FabTime Cycle Time Management Newsletter

Volume 5, No. 3 March 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 this version (6.0) include dynamic x-factor trend and pareto charts, and stacked by owner versions of the moves trend and pareto charts.

Editor: Jennifer Robinson

Contributors: James Ignizio (Intel); Sanjay Jain (Virginia Tech); John Fowler (ASU); Jimmy Martin (Analog Devices); Alexander Schoemig (Infineon Technologies)

Table of Contents

Welcome

Welcome to Volume 5, Number 3 of the FabTime Cycle Time Management Newsletter, and Happy Spring! This month we have a large number of new subscribing companies, due at least in part to the APICS talks that I gave recently in Fremont and Thousand Oaks, CA. We welcome you all, and encourage you to pass along the newsletter to your colleagues, if you find it useful. We would also like to remind you that you can purchase copies of past issues of the newsletter from our Amazon zShop, at www.amazon.com/shops/fabtime. There is still no charge for subscribers to receive the current issue of the newsletter each month, but we do charge for back issues.

Community announcements in this issue include two calls for papers for conference sessions related to semiconductor manufacturing applications. Subscriber discussion topics include wafer holds, cycle time and yield, operator utilization, and dynamic x-factor. In our main article this month, we revisit the topic of dynamic x-factor, a metric that we first described back in issue 4.08. Dynamic x-factor is a point estimate that looks at the total wafers that you have in your fab, divided by the wafers that are currently being processed on tools. In this article, we look further into what dynamic x-factor can tell us about how a fab is operating, with emphasis on evaluation of shift change coverage policies and comparison of relative performance across modules.

Thanks for reading!—Jennifer

FabTime

325M Sharon Park Dr. #219
Menlo Park, CA 94025
Tel: (408) 549-9932
Fax: (408) 549-9941
www.FabTime.com

Community News/Announcements

CALL FOR PAPERS: OR46, York, England

James Ignizio (Intel Corporation) submitted the following announcement: "As Organizer of the Manufacturing Session of the Annual Conference of the British Operational Research Society (OR46) in York, England, I want to extend an invitation to the readers of the FabTime Newsletter to submit a paper for presentation at this session.

The theme of the Manufacturing Session is that of the application of Operational Research (e.g., Operations Research, Management Science, Industrial Engineering, Factory Physics) to problems within today's manufacturing sector. The specific focus is the application of such methods in Semiconductor Manufacturing.

An abstract of no more than 200 words should be submitted to the Conference website (and please copy me at James.P.Ignizio@Intel.com) at: http://www.orsoc.org.uk/conf/or46/main.htm. It is not necessary to produce a formal paper. Your presentation should be of approximately 20 minutes in duration, plus allowing an additional 5 minutes for questions and answers.

The conference will be held on the campus of York University from 7 to 9 September 2004. York, by the way, is a terrific city to visit, particularly in early September. The URL leading to more detailed information is the same as indicated above."

CALL FOR PAPERS: Winter Simulation Conference, Washington, DC

Sanjay Jain and John Fowler submitted the following announcement: "We would like to invite you to contribute a paper to 2004 Winter Simulation Conference. The program includes the following mini-tracks that are relevant to semiconductor manufacturing professionals:

- Simulation Based Scheduling (coordinated by Sanjay Jain, sanjay.jain@vt.edu)
- Semiconductor Manufacturing (coordinated by John Fowler, john.fowler@asu.edu)

Contributed papers are due April 12th. The conference will be held in Washington DC, Dec 5-8. Please visit www.wintersim.org for more details on the conference and the call for papers. Please contact either of us for suitability of your proposed submissions to the specific mini-tracks."

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

FabTime User Tip of the Month

Track Down the Current Status of Lots

A new feature in FabTime is the ability to quickly identify lots that went through a particular tool during a particular time interval, and find out where they are right now. To do this, generate a "Moves Lot List" chart for the time period of interest. Fill in the name of the tool in question in the "Tool:" filter and press the "Go" button. FabTime will display a list of all lots that were processed through the tool

during that time period. To find out where the lots are now, look at the two right-most columns in the data table for the chart, labeled "Current Status". The first column, "Location", gives the operation at which the lot can be found. The second column, "Time", gives the latest time at which the lot has been at that operation (i.e., the latest time that FabTime has received data from your MES). The "Time" column is helpful if you choose to print out the data table, or export it to Excel, ensuring that you always know what time the lots listed were at the operations in question. You might use this capability to follow up on all of the lots that went through a tool that later proved to have some sort of yield problem.

