

Hewlett-Packard Uses Operations Research to Improve the Design of a Printer Production Line

Mitchell Burman

*Analytics, Inc.
101 Rogers Street, Suite 216
Cambridge, Massachusetts 02142*

Stanley B. Gershwin

*Massachusetts Institute of Technology
77 Massachusetts Avenue,
Cambridge, Massachusetts 02139*

Curtis Suyematsu

*Hewlett-Packard Company
PO Box 8906
Vancouver, Washington 98668-8906*

As Hewlett-Packard Corporation installed a system for manufacturing ink-jet printers in Vancouver, Washington, in 1993, it realized that the system would not be fast enough or reliable enough to meet its production goals. At the time, the market for ink-jet printers was exploding, and any incremental printer shipments would translate directly into market share and revenue gains. The company undertook a simulation project to develop recommendations for design changes to improve the system performance but concluded that that project would take too long to be useful. MIT researchers used analytical methods to predict capacity and to determine the sizes and locations of buffers that would increase capacity at the cost of a minor increase in inventory. HP's implementation of this work yielded incremental revenues of about \$280 million in printer sales and additional revenues from ancillary products, replacement ink-jet cartridges, media, and related items. Productivity increased about 50 percent, making the assembly of the print engine cost competitive. Finally, HP developed a method of creating rapid and effective system designs in the future.

The Massachusetts Institute of Technology and Hewlett-Packard Corporation collaborated on a factory improvement project. This work (1) enabled incremental revenues of approximately \$280 million in printer shipments and additional revenues due to ancillary products, replacement ink-jet cartridges, media, and so forth, (2) increased labor productivity by up to 50 percent, making the assembly of print engines cost competitive, and (3) gave HP a method for designing systems rapidly and effectively. The last benefit is potentially the greatest in the long term. It could improve predictability, robustness, and system implementation.

This project has helped to establish the usefulness of these analytical methods. The technology is described in MIT courses and the OR research literature, has been commercialized through a spin-off business (Analytics, Inc.), and is finding applications in the general manufacturing environment.

Reasons for Success

This project was particularly successful because of its timing. During this phase of the ink-jet-technology life cycle, demand for ink-jet printers greatly exceeded the supply from HP and all of its major competitors. Any incremental units produced during this time period resulted in incremental revenue and market share. The technology used in this project helped the Vancouver Division (VCD) of HP to greatly lower production costs and, most important, to increase printer shipments at a time when market share was up for grabs.

A key factor in the success of the project was the university participants' ability to innovate and to adapt technology on a

tight project schedule. The best time to change a design is always early in a project. Later redesign often requires more resources and expense than can be supported. In this case, the project team adapted the tools and performed the analysis early enough so that HP could incorporate the improvements into the system development.

The technology was necessary but not

The efficiency of a machine is the ratio of the working time to the total time available

sufficient. Most important was the willingness of HP managers to structure, sponsor, and then implement design recommendations based on this unproven technology. They took a risk when pressure was high to meet an aggressive product-ramp schedule. They had to manage changes in the hardware, information, and people systems concurrently to realize the business benefits.

Technology

The technology consists of a set of algorithms for analyzing and designing production systems (appendix). These algorithms calculate some of the same performance measures as simulation. However, this technology is very easy to use and very fast: it reduces the multi-month development time and the multi-hour run time of simulations to minutes or less. This is its major advantage over simulation. It enabled us to recommend design changes within the development window of opportunity.

The technology was successful because it allowed the joint HP/MIT design team

to evaluate many designs very quickly. This empowered them by providing them with the flexibility to experiment and test their intuition about system behavior. It reduced the time to market, and it led to a more robust and optimized design.

Business Need

HP is a multinational manufacturer of electronic equipment with 1996 sales exceeding \$38 billion. The Vancouver Division (VCD) is part of HP’s Computer Products Organization (the CPO group), which includes DeskJet products, LaserJet products, and computer systems. VCD is one of two divisions in Vancouver, Washington, and is one of three manufacturing hubs for the DeskJet series printers. Ink-jet printers are a multibillion dollar industry worldwide. This market, as well as the competition for it, has grown rapidly since HP introduced the DeskJet product a decade ago.

