FabTime Cycle Time Management Newsletter

Volume 9, No. 1 January 2008

Information

Mission: To discuss issues relating to proactive wafer fab cycle time management

Publisher: FabTime Inc. FabTime sells cycle time management software for wafer fab managers. New features in this month include a "Change in data rows" home-page-chart alert, and the ability to type in alert object names directly (in addition to selecting them from drop-down lists).

Editor: Jennifer Robinson

Contributors: Dov Kotlar (Tower Semiconductor)

Table of Contents

Welcome

Welcome to Volume 9, Number 1 of the FabTime Cycle Time Management Newsletter, and Happy 2008! We hope that the new year brings you happiness and prosperity. Here at FabTime, we're beginning the new year by starting installation at our 16th customer site (although, sadly, one of our existing customer sites will be closing down this month). We are also beginning, though this is difficult to believe, the ninth year of the newsletter's publication. To those of you who have been with us since the beginning, thanks for your loyalty. And to our newer subscribers, we hope that you're finding the newsletter useful. This is our 80th issue.

We have one brief community announcement this month, about the winter Fab Owners Association meeting. Our FabTime software user tip of the month is about eliminating time spent with particular hold codes or owner codes from Operation Cycle Time Trend and Pareto charts. We also have a subscriber discussion response from Dov Kotlar of Tower Semiconductor (one of our software customers) to some previous questions that we raised about cycle time benchmarking.

In our main article this month we return to a topic that we have discussed before, but that continues to pose challenges for people who manage wafer fabs: single path operations. We review the different types of single path operations, and focus on those that stem from tool dedication. We present a rule of thumb for estimating the potential impact of going from single path to dual path for a given operation, and discuss two particularly insidious forms of tool dedication: soft dedication due to operator preferences; and process restrictions for new operations. In both cases, we recommend strategies for identifying and eliminating the single path operations. We believe that this is one of the highest benefit low-cost changes that an existing fab can make to improve cycle time.

Thanks for reading!—Jennifer

FabTime

Tel: (408) 549-9932
Fax: (408) 549-9941
www.FabTime.com
Sales@FabTime.com

Community News/Announcements

Winter FOA Meeting

The next Fab Owners Association meeting will be held on February 6th and 7th at Intersil's facility in Palm Bay, Florida. Jennifer Robinson will be representing FabTime at the meeting, and hopes to see some of you there. According to the FOA website, "Fab Owners Association (FOA) is an international, non-profit, mutual benefit corporation composed of semiconductor and MEMS manufacturers, along with our industry suppliers. We are headquartered in Cupertino, California, in the heart of Silicon Valley. FOA was conceived to provide a forum for semiconductor manufacturing executives to discuss and act on common manufacturing issues. The association was founded in 2004."

FabTime is an Associate Member of the FOA, and recommends this organization to independent device manufacturers and their suppliers.

FabTime welcomes the opportunity to publish community announcements. Send them to newsletter@FabTime.com.

FabTime User Tip of the Month

Eliminate Certain Time States from Operation Cycle Time Charts

The Operation Cycle Time Trend and Pareto charts in FabTime show, for all lots that moved out of an operation, the average time that lots spent there before the move out. Usually, this is the time from move out of the previous operation to move out of the current operation. Sometimes, however, a lot will remain at the same step, but go on hold or get moved into an extended storage state via an owner code change. One of our FabTime customers requested the ability to subtract time that a lot spends under a specific hold code, or with a particular owner, from the Operation Cycle Time data. We have implemented this using the "-Own" and "-Hold" filters. For example, if a lot changes from Owner "XYZ" to Owner "XYZ:Store" when it goes into an extended hold state, setting "-Own: XYZ:Store" will result in operation cycle times for which all time spent with owner set to "XYZ:Store" is subtracted. Similarly, setting "-Hold: CustDelay" will show operation cycle times with all time spent waiting for "HoldCode=CustDelay" subtracted.

For a more detailed example, suppose that a lot moves out of Operation 900 and is in queue for Operation 1000 for 3 hours, and then (while remaining at Operation 1000) is placed on hold with hold code "CustDelay" due to a customer request. Suppose that the lot remains on hold for 48 hours (with the same hold code). After that the lot is removed from hold, and completes processing at Operation 1000 within 2 hours. The lot will have been at Operation 1000 for 3 + 48 + 2 = 53 hours. That 53 hour figure will normally show up

on the Operation Cycle Time Trend and Pareto charts. However, if we add the "-Hold: CustDelay" filter to the chart, then the operation cycle time recorded for the lot will only be 5 hours. Because these types of holds can be lengthy, this filter has the potential to have a significant impact on the data displayed.

