Manufacturing Operations System Design and Analysis

C. Hilton
Manufacturing Strategic Support
Technology and Manufacturing Engineering, Intel Corp.

Index words: factory, manufacturing, discrete-event, simulation, modeling

Abstract

This paper describes manufacturing operations design and analysis at Intel. The complexities and forces of both the market and the manufacturing process combine to make the development of improved semiconductor fabrication manufacturing strategies (like lot dispatching, micro and macro scheduling policies, labor utilization, layout, etc.) particularly important. We present modeling as an effective way to further this improvement. We describe various categories of models and why they are useful. We present an overview of eight examples of how we are using modeling to improve manufacturing performance and cost. These summaries illustrate how millions of dollars have been saved in direct cost and/or cost avoidance. We conclude with a brief discussion of next steps and cautions for those establishing a manufacturing operations' group.

Introduction

In *A Tale of Two Cities*, Charles Dickens begins, "It was the best of times, it was the worst of times…" [1]. Never in history has more reliable semiconductor power been available to the consumer at such a low price. And never in history have the pressures on the manufacturers of these devices been more severe. In essence, it is the best and worst of times. Product complexity is rising, market and product segmentation is increasingly fracturing the market, lead times are shrinking, available margins are shrinking, the cost to bring a semiconductor fab on line is doubling every three to six years, and the historical avenues of cost improvement—device yields, line yields, and even device shrinking—are all approaching limits.

Intel's product manufacturing occurs in several phases: fabrication of the silicon-based device, testing, assembly and packing of the device, testing again, and sometimes assembly of the computer system or sub-system together with its testing yet again. Further complexity is introduced into each of these manufacturing processes in an effort to meet the demands for special market-driven product features.

The wafer fabs, which produce the silicon-based devices, use complex processes involving two billion dollars of equipment and 300-500 operational steps all housed in an ultra-clean environment. Fabs typically require on the order of a thousand direct staff to operate. The material does not move through the factory in a linear fashion from front to back, but loops back on itself, revisiting some areas sometimes more than twenty times. This creates all sorts of "feedback loops" (known as re-entrancy) in the dynamic response of the factory to perturbations on the factory floor and in the marketplace. These perturbations can be dampened or amplified depending on factory design, operational policies, and the current state of the material being manufactured. The majority of wafer fab cost arises from capital equipment costs. Wafer fabs feed other assembly and test operations where the devices are packaged and tested under a variety of additional constraints. Assembly and test manufacturing flows are generally linear and have several dozen processing steps. To build and equip a factory costs in the order of hundreds of millions of dollars. These factories require hundreds of direct staff to operate.

Most of the high-revenue products being manufactured today did not exist two years ago. Market forces, often not well understood, can drive product functionality as well as product packaging through radical changes within a short period of time. Each new generation of product, introduced every 12-24 months, requires new and even more costly process equipment. Often this equipment is itself on the cutting edge of technology and does not always have the performance or reliability desired for cost-effective operation. This short product lifecycle makes the reuse of equipment and the flexibility of factory use very important.

At the SEMI-Advanced Semiconductor Manufacturing Conference in 1997, Clark Fuhs, Director/Principle Analyst of the Semiconductor Manufacturing Group for Dataquest said, "The next productivity leap in the semiconductor industry will have to come through the implementation of manufacturing science and of industrial engineering practices" [2]. Manufacturing science, or "factory physics" as we like to refer to it, "is a systematic description of the underlying behavior of manufacturing systems" [3]. To understand this underlying behavior of systems of this complexity and cost, we typically do not experiment directly with the manufacturing operation. Rather, validated computer-based models are built to describe the behavior, and these models are used to develop and test factory design and operational practices that will optimize factory performance and flexibility at the lowest cost.

Models: Lies Or Oracles?

From a certain point of view, a model is a lie. From another perspective, it may be an oracle. A model is a lie in the sense that it is a purposeful simplification of a problem with the intent of focusing attention on what are believed to be the critical-few discriminating attributes or salient concepts. A model can be an oracle in the sense that an appropriate model allows us to conveniently manipulate complex systems and find answers to questions we could not approach in any other way. Models may be very detailed or very general (i.e., they are written at different levels of abstraction). Various model evaluation techniques are used within computer-based models, from simple spreadsheets to complex full-factory discrete-event simulations. Table 1 is a comparison of various types of model abstractions used in a study of a representative assembly plant operation. [4]. Note in Table 1, MTTF and MTTR refer to equipment Mean- Time-To-Fail and Mean-Time-To-Repair.