As the ink-jet printer market grew during the early 1990s, VCD faced the following conflicting objectives:

- (1) upholding the HP reputation for quality and service;
- (2) meeting the increasing demand for printers and improving HP’s market share position;
- (3) achieving its targets for profit and revenue growth; and
- (4) sustaining the “HP way” of manage-

ment, which includes stable employment. HP set a target of 300,000 printers per month for this product line at the Vancouver site. To maintain stable employment, HP constrained the size of the workforce. As a result, capacity was limited to approximately 200,000 printers per month using the existing manual methods. HP needed new methods of production. It decided to improve productivity through automation, considering it the best method of meeting the quality, production, cost, and management constraints. HP decided to invest approximately \$25 million in a new automated system for assembling the printer mechanism that would bring output to 300,000 per month while satisfying the other constraints. This automated system was code-named Eclipse.

The Assembly System

Raw materials enter the factory (Figure 1) through an automated material-handling system that includes an automated storage and retrieval system (AS/RS). The AS/RS routes materials to one of two mechanism-assembly systems (Eclipse). Each Eclipse system is structured in a modular fashion with parts feeding into subassembly modules and subassemblies feeding onto the main assembly system. From the Eclipse system, the products go to the mechanism buffer and then to final assembly. The last step in final assembly is testing. After test-




Figure 1: Raw material enters the factory through an automated material handling system and is then routed to the mechanism assembly area. Assembled mechanisms are temporarily stored in a buffer. From the buffer they go to final assembly and then packaging and shipping.




Figure 2: In the original system design, the base is assembled in the upper left subassembly cell (black rectangle). It is attached to a pallet and it moves clockwise on the main loop as various operations are performed. The first subassembly is added at the upper right. After additional operations take place and further subassemblies are added, the completed print mechanism is separated from the pallet. The pallet stays in the system and the completed assemblies are moved downstream. The main loop contains 30 automated work stations. Each subassembly cell does approximately as much work as four main loop stations. Essentially no in-process inventory space was designed in the system.

ing, the printers are automatically routed to packaging and shipping.

In the original Eclipse design (Figure 2), a central automated pallet conveyor was the primary method to move work from station to station. Pallets received an empty product base from an off-line subassembly station. Subsequently, this base traveled on the pallet clockwise around the conveyor, receiving processing at each station along the way. Every few stations, another major subassembly was added to the mechanism. This continued until the mechanism was completed and stored in a separate buffer for completed mechanisms to await final assembly into a finished printer. This would free up the pallet to immediately receive another product base

and begin the assembly of the next mechanism. There was essentially no in-process inventory space within the Eclipse systems themselves.

The *efficiency* of a machine or manufacturing system is the ratio of the working time to the total time available. Alternatively, it is the ratio of the actual production during a time period to what the production would have been if none of the machines ever failed. The efficiency of a system is less than that of its least efficient machine because of the interactions among machines. Because repairs and failures are random, production volume is random. Estimates of efficiency predict the average ratio over long periods. The only forms of down-time that we consider here are ma-

chine failures and idleness due to machine interactions. Some other interruptions, such as worker absence, can be treated similarly. Others, such as setups (in multi-product systems), must be treated differently.

Machine interactions take the form of starvation and blocking. If there is no buffer between a pair of machines and the

Machines can fail only while they are working.

upstream machine fails, the down-stream machine is immediately starved and forced to be idle because it has nothing to work on. If the downstream machine fails, the upstream machine is idle because it is blocked.

A manufacturing system may have none, one, or many buffers. If a production line has buffers and a machine fails, the next buffer downstream loses material while the machine downstream from it continues to operate. If this condition persists long enough, the buffer becomes empty and that machine is starved. The larger the buffer, the longer the time before it is starved. Conversely, the next buffer upstream of a failed machine gains material while the previous upstream machine is still working. If this goes on long enough, the buffer fills up and the upstream machine is blocked. Thus, buffers defer idleness and thereby increase production rate. Larger buffers more effectively decouple machines because they absorb larger disruptions; but they do so at the cost of increased inventory.

Each machine on the main Eclipse line was designed to have a constant cycle

time of nine seconds, or an uptime production rate of 400 units per hour. For both systems, this comes to a total of 800 units per hour. We thought it reasonable to assume that the actual capacity would exceed 540 units per hour and that the plant would achieve this if the main line assembly machines had efficiencies of 99 percent and yields of 99.5 percent and if the subassembly systems had seven-second cycle times and better efficiencies. Based on approximately 685 hours of available production time per month, we estimated that the system could produce 369,900 units per month. This exceeded the target of 300,000 units per month.

