FabTime Cycle Time Management Newsletter

Volume 5, No. 8 September 2004

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 the current version (6.2) include a flexible SQL filter that lets users go beyond the capabilities of the standard set of filters and an HTML archive of all past FabTime tips of the month.

Editor: Jennifer Robinson

Contributors: Allan Ravitch (Peregrine Semiconductor); Gurdeep Singh (Chartered Semiconductor Manufacturing); Walt Trybula (International Sematech); Sean Meyn (University of Illinois at Urbana-Champaign); Brett Brimhall (Maxim Integrated Products); Arnaud Sioën (Atmel); Jeff Lauffer (Fairchild Semiconductor)

Table of Contents

Welcome

Welcome to Volume 5, Number 8 of the FabTime Cycle Time Management Newsletter! This month we have several announcements, as well as a considerable amount of subscriber discussion. Our FabTime user tip of the month describes how to add a chart from a shared home page tab to a user's own home page. Subscriber discussion topics include: capacity planning for time links between process steps, understanding 300 mm cycle time performance, assessing the impact of downtime on fab performance, setting targets for WIP and turns, and defining fab loading in the presence of multiple near-bottlenecks. We also have several responses to our question about the oldest continuously operating wafer fab, and are pleased to give the nod to a Fairchild fab in Mountaintop, PA that began operations back in 1960 (!). Our thanks to everyone who participated in the discussion this month!

Because this issue contains so much subscriber discussion, we have a relatively short article by FabTime. In this article, we discuss real-time alerts sent to fab users to notify them of particular conditions in the fab. We review pros and cons of using these types of alerts at all, and then describe several examples in detail. Finally, we solicit subscriber feedback on the general usefulness of alerts, and on other types of these warning messages that might be useful in fabs. We welcome your feedback.

Thanks for reading!—Jennifer

FabTime

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

Community News/Announcements

FabTime Presentation at Upcoming ISMI Symposium

FabTime's Jennifer Robinson will be presenting the following talk as part of the Statistical Methods track at the International Sematech Manufacturing Initiative (ISMI) Symposium. The Symposium will be held October 25th-28th at the Omni Hotel in Austin, TX. Registration information can be found at www.sematech.org.

J. Robinson and F. Chance, "A Potential Approach for Monitoring Fab Cycle Time Performance through Dynamic X-Factor Control Charts," *1st ISMI Symposium on Manufacturing Effectiveness*, 2004.

Abstract: In this presentation, we discuss a fab performance measure called Dynamic X-Factor. Dynamic X-Factor was introduced by researchers at Yasu Semiconductor in Japan in an ISSM 2002 paper. It measures the speed of the production line on a short-term basis, and gives an early indication of future cycle time problems. Dynamic X-Factor is a point estimate that looks at the total wafers present in the fab, divided by the non-rework wafers that are currently being processed on tools. It can be shown to be equivalent to the traditional cycle time X-Factor (actual cycle time divided by theoretical cycle time) over time. When measured frequently, Dynamic X-Factor can give useful information about short-term, periodic behavior in the fab, such as shift change effects. The purposes of this talk are to first introduce Dynamic X-Factor, and then to solicit audience feedback regarding whether or not it is appropriate to treat fab-level Dynamic X-Factor as a control chart. Our goal is to establish procedures for the automatic identification of control limits. Fabs would then be able to flag observations falling outside of the control limits as early indicators of future degradation in shipped lot cycle time performance.

Job Change Announcement: Allan Ravitch

Allan Ravitch wrote to announce that he is leaving his current position with STMicroelectronics in Phoenix for another opportunity as the Manufacturing Manager for Peregrine Semiconductor in Sydney, Australia. We wish him well in his new position!

Job Announcement: Opening at Synopsys for a Senior DFM Applications Engineer

Synopsys Inc. has an open position for a Senior Corporate Applications Engineer to help drive the adoption of Synopsys Design For Manufacturing (DFM) technologies among Synopsys' semiconductor manufacturing customers and partners. This position will: provide expert-level end-user support for the Synopsys Virtual Stepper System product; work closely with customers to understand their needs and develop integrated flow solutions; and communicate customer and market needs to R&D. Virtual Stepper is the industry-leading software system for automated photomask defect analysis and dispositioning, and addresses a rapidly expanding market.