This paper focuses primarily on describing applications of computer-based models that attempt to represent complex and dynamic manufacturing interrelationships. This class of models is known as discrete-event simulations (DES). DES models are used at Intel to evaluate a variety of system performance, design, automation, and operational issues in a cost-effective, non-disruptive, statistical, and realistic fashion. Figure 1 shows the 12 areas in which we focus our efforts.

DES models can often help in decision making whenever one or more of the following conditions exists [5]:

- equipment utilization is greater than ~80%
- synchronization or merging of separate operations occurs
- manufacturing actions interact with outside events
- operations with widely variable completion times interact
- contention for resources or timing constraints exists

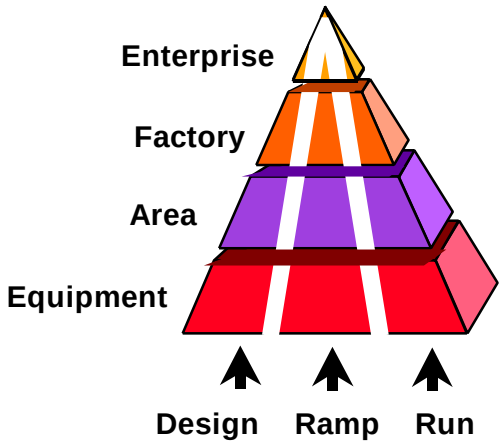

Figure 1: Scope of operations modeling activities

Analytic Method (model abstraction)	Level	Assumptions	Accuracy
Static algebraic analysis of linked systems without buffering	System	No buffering allowed Simple MTTF & MTTR model	Good model, Averages
Linked static analysis with buffering	System	Buffering allowed Simple MTTF & MTTR model	Better model, Averages
Monte Carlo analysis of static models	System Area	Distributions of input parameters are available	Better model, variability
Dynamic discrete-event simulation analysis	System Area Equip.	Buffering Failure & repair distributions	Best model, dynamic, variability
State-space analysis	Equipment	Accurate observation Fast transitions	Best validity (i.e., nature)

Table 1: Example comparison of various types of model abstractions used in solving an equipment linking design problem for an assembly operation

The Design of Experiments

Infinite Possibilities

A number of designed experiments, using static and discrete-event simulation models, are required to determine the expected factory performance under various conditions. To limit the infinite number of possible experiments, fractional and full-factorial designs can be iteratively used to define efficient sets of factory conditions to simulate.

Inputs and outputs

Our simulated factory inputs and outputs are similar to those found in an actual factory. Inputs include equipment count and layout; process time (units per hour); distributions for mean time to fail, repair, and assist; preventive maintenance (PM) time profiles; work in process (WIP) management policies including materials release; setup and batching policies; transportation times; labor availability profiles; and process flow definitions. Likewise, model outputs include equipment utilization, throughput times (TPT), and factory output (often referred to as "outs").

How Outputs are Reported

Simulation outputs are normally reported as differences or ratios of one scenario versus a baseline scenario. For example, rather than report a TPT, one would instead report that the simulation has a TPT that is 2.3 times the theoretical TPT, or that scenario B provides 10 percent more capacity than a baseline scenario A. This type of comparative reporting is useful because it allows the experimenter to focus on the key differences under study while normalizing away model simplifications in areas of less interest. For example, in this report we employ a "TPT Ratio," which is the ratio of a scenario of the simulated TPT, to an arbitrary baseline-theoretic TPT. The intent is for the ratio to be interpreted conceptually as a TPT.

Insight Through Sensitivity Studies

In a sensitivity study, outputs are measured over a range of input variable values. This is an excellent technique for understanding the expected range of the modeled factory performance, for developing intuition about how a specific manufacturing floor will perform, for understanding the impact on the study of questionable input data, and for verifying model performance. We have made extensive use of sensitivity studies in the majority of work reported here.