History of the Project

HP sponsored a simulation project to either solidify confidence in the system design or provide specific recommendations for improvements. However, the scope of the simulation programming was much greater than the vendor had anticipated. It did not expect results until well after the opportunity to affect design changes had passed.

At this point, the HP development engineers and managers sought help from academia. They chose MIT because it had applicable analytic techniques and the ability to review the design and to propose changes to meet the objectives.

Buzacott-Model Estimate of Eclipse Capacity

Mitchell Burman, an MIT PhD candidate, worked on this project. He based his first estimate of the efficiency of the Eclipse system on Buzacott's [1967] zero-buffer formula (appendix). He described his analysis in an HP report. It showed that the system could just barely produce

the targeted 300,000 units per month if it met the assumptions of isolated cycle times and station efficiencies.

The Buzacott formula is based on a set of assumptions. The machines have equal, constant operation times. The first machine is never starved and the last is never blocked. Machines can fail only while they are working. There are no buffers in the system, so that when one machine fails, all other machines are forced to be idle until it is repaired. The efficiency is calculated as a function of the mean time to fail and the mean time to repair.

The basic design of the Eclipse system was sound, but its performance depended on its meeting some design and parameter assumptions. Some were questionable:

- That individual stations would achieve efficiencies of 99 percent,
- That station yields would be 99.5 percent,
- That stations would achieve constant cycle times of nine seconds, and
- That the system was tightly coupled with little buffer space.

HP collected performance data as soon as two of the subassembly cells had been installed and run as isolated operations. This early data indicated that the isolated machine efficiencies were closer to 97 percent than the needed 99 percent. In addition, station cycle times varied much more than anticipated and sometimes exceeded the nine-second design target.

The Buzacott model predicted that, if all stations exhibited this performance, the capacity would be about 125,000 units per month. In this case, HP would have had to add labor to meet the requirements, nullifying many of the benefits of automation.

HP and MIT formed a team to develop recommendations for correcting the problem while staying on schedule and within the current space allocations, and maintaining current labor levels.

Need for Some Inventory

Since the system was already under construction, the team could not change individual machines without disrupting the system’s development cycle. Consequently, the team could not easily change the efficiencies of the components of the production system. Gershwin [1994] describes the relationship between buffer space and system efficiency (Figure 3): when buffer space is small, small increases in buffer space increase system efficiency dramatically. System efficiency asymptotically approaches the efficiency of the least efficient machine in the system as the buffer space approaches infinity. This is because buffers prevent the variability of each machine’s production from blocking or starving the others.

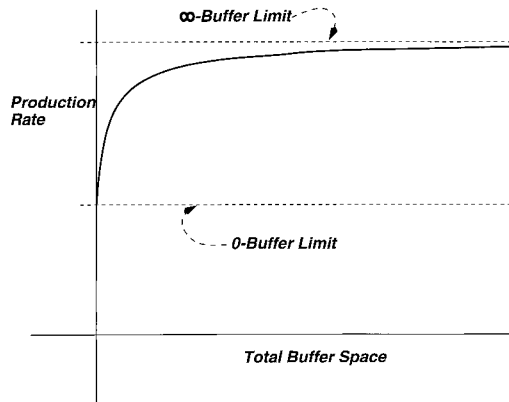


Figure 3: The production rate increases as in-process inventory space increases. This increase is rapid at first and then small. The upper and lower limits are easy to calculate, but the rest of the curve requires the decomposition method.

This relationship suggested that the best way to improve the throughput of the system would be to install limited buffers at strategic points. The goal was to dampen the propagation of the effects of machine failures without expanding inventory space excessively. We needed a method of determining what this curve actually looked like for the Eclipse system and, in turn, how much buffer space we needed to meet the 300,000 unit per month target.

Decomposition

Because of the magnitude and sensitivity of its business investments, HP wanted us to be careful in determining how much buffer space it should install. Burman determined that the quick analytic models found in the flow-line literature [Dallery and Gershwin 1992] were appropriate for the task. He chose the decomposition equations (appendix) developed by Gershwin [1987] and the DDX decomposition algorithm [Dallery, David, and Xie 1988].

Gershwin developed this decomposition method under the assumption that machines have equal, deterministic operation times. He assumed failures and repairs to be random, independent of each other, and independent of the time since any previous event. He assumed buffers to have finite capacities. Except for the assumption of equal operation times, these assumptions are reasonably accurate for the Eclipse system.

