

A FULL-FACTORY SIMULATOR AS A DAILY DECISION-SUPPORT TOOL FOR 300MM WAFER FABRICATION PRODUCTIVITY

Sugato Bagchi
Ching-Hua Chen-Ritzo

Business Analytics and Mathematical Sciences
I.B.M. T. J. Watson Research Center
1101 Kitchawan Road
Yorktown Heights, NY 10598, U.S.A.

Sameer T. Shikalgar
Michael Toner

I.B.M. Systems and Technology Group
2070 Route 52
Hopewell Junction, NY 12533, U.S.A.

ABSTRACT

We describe a discrete event simulator developed for daily prediction of WIP position in an operational 300mm wafer fabrication factory to support tactical decision-making. The simulator is distinctive in that its intended prediction horizon is relatively short, on the order of a few days, while its modeling scope is relatively large. The simulation includes over 90% of the wafers being processed in the fab and all process, measurement and testing tools. The model parameters are automatically updated using statistical analyses performed on the historical event logs generated by the factory. This paper describes the simulation model and the parameter estimation methods. A key requirement to support daily and weekly decision-making is good validation results of the simulation against actual fab performance. Therefore, we also present validation results that compare simulated production metrics against those obtained from the actual fab, for fab-wide, process, tool and product specific metrics.

1 INTRODUCTION

A wafer fabrication factory is a complex manufacturing environment which may consist of hundreds of product routes, thousands of process steps with re-entrant flows through hundreds of tools. Stochasticity is introduced by the inherent variability in processing time, testing time, unplanned tool outages and wafer re-work. Additionally, some modern fabs may be controlled by multiple layers of automation rules for lot dispatching, material handling and processing. In practice, it is difficult to understand their interactions with each other and with the day-to-day decisions made by the fab managers in the interest of maintaining the productivity of the fab.

While factory managers seek to optimize long term (i.e., on the order of months or years) metrics such as cycle time, they also have short term (i.e., on the order of a few days or weeks) objectives for which accurate simulations of fab-wide or sector specific metrics may be

needed. For example, when tools go down unexpectedly and WIP accumulates in certain sectors of the fab, managers may need to consider alternative strategies for clearing the excess work-in-progress (WIP). Alternatively, a fab manager may want to know when a surge of WIP may arrive at a tool, or when an accumulation of WIP can be expected to clear a tool, so that the tool's capacity can be allocated to other jobs or be brought down for preventative maintenance. Occasionally, a customer may need to expedite an order, and fab managers must determine the most efficient method of expediting specific lots through the fab.

Given the dynamic and complex nature of the fab, the best solution to these transient and tactical problems may depend on the current state of the fab. As such, a simulation tool that produces accurate short term predictions of the movement of WIP based on the current state of the fab can be an enormous advantage. As far as the authors are aware, the literature does not report on the validation of full-factory simulation models against short-term (operational) metrics for an actual semiconductor fab.

We have developed a discrete event simulation model for predicting WIP positions and throughput for the I.B.M. 300mm wafer fabrication factory in East Fishkill, NY (Yario 2005). A rather unique aspect of this factory is that it uses a common set of resources to fabricate both high volume *production* wafers representing the relatively mature product lines being manufactured in the fab as well as the lower volume but more diverse *development* lots representing new technologies. Our simulator incorporates all the process, measurement and testing tools being utilized and 90% of the wafers being processed in this factory within a single, integrated model. This simulation tool is capable of automatically accessing the current operational state of the factory for the purpose of calibrating its parameters.

Our key objectives for building this tool include the following:

- Forecast WIP position one to five days into the future to aid day-to-day decision-making

- Optimize the production control rules for implementing tactical decisions
- Provide analytic support for strategic productivity initiatives
- Quantify the potential impact of predicted industry trends (Pillai 2006) on manufacturing productivity (e.g., smaller lot sizes and greater product mix)

To meet these objectives, the requirements of our simulator include the ability to:

1. Model the full scope of the fab with enough detail to provide accurate predictions for fab-wide and localized metrics a few days out from a given state of the fab
2. Perform quickly enough for multiple replications and experiments to be run by the user in an interactive manner
3. Require ‘zero’ data entry for the user
4. Provide automated parameter extraction for the simulation model, based on historical performance data from a rolling time window, so that the model can be re-calibrated to the changing characteristics of the fab.

In this paper, we describe the components of our full-factory simulation model and the production control rules it implements. We describe our approach to addressing the objectives and implementing the capabilities listed above. We also report on the validation of the simulated productivity metrics outputs against the actual data obtained from the factory.

