

# PERFORMANCE ANALYSIS TECHNIQUES FOR IC MANUFACTURING LINES

AT&T TECHNICAL JOURNAL

**David Y. Burman, F. Javier Gurrola-Gal, Ardavan Nozari,  
Shashikant Sathaye, and Joseph P. Sitarik**

*David Y. Burman is a supervisor in the Manufacturing Systems Engineering Department at AT&T Engineering Research Center (ERC) in Murray Hill, New Jersey, and **Ardavan Nozari** and **Shashikant Sathaye** are senior members of technical staff in that department. **Joseph P. Sitarik** is a supervisor in the Manufacturing Systems Engineering Department at ERC in Allentown, Pennsylvania, and **F. Javier Gurrola-Gal** is a senior member of technical staff in that department. Mr. Burman studies the performance of manufacturing lines and recommends and develops new operational procedures. He joined AT&T in 1969 and has a Ph.D. in mathematics from New York University. Mr. Gurrola-Gal, who joined AT&T in 1981, identifies planning and manufacturing problems at AT&T factories, and implements solutions (continued on page 57)*

This paper discusses ways that operations research tools and techniques can be used to analyze the performance of integrated circuit manufacturing lines. It focuses on simulation, queueing, and deterministic capacity models and, through case studies, discusses the applicability of these methods. It also discusses the criteria for selecting the appropriate analysis method for a given problem.

## Background

Manufacturing lines for integrated circuits (ICs) by nature are complex and often not amenable to simple analysis techniques. The operations research tools and methods we describe have been used on several IC lines and some circuit-pack assembly and test lines, and have potential application for an even broader range of manufacturing systems.

IC lines are capital intensive, and getting maximum use from that capital has direct effect on an IC's cost. Usually, 30 percent of the cost of modern electronic equipment is in its ICs. Moreover, the manufacturing interval for ICs is often a significant part of the total time needed to manufacture the equipment.

In this environment, any analysis technique that can improve productivity—e.g., by increasing capital use or decreasing manufacturing interval (which decreases the time to fill a customer's order)—is important.

Here, we focus on simulation, deterministic capacity models, and queueing for studying the performance of IC wafer fabrication lines. For several years, these techniques have successfully helped engineers and managers improve the performance of AT&T IC lines. This paper discusses real examples, including the use of simulation to identify an optimal strategy for reducing work-in-process (WIP), and deterministic capacity models to design new lines.

Although these modeling techniques are not new to manufacturing applications, they have received only limited discussion in the literature on IC manufacturing.<sup>1,2</sup> Indeed, because of the complexity of the IC wafer fabrication process, it was not at all apparent that the techniques could be successfully applied. Here, we describe how these

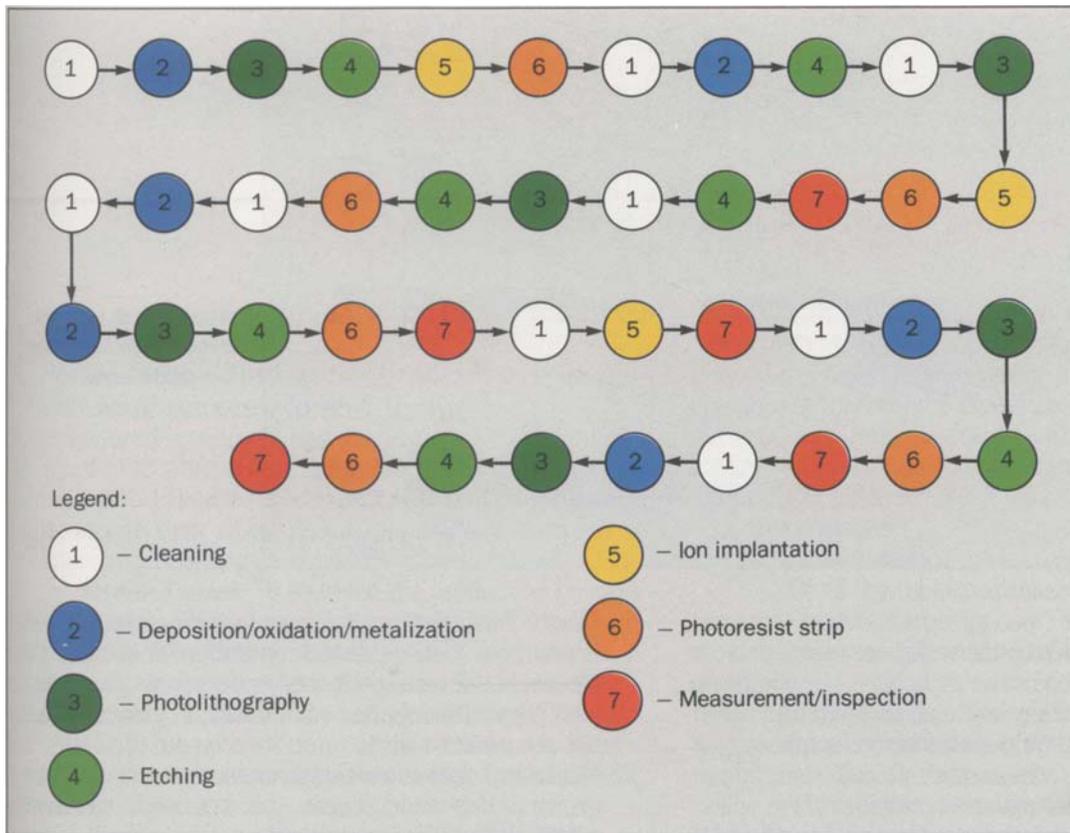


Figure 1. Main steps in an IC process flow.

techniques help answer several questions that arise in studying the performance of complex IC fabrication lines.

