FabTime Cycle Time Management Newsletter

Volume 17, No. 5 October 2016

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 development for FabTime include support for custom X/Y charts for correlation analysis (e.g. cycle time vs. utilization) and ability for users to set preferences for colors in stacked charts.

Editor: Jennifer Robinson

Contributors: Richard Davis (Honeywell)

Keywords: WIP Turns, Metrics and Goals

Table of Contents

Welcome

Welcome to Volume 17, Number 5 of the FabTime Cycle Time Management Newsletter! In this issue we have one announcement, about FabTime team members running the SLO Ultra. Our software tip of the month is about copying chart images to other applications (a modified tip that reflects changes due to our custom JavaScript charting engine). We have a subscriber discussion question regarding the choice between one automatically loaded tool and two manually loaded tools. Long-time readers will be able to guess FabTime's thoughts on that question.

In our main article this month we look at computational issues in reporting and calculating WIP-related metrics, particularly WIP turns. If there is one thing that we've learned in our years of working with fab data, it's that even things that seem straightforward can become complex, when you get down into the nitty-gritty details. Computing average WIP is no exception. As always, 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

Frank Chance, Beth Chance, and Mike Krist run the SLO Ultra @ Wild Cherry Canyon

The inaugural SLO (San Luis Obispo, CA) Ultra — a trail-running running event that includes a 50-mile ultra marathon, marathon, half marathon and 5-mile hike — was run September 10 at Wild Cherry Canyon in Avila Beach. The event, to benefit The Land Conservancy of San Luis Obispo County, was sold out with 1,200 runners. The 50 mile ultra featured 9,500 feet of elevation change, with 125 runners starting, and 93 finishers. FabTime President Frank Chance finished the 50 mile ultra in just over 15 hours. FabTime's Senior Industrial Engineer Mike Krist paced him for the last 19 miles. Beth Chance ran the half marathon, which had an elevation change of 2,750 feet, and finished in 2 hours and 45 minutes.

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

Copy Chart Images to Other Applications

Way back in Tip number 17 we wrote about ways to copy FabTime charts into other applications. The details of this capability have changed a bit over the years (as have web browsers), so today we are providing an updated tip about copying.

Copying Individual Charts (from a home page or a Charts page)

1. Hover your cursor over the chart image and right-click. Depending on your browser (e.g. in Chrome), the pop-up menu may, by default, include options such as "Save Image As..." and "Copy Image." Selecting "Save Image As..." saves the image to a location on your computer as a PNG file (which you can name). You can then insert the image as a picture using other software applications (Word, Excel, email, etc.). Selecting "Copy Image" pulls the image into your computer's clipboard. You can go to another application and select "Paste" or "Paste Special" to paste the image as a device independent bitmap.

2. If the right-click menu for the chart does not include the above options (for example, in Microsoft Edge), this is likely due to the fact that the chart you are looking at is a JavaScript chart. What you need to do prior to copying it is tell FabTime to change from a dynamic chart to a static chart. The quickest way to do this is by clicking the small "D" located in a box in the lower-left corner of the chart (just above a box containing a small "L"). The "D" should then appear with an X through it. This means that the chart is no longer dynamic, and you'll have options such as "Save Picture As", "Copy Picture" and "Share Picture" (in Microsoft Edge). Alternatively, on either the home page or an individual chart page, you can change the "Active" setting under the Format controls from "JavaScript" to "Never." Note that the "Never" option is available only for sites that have purchased ChartFX. This will render the chart (or the entire home page chart) static. You will need to change this back afterward if you want to use mouse-over controls, chart editing, and other features of the JavaScript charts.

3. In either case, you will get a higher image quality if you copy a larger image from FabTime. Making charts smaller after copying them to another application should work (at the same proportional size), but you will not be able to make charts larger without blurring them. You can resize chart images by dragging them (JavaScript only), or by using the "Width" and "Height" controls in the Format section (available on home page tabs and on the individual chart page).

