FabTime Cycle Time Management Newsletter

Volume 8, No. 9 November 2007

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 new formatting for home page charts to allow them be displayed smaller, and to allow the display of chart data tables on the home page

Editor: Jennifer Robinson

Contributors: V.A. Ames (ATDF), Ryan Carlson (Seagate); Tim Skowronski (Intel); Ilia Kaplan (National Semiconductor); Richard Davis (Honeywell Sensing & Control)

Table of Contents

Welcome

Welcome to Volume 8, Number 9 of the FabTime Cycle Time Management Newsletter! We hope that you're all doing well, as 2007 draws towards a close. In this issue we have a brief followup to an earlier announcement, to remind you about the upcoming launch of the Fab Engineering & Operations Magazine, as well as a job change announcement from V.A. Ames. Our FabTime tip of the month is about the use of new formatting controls to enable smaller home page charts. We have subscriber discussion about managing production and development activities in the same fab, loading and managing batch tools, and varying lot sizes in the fab.

Because we have quite a bit of subscriber discussion this month, we bring you a relatively short main article. We discuss some of the challenges of calculating cycle time benchmark data. Specifically, we review the two primary metrics currently used for benchmarking across fabs and technologies, X-factor and days per mask layer (DPML), and discuss specific computational issues that apply to each one. We also discuss the conversion ratio between the two metrics. Our hope is that this article will spur further discussion, which will in turn help people who are looking to benchmark and improve their cycle times.

Thanks for reading!—Jennifer

FabTime

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

Community News/Announcements

Fab Engineering and Operations Magazine

There's still time to subscribe and receive the charter issue of the new quarterly Fab Engineering and Operations Magazine. According to the FEO website (www.feoproject.com/feoproject/):

"Fab Engineering & Operations™ (FEO™) is a new FREE quarterly magazine focusing on mainstream device fabrication plants (fabs) of all shapes and sizes. FEO will be guided, written and read by industry experts – those who are at the forefront of the day-to-day issues and have to make the decisions that affect the running of the fabs. FEO will not focus on the "bleeding edge" of the industry, which only constitutes a small percentage of actual global manufacturing. Instead FEO will concentrate on the majority of fabs, not using this "bleeding edge" technology, and it will look at topics such as cost reduction, business strategies, efficiencies and the supply chain."

There is no cost to subscribe (the magazine is advertiser subsidized, and available via PDF). We think it's going to be a very useful publication, one that will benefit FabTime's newsletter subscribers, and that's why we're telling you about it. We have no formal relationship with FEO, beyond both being affiliate with the Fab Owners Association, but we like to bring good ideas to your attention when we find them.

Job Change Announcement: V.A. Ames

Longtime newsletter subscriber V.A. Ames would like to announce that he is now the Equipment Services Manager at ATDF. His responsibilities include Site Equipment Maintenance, Stores, and Shipping and Receiving. He can be contacted at v.a.ames@atdf.com. We wish him well in his new position.

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

FabTime User Tip of the Month

Display Very Small Charts on your Home Page Tabs

Have you ever wanted to be able to see more home page charts at once, but struggled with being able to see the data once the chart gets too small? We've made some recent enhancements to FabTime that allow you to display smaller charts, but still see the data clearly. These changes are in a new FabTime patch, which may not yet be available at your site. If you don't see the options described below, contact your FabTime administrator to ask about Patch89.

On any home page tab that contains one or more charts you'll find two new

dropdowns and one new text box in the "Format" section (in the lower left-hand corner of the screen). In addition to the dropdown for showing or hiding chart legends, there's now a drop-down to show or hide standard chart titles ("StdTitles"). Standard chart titles are the title rows at the top of each chart, which include the date that the chart was generated, FabTime's name for the chart, and any filters you have defined. If you select "Hide" next to "StdTitles", and then press the "Go" button, FabTime hides all of those title rows. The only title that you'll see for each chart will be the custom title (if any) that you created as part of the chart definition. If you are going to hide chart titles, we do recommend creating your own descriptive custom titles, especially for charts on home page tabs that you intend to share.

The second new drop-down allows you to show or hide "AxisTicks". If you hide axis ticks (and press "Go"), FabTime will no longer display tick marks on the x- and y-axes. Hiding the tick mark labels makes small charts a bit cleaner and easier to read.

