Proceedings of the 1997 Winter Simulation Conference
ed. S. Andradóttir, K. J. Healy, D. H. Withers, and B. L. Nelson

SOLVING ENGINE MAINTENANCE CAPACITY PROBLEMS WITH SIMULATION

Robert Gatland
Delta Air Lines, Inc.
Department 018
Post Office Box 20706
Atlanta, GA 30320-6001, U.S.A.

Eric Yang
Delta Air Lines, Inc.
Department 017
Post Office Box 20706
Atlanta, GA 30320-6001, U.S.A.

Kenneth Buxton
Economic Development Institute
Georgia Institute of Technology
1770 Indian Trail Road, Suite 410
Norcross, GA 30093, U.S.A.

ABSTRACT

For many companies the scheduling of job shops proves very difficult. Competing priorities (due dates, potential profit) from multiple customers confronts the plant manager with many tough questions; which product to work next, can this order wait, and what products to produce together to maximize the machine utilization and minimize setups. These same situations are all too common for the manager of an engine maintenance facility. Most airlines fly multiple aircraft types with a unique engine type for each aircraft. Engine overhauls are usually within a single facility at the airline's main hub. With so many different parts and repair routings, a job shop environment is common for an engine maintenance facility. This paper will provide an overview of an engine maintenance system, clearly define the capacity planning problem, and describe the goal and scope of the simulation model and its applicability. The model development, experiments, results, and future direction will also be discussed.

1 INTRODUCTION

The business strategy of a particular airline may differ depending on the market niche. However, the overriding theme is safety and reliability. Customers want a safe and on-time airline. Depending on the cost strategy, the significance of reliability may vary, but canceling a flight for an unavailable engine is unacceptable. Therefore, the operations strategy of airlines is to provide a quality engine at the least cost. Meeting the operational demand of the airline is paramount to the company. Thus, the engine maintenance department must carefully balance the workload in its shops to generate the greatest throughput in the least amount of time.

Additionally, today's airline industry is becoming more and more competitive. The price of an airline ticket is about the same as it was 10 years ago. To compete, an airline must continually look for ways to reduce cost as well as generate more revenue and do more without increasing capacity. As Richard Cobb (1995) notes, "for today's airline maintenance organizations, there is an increased demand for high-quality work and service at low cost" (p. 25). Cobb (1995) continues noting that "as airlines adjust to the realities of today's competitive environment and the industry's overall poor financial performance, greater attention is being given to controlling cost, both in-house maintenance operations and for maintenance services contracted from outside maintenance vendors" (p. 25). One way previously untapped was the use of maintenance as a source of income. Maintenance has traditionally been a cost center to an airline. In fact, maintenance costs are rarely understood by management (Condom, 1994).

Today, engine maintenance departments are striving to become profit centers. Many airlines are looking for insourcing (revenue generating work from outside the company) to reduce their maintenance cost. Airlines with small fleets can achieve a cost savings by outsourcing to larger airlines (Condom, 1994). Small airlines don't want to invest the capital to perform their own engine maintenance. This allows the larger airlines, like Delta, to insource the smaller airline's maintenance. In fact, the larger airlines are marketing their capabilities to save their own jobs (Condom, 1996). However, insourcing is a complexity previously unknown to airlines. Previously an airline had to worry about only one customer, itself. With insourcing maintenance departments now has multiple customers with competing priorities. Without the right systems in place, an airline can easily overbook the amount of work its maintenance shops can handle. The use of simulation for capacity planning and facility loading will help solve this problem.

At Delta Air Lines, a simulation model of the engine maintenance facility will provide solutions to capacity planning problems. Delta has over 550 aircraft made up of seven different aircraft types and eleven different engine types. The main Technical Operations Center (TOC) in Atlanta, Georgia, overhauls all engine types. However, engine maintenance does not process engines alone. Engine maintenance repairs engines, auxiliary power units (APUs), landing gear, and other miscellaneous parts from the hangar and component shops. Engine parts must compete with all these parts for equipment, personnel, and priority. The addition of insourcing only compounds the problem.