Exportable View

You can click the "Exportable View" link from a home page tab. FabTime will generate a single web page containing just the chart images (with no controls) for that tab. The charts will still be dynamic, however, if they were dynamic previously. You can either click on the "D" for each individual chart or, prior to clicking "Exportable View" change the "Active" setting for the tab to "Never". You'll still need to copy the charts individually to other applications, but this Exportable View can make the process a bit quicker.

Screen Captures

Instead of the above methods, you can always use a screen capture application such as SnagIt (which we like) or any native screen capture capabilities of your computer (Macs make this quite easy). This would allow you to copy a picture of the chart along with filter settings and/or the data table, or to easily take an image of all of the charts from the Exportable View.

Please bear in mind that any fab data that you display in FabTime is most likely confidential to your company. We share this tip because it is something that people

ask about - they want to be able to share FabTime charts with their colleagues in reports, presentations, email messages, etc.

If you have questions about this item, or any other FabTime software questions, just use the Feedback form inside FabTime's software. Subscribe to the separate Tip of the Month email list (with additional discussion for customers only). Thanks!

Subscriber Discussion Forum

Purchasing Question: Two Manual Load Tools or One Auto-Load Tool

Richard Davis from Honeywell wrote: "Here is an interesting diagnostic or thought question. Suppose you need a specialized process or inspection tool. The volume of wafers to run on the tool is only about 1000 to 5000 wafers per year, but you must have this special tool. You have $500k available. A manually loaded version of the tool is only $250k and the auto-load version is twice the price at $500k due to the robot. So, do you buy 2 manual tools to make sure one is always available? Or, do you buy the single auto-load tool to reduce operator handling. Well, as they say, it depends. I tend to side with the former and others who come here from big fabs side with the later. What principles should one follow to make a wise decision rather than an emotional one?

FabTime Response: Taking a cycle time improvement perspective, we would recommend buying the two manual tools, vs. the one auto-load tool. This is because the average cycle time through the two tools will be much less than through the one tool. A rule of thumb is that the cycle time through a toolgroup with two tools is approximately half of the cycle time through a single path tool, even if the toolgroups are at the same utilization. This is because of variability (due to process times or equipment downtime) - when there is redundancy, the chance of any given lot having to wait is much lower. In this case, you would have lower cycle time due to the tools being at a lower utilization AND you would get the benefit of the redundancy. Even if you lost some capacity on the manual-load systems due to the load time, it seems highly unlikely that this would outweigh the larger cycle time impact from the redundancy.

Of course there is the yield risk of the manual handling, but you would have to have fairly high scrap/rework rates for this effect to dominate the effect of the cycle time. Other issues might involve space constraints for the extra tool, and extra maintenance costs for having two tools vs. one. But if it was up to FabTime, we would go for the variability reduction of having that extra redundancy.

Richard added: It appears that you and I are thinking similarly on this. The two manually loaded tools will be much more reliable in the long run especially without a robot to break down. Another way this is particularly illustrated would be at wafer probe. If it takes 5 minutes to probe a wafer (large die without trimming) then a cassette loader might help to keep the wafers moving. If it takes 5 hours to probe a wafer (not uncommon for small die with trimming), then you should buy a second prober/tester rather than buy a robotic loader that sits idle most of the time.

FabTime agrees. What say all of you? Our guess is that those from larger/more automated fab backgrounds would be more likely to lean towards going ahead with the automated tool, while those who have struggled with one-of-a-kind tools in smaller fabs would go for the redundancy.

Computational Issues in Reporting "Average WIP"

Introduction

Something that we've observed repeatedly, in our years working with data from wafer fabs, is that even things that seem simple can be difficult to define and even more difficult to compute. One example of this is line yield, which we discussed in Issue 9.06: Definitions for Short-Term Line Yield Metrics. (The issue here concerns the time window of interest, and the varying lag between starts and shipments by product.) As Facebook recently discovered, and we will discuss in more detail in the next issue, even the computation of simple averages can be problematic. Recently we've been looking closely at the metric average WIP (average work-in-process) and making some changes in how this is computed in FabTime's software.

The point in time measure of WIP for a fab is fairly straightforward, of course. At any given point, we add up all of the wafers that are in the fab, and this is the fab's total WIP. Even here there are sometimes questions about things that should be excluded from this number, such as WIP that is in storage or other extended hold. But it's possible to use hold codes or other exclusion filters to remove this WIP from the calculation if needed, so these issues are not generally a problem.

