FabTime Cycle Time Management Newsletter

Volume 7, No. 7 August/September 2006

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 (2006) include the capability to specify more flexible goals, and the capability to use minutes as the output unit of measure for time-based charts.

Editor: Jennifer Robinson

Contributors: MS Ham (Samsung)

Table of Contents

Welcome

Welcome to Volume 7, Number 7 of the FabTime Cycle Time Management Newsletter! We're clearly in the midst of the last burst of summer holidays, and we have very little in the way of subscriber discussion or community announcements. Newsletter subscriber MS Ham did write in response to last month's issue to provide a brand new reference on making dispatching decisions that account for downstream batch efficiency. Our development team has also been busy, and our FabTime software user tip of the month concerns a new method for configuring flexible goals.

In our main article this month we provide some ammunition to the many people who have asked us for help in estimating financial benefits from cycle time improvement efforts. The dollar benefits from cycle time improvement are not as easy to estimate as those from utilization improvement. However, there are several clear benefits that stem from variability reduction and cycle time improvement. If we can reduce variability in the fab, we have the option of squeezing the existing capacity buffer, and getting some extra throughput out of the same toolset. Alternatively, if we reduce variability in the fab, we can reduce cycle time. Cycle time reduction is tied to several other benefits: improved line yield, decreased WIP carrying cost, decreased cost of engineering change notices, decreased risk of obsolete inventory, and increased revenue from time to market pricing premiums. In this article, we review the first two of these benefits in detail, with numerical examples, and include highlights of the other three. We hope that you find this article useful in justifying and motivating your cycle time improvement projects.

Thanks for reading!—Jennifer

FabTime

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

Community News/Announcements

Internet Demos of FabTime's Digital Dashboard Software

Several people who have been receiving FabTime's newsletter have asked for more information about our web-based digital dashboard software for cycle time management. For those of you who haven't seen it, here is a brief overview, and an opportunity to arrange to see a full version of the software (with demo data).

FabTime helps wafer fab managers get the information that they need, in real-time, to run their fabs effectively, without having to go through the typical cycles of report customization with IT staff. FabTime extracts updated information from your MES every one to five minutes, and presents the information to your staff in an easy to use web-based format. It saves fab managers time by quickly highlighting problems, and by allowing publication of site-specific performance metrics through dashboard and slide show interfaces. You can also use the software to set up alerts, with which the software will notify you via email or text pager if some condition is met, allowing your staff to be notified automatically about critical fab issues. FabTime also includes sophisticated lot dispatching capabilities, complete with flexible, custom dispatch rules, dispatch performance monitoring, and a downstream tool reservation system.

Using the FabTime software, you can improve cycle time in your fab by: identifying sources of variability (which drive up cycle time); more closely tracking individual lots, and their performance to schedule; and improving your dispatching strategy. These improvements can help your fab to cut costs, and to make more money by squeezing additional capacity out of your existing toolset. Improving cycle time can also, as discussed in this month's main article, help your company to increase revenues, because you can charge a premium for products that get to market more quickly. Existing FabTime customers (we are just starting work on our 10th customer site) have reported cycle time improvements on the order of 10-20%, along with significant savings in managers' time. These improvements mean that the payback period for FabTime's software will most likely be less than six months for your fab.

If you work in a wafer fab (or related high-tech manufacturing facility) and would like to see the software in action, please contact me at Jennifer.Robinson@FabTime.com. I'm happy to arrange for an Internet demo with you or, if you are located in the U.S., to arrange to visit your site and show you the software in person. I think that you'll be impressed. Thanks for your interest!

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

FabTime User Tip of the Month

Configure Goals to Use the New "(any)" Filter

Up to now, whenever you have wanted to see a goal line on a chart in FabTime, you have had to be sure that the filters defined on the chart exactly matched the filters set up by your system administrator for a particular goal. The motivation behind this was that when you apply filters to a chart, the goal that ends up displayed should be relevant to the filtered chart. However, this method resulted in the need for duplication of goals in some cases. For example, suppose that you have a 50 day shipped lot cycle time goal that is the same for several different products, and you typically generate charts for one product at a time. Your system administrator will need to set up separate 50-day goals for each product, so that when you filter a chart for one of the products, the appropriate goal line is displayed.