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

Wafer Holds

Jimmy Martin from Analog Devices asked: "We are currently doing work on lot holds in fab and we where wondering would you or any of you subscribers have numbers about what a fab should expect for dcd rate i.e. the number of wafer holds versus wafer moves."

FabTime Response:

We don't have any data for this, but we are hearing holds mentioned more and more lately as a critical cycle time problem (see also the topic below). Perhaps some other subscribers will have something more concrete to add here!

Cycle Time and Yield (Vol. 5, No. 1)

An anonymous subscriber asked: "Have you considered the cycle time vs. yield issue in the context of a lower volume, development fab? (If not, perhaps you could in the future). Specifically, when introducing newer products for development or early production, assets are limited in terms of linefill and there is typically an understandable reluctance on the part of engineers to commit parts to a process that is either not robust or may be judged to be a problem. This often leads to engineering holds, a situation that directly contradicts the objective of increasing cycles of learning for improving the process and yield on the next cycle.

Clearly, nobody wants to move parts if the outcome is to produce scrap, but there is often a huge middle ground that is difficult to navigate. Accepting a lower than optimum yield for the purpose of uncovering other (as yet) unidentified problems seems worthwhile, but where are the diminishing returns? Has this ever been considered or modeled?"

FabTime Response:

The following paper might touch on this issue a bit. In general, we don't see much in the way of papers that apply to low volume development fabs.

■ N. Hatch and D. Mowery, "Process Innovation and Learning by Doing in Semiconductor Manufacturing," *Management Science*, Vol. 44, No. 11, I461-I477, 1998.

We do think that you raise an excellent point, and we know that smaller fabs and/or fabs with short product cycles do struggle with this issue (engineering holds). But we don't know any easy answers. We open this question up to our other subscribers, to see if anyone has comments specific to the cycle time and yield issue in a low volume development fab.

Operator Utilization

Another subscriber wrote, after hearing Jennifer speak at an APICS meeting this month: "Thanks for the interesting talk last night. This morning my boss was talking about how it's hard for him to persuade the V.P. to let him hire more equipment operators when the V.P. sees operators sitting. I thought that you said something last night about an operator utilization of 80% resulting in very high cycle times...do you have any advice as to an approximate optimum operator utilization at wafer fabs (we know that for production equipment around 85% is a good rule of thumb)? Based on your answer, maybe I could swag a good ball-park number for our manufacturing plant here, and maybe persuade the V.P. that unless he's seeing our operators sitting a good, say, 30% of the time, the cost of the extra cycle time is costing us more than additional operators would."

FabTime Response:

We haven't looked into this in much detail. In our course we have a small simulation example that we use in which we have a tool group with 4 identical tools, all 85% loaded. A single operator is used for load and unload, and based on the load/unload times, the operator is about 90% loaded. When we run the simulation, the cycle times are about 6 times theoretical. Without operator constraints, the cycle time is about 2 time theoretical. The tools end up spending nearly 15% of the time waiting for operators, and have almost no idle time.

That's an interesting simulation, and useful for getting the point across about cycle time delays from operator constraints. But it's not very rigorous (no replications, etc.) in terms of providing guidelines. The only thing that we can say about a more general value is that our impression from talking to people is that if they calculate the number of operators this way (to get a "utilization" value), then they aim for something in the 60-70% range for operator loading. But most people don't calculate operator utilizations on that kind of detailed scale – they just use some sort of moves per operator figure. A moves per operator figure will usually depend on the overall factory loading (if the fab is busier, you push the operators for more).

What we suggest is that if you can estimate the percentage of time that your key tools spend idle with WIP waiting, that's a good first pass at where operators could be causing cycle time problems. We're in the process of adding this capability to our FabTime software (where we determine how much WIP is there through the WIP transaction log – we're not depending on anyone logging a tool into a standby no-operator state).

These are our thoughts on this matter. We do have a whole section on operators in our current cycle time management course. If it would help, you could try to go for a one-day version of the course at your site, and get your VP there for this section.

Perhaps some of our other newsletter subscribers will have additional thoughts to contribute on this topic.

Dynamic X-Factor (Vol. 4, No. 8)

Alexander Schoemig (Infineon Technologies) wrote: "I wanted to send you my personal thoughts on the Dynamic X-Factor metric. At first sight, I admit I

was puzzled about the simplicity of this approach.