The ideal candidate has an MS or higher degree in physics (optics), EE, or CS; extensive industry experience in photomask lithography or inspection; excellent communication and organization skills; and familiarity with EDA tools and script programming. Interested candidates should send a cover letter and resume by email to dkohr@synopsys.com. Candidates should be near or be willing to relocate to one of these metro areas: San Jose, CA; Portland, OR; Austin, TX; or Raleigh-Durham, NC. Occasional travel required to customer sites.

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

FabTime User Tip of the Month

Add a Chart from a Shared Home Page to your own Home Page

Suppose that your fab manager has set up a shared home page that contains several charts that you find particularly useful, and which contain very specific filters. You might like to add these charts to your own home page, without having to go back to the chart list to re-create them yourself. To do this in the current version of FabTime (6.2), simply do the following:

1. View the shared home page from within your FabTime account.

2. For each chart that you would like to add to your account, click on the chart, to go down to the detailed chart page view.

3. Click on the "Add" button at the top of the left-hand pane to add the chart to your Default home page. Or, select another home page from the drop-down list (if you have configured multiple home page tabs) and then click "Add". FabTime will add this chart to your home page.

4. Once a chart is on your home page, you can easily modify it by clicking on it to go down to the detailed chart page view, making your changes, and then clicking "Update". This will not change the original chart on the home page tab that someone else had shared, but it will change the chart that you store on your own home page. This can be useful if someone else has configured a chart that is almost exactly what you want.

If you have any questions about this feature (or any other software-related issues), just use the Feedback form in the software.

Subscriber Discussion Forum

Time Links Between Process Steps

Gurdeep Singh (Chartered Semiconductor Manufacturing) sent us the following article: "Capacity Planning For Equipment With Time Link Constraints"

Background

Capacity planning is a particularly complex process when there are operations with time link constraints (time constraints between process steps), especially when there are multiple steps involved. For example, there may be a time link constraint between operation A and B, followed by another time link between steps B and C, or steps C and D. If the capacity loading of these processes is not calculated properly, there can be a devastating effect on cycle time through these operations. There may also be a capacity loss at downstream bottleneck tools.

Introduction

Production efficiency requirements drive fabs towards achieving the best possible cycle times, while also maximizing the utilization of bottleneck equipment. Sometimes the WIP in the fab is not supplied to the actual bottleneck because of artificial bottlenecks created by upstream time link processes. This can result in loss of actual output from the fab. Efficient utilization of the multi-million dollar fab equipment is very important in maintaining profitability. Capacity matching for the equipment is highly complex. Three possible scenarios in

capacity of time linked equipment are outlined below. For simplicity, these examples deal with three operations in series, each of which uses a single tool. The cases can easily be extended to allow for multiple tools at each operation.

Scenario 1

Capacity A > B >C. In this case the bottleneck is tool C (as shown below). Due to scheduled and unscheduled downtime of tool B, tool C could become idle and lose capacity. Because the WIP cannot feed through from tool A (when tool B is down), tool C may end up starved. This is compounded by the time link between tool A and tool B. With tool B down, lots may exceed their A-B time constraint, and have to be sent back to tool A for reprocessing. During this time, tool C will lose even more capacity. This is a typical case of the bottleneck losing capacity due to a time link.

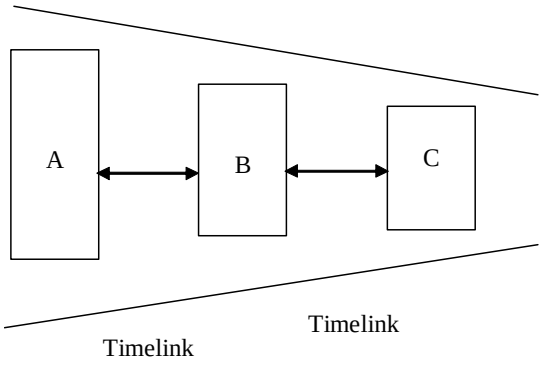

Scenario 2

Capacity of tool B is less than tool A and tool C. Capacity A > Capacity B < Capacity C (as shown below). In scenario 2, any downtime on tool A could cause the bottleneck tool B to become idle. This is because WIP must be processed on tool B within the time link window after being processed on tool A. If tool A goes down, this interrupts the continuous flow of WIP to the bottleneck tool, B. This will result in lost output of wafers from the fab.