We used the method because it was the only one available that could deal with finite buffers and unreliable machines. Overcoming the limitation of equal operation times was one of Burman’s main goals in his PhD work (appendix) and is

one of Analytics’ main software development initiatives. Meanwhile, we accounted for these limitations by using approximations.

These approximations included breaking up the problem into two parts: estimating the performance of the main loop, and determining the sizes of the buffers needed between the subassembly systems and the main loop. In analyzing the main loop, we were able to ignore the subassembly cells because they were faster and more reliable and would therefore only rarely affect production once the buffers were installed between them and the main line. We ignored observed differences in operation times and assumed the common

The model’s estimate of production rates were within 10 percent of the actual observations.

operation time was that of the slowest machine. In addition, we approximated the loop as a line.

We determined the sizes of buffers between the subassembly systems and the main line by treating the material flow into and out of each such buffer as if it were in a two-machine continuous-material transfer line. (The first machine represented the subassembly system and the second represented the main line.) This obviated the need for a long-line decomposition and made it possible to treat the different operation speeds in the different parts of Eclipse. We chose the buffers to be large enough so that the subassembly systems would, in fact, rarely affect main line production.

New System

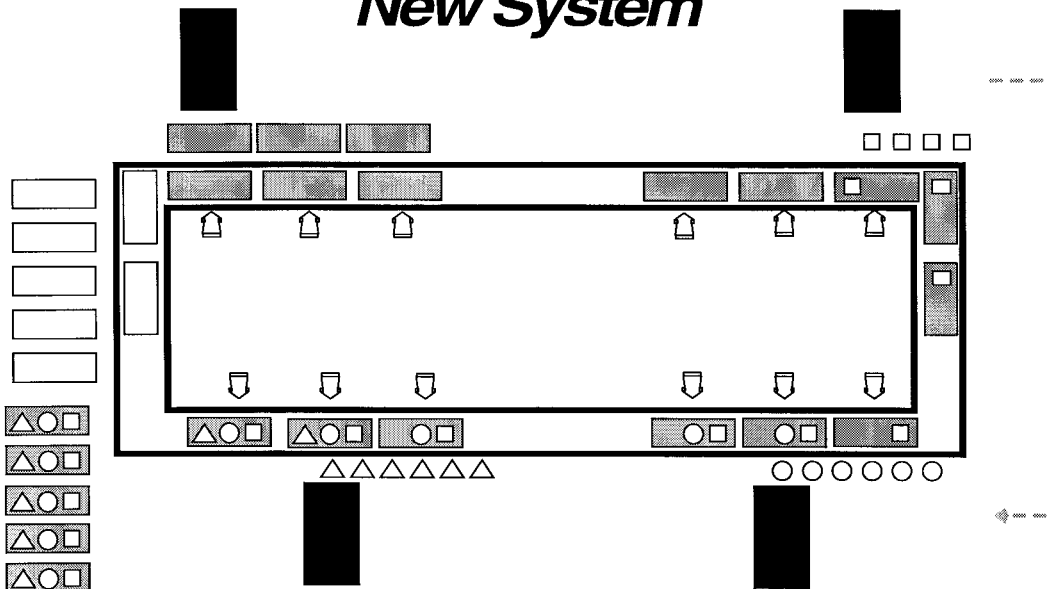


Figure 4: We recommended the addition of (1) an empty pallet buffer, (2) space for in-process inventory between the subassembly systems and the main line, and (3) space on the main line.

Recommendations

We used the decomposition methods to produce a report that demonstrated that the following system improvements (Figure 4) could raise estimated throughput to between 250,000 and 300,000 units per month:

- (1) Adding an empty pallet buffer,
- (2) Adding automated buffers between subassembly cells and the main line, and
- (3) Adding buffers between main line stations.

The system had about 103 locations for pallets and was designed to contain 80 pallets. But this many pallets would lead to congestion and would reduce throughput because the Eclipse is a closed system. When there are much fewer parts than spaces, the machines are frequently starved. But when the number of parts is close to the number of spaces, the ma-

chines are frequently blocked. Increasing the number of spaces (in the form of an empty pallet buffer) increases the production rate because it reduces blocking. We reduced this blockage and starvation by putting a pallet storage area at the end of the system. This storage also served as a buffer to decouple downstream disruptions from upstream part loading.