Engine removals occur for a variety of reasons. First, the engine has parts that are time restricted either by the manufacturer or the FAA. These parts must be removed, inspected, and repaired before their time expires. Second, the engine is boroscope inspected (a tube is inserted into the engine for viewing inner parts either by video or eye) on given intervals to determine wear. If the wear of particular parts is beyond limits, the engine is removed and overhauled to prevent a failure. Additionally, inspection of the engine occurs if performance is becoming deficient. The engines EGT (exhaust gas temperature) margin may force removal of the engine. The goal is to remove the engine before a failure. The final reason for removal is an engine failure. This is a highly undesirable option. A failure may triple the cost of overhaul.

Upon arrival in Atlanta, the engine is workscoped (preliminary inspection) to determine the extent of repairs required. The workscope determines which modules to remove. When a bay (work area) becomes available, the next engine in queue will begin disassembly. At disassembly, the engine will subdivide into modules, subassemblies of an engine. The modules disassemble into piece parts. The engine manual dictates the parts repair routing. Production Control generates a Bill of Material (BOM) based on the parts removed. The scheduling system uses the BOM, inventory and work in process (WIP) to assign parts to a shortage list and derive priorities for part repair. Certain parts are processed through a cellular manufacturing shop developed based on group technology. Blades and vanes, stators, and thrust reverser parts have dedicated focused shops exclusively for these type parts. In addition, cases and landing gear parts are repaired in virtual manufacturing cells. These cells have some dedicated equipment and other equipment located throughout the facility dedicated to these types of parts. Parts not dedicated to a focused or virtual cell are routed through functional areas (including machine shop, welding, heat treat, plasma spray, plating, cleaning, and non-destructive testing).

2 PROBLEM DEFINITION

The engine maintenance department meets the demand of the airline, but at a very high cost. Frequently, parts and engines are produced using overtime. Additionally, excess parts are purchased to ensure engines are completed on schedule. Furthermore, excess capacity is built into the system to repair parts rerouted for rejects or other reasons. All these factors add to the cost of repairing engines. To remain competitive with other airlines and third party repair vendors, engine maintenance costs must be reduced.

Insourcing work from other airlines can produce revenue to offset engine maintenance expenses. However, to efficiently insource additional work the maintenance department must know its available capacity. To bring work into a maintenance department without any forethought will overload the system. This will lead to expenses greater than the income from insourcing. Also, this results in a reduced service level to the airline. To better plan the level and timing of insourcing, the engine maintenance department needs a model that can review all current engines and parts in repair and their routings and be able to compare available capacity with the requirements for the insourced work. Additionally, the simulation model must be able to test multiple scenarios of insourced work from the engine to the piece part level.

Unpredictable engine removals are another problem for the airline. Engine removals are classified by the amount of hours or cycles the engine has flown. While the engine removals portray a stable system, the variation of removals is rather large. It's this unpredictability that causes "waves" in the removal of engines. The model will reveal the interactions that cause these waves by testing multiple input scenarios. The engine maintenance goal is to meet the flying needs of the airline and have replacement engines available at the time of removal. Additionally, the engine maintenance department must be able to handle the sudden influx of one type of engine. To understand these problems, the model will measure customer satisfaction based on service level.

A major difference between manufacturing and remanufacturing is the routing of parts through the system. In general, a manufacturing facility will have one routing for production. The process is predictable. However, in a remanufacturing environment the part routing is dependent on a number of conditions. As Guide (1993) points out, "the number of operations required and the time required at each repair operation may be unique for each part since each unit may have varying amounts of wear" (p. 908). The length of flying time, the flying conditions, and the nature of earlier repair work will influence repair requirements. Considering these factors, the part routing will vary. Guide (1993) notes, "part mixes are probabilistic because what part needs repair is not known until the product is inducted, disassembled, and inspected" (p. 911). Some parts will not need repair. One part may require only superficial work, while another part requires a full overhaul. The uncertainty of requirements is more pronounced in the remanufacturing environment (Guide, 1993). It's this stochastic nature of parts routing that adds an extra level of complexity to the engine maintenance capacity planning problem.