I had a close look at the derivation of the formulae. I found a few small gaps, but nothing serious and I consider the algorithm to be basically sound.

Still, there was something nagging in the back of my head that made me feel uneasy about it. In the article quite a few shortcomings (or let's say, things to be better aware of) are mentioned. I believe they are not complete, or not thoroughly elaborated. Especially the idea that "Dynamic X-Factor is a point estimate". It certainly is, as the current WIP is a discrete number in time as well as the # of wafers moved in or in Process are. My concern is: What is the distribution of this estimate? E.g., as we know the mean of a sample is a point estimate and it is normally distributed. But we cannot be sure that the number of wafers in process does not change dramatically within an hour or so time frame.

The article already mentions considering non-productive wafers in process and shift effects. I would be more concerned about the availability of machines, which also plays a major role in this context.

The consequence for practical use: You need to sample! I believe, this takes away one of the biggest advantages of this new algorithm. You cannot just get one estimate a day and that's it, you very well need 10, 20, maybe a hundred. And as it is in the article rightfully stated, you better not measure at the same time during the day (and let all the operators know) since there will be some tweaking going on soon. So then you need to randomize your sampling. It is also a necessity if you look down at a workcenter or a tool group, since you need to consider the arrival process!

These would be my major concerns about this new metric.

There is one aspect, however, why I particularly like this new method: It does not depend on your definition of raw process time (RPT). My experience is that if you asked 3 people about their notion of RPT, you would yield 5 answers. Why? Because before you're finished interrogating the third person, the first two persons will have already changed their mind. (And maybe gotten into some argument... ;-)"

FabTime Response

Thanks for taking such a close look at Dynamic X-Factor. The lack of reliance on raw process time data was one of the things that we really liked about it, too. We agree with your points about sampling, and about DXF as a point estimate. We've started looking at this for some actual fabs, and have concluded that you really need to report it on something like an hourly basis.

We see it not so much as a single value that you measure, but more like a control chart that you plot over time. Then you use it to a) identify systematic issues in your fab, and b) to get an early indicator of where cycle time may be drifting up (if DXF appears to be trending upward over time). We think that if your reporting system allows you to measure this on a regular and very short-term basis, you avoid most of the sampling issues, especially if looking at the fab-level data. If you can automate this, we don't think that the need to generate a lot of points is a big negative – the data is easy to collect, and not a victim of the subjectivity of raw process time estimates. We discuss this further in the article that follows.

Dynamic X-Factor Revisited

Introduction

Back in Issue 4.08, we described Dynamic X-Factor, a performance metric introduced by S. Johnishi, K. Ozawa and N. Satoh at ISSM 2002. Dynamic X-Factor is calculated by taking the total number of wafers in the fab and dividing by the number of wafers actually being processed at tools. We showed in the earlier article that Dynamic X-Factor works out to be the same as the regular cycle time X-Factor (cycle time / theoretical cycle time) on a long-term basis. However, Dynamic X-Factor is easier to calculate, and is more forward-looking than an X-Factor based on shipped lot cycle times. If we see Dynamic X-Factor drifting upwards, we know that shift lot cycle time will also drift upwards, if the situation is not changed.

As pointed out by Alex Schoemig in the above subscriber discussion topic, it's not clear what the theoretical distribution of Dynamic X-Factor really is. As a consequence, samples must be taken frequently to avoid biased results. However, if generated on a regular basis, we think that Dynamic X-Factor gives some useful clues about how the fab is operating. Two particular applications are evaluation of shift change coverage policies and comparison of relative performance across modules (or areas). Both are discussed below.

Shift Change Coverage

If you generate Dynamic X-Factor values on an hourly basis for the entire fab, or for an individual module, you may see interesting behavior. An example that mimics behavior we have observed in a real fab is included below. This chart shows that for most hours the Dynamic X-factor for the fab was between three and four, indicating that for every lot being processed, two to three lots, on average, were waiting in queue. This data suggests that over time the actual average cycle time of lots through the fab will be between three and four times the theoretical process time.