Example Model-Based Manufacturing Studies

Ergonomic Simulations

As unit volumes increase in the back-end, and as wafer size and weight increase in the front-end, our factory personnel experience ever increasing physical demands. In an industry unaccustomed to dealing with heavy and/or sustained physical labor, this poses new challenges for factory operation and design. We applied a specialized simulation-modeling environment that combines advanced software and computational techniques with standard ergonomic metrics and detailed full-motion models of humans [8]. The DES models, which define task type and frequency, are seamlessly merged with ergo models, which evaluate the physical impact of programmed actions on the worker. To enable visualization of the stresses, a display environment capable of stereoscopic 3-D showed worker actions in real-time from any perspective (including that of the wafer if you are so inclined!). In this environment, the upper limbs of virtual workers dynamically changed color to express the degree of strain to which they were subject during their activities. Ergonomic evaluations included reachability, field of view, RULA (rapid upper limb assessment), posture analysis, NIOSH lifting guidelines, energy expenditure, and activity timing. These models, together with other factory, cost, and strategic models were used to support our selection of an optimized wafer-lot size of 25 wafers to be used in our next generation 300mm wafer fabs.

Design for Environment

Intel is active in assuring safe and high-quality environmental conditions for its workers and the families that live in the communities where it operates. In addition, the environmental permitting process, required by our operations, often requires more lead time than it does to build and start up a manufacturing facility.

Design for Environment is integrated into our process introduction business plans and has targets that must be met, just like other process goals. Key to a Design for Environment program is the use of a variety of models including an integrated mass and energy balance model. This static model incorporates the best data and models from throughout the corporation and is interfaced with our factory design models. These models allow us to project, years in advance, the effluent and energy demands of our processes. Where needed, process changes can be made to assure environmental quality. This also allows us to effectively target R&D efforts with suppliers, universities, and national labs to assure more benign processes.

This approach is working successfully. As shown in Figures 2 and 3 we continue to use less water / (silicon area manufactured), and our air emissions are better from generation to generation.

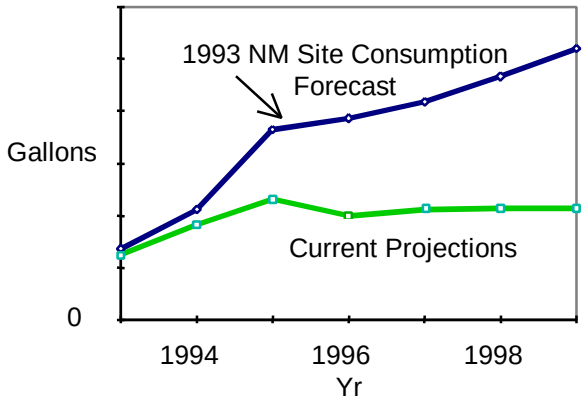

Figure 2: NM site water consumption per amount of silicon manufactured

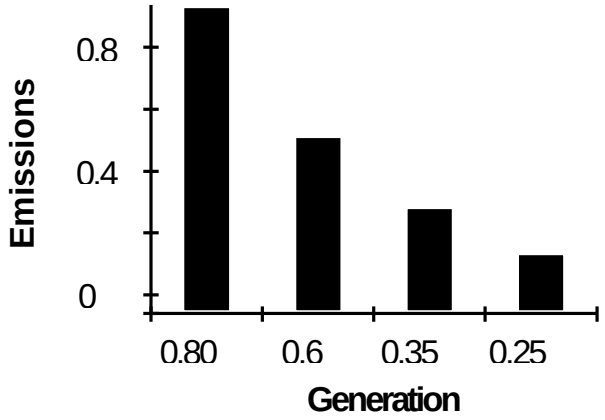

Figure 3: Targeted VOC emissions as a function of process generation (microns). Emissions are normalized by square centimeter of silicon devices manufactured.

