FabTime Newsletter

Volume 23, No. 3 **June 2022**

Information

Publisher: FabTime Inc. FabTime sells cycle time management software for wafer fab managers. FabTime's mission is to help the people who run fabs improve performance by 1) helping them to understand the factors that drive fab performance and giving them the data to identify current improvement opportunities; 2) letting them control that data by setting parameters for their own charts, so they don't have to go back to IT every time they want a different piece of information; and 3) including them in a community of people around the world who are all working to drive better fab operations.

Editor: Jennifer Robinson

Contributors: Thomas Quarg (ams OSRAM Group); David Carmichael; and Hani Ofeck

Date: Wednesday, June 1, 2022 – Vol. 23, No. 3

Keywords: Delivery Performance; Dispatching; Predicted Cycle Time; Time Constrained Processing; Factory Ramp-Up; Capacity Planning

Table of Contents

Welcome

Welcome to Volume 23, Number 3 of the FabTime Cycle Time Management Newsletter. In this issue we have an announcement about a workshop we are holding for our software customers, as well as several articles from Jennifer's LinkedIn about the chip shortage and the high-tech labor shortage. Our FabTime tip of the month is about using fiscal calendars when configuring FabTime charts.

Our subscriber discussion forum is brimming with questions and responses on topics from queue time limits to WIP linearity to capacity planning to operator productivity. We also have a response to the previous issue on managing cycle time while ramping starts. Our main article this month was inspired by a discussion with Hani Ofeck about meeting delivery performance targets in overloaded fabs. We welcome your feedback on that, as well as on the other subscriber discussion topics.

Thanks for reading! – Jennifer, Frank, Lara, and the FabTime Team

Community News/Announcements

Workshop on Problem-Solving with FabTime (Software Customers Only)

Our June 14th User Group meeting will be replaced with a special two-hour session for software customers dedicated to problem-solving with FabTime. Participants should come with something to share, such as:

- A problem you solve using FabTime.
- A problem you are struggling to solve.
- A chart you would like to see in FabTime.

This session is primarily intended for regular participants in our User Group meetings (our user sponsors at each customer site). However, if you are a FabTime power user and would like to attend this meeting, please contact Frank Chance or Lara Nichols for further information. This meeting is only for people who use FabTime's software.

A Few Highlights from Jennifer's LinkedIn

Jennifer continues to share articles about business management, the semiconductor industry, and productivity improvement on her LinkedIn feed. Recent links have included:

- A WSJ article indicating that the chip shortage "has a new pain point: a lack of chips needed for the machines that make chips, industry executives say." This has significantly extended the lead time for equipment purchases, dimming hopes of a quick resolution to the chip manufacturing shortages. [LinkedIn post.]
- An article from SEMI about a neat and timely program extending SEMI's workforce development efforts into K-12 classrooms. "The new program – High Tech U in the Classroom – will provide teachers with engaging, hands-on STEM kits and curriculum with clear relevance to the microelectronics industry. Students will not only learn how to code a microcontroller to test and prototype a medtech innovation, but they will also connect what they learned to contemporary real-world innovations powered by the industry, such as the fall detection feature on their grandparent's smartwatch." [LinkedIn post.]
- An interesting presentation by Bill Wiseman about the future of the US semiconductor industry, including the expected skilled labor shortage and a predicted drawing out of the chip shortage in more mature nodes. The presentation was shared by Thomas Beeg on his Factory Physics and Automation blog. [LinkedIn post.]
- A deep dive by Adele Hars in an article from Semiconductor Engineering into the wafer shortage that is going to continue influencing the chip shortage in the coming years, especially at the 200mm level. This caused Jennifer to muse that "maybe some of the wafer manufacturers like Okmetic could use a session of FabTime's remote cycle time management course for some tips on reducing variability to allow increased throughput." [LinkedIn post.]
- A WSJ article about TSMC considering building a new advanced fab in Singapore. [LinkedIn post.]
- A report by Marie Baca for Semiconductor Engineering that asks whether the recent push by big-name tech companies like Apple and Google to bring their chip design efforts in house will ultimately help ease staffing shortages by attracting more workers to the semiconductor industry. [LinkedIn post.]

For more industry news, connect with Jennifer on LinkedIn. FabTime welcomes the opportunity to publish community announcements, including calls for papers. Send them to newsletter@FabTime.com.