More tricky questions are these:

1. What WIP value should we use for calculating WIP turns, which are defined as moves divided by WIP? Should we use starting WIP, ending WIP, or average WIP?

2. For cases where we want to use average WIP, how should we compute that average? Simply taking the average of starting and ending WIP for a period can lead to distortions, particularly for long time periods, and for smaller segments of the fab. For example, WIP at a batch toolgroup might be high for most of a shift, and then drop down right before the end of the shift, because a couple of large batches were moved out. The average taken from starting WIP and ending WIP would not accurately reflect this situation.

Using Starting WIP for Calculating Turns

When we first began reporting WIP turns in FabTime, we defined turns as moves / starting WIP for each time period. This made some intuitive sense, since the idea of turns was to look at how many times each wafer moved during the day or shift. We can think about this fairly easily if we consider how many wafers were there at the start of the shift, and then we report how many times each of those wafers moved.

However, we found that this could lead to distortions for longer time periods. For instance, you might have a high capacity toolgroup that had a lot of WIP at the start of the shift, worked that off quickly, and then was starved during the rest of the shift and didn't have any more moves. The turns rate would end up low, reflecting the low overall moves rate AND the high starting WIP. This would be true even though the operators were processing WIP when there was WIP there to process. The longer the time period reported, the worse this effect could be.

An example showing the turns rate for a tool group by week, over a five week period, is shown at the top of the next page. In this example, starting WIP is used in the turns calculation. We can see high turns rates recorded where starting WIP is low, and vice versa. The turns value is shown on or above each bar.

We concluded from examples like this that using starting WIP was insufficient. Similarly, using ending WIP would not have helped, and might have also led to undesirable incentives.