2 BACKGROUND

The use of simulation for modeling semiconductor manufacturing is abundant in the literature. Simulation studies are typically focused on studying the longer term, steady state behavior of the fab. For this purpose, one stream of research has focused on developing simple but efficient models of the fab for estimating impact of tool utilization or throughput on cycle time (Johnson et al. 2005, Rose 2007). Another stream of research has focused on the use of discrete event simulators that allow for more detailed modeling of inputs and outputs (DeJong and Fischbein 2000). For example, Pillai et al. (2004) study, among other things, the impact of priority lots, new product introductions, and lot sampling policies. Bureau et al. (2007) use simulation to compare alternative WIP management policies. Sivakumar and Chong (2001) quantify the effects of lot size, lot release controls and machine dispatching rules, among other things, on throughput, cycle time and cycle time spread, for backend manufacturing steps. Klein and Kalir (2006) develop a simulation model for studying transient behavior

of a factory as new products ramp up and old products ramp down.

While the previously cited simulation efforts are primarily focused on studying the impact of various factors from a planning perspective, our focus is on producing near term predictions, based on the current state of the fab, for day-to-day decision making. The reasons for this are described in the introduction. Therefore, successful validation of our model pertaining to short term movements in WIP is of considerable importance, and we present several validation results in this paper. Much of the prior literature provide limited validation results focusing on high-level metrics such as fab cycle time or throughput. In contrast, the users of our simulation model are interested in metrics at the sector, process or product flow level, in addition to the fab level. Therefore, we present validation results at various levels of granularity.

We chose to develop our discrete event simulator in-house rather than using commercially available packages so as to allow ourselves the greatest amount of flexibility in modeling and reporting. Simulation efficiency for a factory-wide model and the ability to integrate the simulator with in-house optimization plug-ins were also high priorities. Additionally, visualization of the simulation, one of the main advantages of commercial simulation packages, was not a priority. For readers interested in commercial simulation packages, Mason and Jensen (1996) performed a benchmarking study of some of the most popularly used packages.

3 SIMULATION MODEL

3.1 Fab Model Components

In order to model fab-wide scope of the production control rules, the simulation model covers the entire fab. It includes over 90% of the wafers being fabricated, comprising over 50,000 wafers organized in lots of up to 25 wafers. The types of lots being modeled are *production* which is high volume, *development* which is lower volume with fewer wafers per lot, and *engineering*. Each lot is assigned to a *route* or main process description which specifies the sequence of *operations* on the lot. An operation describes a *process* to be performed on a wafer using a *recipe* that depends on the product being made. The chosen recipe determines a set of *tools* qualified for it. The simulation model includes all the routes that are required to model the targeted 90% of the WIP (about 90 routes), and all of the operations (60,000 over all the modeled routes), processes (8000), recipes (30,000), and tools (600) associated with these routes. Figure 1 shows the relationship between the various components in the simulation model.

The fab contains hundreds of process tools with a variety of wafer processing behavior that must be modeled. Most

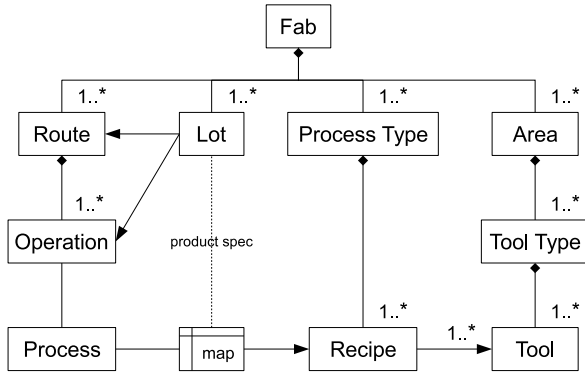


Figure 1: Fab Model Components

tools process wafer by wafer, although they may often have multiple chambers for parallel or pipeline wafer processing. Some tools process a batch of lots at a time while others process a small batch of a few wafers (within a lot) at a time. The simulator also models sampling tools where a few wafers from each lot are measured or tested. Tools are organized into tool groups that perform similar types of processes and then further grouped into areas or sectors.

The flow of WIP through the fab is controlled by various factors that are modeled in the simulation. Foremost among them are the production control rules that are described in detail in §3.2. The other factors modeled are lot holds, lot sampling, and tool downs.

Lots can go on hold due to certain pre-defined or event-based triggers. Held lots are taken out of the automated workflow so that process managers may decide their future course. Holds are more prevalent on lots on development routes where the manufacturing process is still being perfected.

Lot sampling is performed at several operations in a route which are non-mandatory. These are usually test and measurement steps. Various rules based on quality control determine the sampling rate of the number of lots that perform this operation as well as the number of wafers sampled from these lots.

Tool down-time can be grouped into actual equipment maintenance event that may be scheduled or unscheduled. In addition, tools are often not available for processing wafers because they are either inhibited, preparing monitor wafers, or being re-qualified for processing.