### Choosing the Technique

Choosing the appropriate technique for a specific problem naturally involves studying the tradeoffs of the techniques and a specific problem's requirements.

Discrete-event simulation models are developed either as large-scale models of the total production area that focus on answering a wide variety of potential questions, or smaller models that address specific issues. Either requires significant time to develop a model, verify that it does what is intended, validate that it tracks the real system, and analyze the output. Simulation's main advantage is that it allows us to address a wide variety of complex problems that would otherwise defy analysis.

Deterministic capacity models address issues of capacity estimation, are easy to develop, and are quick to run. However, the range of questions that these models can address is limited. For example, estimating the WIP or throughput time is completely beyond their scope.

Queueing models tend to lie between the other two. Efficient models can be developed reasonably fast to estimate a wider range of parameters than the deterministic capacity models. But because they tend to be extremely rough, queueing models must be used with caution.

Later, we will present a series of examples motivated by real applications. They show the applicability of the tools, and the gains that may be realized from using them. To help readers understand the subtleties of the models, first we describe the IC manufacturing process with particular emphasis on wafer fabrication.

### Wafer Fabrication Process

A complex multistage process transforms silicon (or, less typically, gallium arsenide)—in the form of thin, highly polished disks, called wafers—into ICs.

1. *Raw wafer manufacturing.* Crystalline ingots are drawn from molten silicon, sliced into wafers, and polished.
2. *Wafer fabrication.* Many identical ICs are formed on the same wafer using a complex sequence of

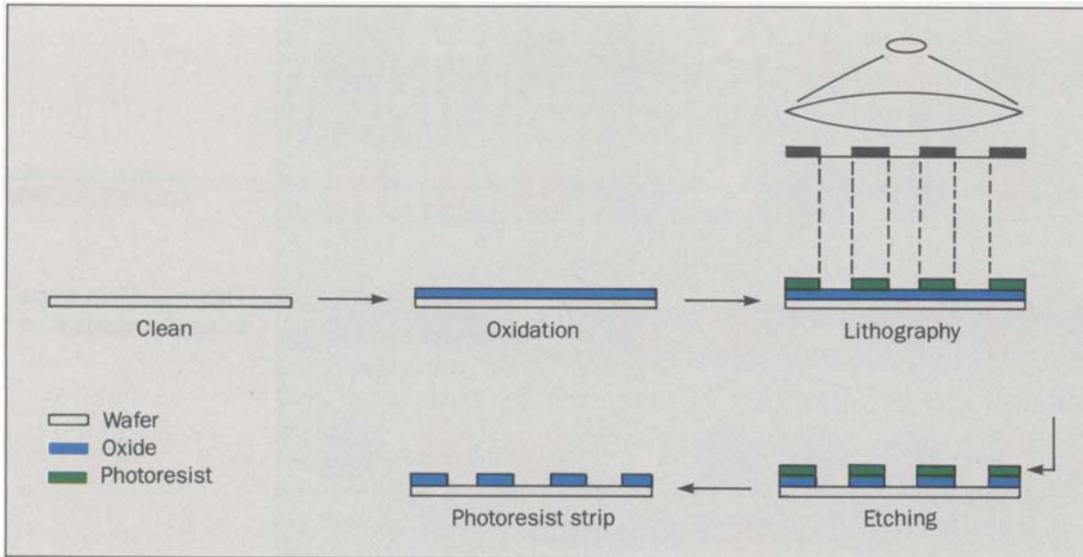


Figure 2. A sample process to fabricate an oxide layer.

processes.

3. *Probe and dice.* The ICs on the wafer are tested and sawed into individual chips.
4. *Package and test.* The chips are encapsulated into plastic or ceramic packages and tested before shipping.

In this paper, we focus on the (silicon) wafer fabrication stage because of its length and complexity. We will describe the wafer fabrication process from an operational point of view and discuss some of the more critical modeling issues raised by this process. (For a detailed analysis of the total process, see Reference 3.)

**Operational Issues.** The manufacture of complex IC circuitry uses photolithography to superimpose patterned layers of different materials on a wafer to produce the final chip. A wafer contains many identical chips, and wafers that correspond to the same product travel in a *lot*. The lot's size is determined by the production line and the type of product. (Usually, it is a multiple of the size of a standard cassette that is often used for loading wafers onto certain machines.)

The recipe for manufacturing the particular product is defined by a *process log*. This log completely describes the process steps (which can be grouped as described below) and the particulars of the machine, its settings, and other pertinent information. The diagram in Figure 1 (derived from a process that S. Sze described<sup>3</sup>) gives some feel for the process flow's complexity.