However, the chart also shows spikes at regular intervals, at which the Dynamic X-Factor rises to as high as 7.5. Now, we know that the total WIP in the fab (the numerator of Dynamic X-Factor) is unlikely to change significantly from one hour to the next. Therefore, these spikes indicate changes in the denominator of

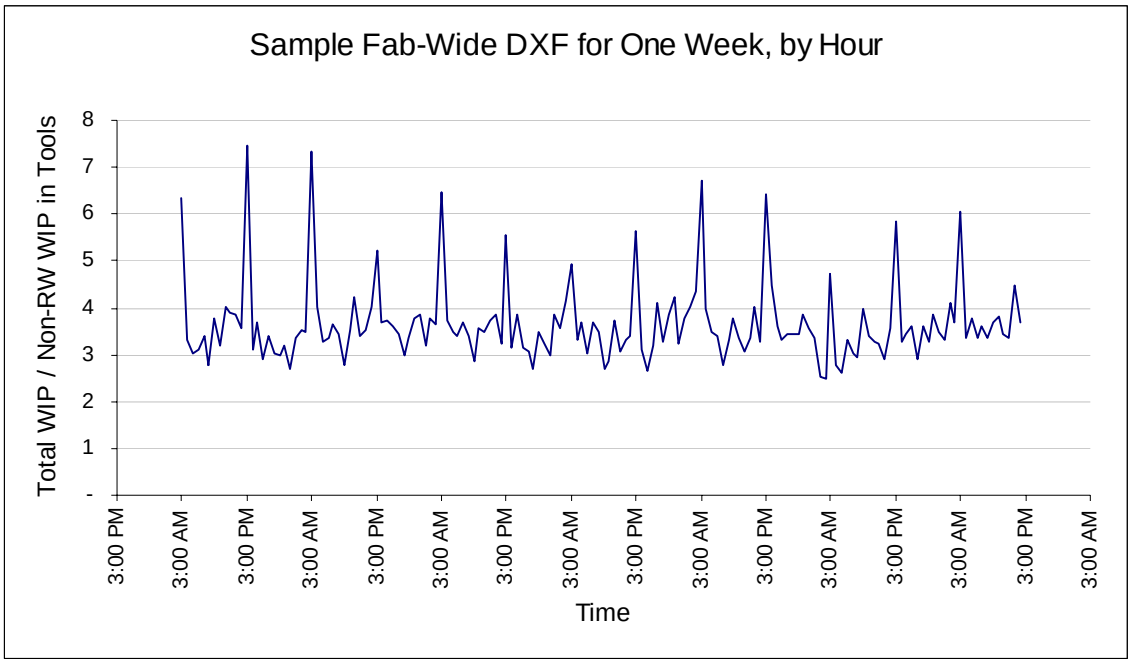

Dynamic X-Factor – the number of non-rework wafers in process at tools. Specifically, the spikes are times at which significantly fewer wafers are in process. A closer look at the chart reveals that the spikes are occurring at shift change.

This begs the question: are these spikes acceptable? Is it reasonable that at shift change, only half as many wafers as usual are in process at tools? Of course the answer depends on the situation in your fab. Perhaps there are tools that can't be run during shift change for yield-related reasons, or for accountability reasons (we must have one person responsible for making the entire run on a critical tool), or for cost control reasons (we choose not to allow any overlap in staffing). We can't give you an answer as to what's right or wrong for your fab. But certainly these spikes indicate a capacity loss that is occurring during shift change (tools are sitting idle). Dynamic X-Factor makes it easy to see this behavior, and to see the effects of any operational changes that you might put into place to mitigate it.

Another question that you might ask is: how is this any different from looking at moves by hour, and monitoring for dips around shift change? Certainly the two metrics are similar. However, a spike in Dynamic X-Factor will be visible slightly earlier (you don't have to wait until the moves are completed). Also, because of the close relationship between Dynamic X-Factor and Cycle Time X-Factor, Dynamic X-Factor emphasizes the point that these dips in moves directly inflate cycle time. Similarly, you could look at just the denominator of Dynamic X-Factor, hourly measures of WIP-in-tools. But here again, Dynamic X-Factor speaks in terms of cycle time. For the example shown previously, the behavior at shift change suggests cycle times of five or six times theoretical, instead of the three to four times theoretical performance during the shifts. That is, if the fab operated all the time the way it does during shift change, actual cycle times could end up being nearly twice what they are now – clearly not an acceptable outcome.