We do not simulate the details of the automated material handling system that is used to transport lots between tools. Instead, we derive travel time parameters from historical logs.

3.2 Production Control Rules

The production control rules influence the scheduling and dispatching of lots. They are defined over a wide range of time horizons, from monthly to ‘real-time’. At monthly, or weekly, intervals, a set of rules assign ‘daily takt rates’ (DTRs) to product groups based on planned wafer starts. Meanwhile, daily, or intra-day, throughput targets are set according to a set of ‘range target’ rules based on actual work-in-progress (WIP) positions. Finally, ‘real-time’ lot dispatching rules assign priority ‘tags’ to individual lots. The range target and lot tagging rules are used to align the movement of WIP with the DTRs. targeted throughput levels. Table 3.2 summarizes the basic characteristics of the DTR, range targets and lot tags.

Table 1: Summary of Elements of Production Control

Elements	Description
DTR	Desired daily throughput by range Updated weekly
Range Target	Guides lot tagging by range Updated every 4 hours
Lot Tag	Influences dispatch priority of a lot Updated every 30 minutes

DTRs are collectively set for a *flow*, which is a group of routes with similar operations. The DTR levels for all the flows should account for the capacity of the fab and be based on the planned wafer starts, so that the throughput targets are realistically set. Within a flow, different DTR levels may be specified for different segments along the flow, depending on how the wafers are distributed from the front-end to the back-end of the routes in the flow.

To manage the wafers within a flow, the flow is divided into *ranges*, where each range consists of a set of consecutive operations that are expected to complete in 24 hours. A DTR may be specified for each range, though adjacent ranges may often be given the same DTR. Figure 2 provides an example of how DTR levels may be set for ranges in a given flow. In this figure the WIP in a range is the sum of the WIP in all operations within that range. The DTRs are typically reviewed weekly, and modified if necessary.

Given the current WIP and DTR levels for each range, production control rules are used to help achieve the DTR by influencing how lots are dispatched. These rules are implemented in two stages. The first stage defines a ‘range target’ for each range in a product flow. Unlike the DTR, the range target is re-calculated every 4 hours and is meant to be a short term throughput target for the range to bring the WIP and DTR into closer alignment. The second stage defines a ‘tag’ for each lot. The lot tags are updated every 30 minutes.

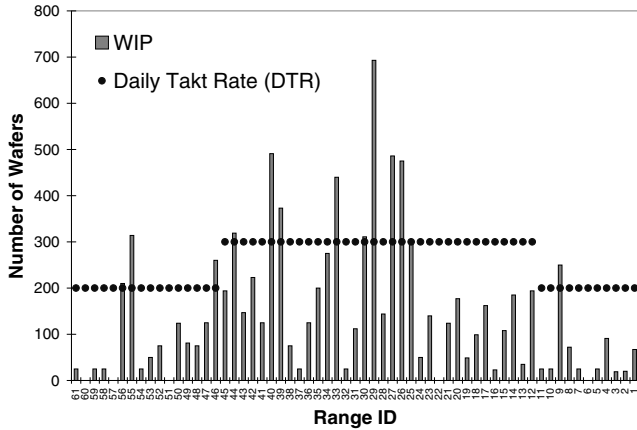


Figure 2: Example of Daily Takt Rate Levels By Range for a Given Flow

In the first stage, the rules for setting the range target for a given range, i , reflect the following general guidelines, where ‘low’ and ‘high’ WIP are defined relative to the DTR level:

- If the WIP level in range $i - 1$ (i.e., the range immediately downstream from range i) is ‘low’, then if the WIP in the range i is high enough, range i will be given a ‘high’ range target.
- If the WIP level in range $i - 1$ is ‘low’, then if the WIP in the range i is *not* high enough, range i will ‘pull’ WIP from range $i + 1$ (i.e., the range immediately upstream from range i).
- If the WIP level in range $i - 1$ is ‘high’, then the range target for range i is set to zero (i.e., range i is ‘stopped’).

In the second stage, the rules for specifying lot tags reflect the following general guidelines, where a tag of lower value receives a higher priority in dispatching:

- All other things remaining equal between a pair of lots:
 - Lots classified as high priority will be given a lower tag value.
 - Lots being ‘pulled’ (see range target rule guidelines) into the range immediately downstream will be given a lower tag value.
 - Lots nearer to completing a range will be given a lower tag value.
 - Lots arriving in a queue earlier will be given a lower tag value.

- Lots that must be processed through the range that day in order to achieve the range target will be given a lower tag value.

When a lot with a sufficiently high valued tag (i.e., sufficiently low priority) reaches the end of a range, it will be stopped if the range target has already been met.