We also have comprehensive programs to reduce, recover, and/or recycle a variety of our processing chemicals. Modeling also plays a role in these efforts. For example, one of the difficult challenges we face is dealing with the changes in the chemical content in the emissions from generation to generation. Our ability to effectively measure and model individual tools is key to determining the correct solution path. For example, in Figure 4 we can see that the total quantity and chemistry is changing for Per Flouro Compound (PFC) going from our 0.35 micron generation to our current predictions for the 0.13 micron generation by a factor of nearly 3X. In addition, the early generations contained relatively greater quantities of C_2F_6 and CF_4 when compared to the other effluent gases. The optimal technologies to deal with these kinds of emissions are fundamentally different. Models help us understand these requirements and assure appropriate technology is in place [6].

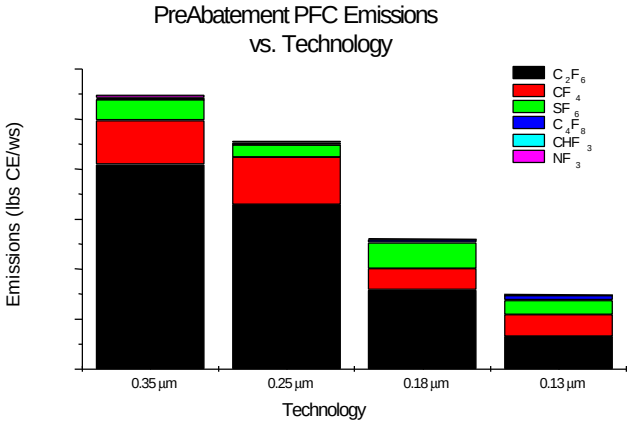

Figure 4: Pre-abatement PFC emmissions vs. technology

Assembly Floor Layout and Operation [7]

This project sought to integrate a variety of views about Pentium® processor assembly line layout from our manufacturing experts and evaluate performance of specific options as part of Intel's Copy EXACTLY! program to assure consistent implementation of our manufacturing operations policy across our plants worldwide. At one philosophical extreme is a flow layout that mimics the process flow, with process equipment hard-linked together for streamlined processing. This approach may be justified by the expectation of improved TPT and the immediate visibility provided to any equipment failures. At another extreme is the functional or grouped layout where like machines are grouped together on the floor, and product is routed to the first available machine. A justification for this approach is the expectation of improved capacity flexibility and robust performance in the presence of equipment service needs. In this study we evaluate these extremes and hybrids of them.

This study involved a variety of models both simulation and static and represented cost, product mix, and operational concerns. The output shown here is a snapshot of some of the results from the DES models.

As shown in Figure 5, factory scenarios are distinguished one from another according to their layout (flow versus functional) and the WIP machine assignment strategy used (tied versus untied). Notice there is no re-entrancy in the assembly process flows (in contrast to the highly re-entrant flows of wafer fabs).

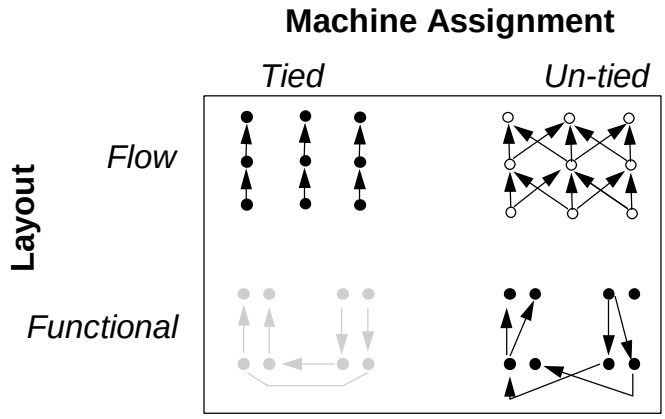

Figure 5: A conceptual view of the Factory Scenario Matrix showing the different conceptual cases considered (the functional-tied case was not considered)

A wide variety of experiments were performed assessing and ranking various WIP management strategies and understanding the impact of equipment reliability on overall factory performance. The "flow" test factory is laid out as four distinct manufacturing lines, with machines laid out sequentially in the order in which steps are performed in the manufacturing process.

In all models, a lot arriving at a step will not begin processing unless a "kanban space" at the output of the step is vacant. All models contain the exact same equipment set and have the same total number of kanban spaces available, so the maximum possible factory WIP is the same in every case. A common batching policy intended to maximize factory output while minimizing TPT is incorporated into all models where batched operations exist.