Figure 2 shows a sample five-step process to form one layer. In general, wafer fabrication can be grouped roughly into seven types of processes:

1. *Clean.* Particulate matter is removed before the layer is produced. This is a batch operation on one to eight cassettes. Generally, cleaned wafers must go to the next step within a prespecified time to prevent additional contamination.
2. *Oxidation, deposition, metalization.* A material layer is grown or deposited. These, too, are batch operations where the batch size varies dramatically (from 10 to 200 wafers) depending on the machine and process. Because of extensive setup times, machines are dedicated to a limited number of operations.
3. *Lithography,* the most complex operation. A photoresist (PR) is deposited on the wafer and the circuit pattern is defined using photography. First, the wafers are coated with the PR and baked. Next, the PR is exposed (on a printer) to ultraviolet light through a mask that contains the desired circuit pattern. This step requires extreme care to guarantee proper alignment with the previous layers and prevent contamination on the mask. Finally, the exposed wafer is developed and baked.
4. *Etching.* To define the circuits, the exposed part of the material layer is etched away. Several types of etchants are used (wet chemical and plasma), and the batch size varies significantly within a type.
5. *Ion implant.* Selected impurities are introduced in a controlled fashion to change the electrical properties of the exposed portion of the layer. Ion implanters require setups that range from minutes (to change settings) to hours (to change impurities) and, depending

---

on the machine, can process cassettes automatically or must be hand loaded.

6. *Photoresist strip*. The remaining photoresist is stripped away using a process similar to etching.
7. *Inspect and measure*. The layer is inspected and measured to identify defects and guide future processing of the wafer. Failed wafers are either scrapped, reworked along with their whole lot (*lot rework*), or stored for future reprocessing as *reconstituted lots*.

**Modeling Issues.** To examine the production process from a modeling point of view, let us divide the process into two parts. The first deals with product flow, which includes not only flow between processes, but also flow within a process (as dictated by scheduling rules). The second deals with the process times at the various machines. In both cases, we consider the deterministic (nonrandom) and random aspects.

The predominant factor that affects flow is the process log (Figure 1) that dictates the order in which various machines are visited. The order varies from product to product and, generally, specifies multiple visits to the same machine. Flow between processes is further complicated by lot reworks that effectively cause a lot to visit the same steps a random number of times.

Within a process, we must determine which lot is to be processed next and on which machine. This affects the yield (for example, the timing of the clean step), the machine's efficiency for large setups, and the output for heavily loaded rooms.

Process times depend largely on the operation being done. For batch operations (e.g., clean), process times generally do not vary with lot size. For other operations (e.g., lithography), process times depend on the number of wafers in the lot. Besides the prescribed process times, it is important to realize that, from a modeling point of view, machine breakdown and operator unavailability add variability to the effective lot process time. Scrapped wafers (e.g., from inspections), which reduce the lot size, also affect certain process times.

Thus, to understand an IC wafer fabrication line's

performance we must appreciate the operational aspects of the production process. These aspects decompose into flow- and process-related issues. Each has a deterministic component, dictated by the process log and process engineering, and a random component that results from many causes, including failed inspections, rework, and machine availability.

### Computer Simulation

Of all the tools available for analyzing the performance of manufacturing lines, simulation offers the most flexibility and allows the modeler to answer the most questions. Its main drawbacks are: Simulations tend to be time-consuming, and analysis of the output requires a degree of sophistication that, on the surface, may not be required of other techniques.

This section discusses the application of discrete-event simulation to IC manufacturing.<sup>1,2</sup>

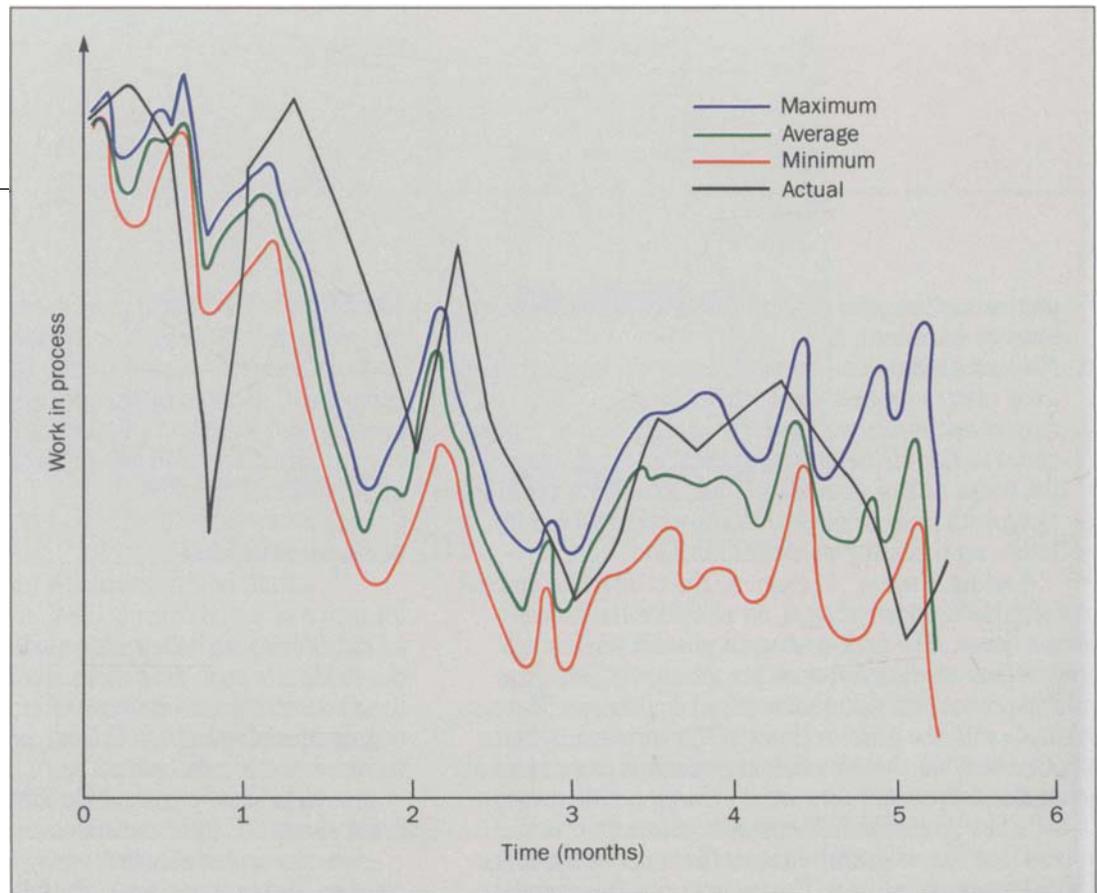
Computer simulation may be divided into two categories. The first are large, flexible models of the entire manufacturing facility that readily address many operational issues. For example, they can estimate capacity, identify bottlenecks (both labor and equipment), and estimate expected WIP and throughput time. The second category includes models developed to address a specific issue, such as a machine's schedule.