The system was designed with subassembly systems that ran faster than the main line. However, since there were no buffers between the subassemblies and the main line, when a subassembly cell stopped, the main line stopped. By installing subassembly buffers that could hold 30 minutes of subassembly (approximately 200 units of in-process inventory), HP isolated any disruptions in the subassembly cells from the main line. We expected these buffers to be full most of the time

because the subassembly systems had shorter operation times than the main line. Although this might appear to be a great deal of inventory, 30 minutes was needed to guarantee that subassembly failures (that have a mean duration of several minutes) would rarely interrupt main-line production.

We suggested that HP install three other buffers of approximately 12 units each one quarter, half, and three quarters of the way through the line.

Based on the performance improvements we estimated with the analytic models, HP made these modifications. The total capital costs for the changes to both Eclipse systems were \$1,400,000.

System Improvement Prioritization

During Ramp-Up

By mid-1994, HP had finished the hardware changes to the system. It set new target levels for station efficiencies based on more realistic performance levels. It still had to achieve these efficiencies and sustain the modeled cycle times.

In parallel to the hardware changes, we developed and implemented a more accurate and reliable information system on

The increased throughput resulted in incremental revenue of \$280 million.

the Eclipse system. This system provided the information engineering and operations needed to know precisely how the system was performing as measured by its critical parameters. This would pinpoint the exact location of any deviations and would signal when repairs were adequate. It provided both historical and real-time

data for such items as cycle time, uptime, and downtime, and it was used extensively by engineering and operations in achieving the Eclipse objectives. Using the decomposition model, we performed sensitivity analysis to determine which action would have the greatest impact on system performance.

Simultaneously, Burman was asked to develop extensions to the Gershwin decomposition that

- (1) would handle different machine operation times,
- (2) could handle very long lines,
- (3) would run in seconds,
- (4) had a better user interface,
- (5) had simple data requirements, and
- (6) converged more reliably.

Burman completed this work successfully by March of 1995 and described it in his PhD thesis [Burman 1995].

Validation

Using system data from May and June of 1995, Burman compared actual system performance to the performance predicted by the model he developed [Burman 1995]. The model's estimates of production rates were within 10 percent of the actual observations.

Actual system production was raised to between 250,000 and 300,000 per month with the expected productivity. When called upon, the system has demonstrated the ability to sustain production greater than 300,000 per month.

HP Impact

This technology contributed greatly to HP.

The technology we recommended for system development helped the Vancouver Division to meet its business commit-

ments during this period. The recommendations took into account the real-world constraints of schedule, space, and resources. The approach to improvement

Supply will exceed demand, prices will drop, and the focus will shift to cost control.

had a great value to HP because the recommendations were based on rigorous analytic methods. This increased HP's willingness to take an enormous risk based on the recommendations. The changes resulted in substantial incremental revenue.

HP expects the technology to have an impact on its future business. The potential long-term benefit of this project at HP is the establishment of a methodology for designing manufacturing systems that have improved throughput, robustness, and predictability. The Vancouver Division has leveraged this design approach for next-generation manufacturing systems. It estimates that this method can increase the throughput of all its future manual and automated systems.

This project helped HP to design and improve the Eclipse automated print-engine assembly system. It had an immediate impact because of the timeliness of the solutions. It also helped HP to set a clear management focus for system improvement. It made the issues comprehensible and quantifiable, and it provided a useful discipline and language for discussing and evaluating manufacturing-system designs.

During the period when demand exceeded supply, the increased throughput resulted in incremental revenue of \$280 million. We base this figure on the produc-

tion achieved in excess of the predicted manual assembly capacity of 200,000 per month. HP also realized additional revenues from ancillary products, replacement cartridges, media, and so forth.

Because productivity was increased by up to 50 percent, assembly of the print engine was cost competitive. This is a significant benefit for the future, because the business has left the period of explosive growth. Supply will exceed demand, prices will drop, and the focus will shift from revenue to cost control.

Other Commercial Impacts

The success of this project led Burman to found Analytics, Inc. Analytics provides manufacturing-systems design services that are, in part, based on the principles of the work described here. Analytics has used this methodology in a project for Johnson and Johnson, a project for Boeing, and an unrelated project for Hewlett-Packard in Corvallis, Oregon.

Analytics has continued to improve the technology and has added features including assembly modules (appendix) and a graphical user interface. The speed and flexibility of this technology make it an ideal complement to simulation. By using this approach, one can make robust strategic-design decisions and fine-tune with simulation after solidifying the major system architecture.

