SIMULATION FOR PLANNING AND SCHEDULING

Russell M. Tahmoush
John Price Rudy
Raytheon Company
Andover, Massachusetts

ABSTRACT

This paper discusses a simulation approach to production planning and scheduling utilized at Raytheon's Missile Systems Plant in Andover, Mass.* A GPSS model and computer program has been developed to simulate the production of some 35,000 printed wiring assemblies. There are a number of different types of assemblies, each of which follows a fixed sequence of operations performed at different work stations. This situation results in a large-scale queuing problem which must be solved by a dynamic technique like simulation. The output from the simulation is currently being used to determine lot size, manpower requirements and production schedules.

Introduction

As a result of a new contract there were to be significant manufacturing changes in the Printed Wiring Assembly (PWA) area where all circuit boards are assembled. These circuit boards are used in defensive missiles and related ground support equipment which have specific end-item (missile or radar unit) delivery dates. Certain decisions regarding manpower and capital equipment had been provisionally made but the scope of the workload made it impossible to mathematically determine the exact requirements for this area. Specifically:

Is the capital equipment sufficient?
How many workers are required and when should they be hired?
What lot size can be run while still maintaining schedule?
When can we expect to deliver the completed assemblies to the next using area?

Also of concern were worker and machine utilization as they are affected by the above conditions. Basically, the logistics problem revolved around the small quantities of hundreds of different types of assemblies.

The decision to apply computer simulation techniques to these problems was partially by default and partially spontaneous. It was quickly learned that the extent of the interactions between work centers precluded our obtaining reliable results using static analysis. When the simulation study was begun it was considered an experiment and that any results would be a bonus. Manufacturing in the area was scheduled to begin within several weeks and so functional managers were made cognizant of the fact that the simulation might not be completed in time to affect decisions for the current contract. However, it soon became apparent that management was relying heavily on the simulation to provide meaningful data to be used as the basis of manufacturing decisions.

Data Collection

One of the first steps was to learn exactly how the circuit boards were to be assembled. The results of this study were six (6) flow charts that describe the sequence of operations pertaining to each of the six basic types of printed wiring assemblies. (Figure 1) The standard work time for each operation was then provided by the Labor Standards Section. This information was reduced to the production time, in minutes, that each board requires at each operation or station. To make the standard hours more realistic, learning curves were developed to account for initial slowness and error until the operator becomes familiar with the equipment and the process.

In order to construct the model, the lot sizes and start-up dates for each major assembly had to be known. This information was obtained from the Production Control Department. Other inputs, such as the capacity of each work station, were readily available from a station loading computer program.

Model Construction

The most important part of simulation is the construction of a model that satisfactorily represents the system under study. The analyst must strive to make the model as realistic as possible without making it so complex as to seriously reduce the probability of its completion.

The PWA model consists of a large block diagram and a set of operating rules. The model generates "imaginary" printed wiring boards that require a certain amount of "servicing" at each work center. Boards are "generated", or released, into the model according to the predetermined lot size and start dates. In the case of a missile board it is split into three parts (representing the

*Certain of the data used in the simulation is classified and thus fictitious numbers are occasionally used.

board itself and two types of components) and simultaneously directed to the first three operations: pre-mechanical, automatic sequencing, and manual preparation (Figure 1). This is done because these three operations can be performed concurrently. These three parts will be joined later in the manufacturing process (and the simulation) when the components are inserted into the printed circuit boards. The pre-mechanical station has the capacity to engage nine boards at a time. If there are, for example, only six boards at the station and a seventh arrives for servicing, it can enter immediately because there is idle capacity (Figure 2a). However, if there are nine boards already in process at pre-mechanical, any new board arriving must wait in a queue (Figure 2b). The queue is serviced on a first in, first out basis unless the analyst assigns priorities to the boards. The highest priority was assigned to the first lot of boards and lower priorities to later lots. As soon as a board is finished with the pre-mechanical phase of production it leaves that work center freeing a machine and worker. The first board in line can then leave the queue and enter the work center.