3.3 Model Parameter Estimation

The fab-wide scope and the daily frequency of running the simulation model requires that the data inputs to the simulator be automated. The model data can be grouped into two categories. First, data such as the current WIP, routes, list of tools, can be queried directly from databases and loaded into the simulator. The second category consists of parameters that need to be estimated, such as tool performance. This section describes how we estimate these parameters based on historical event log data over a rolling time window. These estimations are implemented in software and refreshed once a week to keep track of changes in the fab.

3.3.1 Tool Processing Parameters

Our objective is to model the tools in the fab as accurately as possible without explicitly modeling the internal details of these tools. For each tool, our model predicts the timing of two events:

1. the *process start time*, given the time when a lot arrives at a tool load port and the lots currently being processed by the tool.
2. the *process end time*, given the process start time and the recipe used to process the lot.

The process start time prediction depends on the type of tool. For pipeline tools such as photolithography, wafers from each lot go through a series of steps inside the tool. The first wafer of a lot starts processing on the tool once the last wafer from the previous lot, if any, in the tool completes a processing step. This logic is shown in Figure 3(a), which shows the load complete (LC), process start (PS), and process end (PE) events along the processing timeline of two lots. We denote the time gap between two successive starts as the *time between batch ins (TBBI)*. TBBI is a linear function of the number of wafers in the previous lot. Therefore, we use linear regression to estimate this parameter from tool processing event logs for each pipeline tool-recipe pair.

Other tools process lots or batches of lots in a sequential manner. The trigger for the second lot to enter process start (PS) depends on the process end (PE) event of the previous lot. This trigger is often received a fixed time interval prior to the process end of the previous lot, denoted as *overlap* in Figure 3(b). In some tools, this overlap is negative,

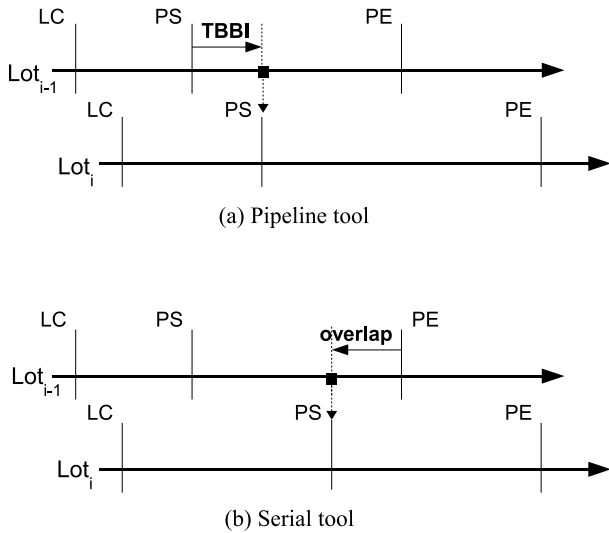


Figure 3: Process Start Logic

signifying a gap between processing successive lots on these tools. We estimate the overlap parameter for each serial tool by taking an average over the observed overlaps or gaps in all the cases where the second lot had to wait at the tool prior to processing.

The processing time for a lot is another critical factor to be estimated from tool processing event logs. This is the time between process start (PS) and process end (PE), as shown in Figure 3. It could depend on various factors such as the number of wafers, the recipe being used, and the number of parallel chambers in the tool. We categorize tools according to their processing logic (i.e., wafer, batch, small-batch and sample):

For wafer tools, processing time is linearly proportional to the number of wafers in a lot. We estimate the fixed setup time and variable time per wafer using linear regression over all tool processing events for a given recipe (see Figure 4).

For batch tools, processing time is fixed for a set of lots that can be processed in a batch. We estimate the average batch processing time for each recipe from the tool processing event logs (see Figure 5).

For small-batch tools, a small batch of wafers (e.g., 13) within a lot can be processed in constant time. The total time required for a lot depends on the multiples of small batches (1 or 2) that can be constructed for a lot. We estimate the average time required as a function of the number of batches per lot and the recipe used from the tool processing event logs.

For sampling tools, measurement and testing tools often sample a fixed number of wafers (e.g., 3) from a lot. The time required depends on the number of wafers sampled.