FabTime now (once a new patch is installed for your site) allows the system administrator to set up a goal that lists "(any)" in one or more of the chart filters. So, in the example above, your system administrator would set up a goal for 50-day cycle time, and put "(any)" in the product filter. You could then generate shipped lot cycle time charts, and have the goal of 50 days displayed, regardless of which product or products you chose to list in the product filter (provided that the other filters on the chart matched those configured for the goal). You can still have other, more specific cycle time goals that apply for certain products. The "(any)" goal will only be used if no other matching filter applies for that field.

For example, your site might have two shipped lot cycle time goals, as follows:

Goal: 50, Filters: Products=(any)
Goal: 60, Filters: Products=XYZ

If you create a chart that is filtered to display only product XYZ, the goal line of 60 days will apply. If you create a chart that is filtered to display product JKL, or products A, B, and C, the goal line of 50 days will apply. The "(any)" filter will also apply in cases where you don't specify any product filter on your chart (i.e., you display the average cycle time across all products).

This "(any" filter can be used on all goal filters, not just the cycle time filter. Your site can use the "(any)" filters to set up default goals which will display whenever no more specific goal applies. We think that this will greatly enhance your ability to display goals in the software.

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

Pre-Furnace Process Batching Optimization

In response to last month's question about making dispatching decisions that account for downstream batch efficiency, specifically in reference to wet benches, newsletter subscriber MS Ham from Samsung wrote to let us know of a paper that he recently had published in the 2006 Advanced Semiconductor Manufacturing Conference. Here is the reference:

Myoungsoo Ham; Raiford, M.; Dillard, F.; Risner, W.; Knisely, M.; Harrington, J.; Murtha, T.; HyungTae Park, "Dynamic Wet-Furnace Dispatching/Scheduling in Wafer Fab," *Proceedings of the 2006 Advanced Semiconductor Manufacturing Conference (ASMC 2006)*, page(s): 144-147.

FabTime welcomes the opportunity to respond to your questions, and to include your responses to topics raised in previous newsletter issues. Simply send your questions or comments to newsletter@FabTime.com

Financial Justification for CT Improvement Efforts

Have you ever wondered:

- How do I cost-justify cycle time improvement efforts?
- What is the impact on the bottom line from cycle time improvement?
- What difference does having better cycle times make? Shouldn't we just focus on throughput?

These are the kinds of questions that newsletter subscribers and prospective customers ask FabTime on a regular basis. We've addressed dollar benefits from cycle time improvement in two previous newsletter issues (Issue 2.6 and Issue 3.5), but thought that we could all benefit from a fresh, nuts-and-bolts look at the question, with numerical examples.

This topic follows naturally on the heels of last month's article (Issue 7.06) about a fundamental conflict in wafer fabs: the pressure to simultaneously increase tool utilization and decrease cycle time. As we discussed, fabs are under cost pressure to increase utilization, so that they can get more throughput from the same toolset. At the same time, there is pressure to reduce cycle times, to please customers and introduce new products quickly. These two pressures are at odds with one another, because cycle time tends to increase with increasing equipment utilization. What makes this conflict particularly difficult to balance is the fact that there is no one straightforward equation for quantifying the dollar value of cycle time reduction. So, on the one hand, we have a clear financial benefit that is tied to utilization increases. On the other hand, the financial benefit from cycle time reduction is much murkier. However, there are some relationships that we can quantify, as discussed below.

Increasing Throughput through Improved Management of Cycle Time

As we have discussed previously in this newsletter, every fab has an operating curve, which is the graph of cycle time x-factor (cycle time divided by theoretical process time) vs. fab utilization percentage. The operating curve generally looks like a hockey stick. It starts out low and flat, at low utilization values, and then increases rapidly and non-linearly at higher utilizations. When the fab utilization (generally defined as the utilization of the bottleneck), approaches 100%, the cycle time gets very large. This is because the bottleneck doesn't have any catch-up capacity, and once a queue starts to build up, there's no way to ever work that queue off. To avoid this, most fabs plan their capacity such that the bottleneck tool group (the tool group with the highest utilization) is loaded to no more than 85% or 90% of the maximum amount that could be run on the tools. The remaining 10%-15% is called spare capacity, catch-up capacity, slack capacity, and other names. But the idea is to provide a buffer to keep cycle times from getting out of hand. Other tool groups in the fab have the same buffer, or one that is even larger. This allows a fab to avoid the steepest part of the operating curve.