FabTime® User Tip of the Month

Use Fiscal Calendars to Set Dates for FabTime Charts

One of the earliest capabilities of FabTime is the ability for end users to select their own date range and period length for any chart. You can do this in the current version of FabTime by using the "Custom" setting in the "Dates" section to select "From" and "To" dates and, if applicable, "Period" and "Sub-Period" lengths. Normally, range-based charts default to show the most recent week, with 24-hour periods. Clicking on "From" or "To" brings up a calendar you can use to change that date. Select your date of interest and click "Done" to exit the calendar. Press "Go" to update the chart.

If fiscal period data is integrated with FabTime at your site, you can also use fiscal calendars to specify the time range for a chart. You can select the current period, the previous period, some range of periods, or a range defined relative to the current period (e.g., 3 weeks ago to 2 weeks ago). One benefit of this is that if you use fiscal periods, you may not need to set up Auto Slide for your home page tabs (as you do with legacy/custom dates). You can create a chart that shows, for example, today's moves by hour and add it to a tab. When you come in tomorrow, or any other day, the chart will automatically show that day's moves.

Here are some examples:

To create a Moves Trend chart showing hourly moves for the current day:

1. Select the Moves Trend chart from the Charts page (or use the Search box).
2. Change the top drop-down within the Dates section in the left-hand pane to "Current". Once you select "Current" another drop-down appears to the right, defaulting to "Shift". Set that to "Day". Modify the start time in the next drop-down that appears to the right, if needed (this defaults to your standard start time of day).
3. Modify the period length in the "Len:" box, changing the units to specify 1 hour. Add a sub-period if desired. On the Moves Trend chart, the sub-period is used for point-average WIP calculations and is optional. The sub-period must be less than or equal to the period length.
4. Press "Enter" after entering the period length or press the "Go" button to update the chart. If added to a Home Page tab, this chart will always display data for the current day. No Auto Slide is needed.
5. To instead display the same data for the prior production day, change the top drop-down within the Dates section from "Current" to "Prior" and press "Go."

To create a Shipments Trend chart showing shipments by week for a specific range of work-weeks:

1. Select the Shipments Trend chart from the Charts page (or use the Search box).
2. Change the top drop-down within the Dates section from "Custom" to "Range" and then set the drop-down that appears to the right to "Week". The drop-downs immediately below (for setting the "From" and "To" dates) will update to display work-weeks (as defined in your fiscal calendar data). Select the "From" workweek from the left-hand drop-down and the "To" workweek from the right-hand drop-down.
3. Modify the period length by selecting "Weeks" from the "Len" drop-down. A value is only needed if you select "Hours". If you select "Weeks", FabTime just displays one week of data at a time. Press the "Go" button to update the chart.

Note: If your fiscal calendar data includes future periods, you can select those. When you add the chart to your home page, the data for the future periods will gradually fill in over time, as available. The date range will remain fixed for the chart, unless you use a manual or Auto Slide for the home page tab. This setting is useful for showing things like weekly shipments to date within a fiscal quarter.

To create a Shipments Trend chart that always shows shipments by week for the current work-week and each of the previous three work-weeks:

1. Select the Shipments Trend chart from the Charts page (or use the Search box).
2. Change the top Dates drop-down to "Relative" then set the next drop-down to "Weeks". The drop-downs in the next row will populate with "-3 week", "-2 week", etc. "0 week" is the current fiscal week. To include the current week and the three previous weeks, set the left-hand drop-down to "-3 week" and leave the right-hand drop-down at "0 week". Set the "Len" drop-down to "Weeks". Press the "Go" button to update the chart, as shown to the right.
3. When added to a home page tab, this chart will always display the current week together with the previous 3 weeks. No AutoSlide is needed.