Each category may be further classified into two types: *steady-state* (nonterminating), and *transient* (terminating). Because the statistical analysis of the simulation output depends on the type of study, it is important to distinguish between the two.

In a transient problem, natural events define the beginning and end of the problem. An example is estimating the expected average WIP in a clean room over the next six months, given a starting WIP and a production schedule. On the other hand, a steady-state problem is one where no natural events can define the beginning and end. Estimating the expected throughput time, given a constant, daily input rate is an example.

No attempt is made here to discuss the important

**Figure 3. Comparison of simulated and actual work in progress in a clean room for a 5-month period.**



issue of output analysis. (See Reference 4 for an overview of output analysis, and References 5 and 6 for techniques more specialized to our application.)

We discuss the structure of large-scale models and their verification and validation, and present two examples of their use to resolve production issues. To show how simulation is used for answering specific questions that arise from IC manufacturing lines, we also describe a study of the operational issues related to running direct step on wafer (DSW) printers.

**Large-Scale Models.** The value of large-scale models is in their ability to be adapted easily to answer a variety of questions. Several requirements are essential to meet this goal effectively:

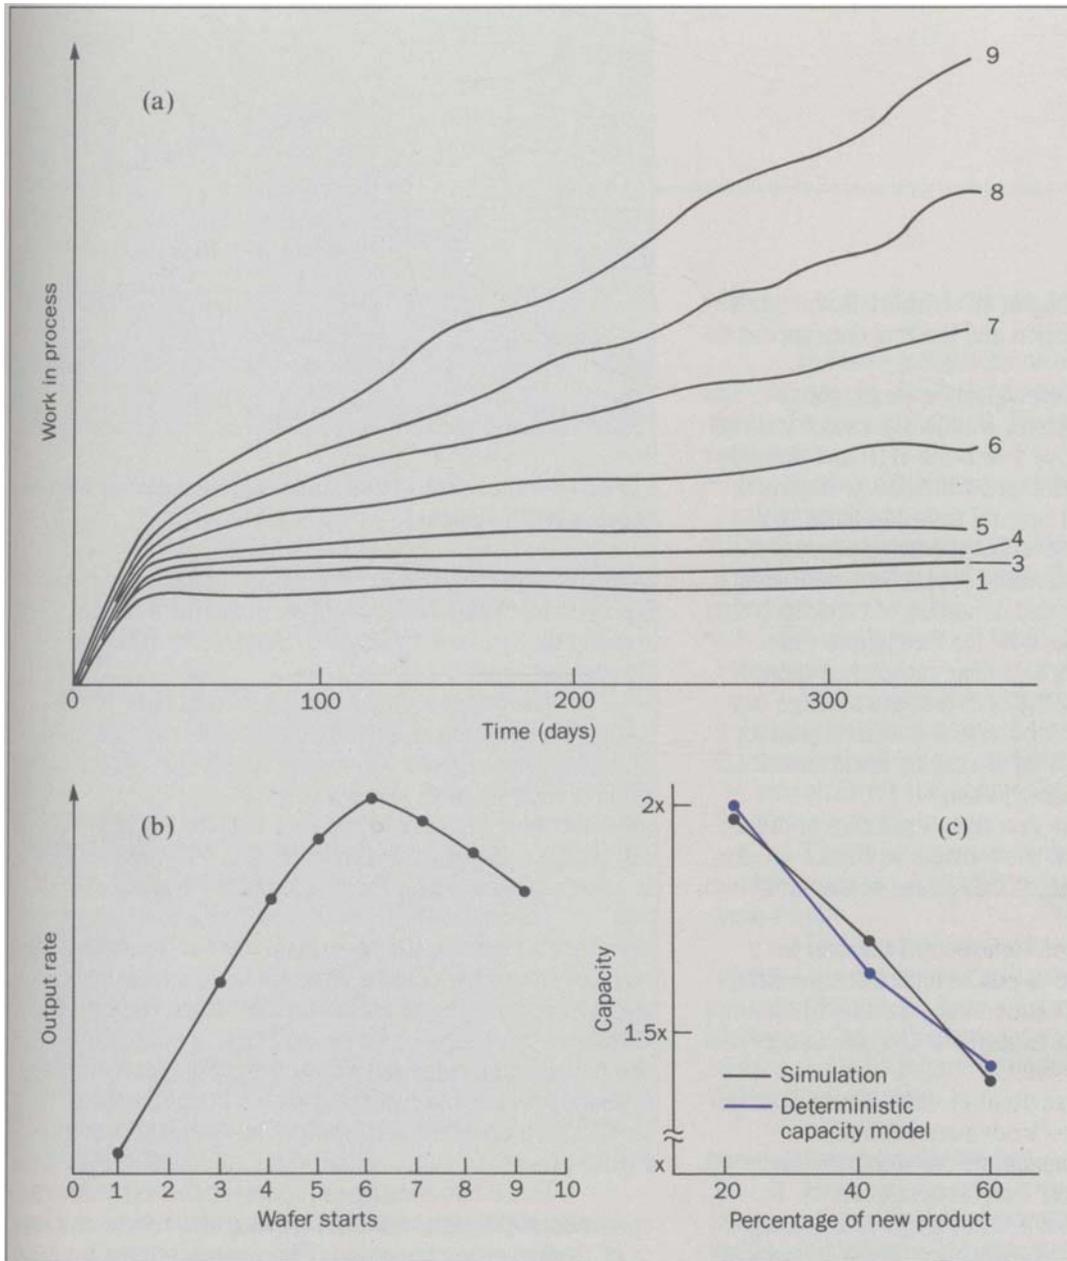
1. *Scope.* Because of the high correlation between processes that complex process flow causes, the models must describe the total production process.
2. *Maintainability.* Changes to the room (such as new process logs, changes in the machines, and new oper-

