FabTime Cycle Time Management Newsletter

Volume 12, No. 3 May/June 2011

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 software this month include increased support for maximum number of site-specific WIP attributes and site-specific cap on the maximum number of data table rows displayed.

Editor: Jennifer Robinson

Contributors: Bob Kotcher (Philips Lumileds)

Table of Contents

Welcome

Welcome to Volume 12, Number 3 of the FabTime Cycle Time Management Newsletter! In this issue, we begin with a call for papers for the ISMI Symposium on Manufacturing Effectiveness. Our FabTime user tip of the month is about using a PowerPoint add-in to display live FabTime charts (mixed with other content) on monitors. In our subscriber discussion forum we have inputs on analyzing staffing productivity, embracing the downturn, and scheduling in the lithography area.

In our main article this month we discuss the application of queueing models to wafer fabs. We begin by outlining the benefits and drawbacks of queueing models (as compared with static models and with simulation). We then discuss toolgroup-level models, as implemented in FabTime's operating curve spreadsheet, as well as different approaches for constructing fab-level models. We conclude by discussing the simplified approach of using aggregated fab-level inputs in a simple G/G/c queueing model, and where this approach might, and might not, be useful. If any readers would care to share their experiences in applying queueing models to fab planning or operations, we will post those in a followup article. 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

Call for Papers: ISMI Symposium on Manufacturing Effectiveness

There's only a little time left to submit abstracts before the June 1st deadline, but this is an excellent conference, and we wanted to share the announcement with you. If you can't submit a paper, you might still consider attending the conference.

ISMI Manufacturing Week 2011: October 17–21, 2011 • Austin, Texas

Don't miss ISMI's Manufacturing Week — the most informative week of the year for the IC industry. Manufacturing experts, managers and engineers will come together to share their knowledge of efficient manufacturing, factory and equipment productivity improvements, and cost reduction methodologies.

Call for Papers

Increased manufacturing productivity—including advanced equipment and process control—as well as reduced operational costs are absolutely critical for a profitable manufacturing facility. The ISMI Symposium on Manufacturing Effectiveness is the semiconductor industry's most valuable event for exchanging ideas about new cost-saving solutions to help your equipment, your factory and your company to become more productive. Submissions for the ISMI Symposium on Manufacturing Effectiveness are currently being accepted.

For more information, visit the ISMI Manufacturing Week website:

http://ismi.sematech.org/ismisymposium

FabTime welcomes the opportunity to publish community announcements, including conference notices and calls for papers. Send them to newsletter@FabTime.com.

FabTime User Tip of the Month

Use PowerPoint to Display Live FabTime Charts on Monitors

This month's tip comes to us courtesy of Site20. The site recently started using the LiveWeb add-in for Microsoft PointPoint to create mixed slide shows that include FabTime charts as well as other, non-FabTime graphs and charts. They thought that other FabTime sites might be interested in this solution, which is free (provided you already have PowerPoint). Here's the procedure:

1. Install the free LiveWeb Add-In from http://skp.mvps.org/liveweb.htm (follow installation instructions on that page to enable the add-in). This application works with PowerPoint 97 and later (though FabTime has only tested it with PowerPoint 2010). Note that this product is offered by a third-party company, and not directly from Microsoft.

2. Open a PowerPoint file, add a new slide, and select Insert|Web Page (LiveWeb will add this to your Insert Toolbar).

3. Enter the detailed web address for the FabTime chart or data table of interest, including a user login name and password. Note that your site must have Single Page Access allowed for FabTime (refer your site system administrator to this help page: http://fabtime.editme.com/ExternalReference if needed). The username and password will be displayed in clear text. However, links will only work from inside your site's firewall in any case.

4. FabTime charts or data tables will be displayed live when you run a PowerPoint slide show (which you can do during a meeting or from a monitor anywhere inside your firewall).

If you have been using the web screen saver solution outlined at http://fabtime.editme.com/A0034, you can use the same syntax for your chart entries. The difference is that you will also have full PowerPoint capability for formatting and adding other elements (and the LiveWeb application is free, while the web slideshow application required a small license fee).