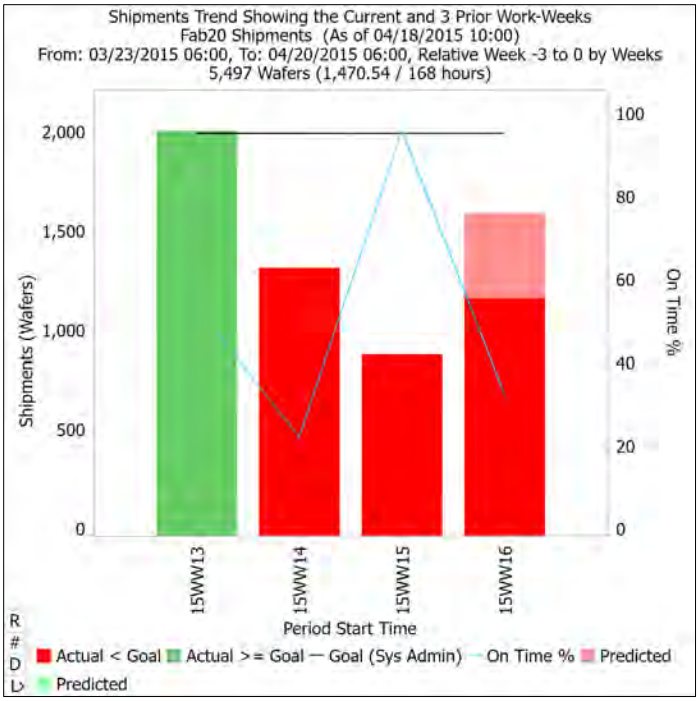

Note that for point-in-time charts such as the WIP Lot List, the fiscal calendars (work-weeks, etc.) are not displayed. When you place your cursor in the single "Date" field, a calendar does pop up to let you select a specific date. If you leave the Date field blank, the chart will always show the most recent data, including when it is on a home page tab. Again, no Auto Slide is needed in that case.

We hope you find this tip useful.

Subscribe to the separate Tip of the Month email list (with additional discussion for customers only) here: http://www.fabtime.com/tip-of-the-month.php. Thanks!

Subscriber Discussion Forum

WIP Linearity

David Carmichael wrote in response to last month's subscriber discussion question on WIP linearity. "I was interested in the section about line linearity because it is in such direct conflict with many other line requirements.

I think the desire for line WIP linearity is a gut-feel management thing (like keeping all tools busy or running tools at > 100%) and probably stems from a belief that a fab is like a car assembly line.

It isn't.

WIP linearity in a fab is a goal to be avoided at all costs in my opinion. The exact opposite is in fact required – excess WIP before each of the constraints (or near-constraints), dependent upon mix, is essential for maximum throughput. The complete absence of WIP where there is extreme excess capacity should not cause anyone any heartache.

True linearity is actually a combination of having enough WIP to keep constraint tools busy coupled with as linear as possible a **Critical Ratio** for each lot, with as little wild divergence from 1.0 as is possible.

The CR part of this is trivial to graph and shows with dramatic clarity where lots have been held up in the past or have been moving too fast, and this is throughout their lives to date, not just where they are now.

Of course, the CR of a lot can change due to external forces. In a make-to-order fab, a customer changes an order quantity or required date and bang, all of the lots pegged to that order change their CR. This ought to be of great interest to any fab (though I can't remember anyone in fab management or industrial engineering ever talking about it).

Anyway, these are my thoughts on the subject of so-called linearity."

FabTime Response: We take your point about the importance of keeping sufficient WIP in front of constraint and near-constraint tools. We would never advocate for balancing the WIP by tool in the fab. However, we believe that balancing WIP by segment of the line can be useful. Segments are linear chunks of the line, often about a week long. Balancing the WIP across segments helps a fab to maintain a steady flow of output, and doesn't require balancing the WIP across the tools within any segment. In fact, most constraint tools will have WIP waiting from multiple segments, because of the reentrant nature of the line. But your reminder about keeping WIP in front of constraint tools does suggest a need to refrain from particularly small sub-segments for this type of analysis. We will also certainly pass along your suggestion of using Critical Ratio by lot as a proxy for line linearity. Thank you so much for taking the time to respond to our anonymous subscriber's question.

Managing Cycle Time while Ramping Starts

Thomas Quarg from ams OSRAM Group wrote in response to last month's main article. He said "I did in my former work several analyses about start and ramp rates. At all times this was a highly political discussion for new fabs, as also for "filling up" fabs, based on installed capacity. Here are some short notes:

It starts with the definition of a ramp rate (RR). If the ramp rate is 100/week the incremental starts volume after one quarter would be 13*13*100 = 16.9k per quarter (shown in yellow). During the first quarter the additional starts would be 9.1k, as shown.

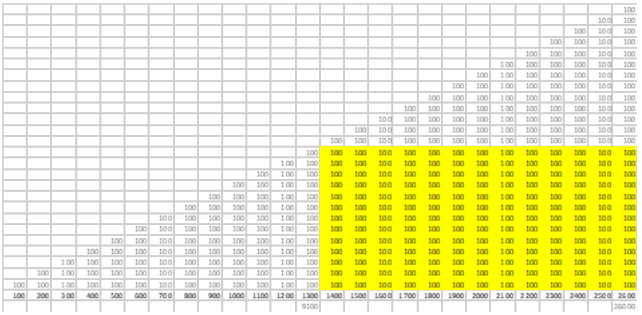

If the fab is already filled and we need to wake up the tools, then we need to consider the fact that a homogeneous starts profile in a fab with different technologies will not generate a homogeneous out profile. Do you agree? [**FabTime Response:** Yes, because turning on of different tools results in an uneven change to fab capacity and because of the different length of the different flows.] Therefore we can decide what behavior we would like to have, based on our customer requests and back-end capacity.

