FabTime Cycle Time Management Newsletter

Volume 8, No. 5 June 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 a "ToTool" input on the Forecast Outs Charts (allowing you to generate a list of which lots are likely to complete processing at a given tool within a future time window)

Editor: Jennifer Robinson

Contributors: Dov Kotlar (Tower Semiconductor)

Table of Contents

Welcome

Welcome to Volume 8, Number 5 of the FabTime Cycle Time Management Newsletter! I hope that you're having a nice summer, wherever you are. Things remain busy with FabTime, as we work on installation for our twelfth and thirteen customer sites, and on holding various sessions of our cycle time management course. We're also working on some great new features of the software, which you'll hear about later in the year. Our FabTime software tip of the month this issue is about using FabTime's new tool qualification charts to identify single-path operations in the fab. In our subscriber discussion forum, we have an extended discussion with Dov Kotlar of Tower Semiconductor about metrics for measuring fab utilization.

In our main article this month we tackle the subject of WIP bubbles. People ask us occasionally: "how do I manage WIP bubbles in the fab?". A WIP bubble is a large pile of WIP, usually in queue at a particular tool-group or small set of tool-groups. WIP bubbles occur due to a variety of causes, the most notable of which is extended downtime on a one-of-a-kind tool. In this article, we discuss common causes of WIP bubbles, methods for avoiding them, early WIP bubble indicators, and potential methods for mitigating their effect. Several of the latter involve making dispatching decisions that encompass information about downstream operations. We hope that you'll find this article useful, and we would love to hear your feedback.

Thanks for reading!—Jennifer

FabTime

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

Community News/Announcements

FabTime welcomes the opportunity to publish community announcements. Send them to newsletter@FabTime.com. We have no announcements at this time.

FabTime User Tip of the Month

Identify Single-Path Operations in Your Fab

A relatively new chart in FabTime is the Tool Qualification WIP Detail chart, available under "Tool Qualification Charts" on the FabTime chart list. This chart displays a snapshot of WIP at all flow/step combinations in FabTime, and lists the number of qualified tools that the FabTime database shows as qualified to run that WIP. A flow/step combination is a unique operation in FabTime (a step within a particular process flow). FabTime maintains a list of exactly which tools, according to data we've read from your MES or other databases, are qualified to run WIP at that step. Of course the accuracy of this data depends on how (or if) the data is maintained in your systems.

The chart shows the amount of WIP at each step, color-coded according to how many tools are available – black for no qualified tools, red for one qualified tool, yellow for two tools, and green for three or more tools. The data table spells out the set of tools qualified for each step. Note that this chart reflects the number of tools that are qualified to run WIP at a particular step, and does not necessarily mean that those tools are available at a given moment.

By default, this chart shows a snapshot of the current WIP, but you can enter a specific date in the "Date" field to look back in time. You will also need to enter your usual WIP filters, in order to filter down the WIP to only include lots that are currently eligible to be processed (e.g. to exclude monitor wafers, wafers in storage, etc.).

This chart can be a useful indicator of single path operations, which are known to have a significant impact on fab cycle times. We hope that you find it useful.

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

Fab Utilization

Dov Kotlar from Tower Semiconductor emailed us the following question: "I would like to discuss how to measure fab utilization. There are some aspects to the question, such as:

- Should we measure the utilization out of the total time (OEE) or the available time?
- On what time frame should it be measured and averaged? Shorter time frames may lead to a higher figure (due to variability), moving between different tools (temporary bottlenecks). And then – how should we generate the periodical result from the several tools (in the several short periods) utilization? Longer periods of measurement may hide significantly high utilization figures for tools that are not highlighted for the total long period.
- On how many tools should we measure the whole fab utilization? Only the bottleneck, across some leading utilization tools, or across all tools.

For me, and maybe for others too, it may be very interesting to be exposed to the different approaches in practice."

FabTime Response:

At FabTime, we think that the right thing to do is measure utilization of available time, rather than utilization of total time. That way, you're recording how much time you spend processing wafers out of the total time that the tool is available to manufacturing. This is what drives cycle time. For a fab, it would be the utilization of the bottleneck (most heavily utilized) tool. Yes, this can shift over time (which tool is the bottleneck), but I would argue that at any point in time, you do know which tool is the bottleneck, and it's the capacity of that tool that gates the total throughput of the fab.

