average	86.2	88.7	87.7	88.7	88.0	88.7	88.66	88.7
	95% C.I.	-		87.5 – 88.0		87.3 – 88.7		87.7 – 89.5	
England, Spain, Germany	Average	86.2	88.7	87.5	88.7	88.7	88.7	89.0	88.7
	95% C.I.	-		87.1 – 87.8		87.7 – 89.8		87.7 – 90.2	

Average Time in Location (hours)

Scenario		a.1		a.2		a.3		a.4	
		Simul.	Appr.	Simul.	Appr.	Simul.	Appr.	Simul.	Appr.
APOE	Average	9.9	10.0	9.9	10.0	10.0	10.0	9.9	10.0
	95% C.I.	-		9.8 – 10.0		9.9 – 10.0		9.9 – 9.9	
APOD	Average	9.9	10.0	9.9	10.0	9.9	10.0	9.9	10.0
	95% C.I.	-		9.8 – 10.0		9.8 – 10.0		9.8 – 10.0	
Queue before APOD	Average	0.9	0.7	1.5	0.9	1.9	3.2	2.4	3.8
	95% C.I.	-		1.3 – 1.6		1.8 – 2.0		2.0 – 2.9	
England, Spain, Germany	Average	4.9	5.0	5.01	5.0	4.9	5.0	5.0	5.0
	95% C.I.	-		4.96 – 5.07		4.9 – 5.0		4.9 – 5.0	
Queue before England, Spain, Germany	Average	0.5	1.6	1.9	2.5	2.6	6.8	2.9	8.6
	95% C.I.	-		1.8 – 2.0		2.4 – 2.8		2.7 – 3.1	

Average Content of Location (units)

Scenario		a.1		a.2		a.3		a.4	
		Simul.	Appr.	Simul.	Appr.	Simul.	Appr.	Simul.	Appr.
APOE	Average	25.8	27.0	26.3	27.0	27.2	27.6	26.5	26.6
	95% C.I.	-		26.1 – 26.6		26.7 – 27.8		26.3 – 26.8	
APOD	Average	25.8	27.0	26.4	27.0	27.1	27.6	26.5	26.6
	95% C.I.	-		26.1 – 26.6		26.6 – 27.5		26.3 – 26.7	
Queue before APOD	Average	2.5	3.4	3.9	5.6	5.2	17.6	6.5	19.9
	95% C.I.	-		3.6 – 4.3		4.7 – 5.6		5.5 – 7.5	
England, Spain, Germany	Average	8.6	9.0	8.8	9.0	9.0	9.2	8.8	8.9
	95% C.I.	-		8.7 – 8.9		8.9 – 9.2		8.7 – 9.0	
Queue before England, Spain, Germany	Average	0.8	3.0	4.7	4.5	4.7	12.5	5.3	15.1
	95% C.I.	-		4.3 – 5.0		4.4 – 5.0		4.8 – 5.8	
Other	Average	72.3	75.6	73.3	75.6	76.5	75.6	75.0	74.4

This first experiment gives several insights:

- Increasing SCV tends to increase airlift completion time and flow times. As would be expected, when there is more and more randomness in the system it takes longer to complete the operation.
- Utilizations remain roughly the same, since only the SCV of flying times and ground times are changed and not the mean. Similarly, average ground times and average number of aircraft at airfields do not change significantly.
- However, increasing SCV tends to increase the waiting times aircraft experience before landing at airfields. Also, the waiting queues before airfields with limited ramp capacity contain more aircraft and for longer times. Both simulation and network approximation models indicate that variability can cause increases by up to a factor of 6 in waiting times and average number of aircraft waiting for ramp space, though the divergence between their predicted waiting times increases with SCV. A deterministic model would miss this significant impact of stochasticity on airlift performance.

3.2 Impact of Fleet Size with Stochastic Ground Times

To study the impact of fleet size on system performance, we ran the models with 20, 50, 80 and 100 aircraft in each fleet, comparing airfield utilization, flow times and numbers of aircraft in process. For each of these scenarios the numerical data are as in Table 1. Ground times are stochastic with SCV of 5.0. Because the network is symmetric, results for intermediate airfields are aggregated. Tables 3a and 3b show the results of this experiment.