As we have also discussed, the exact shape of a fab's operating curve is heavily dependent on the amount of variability in the fab. That is, the more sources of variability a fab demonstrates, the higher the cycle time will be at a given utilization, and the higher the curve will appear. Here we're talking about variability in how lots arrive to tools (do they arrive at evenly spaced intervals, or in burst?), and in how lots are processed at tools (is the rate at which lots leave a tool consistent, or do we sometimes have downtime and setups and operator delays that make things more variable?). If we can do things in the fab to better manage cycle times, by reducing the amount of variability, we can actually move the fab onto a different, more favorable, operating curve. This gives us two choices:

1. Reduce cycle time, while maintaining the same throughput rate (by moving straight down from the old operating curve to the new, lower one).

2. Increase the throughput rate, while maintaining the same cycle time (by moving across horizontally to the new operating curve).

Let's look at that second option again. By reducing variability in the fab (by better managing operation cycle times) we can choose to increase throughput slightly, and still get the same overall fab cycle time. If we manage variability well, we can reduce the size of the capacity planning buffer described above. This means that we can get more throughput out of the same toolset. The financial benefit from increased throughput in an existing toolset is straightforward to quantify.

*Current Annual Profit * Percent Increase in Fab Throughput = Annual Profit Increase*

The above formula assumes that your fab will be able to sell the additional wafers produced at approximately the same profit margin.

Throughput Increase Example:

Here is a simple example, which you can adjust to make relevant for your own fab:

Assume a capital equipment base of $100 million and five-year straight-line depreciation. (Substitute here for whatever is relevant for your fab).

In the above case, the fab must generate at least $20 million annually in profit simply to cover depreciation. (Again, the idea is to use a ballpark estimate of annual profits, regardless of whether you look at the depreciation, or some other index to get the number).

If a cycle time improvement project allows you to squeeze your capacity buffer, and increase your wafer ships by 1% (with the same equipment), and your fab sells them at the same profit margin, then your profit increases by $20M * 0.01 = $200,000/year = $16,667/month.

Financial Benefits from Cycle Time Reduction Efforts

Instead of using variability reduction to drive throughput increases, you can also, of course, choose to reduce cycle time in your fab. The financial benefits that you may observe from this change include:

1. Improved Line Yield
2. Reduced Cost of Carrying WIP
3. Reduced Cost of Engineering Change Notices
4. Reduced Risk of Writing Off Obsolete Inventory
5. Increased Revenue Due to Pricing Premiums from Getting Products to Market More Quickly

We will discuss the first two of these in detail, as these are the easiest to agree upon and quantify for many fabs, and will briefly discuss the other three. You can find most of the methods discussed here included in FabTime's Bottom Line Benefits Calculator, a free Excel tool available for download from www.fabtime.com/-bottomline.shtml. See especially the formulas outlined on the *Details* page.

1. Improved Line Yield:

It is generally accepted that the longer a lot is in the fab, the higher the probability that the lot will encounter some sort of yield problem. While there are no hard and fast numbers here, you can make an assumption that seems reasonable for your fab. For example, a 10% cycle time reduction might correspond to a 0.5% increase in line yield. Improving line yield means that you can either start fewer wafers to obtain the same number of outs, or produce some additional good wafers out.

Starting fewer wafers results in a straightforward savings in raw material cost. Also, if you start fewer wafers, you get a small decrease in bottleneck utilization, which in turn will tend to lead to further cycle time improvements. These cycle time improvements may lead to further yield improvements, and your fab will experience a positive improvement cycle. Looking only at the savings from decreased raw material requirements, we have:

*Cost Reduction Due to Yield Improvement = (Previous Number of Wafer Starts per Year - Revised Number of Wafer Starts per Year) * Raw Wafer Cost.*

Example:
Suppose that you currently start 1000 wafers per week, and have a 95% line yield. This means that your throughput rate is 950 good wafers out per week. If you improve line yield by 0.5%, to 95.5%, then you only need to start 950/0.955 = 994.8 wafers/week. This means that there are five wafers per week that your fab will not need to start. This multiplies out to 5*50 = 250 wafers per year. At a raw wafer cost of $40/wafer, this is a savings of $10,000.