Thank you Site20 for taking time to share this solution with other sites. 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

Staffing Productivity

Bob Kotcher from Philips Lumileds wrote: Hey, one of the topics you suggested people discuss was "staffing productivity." I presented a paper at WinterSim in 2001 on a simulation analysis I did on staffing levels in our litho area at Headway Technologies (now TDK). In that analysis, I found that increasing our staffing level would be a much cheaper way to increase throughput than buying more equipment. This was surprising given that our operators were pretty lightly loaded. But the machines were very expensive, the operators comparatively cheap, and the machines paused frequently, waiting for quick manual alignments, which meant that even brief waits for operators added up to a big percentage loss in machine throughput. Interestingly, I just finished a similar simulation analysis here at Philips Lumileds, where we make LEDs, only this time I was looking at repair/ maintenance technician staffing levels. I was expecting similar findings to my TDK analysis, but I was surprised to find that the simulation showed that our per-wafer costs would actually be lower if we decreased technician staffing levels. This would add several percentage points of downtime to each machine, but still reduce costs per wafer out. This was because, compared to TDK, these machines are cheaper and have longer service lives, the need for technicians is much less frequent (though of much longer duration), and the technicians are quite a bit more highly paid. It goes to show the value of a simulation model in determining cost-minimizing staffing levels—imagine if a

company did these sorts of analyses throughout its entire fab! The savings could be in the tens of millions.

Embracing the Downturn

An anonymous subscriber wrote: I would like to suggest that your subscribers "Embrace the Downturn". I have been part of the semiconductor industry for 35 years and not once have I been able to convince anyone that downturns should be considered a blessing instead of a time to lay people off and cut budgets. To me a downturn is the time to ensure all operating specs are in order, operators and technicians are current with their training, equipment maintenance and upgrades are performed, and the biggest one of all "expansion of the manufacturing facility". There is no better time to purchase fab equipment at bargain prices and have the time to do the installs properly and well thought out than in the downturn. If your company has sound management (top-down), cash in the bank, and is in the top 5 for the products you produce then you are going to be poised to grab market share from your competition when the recovery happens, and the recovery will happen! If it ain't broke then break it and quit following conventional wisdom "Embrace the Downturn"! That's my 2 cents worth.

FabTime Response: We actually wrote back in newsletter issue 2.3 (in 2001) about ways to improve cycle time during a downturn. We said that a downturn offered at least one potential benefit: time to think. So, we do tend to agree with this subscriber (although we're hoping that the next downturn is still a ways off).

Benefits of Scheduling in the Lithography Area

Another anonymous subscriber submitted a question about the benefits of scheduling. He said:

"I am the owner of my fab's simulation model. We are considering adopting a concept of mathematical based scheduling for the Litho area at our fab. The model's final output should be a plan how to run Litho tools:

- Which tools to run specific steps from specific products?
- Which lots will run on which tool?

The model should take into consideration all relevant parameters:

- Incoming WIP
- Existing Reticles
- Tool resists
- Number of tools
- Etc.

This is close to the concept that was presented at the 2010 Winter Simulation Conference by Infineon Technologies.

I am contacting you because we wonder what is the benefit from this project? It will required major development resources and changing work methods (which work well now). What is the magnitude of the benefit? How much will tool utilization and cycle time improve? I would be happy to get some knowledge on that from other sites that have implemented this method.

Could I use the FabTime platform to discuss this issue?"

FabTime Response: We can't answer this question in detail in any abstract sense. The magnitude of the benefit from such a project would depend on where the fab is starting, in terms of cycle time and utilization, and how the fab is already doing in terms of scheduling. We do believe that if one is going to implement scheduling (as opposed to pure dispatching), it makes sense to start with a single area of the fab, like litho, rather than attempt to schedule the entire fab at once.

In our product we've chosen to focus instead on taking information from other parts of the fab into account while making local dispatching decisions. The dispatch decision is still made locally (to order the lots in queue). But it might use information about which lots will feed a downstream

bottleneck, and the current tool status of that downstream bottleneck, etc.

But we would be interested to hear if any other subscribers have feedback to share on the benefits of a more mathematically-based scheduling approach. For more on the difference between scheduling and dispatching, see FabTime Newsletter 6.04 (available on request to current subscribers from newsletter@FabTime.com).

FabTime welcomes the opportunity to publish subscriber discussion questions and responses. Simply send your contributions or questions to Jennifer.Robinson@FabTime.com.