Here are three choices: Outs Driven Starts, Starts Driven Outs, and Flex Cycle Time Shifted Starts.

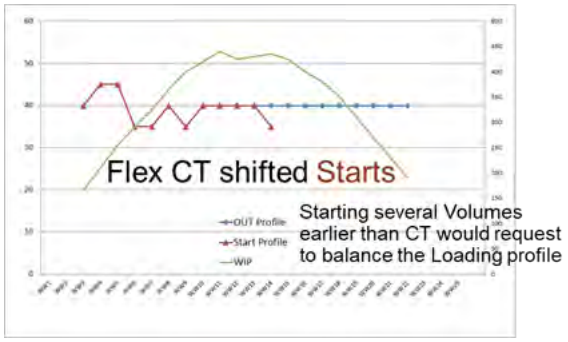

When increasing the start rate, we must always consider the capacity of the whole fab, including re-entrant flows. Dynamic simulation can help here. (Qualification, Test Wafers, etc. need to be planned as well – as you described.)

Ramping a new fab, is much more complicated, because the capacity of the tool installation and personnel for performing qualification need to be calculated.

Why is a ramp rate higher than 100 WSW on average in mode 1A (1A= new tools required and continuous ramp rate) critical? (Q's law 😊)

- Frequency of process tool installation.
- Frequency of metrology and support tool installation.
- Tool logistics
  - Transportation path / elevator / "black level" maneuver and support tool density, craft priorities and sequences (air-channels vs. gas pipes, vs. water pipes)
- Tool qualification
  - Providing test wafer and test wafer volume
  - Measurement capacity
  - Lab capacity
- Manpower: hook up, suppliers, internal specialists
- Risk assessment, concomitance work, removal of obsolete resources afterward

There is also the so-called overlapping effect from different teams. This means that you would need 3-4 hook up teams to work at the same time as well as qualification teams from suppliers and your own qualification teams. That is a short overview only with some points to consider and not completed – have fun!

FabTime Response: This is great stuff, Thomas. Thank you! We don't have much experience with the implementation details of a capacity ramp and appreciate you sharing your experience.

Queue Time Limits for Yield Protection

An **anonymous subscriber** wrote: "Thanks for sharing the FabTime cycle time management newsletter. Actually, lots of topics in the newsletter are similar to what we are doing. I have another question that has been bothering us for a long time. It's about the ROI of a newly added queue time limit. Sometimes the factory's technical team will ask us to add new queue time limits within the process flow to control process margin in order to mitigate specific yield loss. For example:

Step A (tool A) occurs prior to Step B (Tool B)

- Before. No queue time limit going from A to B: more lots can be pushed to step B and more WIP can be waiting in front of tool B.
- After. Let's say, 5-hour queue time limit is inserted between A and B. We must slow down lots leaving tool A since lots can't be waiting too long in front of tool B due to queue time concerns.

This results in lots moving slowly before A and B, increases lot cycle time and even decreases outs in the worst case.

Do you have experience about how to measure the cycle time loss or wafer outs loss by adding queue time limits? Is there any formula that can be shared? Thanks."

FabTime Response: We are certainly familiar with these types of time constraints in wafer fab production. We've seen them called time bound sequences, time links, and time constraints between process steps. What we generally see is people doing one of two things:

1. Hold the WIP in front of Step A until Tool B is free, or nearly free. Often a reservation on Tool B is required in the dispatch system before the lot is released to run on Tool A. This, of course, can lead to capacity loss on Tool B, if the WIP is not released soon enough from Step A.
2. Use an alert to warn operators when a lot waiting for Tool B is in danger of exceeding the time limit, so that they can process that lot right away.

It makes sense to be conservative in one's operational approach here, to avoid the risk of exceeding the time limit. This is because, in addition to the yield issues, if you end up missing the time limit and having to re-process the lot on Tool A, there will be capacity loss at Tool A (non-value-added time reworking the wafer).