Another problem in remanufacturing is the variable processing times of each step in the part repair routing. While almost every industry has some level of variable process time work, the remanufacturing repairs are more widely distributed. As noted earlier, the routing of the part will vary based on the part condition. This varying level of damage has an effect on the processing time of the part. If severely damaged, the same repair step may vary significantly. These large repair time distributions must be taken into account in any capacity model. Again, the system is more complex than a normal manufacturing facility.

3 GOAL OF SIMULATION

The goal of this simulation model is to better understand the production capability of the engine maintenance facility by investigating the affects of facility loading on turn time, throughput, and capacity. The simulation will provide data on the number of engines produced over a year's time-frame, the number of engines produced each week, and the time necessary to produce each engine. Additionally, the utilization rate for each person and machine will be available for comparison purposes. The simulation model will be flexible enough to handle a multitude of scenarios. Scenarios may include: engine mix, engine repair levels, level of insourcing, type of insourcing, and additional non-engine work from within the airline.

4 APPLICABILITY OF SIMULATION

Over the past few years it became apparent that better decision support tools and methods were needed in the maintenance department. Simulation is a valuable tool because it can handle complex requirements and stochastic processing times (Harvey et al., 1992). Also, simulation works especially well in diagnosing how systems respond to changes in flow patterns (Cobb, 1995). For these reasons, Delta chose the ARENA simulation package as its analytical tool. The ease of model development without a knowledge of the program language was a key selling point.

For engine maintenance capacity planning, the use of another analytical tool, such as a spreadsheet, would be too cumbersome. Simulation in this type of complex environment provides the capability of changing many variables simultaneously. The spreadsheet approach is too limited. Responsibility for over ten thousand different parts, all with somewhat unique routings and repair times requires a more complex model. Profiling the potential routings, repair times, and equipment availability can be done only using simulation.

5 MODEL DEVELOPMENT

This section will provide an overview of the use of simulation to model an engine maintenance facility. The logic behind BOM generation and piece part routings will be explained, as well as the control mechanisms and assumptions.

The model was subdivided into shop segments. This was done to make development and validation and verification easier. Measuring throughput and resource utilization to validate the model (Harvey et al., 1992) provides the best method to verify the accuracy of each individual shop before the models are combined. Within the shop model, each piece of equipment was given a unique resource (ARENA element). Past records of equipment downtime were gathered and used to calculate an equipment failure rate. Similar machines were organized together under a set. Therefore, any part being routed to a specific type of machine could choose the next available similar machine without waiting for a specific machine. In addition, selected series of operations using a single workstation were combined for model simplicity.

All shop personnel are included in the model. The personnel were added to resource sets based on their capabilities. A set of personnel was developed for each machine or set of machines. Thus, only those qualified to run a particular type of machine could work a part at that work area. The shift schedules for each employee along with their vacation and sick days were included in the model. Some personnel work 5-day weeks, and others work 4-day weeks. There is a large first shift operation with a minor second shift and skeleton third shift.

A singular engine shop was developed for simplicity purposes. As Law (1990) notes, “it is generally not necessary or feasible to have a one-to-one correspondence between every element of the system and every element of the model" (p. 16). In the future, this will be further expanded to allow for individual engine shop areas based on the engine type. This singular shop will process the engine (parent) through the disassembly, create parts (children), send parts to repair, and assemble the engine. The quantity of engines in this area will be maintained continuously. When one engine is completed, another will enter the system. The engine will be the main entity (parent), with a split into children based on the workscope of the engine. If the engine requires major repairs, the engine will produce a larger

BOM. Minor repairs will translate into smaller BOMs. Parts will have a sequence of steps routing them through the resources of the repair facility. Once all parts for the engine are complete, the engine will begin assembly. While in reality the engine would begin assembly before all parts are available, this model assumes assembly will not begin until all parts are available and the assembly time is shortened respectively. Future models will include the capability to model engine areas separately and to start assembly before all parts are available.