Figure 6 is an example of some of the output from the models. It shows the throughput time (relative to a base value) as a function of output for a variety of layout and operational scenarios. (For further information regarding the background of the ToC or Theory of Constraints operations point, see Karl Kempf's paper "Improving Throughput Across the Factory Life-Cycle" also in this issue of the *Intel Technology Journal*.)

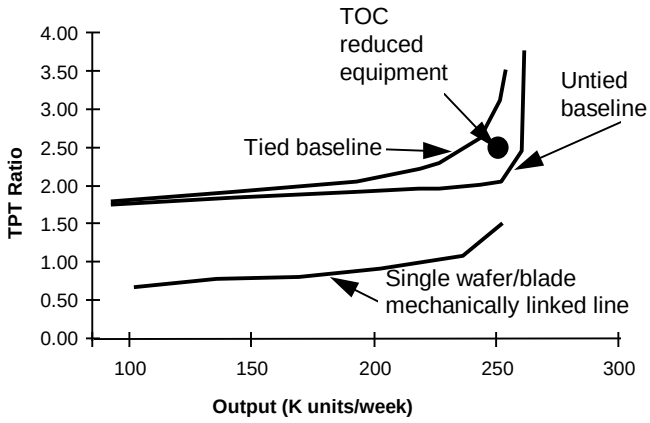

Figure 6: A comparison of baseline cases and ToC and mechanically linked-line scenarios

An examination of Figure 6 might suggest that one should operate the floor as a single blade mechanically linked line (continuous flow manufacturing). It turns out, however, that this option was much more expensive and resulted in a poor overall strategy because it did not have the required flexibility to handle the variety of expected product mix.

Other key learnings included the following:

- WIP management policy is more important than layout for this class of operation.
- Material transport, while affecting TPT and output, is not a primary modulator of performance unless it becomes very protracted.
- For the conditions studied, the functional approach outperforms the tied-line approach even when the number of equipment failures is cut in half. This suggests that simple improvements in equipment would not substantively alter the model-based recommendations.

Operation-Robot WIP Flow Interaction

A suite of three models was developed to allow assessment of WIP flow, TPT, robotic link performance and capacity issues in a key factory operation. The model quantified the impact of changes to the lot cascading ratio (a form of batching) to be about 20%. The model also highlights changes in PM and logic strategy resulting in a 10% capacity gain. These models were used to propose modifications to existing control software algorithms used by the robotics vendor. The models were used to quantify gains from eliminating specific steps and other conservative monitoring practices. This model showed that the suggested policy of putting wafers "at-risk" by processing them while waiting for monitor wafers to be read needlessly exposed the factory to line-

yield risk without any compensating productivity gains. Model use allowed us to propose novel methods for managing equipment setups.

Implant Area WIP Management Simulation

A DES model for an ion implantation processing area was based on the ability of a variety of diffusion and implant engineers, supervisors, and operators to determine the best area-WIP management rules. The model allowed implant staff to understand the impact of various WIP policies on productivity. It has highlighted a future equipment shortfall that no "creative" WIP management policy would alleviate. This model was also used as a decision-support tool to guide selection of current and future WIP management rules, and finally, it was used to guide prioritization of floor activities in the face of serious WIP surges caused by volume-ramping interactions.

Capacity Analysis Based on the Theory of Constraints—Simulation Studies

As mentioned earlier, our industry is challenged by the enormous capital costs of new factory equipment. Historically, Intel has estimated the amount of equipment needed based on a static and constant utilization-target-driven analysis (known as balanced line analysis). DES models were used to show that certain sets of equipment should be chosen as the factory constraints and that various other utilization targets should be used to assure a surplus of other equipment to optimize factory performance for a given cost, automation, and operations strategy. Based on the model findings, appropriate targets were selected for key equipment, with the expectation of a 15% reduction in TPT and a 50% reduction in variability of output for a 3% increase in capital cost over the balanced design [9].