Please note that using the "-Own:" and "-Hold:" filters is NOT the same as using a not filter (the ~ character) in the regular "Own:" and "HoldCode:" fields. The not filter tells FabTime to exclude any move transactions for which a particular owner or hold code applied. However, it doesn't look back at what the hold or owner codes were during the time at that operation. So, for example, setting "-Hold: CustDelay" tells FabTime to subtract out any time spent with HoldCode=CustDelay for any lots. This is very different from setting "Hold:~CustDelay", which tells FabTime to only include operation move outs for which the hold code specified at the move out was not equal to CustDelay. Since the hold code is usually not still in place by the time the lot moves out, the latter will rarely have any effect.

For a related example, the WIP Lot List chart, which shows the time that each lot has been at it's current operation, has "-Own:" and "-Opn:" fields, as well as "Own:" and "Opn:" fields. Setting "-Own: Store" will display all of the lots, but subtract out any time that any of those lots has spent owned by "Store". By contrast, setting "Own:~Store" will display all of the lots that currently do not have owner code Store. If a lot previously had owner code Store, but has a different owner code now, all of that time that the lot was owned by Store will show up in the time at the operation.

"Not" filters and "subtract" filters can each be useful in different circumstances. We hope that this tip has helped to make their different definitions more clear. If you have any questions about this feature, or any questions about the software, just use the Feedback form inside FabTime. Thanks!

Subscriber Discussion Forum

Issue 8.09: Cycle Time Benchmarking

Dov Kotlar from **Tower Semiconductor** sent us the following response to the main article from Issue 8.09, about cycle time benchmarking and measurement:

"**1X measure** – we measure X as the theoretical time to process a lot through the entire flow. We include load, process and unload times, using a flow data base. Smaller, hand carry lots can run faster than X.

DXF comparison to actual CT – We do this calculation occasionally but not as a common practice. The values we get from both calculations are quite similar.

Shipped lots CT – We measure both: "Fab CT" - which is the net time from start to the last operation before the ET (electrical test), and "Gross CT" - which is the time from start to ship (including ET, sort, and cage time)

Dynamic DPML calculation – We measure closed lot CT once a month for all the lots shipped within the month. We also measure DPML dynamically, as the inverse value of the WIP turn. The WIP turn is measured constantly. Once a day, we invert it to the DPML value for the entire Fab.

We mainly use the **DPML**, secondly the **closed lot CT**, and lastly the **X-factor** (occasionally)

Theoretical CT per mask layer – We do know the average theoretical CT per layer, but we don't differentiate between the layers (until recently). Lately we started analyzing the actual CT per layer, through the FabTime dashboard software. This practice is very easy and yields up interesting results. What you have to do is: choose slice variable "layer" on the Summed Operation CT Pareto chart. We have a problem with our fill in data, so (until we fix that up) we take a large range of time, in order to minimize the fill in part, and what we actually see is the average time a lot is spending in every layer. We repeat that analysis for different flows, technologies, and any other segment we are interested in. The results can highlight the layers where the biggest opportunities are, and those that need to maintain the best performance, to get the overall CT improvement.

Tackle Single Path Operations

Introduction

As the New Year begins, we would like to focus on what we believe is the number one low-cost thing that existing wafer fabs can and should do to improve cycle time: identify and eliminate single path operations. We have talked about this before, back in Issue 1.8: Understanding the Impact of Single-Path Tools and Issue 3.3: How Much Does Tool Dedication Inflate Cycle Time. We also included "number of qualified tools per tool group" as one of our Three Fundamental Drivers of Cycle Time in Issue 6.5. However, we feel that this topic is important enough to bear repeating, because it is something that we continue to hear about at every fab that we visit. We have also continued to learn more about this, as we conduct our cycle time classes at various sites, and we would like to share with you our current understanding.

In the simplest interpretation: every time you have a single-path operation in your fab, you roughly double cycle time through that step, relative to having a secondary path. We're not talking about the utilization benefit that you get in going from one tool that's 90% utilized to having two tools that are 45% utilized. We're talking purely about what happens when you go from having one relatively highly loaded tool that can perform a step to having two tools, at similar utilization levels, that can perform the step. We are talking about tool dedication, including tool qualification for new operations and soft dedication. Each of these will be discussed below.