The engine shop is also responsible for many non-engine parts. These include landing gear, APUs (auxiliary power units), component parts, hangar parts, and potential insourcing. Except for the insourcing of complete engines, these parts will be routed though the facility as generic parts, and their turn times will not be captured. The reason for including these parts is to factor in any weight they may have with regard to the capacity usage and requirements. Without these parts the model would be invalid.

Data collection is probably the hardest part of any model development. This was a major problem in this model development. Past systems contained little applicable data. Data was required for:

- <u>equipment downtime</u> - system data from equipment maintenance department provided the number of equipment failures. However, the amount of time the equipment was unavailable was not captured in the database. This data was developed by consultation with the shop foreman, equipment users, and equipment maintenance personnel.
- <u>personnel on vacation</u> - vacation time based on years of seniority were added to the model.
- <u>personnel out sick</u> - the average sick time for all employees in a department was added to the model.
- <u>part routings</u> - complete part routings were captured from developed job planning cards maintained by engineering.
- <u>part process times for each individual step</u> - some information was attainable from a production count system. Where this data was not available, discussions with shop personnel led to a minimum, maximum, and average time for each operation. These numbers were later verified in test runs of the model.
- <u>routing step frequencies</u> - the percent performed of each operation was estimated by the inspector developing the part routing since no system data was available.
- <u>part reject rates for each individual step</u> - reject percentages for each step were gathered from a count of historical reject cards divided by the number of times the step was completed.
- <u>part removal rates from next higher assembly</u> - part removal rates were calculated from the tracking system based on the number of tracking numbers created divided by the number of next higher assembly removals.
- <u>group technology module rates based on engine workscope</u> - past engine BOMs for modules were compared with workscopes to develop a group technology base for modules removed from engines based onworkscope. This data developed a percent removed of modules based on the workscope.
- <u>engine removal workscope</u> - the type of engine repair required was derived from past workscope data.

The generation of demand and the BOM explosion were generated by:

1. Determine level of engine repair (workscope)
2. Determine module removals
3. Determine part removals
4. See Figure 1 for BOM explosion method.

Those parts designated for removal would then be compared against the repair routings database. The database contains the complete part routing for every part. The database contains part number, step number, workstation, step frequency, process time distribution, reject percentage, and reject step (step the part returns to if rejected). See Table 1 for example of part routing database.

The domain of the capacity planning model is within the repair areas of engine maintenance. The assembly and disassembly areas are constrained to a specific number of engines. A push environment is assumed (all parts are released into rework immediately after disassembly). The prioritization of parts in repair is based on first-in first-out (FIFO). No allowances for tooling availability have been made at this time. The transportation time and computer transaction times are considered constant at this time for model simplicity, but will be reviewed at a later time.