Table 3a: Experiment (b) – Impact of Fleet Size on Network Performance

Scenario	b.1	b.2	b.3	b.4
Fleet size	20	50	80	100

System Performance Measures (hours)										
Scenario			b.1		b.2		b.3		b.4	
			Simul.[1]	Appr.[2]	Simul.	Appr.	Simul.	Appr.	Simul.	Appr.
Completion Time			1,472	NA[4]	616	NA	533	NA	525	NA
Flow Time	Delivery	Average	29.0	29.0	30.7	34.0	45.2	60.7	56.9	64.2
		95% C.I.[3]	28.8 – 29.2		30.4 – 31.1		44.2 – 46.2		54.4 – 59.4	
	Recovery	Average	28.9	29.0	30.2	34.0	38.4	60.7	46.0	64.2
		95% C.I.	28.6 – 29.1		30.0 – 30.4		37.3 – 39.4		43.9 – 48.2	

[1]. Simulation [2]. Approximation [3]. 95% Confidence Interval [4]. Not Applicable.

Utilization of Location (in %)									
Scenario		b.1		b.2		b.3		b.4	
		Simul.	Appr.	Simul.	Appr.	Simul.	Appr.	Simul.	Appr.
APOE	Average	5.1	5.2	12.1	12.2	14.2	14.1	14.1	14.2
	95% C.I.	5.1 – 5.2		12.0 – 12.1		14.0 – 14.4		13.9 – 14.4	
APOD	Average	34.1	34.4	81.0	81.4	93.6	94.2	95.7	94.7
	95% C.I.	33.9 – 34.3		80.5 – 81.5		92.6 – 94.6		94.6 – 96.7	
England, Spain, Germany	Average	34.3	34.4	81.5	81.4	94.3	94.2	95.5	94.7
	95% C.I.	33.9 – 34.5		80.4 – 82.5		93.1 – 95.6		94.1 – 96.9	

Average Time in Location (hours)									
Scenario		b.1		b.2		b.3		b.4	
		Simul.	Appr.	Simul.	Appr.	Simul.	Appr.	Simul.	Appr.
APOE	Average	10.0	10.04	9.9	9.9	10.0	10.0	9.9	10.0
	95% C.I.	9.9 – 10.1		9.8 – 10.0		9.9 – 10.1		9.8 – 10.0	
APOD	Average	9.9	10.04	9.9	9.9	9.9	10.0	10.0	10.0
	95% C.I.	9.9 – 9.9		9.8 – 10.0		9.8 – 10.0		9.8 – 10.1	
Queue before APOD	Average	0.0	0.0	0.5	1.2	8.6	10.9	16.7	12.4
	95% C.I.	-		0.4 – 0.7		7.3 – 9.9		12.7 – 20.5	
England, Spain, Germany	Average	4.9	5.08	5.0	4.9	5.0	5.0	4.9	5.0
	95% C.I.	4.9 – 5.0		4.9 – 5.0		4.9 – 5.0		4.9 – 5.0	
Queue before England, Spain, Germany	Average	0.0	0.12	1.2	3.8	8.4	20.0	14.0	23.2
	95% C.I.	-		1.1 – 1.3		7.7 – 9.2		12.0 – 15.8	