Definitions

A one-of-a-kind tool is, as is clear from the name, a tool of which there is only one like tool in the fab. One-of-a-kind tools are primarily found in smaller fabs, though larger fabs may also have one-of-a-kind tools for pilot lines, and/or during a wafer size or technology transition. Most people who work in fabs that have one-of-a-kind tools are painfully aware of their impact. One-of-a-kind tools tend to have high per-visit cycle times (unless they are very lightly utilized), and to contribute significantly to fab variability. We've been in fabs in which a critical one-of-a-kind tool has gone down for two to three weeks. Needless to say, this wreaks havoc on fab performance. If a fab happens to have money to spend on capital purchases, elimination of one-of-a-kind tools is an excellent place to start.

Tool dedication is when a specific operation or recipe is restricted to run on a sub-set of the possible tools in the fab. Some people call this "process restriction". Tool dedication occurs for a variety of reasons:

- **Process reasons** (only some of the tools in a group have been qualified to perform certain operations). Tool dedication for process reasons occurs especially in processes that are newly developed. When a new operation is brought online, only one tool will be qualified to run that operation.
- **Yield improvement**
- Tool capability (only new tools can perform some operations, although others can also be done on the older tools)
- **Setup reduction**. Implanters, for example, are commonly dedicated to particular recipes, to minimize the relatively long setups.
- **Customer service reasons** (to get hot lots through quickly, or to reserve certain tools for high-priority customers in a foundry-like environment)
- **Layout reasons**. This can happen especially in older fabs, where capacity has been added in various locations, due to space constraints.

Tool dedication and one-of-a-kind tools can each lead to "single path operations". A single path operation is an operation which can only be performed on one tool in the fab. Regardless of their cause, single path operations have the same effect: high per-visit cycle times.

Impact of Tool Dedication and Single Path Operations

Tool dedication causes two problems. The first is that the smaller your tool groups are, the more likely you are to have tool groups with high utilizations. For example, suppose you have five equal volume recipes, and you decide to dedicate, and apply them to two tools. You have to assign three recipes to one tool, and two to the other. The tool with three recipes is going to have a higher utilization than you would get if you averaged all recipes across both tools (probably, depending on the magnitude of any setups).

The bigger problem with tool dedication is that smaller tool groups, especially single path situations, are much more heavily affected by variability than larger tool groups. You have experienced this in your day to day life many times. When you drive on a one-lane road, you are completely at the mercy of the variability in front of you. A slow truck will slow you down, just as a downtime event will slow down all of the lots passing through a single path step. By contrast, when you have two lanes on the highway, chances are there will be a chance for you to pass the truck, just as, if you have two tools, your lots can often be processed on the other tool, if one is down.

As another example, once you commit to a line at the grocery store, and take your things out of your cart, you are subject to the variability of the customers in front of

you. If someone has a stack of coupons, or a price check, you have to wait. When you're in line to check in at the airport, however, usually there are multiple check-in people, so that even if one person has a complex transaction, the overall line continues to move. The more check-in people there are, the faster the line will tend to move (although it still probably seemed slow over the holidays).

For another intuitive explanation, pretend that you're a lot, arriving to a bank of ten tools. If the tools are fully cross-qualified, then whenever you arrive you can look for an idle tool. Even if the tools are 90% loaded, chances are still pretty good (about 65%) that whenever you get there, one of those ten tools will be free. But now consider that the tools are dedicated, and when you arrive, you have to go to one particular tool, which is 90% loaded. This means that 90% of the time, when you arrive, the tool will be busy, and you'll have to wait. On average, there's only a 10% chance that you'll arrive and find the tool free. This is true even in the absence of downtime, simply due to arrival and process time variability.

Downtimes simply compound the situation. If you have two tools, each down for a random 20% of the time, the odds of both of them being down at the same time, if they are truly independent, are 4% (0.2 * 0.2). The more tools you have, the smaller the probability of all of the tools being down at the same time. With three tools, the probability goes down to less than 1%, and diminishes from there.