Alternatively, if cycle time improvement leads to a yield improvement, your fab may elect to simply get more wafers out, for the same start rate. The additional cost for getting these wafers out is very small (mostly extra consumables). So, you can estimate the increased revenue as:

*Revenue Due to Yield Improvement = (Previous Number of Wafer Outs per Year - Revised Number of Wafer Outs per Year) * Selling Price/Wafer*

Example:
Using the above example, if the fab initially gets 950 wafers out per week, and yield is improved to 95.5%, then the fab will start to get out 955 wafers per week, or 5 additional wafers per week. This multiplies out to 250 wafers per year. If your selling price is $1000/wafer, this is a revenue boost of $250,000/year.

Note that your fab has a choice of one yield improvement benefit or the other here. You can either start fewer wafers, to get the same throughput, or you can get more wafers out, at the same start rate. If you choose the latter, there may be a slight negative impact on cycle time. This is because for the operations later in the line, the extra wafers that have not been scrapped will cause a slight utilization increase. This may result in some circular behavior by which cycle time goes down, fewer wafers are scrapped, utilization increases on back-end tools, and then cycle time increases slightly. However, it is unlikely that cycle time will increase to your previous level, and you will in any event have the revenue from the extra outs to make up for this inconvenience.

2. Reducing WIP Carrying Cost

Another dollar benefit from reducing cycle time comes from reducing the cost of carrying the WIP in your fab. Because your fab has inventory sitting on the floor in the form of partially processed wafers, your company is not able to use the value of that WIP for other investments. If you reduce the WIP in the fab, you reduce the associated carrying cost of the WIP, and you free up some money for other things. This is a one-time benefit, but can be substantial.

WIP Carrying Cost Reduction = Original WIP Carrying Cost - Revised WIP Carrying Cost

where

*WIP Carrying Cost = Average WIP * Mid-Line Value per Wafer * Company's Internal Cost of Capital*

and we know from Little's Law that

*Average WIP = Start Rate * Cycle Time * Yield Correction*

where the standard yield correction is:

Yield Correction = (1 + Line Yield) / 2

(This assumes that scrap occurs linearly across the line)

Example
Suppose, using some of the numbers from above, that a fab starts 1000 wafers per week, has a cycle time of 5 weeks, and has a line yield of 95%. Little's Law tells us that the average WIP in this fab will be 1000 wafers/week * 5 weeks * (1+0.95)/2 = 4875 wafers. We said earlier that the raw

wafer cost was $40, and that the per wafer selling price was $1000. This means that the average value of each wafer in WIP (assuming that value accrues linearly throughout the line) is (1000+40)/2 = $520/wafer. This means that the average value of the WIP in the fab is 4875 wafers * $520/wafer = $2,535,000. If the fab's internal rate of return for investments is 15%, then the cost of carrying this WIP, at any point in time, is $380,250.

If this fab makes improvements to reduce the cycle time by 20%, from 5 weeks to 4 weeks, then the WIP in the fab decreases by the same 20%, and we have a revised average WIP of 3900 wafers (neglecting any possible changes to the yield rate). The average value of this WIP is 3900*$520 = $2,028,000, and the cost to the fab of carrying the WIP is $304,200.

If this fab can reduce cycle time by 20%, then the cost of carrying the WIP will decrease from $380,250 to $304,200, a one-time difference of $76,050.

3. Reduced Cost of Engineering Change Notices

The idea here, as with the line yield discussion, is that the longer your WIP is in the fab, the greater the risk of it being subject to an engineering change notice (ECN). Usually some percentage of WIP is subject to ECNs. Therefore, if you have the numbers available, you can compare the current cost of ECNs for your fab to a likely ECN cost if the average WIP decreases by some percentage.