Another way that people report fab utilization that's relatively straightforward is to take:

Total Wafer Outs per time period (usually a month) / Maximum Wafer Outs possible, for the current toolset and product mix.

The maximum wafer outs possible would be the amount that you could run that would drive the bottleneck tool to being 100% busy, after accounting for any downtime. So, still based on utilization of available time (or at least based on utilization of planned available time), and limited by the throughput of the top bottleneck tool only.

As far as time frame goes, I think that it depends on what you want to do with the data. Looking at weekly utilization reports for bottleneck tools is helpful in understanding which tools are likely to be contributing a lot to cycle time. But if all you want to know is: what am I producing relative to the most that I could be producing?, then you probably want to use a monthly definition, like the one above. Your relative use of the different tools will be accounted for somewhat because that will affect that "what you are producing" in the numerator. The top bottleneck tool only will be the one that determines the denominator.

Again, it depends on what your purpose is. It's helpful to know the utilization of several top bottleneck tools, especially if product mix changes cause the bottleneck to shift between tool groups. But the utilization of the fab at any point in time is limited by the bottleneck utilization.

Dov's Response:

I agree with your approach, but then the down side is that the actual utilization is subjected to the availability variance, which is not always what you may want to do. Regarding the tool that gates the capacity of the fab, I think this may be another interesting topic to discuss in the forum - which is the bottleneck tool? Is it the

temporarily most utilized tool (for any given period), or is it the one that is the most loaded according to capacity per mix calculation? How often should you ask the bottleneck question? When there is a change of mix, a change of process...? Say some tool is suffering from low available time (significantly lower then its planned goal), will this tool become the fab bottleneck, if it is the most utilized tool for a certain period? (This is of course if you measure utilization out of the actual available time). I also still think that some kind of average across the top utilized tools can also be a good reference for fab utilization. Looking only at the bottleneck is somewhat not sufficient, as all the rest of the fab can be either very loaded or very idle, which are totally different circumstances.

FabTime Response:
It sounds to us like what you're looking for is a broader measure of how relatively busy your fab tools are. You certainly could use some average value for that, across several bottlenecks. We would still argue that the bottleneck utilization is that one that reflects how busy the fab is as a whole, in the sense of what you're processing relative to the most that you could be processing. But if you want some other metric around the variability of the utilization across your top bottlenecks, you could certainly come up with something like that. It's not clear what the right thing is, though, because you don't want to have 100 tools that are each 90% utilized. It's much easier to run the fab if there are only a few very highly utilized tools, and the others are at something less.

What do the rest of you think?

Conquering WIP Bubbles

A question that FabTime is asked occasionally is "how do I manage WIP bubbles in the fab?". A WIP bubble is a large pile of WIP, usually in queue at a particular tool-group or small set of tool-groups. WIP bubbles occur due to a variety of causes, the most notable of which is extended downtime on a one-of-a-kind tool. In this article, we discuss common causes of WIP bubbles in wafer fabs, and propose methods to mitigate their effect.

Causes of WIP Bubbles

The primary cause of WIP bubbles, as indicated above, is extended unavailable time, especially on one-of-a-kind tools. This unavailable time is usually caused by unscheduled downtime, and might, for example, include time spent waiting for a part to be flown in from another location. It is also possible to create WIP bubbles through grouping of preventive maintenance events (such that the tool is unavailable due to PM for a protracted period). We have seen situations in which tools are unavailable to manufacturing for two or three days at a time because they are leased to outside customers in some way, or given over to extended engineering experiments. We have also seen cases where a fab chose to run a high-capacity

tool only on every other shift, creating small WIP bubbles every day, when the tool came back online.

Sharing a one-of-a-kind tool between 200mm and 300mm manufacturing (or 150mm and 200mm, for that matter) can also lead to WIP bubbles. For the wafer size that is not set up on the tool, the time during which the tool is set up for the other wafer size looks like downtime. No lots of that wafer size can move past the tool. Setup avoidance policies (in which a tool is only changed over if there are no lots in queue with matching setup ID) can have a similar effect.