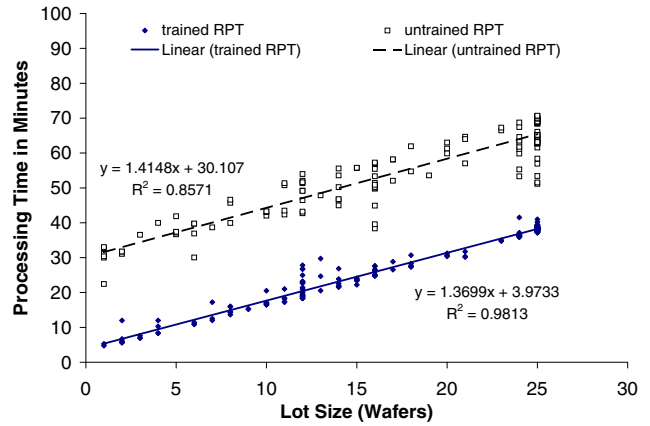


Figure 4: Processing Times For a Wafer By Wafer Tool with Training Effects

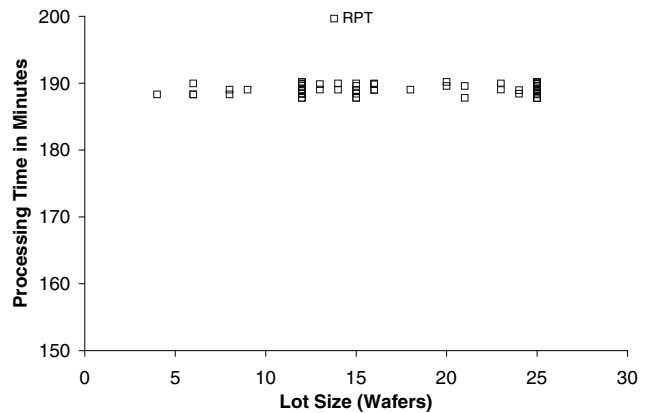


Figure 5: Processing Times For a Batch Tool

Another aspect of tool processing time that is critical for the accuracy of the simulation is ‘training’. A tool can reduce setup time when it processes two successive lots with the same recipe. To benefit from this, lot dispatching rules often create ‘trains’ (i.e., a back-to-back sequence of lots with similar recipes). When estimating the processing times from event logs, it is important to detect the effect of such trains. Figure 4 shows the raw processing time data plotted as a function of the number of wafers in a lot for a single tool. In this example, the effect of training can be seen as a 25 minute reduction in setup time. Based on the train or non-train classification of each lot, we generate two linear regressions for each recipe. During simulation we apply the appropriate regression parameters depending on the recipe of the current and previous lot processed in the tool.

3.3.2 Tool Availability

We consider the following tool states in the simulator: unplanned down, planned down, idle with WIP, idle without WIP and processing. Unplanned down states refer to tool down states that result from unexpected events such as machine failure or failed quality tests. Planned downs are down states that are the result of scheduled maintenance, or anticipated ‘trigger’ events (e.g., a certain number of wafers have been processed on the tool) that regularly cause the tool to go down for maintenance. While planned downs are certainly more predictable than unplanned down states, these ‘trigger’ events certainly lend an element of uncertainty to even the planned down states. Additionally, even for planned down states, the duration of the down is not necessarily certain. If a tool is either idle or processing, it is considered to be ‘up’, or available.

We use the same approach for simulating unplanned and planned down states. Using historical tool state data, we generate discretized empirical distributions for the time between ‘failures’ for individual tools as well as for the duration of a ‘failure’. During the simulation, down events for each tool are generated from these empirical distributions. If sampling from these distributions results in an overlapping planned and unplanned down state, we simulate these as back-to-back down states.

3.3.3 Lot Holds and Lot Sampling

In the fab, a lot may go on hold (see §3.1), in which case it does not undergo any processing. We simulate lot holds similar to the manner in which we simulate tool down states. That is, we collect historical data regarding the time between hold events and the duration of hold events for lots performing a particular process and generate empirical distributions from this data.

When an operation in a route involves a sampling tool (i.e., a non-mandatory operation), not all lots on the route will have that operation performed. Since the sampling rates for these operations are not explicitly recorded in the fab, we estimated the sampling rates from historical data. In particular, for each non-mandatory operation, we observed the number of lots that were processed at the last mandatory operation that preceded the operation. We then compare that number with the number of lots that were processed at the non-mandatory operation. The ratio between these two numbers is used to estimate the lot sampling rate for the non-mandatory operation.

3.3.4 Implementation

The simulator was implemented as a Java application. It takes less than 1 second to simulate a single replication of the flow of 50,000 wafers for a day on a 2.16 GHz

personal computer. The time it takes to run the simulation increases linearly with the number of days and replications simulated. All the inputs for a simulation are read in from networked drives containing current and historical snapshots of the fab state as well as the simulation parameters that are periodically extracted from event logs. The user specifies the starting date for the simulation and the number of days and replications to run. Users can run simulations using current fab data or historical data.

3.3.5 Simulation Initialization

In order to predict hourly WIP positions while running short term simulation spanning one to three days, the simulator cannot enjoy a substantial warm up period. Our experiments showed that the warm-up times were significantly affected by the tool states. We start the simulation with a 6AM snapshot of the WIP in the actual fab, the lot holds and tool states and initialize the simulated lots and tools accordingly.