4. Reduced Risk of Writing Off Obsolete Inventory

Most fabs make at least some of their WIP to stock, providing a safety stock amount as a buffer against uncertainty in planning. The problem with this is that the more WIP you have in your safety stock, the greater the probability that the industry will take a downturn, and you will have to write off some of that WIP. There are formulas for estimating the required safety stock, given a fab's cycle time. By reducing cycle time, your fab can afford to hold less safety stock, decreasing your risk of writing off WIP. You can find more details in Issue 3.5 (email newsletter@fabtime.com for a copy).

5. Increased Revenue Due to Pricing Premiums from Getting Products to Market More Quickly

Increased sales revenue is potentially the highest lever, in terms of justifying the cost of cycle time improvement efforts. However, it is also the lever that is hardest to quantify, and get people to agree to, because it requires assumptions about what customers will pay in the future. However, as a very simple example, suppose that by reducing cycle time by 20%, your company can get a key new product to market more quickly, and can charge a 30% price premium (for some limited time). You can easily do the math to estimate what that might be worth for your company.

In the semiconductor industry, selling prices drop rapidly over time, with the peak price being charged when a product first comes to market. This is particularly pronounced in the memory chip market. There is a published paper by Robert Leachman (reference below) about a cycle time improvement project that: "reduced manufacturing cycle times to fabricate dynamic random access memory devices from more than 80 days to less than 30. Considering the decline of selling prices for dynamic random access memory devices, (the project) enabled Samsung to capture an additional $1 billion in sales revenue compared to the revenue it would have realized had cycle times not been reduced." Several potential methods of quantifying the increased revenue from cycle time reduction are outlined in Issue 3.5 (email newsletter@fabtime.com for a copy).

Conclusions

The financial benefits from cycle time improvement are not as easy to estimate as the dollar benefits from utilization improvement. However, there are several clear and quantifiable benefits that stem from variability reduction and cycle time improvement. If we can reduce variability in the fab, we have the option of squeezing the existing capacity buffer, and getting some extra throughput out of the same toolset. This has a clear financial benefit, as outlined above. Alternatively, if we reduce variability in the fab, we can reduce cycle time. Cycle time reduction is tied to several other benefits: improved line yield, decreased WIP carrying cost, decreased cost of engineering change notices, decreased risk of obsolete inventory, and increased revenue from time to market pricing premiums. In this article, we have reviewed the first two of these benefits in detail, with numerical examples, and included highlights of the other three. We hope that you find this article useful in justifying and motivating your cycle time improvement projects.

Closing Questions for FabTime Subscribers

Do you think that these calculations are reasonable? Would you like to see examples for the other methods not described in detail here? How does your fab quantify dollar benefits from cycle time improvement?

Acknowledgement

We would like to thank Ken Beller and Stuart Carr for past discussions on these topics, which have contributed to our understanding in this area. Special thanks to Ken also for suggesting that we write this current article.

Further Reading on Cost Analysis for Wafer Fabs

■ F. Chance, S. Carr, and K. Beller, "What Is One Day of Cycle Time Reduction Worth?", *FabTime Newsletter*, Volume 2, No. 6, 2001. Because this past newsletter issue is directly related to the topic of the current issue, we are offering to provide it to you at no cost. If you would like a copy of Issue 2.6, simply email newsletter@FabTime.com to request it. This offer will be in place for one month following publication of Issue 7.07

■ J. Robinson and F. Chance, "The Bottom Line Benefits of Cycle Time Management," *FabTime Newsletter*, Volume 3, No. 5, 2002. Because this past newsletter issue is directly related to the topic of the current issue, we are offering to provide it to you at no cost. If you would like a copy of Issue 3.5, simply email newsletter-@FabTime.com to request it. This offer will be in place for one month following publication of Issue 7.07

■ R. J. Baseman, W. Grey, S. J. Hood, C. A. Kovac, and R. C. Brilla, "Cycle Time Driven Inventory Cost Analysis," *Proceedings of the IBM International Manufacturing Productivity Symposium*, IBM East Fishkill, New York, October 12-15, 1993.

■ R. Carnes and M. Su, "Long Term Cost of Ownership: Beyond Purchase Price," *Proceedings of the IEEE/SEMI International Semiconductor Manufacturing Science Symposium*, 39-43, 1991.