Relative Performance

Other useful glimpses into fab performance can come when you pareto Dynamic X-Factor by module. We would generally expect to see different values for diffusion than for photo, for example. Unless it's a bottleneck, the diffusion area will probably run at a lower Dynamic X-Factor value, because of longer theoretical cycle times and because of the way that batch tools perform (cycle time is flatter up to the point the furnaces are overwhelmed). Based on how you plan capacity for your fab, you likely have an expected ordering of Dynamic X-Factor performance by module (you expect more WIP in queue in the bottleneck areas, and less WIP in queue in areas with more spare capacity). If this ordering changes dramatically, this indicates a fairly radical change in the cycle time performance among the different modules. This can be an early indicator of some serious problem.

Note, however, that if you end up in a situation in which most of the WIP in the fab is currently in one module (e.g. a WIP bubble caused by some dramatic tool downtime problem), then you will see large swings in Dynamic X-Factor for the modules that do not have much WIP. This is the same problem that you have with using WIP Turns by module, when some modules don't have very much WIP in them. Theoretically, you can also pareto Dynamic X-Factor by tool or tool group. However, then you are more likely to run into computation problems due to down tools (no WIP in tool), or WIP shortages.

One other way in which you can use Dynamic X-Factor to indicate relative performance is to look at a pareto by product family, or by priority (an example is shown below). Here again we tend to have an expected ordering, which corresponds to the relative cycle time X-

factor performance that is planned for lots with different priorities. For example, we plan that higher priority lots will spend less time in queue than other lots, so we expect them to have a lower Dynamic X-Factor. If this changes for some reason, it may indicate dispatch compliance issues or other problems.

Summary

Although we don't know the exact distribution of Dynamic X-Factor, it can still give useful clues about how the fab is running. Looking at fab-wide Dynamic X-Factor values by hour helps us to a) identify systematic problems such as lack of shift change coverage; and b) get an early indication of when cycle time is drifting upwards. Generating a Dynamic X-Factor pareto by module, or by priority, can also give useful information, especially when the relative ordering of the values is different from what we expect. Dynamic X-Factor remains relatively easy to calculate (provided you can automate the calculations), and requires no up-front data about theoretical cycle times. After looking at Dynamic X-Factor values for actual fabs, we continue to think that it is a useful metric, and we recommend that you consider using it.

Closing Questions for FabTime Subscribers

Have you tried measuring dynamic x-factor for your fab? If so, have the results told you anything interesting about how your fab operates? Have you run into any problems with this metric? Have you compared it with Cycle Time X-Factor values for your fab?

Further Reading

- S. Johnishi, K. Ozawa and N. Satoh, "Dynamic X-Factor Application for Optimizing Lot Control for Agile Manufacturing," *Proceedings of the 2002 International Symposium on Semiconductor Manufacturing (ISSM2002)*, Tokyo, Japan, 2002. This is the paper that introduced the metric Dynamic X-Factor.

- F Chance and J. Robinson, "Dynamic X-Factor," *FabTime Cycle Time Management Newsletter*, Volume 4, Number 8, 2003. This issue is available for purchase from FabTime's Amazon zShop, at www.Amazon.com/shops/fabtime, by following the link for "Volume 4 – Single Issues".