The third new format control is a text box labeled "FontSize". You can use this to shrink the font for chart titles and x- and y-axis labels. Using smaller fonts allows you to see more axis labels in some cases, and usually increases the displayed size of the chart itself. You can easily experiment with this on your own. For instance, if you hide legends, standard titles, and axis ticks, and change the font size to 7, you can probably (depending on your screen resolution) make your charts 220 by 200 and fit four columns to a row. We've even seen five charts to a row for high-resolution monitors. These small charts may not be for everyone, but for people who do like to see lots of data at once, they can be quite useful. See an example on page 11 of this issue.

Please note that we made one other change in conjunction with adding this new home page chart formatting. Home page chart formats are now tab-specific. This allows you to have, for example, lots of small charts on your daily production summary page, but have larger charts on some other tab. FabTime will, of course, remember the last format settings that you specified for each tab. 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

Managing Combined Production/Development Fabs

Ryan Carlson of Seagate Technology submitted the following question for subscriber discussion: "How do fabs that run both development and production wafers maintain learning with fewer moves devoted to R&D lots?"

FabTime Response:

We don't have direct experience with Ryan's question, so we are opening it up to other subscribers. In you work in a fab in which you're juggling production and development WIP, do you have any strategies for maintaining learning with fewer R&D moves? We would value your

input. Your response can be kept confidential by FabTime if you choose.

Issue 8.06: Cluster Tools

Ilia Kaplan from National Semiconductor wrote: "I just read the Vol. 8 No. 7 issue of FabTime and would like to try and help answer 'Issue 8.06: Cluster Tools' from Bob Kotcher (Western Digital) regarding how to model cluster tools. Bob, if you're trying to determined theoretical UPH, or how to optimize your robot, I would suggest developing a speed model based on observations (time studies). In my experience, the most effective way to determine robot efficiency while trying out 'what-if' scenarios is to use Gantt charts in Excel. It's a simplistic, yet effective, method to accomplish what you're looking to do. You can even determine robot utilization and % of non-value add activity."

Batch Tool Loading Analogy to Setups (response to ongoing discussion topic)

Tim Skowronski of Intel Corporation wrote in response to Ulrich Dierks and Thomas Quarg's subscriber discussion article in last month's issue: "I have read the comments from the AMD guys about batch tools becoming factory bottlenecks while waiting for full-loads. I think this is a subset of a larger problem that involves conversion/set-up time. I mean, if you have a limited amount of equipment and that equipment is set-up to run a given product (assembly) or layer (fab) and you let the equipment run idle because the WIP that is waiting is not sufficient to warrant a conversion (per manufacturing policy) ...this is the same problem, no?

If you agree, then I have witnessed a non-constraint toolset become a high CT contributor and factory output bottleneck first hand in an assembly environment. Here, running smaller batches (equivalent to short-loads in fab) eats too much capacity and output suffers while CT suffers when WIP waits to accumulate batches large enough to warrant a conversion.

This is a problem that I think fabs are struggling mightily with. Past focus has been to determine optimal batch sizes, or lot size increases. I think that it makes sense to shift the focus to the real problem of the conversion itself (e.g. reducing it). Do you have any recommendations for material on set-up / conversion time reduction in a semiconductor / assembly environment?"

FabTime Response:

We think that you make an excellent point. We would agree with you that the setup question and the batch loading question are similar. We think that the decisions are analogous to one another, with both being a subset of a more general question about dispatching in the presence of material with different loading characteristics. Your question about setup is certainly one that we hear about from many people.

The most common approach that we've seen in practice is something called a Setup Avoidance dispatch policy. This rule states that as long as there's something in queue with a matching setup ID, you keep processing that setup ID. However, if there is nothing in queue with a matching setup ID, and there are other lots in queue, then you go ahead and do the setup. You don't hold the tool idle to avoid the setup (at least in theory, though we know that people do hold the tool idle sometimes in practice).

We think that, as with the case of running a greedy batch policy (where you start the tool as long as there's anything there), it's generally going to be better for cycle time if, whenever there is WIP in queue, if there's no matching lot, you go ahead and perform the setup. The capacity that you lose by holding a tool idle because you don't want to do a changeover will often balance out the capacity that you lose by doing the setup, and you'll keep lot moving

through the fab, instead of held up waiting for a setup. Of course this is not always the case. If you have a lot with a matching setup ID arriving in 10 minutes, and the setup takes four hours, you'll want to wait for the lot that's coming in 10 minutes.

With any setup avoidance policy, it's also necessary to have a cap on the maximum time that any lot can wait in queue. This prevents your low volume lots from sitting ...forever. What that cut-off should be is very case-specific, and we've seen no general results. And of course we agree with you that anything that you can do to reduce the magnitude of the setups will help you to increase your factory's flexibility, and thus cope better with this setup problem.