Production control and management issues can also lead to WIP bubbles. If too much material is released into the line at once, it can take time for the WIP to spread out into the fab. Problems can also occur when a large amount of WIP is brought back into the line after an extended storage period (referred to as crib, bank, store, etc.) during which time the WIP was not manufacturing's responsibility. A change in the desired product mix can cause such WIP to re-appear, suddenly. Significant yield problems can also lead to starting too much WIP in a short time period, when lots must be released quickly to make up for the scrapped wafers.

Methods for Avoiding WIP Bubbles in the Future

The above causes of WIP bubbles imply the following suggestions for avoiding future WIP bubbles:

- Focus downtime improvement programs on bringing down the duration of unscheduled downtime events. This may require additional budget for spare parts or service contracts, but can be money well-spent, especially for tools that also limit fab capacity.

- Break up scheduled downtime events where practical (subject to limitations from qualification runs, etc.), especially on one-of-a-kind tools. Similarly, look at engineering time, or other times that the tool is taken out of the hands of manufacturing, and try to break those up into smaller time periods.

- Avoid having single path operations due to tool dedication, and, if at all possible, work to eliminate one of a kind tools. The latter is especially true for one of a kind tools that are shared between two wafer sizes.

- Keep starts into the line flowing as smoothly as possible.

We realize that the above changes may not be feasible (especially adding capacity at one of a kind tools), and so we proceed with discussion of identifying and recovering from WIP bubbles.

Indicators of WIP Bubbles

How do we know when we have a WIP bubble? Serious ones are quite obvious. There's a big pile of WIP waiting for a particular toolgroup, and WIP reports for that toolgroup, compared to historical data, show a significant spike. People from other areas complain about missing their moves goals, because they don't have enough WIP. Often we know to expect a WIP bubble because of something that happens in the fab (the start rate was tripled this week, or our one of a kind bottleneck has been down for three days, etc.) But it could also be useful to have an early, data-based of a developing WIP bubble, to allow earlier application of recovery measures.

The two primary indicators for an incipient WIP bubble at the tool group level are WIP in queue and inventory age. Inventory age for a lot at an operation is the time since that lot moved out of the previous operation. Inventory age can be aggregated across all lots in queue for a particular tool group, by either adding up or averaging the individual lot inventory ages. Rising average inventory age, in the presence of a constant amount of WIP, indicates that some WIP is not being moved. Rising

summed inventory age, regardless of WIP level, indicates that significant queue time is accumulating for the tool group. An example showing rising WIP at a tool group is shown below.

We can use Little's Law (discussed way back in Issue 1.3) to set targets for the average amount of WIP expected at a tool group. Little's Law is a mathematical relationship between cycle time, WIP, and start rate (or throughput rate, given line yield assumptions). Little's Law states that average WIP = cycle time * start rate (times a yield correction factor).

At the operation level, average WIP will equal historical average cycle time for the operation multiplied by arrival rate to the operation. For example, if we have an operation that usually takes 1 day for each lot to get through (including queue time and process time), and 100 lots per day arrive to that operation (on average), then we normally expect the WIP to be approximately 100 lots. Adding up across all of the operations processed on the tool group (under the current product mix) gives an indicator of how much total WIP we expect at the toolgroup. When the actual WIP exceeds this expected WIP by a significant percentage, a WIP bubble may be brewing. How significant, you ask? Well, you'll need to look at historical data for the tool group, and see how much the WIP ordinarily fluctuates

Methods for Managing WIP Bubbles

But in general, it's relatively easy to find WIP bubbles. What we need are techniques for managing them. This is not an easy question, because a wafer fab is such a complex system. However, we have a few suggestions:

1. Temporarily overstaff critical tool groups:

What you don't want, when trying to recover from a WIP bubble, is to lose any capacity on the toolgroups through which the WIP bubble is passing. This is a time for re-assigning operators, if at all possible, to ensure that no capacity is lost due to tools waiting for an operator to load or unload. Whenever you have a single operator running more than one tool, some capacity loss is likely due to both tools needing the operator at the same time. Therefore, you want to make sure