Queueing Models for Wafer Fabs

Introduction

In the last newsletter issue, we raised a series of potential discussion topics and asked for feedback from our subscribers on which ones people would like to see addressed in the newsletter. The topic that generated the most interest was queueing models for wafer fabs. This is a huge and potentially complex topic - there are many dissertations out there on various applications of queueing models for fabs (including Jennifer's). In this article, we offer an introduction to the use of queueing models for wafer fabs. We welcome your feedback.

Background: Where Queueing Models Fit In

Queueing models are mathematical models that predict the long-term, steady-state behavior of a system. Unlike static models (such as simple spreadsheet calculations), queueing models can predict dynamic outcomes such as cycle times and WIP levels. Unlike simulations models (in which a series of potential outcomes are mapped in detail, and the results averaged), the outcome of a queueing model is not a range of possible results. The outcome from a queueing model is a single value.

Queueing models are equations, rather than the samples used in simulation models, and are thus very quick to run. They can be embedded in spreadsheets, or coded into other programs. This combination of speed, clear outcomes, and ability to estimate dynamic performance makes queueing models very appealing for predicting fab performance.

However, there are significant drawbacks to the use of queueing models for fab analysis. They are mathematical models built on a series of assumptions. While they can be very accurate for modeling simple systems (like a single tool with an exponential arrival process), the calculations become prohibitively complex for more sophisticated systems (like a reentrant batch system with operator delays and a mix of scheduled and unscheduled downtimes, for instance). In some cases, approximations can be used. In other cases, no accurate, closed form queueing model exists, and simplifying assumptions must be made about the system being studied. And while interpretation of the results is fairly easy (since the outcome will be a single value), making changes to the detailed calculations may require specialized training.

What these strengths and weaknesses mean is that queueing models have a place in fab analysis, but that they tend to be better at giving relative answers (e.g., by what percentage might we decrease cycle time if we cross-qualify this dedicated step?) than absolute answers (e.g., what is the exact cycle time going to be through the fab for this product line?). We feel that queueing models have more of a place in validating the behavior of individual toolgroups, rather than understanding the behavior of the fab as a whole.

Tool-Level Queueing Models

FabTime has used queueing models in our newsletter and our cycle time management class to help people understand how various factors affect expected dynamic performance of toolgroups. We have collected a series of relevant queueing formulas on our website (with references) at http://www.fabtime.com/formula.htm. We also have a spreadsheet tool that we built that allows one to look at the impact of several variables on the operating curve (cycle time x-factor vs. utilization) of a single toolgroup. Inputs include:

- Number of tools in the toolgroup
- Average processing time
- Coefficient of variation of process times
- Coefficient of variation of time between arrivals
- Downtime (MTBF, percent downtime, and repair time variability)
- Arriving batch size in lots, and variability of batch sizes
- Percentage of hot lots

A simpler version of this spreadsheet is available for download from FabTime's website. The full version is only available to customers of our software or our cycle time management course. A screen snapshot of the full version is shown at the top of the next page.

The output of the spreadsheet is a series of operating curves showing cycle time x-factor at various utilizations, for up to three scenarios. Queueing models underlie the calculations (and are documented on Calculator Details and Notes pages of the spreadsheet). This spreadsheet tool is quite useful for showing the impact of tool utilization on cycle time, as well as the relative impact of the above variables on the cycle time of a particular toolgroup.

Fab-Level Queueing Models

We also used the calculations described above to build what we called a Cycle Time Entitlement Calculator in Excel. The spreadsheet had a row for each step of a process flow. Users could enter the above variables, plus the average tool utilization for the step, and the spreadsheet would use the queueing models to predict the cycle time for each step. These could then be added up across all of the steps in the flow, to predict an overall cycle time.

This approach was much less successful. There were several reasons for this:

- Often a step could be done on tools from several different toolgroups, with different downtime characteristics. The spreadsheet required aggregate inputs across these disparate tools for each step, reducing accuracy.
- Although it allowed for batch arrivals, the spreadsheet didn't account for batch processing at tools.
- The spreadsheet didn't incorporate rework, holds, dispatch rules, or operator delays, among other fab complexities.
- Data collection, particularly for the variability data, was difficult (particularly because of the above-mentioned issue of each step being done by multiple tools).

While some of these issues could have been resolved by researching other approximations, and using the FabTime database for calculating inputs at FabTime customer sites, the central issue was that while the queueing models that we were using were informative at the step level, they simply weren't accurate enough for

FabTime Cycle Time Operating Curve Generator

Calculates regular lot (non-hot-lot) cycle time for multiple-tool machine groups with general arrival and service processes. Each tool has a single failure process, indicated by MTBF and percentage downtime.