We did discuss this question somewhat back in newsletter issue 6.07. There's plenty of literature about scheduling in the presence of sequence dependent setups, but much of it is more theoretical than applied. We did find two papers that looked like they might be helpful:

- S. Park, J. W. Fowler, M. Hickie, and M. Carlyle, "Assessment of Potential Gains in Productivity Due to Proactive Reticle Management Using Discrete Event Simulation," Proceedings of the 1999 Winter Simulation Conference, 1999.

- R. Sunkara and R. Rao, "A Heuristic to Determine Equipment Setup Changes Based on Estimated Lot Arrivals in a Semiconductor Fab," Proceedings of the 2004 Winter Simulation Conference, Washington, DC, Dec. 5-8, 2004.

Both are available for download from the Winter Simulation Conference website. We are including this discussion in the subscriber discussion forum to see if any other readers would like to contribute. How do you manage this setup question? Do you ever leave tools idle when there is WIP there, to avoid doing a setup? Do you have a trigger that forces a setup, if a lot waits more than some interval of time? We look forward to your responses.

Lot Size and Batch Loading

Richard Davis from Honeywell Sensing & Control wrote in response to the ongoing batch size discussion: "There seems to be a good bit of discussion concerning proper batch size as it relates to cycle time. One aspect that I have not seen discussed is the following. There is some general consensus that cycle time is a function of percent capacity. In other words, your cycle time is usually lower at 50% capacity than at 95% capacity because the queue times will probably be shorter. Now, within certain bounds, larger batch sizes will effectively create larger capacity for your fab without added people or equipment. In other words, running a 24 wafer lot does not cost much more than running a 12 wafer lot. If your fab has a lengthy setup time (e.g. confirming CD on a stepper before running a lot) or has quite a few batch tools (e.g. oxide wet etch that holds 2 cassettes or a diffusion furnace that holds 4 cassettes) then it would seem that running larger batch sizes would not only lower the cost, but might actually reduce cycle time by increasing capacity. In other words, capacity would increase while total volume remained constant thereby reducing per cent capacity.

A modification of this idea is to divide the fab into two areas: a) photo, wet etch, and diffusion and then, b) metal and plasma. The first part is best modeled as a hub and spoke system in which 24 or 48 wafer batches makes more sense. The latter portion is somewhat more serial in flow and so smaller batches would likely work better. So, a fab might start 48 wafer lots and then split them into 4 each of 12 wafer sublots at metal dep or some other appropriate spot. I have never seen anything like this modeled before. In general, batch and/or lot size is not a straightforward and simple affair. One size does not fit all. There are different factors involved such that one lot size is not always ideal for each fab.

Also, for those who are interested in processing, I recommend a show called "Unwrapped" which is on the Food Network. It shows how all kinds of food products are processed and packaged and gives one a different perspective on "batch" when you see all the tons of flour and vats of sugar that are processed."

FabTime Response:

Thanks for the feedback, Richard. We agree that the question of lot size as it relates to cycle time is not straightforward. We have written about this, but it was back in Issue 2.2. While having larger lot sizes will help with capacity on certain types of tools, such as sinks, where the process time is independent of the lot size, there are other tools where processing is on a per-wafer basis. Having larger lot sizes means that the first lot will always wait longer for the entire lot to be processed. Many sites have tried running smaller lot sizes, in order to improve cycle time, though this can lead to capacity issues on your per-lot tools. In general, we think that lot size changes should be approached very carefully, but can certainly be helpful in specific situations. We agree that we haven't seen an article that look at exactly what you describe, a fab with the larger lot sizes in the front end split into smaller lot sizes in the back end. Perhaps your comments will inspire someone to model this.

Definitions for Cycle Time Benchmarking

This article was written by Jennifer Robinson from FabTime

Introduction

Because we have quite a bit of subscriber discussion this month, we bring you a relatively short main article. Our hope is that this article will spur further discussion, which will in turn help people who are looking to benchmark and improve their cycle times. As regular readers of this newsletter know, FabTime is a proponent of benchmarking a fab's cycle time performance based on X-factor. X-factor is defined, in its simplest terms, as total fab cycle time divided by theoretical fab cycle time. We like X-factor as a metric because a) X-factor tells you how you are doing compared to the best that you could be doing; and b) there are queueing models to predict expected X-factor values, and these tend to be helpful for intuition-building. However, many fabs shy away from using X-factor for benchmarking, and instead compare themselves to other fabs based on days per mask layer (total cycle time divided by number of layers). Days per mask layer seems to be in more widespread use as a benchmarking metric – people are familiar with it, and they find it easy to calculate.