Impact of Constraint Equipment Dedication Strategy on Fab Performance

As mentioned earlier, a unique feature of wafer-fab processing is the re-entrant nature of the material flow. For some of our fabrication processes, we require lot-level dedication of some manufacturing equipment at certain operations. This means that, when a lot returns to the operation, it must be post-processed on the identical piece of equipment on which it was pre-processed. An equivalent piece of equipment cannot be used. The intuitive expectation of the process development and operations staff was that significant factory capacity would be lost through the "interference" caused by these requirements. Additional factory space and very expensive equipment was placed into the purchasing cycle to compensate for these effects.

We built a detailed full factory model to assess the impact of this unusual constraint upon the re-entrant manufacturing line. It was found that there was no deleterious impact on factory performance, and we thoroughly understood why. While it is true that the amount of material in the queue at the dedicated steps increased and batching strategies and other resource issues needed to be optimized, we demonstrated that overall factory output and throughput time were not affected due to compensatory effects in other areas of the fab. In fact, we realized a savings of many millions of dollars by intercepting the purchase and installation of equipment for these fabrication lines.

It has often been remarked that hindsight is 20/20 and that we should have quickly realized this would be the case. One of the powerful benefits we derive from having complete and validated DES models is that we can develop correct intuition for cases like this where our own assessments fall short.

Factory Layout and Operation of the Single-Edge Connector Card Factory

Intel's Single-Edge Connector Card operation inherited much of its initial technology from our Systems (motherboard manufacturing) group. The Systems group had found very large gains when, years ago, they implemented mechanically linked continuous-flow manufacturing processes. It was assumed that as single-edge connector cards are like motherboards, their manufacturing operations should be designed in a similar fashion.

As part of our internal due diligence evaluation, we built complete DES and static models of the card operation that allowed us to study the performance of layout options similar to those shown above in Figure 5.

We found that the card operation would benefit greatly from not following in the footsteps of the motherboard line, and we found out why. The product mix of the card operation was significantly different from that of the motherboard line. Product differentiation in the market place drove constantly changing product requirements and time frames that could not effectively utilize all the capital equipment and labor that would be put in place to support a balanced and hard-linked line. We identified the optimal layout, batch sizes, buffering strategies (where needed), and operating policies. The verified models allowed us to thoroughly study the operations space and intercept factory build and fit-up plans, which resulted in a savings of many tens of millions of dollars in direct costs. Moreover, they gave us greater operational capacity and flexibility while still meeting our aggressive volume-ramp goals.

Operations Modeling Next Steps

As valuable as our model suites are and as useful as our teams of model-savvy engineers are, there is still much that remains to be done.

Quality and Voluminous Data

One of the largest hurdles inhibiting the development and application of DES models within a factory is the difficulty in obtaining detailed and accurate equipment and operations performance data in a timely fashion. These data must often be distribution based and not just a simple summary of means. The efforts of the many people within Intel working on common data definitions, common databases, and improved data gathering are of primary importance to the growing and widespread use of models.

Time Required to Create Models

There is a characteristic dissonance or tension that seems to always exist between model developers and those ultimately responsible for the tactical and strategic decisions. The decision makers want a model yesterday to help answer a problem they will first articulate tomorrow. The modelers, in turn, want months to develop validated models and gather certified data upon which multi-million dollar decisions can be confidently based. We are resolving this dissonance in two ways.

First, we are prebuilding a variety of models that, generically, we know will be of value. All future factory processes, when they leave our technology development site, will be transferred to high-volume manufacturing with a validated full-factory model already in place.

Second, as models always seem to take longer than expected to develop, we have an active program underway to ensure that past validated models can be rapidly put to new uses. We use standardized model languages. We are in the process of developing a common modeling framework and language interfaces to maximize our ability to plug-and-play various models and to model modules with each other. This framework must allow, at a minimum, the interfacing of models from different levels of abstraction. Our work to date has allowed us to solve problems in days and weeks that a year ago would have taken months to complete. More work is needed, however.

Advances in DES model performance and application areas, especially those associated with visualization, are expected to accelerate. These applications too must integrate into common data and language frameworks where possible. There are no industry standards to guide this work today.

The Speed of Model Experiment Execution

Model-based experimentation, while very fast compared to working directly with a factory, is still a bottleneck for various reasons.