Table 3b: Experiment (b) – Impact of Fleet Size on Network Performance

Average Content of Location (units)									
Scenario		b.1		b.2		b.3		b.4	
		Simul.	Appr.	Simul.	Appr.	Simul.	Appr.	Simul.	Appr.
APOE	Average	10.3	10.3	24.2	24.4	28.4	28.3	28.3	28.4
	95% C.I.	10.2 – 10.4		24.0 – 24.3		27.9 – 28.8		27.9 – 28.8	
APOD	Average	10.2	10.3	24.3	24.4	28.0	28.3	28.7	28.4
	95% C.I.	10.1 – 10.3		24.1 – 24.4		27.8 – 28.3		28.4 – 29.0	
Queue before APOD	Average	0.0	0.0	1.4	6.2	24.4	61.8	48.1	71.1
	95% C.I.	-		1.1 – 1.7		20.8 – 28.0		37.1 – 59.1	
England, Spain, Germany	Average	3.4	3.4	8.1	8.1	9.4	9.4	9.5	9.5
	95% C.I.	3.3 – 3.4		8.0 – 8.2		9.3 – 9.5		9.4 – 9.6	
Queue before England, Spain, Germany	Average	0.0	0.1	2.0	6.2	15.9	39.1	26.8	44.2
	95% C.I.	-		1.8 – 2.2		14.4 – 17.4		23.3 – 30.3	
Other	Average	28.6	28.8	67.6	68.4	78.2	79.2	79.2	79.8

This second experiment gives several insights.

- Very small fleet sizes tend to increase airlift completion time. In this case there is no competition for ramp space between aircraft. On the other hand, very large fleet sizes could reduce airlift completion time (by 64% for our data set), but flow times could inflate (by about 100% for our data set).
- Increasing fleet size tends to increase utilization of ramp space. After a point, increasing fleet size also results in increasing wait for ramp space. Correspondingly, waiting time and number of aircraft waiting for ramp space increase considerably.

Since both the simulation and network approximation models incorporate stochastic elements, we are able to demonstrate the significant impact fleet size has on delivery flow times.

3.3 Combination of Network Approximations and Simulation in an Improvement Process

Experiments (a) and (b) seem to indicate the importance of including stochastic effects in any modeling tool used for analyzing or designing airlift networks. In addition, they show that network approximations require much less computation time (a 98% reduction in this case), while simulation enables us to estimate performance with better accuracy. In experiment (c) we combined network approximations and simulation in order to improve an airlift network. Since network approximations are faster we used a simple design of experiments to predict the impact of different variables on network performance. We then confirmed the predicted improvement with one simulation run on the variable set that promised best system performance.

Specifically, in the network approximation model we used a 2^4 factorial design, with four quantitative variables: SCV, ramp space at APOD and intermediate airfields, flying times and ground times at airfields, and fleet size. The responses are the delivery flow time and the utilization rates of intermediate airfields. After analyzing the results of our experiment, we adjusted one parameter accordingly in the airlift network and made one simulation run to see whether the improvement predicted by the network approximation model would actually happen. Table 4 shows the input data for the two levels of the factorial design.

Table 4: Experiment (c) – data for the 2^4 factorial design

	SCV	Fleet Size (units)
Low	0.5	40
High	5.0	60

	Ramp Capacity (units)		Average Run Times (hours)		
	APOD	England, Germany, Spain	APOE, APOD	Flying time	England, Germany, Spain
Low	48	16	8	5.6	4
High	72	24	12	8.4	6

Figures 4 and 5 show the main and two-factor interaction effects on the utilization rates of intermediate airfields and delivery flow times.

Ramp space has a large main effect on intermediate airfield utilization rates, and some main effect on delivery flow times. Therefore, if we consider scenario (b.3) we can hypothesize that increasing APOD and intermediate airfields ramp space will reduce (i.e. improve) these performance measures. Indeed, Table 5 gives the main performance measures of the airlift network, obtained from the network approximation and simulation models, where the data are as in scenario (b.3) except that the APOD and intermediate airfields have ramp space of 45 and 15 respectively. Utilization decreased by 9.1% and delivery flow time by 31.6%. The completion time for MPX in Table 5 is deduced from the delivery flow time and number of trips needed to achieve a cargo requirement of 500.

As a check, we ran the same factorial design using simulation instead of network approximation methods. As might have been expected, we found substantial differences, reflecting the inaccuracy of some of the network approximations. However, the key point is that the approximations proved accurate enough to guide network improvement, at a 98% reduction in computation time.