ating parameters) have to be easy to incorporate.

3. *Verifiability.* It must be easy to verify that the model behaves as intended and validate that it agrees with the real system.<sup>7</sup>
4. *Flexibility.* New studies should require minimal reprogramming of the model.

To meet these requirements, one needs a structured, modularized approach to constructing the large-scale models. It consists of a series of individual modules, each corresponding to a facility group with several similar machines. Entities that correspond to wafer lots are routed through these modules, as prescribed by the product's process log.

This approach has greatly simplified the development of large-scale models and reduced development time. Initially, the modules are extremely simple and consist of only two activities. The first represents the involvement of labor and the equipment, and the second the equipment without the labor. As the model is validated, additional



**Figure 4. Analysis of a clean room's capacity. (a) Work in process versus time at different start levels for a fixed product mix. (b) Steady-state output versus start level for a fixed product. (c) Manufacturing line capacity versus product mix in the line.**

details—such as machine breakdown, operator breaks, and scheduling rules—are included as needed.

**Model validation.** Validation involves several difficult and important philosophical questions<sup>8</sup> that are well beyond the scope of this paper. A common and intuitively appealing approach is to compare the model's response with that of reality under exactly the same conditions.

Figure 3 shows the WIP for a 5-month period in a room and for 10 runs of a simulation that used different random number seeds. For each simulation run, the initial WIP and the daily lot starts were identical to those in the room for the 5-month validation period. During this time, several different products were being produced in the room, and the daily lot starts were not constant.

Despite the high degree of volatility in the room during this time, the simulation and the real data appear to agree.

**Capacity analysis.** Determining a clean room's capacity is not a trivial problem. Intuitively, capacity is the number of wafer starts above which the WIP and throughput time become unstable and grow to infinity. Figure 4a shows WIP versus time for several runs of a simulation, each with a different start level but the same product mix. The graph shows that the capacity is reached about level 6.

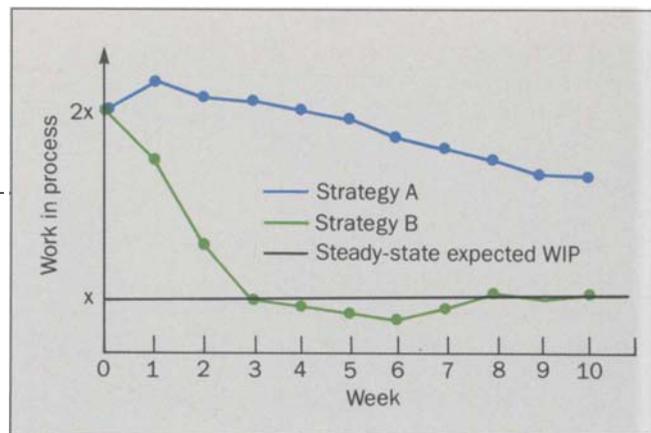
One problem with this definition of capacity is that steady-state behavior of the WIP (or throughput time) becomes more difficult to detect near capacity. A more direct way to calculate capacity is to plot the average output rate as a function of wafer starts (for a fixed product mix). Below capacity, this is an increasing linear function, and above capacity, it is nonincreasing.

Figure 4b gives the average output rate at different wafer start rates for the same model as Figure 4a. As we can see, the average output rate peaks at the same level, about 6.

Figures 4a and 4b determine the capacity for a fixed product mix. Because of process variation for different products, the capacity varies considerably with the mix of products. Figure 4c plots capacity as the percentage of one product changes.