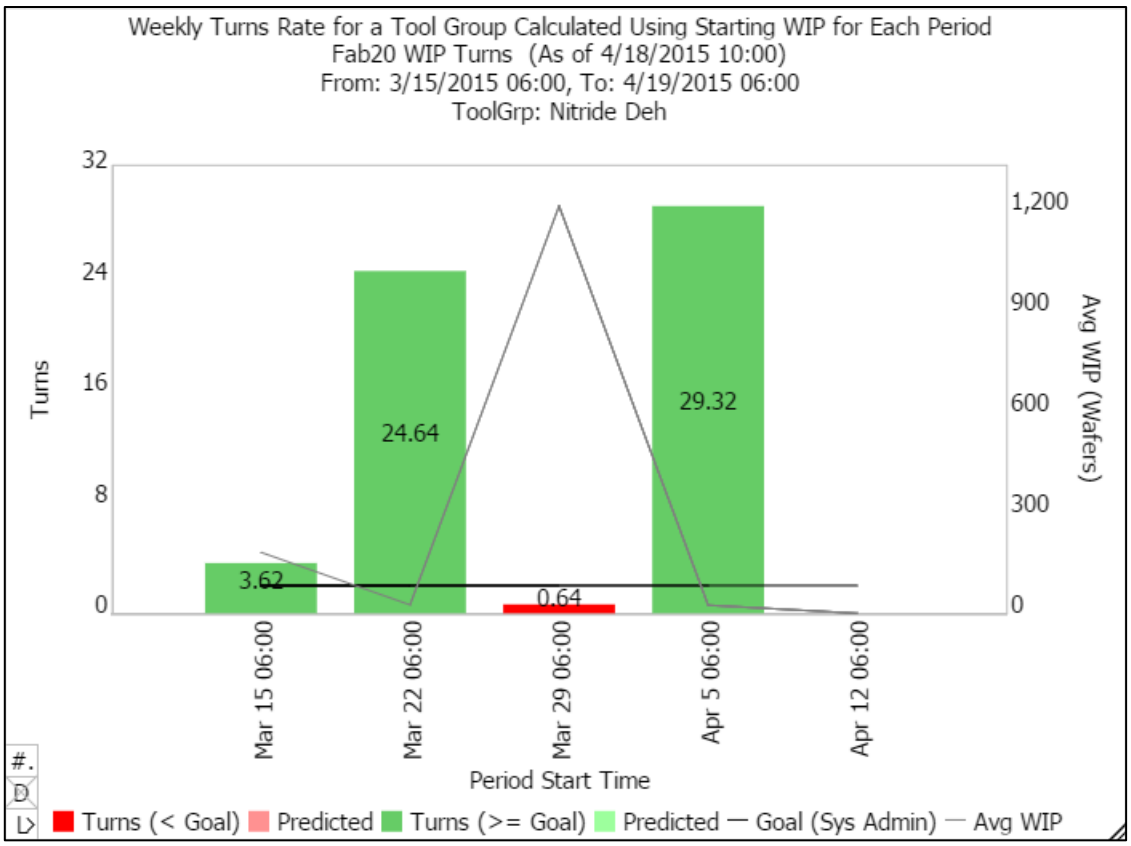

Using Point-Average WIP for Calculating Turns

We decided that it was necessary to at least add the option of using an average WIP value.

The way we did this initially was to add the concept of a sub-period. We then averaged the WIP values from the start of each sub-period, calling this point-average WIP. So, for example, if the period length was 12 hours (one shift), and the sub-period length was 6 hours, the WIP value used would be the average of two values: the WIP at the start of the shift and the WIP half-way through the shift. Using smaller sub-periods would give a more granular response.

The nice thing about this approach was that for those who preferred to continue using starting WIP, they could just set the sub-period length equal to the period length. This was the way that we set the sub-periods by default, so that no values would change on anyone's charts unless they took the step of setting a smaller sub-period length.

The same example shown previously, using a sub-period length of 6 hours, instead of using starting WIP, is shown at the top of the next page.

However, this point-average approach did have several drawbacks:

1. The need to specify a sub-period length added complexity.

2. For longer sub-periods, the point-average could still be an inaccurate representation of the behavior of the WIP during that time.

3. For shorter sub-periods, the point-average method could be slow, and sometimes used redundant data (if the WIP value stayed the same for some longer period of time).