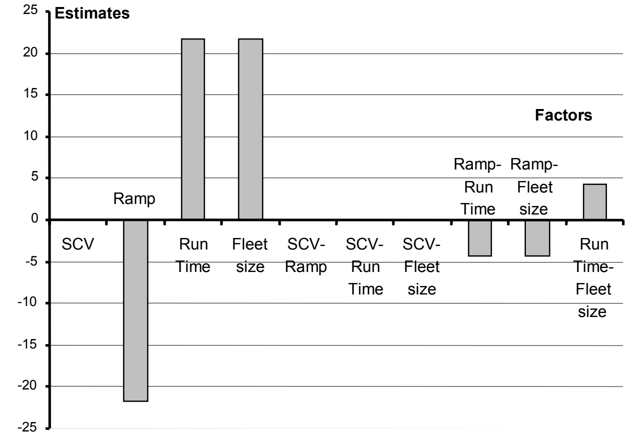

Figure 4: Experiment (c) – Main and Interaction Effects on Utilization (in %) of Intermediate Airfields

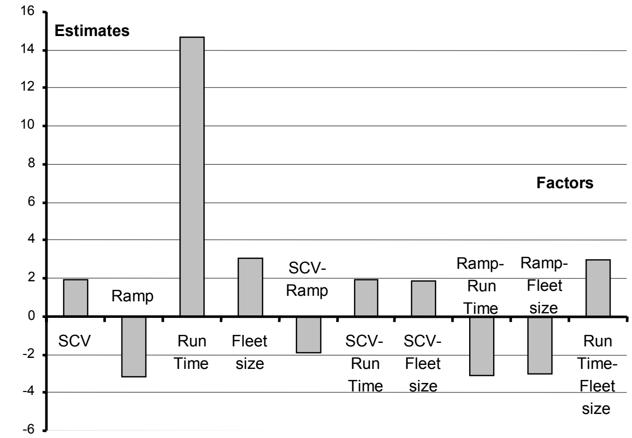

Figure 5: Experiment (c) – Average Main and Interaction Effects on Delivery Flow Time (in hours)

Table 5: Performance Measures of Improved Scenario (b.3)

System performance measures (in hours)			
		Approximation	Simulation
Delivery flow time	Average	31.9	30.9
	95% C.I.		30.6 – 31.2
Completion time		398	391
Utilization of Location (in %)			
		Approximation	Simulation
APOD	Average	82.3	84.5
	95% C.I.		83.5 – 85.5
England, Spain, Germany	Average	82.3	85.7
	95% C.I.		84.7 – 86.7

One might wonder why we did not use a more sophisticated optimization method than a factorial design. In particular, since much of the computational work to obtain network approximations involves solution of equations, one would expect better results from incorporating these equations into the improvement process, rather than just using a few outputs. We entirely agree, but because this was a preliminary experimental effort we used off-the-shelf software, and MPX® permits access only to outputs, not to the equations that produce them. In future developmental work we hope to use the structure of the approximations to much better advantage than we were able to do here.

4 CONCLUSION

This experimental study on a very simple airlift network gives several insights. First, we see that increased stochasticity of ground times results in a considerable increase in waiting time at airfields. Therefore, performance in a given scenario could be quite different from what would be predicted by a deterministic model. Second, our experiments indicate that fleet size has opposing impacts on completion times of airlift operations and on delivery flow times. Third, we found that network approximation models could be run in about 2% of the time required for simulation models, though they could also result in loss of accuracy. Our experiments indicated that when utilization rates of airfields were not too high, the estimates of all performance measures matched satisfactorily. At high utilization rates, although the average content in the whole network matched closely between simulation and network approximation models, they differed in their estimates of average content and average time in waiting queues. There could be several reasons for this difference. One is that our choice of commercially available network approximation software, namely MPX, required modeling the airlift network as an open network of queues. Another is that network approximations may not be very accurate when a network operates under conditions of heavy traffic. Nevertheless, for reasonable network loads, network approximations model key performance features with sufficient accuracy, while requiring only a small fraction of the computational work that would have been needed had we solely relied on simulation.