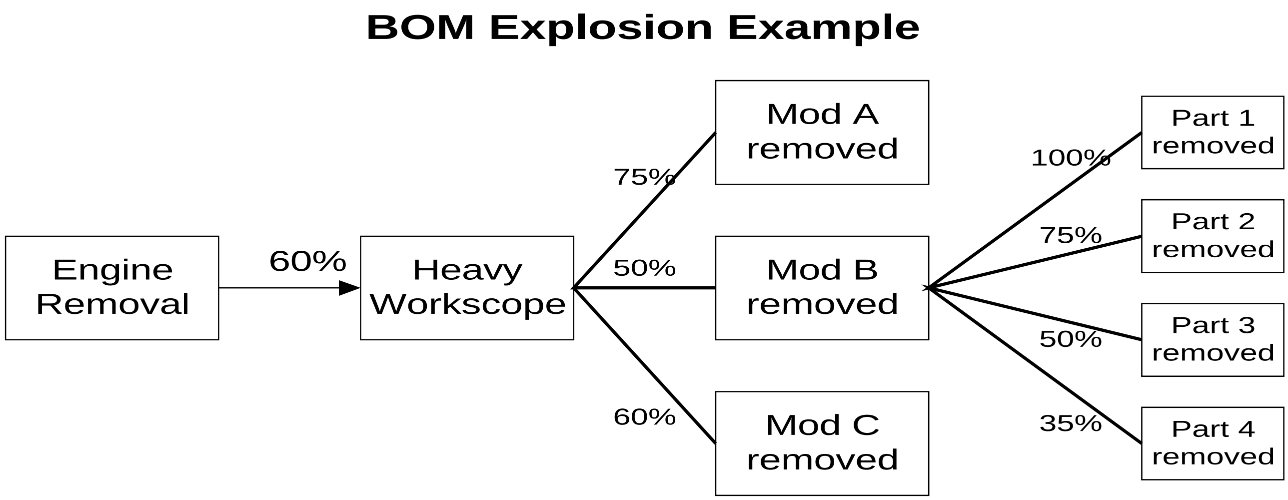

Figure 1: Bom Explosion Example

Table 1: Part Routing Database Example

Part	Step	Station	Frequency	Process Time	Reject %	Reject Step
A	1	50	1.00	Norm (45,12)	0.00	1
A	2	51	0.80	Tria (30,45,60)	0.10	1
A	3	55	1.00	Tria (20,25,30)	0.15	1
B	1	12	1.00	Norm (25,9)	0.00	1
B	2	30	1.00	Tria (120,145,170)	0.05	1
C	1	50	1.00	Tria (30,45,60)	0.00	1
C	2	12	0.30	Norm (30,5)	0.05	1
C	3	55	0.80	Norm (45,12)	0.03	2
C	4	30	1.00	Norm (60,15)	0.07	2

6 EXPERIMENTS

To determine the capacity of the engine maintenance facility, a number of experiments were run:

1. Vary Engine Removals
   - Engine removals at deterministic times for all engine types
   - Engine removals at stochastic times based on past history
2. Vary Engine Disassembly Start Times
   - Engine disassembly start times evenly distributed
   - Engine disassembly start times bunched on certain dates
3. Vary Disassembly Work Schedule
   - Engine disassembly working 7 days a week
   - Engine disassembly working 5 days a week
4. Vary Engine Workscope Mix
   - Randomly mix engines
   - Process all heavy repair engines first
   - Process all light repair engines first

Varying engine removals will show if there is enough inventory of each engine type and if the delay of available inventory produces a wave through the system. Varying start times will demonstrate any effects of improper loading of the system. Varying work schedules will show if the production of the repair shops improves with a uniform disassembly schedule. Finally, varying the product mix will allow analysis on the loading of the repair shops. To understand the implications of each alternative, the following measures were gathered:

- Engine turntime (time from start of disassembly to assembly completion)
- Engine throughput (the number of each engine type produced during experiment)
- Engine service level (the percentage of engines serviceable to the number required)
- Machine utilization rates
- Personnel utilization rates
- Part turntime (time from generation of part to serviceable)
- Part throughput (number of parts produced)
- Average queue time (time each part spends in queue by station and part number)
- Average WIP (work in process) level

Additionally, the variations for the turntimes, utilization rates, throughput, and queue time will be measured to understand the complete effects of the experiments.

7 RESULTS

The results of the first experiment showed no effects on the capacity of the engine maintenance department. Engines were always available for entry into the system in both scenarios. However, this demonstrates a negative characteristic of engine maintenance. Too high an inventory of engines is being maintained. In fact, both scenarios showed that no fewer than 30 engines were constantly in the holding area awaiting entry into disassembly. These extra engines imply an opportunity cost of over $50 million. Also, both scenarios demonstrated the same level of engine turnaround time.