Each board is "delayed" at the station a length of time computed by multiplying the standard minutes by the appropriate K-factor from the assigned learning curve. The K-factor, a function of the number of units completed, decreases as the operator gains proficiency in his job. When this amount of operation time has passed the board is released and directed to the next work center where it either joins a queue or gets serviced immediately depending on both the number of boards already at the station and on worker availability. As each board completes final inspection, it is tabulated and exists from the model. The computer thus simulates production of printed circuit boards entirely within its memory without any need for physical boards or equipment. Throughout the simulation the computer is compiling statistics that will be printed out at the end of the run. The computer operates at a very high speed so that it can simulate one month's production in five to ten minutes.

Computer Programming & Debugging

The analyst has a choice of computer languages for his simulation. He can select an all-purpose scientific language like FORTRAN or a language developed strictly for simulation like General Purpose Simulation System (GPSS) or SIMSCRIPT. GPSS was selected because of manpower and time restrictions and because of our previous experience with it. FORTRAN would have been more flexible but would have required a very large effort for a model of this magnitude. Previous experience has convinced us that a FORTRAN program of this complexity would be far more difficult to construct and far more difficult to change. GPSS is ideally suited to a simulation of this sort. However, we were constrained by both the large core requirements of GPSS and by the relatively slow running speed. Because GPSS is a general purpose program it provides many features and accumulates many statistics not needed in all applications. Additionally, the availability of these extra features tends to encourage the user to utilize them thus expanding his program far beyond what might have been done in FORTRAN.

The program for the model consists of 600 cards, or GPSS statements, and was first written in two man weeks. Most of this two week period was spent in flow charting the system and gathering the statistics. The program was rapidly debugged and was running smoothly at the third run. An IBM 360 Model 65 with half a million bytes of storage was used though only 300K were available for the program.

Since some 35,000 circuit assemblies had to be assembled for the contract it was apparent that due both to core constraints (approximately proportional to the maximum number of transactions in the system at any one time) and running time limitations (a function of the absolute number of transactions, the number of transactions "alive" at any one time and the number and kind of blocks through which each transaction must go) we could not simulate each assembly individually. By grouping the boards into lots of fifteen it was possible to fit into the available core and to run the simulation in a reasonable amount of time.

Whenever the quota of boards required to complete each end-item is finished the time, in manufacturing minutes, is printed out. Also, periodic print-outs of facility utilization, worker efficiency and queue lengths are made.

In the first run it was assumed that each worker is fixed to an assigned work station. Under these rules he would not leave his station even if there were no work at his station while other stations had backlogs. This produced results that showed very low worker efficiency during the start up period when the volume of work was low.

After discussing this problem with the area foreman, it became evident that workers should be capable of performing more than one task. It was decided to divide the workers into two major job classifications. A worker in classification 1 could perform most of the jobs involving large pieces of capital equipment while those in classification 2 would be trained to perform all the manual operations. This situation was modelled in GPSS by having two SAVEVALUE blocks (counters) maintain a running total of the number of workers available in each classification. Whenever a unit of work attempted to enter a work station the model had to certify that a worker in the appropriate class was available and that there was sufficient capacity at the work station to accept the new unit of work. If both these conditions were satisfied, the work was admitted to the station and the number of available (idle) workers in the classification was decreased by one. Whenever a unit of work left a station the SAVEVALUE was increased by one, since the worker was now available for the next piece of work. By using the GPSS TABLE option and periodically tabulating the values of the two SAVEVALUES a statistical print-out of worker utilization in both job classifications could be obtained.

This change resulted in vast improvements in work force utilization, much to the delight of management and the analysts. To maintain a utilization around 95% the simulation was started with a small number of workers. The number of workers was increased as the work filtered through the system. The peak work force was maintained for several manufacturing months until "lay offs" were necessary as the end of the contract approached and the work load diminished.