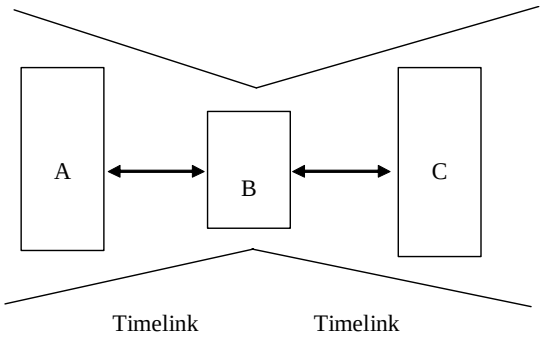

Scenario 3

Capacity of A<B<C, when the bottleneck equipment is tool A (as shown below). The capacity planning of tool C should take into consideration the cycle time of process A and process B. When tool A is down, there will be a capacity loss at B and C, because A is the bottleneck. Due to the time link constraints, B and C will run out of WIP and will be idle during that period. This is again because the time link constraint requires a fairly continuous flow of WIP to tool B and tool C, to avoid having the time constraint exceeded for any lot.

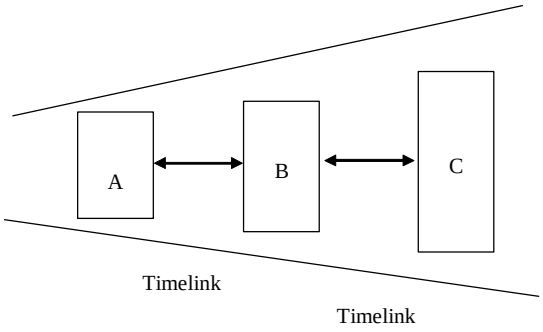

Summary

In a time link operation with more than one operation in series, if the bottleneck tool is not the first tool in the series, it may lose capacity due to upstream equipment downtime. This loss in capacity must be factored in during the capacity planning calculations."

300mm Cycle Times

Walt Trybula of SEMATECH wrote: "There are some indications that the 300mm fab cycle times are longer than 200mm cycle times. This is true even when the processes would be identical in terms of number of levels, time, and equipment speed. Is this something that your newsletter community has observed, and if so, what is the cause? I have some strong opinions on the cause, but would like to hear what others in the industry are finding. Historically, the goal has been one day per level with 1.5 days per level being achievable. If we can identify the cause, it would be a service to the industry."

FabTime Response: We have heard some of the same types of comments - that 300mm fab cycle times are turning out to be higher than corresponding 200mm fab cycle times. Possible reasons might include increased complexity of dispatching (from either increased automation or from mixing multiple products in the same lot), or reliability issues in the presence of increasingly linked toolsets. However, we haven't seen anything formally published to this effect. We would like to open this as a new subscriber discussion topic.

Issue 5.07 – Impact of Downtime on Fab Performance

Sean Meyn wrote in response to last month's issue about the impact of downtime on fab performance. "Speaking of downtime, the impact is of course dependent on policy (!). There are some nice illustrations in: M. Chen, R. Dubrawski, and S.P. Meyn, Management of Demand-Driven Production Systems, IEEE Trans. Automatic Control, Vol. 49, pp. 686-698, May 2004. This paper can be downloaded from black.csl.uiuc.edu-/~meyn/pages/CDM04.pdf."

Impact of Downtime on Fab Performance, WIP Targets, Turns Targets

Brett Brimhall of Maxim Integrated Products wrote in response to last month's article on Quantifying the Effect of Tool Downtime. "We have just gone through our 8" fab startup and ramp over the last 2 1/2 years. It has been interesting to note that in the early days all of the 8" tools were single path tools. We would see huge variations in moves and WIP due to starvation from single path tool downtimes, with tremendous impact to cycle time. Now that our fab is mature the variation due to downs is much less pronounced."

Brett also had two new topics to introduce. Regarding WIP Targets, he wrote: "Is there a generally accepted way for calculating WIP targets? Currently, I use historical data to derive WIP targets for an area or sub area in the fab. But our management asked about calculating this based on theoretical time, number of tools in the toolset, average batch size, etc. They would also like to see the WIP target for each operation in a flow. I know some of this can be answered with queuing theory but was wondering if there is some relatively simple method that is generally accepted in the industry."