■ R. C. Leachman, J. Kang, and V. Lin, "SLIM: Short Cycle Time and Low Inventory in Manufacturing at Samsung Electronics," *Interfaces*, Vol. 32, No 1, 2002

■ N. G. Pierce and T. Yurtsever, "Value-Based Dispatching for Semiconductor Wafer Fabrication," *IEEE 2000 Advanced Semiconductor Manufacturing Conference (ASMC '00)*, 245-249, 2000.

■ R. Plieninger, U. Muller, H. Ehm, and W. Reczek, "Cost Reduction using Systematic Target Setting of the Reference Fab Methodology," *IEEE 2001 Advanced Semiconductor Manufacturing Conference (ASMC '01)*, 17-20, 2001.

■ R. Sandell and N. G. Pierce, "A Hierarchical Approach to Cost Analysis for Next Generation Semiconductor Processes," *Proceedings of the International*

Conference on Modeling and Analysis of Semiconductor Manufacturing (MASM 2002), Editors G. T. Mackulak, J. W. Fowler, and A. Schoemig, Tempe, AZ, April 10-12, 2002. 169-174.

- C. K. G. Thing, Y. S. Chang, and C. B. Bezant, "Manufacturing Cost Analysis System in Semiconductor Enterprises," *International Journal of Industrial Engineering – Theory, Applications, and Practice*, Vol. 8, No. 2, 131-141, 2001.
- See also Wright Williams & Kelly's Cost of Ownership Bibliography (www.wwk.com/resources.html).

Subscriber List

Total number of subscribers: 2161, from 420 companies and universities. 23 consultants.

Top 10 subscribing companies:

- Intel Corporation (133)
- Analog Devices (75)
- Atmel (66)
- Micron Technology (65)
- Infineon Technologies (64)
- Freescale Semiconductor (58)
- STMicroelectronics (56)
- Philips (52)
- Texas Instruments (51)
- TECH Semiconductor (46)

Top 3 subscribing universities:

- Virginia Tech (11)
- Arizona State University (7)
- Ben Gurion Univ. of the Negev (7)

New companies and universities this month:

- Ipek Kagit
- Net Profits Inc.
- Records RSA
- Semtech
- Soitec

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 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. FabTime also now includes **lot dispatching** (via dispatch rules) and static **capacity planning**.

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

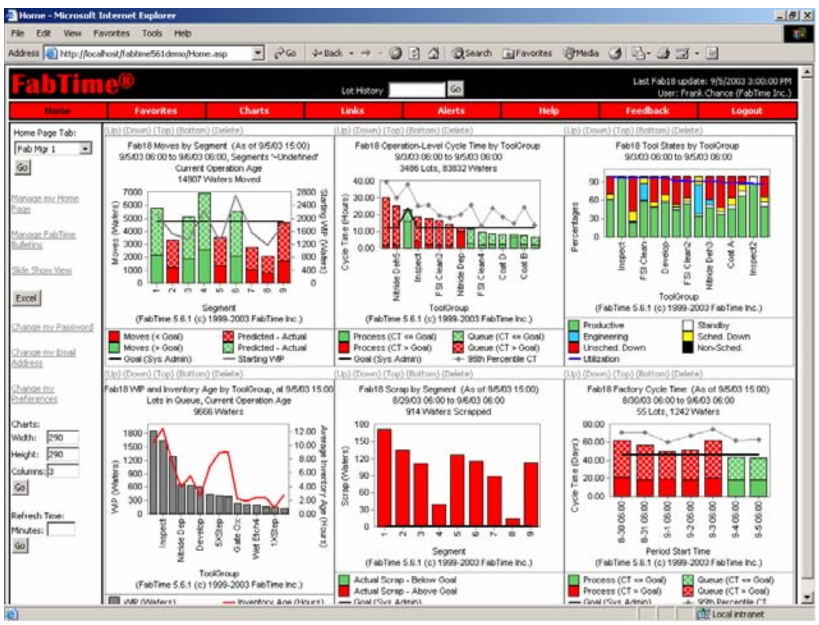

FabTime Benefits

- 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%.