From an operational point of view, even this calculated definition of capacity is inadequate. A more operational definition for capacity is the start rate that guarantees a reasonable WIP (or throughput time). In practice, the WIP level predicted in Figure 4a is plotted against that start rate (for values where the WIP is stable) and a reasonable capacity is determined from the plot.<sup>2</sup>

**WIP reduction.** Complexity of equipment results in lengthy machine downtimes, while complexity of processes results in periods when yield remains abnormally low. But this complexity and the difficulty of estimating capacity also make clean rooms particularly subject to extremely high levels of WIP (and, consequently, throughput time). As a result, real throughput time is commonly several



**Figure 5. Comparison of two strategies for reducing work-in-process (WIP) inventory.**

times its expected value. Thus, after a problem's cause has been identified and corrected, some form of control on running the line can dramatically reduce the WIP and throughput time.

The large-scale simulation models have proved extremely effective at quantifying the different WIP reduction strategies. Figure 5 gives the results for two inventory reduction strategies:

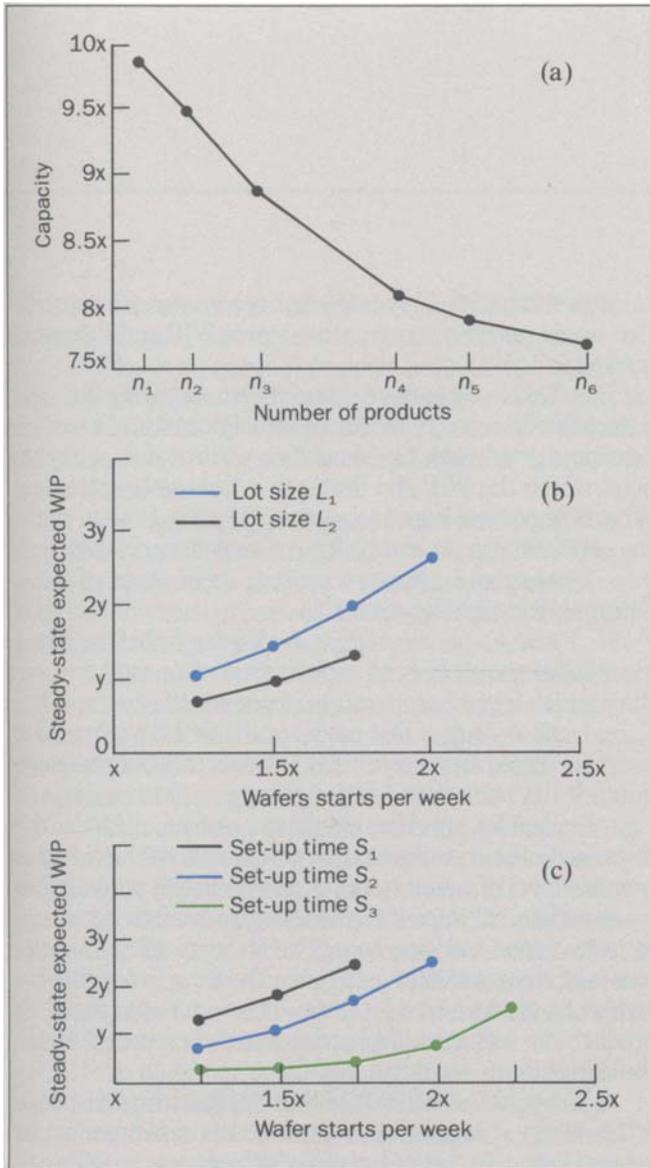
- A. Maintain the input at capacity despite the high WIP
- B. Reduce the input to zero until the WIP reaches the steady-state level; then, resume at capacity.

The initial WIP was distributed among the processes *as they were* in the room, and the two strategies were run for ten weeks. (This is an example of a transient study.) Both strategies will eventually reach the target WIP level. Strategy A takes a long time; hence, the throughput time remains excessive for that time. Strategy B sharply reduces the WIP but undershoots the desired level. This also reduces the output level (not shown in Figure 5).

These two strategies represent the extremes of a spectrum of potential recovery strategies.

**DSW Printer Throughput.** The photolithography process for the most advanced ICs uses direct step on wafer printers. These printers require extensive setup and verification procedures whenever a new mask is used. Because of their size and cost, they represent a potential bottleneck in the clean room.

The objective of this throughput study was to determine the effect of lot size, number of products, and set-up time on the capacity and expected WIP.



**Figure 6. Analysis of wafer printer capacity. (a) Manufacturing line capacity versus number of products in the line. (b) Effect of lot size for set-up time  $S_2$  and number of products  $n_4$ . (c) Effect of set-up time for lot size  $L_1$  and number of products  $n_4$ .**

A detailed simulation model was made of the DSW printer operations. To approximate the flow of product between the DSW printers and the rest of the clean room, large-scale simulations were used to obtain empirical dis-

tributions of the time between arrivals to the lithography area.

Because many runs were planned for sensitivity analysis, special simulations were developed (rather than use the large-scale model). We made the following assumptions:

- The number of printers and the sequencing and set-up rules were fixed.
- Once a printer was set up, lots were processed on a *first-come, first-served* basis. All lots that could be were processed.
- A DSW was always set up for the product with the most lots waiting.

We consider three cases here:

1. two values for lot size,  $L_1 > L_2$
2. three values for set-up times,  $S_1 > S_2 > S_3$
3. six values for the number of products,  $n_1 > n_2 > \dots > n_6$ .

This is an example of the application of experimental designs in computer simulation.<sup>9</sup> Figure 6 shows some results.