Using Time-Average WIP in FabTime

What we've done most recently is add a new time-average WIP computation. The

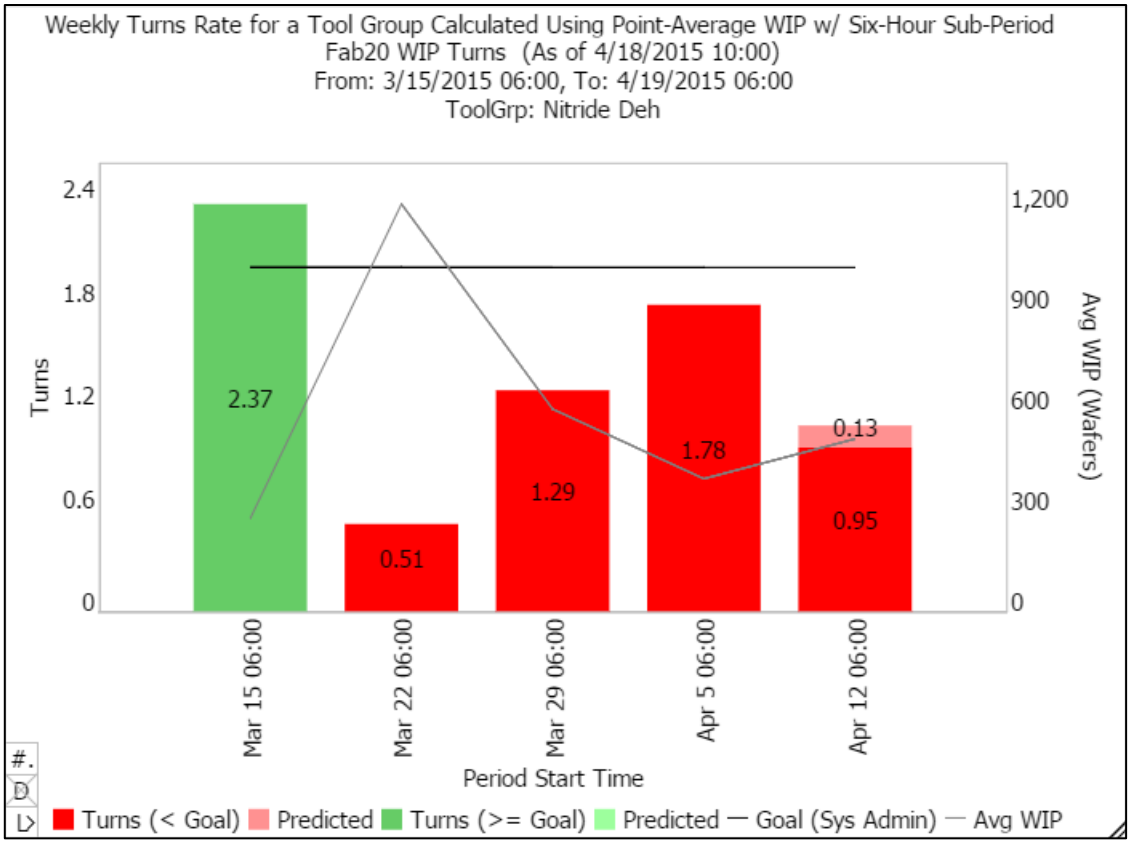

time-average method looks at how long the WIP is at each level, and computes a weighted average of those values. Setting the sub-period length to zero, or leaving it blank, tells FabTime to use the time-average method. Specifying a non-zero sub-period tells FabTime to use the point-average method. Where the sub-period length is equal to the period length, the point-average method defaults to starting WIP.

This time-average or point-average WIP capability is available on the WIP turns charts, as discussed above, and on other charts that display WIP, like the WIP trend and moves trend charts. We've also replaced starting WIP with average WIP on the OEE charts.

The chart above looks quite similar when calculated using time-average WIP, as shown at the top of the next page. The shape of the graph is the same, though small differences in the turns values can be seen.

Below are numerical examples of average WIP calculations using both the point-average and time-average methods.

Examples – Sample Data

Suppose the period of interest is Monday 5pm to Tuesday 5pm.

Suppose the WIP during this time is as follows:

- Monday 5pm: 300 wafers
- Tuesday 1am: 200 wafers
- Tuesday 9am: 100 wafers
- Tuesday 10am: 800 wafers (e.g., a WIP bubble arrives at 10am)

Point-Average WIP

Point-average WIP is computed by dividing each period into sub-periods (points), and computing the average starting WIP for each sub-period within the period. Suppose we choose a sub-period length of 8 hours.

Then our three WIP snapshots are as follows:

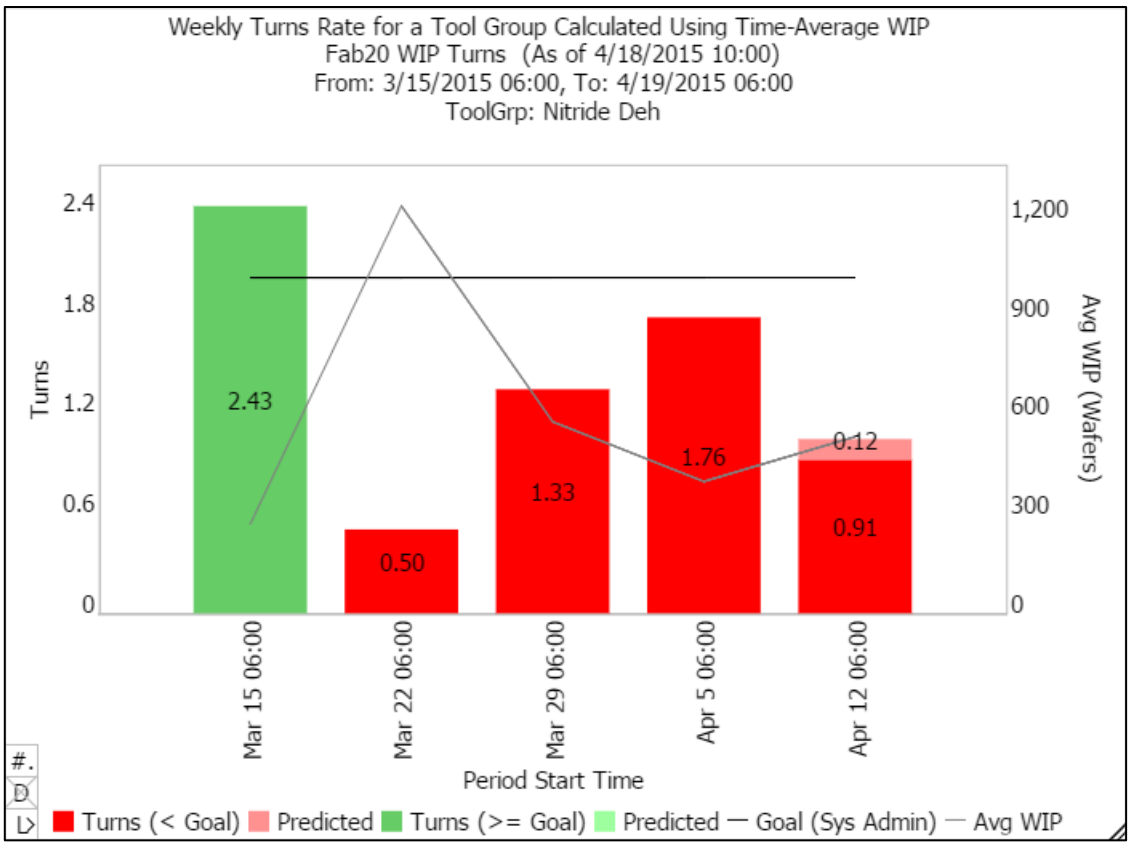

- Monday as-of 5pm: 300 wafers
- Tuesday as-of 1am: 200 wafers
- Tuesday as-of 9am: 100 wafers

The point-average WIP is (300+200+100 wafers)/3 = 600 wafers / 3 = 200 wafers.

Note that the WIP value at Tuesday 5pm is not used in point-average WIP, just as the ending WIP value is not used in time-average WIP. And our computations ignore completely the WIP bubble that arrived at 10am.

Time-Average WIP

Time-average WIP is computed by summing the amount of time spent at each WIP value, then dividing by the total time.

Using our sample data, time-average WIP is computed as:

- 8 hours (5pm-1am) * 300 wafers = 2400 wafer-hours
- 8 hours (1am-9am) * 200 wafers = 1600 wafer-hours
- 1 hour (9am-10am) * 100 wafers = 100 wafer-hours
- 7 hours (10am-5pm) * 800 wafers = 5600 wafer-hours

Total wafer-hours = 2400 + 1600 + 100 + 5600 = 9,700 wafer hours

Total hours = 24

Time-average WIP = total wafer-hours / total hours = 9,700 / 24 = 404 wafers

Note that the WIP value at Tuesday 5pm is not used in time-average WIP, just as the ending WIP value is not used in point-average WIP.

Benefits of Using Time-Average WIP vs. Point-Average WIP

There are several benefits to using time-average WIP instead of point-average WIP. Most notable is that this time-average method doesn't throw away any data. Sampling every hour (or every 8 hours) throws away everything that happened between those sample times – and in the example above, point-average WIP misses the WIP bubble that arrived at 10am. We could switch to a shorter sub-

period length with the point-average WIP, but someone has to make that decision, and switching a very small sub-period length leads to slower computations. But with the time-average method, we note every time that the WIP changes, thus using all of the available data. The time-average calculation is also simpler (not requiring users to choose a sub-period length), and can be faster than the point-average calculation in some cases. In cases where the WIP level changes very frequently, however, the point-average method will likely be faster, and make more sense to use, particularly where the WIP level is changing only by small amounts.

Conclusions

WIP is one of the most common variables that people use in monitoring fab performance. We look at the linearity of WIP throughout the line. We look for trends in WIP over time. And we use WIP as a variable in other calculations, such as turns. But even with such a commonly used and well understood variable, there are decisions to be made. Do we report and use WIP values from the start of each shift, or do we use averages? And where we are using average WIP, is it better to use a point-average method (sampling, in essence), or to go to the additional level of detail of using a time-average (tracking how long the WIP remains at each level and aggregating that)?