FabTime Response (WIP Targets): Many people use Little's Law to calculate fab-wide WIP targets (Average WIP = Throughput Rate * Cycle Time - see newsletter issue 1.3). However, if you are looking at areas and sub-areas this is a bit trickier, because you need cycle time estimates. You can use queueing formulas to estimate per-visit cycle times at the tool group or operation level, and then use Little's Law to translate those estimates to WIP estimates. However, in order for these estimates to be accurate, you really need to have estimates for the coefficient of variation of both process times and times between arrivals (in addition to theoretical time, number of tools, and average batch size). Our experience has been that collecting and maintaining this type of data can be burdensome. We do have a spreadsheet tool that we have developed to help estimate toolgroup-level

cycle times from queueing formulas. A free version is available for download from our website at www.fabtime.com/char-curve.shtml. This version is described in newsletter Issue 2.7. We have a more detailed version (which includes multiple tools per tool group, batch arrivals, and hot lots, and allows you to roll up the estimates across a route operation list). However, the more detailed version is only available to people who take our cycle time management course. Perhaps some of our other subscribers will have something more to add.

Regarding Turns Targets, Brett wrote: "Also, as I've explained before, our 8 inch fab has ramped up and is quite mature now. We have noticed, however, that our turns rate is quite a bit less than it was for our 6 inch fab. We are scratching our heads trying to understand what fundamental differences there are. We know that our furnace batch sizes are smaller with 8 inch furnaces, but automation of other equipment is better, and should offset the furnace batch size. For instance, all of our tracks are linked to steppers now and some wet sink operations are combined in an automated fashion also. Does the industry generally see a drop in turns rate when going from 6 inch to 8 inch platforms?"

FabTime Response (Turns Targets): Turns reflect how often each wafer moves during a given shift, and are usually calculated as (operation moves) / (starting WIP). If your turns rate went down during the transition, then either you have more WIP in the 8 inch fab, or you are doing fewer moves per shift. If you record move outs only at the last operation of the linked tools in the 8 inch fab, that will reduce the number of moves that you report, and so could reduce turns rates. At any rate, we don't have any data about what the industry generally sees when going from 6 inch to 8 inch platforms, and so we are opening the question up to our newsletter subscribers.

Issue 5.06 – Increasing Fab Cycle Time Constrained Capacity

Arnaud Sioën of Atmel wrote in response to Issue 5.06: "After reading last month's issue on characteristic curves, and doing some bibliography searching on my own, I have been unable to find a definition for "fab loading". Is it defined as the load of the bottleneck tool with the highest utilization? In this case, having one bottleneck at 86% and the following one at 80% would be considered more critical than having four tools at 85%. I have tried several solutions, such as considering the average utilization of all of the tools that are loaded to more than 85%, and also considering the average utilization of the twenty main tools in the fab (where main means either a tool with a high number of moves or a single tool). My feeling is that these solutions don't accurately reflect growth in the loading of the fab. Have you heard of any metrics or benchmarks on this topic?"

FabTime Response: We think that you raise a good point. We've always understood fab loading to be defined as the utilization of the bottleneck, which is the most heavily loaded toolset. The reason for this is that it is this bottleneck utilization which limits fab throughput. However, for cycle time, you are absolutely correct that having many tool groups with high utilizations values is more of a problem. The number of tools in each toolset is also a critical factor, because having a single tool that's loaded at 80% is likely to be much more of a cycle time problem than having a toolgroup with four tools each loaded to 85%. However, we haven't heard much about aggregate metrics that capture this, beyond of course measuring the cycle time. There is a metric that was under development by SEMI last year called Overall Factory Efficiency (OFE), but our understanding is that OFE is an overall factory metric that involves all of the tools. We would like to open this question up to our other subscribers, to see how people address this issue."

Oldest Continuously Operating Fab

We heard from one fab that has been open since 1962, and has changed ownership many times, but the company did not wish to see details reported publicly.

Another subscriber wrote: "I used to work in STMicro Carrollton, TX. This site was the original Mostek plant. When I was there in 89-93, we had two fabs, Fab4 and Fab6. Now I think they have merged into one multi-floor fab, and upgraded to 6". I believe the original floor space of Fab4 dates back to the early 80's, maybe even 78-79?"

And finally, a winning entry from **Jeff Lauffer of Fairchild Semiconductor** in Mountaintop, PA. He wrote: "Our HR department sent me a copy of your last newsletter and asked me to respond to your survey question on the longest continuously operating wafer fab. How about 44 years and counting? I'd like you consider Fairchild Semiconductor in Mountaintop, PA. Its longevity is due to the dedicated, flexible, and innovative workforce with a very good relationship between the unionized workers (IUE/CWA Local 88177) and management.