Impact on MIT

The immediate impact at MIT was in the PhD research of Mitchell Burman, which he completed in 1995. He based this research on the problems he observed directly at the HP printer factory. Some of these problems could be treated with methodology already in the literature, but

some required new research. His results were used in the master’s thesis of James Schor [1995] and the PhD thesis of Asbjorn Bonvik [1996]. In addition, this project has led to opportunities for other MIT personnel to visit and work with HP.

MIT will use the new methods and software—as well as the HP case—in manufacturing-systems courses, such as “Manufacturing systems analysis” in the mechanical engineering department and “Operations management models and applications” in the Sloan School of Management. This work provides important results that will further research in systems design and operation.

The success of this university-industry interaction demonstrates the benefits that can result, and it should encourage more collaborative efforts of this sort.

Impact on the Operations Research Community

The OR community will benefit from this project’s dramatic financial and other benefits. The methods are easy to use (though challenging to derive) and will therefore add to the community’s credibility. This project can serve as a model for industry-academia interaction.

Academia and the OR community can continue to benefit by applying research to enable such business benefits. By commercializing such technologies, we can realize social and economic benefits.

Acknowledgments

We are grateful for all the support we received from HP personnel. The success of the Eclipse project was made possible by the dedication and hard work of many individuals throughout the HP organization.

APPENDIX—OPERATIONS RESEARCH TECHNOLOGY

Buzacott’s Zero-Buffer Model

Assume a line has k machines that have equal, constant operation times. The first machine is never starved and the last is never blocked. The operation time is chosen to be the time unit. Machines can fail only while they are working. The probability of Machine M_i failing during a time unit when it is operating is p_i , and the probability of Machine M_i being repaired during a time unit when it is down is r_i . We define $F_i = 1/f_i$ and $R_i = 1/r_i$, the mean time to fail and the mean time to repair.

Lines with buffers of size zero have no decoupling between stages. Consequently, as soon as any machine fails, the whole line is forced to wait. Using these assumptions, Buzacott [1967] showed that a good approximation of the efficiency of the line is

$$E = \frac{1}{1 + \sum_{i=1}^k \frac{f_i}{r_i}} = \frac{1}{1 + \sum_{i=1}^k \frac{R_i}{F_i}}$$

In reality, each machine’s work area acts like a buffer of size 1. However, we felt the Buzacott approximation was appropriate because the operation time was small relative to the failure time.

The efficiency of the line is the ratio of the number of parts produced to the number that would have been produced if there were no failures. Here, it is the same as the production rate since it measures the average number of parts produced per time unit.

Buzacott and Shanthikumar [1993] generalize the formula to include systems with different operation times.

Decomposition of Synchronous Model of Line

Figure 5 shows a five-machine, four-buffer production line. For a system in which all the machines had constant,




Figure 5: In this representation of a production line, the squares represent machines and the circles represent buffers. Material moves in the direction indicated from machine to temporary storage buffer to machine. Machines fail at random times and stay down for random lengths of time. The buffers are finite, so disruptions are propagated in both directions in the form of starvation and blockage. Because of the complexity of its behavior, this system can be evaluated only by simulation or by a decomposition approximation.

equal operation times and geometrically distributed repair and failure times, Gershwin [1987] showed that the production rate and average in-process inventories could be approximated by those of a set of two-machine lines (Figure 6).

Each of the two-machine lines is constructed with a buffer that is the same size as that of one of the buffers in the original line. The machines are chosen so that an observer in the buffer, who sees only the

arrival-and-departure processes, sees almost the same processes as a similar observer in the corresponding buffer of the original line. Gershwin [1987] developed the equations that determined the parameters of the machines; Dallery, David, and Xie [1988] provided an efficient algorithm for solving these equations (called the DDX algorithm). Gershwin [1994] describes this.

We used this algorithm in the early stages of the Eclipse project. Its advantage is its speed. A line can be analyzed in less than a second on a personal computer; writing a simulation could take months, and running it could take hours. Such long times make simulation awkward for analysis and design.

This method is an approximation. The actual behavior of material entering and leaving buffers of larger systems is subtly different from the behavior of material in two-machine lines. This difference increases for systems in which mean time to failure or mean time to repair are very different among the machines.