We've done simulation (Issue 1.8) and queueing analyses (Issue 3.3) of various tool dedication scenarios. What we're found (and the queueing analysis backs this up) is that for tools that are relatively heavily utilized, say above 60% or so, as you go from having a single path operation to having a backup tool (for a total of two tools that can perform that operation) cycle time drops by approximately 50%. If you go to three tools, the cycle time drops a bit further, but after that things start to level off. The message here is very clear. Going from single path to dual path, at the same utilization, cuts per-visit cycle times in half. Going to three qualified tools is even better, though not as critical as going from one to two. While we're not saying that everything needs to be 100% cross-qualified (the benefits do drop off), we are strongly recommending that anywhere you can, you go from single path to dual path.

This 50% cycle time reduction rule is actually a lower bound. It describes the impact on the cycle time at that one operation. Single path operations also tend to send more variability downstream, thus contributing to cycle time at other operations, also. When a tool goes down that is running a single path operation, that operation is stopped completely. When the tool comes back up, depending on the priority of that operation, you may get a large burst of like WIP, all sent downstream in a short period of time.

Of course there are exceptions and other considerations. If you have significant setups, you may need to perform a simulation analysis, to explore trade-offs. Contamination issues may require you to maintain certain levels of dedication. And, of course, sometimes you just have one-of-a-kind tools, and can't do anything about that. But wherever you can do it, going to dual path will reap cycle time significant benefits.

Soft Dedication

An important subset of tool dedication is soft dedication. Soft dedication occurs when your fab's MES and capacity model say that an operation can be performed on some set of possible tools. However, when you explore the move-out data in detail, you find that only a sub-set of the possible tools are actually being used for that operation. The operators, for a variety of possible reasons, are choosing not to use all of the qualified tools for that step. Usually this occurs because of speed issues

with the tools (the operator can log more moves by mainly using the faster tool), or because of layout issues (one or more tools is in a less convenient location, and the operators choose not to incur the transport time to get there).

Soft dedication can be a difficult thing to conquer. The operators don't perform soft dedication arbitrarily - they generally have a good reason, one that is focused on improving their local performance metrics. However, even if you decide to accept a certain amount of soft dedication, it's important to know where it is occurring, so that planning models can be modified to reflect the true situation. Where soft dedication leads to single path operations, an educational campaign may help. Alternatively, modifications to performance metrics, or the application of a more strict dispatching system, may be required.

Soft dedication can also be an indication of a fab that is operator constrained, instead of being equipment-constrained. This is done on purpose sometimes, of course, but usually it is not cost effective, given the depreciation costs of the tools. If you don't have enough operators to use all of the available tools for a given recipe, your fab might benefit overall by adding or reassigning operators.

Soft dedication can be a hidden source of cycle time problems, one worth investigating and tackling.

New Process Restrictions

Another subset of tool dedication that merits further discussion is tool dedication for new processes. Typically when a new operation is first introduced in a fab, only a single tool is qualified to run that operation. As the process matures, the question arises of when (or if) it makes sense to take time and resources to qualify additional tools. Our answer is that yes, in most cases it IS worth the time and resources to qualify a second tool. You can get a sustained 50% reduction in cycle time per visit for lots going through that step. Usually this is worth a few hours of lost capacity to qualify the second tool. This is an issue in which manufacturing and process engineering personnel have to work together. Sometimes the process engineers might lack the time to do the qualification work. Sometimes manufacturing won't want to give up the second tool for the time it takes to do the qualification. But the ultimate goal of reducing cycle time and improving variability in the fab is generally worthwhile, and an education campaign regarding the cost of single path operations may be helpful.

Another thing that can happen is that even if people aren't resistant to qualifying a second tool, the single path operation can slip through the cracks. Unless you have a good system that reports single path operations on a regular basis, and warns of the impact of any new process restrictions, you can miss these. This is a problem in many fabs because of ever-growing levels of product mix, and ever-shorter new product introduction cycles. In FabTime, we have a chart that displays, for every route-step combination that has active WIP, the number of qualified tools for that step. An example is shown at the top of the next page. We highly recommend, whatever system you use, that you find a way to make the collection and display of this data easy, so that your team can analyze it frequently.