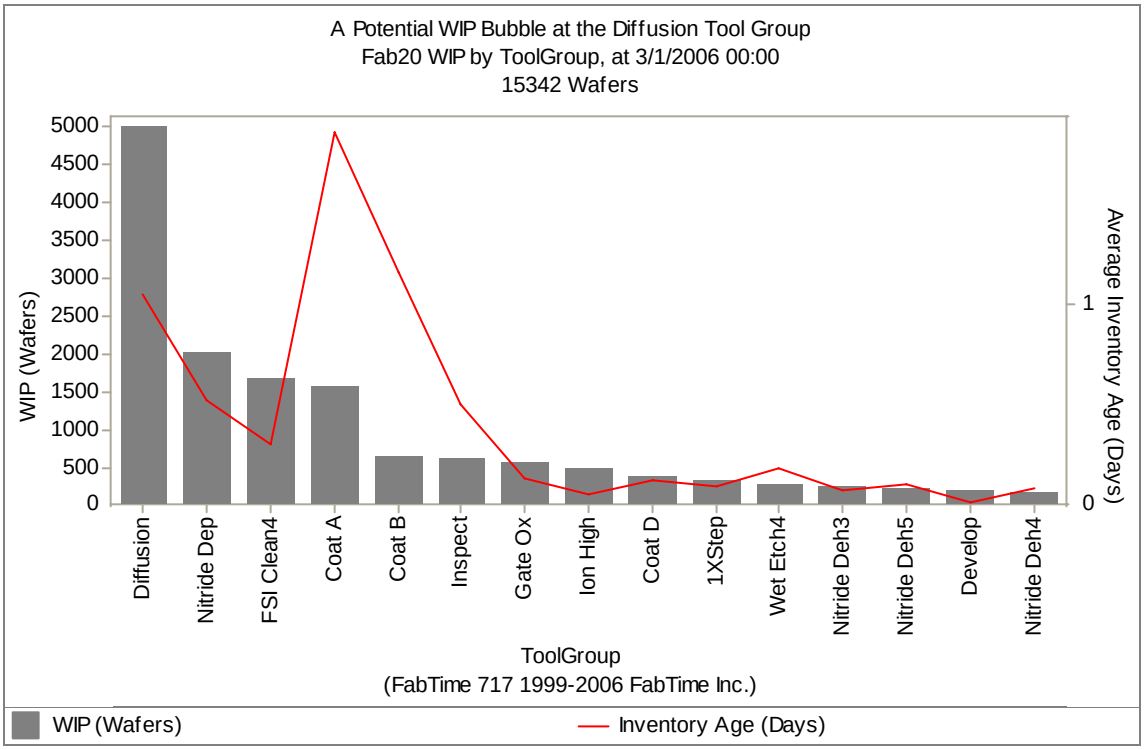

that your operators have good backup during critical recovery periods. This is also a good time to focus on OEE for these critical toolgroups – in addition to 100% utilization of available time, you want to ensure that OEE is as close to 100% as possible.

2. Look downstream based on amount of WIP before deciding which lots to process next.

One idea is to look down the path of the next several operations for each lot in queue, and choose lots to process that have little WIP ahead of them over the next few steps. This can be a tricky balancing issue, however. On the one hand, you want to avoid setups, so that you don't lose capacity during the bubble recovery period. This drives you to keep processing lots of the same recipe type. On the other hand, it's better to process lots that are going to different tools downstream (to spread out the WIP bubble). This implies processing lots of different recipe types, which could lead to additional setups.

3. Look downstream based on utilization of downstream tools.

One major problem with WIP bubbles is that they can lead to starvation of bottleneck tools. Because such tools have little spare capacity, it can be very difficult to make up such lost time. Therefore, another approach in making dispatch decisions during WIP bubble recovery is to prioritize the lots in queue by which ones will be headed to the bottleneck (or bottlenecks) soonest. While we don't want to send the whole WIP bubble along to a single bottleneck tool, if we can avoid it, we do want to keep an eye on the queue in front of the bottleneck, to prevent future starvation.

4. De-prioritize lots that will return to the same tool soon

Another thing to consider is that if you have a WIP bubble at a bottleneck tool, processing lots that will return to the tool very quickly (e.g. there's a clean operation, and then right back to this tool) might not be helpful. It's better to look at lots that will go to other tools downstream that are currently without WIP.

5. Turn off starts

In the case of a severe WIP bubble, it may be necessary to turn off starts into the fab. This allows the operators to work on the methods described above, and get WIP flowing more smoothly throughout the fab. Once the situation calms down, and the WIP level in the fab decreases, the cycle time will most likely decrease, too. Meaning that the lots started a bit later can still be processed in time, and with less risk of yield problems. This option requires considerable management fortitude, but is sometimes necessary.