Decomposition of Continuous-Material Model of Line

Because the machines in the Eclipse system do not have the same operation times, we needed a better approximation. However, extending Gershwin [1987] to systems with machines that have different,

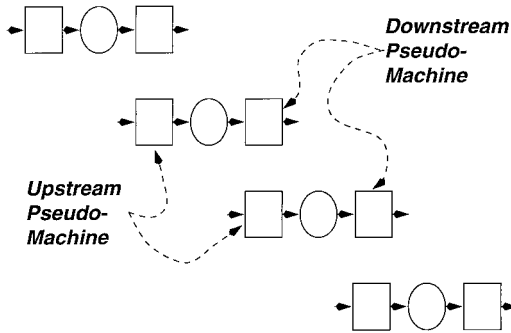


Figure 6: In decomposing a production line, we create two pseudomachines for each buffer in the original line. We form a two-machine line with those machines and a buffer that is the same size as the corresponding buffer in Figure 5. We choose the pseudomachines so that the material-flow behavior in the buffer of each two-machine line is nearly the same as that in the corresponding buffer of the original line.

deterministic processing times appeared to be prohibitively difficult, so Burman [1995] used the continuous material two-machine line of Gershwin and Schick [1980] in a similar decomposition for long lines with continuous material and different processing rates. (The differing processing rates made the decomposition equations considerably more difficult to derive than those of earlier systems.) He also adapted the DDX algorithm for this case (and called the new algorithm the ADDX algorithm). It was also extremely fast.

Decomposition of Continuous-Material Model of Assembly-Disassembly System

Since the Hewlett-Packard production system included assembly, we needed more than the ADDX algorithm for production lines. Consequently, Gershwin and Burman [1997] extended the ADDX algorithm (as well as Gershwin's [1991] earlier assembly-disassembly system work for synchronous systems and Di Mascolo, David, and Dallery's [1991] work for continuous-material systems with equal processing rates) to assembly-disassembly systems with continuous material and different processing rates. Again this algorithm proved to be fast and practical. It has been implemented with a graphical user interface, and Hewlett-Packard is using it as part of its standard methodology for designing new printer production systems.

References

- Bonvik, A. M. 1996, "Performance analysis of manufacturing systems under hybrid control policies," PhD thesis, Massachusetts Institute of Technology, Cambridge, Massachusetts.
- Burman, M. 1995, "New results in flow line analysis," PhD thesis, Massachusetts Institute of Technology, Cambridge, Massachusetts.
- Buzacott, J. A. 1967, "Automatic transfer lines with buffer stocks," *International Journal of Production Research*, Vol. 5, No. 3, pp. 183–200.
- Buzacott, J. A. and Shanthikumar, J. G. 1993, *Stochastic Models of Manufacturing Systems*, Prentice Hall, Englewood Cliffs, New Jersey.
- Dallery, Y.; David, R.; and Xie, X.-L. 1988, "An efficient algorithm for analysis of transfer lines with unreliable machines and finite buffers," *IIE Transactions*, Vol. 20, No. 3 (September), pp. 280–283.
- Dallery, Y. and Gershwin, S. B. 1992, "Manufacturing flow line systems: A review of models and analytical results," *Queueing Systems Theory and Applications*, Special Issue on Queueing Models of Manufacturing Systems, Volume 12, No. 1–2 (December), pp. 3–94.
- Di Mascolo, M.; David, R.; and Dallery, Y. 1991, "Modeling and analysis of assembly systems with unreliable machines and finite buffers," *IIE Transactions*, Vol. 23, No. 4 (December), pp. 315–331.
- Gershwin, S. B. 1987, "An efficient decomposition method for the approximate evaluation of tandem queues with finite storage space and blocking," *Operations Research*, Vol. 35, No. 2 (March–April), pp. 291–305.
- Gershwin, S. B. 1991, "Assembly/disassembly systems: An efficient decomposition algorithm for tree-structured networks," *IIE Transactions*, Vol. 23, No. 4 (December), pp. 302–314.
- Gershwin, S. B. 1994, *Manufacturing Systems Engineering*, Prentice Hall, Englewood Cliffs, New Jersey.
- Gershwin, S. B. and Burman, M. H. 1997, "Decomposition equations and algorithm for A/D systems—Derivation of equations," working paper.
- Gershwin, S. B. and Schick, I. C. 1980, "Continuous model of an unreliable two-stage material flow system with a finite interstage buffer," Massachusetts Institute of Technology Laboratory for Information and Decision Systems Report LIDS-R-1039, September.
- Schor, J. E. 1995, "Efficient algorithms for buffer allocation," MS thesis, Massachusetts Institute of Technology, Cambridge, Massachusetts.