Pilot lines, which take place on a whole new set of tools, are a special case of new process restrictions. Often fabs will make the transition between wafer sizes or technologies slowly, by buying one of each type of tool for the new major technology, and ramping up gradually over time. There's not much that can be done about the single path steps in this case - you are basically running a small, one-of-a-kind fab within another fab. Our only advice is to watch carefully for the point when it is

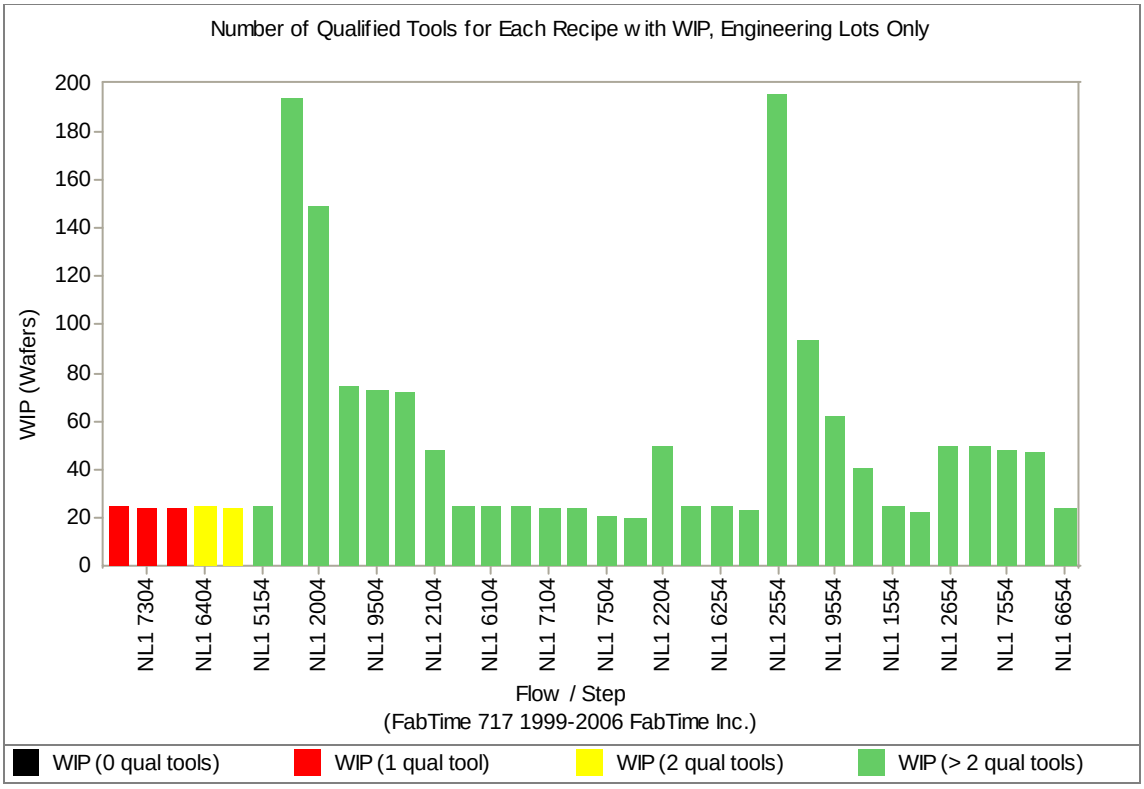

possible to add redundant tools, or convert less-used tools from the existing product line. As a general rule, and because of the behavior discussed above, you need to run one-of-a-kind tools at a lower utilization than you run tool groups with multiple tools, in order to achieve reasonable cycle times.

Conclusions

We've said this before, and we will doubtless say it again: if you are looking for low-cost ways to reduce cycle time in an existing fab, our number one recommendation is to identify and eliminate single path operations. You can reduce cycle time by roughly 50% for all lots that go through each operation that you transform from single path to dual path, as well as reducing variability in the arrival process to downstream tools. Single path operations that occur on true one-of-a-kind tools, of course, are difficult to eliminate, as they require capital purchases. However, it has been our experience that many single path operations in fabs occur due to tool dedication. Some tool dedication will probably always be necessary, especially where it results in significant yield benefits or dramatic setup reductions. However, where tool dedication can be relaxed, especially to go from single path to dual path, the potential benefit is significant.

Two particularly insidious forms of tool dedication are soft dedication due to operator preferences, and process restrictions for new operations. In both cases, we recommend education about the cycle time penalty of single path operations (education for operators and process engineers, respectively). We also recommend implementation of data systems that will highlight single path operations, so that they don't slip through the cracks. The cycle time penalty that you are paying for single path operations in your fab may be higher than you realize. A cycle time improvement program focused on the identification and elimination of such single path operations is likely to yield significant benefits.

Closing Questions for FabTime Subscribers

Do you have a formal procedure for identifying and eliminating single path operations? Do you have any systems for identifying soft dedication in your fab?