4 RESULTS

Given our objective of using the simulator as an operational decision-support tool, we validate its predictions against the actual data observed in the fab on a daily basis. In this section, we describe the various types of validations we perform.

4.1 Validating Tool Throughput

To validate our model of how tools process lots, we independently considered each process and sampling tool in the fab. Using the models described in §3.3.1, we simulated each tool with a stream of job arrivals corresponding to the actual stream of jobs that the real tool encountered at some point in history. We then compared the simulated throughput of the tool with the actual throughput of the tool. Figure 6 shows the comparison of hourly throughput of a given tool over a 7 day simulation horizon.

4.2 Validating Fab Throughput

A popular metric for measuring productivity in a fab is fab throughput, or the total number of daily wafer moves (i.e., the total number of operations completed on all wafers) in the fab. Once the throughput for individual tools was validated, the obvious next step was to validate the daily wafer moves in the fab. In Figure 7, we compare 3 data series over eight, 24 hour periods. The first data series charts the actual fab throughput on each of the given dates. The second and third data series were generated from daily simulations over 24 hour periods that was initialized with the WIP and route data that was current at the start of each period. Each data series in Figure 7 is missing a data point

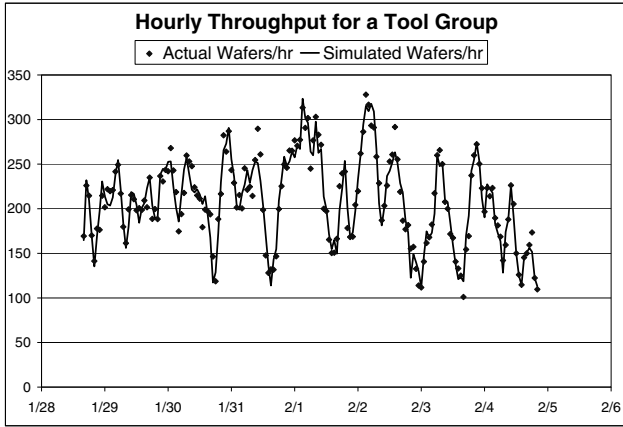


Figure 6: Validation of Tool Throughput

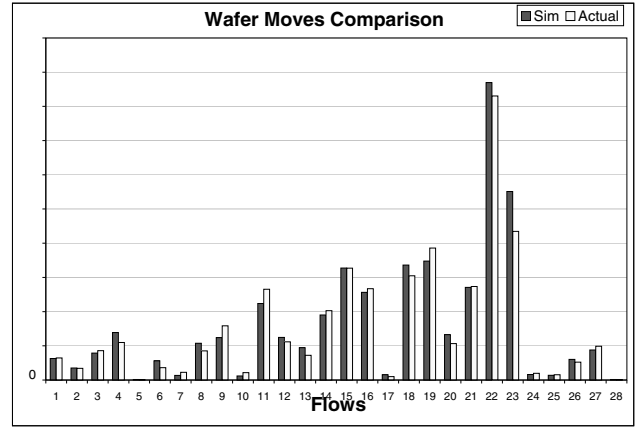


Figure 8: Validation of Moves by Flow

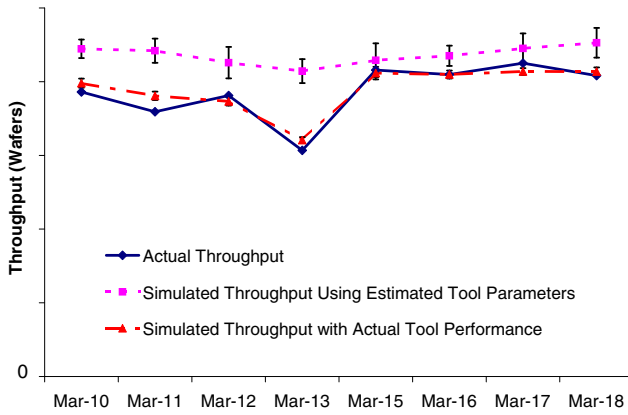


Figure 7: Validation of Overall Moves

for March 14 due to incomplete input data, which prevented us from simulating fab throughput on that date. The error bars provided for the simulated data series provide lower and upper bounds that are within two standard deviations of the average fab throughput. 20 replications were run for each simulated date.