Source	Description	Set 1	Set 2	Set 3
Operations	Average process time (hours)	3	2	1
Operations	Hot lot percentage	0%	0%	0%
Operations	Number of tools	1	1	1
Operations	Process time variability	0.5	0.5	0.5
Operations	Inter-batch arrival variability	1	1	1
Operations	Average arriving batch size (lots)	1	1	1
Operations	Arriving batch size variability	0	0	0
Operations	Repair time variability	1	1	1
Operations	Mean time between failures (hours)	24	24	24
Operations	Percentage downtime	15%	15%	15%

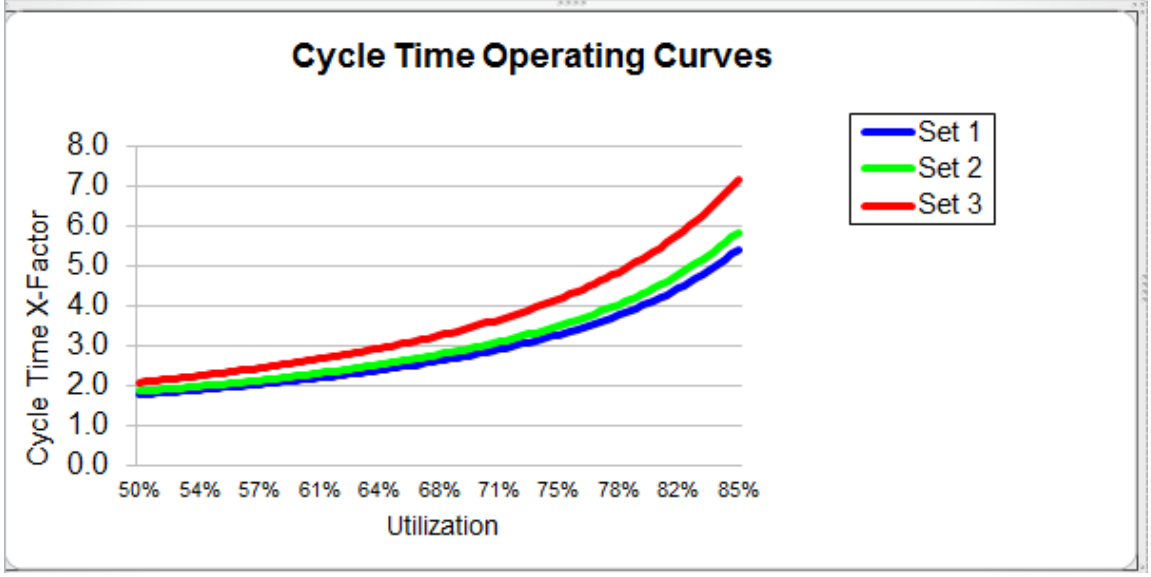

such detailed analysis of a full factory model. The approach of looking at each step separately and adding them up was also a fundamentally inaccurate way to look at the fab as a whole.

Quite a bit of research has been done on queueing network models, which take into account the fact that the outputs from one step are the inputs to the next step, and should be analyzed together. See the references outlined below for some examples. Queueing network models, however, do suffer from the same basic limitations as more localized models – it's difficult to take into account secondary-level effects, like needing to acquire first an operator, and then a tool. We don't know of anyone who is using a queueing network model to analyze their whole fab on a day-to-day basis (though we would certainly be interested to hear from anyone who is).

Isn't There a Simpler Way?

Some of the people we have worked with over the years have asked: well, can't we just simplify this? Can't we use the single tool approximation in your first spreadsheet and just enter overall inputs, to get a rough idea of the operating curve for the fab? The answer to this is a cautious yes. One can come up with average values across the fab, and enter them into the operating curve spreadsheet, and it will show the fab's cycle time increasing with average utilization and overall variability. But for sure you will be looking at relative answers, not absolute answers, because you're glossing over a huge amount of detail.

Still, this can be a reasonable approach if what you want is, say, a way to show that yes, reducing arrival variability will drive down cycle time, and yes, reducing downtime variability will drive down cycle time, etc.