Another section of the model underwent major revision after the results of the first run were analyzed. Since a separate learning curve had already been developed for each end-item, it was suggested that these curves be used in the simulation. However, erratic equipment utilization in the first run indicated that because of these learning curves the model was not behaving realistically. Since circuit boards for all end-items are similar and are manufactured on the same equipment over the span of the contract, one would expect the same K-factor to apply to all boards at any one time on a given machine regardless of end-item. With the help of the foreman and Industrial Engineers, a learning curve was defined for each operation in PWA. The structure of the model had to be altered considerably to accept this new type of curve.

Validation & Analysis

The output of the computer simulation was validated to certify that the model was working as planned. A typical output from an early run (Figure 3) shows days 71-80 of the model. For instance at the Ragen Insertion work center, which contains five numerically controlled machines used for component insertion, an average of 1.96 machines were busy during the ten day period. This yields a utilization of 39% for the five Ragens. The simulation indicates that 1668 boards had passed through the Ragen station at an average rate of 5.61 minutes per board. At the time of the print-out (day 80) there were 3 boards in the work station, while the maximum number of boards worked on simultaneously was five, the station capacity. In manpower planning, two (2) girls should be made available for the Ragen station during these 10 M-days since the average workload is predicted to be 1.96 boards. Three Ragens will be idle for most of these ten days, although at some time the workload is sufficient to utilize all five machines. At those times some workers from another work station can be temporarily transferred to the Ragen work center. Figure 3 also indicates that the automatic sequencing machine, the flat pack assembly center, and the hand solder center will operate at full capacity during these ten days.

Statistics are also available on queue size and delay times. For each operation, the maximum, average, and current length of the queue are calculated. In addition, the average time a board has had to wait for a facility is printed out. These figures are useful in pinpointing bottlenecks in the production line.

Model Revision & Implementation

From the volumes of figures printed out, it was possible to develop a week by week manpower requirement schedule and an overall production schedule which indicated when each end-item will be completed (Figure 4). If the schedule or the efficiency proved to be unsatisfactory, the start dates of the end-items, the lot sizes, and/or the number of workers in the model would be revised and another run made. Revisions and new runs were made until an acceptable schedule was developed. This is a heuristic approach to problem solving. The analyst makes changes in the program and analyzes the results to determine if these changes result in improvements. When he is finally satisfied with the answers, he has his solution, which may not be the optimum but which is satisfactory to the analyst and user groups. If the analyst had constructed an accurate model, this solution should prove to be more realistic then solutions developed by other methods. Simulation does not select the best solution, it just tells you what will happen under the conditions that you specify in the model.

Once the model yields acceptable results, a report can be published with recommendations for manpower requirements, scheduling, capital equipment requirements and bottleneck prevention.

Feedback

As work on the contract continues data for final validation as well as input changes for new runs are being received. It has been found that one of the biggest problems in simulating the production of a small number of large complex units is that process or product changes are occasionally made during the initial phases of production. These changes cause further deviations from supposedly fixed and uniform schedules and lot sizes. Part shortages create another problem. Certain parts are not available when their assembly is scheduled by the simulation, which leads to invalid simulation results. The resultant deviations from planned schedules reduce the effectiveness of the simulation considerably and make validation difficult. However, attempts have been made to keep the model reasonably updated. Whenever major schedule changes are introduced, a new run is made for manpower planning.

The initial series of runs indicated that with a peak of thirty direct laborers an acceptable balance between worker utilization, equipment utilization, and manufacturing time could be achieved. Since production started slightly behind schedule, overtime has been required to make up for this lost time. With thirty people working in the area approximately one-fifth of the way through the contract, the scheduled amount of work is being completed during the first shift, but the foreman must occasionally schedule overtime.

Conclusion

At this point, some readers are probably curious about the cost of our simulation. The cost of a simulation varies widely with the size and type of the system to be simulated. The PWA simulation is fairly small. A run to simulate the recent contract in PWA covering approximately seven real-time months took, on the average, 45 minutes on the IBM 360/65 (280K Core Region under MVT). The analysis required two to three man months of work, principally because the data was in raw form and changed frequently as the start of production approached. Several complete runs have been made to date.