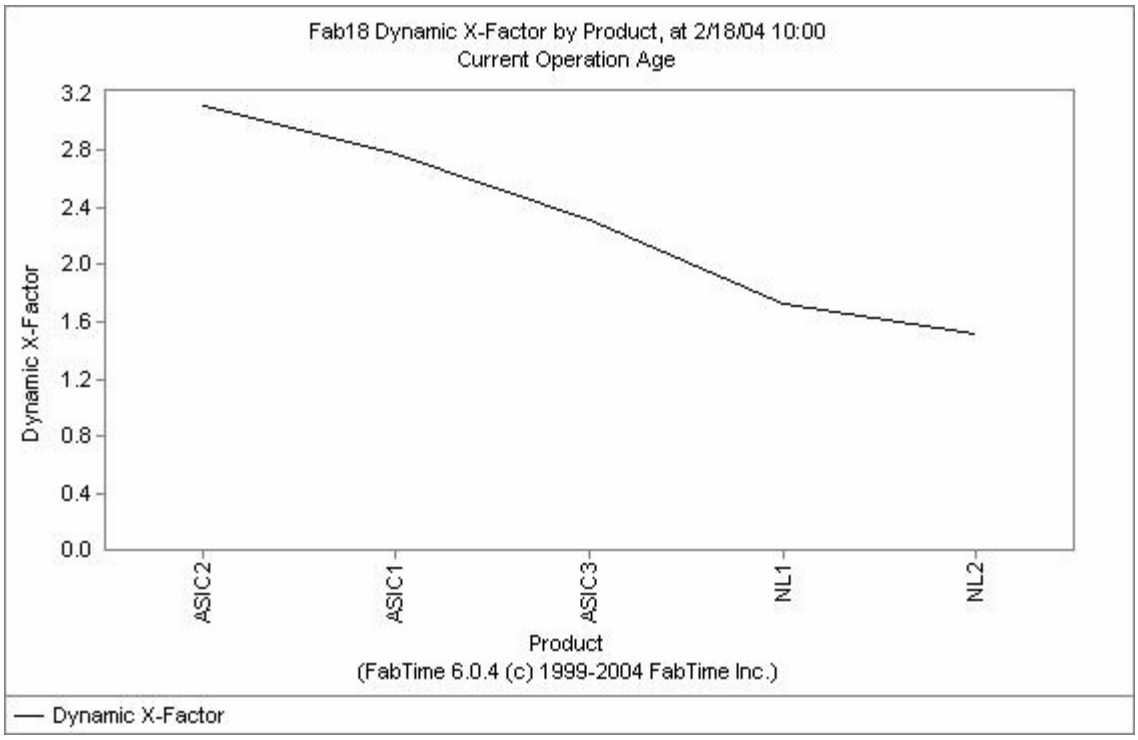

Subscriber List

Total number of subscribers: 1551, from 386 companies and universities. 27 consultants.

Top 10 subscribing companies:

- Intel Corporation (69)
- Motorola Corporation (57)
- Analog Devices (51)
- Infineon Technologies (47)
- Philips (46)
- STMicroelectronics (46)
- Micron Technology, Inc. (42)
- Seagate Technology (42)
- Advanced Micro Devices (36)
- Texas Instruments (36)

Top 3 subscribing universities:

- Arizona State University (11)
- Virginia Tech (9)
- Technical University of Eindhoven (7)

New companies and universities this month:

- Agile Materials & Technologies, Inc.
- Alpha & Omega Semiconductor
- Apple
- Aquaria, Inc.
- BioSource International
- Clayton Consulting
- Colibrys
- Entegris, Inc.
- Ettore Products Company
- Lam Research
- Pacific National
- Semiconductor Equipment Corp.
- Siltronic Corporation
- Structural Integrity
- Technical University Berlin
- Western Digital Corporation

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.
325M Sharon Park Drive #219
Menlo Park, California 94025
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

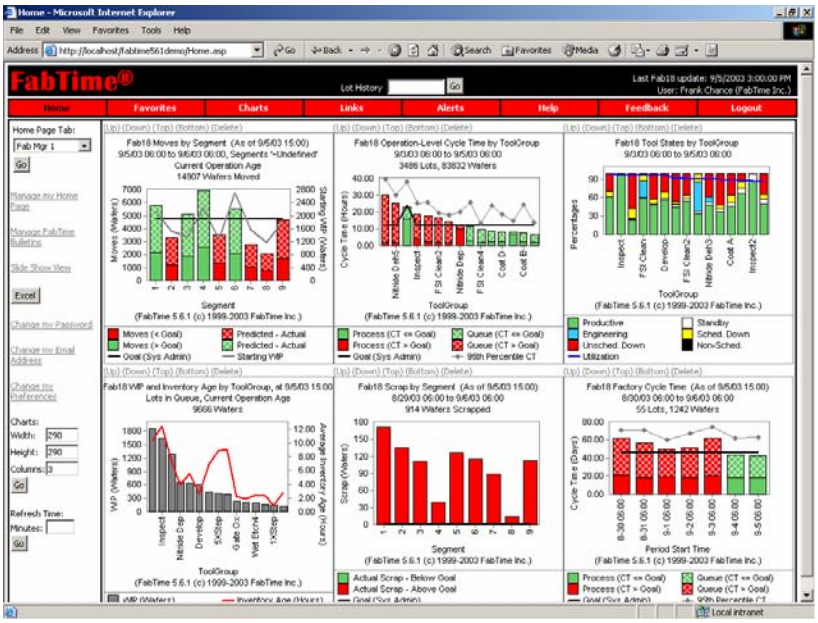

FabTime Benefits

- 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%.