First, it is difficult to determine the right level of abstraction to be used in a proposed model. This is largely an art now that is best done by our most experienced modeling engineers. Large DES models (due to detailed abstractions) take many hours of CPU time to run, require very large data sets, and may be complex to write. One would like to only use them when necessary and instead use fast analytic models, or models with less detail, whenever possible. We have studied this problem and for the time being have opted to model all the detail we can afford. We set priorities as to where detail is needed and we do all that the budget or time frame allows. Our experience when improving already highly effective wafer-fabs, running at full volume, has been that everything matters and interacts with everything else, thus requiring high-quality DES models. Linear flow assembly processes are not so sensitive.

The problem of long execution times can also be addressed by improving analytical models. In the past, queuing theory analytic solutions have not readily handled the re-entrant constraint of wafer fabs. Important work [10] in this area is now being tested in various proprietary settings.

Second, the synthesis approach of using the try-test-analyze and try again method of simulation model use is slow. Improved optimization techniques are required that take into account the classes of constraint problems that must be posed and at the same time can be integrated into our modeling tools.

Third, when one analyzes the results of a full-factory simulation, particularly when validating a model or studying dynamic changes, one may be analyzing an enormous volume of data. The analyst must have a thorough understanding of factory physics, be sensitive to subtle interactions, and be able to deal with mountains of numbers. Tools to assist in this process are needed.

The Chaotic Factory

Although we refer, tongue in cheek, to our jobs within manufacturing as challenging and chaotic, we may be more correct than we know. One often makes the assumptions that models and the factories they model are well behaved. That is, that small changes in input produce, generally, small changes in output.

Work by Beaumariage and Kempf [11] has started to carefully explore the notion of chaotic behavior within semiconductor fabs. It turns out that when factories

become very heavily loaded, performance may become unpredictable and strange in unexpected ways. DES models are excellent for identifying and studying this behavior.

Their study is founded on the mathematical notion of chaos theory, which says that a complex system can appear to randomly jump between a number of stable states with very little provocation. This system instability is a consequence of nature when large- and small-scale phenomena interact (for example, when multiple re-entrant flows interact with local equipment and operator issues).

The researchers began to suspect that formal chaos was present in fabs when they observed in simulation models that optimum schedules for heavily loaded Intel fabs changed dramatically with only slight changes in input. They developed specialized small DES models that capture key fab behavior. With these models, the team was able to observe and study transitions between multiple stable states, each with its own performance profile. They found, for example, that changing the order of one lot in a queue was sufficient to move the model to a state where average TPT was increased by up to 50%. In other cases, they demonstrated transitions between states with similar average TPT performance, but with widely different variability week to week. At least part of the chaotic nature seems to be aggravated by very high equipment utilization and by issues surrounding re-entrant flow. It is not yet clear at what level these findings are relevant in a real operating factory.

Labor

Most model environments treat labor as though it were a machine or a machine resource requirement like a jig or fixture. This is a large shortcoming when modeling loaded factories in detail. People are not automatons. They plan, they create, they locally optimize, they preempt, and they "glue" together situations that might otherwise reduce a factory's performance. An increased understanding of the human role and the degree to which it affects a factory's bottom line is needed. Kempf [12] has done some recent work in this area.

Operations Modeling Cautions

A few words of caution are in order for those who are establishing modeling centers within their companies. While it is true that having validated full-factory models and other models available for use is of tremendous value, it is not without risk. These risks include:

- *Incorrect understanding of the nature of the problem*: It is very easy for us "carpet dwellers" (as we affectionately refer to ourselves) to easily lose touch with what is really happening on the factory floors of the operations we seek to model. We find it essential to have floor operations' people intimately involved in all our projects and to ensure that our modeling engineers spend time on the factory floor. At the same time, we find it is very easy for factory and other personnel to not understand the full interactive physics of the factory, to create models that are simplistic, and to interpret the output of correct models in an incorrect fashion.
- *Non-validated models*: It is very tough to create a validated model. A validated model means that there is a preponderance of evidence to suggest that the model actually behaves like the real world in the area of interest. We have been involved in the "retreading" of a number of models from other companies or organizations due to validation problems with their work. Validation practices must be carefully thought out and religiously adhered to.
- *Controlling expectations and resourcing of projects*: To many people, good models may be magic. There is an almost universal lack of understanding of how a model can be correctly developed and used. Its use depends on the assumptions with which it was constructed, on the data available, and on the skill of the analyst applying the tool. Its development may take several people many months of effort. The temptation is always present to cut corners and deliver results sooner. This is almost always a mistake. To provide a model-based answer that is significantly in error is to perhaps fatally undermine the organization's modeling efforts. We find it very helpful to align our modeling efforts with the multi-year strategic roadmaps of our customers so that when their need arises we are already ready.
- *Starting with a focus on tools rather than on understanding the problem at hand and its business implications*: We use a formal contracting process with our internal customers to define clearly the business need, expected value, data and other information required from the customer, timelines, and deliverables. This agreement must be signed by the manager of the responsible customer organization and by the manager of the central modeling group.