Subscriber List

Total number of subscribers: 2818, from 480 companies and universities. 21 consultants.

Top 20 subscribing companies:

- Maxim Integrated Products, Inc. (251)
- Intel Corporation (161)
- Micron Technology, Inc. (86)
- ATMEL (70)
- Analog Devices (68)
- Infineon Technologies (66)
- Freescale Semiconductor (64)
- X-FAB Inc. (63)
- Texas Instruments (59)
- International Rectifier (58)
- STMicroelectronics (57)
- Cypress Semiconductor (55)
- TECH Semiconductor Singapore (54)
- Chartered Semiconductor Mfg (51)
- NXP Semiconductors (50)
- ON Semiconductor (50)
- IBM (46)
- Spansion (37)
- Seagate Technology (33)
- BAE Systems (30)

Top 3 subscribing universities:

- Virginia Tech (11)
- Ben Gurion Univ. of the Negev (7)
- Nanyang Technological University (7)

New companies and universities this month:

- NanoMedical Systems, Inc.
- Tolomatic

Note: Inclusion in the subscriber profile for this newsletter indicates an interest, on the part of individual subscribers, in cycle time management. It does not imply any endorsement of FabTime or its products by any individual or his or her company.

There is no charge to subscribe and receive the current issue of the newsletter each month. Past issues of the newsletter are currently only available to customers of FabTime's web-based digital dashboard software or cycle time management course.

To subscribe to the newsletter, send email to newsletter@FabTime.com, or use the form at www.FabTime.com/newsletter.htm. To unsubscribe, send email to newsletter@FabTime.com with "Unsubscribe" in the subject. FabTime will not, under any circumstances, give your email address or other contact information to anyone outside of FabTime without your permission.

FabTime® Cycle Time Management Training

"It was helpful to see best-in-class methods for wafer fab cycle time management. Discussing these matters in-depth with you was quite valuable, as we could ask questions specific to our fab and processes."

Shinya Morishita
Manager, Wafer Engineering
TDK Corporation

Course Code: FT105

This course provides production personnel with the tools needed to manage cycle times. It covers:

- Cycle time relationships
- Metrics and goals
- Cycle time intuition

Price

$7500 plus travel expenses for delivery at your site for up to 20 participants, each additional participant $300.

Interested?

Contact FabTime for a quote.

FabTime Inc.
Phone: +1 (408) 549-9932
Fax: +1 (408) 549-9941
Email: Sales@FabTime.com
Web: www.FabTime.com

Do you make the best possible decisions?

- Do your supervisors possess good cycle time intuition?
- Are you using metrics that identify cycle time problems early?
- Can you make operational changes to improve cycle time?

FabTime's Cycle Time Management Training is a one-day course designed to provide production personnel with an in-depth understanding of the issues that cause cycle time problems in a fab, and to suggest approaches for improving cycle times. A two-day version is also available upon request.

Prerequisites

Basic Excel skills for samples and exercises.

Who Can Benefit

This course is designed for production personnel such as production managers, module managers, shift supervisors, hot lot coordinators, and production control.

Skills Gained

Upon completion of this course, you will be able to:

- Identify appropriate cycle time management styles.
- Teach others about utilization and cycle time relationships.
- Define and calculate relevant metrics for cycle time.
- Teach others about Little's law and variability.
- Quantify the impact of single-path tools and hot lots.
- Apply cycle time intuition to operational decisions.

Sample Course Tools

Excel Cycle Time Simulator

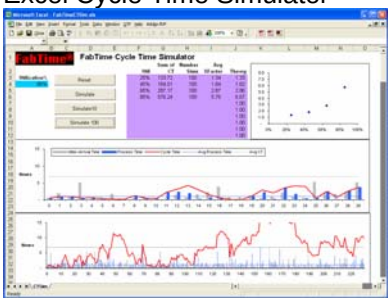

Staffing Delay Simulator

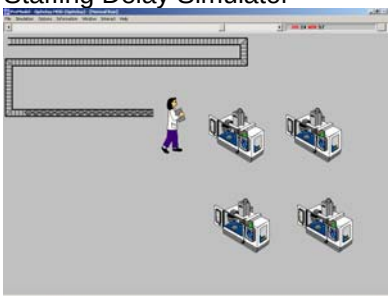

Additional Half-Day Modules

- Executive Management Session.
- Site-Specific Metrics Review.
- Capacity Planning Review and Benchmark.