The work reported above leads us to suspect that a combination of simulation and network approximations should yield substantially better performance than either one of these alone. Roughly speaking, one would use the network approximations to explore the variable space and identify parameter values that promise improvements in system performance, then validate these using simulation. In our study, network approximations required only about 2% of the time needed to generate the same quantities by simulation, and we suspect that the advantage in speed would be even more pronounced for a larger and more complex network. Thus, using approximations for exploration should save considerable amounts of time. We expect that using the structure of the approximations explicitly (rather than outputs only, as we did here) should make the technique more effective. Overall, we concluded that it should be worth the effort to develop a combination of these two methods, and that such a combination could have significant potential for practical benefit.

ACKNOWLEDGMENTS

We thank Professors Richard E. Rosenthal and Laura M. Williams of the Naval Postgraduate School for their assistance, which included furnishing valuable airlift scenario data.

Support provided by the Center for Quick Response Manufacturing, University of Wisconsin-Madison, is also gratefully acknowledged.

This material is based on research sponsored in part by the Air Force Research Laboratory under agreement number F49620-01-1-0040, and in part by the U.S. Army Research Laboratory and the U.S. Army Research Office under grant numbers DAAG55-97-1-0324 and DAAD19-01-1-0502. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright notation thereon. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the sponsoring agencies or the U.S. Government.

REFERENCES

Baker, S. F., D. P. Morton, R. E. Rosenthal, L. M. Williams. 1999. Optimizing Strategic Airlift. Technical Report NPS-OR-99-004, Naval Postgraduate School, Monterey, California.

Bitran, G. R., and R. Morabito. 1996. Open Queueing Networks: Optimization and Performance Evaluation Models for Discrete Manufacturing Systems. *Production and Operations Management* 5:163-193.

Brown, E. 1988. IBM combines rapid modeling technique and simulation to design PCB factory-of-the-future. *Industrial Engineering,* June 1988, 23 ff.

Network Dynamics, Inc. 1999. *MPX User Manual, Software Version WIN 3.3.* Framingham, MA: Network Dynamics Inc., `<http://www.networkdyn.com>`.

Niemi, A. 2000. Stochastic modeling for the NPS/RAND Mobility Optimization Model. Technical Report, Department of Industrial Engineering, University of Wisconsin-Madison, Madison, Wisconsin. Available online via `<http://ieserv3.ie.wisc.edu/robinson/Niemi.htm>`.

Rousseau, G. 1998. Airlift Flow Model (AFM) Testing Process. Unpublished PowerPoint presentation, Air Mobility Command Studies and Analysis Flight, Scott AFB, IL.

Suri, R. 1998. *Quick Response Manufacturing: A Company-wide Approach to Lead Time Reduction.* Portland, OR: Productivity Press.

Suri, R., J. L. Sanders, and M. Kamath. 1993. Performance Evaluation of Production Networks. In *Handbooks in OR & MS, Vol. 4,* Eds. S.C. Graves et al., 199-286. Amsterdam: Elsevier.

AUTHOR BIOGRAPHIES

JULIEN GRANGER is Research Assistant in the Department of Industrial Engineering at the University of Wisconsin-Madison. His e-mail and web addresses are: <granger@cae.wisc.edu>and <http://www.cae.wisc.edu/~granger/homepage.shtml>.

ANANTH KRISHNAMURTHY is a Ph.D. candidate in the Department of Industrial Engineering at the University of Wisconsin-Madison. He is currently working on performance analysis of manufacturing systems. His e-mail and web addresses are: <akrishna@cae.wisc.edu> and <http://www.cae.wisc.edu/~akrishna>.

STEPHEN M. ROBINSON is Professor of Industrial Engineering and Computer Sciences at the University of Wisconsin-Madison. His research interests are in quantitative methods in managerial economics, especially deterministic and stochastic optimization and methods to support decision under uncertainty. He is a member of INFORMS, the Mathematical Programming Society, SIAM, and IIE. His email and web addresses are <smrobins@ fac-staff.wisc.edu> and <http://www.engr.wisc.edu/ie/faculty/robinson_stephen.html>.