The second data series charts the simulated average throughput when all model parameters are estimated from historical data as described in §3.3. The third data series charts the simulated average throughput when assuming actual tool availabilities instead of the estimated tool availabilities. Additionally, the third data series was generated using tool processing parameters estimated from the most recent 2 weeks of historical data as opposed to 8 weeks of historical data, as is the case for the first data series. Therefore, the uncertainty in the simulation used to generate the third data series is more or less limited to lot holds and lot sampling rates. We expect the third data series to match

more closely with the actual fab throughput than the second data series, and this indeed what is observed. The low error in the third data series indicates success in modeling the fab processes, the production control rules, and our data estimation approach. The higher error in the second data series reflects the daily deviations of tool performance from their historical averages that result in temporary bottlenecks. For predictions of future fab performance, the simulation accuracy for fab throughput is likely to be similar to that obtained by the second data series. However, if one uses the simulator to study how fab performance for past time periods could have been impacted by alternative scenarios or decisions than what had actually occurred, then the accuracy of the simulator is expected to be closer to that associated with the third data series.

In addition to overall fab throughput, engineers are also often interested in the wafer moves by flow or lot type (i.e., production, development, engineering). Figure 8 provides a comparison of our simulated wafer moves by flow with the actual moves by flow, over a 24 hour period.

4.3 Validating Lot Trajectories

Another form of forecast useful to the product and development managers of the fab is the expected movement of ‘focus’ lots. Figure 9 shows multiple replications of the simulated trajectory of a lot over several days. Specifically, a single trajectory charts the range that the lot is in on a given day. Recall that lots move from higher range numbers to lower range numbers over time. Overlaid on the simulated trajectories is the actual trajectory of this lot, as recorded in the historical data. As expected, the error in the predicted trajectory increases farther into the simulation horizon. The error is also a function of the type of route the lot is on. The trajectory of lots on production routes is

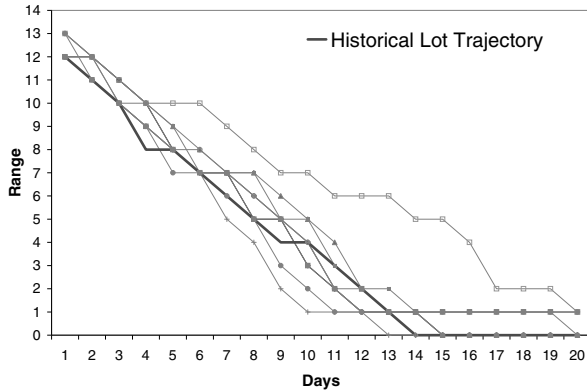


Figure 9: Example of Lot Trajectory Prediction for a Single Lot

easier to predict than development routes due to the higher chances of branching and holds in the latter.

4.4 Validating X-Factor

X-factor is a common semiconductor manufacturing metric that measures the total time that that a wafer spends queueing and processing, relative to its raw processing time. Different product flows may have different X-factor targets. Typically, the fab managers at Fishkill are interested in tracking X-factor by lot type (e.g., production, development). Table 2, shows the simulation error in the expected X-factor for production and development lots, in 4 different weeks. That is, the simulator was run with a 7 day simulation horizon for 4 different start dates. According to the results in this table, the average error over the 4 weeks is under 5%.

Table 2: Percentage Deviation of Simulated Weekly X-Factor from Actual

Week	Production Lots	Development Lots
1	4.3%	-5.1%
2	4.3%	-4.1%
3	1.2%	1.5%
4	9.2%	0.2%

4.5 Examples of Daily Projections by Process Type

Recall from Figure 1 that similar recipes belong to the same process type. A process type is confined to one area, and utilizes a subset of the tools in that area. A tool could be assigned to more than one process type. In the Fishkill manufacturing plant, WIP is regularly reported by process




Figure 10: Example of Daily WIP Projection for a Process Type




Figure 11: Example of Daily Throughput Projection for a Process Type

type because, for most purposes, it provides a convenient summary of data that would be considered too detailed if provided at the individual tool level. Additionally, process types are more easily correlated to product flows than a simple tool grouping (i.e., a partition of the set of tools) would be. Figure 10 shows a daily prediction of average hourly WIP for a process type and its comparison with the actual. The error bars around each data point capture 2 standard deviations around the simulated average hourly wip.

Figure 11 shows a daily projection of the throughput expected for a given process type. Given the throughput and WIP for a process type, one can also derive the expected number of incoming wafers for a process type. An example of daily in-gate projections is provided in Figure 12. Daily WIP, in-gate and throughput projections provide a valuable forecast of the load expected at a process type and allows equipment engineers to prioritize tool maintenance events and allocation of tool capacity accordingly.