The second experiment showed substantial waves are created in the system if engines are disassembled close together. While each engine has a unique make-up, the routings of parts are similar enough to create this phenomenon. This may be further compounded by the use of certain machines for many types of jobs. The machine shop showed bottlenecks on the CNC machines while the non-CNC machines had available capacity. This clearly shows that the shop's desire to work parts on the CNC machines have exceeded their capacity. Management needs to determine which parts require these machines and transfer the other parts to non-CNC machines. The waves were small or non-existent when the disassembly times were uniformly distributed. Furthermore, this actually reduced the overall turnaround time of the engines and parts. This would result in inventory reductions. The variations were also more controlled in this scenario.

In the third experiment, the longer work period revealed a more even workload. The same effects of the second experiment were demonstrated. The stretching out of the work period to seven days helped evenly distribute the work. There was still some bunching, but not as severe as the five-day work week. Considering these results, another experiment was tested to determine if a seven-day, evenly distributed work schedule would produce even more significant results. As expected, the turnaround time of the engine was reduced and the waves seen previously were rare. While there were some disadvantages to the seven-day work week, it clearly demonstrates that some compromise must be made.

In the fourth experiment, the use of mixed engines produced the best results. Obviously, using all heavy repair engines produced great bottlenecks in the system. The turnaround time was high and the variations were extreme. The throughput was also less than the other two. The use of all light engines reduced the turnaround time, but it still had some waves. This example clearly showed that many finishing operations became overloaded while some machines sat idle. The system was unbalanced. The use of mixed engines, one heavy then one light, provided a more evenly distributed workload. The turntime was slightly greater than all light engines, but the variation was minimized. Also, the throughput was almost identical. This probably has a greater effect on the workforce. An unbalanced workload creates stress in a company. With the mixing of repair levels a balanced environment can be achieved.

All these experiments clearly showed that the entry procedures have a great effect on the capacity of the engine repair area. In fact, the order release procedure has a great effect on the system performance (Rogers and Flanagan, 1991). The WIP flowed in waves if the releases were not evenly distributed. A longer work week in disassembly may solve this problem. Heavy repair engines compounded this problem.

8 FUTURE EXPERIMENTS

As the results show, the capacity of the engine shop can show wide fluctuations based on the input of specific engines into the overhaul facility. Disassembling many engines at one time also creates waves in the system. In today's environment these entries are uncontrolled. These uncertainties lead to high inventories of engines and piece parts in an attempt to reduce the engine turntime. Therefore, plenty of unserviceable engines and piece parts (i.e., raw materials to the remanufacturing environment) are available to choose from to enter the rebuild system. Working with many different engine types in one facility presents many scheduling problems. The master scheduler asks, how do I best load the system to meet the operational needs of the airline? The use of this model combined with a control mechanism for engine entries will solve this problem. As Rogers and Flanagan (1991) note, "the … benefit of on-line simulation comes from using it to support real time scheduling decisions" (p. 38). Rogers and Flanagan (1991) continue by saying, "in dynamic manufacturing environments, where uncertainty prevails, it may be advantageous to change the way the shop is controlled at certain points in time" (p. 38). Using simulation, the master scheduler can select the best options based on the current environment (Rogers and Flanagan, 1991). By testing multiple engine releases, the master scheduler can determine which engine will fit best in the current workload to meet the customer demands. Knowing the requirements of the engines for release, the model can simulate the entry into disassembly, repair, and assembly to determine the effects to the entire system. This facility loading capability provides the strongest use of this model for customer satisfaction. By simulating the loading of individual engines, the master scheduler can plan which engines to introduce into the system to meet demands of the fleet. Additionally, this off-line analysis will provide a proactive view of potential constraints in the system.

The model can also test different levels of WIP (work in process) to determine what level produces the best engine turntime (Buxton and Gatland, 1995). The facility loading capability actually tests to determine the proper level of WIP. Too much WIP increases turnaround time through larger queues at each workstation and increased confusion (Buxton and Gatland, 1995). Too little WIP can also increase product turnaround time and delivery time through human factors, such as people holding work to look busy and lost production time at constrained work stations (Buxton and Gatland, 1995).