This site has had a long history and several owners and has manufactured power discrete devices throughout its lifetime. It was built in 1960 as an RCA facility making germanium power transistors, NPN and PNP bipolar silicon power transistors, photo & solar cells. The site was purchased by GE in 1986 as part of the GE/RCA mega merger, and purchased by Harris Semiconductor in 1989. We built the world's first 8" SMIF'd power discrete fab at the site in 1996, with the time from ground breaking to first silicon out only 13 months. The site was spun off as part of the Intersil IPO in 1999, purchased by Fairchild in 2001, and continues to manufacture power discrete MOSFETs & IGBTs."

Real-Time Alerting based on Fab Conditions

In our FabTime software, we include functionality for alerting. The alerts allow end users to set triggers by which they will be alerted, in near real-time, to various conditions in the fab. The alerts are checked whenever new data is loaded into FabTime (usually every five minutes), and are sent as short text messages to email, pager, and/or cell phone. In this article, we discuss some types of alerts that we think may be useful in improving fab performance (regardless of what type of software you use to set them). We also solicit subscriber feedback regarding other types of real-time alerts.

The Case for Alerts

To make sure that we're all on the same page, let's begin by defining alerts. An alert is a real-time message sent to someone who works for a fab, to notify him or her about some current situation in the fab. Alerts can be automatically configured, based on some percentage discrepancy between actual and planned performance. Alternatively, alerts can be configured by individual users, according to very specific sets of parameters. The motivation behind sending alerts is simple. A fab is a highly complex environment, with a lot of things going on at the same time. Individuals

looking at historical reports can easily miss improvement opportunities, or simply notice them a bit too late. If you set up a system by which individuals receive warnings about specific situations, they have the opportunity to respond more quickly.

The Case Against Alerts

There is a case against alerts, of course, especially against alerts that are configured automatically by a software system. If a person receives too many alert messages, he or she will start to find them annoying, at best. At worst, the person will start to ignore the messages altogether, thus defeating the purpose of setting up the alerts in the first place. This is particularly true of alerts sent to pagers. After all, how difficult is it to just turn your pager off? That said, we believe that alerts can be useful in moderation. One path to their acceptance is to allow end users to establish their own alerts, rather than forcing people to receive alerts based on conditions that someone else (or, even worse, some rule in the software) felt were important.

Examples

Some types of alerts that we have seen used successfully, or believe could be used successfully, are described below.

Hot Lot Queue Delay: An end user responsible for a hot lot might configure an alert to notify the user if the lot ever waits in queue for more than 30 minutes. If the alert is triggered, then the user can notify someone of the problem, and get the lot moving again as quickly as possible.

Early Warning of Lots Due to Reach Time Limit: As described above in the subscriber discussion forum, sometimes there are time constraints between process steps. If a lot is not processed in time at the second operation of a sequence, it must be sent back to be re-processed at the first operation (e.g. a pre-clean). An alert can be configured based on the queue time of lots waiting for the second operation, so that an operator can be notified of lots that are approaching their time limit. An early warning about these lots can reduce the need for reprocessing, and improve the usage of the upstream tool.

Critical Tool Idle with WIP Available: We talked in Issue 5.05 about reducing time when a tool is available, but not processing, and has WIP waiting. For critical tools, it may be worth establishing an alerting system whereby someone gets paged whenever the tool spends more than 10 minutes in standby time with WIP available. Other tool state alerts can also be useful, as when a critical tool goes either into or out of a tool down state.

Some other possibilities:

- An alert to generate a list of inactive lots at a particular toolgroup (lots in queue, and not on hold, for more than some threshold).
- For critical tools and/or critical lots, an alert to indicate when a lot is finished processing, and needs to be removed from the tool.
- For bottlenecks, an alert for when the WIP falls below some threshold, or exceeds some other threshold. This can be useful for managing upstream tools that feed the bottleneck.
- For tool groups or operations, an alert for when the average inventory age exceeds some threshold (indicating that lots are sitting too long at this tool group or operation).
- For tools, an alert to indicate when the move rate per productive hour (the tool speed) drops below some target.
- For lots, an alert for when the projected completion date for the lot exceeds the due date.

Conclusions

Real-time alerting based on fab conditions can give useful early warning information

to the people who operate wafer fabs. These alerts can help to ensure the rapid flow of hot lots, to minimize the level of reprocessing at time constrained operations, and to maximize the WIP utilization of bottleneck tools, among other benefits. However, we believe that alerts should be used in moderation, to keep people from becoming de-sensitized to them, and that people should be able to maintain individual control over the types of alerts that they receive. Used in this manner, alerts can help to enable significant improvements in fab performance.