What did we get for our money? Probably the main benefit from the simulation has no relation to PWA or even production. Production simulation is a new technique of which very few managers and engineers are aware. We therefore consider our prime benefit to be educational; not only has the capability to simulate most situations in the plant been developed but also a lot of people have been made aware of the potential of simulation and have become agreeable to its use. Of course, the PWA simulation will be of some practical value. It was begun too late to influence the purchase of capital equipment, but it is being used to help determine lot size, production schedules, and manpower requirements. Furthermore, we anticipate that simulation will be an integral part of planning for the coming contract. One immediate effect is that several other smaller simulations already have been utilized in the plant.

In the future Raytheon may use simulation in combination with other techniques to schedule the entire production activity of a plant. Some companies are now using simulation schedulers to load and sequence work in a job shop.[1] With such a system daily reports can indicate, for each work station, all the work orders at the station at the beginning of the day, and the work orders that are expected to arrive during the course of the day. The system will sequence these work orders so as to minimize the number of late deliveries and optimize machine utilization. One company in California reports the following benefits from a simulation scheduler.

A 10% increase in orders completed by their scheduled due date.

A one week reduction in average order-cycle time.

A 60% decrease in expediting effort. With such impressive results as these there is little doubt that simulation will eventually eliminate much of the guess work in production scheduling.

References

1. Buffa, E.S., Readings in Production and Operations' Management, Wiley, New York, 1966.

2. GPSS Users Manual H20-032600, IBM Technical Publications Dept., White Plains, N. Y. 1967.

Biography

Russell M. Tahmoush is currently an Operations Research Engineer in the Industrial Engineering Department of Raytheon Company's Missile Systems Division. He received a B.A. in Applied Physics from Harvard College in 1966 and an M.B.A. in Production Management and Operations Research from Columbia University in 1968.

Mr. Rudy received his B.S. from Massachusetts Institute of Technology in 1967 and his M.S. from M.I.T. in 1968. He is an Operations Research Engineer in the Industrial Engineering Department at the Raytheon Company. Primary interests are in simulation, computerized scheduling and M.I.S. He is an instructor at Fitchburg State College of various computer courses. He is a member of T.I.M.S.

FLOW CHART

MISSILE PRINTED WIRING ASSY.

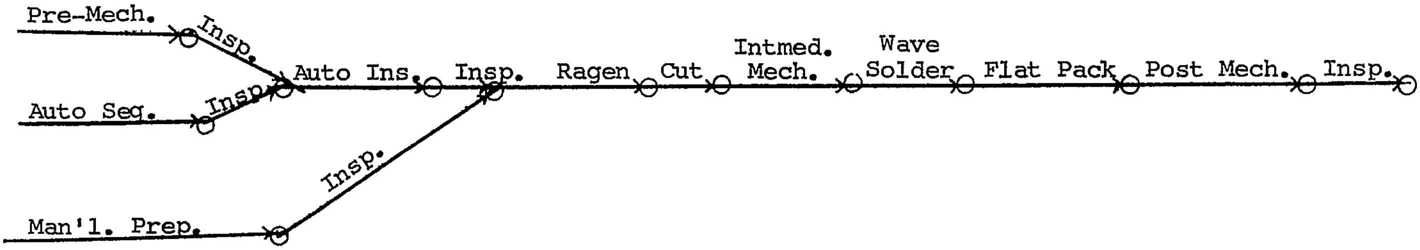

FIGURE 1

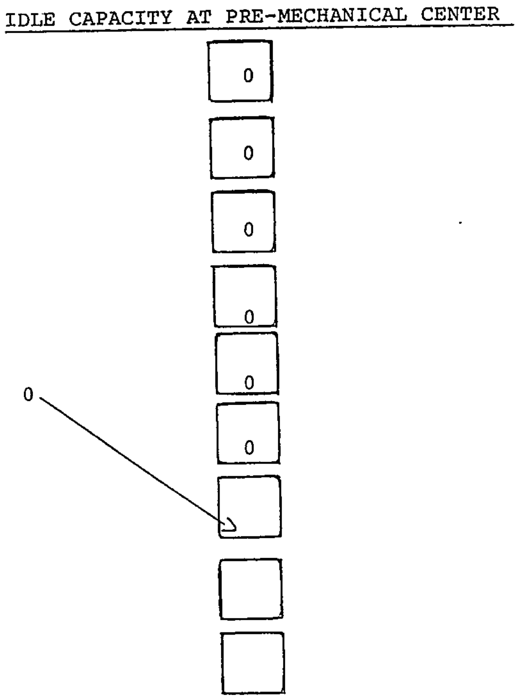

When a board arrives for servicing at the work center, it enters immediately if there is a free machine and worker.