Insourcing will become more important to airline maintenance in the coming years. With increased competition from no-frills carriers, the larger airlines will be looking for opportunities to cut costs. The smaller carriers not wanting the major capital expense required to establish an engine maintenance facility will be outsourcing their work. The simulation model will provide a valuable tool to ensure this excess work can be scheduled into the current facility with little or no negative affect. Simulation provides the sales department with reliable information (Rogers and Flanagan, 1991). The insourcing of engines or single-entity pieces will be modeled the same as other required work. A comparison of models run with the insourcing versus those without will provide insight into the affects on the system. Multiple scenarios can be tested quickly to determine the proper timing of entry into the system and which engines to load with the insourced work. Additionally, the simulation model will provide an accurate estimate of the completion time for contractual purposes.

9 CONCLUSION

The simulation model was developed to provide a better understanding of the available capacity of the engine maintenance facility versus the current realized capacity. The model clearly demonstrated that the loading of engines into the repair cycle has a great effect on the capacity of the facility. If the engines are loaded incorrectly, the amount of engines produced lessens. Additionally, the turntime increases as well as the variability. Simulation proved to be a valuable tool because it accurately depicted the interactions between different parts and the resources required. In the future, the model can be expanded to include insourcing. The model will provide insight into the timing and amount of insourcing available based on the current engines in repair.

REFERENCES

Buxton, Kenneth V. and Gatland, Robert G. (1995). Simulating the effects of WIP Levels on Customer Satisfaction in a Manufacturing Environment. In: *Proceedings of the 1995 Winter Simulation Conference*, 940 - 944. LaJolla, California: Society for Computer Simulation.

Cobb, Richard (1995, March). "Simulation supports maintenance marketing efforts". *Air Transport Management*, pp. 25 - 32.

Condom, Pierre (1994, May). "Is outsourcing the winning solution?". *Intervavia Business and Technology*, pp. 34 - 36.

Guide, V. Daniel R. (1993). A Simulation of Synchronous Manufacturing at a Naval Aviation Depot. In: *Proceedings of the 1993 Winter Simulation Conference*, 908 - 915. LaJolla, California: Society for Computer Simulation.

Harvey, Garrison R., McElveen, David V., Miyares, Peter H. and Schuppe, Thomas F. (1992). C-141 Depot Maintenance: Using Simulation To Define Resource Requirements. In: *Proceedings of the 1992 Winter Simulation Conference*, 1145 - 1152. LaJolla, California: Society for Computer Simulation.

Law, Averill M. (1990, April). "How to build valid, credible and appropriately detailed simulation models". *Industrial Engineering*, pp. 16-17.

Rogers, Paul and Flanagan, Maureen T. (1991, December). "On-line simulation for real-time scheduling of manufacturing systems". *Industrial Engineering*, pp. 37 - 40.

AUTHOR BIOGRAPHIES

ROBERT GATLAND is a Specialist in Corporate Planning for Delta Air Lines, Inc. in Atlanta, Georgia. Robert earned a Bachelor's degree in Industrial Engineering from Georgia Tech and is currently working on a Master of Science in Management degree specializing in Production and Operations Management at Georgia State University. At Delta Robert is responsible for development and review of business plans and process design projects. Robert's work interests include simulation, operations management, production scheduling, cellular manufacturing, and process design.

ERIC YANG is a Specialist in Decision Technologies for Delta Air Lines, Inc. in Atlanta, Georgia. Eric earned his Master's and Ph.D. from Purdue University. At Delta Eric is responsible for development of simulation projects and other special projects. Eric's work interests include simulation and production scheduling.

KENNETH BUXTON is a Research Engineer for the Georgia Manufacturing Extension Alliance (GMEA) where he consults with Georgia business and industry. Prior to employment with GMEA, he worked four years as a consultant with the Georgia Productivity and Quality Center and three years in Engineering at Delta Air Lines, Inc. Ken holds a bachelor's degree in industrial engineering from the Georgia Institute of Technology and an MBA in operations management from Georgia State University. His current interests include computer simulation, lean manufacturing, theory of constraints, ISO 9000, and strategic planning.