Conclusion

Modern semiconductor factories are far too complex and costly to be optimized without the use of validated models. Models can be effectively applied to any level,

from the enterprise level down to the level of a specific robot within a piece of process equipment. Discrete-event simulation models are able to represent the richness of an operating factory and help provide insight into the dynamic response and optimal operation of the factory floor. Intel's use of manufacturing operations models saves millions of dollars in direct and avoided costs each year.

Acknowledgments

I am grateful for the excellent work of the many model developers, engineers, and manufacturing personnel whose work is referenced here. It is a pleasure to report on their vision and contribution. In particular the work of Sean P. Cunningham, Lance I Solomon, Edward J. Yellig, Erik A. Stewart, Karl G. Kempf, Gregory J. Mazenko, Rich Polliak, Mark F. Schaeffer and his Fab 11 team, KC Yoon and his Penang team, Bill Hendricks and his Chandler team, is gratefully acknowledged.

References

[1] Dickens, C., *A Tale of Two Cities*. Book the First, Chapter 1, line 1 (Gutenberg e-text January 1994, used as source)

[2] Fuhs, C., *SEMI-Advanced Semiconductor Manufacturing Conference*, Boston, September 10-12, 1997.

[3] Hopp, W., and Spearman, M, *Factory Physics*. Irwin (a Times Mirror Company), 1996, p. 1.

[4] Modified from: Cunningham, S. P., Internal communication.

[5] Miller, R., *Manufacturing Simulation*, The Fairmont Press, Inc., 1990.

[6] Poliak, R., Internal communication.

[7] Hilton, C., Mazenko, G., Solomon, L., and Kempf, K. *Assembly Floor Layout and Operation: Quantifying the Differences*, ISSM 1996, Tokyo.

[8] Deneb Robotics, Inc., 3285 Lapeeer Rd. West, PO Box 214687, Auburn Hills, MI 48321-4687 USA, 810-377-6900.

[9] Srivatsan, V. et. al., Internal communication.

[10] Kumar, P.R., "On the guaranteed throughput and efficiency of closed re-entrant lines" to appear in *Queueing Systems: Theory and Applications*. Paper is at http://black.csl.uiuc.edu/~prkumar/postscript_files.html

[11] Beaumariage, T., and Kempf, K., "The Nature and Origin of Chaos in Manufacturing Systems," *5th Annual !EEE/SEMI Advanced Semiconductor Manufacturing Conference and Workshop*, 1994, p. 169, Cambridge, MA, USA.

[12] Kempf, K., and Spier, J., "Simulation of Emergent Behaviour in Manufacturing Systems," *6th Annual IEEE/SEMI Advanced Semiconductor Manufacturing Conference and Workshop*, 1995, p. 90, Cambridge, MA, USA.

Author's Biography

Courtland M. Hilton joined Intel in 1981 after receiving his Ph.D. in chemical engineering from the University of Illinois at Urbana-Champaign and prior degrees from Brigham Young University. Court has worked in a variety of positions from research and development, factory engineering management, knowledge applications and strategic information technologies to his current position as manager of Intel's Manufacturing Operations Modeling and Simulation Group. He is active in the SRC (Semiconductor Research Corporation) and is a popular and invited guest lecturer on many topics related to manufacturing science. He has served as an adjunct professor of statistics and is a recipient of Intel's Individual Achievement Award.
His email address is courtland.m.hilton@intel.com.