There are some technical issues to consider, if you're trying to come up with overall inputs for the fab to use in the queueing model. You have to decide what to use for average process time (what does that even mean across a fab?) and number of tools per toolgroup. The most likely usefulness of these inputs would be in comparing two different fabs. If you had one fab that had mostly one of a kind tools, and you entered one for average number of tools, you would see a very different operating curve from a fab where you entered four or five.

Even trickier is coming up with downtime and variability inputs. One concrete suggestion that we have is that if you're calculating coefficient of variation of process times, aggregated across the fab, the way to do this is to first calculate the CV for each toolgroup, and then average the CVs. You don't want to simply look at the entire string of process times across the whole fab, and calculate the CV for that set of values. This is because there's a lot of variation between tools. Some have long process times, and some have short process times. This variation isn't a problem. It's only when you have long and short process times happening on the same toolset that you need to worry about that variability. Similarly when looking at arrival variability, you want to look at the CV of arrivals to each toolgroup separately, and then aggregate. Otherwise you'll underestimate the variability (if you try to look at the time between arrivals to any step, regardless of the tool).

One temptation, instead of coming up with some sort of representative number for the fab, is to use the inputs for the bottleneck. By definition, however, the bottleneck has the highest utilization in the fab. It will likely have one of the highest cycle times, too. The x-factor for the fab as a whole is usually quite a bit lower than the x-factor for the bottleneck, because you are able to offset long cycle times at the bottleneck with shorter cycle times at lower utilization tools. Also, the impact of arrival variability to the bottleneck is usually mitigated by the fact that lots arrive to the end of a large queue. Arrival variability can matter less at the bottleneck than it does at lightly utilized tools. For these reasons, the inputs for bottleneck are not representative of the fab's overall operating curve.

A Note on Distributions

The simplest queueing models assume exponential distributions. The exponential distribution has a coefficient of variation of 1.0. Assuming exponential distributions tends to make the math work out well, removing a great deal of the complexity in a queueing model. It's not uncommon to see research that assumes that arrival times in fabs are exponentially distributed, while process times are deterministic (with no variability - which also simplifies things). In our experience (and we measure these variables in FabTime, so we're basing this on actual data), the coefficient of variation of arrival times to toolgroups in wafer fabs is often greater than 1. Sometimes much greater. The coefficient of variation of process times to toolgroups varies, but is usually somewhere between 0 and 1. It's not zero because, at a minimum, you have multiple operations taking place on the same tool, and these can have different process times. What this means is that exponential queueing models are probably not accurate enough to use for understanding wafer fabs. What you want are queueing models based on generalized distributions (G/G/c models), for which you measure and input the arrival and process time variability. Of course these are more complex to apply.

Conclusions

Queueing formulas can be used to estimate manufacturing system performance measures such as average cycle time and throughput. Where available, queueing models offer closed-form solutions that can be easily coded into spreadsheets and other programs. While the complexities of fabs make it difficult to apply queueing formulas to the fab as a whole, queueing models can be quite useful for understanding and validating toolgroup-level behavior. FabTime has coded queueing approximations into a spreadsheet-based operating curve generator, for use in understanding the relative impact of several variables on cycle time. This approach does not scale well for building detailed fab-level queueing models. However, the single toolgroup model can be used, with aggregated inputs, to give a very general idea of overall fab behavior.

If any subscribers would like to share their experiences in applying queueing models to fab planning and management, we will share those in the next newsletter issue.

Closing Questions for FabTime Subscribers

Do you use queueing models for any day-to-day applications in your fab? Do you use toolgroup-level models, or more of a queueing network approach that takes the whole fab into account? Have you found queueing models to be sufficiently accurate for your purpose? Do you have any papers on this topic to share with the newsletter community?

Further Reading

There are hundreds of published articles about the application of queueing models to particular aspects of wafer fabrication. Many of them are university studies based on test data, but some apply to actual fab data. Here are a few that may be of interest:

R. Akhavan-Tabatabaei, S. Ding and G. Shanthikumar, "A Method for Cycle Time Estimation of Semiconductor Manufacturing Toolsets with Correlations," *Proceedings of the 2009 Winter Simulation Conference, Modeling and Analysis of Statistical Methods (MASM) Track*, 2009. Full paper available here.

F. Chance. *The Factory Explorer User Manua*l. Wright Williams & Kelly. 1999.

H. Chen, M. Harrison, A. Mandelbaum, A. Van Ackere, and L. Wein, "Empirical Evaluation of A Queueing Network Model for Semiconductor Wafer Fabrication," *Operations Research*, Vol. 36, No. 2, 202-215, 1988.