What we have learned in working through these issues is that while there are pros and cons to the different methods (particularly in computation speed), using time-average WIP provides the most accurate picture of what's going on in the fab. We are grateful to our customers and our employees for their input as we have gone through this learning process.

Closing Questions for FabTime Subscribers

What value of WIP do you use in your WIP turns calculations? Starting WIP, point-average WIP, time-average WIP, or something else?

Subscriber List

Total number of subscribers: 2762.

Top 20 subscribing companies:

- Infineon Technologies (inc. International Rectifier) (144)
- Micron Technology, Inc. (138)
- Intel Corporation (129)
- ON Semiconductor (108)
- Maxim Integrated Products, Inc. (106)
- GLOBALFOUNDRIES (98)
- Carsem M Sdn Bhd (71)
- Fairchild Semiconductor (69)
- Texas Instruments (65)
- STMicroelectronics (62)
- X-FAB Inc. (56)
- Skyworks Solutions, Inc. (53)
- Seagate Technology (52)
- Freescale Semiconductor (51)
- Western Digital Corporation (49)
- Analog Devices (45)
- Microchip Technology (44)
- TDK (inc. Epcos) (41)
- Atmel Corporation (38)
- NXP Semiconductors (33)

Top 4 subscribing universities:

- Ecole des Mines de Saint-Etienne (EMSE) (18)
- Arizona State University (8)
- Nanyang Technological University (7)
- Virginia Tech (7)

New companies and universities this month:

- Bruker BNS
- Dow Chemical
- Enovix Corp.
- Google
- VDL ETG T&D

Sampler Set of Other Subscribing Companies and Universities:

- Adams Associates (1)
- BAE Systems (12)
- bTendo (1)
- COM DEV Space (1)
- Delphi Delco Electronics Systems (1)
- Dynacraft Industries Sdn Bhd (2)
- Ecole des Mines de Nantes (1)
- Front Line Performance (1)
- Georgia Tech (1)
- Integrated Technologies Company (1)
- McKinsey (1)
- Mentor Graphics (1)
- Middlesex General Industries, Inc. (1)
- Nimble Consulting Services (1)
- PerkinElmer (1)
- Sirris Belgium (1)
- SunPower Corp. (1)
- Tower Jazz Semiconductor Ltd. (18)
- Westcode Semiconductor (1)
- Xilinx Semiconductors (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® Software for Assembly and Test

"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. 4 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

FabTime's Web-Based Dashboard is Fully Applicable for Assembly & Test Facilities

- Do your customers (internal or external) want more visibility into your factory?
- Is it difficult to look at trends in equipment performance, or tie equipment performance to throughput and cycle time?
- Does your factory lack real-time reporting?

FabTime can help. FabTime saves your management team time daily by turning 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. Most importantly, FabTime can help your company to increase revenue by reducing cycle times up to 20% for regular lots, and even more for high-priority lots.

Although FabTime was originally designed for front-end manufacturing, you can use FabTime for your assembly or test facility. You simply need to have a transaction-based manufacturing execution system. FabTime can link to all commercial systems commonly used in the industry (e.g. WorkStream, Promis, Eyelit, Mesa, FactoryWorks) or can link to internally developed systems. FabTime can pull data from multiple databases if needed (e.g. WIP transactions from the MES, tool transactions from another system). FabTime is currently being implemented in two assembly and test facilities, with no major technical hurdles.

FabTime Applicability for Back-End Factories

- FabTime handles lot merging and splitting, with full tracking of overall cycle times.
- All chart quantities (moves, WIP, etc.) can be displayed as die, with data tables formatted for readability of large quantity values.
- Custom assembly and test parameters (applicable to WIP or tool state transactions) can be mapped.
- Custom site-specific reports for wire bond area have been developed for customers (die and component placements, etc.).
- Custom dispatch factors allow for incorporation of back-end-specific data used in dispatch decisions (e.g. availability of boards, and minimization of sequence-dependent setups).