Figure 2A

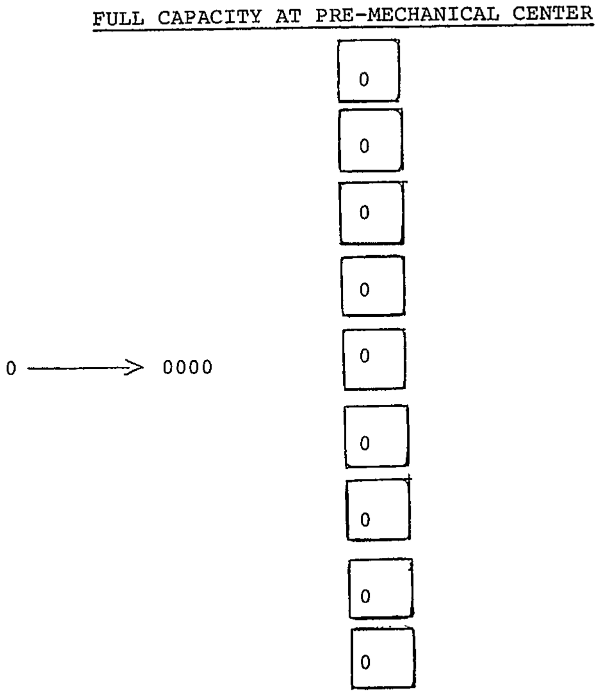

In this case, since all nine work stations are occupied, a board just arriving must join the queue of boards waiting to get service.

Figure 2B

SAMPLE OUTPUT FOR MANUFACTURING DAYS 71-80

Operation	Capacity	Average Content	Average Utilization	No. of Bds. Worked On	Avg. Oper. Time per Bd.	Current Contents	Max. Contents
Pre-Mechanical	9	2.28	25%	2539	4.31	7	5
Auto. Seq.	1	1.00	100	1708	2.81	1	1
Manual Prep.	14	5.62	40	2448	10.98	6	10
Auto. Insertion	1	.57	57	1503	1.82	0	1
Ragen Insertion	5	1.96	39	1668	5.61	3	5
Wave Solder	1	.46	46	1104	2.00	0	2
Flat Pack	8	8.00	100	911	42.12	4	8
Post-Mechanical	11	1.63	15	431	18.15	6	3
Intermed-Mech.	5	.62	10	353	6.80	1	5
Cut	3	1.56	52	1437	5.21	2	3

FIGURE 3

NOTE: Numbers have been changed to protect classified information.

	RADAR A	RADAR B	RADAR C	RADAR D	RADAR E	[illegible]
Lot 1						
Lot Size	5	5	5	5	6	20
Start Date	0	0	15	15	30	53
Finish Date	68.1	44.3	75.4	51.1	58.8	101.9
Production Time	68.1	44.3	60.4	36.1	28.8	48.9
Lot 2						
Lot Size	6	6	6	6	7	20
Start Date	50	50	65	65	80	73
Finish Date	116.8	88.4	124.7	102.3	117.1	113.2
Production Time	66.8	38.4	59.7	37.3	37.1	40.2
Lot 3						
Lot Size						20
Start Date						93
Finish Date						123.7
Production Time						30.7
Lot 4						
Lot Size						20
Start Date						113
Finish Date						133.7
Production Time						20.7
Lot 5						
Lot Size						20
Start Date						133
Finish Date						153.8
Production Time						20.8

Start Date and Finish Date are in M-Days from the start of FY69 Production. Production time is the number of production days the lot required in the simulation.

FIGURE 4