Acknowledgement

We would like to thank our FabTime customers for inspiring several of the examples described above. Their creativity in the use of alerts far exceeds our own.

Closing Questions for FabTime Subscribers

Are there other types of alerts that you have found to be useful in improving fab performance? Or that you think would be useful? Do you think that alerts for manufacturing personnel are a good idea, or more of a distraction?

Subscriber List

Total number of subscribers: 1656, from 395 companies and universities. 25 consultants.

Top 10 subscribing companies:

- Intel Corporation (81)
- Analog Devices (79)
- Freescale Semiconductor (54)
- Infineon Technologies (53)
- STMicroelectronics (50)
- Philips (42)
- Micron Technology (41)
- Seagate Technology (40)
- Texas Instruments (39)
- AMD/Spansion (35)

Top 5 subscribing universities:

- Arizona State University (10)
- Virginia Tech (10)
- Georgia Tech (6)
- Nanyang Technological University (6)
- University of California – Berkeley (6)

New companies and universities this month:

- Peregrine Semiconductor
- VIA Telecom

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 available for a small fee from FabTime's Amazon zShop, at www.amazon.com/shops/fabtime.

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 Software

"Instead of spending time preparing reports, shift facilitators can get the data they need quickly from FabTime, and then spend their time making real improvements."
Mike Hillis
Cycle Time and Line Yield Improvement Manager
AMD Fab 25

FabTime Installation

One fixed price includes

- Site license, unlimited users.
- Implementation & training.
- Software maintenance.

Pilot Project – Analyze your data with FabTime

For $4950, FabTime will

- Identify key contributors.
- Benchmark common metrics.
- Review results at your site.

Interested?

Contact FabTime for technical details or a pilot project quote.

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

Do you have the best possible information?

- Are your supervisors swamped with daily reports, but lacking real-time information?
- Is it difficult to link equipment performance to cycle time?
- Does each new cycle time analysis require IT resources?

FabTime is a digital dashboard for your fab. In real-time, it provides a comprehensive view of fab performance data – everything you need for proactive management of cycle time. FabTime is designed for hands-on use by managers and supervisors, unlike traditional reporting tools, which were designed for programmers.

A Web-Based Digital Dashboard

"I use FabTime every day, and so do the supervisors who report to me. The data that I need is right on my home page where I need it when I come in every morning."
Jim Wright
Production Manager
Headway Technologies

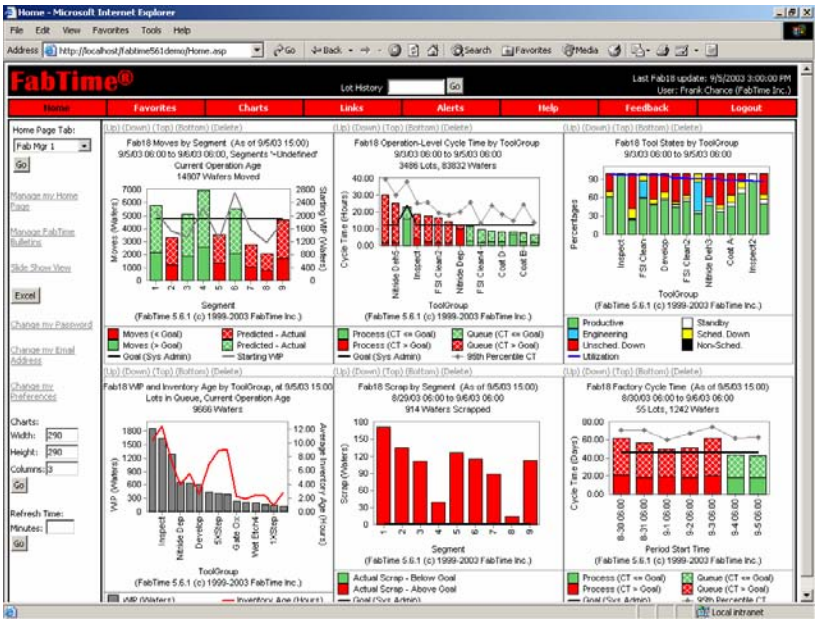

FabTime Benefits

- Cut production cycle times by 10%, hot lot cycle times by 20%.
- Focus improvement efforts on the tools that inflate cycle time.
- Improve supervisor productivity – cut reporting time by 50%.