This study was able to quantify several intuitive results. Capacity is inversely proportional to the number of products (Figure 6a) because the number of setups increased. Capacity increases with lot size (Figure 6b), but so does the WIP. Figure 6c quantifies the tradeoff between WIP and set-up time.

#### Deterministic Capacity Models

Here, we describe an approach that was developed to respond quickly to questions about clean room capacity. Unlike simulation, it calculates capacity directly with such speed that several runs can be made in a fraction of a second. Its main disadvantage is that it cannot calculate statistics, which depend on the distribution of random events (e.g., the throughput time and WIP).

**The Model.** As mentioned before, the main modeling issues in estimating clean room performance decompose naturally into two categories: those that pertain to *product flow*, and those that pertain to *process time*.

---

Each has deterministic and random components.

In the deterministic capacity model, we calculate the mean number of visits to a process step (obtained from the process log and the lot rework rates), and the mean process time at each step.

The model assumes that the process time of a lot at a step is the sum of a constant term and a term that is proportional to the number of wafers in the lot at that step. The constant term represents either a set-up time or a time to process a batch that consists of a lot or more. The proportionality constant represents the time to process a single wafer.

From step to step in the process log, we assume that the number of wafers per lot decreases because of scrapped wafers and wafers held for reprocessing in reconstituted lots. These reconstituted lots also increase the number of lots in the room. Hence, they account for additional load on the machines.

These data then are used to calculate the load placed on each machine in the clean room to produce one wafer of a particular product. Alternatively, we can view this as the resources required to produce a wafer of that product. For a particular product mix and a particular input rate, then we can get the resources required to meet that program of starts.

To determine if ample capacity exists, we need to estimate the amount of each resource available. For each resource, we could multiply the number of machines and the machine availability, which yields the theoretical capacity available. If this exceeds the resources required to meet the program (as calculated above), then the program is feasible. The drawback with this approach is: If a particular machine is being used at, or near capacity, then excessive WIP and throughput time will result at that machine.

As we mentioned, this approach cannot directly estimate the mean WIP or throughput time. To circumvent this problem, the average use of the machines (fraction of time a machine is busy) is estimated at a reasonable WIP and throughput time level and multiplied by the previously

calculated capacity. This leads to a recommended capacity that would result in a reasonably stable WIP and throughput time.

Figure 6 compares capacity predicted by the deterministic capacity model and by simulation. As we remarked previously, the simulation estimated capacity at a point where the WIP and throughput time were stable. What is important here is that the computer time to run the simulation study was orders of magnitude longer.

**Applications.** Now we present applications of a deterministic capacity model.

**New product introduction.** As Figure 6 shows, this model allows engineers to estimate a room's capacity using the process logs of products proposed for the room.

**Room design.** This model predicts the minimum number of machines required to produce a proposed program. It has been useful in designing new clean rooms.

Besides checking equipment orders, it is desirable to evaluate the mix of products that would yield the highest possible level of wafer starts. A simple linear program can be formulated to maximize the total production, if we assume that no machine capacities are exceeded. The elements of the constraint matrix are the time to produce a wafer of a product on a machine. The resulting optimal product mix and start level are now standard output from the deterministic capacity model.

**Program allocation.** Another application of the model arises where it is possible to produce the same product on several lines. Because of capacity's dependence on the product mix, answers to several questions are not obvious. Where to produce a product? How much to produce in each room? How much total production can all the facilities support?

We have developed a linear programming model that uses the deterministic capacity model's outputs to determine the optimal allocation.

### **Queueing Network Analysis**

Queueing network models offer a compromise between the simulation and deterministic capacity models.

---

They provide stochastic information, such as WIP and throughput time, and require minimal computer time.

However, queueing models have limited capabilities, and extensive analytical and subjective work must be done to model the IC lines accurately. Because of the subjective decisions, the predictive capabilities of the models is also questionable.

**Our Method.** Queueing network models cast the manufacturing line into a rigid class of models.<sup>10-12</sup> The modeling approach is similar to what simulation uses, because individual processes are modeled first and then put together using legal routings.

The modeler describes process times, breakdown and repair times, interarrival times, and other random events using distributions of random variables. The outputs include information about the use and steady-state distribution of throughput time and WIP.

Can queueing networks model IC manufacturing processes? To test this, we developed a queueing network model of the lithography area and compared its outputs against a validated simulation model's predictions for the same area. The model building process is similar to that of validating a simulation model. One makes many subjective decisions about the best way to approximate a feature that does not fit the standard queueing constructs.

To meet our needs, we had to enhance the software packages—QNA<sup>13</sup> and PANACEA<sup>14</sup>—we used for evaluating the models to achieve reasonable approximations. After considerable work, we found that the queueing model's expected WIP and throughput times deviated by 7 to 20 percent from those predicted by the simulation model. However, the run times were one-tenth those of the simulation; these results suggest that queueing network models may be useful for several applications.

**Essential Features.** We found certain features were essential for accurately modeling the IC process. Without any of them, we could not have achieved the accuracy that we have reported.

**Multiserver nodes.** Each node in the queueing network represents a work center. We used a work center

with one machine working at  $N$  times to approximate a work center with  $N$  machines, but this level of approximation proved to be too inaccurate for our needs.

**General service times.** The service time is the time needed to process a customer's order (or a lot, in our applications) on a machine. The standard assumption is that this is a random variable with an exponential distribution. Some of the new analysis software includes approximations for the case where service time has a different distribution.