D. Connors, G. Feigin, and D. Yao, "A Queueing Network Model for Semiconductor Manufacturing*," IEEE Transactions on Semiconductor Manufacturing*, Vol. 9, No. 3, 412-427, 1996.

W. Scholl, E. S. Gel, K. Khowala, and J. W. Fowler, "Use of Analytical Queueing Approximations to Set Processing Step Performance Targets at Infineon Technologies Dresden," *Proceedings of the International Conference on Modeling and Analysis of Semiconductor Manufacturing (MASM 2002)*, Tempe, AZ, April 10-12, 2002. 206-210.

K. Wu, . F. McGinnis, and B. Zwart, "Queueing Models for Single Machine Manufacturing Systems with Interruptions*", Proceedings of the 2008 Winter Simulation Conference*, Miami, FL, December 7-10, 2008. Full paper available here.

H. Zisgen, I. Meents, B. R. Wheeler, and T. Hanschke, "A Queueing Network Based System to Model Capacity and Cycle Time for Semiconductor Fabrication", *Proceedings of the 2008 Winter Simulation Conference*, Miami, FL, December 7-10, 2008. Full paper available here

See other references from FabTime here and here.

Subscriber List

Total number of subscribers: 2787, from 453 companies and universities.

Top 20 subscribing companies:

- Maxim Integrated Products, Inc. (175)
- Intel Corporation (146)
- Micron Technology, Inc. (104)
- GLOBALFOUNDRIES (96)
- Western Digital Corporation (70)
- Carsem M Sdn Bhd (67)
- X-FAB Inc. (67)
- Texas Instruments (66)
- International Rectifier (63)
- TECH Semiconductor Singapore (61)
- ON Semiconductor (59)
- STMicroelectronics (56)
- Analog Devices (53)
- Freescale Semiconductor (53)
- IBM (49)
- NEC Electronics (46)
- Skyworks Solutions, Inc. (45)
- Infineon Technologies (41)
- Cypress Semiconductor (37)
- Seagate Technology (35)

Top 4 subscribing universities:

- Ecole des Mines de Saint-Etienne (EMSE) (12)
- Arizona State University (8)
- Ben Gurion Univ. of the Negev (8)
- Nanyang Technological University (8)

New companies and universities this month:

- Cisco
- HLI-HUME Mgmt Co Sdn Bhd
- Portland State University
- Ralls Construction Corporation
- Transform Inc.

Sampler Set of Other Subscribing Companies and Universities:

- Adams Associates (1)
- AltF5 Software LLC (2)
- Aviza Technology (1)
- Booz Allen Hamilton (1)
- China Electronics Engineering Design Institute (1)
- Comlase AB (1)
- Crocus Technology (1)
- Harris Stratex (1)
- Hynix-ST (1)
- Integrated Technologies Company (2)
- Microchip Technology (24)
- Mikron Corporation (1)
- Panasonic Semi. Singapore (2)
- PRTM (2)
- Raytheon (10)
- Suss MicroOptics (1)
- Sygentrics (1)
- Tara Technologies (1)
- TDK (6)
- University of Alabama – Huntsville (1)

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
Spansion Fab 25

FabTime Subscription

One low monthly price includes

- Software installation and real-time connect to your MES
- End user and system administrator training
- Unlimited users via your Intranet.
- Software maintenance and regular upgrades (approx. 6 per year, via our no-downtime patch system)
- Add-on dispatching and planning module for a slightly higher monthly fee

Interested?

Contact FabTime for technical details and/or a web-based demonstration.

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

Turn fab MES data into information and save time and money

- 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 can help. FabTime saves your management team time daily by turning fab MES data into information, via a real-time web-based dashboard that includes lot dispatching. FabTime saves your IT staff time by breaking the cycle of custom-developed reports. With FabTime, the end user can filter for exactly what he or she needs, while staying in a comprehensive framework of pre-defined charts. Most importantly, FabTime can help your company to increase revenue by reducing cycle times up to 20%.

"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 cycle times by up to by 20%.
- Focus improvement efforts on the tools that inflate cycle time.
- Improve supervisor productivity – cut reporting time by 50%.
- Improve IT productivity – eliminate need for custom reports.