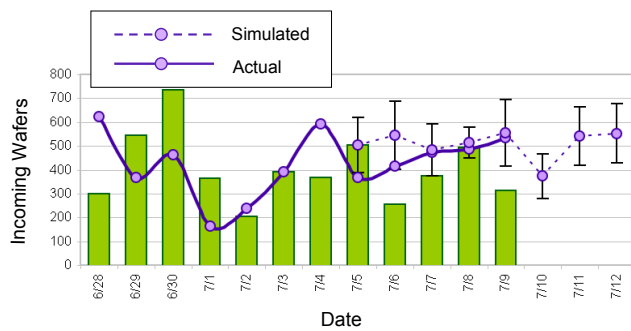


Figure 12: Example of Daily In-Gates for a Process Type

5 CONCLUSIONS

In this paper we described the components of a discrete-event simulator that we used to model the IBM 300mm semiconductor manufacturing plant in East Fishkill, NY. We also presented results validating the accuracy of the model. The simulator is capable of providing daily forecasts of the average WIP, incoming WIP, and throughput to process managers and equipment engineers to support their daily decision-making. It can also be used to analyze the potential impact of tactical changes in the production control rules or tool maintenance schedules before implementing them in the fab. Other useful applications of the simulator include the analysis of the potential impact of the high-priority lots, WIP imbalances, lot sampling policies, tool availability, lot and batch sizes and WIP balancing.

REFERENCES

- Bureau, M., S. Dauzere-Peres, C. Yugma, L. Vermarien, and J.-B. Maria. 2007. Simulation results and formalism for global-local scheduling in semiconductor manufacturing facilities. In *Proceedings of the 2007 Winter Simulation Conference*, ed. S. G. Henderson, B. Biller, M.-H. Hsieh, J. Shortle, J. D. Tew, and R. R. Barton, 1768–1773.
- DeJong, C. D., and S. A. Fischbein. 2000. Integrating dynamic fab capacity and automation models for 300mm semiconductor manufacturing. In *Proceedings of the 2000 Winter Simulation Conference*, ed. J. A. Joines, R. R. Barton, K. Kang, and P. A. Fishwick, 1505–1509.
- Johnson, R. T., J. W. Fowler, and G. T. Mackulak. 2005. A discrete event simulation model simplification technique. In *Proceedings of the 2005 Winter Simulation Conference*, ed. M. E. Kuhl, N. M. Steiger, F. B. Armstrong, and J. A. Joines, 2172–2176.
- Klein, M., and A. Kalir. 2006. Improved simple simulation models for semiconductor wafer factories. In *Proceedings of the 2006 Winter Simulation Conference*, ed. L. F. Perrone, F. P. Wieland, J. Liu, B. G. Lawson, D. M. Nicol, and R. M. Fujimoto, 1836–1839.
- Mason, S. J., and P. A. Jensen. 1996. A comparison study of the logic of four wafer fabrication simulators. In *Proceedings of the 1996 Winter Simulation Conference*, ed. J. M. Charnes, D. J. Morrice, D. T. Brunner, and J. J. Swain, 1031–1038.
- Pillai, D. D. 2006. The future of semiconductor manufacturing. *IEEE Robotics Automation Magazine* 13 (4): 16–24.
- Pillai, D. D., E. L. Bass, J. C. Dempsey, and E. J. Yellig. 2004. 300-mm full-factory simulations for 90- and 65-nm ic manufacturing. *IEEE Transactions on Semiconductor Manufacturing* 17 (3): 292–298.
- Rose, O. 2007. Improved simple simulation models for semiconductor wafer factories. In *Proceedings of the 2007 Winter Simulation Conference*, ed. S. G. Henderson, B. Biller, M.-H. Hsieh, J. Shortle, J. D. Tew, and R. R. Barton, 2708–2712.
- Sivakumar, A. I., and C. S. Chong. 2001. A simulation based analysis of cycle time distribution, and throughput in semiconductor backend manufacturing. *Computers in Industry* 45 (1): 59–78.
- Yario, J. 2005. 2005 top fab: IBM. *Semiconductor International*.

AUTHOR BIOGRAPHIES

SUGATO BAGCHI is a Research Staff Member in the Business Analytics and Mathematical Sciences group at the I.B.M. T.J. Watson Research Center in Yorktown Heights, NY. His email address is <bagchi@us.ibm.com>.

CHING-HUA CHEN-RITZO is a Research Staff Member in the Business Analytics and Mathematical Sciences group at the I.B.M. T.J. Watson Research Center in Yorktown Heights, NY. Her email address is <chenritzo@us.ibm.com>.

SAMEER T. SHIKALGAR is an industrial engineer and 300mm material logistics manager in the I.B.M. Systems and Technology Group, in Hopewell Junction, NY. His email address is <sshik@us.ibm.com>.

MIKE TONER is an industrial engineer responsible for 300mm range management and continuous improvement in the I.B.M. Systems and Technology Group, in Hopewell Junction, NY. His email address is <tonerm@us.ibm.com>.