In manufacturing environments, the real process times are closer to constants. Variability is introduced largely by wafer yield and machine breakdown. Usually, the effective distribution to process a customer is not exponential.

**General interarrival times.** Customer arrivals to the network correspond to the start times of lots in the production line. Because arrivals generally occur on a schedule, the standard assumption of exponential interarrival times of successive customers is inadequate for modeling.

**Customer routing.** As described previously, customer routing in IC wafer fabrication lines is predominantly deterministic although rework introduces some randomness. In addition, the routing can depend heavily (both in the deterministic and random portion) on the type of product being produced.

**Batching and splitting.** Batching customers before service and splitting customers after service is not a standard operation in queueing-network-analysis software packages. But it is crucial for certain batch-service processes and certain dicing (or sawing) operations, especially those that represent heavily used critical resources.

### Concluding Remarks

We have presented some techniques that we found effective for analyzing the performance of IC wafer fabrication lines. We have also discussed the considerations for selecting the techniques. The common wisdom is "when in doubt, simulate"; this is not without merit. Simulation is a

---

technique that offers the modeler the greatest flexibility and surest path to success.

The many simulation languages that are specially adapted to manufacturing applications have greatly eased the excessive time to develop a model and the difficulty of interpreting the output. However, considerable time still is required to develop models, and the modeler's judgment must be exercised. Simulations also require long run times and considerable output analysis.

Often, other techniques—such as the deterministic capacity model or the queueing network models—can answer our questions accurately and quickly. Even when these techniques are rough, we can obtain first-cut answers with them to guide simulation studies and then rely on the simulation for further verification. Consequently, these techniques are also essential for industrial engineers.

#### Acknowledgments

The work described in this paper was done over several years and involved several people. We would specifically like to acknowledge Horacio Marcos and Esfandiar Lohraspbour, who were instrumental in developing many of the large-scale simulation models of the IC clean rooms. We also acknowledge Randall V. Sandt of Rutgers University, who worked with us on the queueing network studies and identified several enhancements to the queueing network packages.

Largely because of the effectiveness of this work, several groups have been established in AT&T's IC factories to carry on this type of analysis. Without their support, this work would never have advanced to its current level.

#### References

1. J. E. Dayhoff and R. E. Atherton, "Simulation of VLSI Manufacturing Areas," *VLSI Design*, Vol. V, 1984, pp. 84-92.
2. E. Lohraspbour and S. Sathaye, "Simulation of IC Wafer Fabrication Lines," *Semicon/West Technical Proceedings*, 1984, pp. 93-97.
3. S. M. Sze, *VLSI Technology*, McGraw-Hill, New York, 1983.
4. A. M. Law, "Statistical Analysis of Simulation Output Data," *Operations Research*, Vol. 31, 1983, pp. 983-1029.
5. L. Schruben, "Confidence Intervals Using Standardized Time Series," *Operations Research*, Vol. 31, 1983, pp. 1090-1108.
6. A. Nozari, "Confidence Intervals Based on Continuous-Time Statistics," *Operations Research Letters*, 1985.
7. G. S. Fishman and P. J. Kiviat, "The analysis of Simulation Generated Time Series," *Management Science*, Vol. 13, 1967, pp. 525-557.
8. R. E. Shannon, *Systems Simulation: The Art and Science*, Prentice-Hall, Englewood Cliffs, New Jersey, 1975.
9. J. P. C. Kleijnen, *Statistical Techniques in Simulation*, Part II, Marcel Dekker, New York, 1975.
10. C. H. Sauer and K. M. Chandy, *Computer Systems Performance Modeling*, Prentice-Hall, Englewood Cliffs, New Jersey, 1981.
11. E. Gelenbe and I. Mitrani, *Analysis and Synthesis of Computer Systems*, Academic Press, New York, 1980.
12. F. P. Kelly, *Reversibility and Stochastic Networks*, J. Wiley and Sons, New York, 1979.
13. W. Whitt, "The Queueing Network Analyzer," *The Bell System Technical Journal*, Vol. 62, No. 9, November 1983, pp. 2779-2815.
14. K. G. Ramakrishnan and D. Mitra, "An Overview of PANACEA, A Software Package for Analyzing Markovian Queueing Networks," *The Bell System Technical Journal*, Vol. 61, No. 10, December 1982, pp. 2849-2872.

---

Biographies (continued)

*and applications of operations research techniques. He has a B.S. in electrical and mechanical engineering from Universidad Nacional Autónoma de México and an M.S. in industrial engineering and operations research from Pennsylvania State University. Mr. Nozari joined AT&T in 1984 and has a Ph.D. in industrial engineering from Pennsylvania State University. Mr. Sathaye is the material-flow team leader on a project to improve System 75 manufacturing by implementing JIT/TQC concepts. He joined AT&T in 1980, and has a B.Tech. in mechanical engineering from the Indian Institute of Technology in Bombay, India, and an M.E. and Ph.D. in industrial engineering and operations research from Texas A&M University. Mr. Sitarik is responsible for methods and tools for designing and diagnosing manufacturing lines. He joined AT&T in 1963 and has a B.S. from Carnegie Mellon University, an M.S. from Ohio State University, and a Ph.D. from Lehigh University, all in electrical engineering.*

*(Manuscript received February 19, 1986)*

JULY/AUGUST • VOLUME 65 • ISSUE 4