In some informal discussions at a recent industry meeting, certain detailed questions arose about the calculation and use of both x-factor and days per mask layer. In this article, we will review some of these issues and some others that we have noticed over the years. We also seek subscriber feedback. If we receive enough feedback, we will write a follow-on article about this for a future newsletter issue. Please send your feedback to newsletter@fabtime.com, and please specify whether you would like your name associated with your feedback, or whether you would like your contribution to be anonymous. Thanks!

X-Factor

As stated above, X-factor for shipped lots is defined as overall manufacturing cycle time divided by theoretical cycle time (also called raw process time). Usually a weighted average X-factor is calculated for a fab, based on the current product mix. A fab that reports a 3X cycle time has a cycle time that is, on average, three times the theoretical cycle time. Lots spend two hours in queue/transit for every hour they spend in process.

The primary issue with X-factor is: how do you measure X (the theoretical cycle time)? Some fabs use the cycle time of a hand-carry lot, and call that the best-case cycle time. Our experience, however, has been that even fabs that are very good at processing hand-carry lots will still have hand-carry lot cycle times on the order of 1.2X, instead of 1.0X. Thus using a hand-carry lot cycle time as a proxy for X will tend to give a slightly optimistic estimate of your fab's performance. There's also an issue with lot size. Often hand-carry lots have a smaller lot size than regular production lots. Thus the best-case cycle time achieved for a hand-carry lot might be less than the best-case cycle time that you would achieve for a regular production lot. Using the smaller lot size hand-carry cycle time as a proxy for X will in this case tend to give a pessimistic estimate of your performance. The other issue with using hand-carry lots is that X is product-specific. Do you run a hand-carry lot for every new product (or process)? If you have more than one or two hand-carry lots in your fab at one time, their cycle time will likely be higher than 1.2X.

Some fabs take the raw process time for each step, and add that up across all of the steps in a process flow. But there are issues here, too. Do you use the actual best case process time for lots that have been run through an operation? Do you use an average process time? Or do you use a planned cycle time value for each step, which is generated based on process engineering requirements and/or UPH rates? With the former (using actual process time values) there are potential transaction logging issues. For instance, there is the occasional practice of doing track in and track out back to back, instead of doing track in before the lot is processed. That can bias your process time estimates low. Also, for fabs that are not highly automated, recorded process times may be inflated when there is no one there to unload the lot from the tool, and log the track out. If you're going to use calculated planned cycle time values, you need someone keeping planned process time values up to date (no small task). And you have to decide how you're going to treat load and setup times, especially on cluster tools and linked steppers. You also have to decide if you're going to include any travel time. And you still have the issues of needing an X value for each process flow, and of X being lot-size dependent.

Some fabs that we work with use Dynamic X-Factor (DXF) to get a ballpark estimate of the fab's weighted average X-factor, without needing to know X. DXF (as described in past newsletters) is a point estimate measured frequently over time. Each observation records total WIP in the fab divided by non-rework WIP that is currently in process on tools. A DXF of 3 means that for every lot in process there

are two lots in queue/transport. If you measure DXF frequently over several days, so that you have at least 30 observations, (e.g. hourly over 2-3 days) you can average the results to get an estimate of your overall DXF. This may track fairly closely with your actual X-Factor, though it may vary due to issues with completed lots not being moved out from tools, and how you handle hold and rework WIP.

What we haven't seen is anyone comparing their observed DXF values with observed X-factor values (based on actual weighted average cycle time over theoretical cycle time, as calculated above). Has anyone out there looked at this?

Days per Mask Layer

On the surface, days per mask layer (DPML) is an easier metric to calculate than X-factor. For each lot, you take the shipped lot cycle time and divide by the number of mask levels for that lot's process. Then you average across all of the lots. However, there are issues with DPML, too. Do you count rework layers as a layer? (We would think not, but we've never asked this before). How do you handle comparisons across technologies when a new technology might have much more complex layers? Is there a way that you can measure DPML for lots that are still in the fab, or do you always have to wait for the lots to ship before you can calculate? Is there a dynamic version of DPML that can be used to get a quick estimate of the weighted average value, without having to calculate separately for each lot? How do people handle this in practice?