Conclusions

WIP bubbles seem to be a fact of life in wafer fabs, especially for small fabs that have one of a kind toolgroups. While there are things we can do to try to forestall them, and to identify them as quickly as possible, fab attributes like long, unscheduled downtimes and rapidly shifting product mix make them hard to avoid completely. In this article we have reviewed some of the causes of WIP bubbles and suggested data-driven techniques for identifying them (in addition to the time-tested methods of looking for the pile of WIP, or waiting for the people running the fab to complain). We have also made a few suggestions for smoothing out WIP bubbles as quickly as possible. Three out of five of our suggestions involve making lot dispatching decisions that take into account future operations. We've come to believe that this type of forward-looking dispatching, though sometimes complex to implement, is necessary for today's fabs. The reentrant nature of the fab, the rapid product mix changes, and the presence of WIP bubbles create an environment in which purely local dispatching is probably not sufficient.

We hope that the discussion and suggestions in this article help you in managing (and avoiding) WIP bubbles in your fabs.

Closing Questions for FabTime Subscribers

How do you manage WIP bubbles in your fab? Do you have a data-driven early warning system, to notify you of out-of-balance WIP? Or do you rely on human knowledge and experience? Are you able to look downstream when you make your dispatching decisions, to smooth the WIP bubbles back out? Or is that a manual process, too?

Subscriber List

Total number of subscribers: 2674, from 465 companies and universities. 22 consultants.

Top 20 subscribing companies:

- Maxim Integrated Products (205)
- Intel Corporation (155)
- Micron Technology, Inc. (88)
- ATMEL (73)
- Analog Devices (70)
- Freescale Semiconductor (66)
- Infineon Technologies (65)
- STMicroelectronics (58)
- International Rectifier (57)
- Cypress Semiconductor (56)
- Texas Instruments (56)
- X-FAB Inc. (54)
- NXP Semiconductors (51)
- ON Semiconductor (50)
- TECH Semiconductor Singapore (50)
- Chartered Semiconductor Mfg (49)
- IBM (44)
- Spansion (37)
- Seagate Technology (33)
- BAE Systems (30)

Top 3 subscribing universities:

- Virginia Tech (11)
- Ben Guiron Univ. of the Negev (7)
- Arizona State University (6)

New companies and universities this month:

- Aquest Systems Corporation
- Fico Asia Sdn Bhd
- Michigan State University
- Tokyo Electron Ltd.

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 $7500, 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

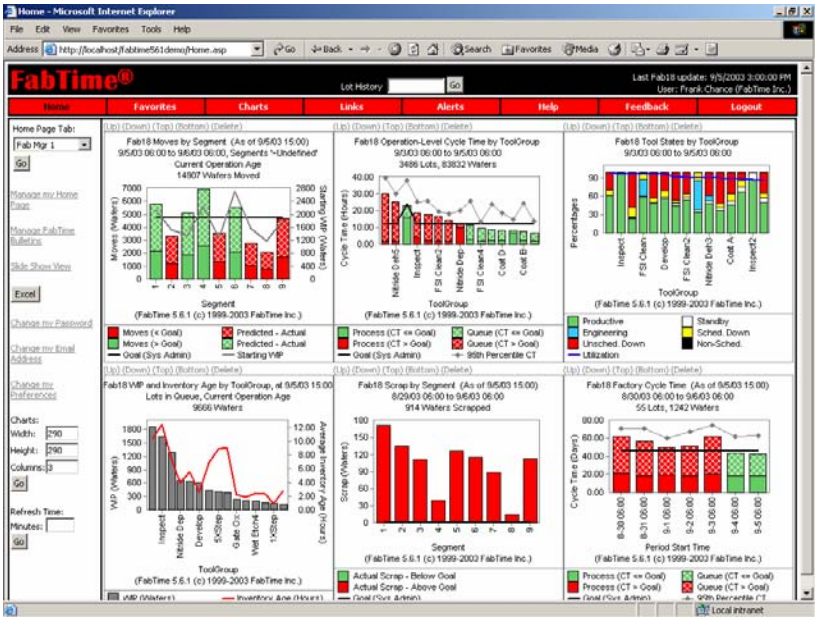

FabTime Benefits

- 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%.