Shipped Lot Cycle Time

X-Factor and DPML have the same numerator: shipped lot cycle time. Even here there are questions about how to calculate shipped lot cycle time for a lot. Your fab has to agree on a start operation and a shipment operation for cycle time purposes. You have to decide whether you want to include electrical test at the end. What about backgrind and die sort? The important thing here is to be consistent with other fabs that you might be benchmarking against. You also have to decide if anything will be excluded. It's not uncommon for fabs to exclude time that lots spend in extended storage (as when a lot is taken off the manufacturing floor, because of changing market demands, and later re-started). This is valid if you are assessing the manufacturing organization's performance. However, in real time, for your company as a whole, that time was incurred as cycle time. Apart from the possible exclusion of extended storage time, FabTime's definition has always been that everything else is included in cycle time: queue time, process time, hold time, setup time, travel time, etc. Is this universal?

Converting between X-Factor and DPML

Another question that came up in our informal discussions recently concerned the possibility of converting between X-factor and DPML. A conversion like this could be helpful for fabs that want to understand their X-factor, but are leery of the computation issues in measuring X. Some fabs, at least, seem to know a ballpark conversion ratio, but this is not universal. As an example of a conversion ratio, suppose your fab has a cycle time goal of 3X. If the same fab is targeting 1.5 DPML, then they have a DPML to X-factor conversation ratio of 0.5.

Conversion Ratio = DPML / X-factor

or equivalently, since both have the same numerator:

Conversion Ratio = (CT / Mask Levels) / (CT / Theoretical CT) = Theoretical CT / Mask Levels

The conversion ratio is thus nothing more than the average raw process time required to complete a mask level. We haven't seen any published figures, or even heard very

much in the way of discussion, on what this ratio might be in practice. It seems likely to be technology dependent. But we'd be happy to collect some values, in confidence, to get a sense of the range, if that would be helpful for people.

Conclusions

In this article, we have discussed some of the challenges of calculating cycle time benchmark data. We have reviewed the two primary metrics currently used for benchmarking across fabs and technologies, X-factor and days per mask layer (DPML), and discussed specific computational issues that apply to each one. We have also discussed the conversion ratio between the two metrics. We have raised these issues, and asked several questions, in an open format, in the hope of generating some useful discussion among members of the newsletter community. If you believe that these questions are of value, please take a few moments to send us your thoughts on the questions below. If we receive enough responses, we will write a followup article for a future issue. Otherwise, we will simply include responses in the subscriber discussion forum. Either way, you are welcome to include your response as attributed or anonymous comments. In the latter case, no one outside of FabTime will no know that your response came from you. You have our promise of that. We welcome your feedback.

Closing Questions for FabTime Subscribers

- How does your fab measure X (theoretical CT)? Do you use hand-carry lots, or process time estimates? Do you include travel time? Are your measurements lot-size dependent?
- Have you ever compared DXF values to shipped lot cycle time X-factors for your fab? Were they representative?
- How does your fab measure shipped lot cycle time (through electrical test? die sort?)?
- Does your fab use a dynamic version of DPML, or do you only compute DPML for shipped lots?
- Does your fab use X-factor, DPML, or both?
- Do you know the theoretical cycle time per mask level in your fab?

Acknowledgement

Parts of this article were inspired by informal discussions with Guy Gandenberger from Micrel and other members of the Fab Owners Association (http://www.waferfabs.org/).

Subscriber List

Total number of subscribers: 2794 from 479 companies and universities. 21 consultants.

Top 20 subscribing companies:

- Maxim Integrated Products, Inc (252)
- Intel Corporation (160)
- Micron Technology, Inc. (87)
- ATMEL (71)
- Analog Devices (67)
- Infineon Technologies (66)
- Freescale Semiconductor (65)
- X-FAB Inc. (62)
- Texas Instruments (58)
- International Rectifier (57)
- STMicroelectronics (56)
- Cypress Semiconductor (55)
- ON Semiconductor (51)
- TECH Semiconductor Singapore (51)
- Chartered Semiconductor Mfg (50)
- NXP Semiconductors (49)
- IBM (46)
- Spansion (38)
- Seagate Technology (32)
- BAE Systems (30)

Top 3 subscribing universities:

- Virginia Tech (11)
- Ben Gurion Univ. of the Negev (7)
- Nanyang Technological University (7)

New companies and universities this month:

- Advent Solar
- Alfalight
- CEITEC
- CH2M Hill
- Ching Yun University
- Diodes-FabTech
- Hendon Semiconductors
- K.U. Leuven University
- Microwave Device Technology
- Nitronex
- Rohm and Haas

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

Fixed price includes

- Software installation and real-time connect to your MES.
- User and admin training.

FabTime Subscription

Low monthly price includes

- Unlimited users via your Intranet.
- Maintenance and upgrades.
- Support for add-on dispatch and planning modules (one-time configuration fees for each)

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

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%.