Unfortunately, modeling the behavior of these types of systems is quite complex (especially in the case where there is another operation between operations A and B, as happens sometimes in practice). There is no simple formula that we know of that captures the impact of time constraints on overall fab cycle time or wafer outs. You could do some analysis of time constraints using simulation models, but the outputs will likely be quite case-specific.

Other subscribers, do have any insights or experiences to share regarding time constraints / queue time limits between process steps in the fab?

Operator Productivity

An anonymous subscriber wrote: "Just wondering if you have information on Operator Productivity and how to increase it? I would also be interested if you have anything on Operator Staffing models."

FabTime Response: We do have a bit of information about how people measure operator productivity. We wrote an article about this many years ago (see Issue 4.06, In-Depth Guide to Operators and Cycle Time). We also wrote about the impact of operators on cycle time recently, in Issue 22.05.

We also discussed a metric designed to overcome productivity losses during shift change in Issue 14.01, with some subscriber discussion about that in 14.02. We implemented that metric (Earned Plan Hours) in our software, and some of our customers do use that. [See the FabTime Newsletter Archive to download all issues mentioned. The current password is "FabTimeCommunity".]

We do not, however, have any staffing models. There is certainly considerable interest in making operators more productive these days, because of staffing shortages. We wonder if any other subscribers have thoughts to share on measuring and improving operator productivity or using staffing models. We would be happy to share more about this topic in the next issue.

Capacity Planning for Cluster Tools

An anonymous subscriber wrote: "I have a question regarding capacity planning. We have a cluster tool for which we currently have three process chambers. Each chamber handles its own operations or process steps. We have two load locks which prevent us from running all three chambers at the same time. We have two chambers that run parallel processing (M1 and M2) with each other and then the remaining chamber is on its own (M4). The question is whether we count each of these chambers as a separate tool or count M4 as its own tool and then ratio out M1 and M2 to equal a tool.

FabTime Response: Cluster tools are one of the biggest challenges in capacity planning. When people include cluster tools in our software (which is a reporting system but did at one time have a capacity model), they typically include the chambers as individual tools and include the load locks as tools. Moves get tracked on the chambers, but downtime is sometimes tracked on the load locks, as representative of the entire tools.

In this case, we would treat M4 as a separate tool. We would probably treat M1 and M2 separately in the model, too, but downrate them using a loss factor to account for the lost capacity due to sharing of the load lock. You would know how much time you need on each of those chambers to process a lot, summing across the operations, but would set the available time of each chamber lower to reflect lost time when there's no load lock available. However, we have not worked on capacity planning at this level of detail in many years. Do any other subscribers have input on handling this issue?

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

Delivery Performance in Overloaded Wafer Fabs

Introduction

Some fabs are struggling right now with using delivery performance as a metric in highly capacity-constrained environments. Often, fabs that have good cycle time performance also have good on-time delivery (OTD) performance. When a fab has sufficient capacity, as well as a reasonably tight distribution of cycle times, cycle time estimates used for commit dates can be set to the 95th or 98th percentile of actual cycle time values, resulting in few lots missing their delivery dates.

In the current capacity-constrained environment, many fabs are operating without their usual capacity buffers (close to or at 100% effective utilization on some tool groups). Actual cycle times (and the width of the distribution of those cycle times) are increasing significantly (as happens when fabs get into the steep portion of the operating curve). Unfortunately, due to pressure from customers, cycle time estimates for commit dates are not necessarily being increased to match the situation in the fab. As a result of this, as well as of longer cycle times in general, more lots are missing their due dates.

This results in unhappy customers, of course. But it also poses a production management problem in many fabs, especially those running a Critical Ratio dispatch rule. Critical Ratio (CR) prioritizes lots based on their current performance relative to due date. What happens to CR when a significant portion of the lots are behind schedule? We've done prior research that found that in the presence of many late lots, critical ratio can lead to oscillating WIP bubbles in the fab (see Issue 15.01). We've also discussed in the newsletter's subscriber discussion forum (see Issues 22.04 and 22.05) the fact that in the reentrant environment of a fab there may be higher variance in CR values for lots that are further along in the line.

Other issues to keep in mind are that a fab can have many different technologies that have different cycle times, and that fabs have many different customers who may receive different commit dates. Some customers may exert pressure that leads to their lots being made hot, increasing variability in the fab. These factors increase complexity under the best of circumstances, and especially right now. In this article, we discuss delivery performance in capacity constrained fabs, and how fabs may be able to better manage this issue.

What Happens When You Use CR with Overly Optimistic Commit Dates?

Oliver Rose wrote about issues relating to the use of CR in wafer fabs in a 2002 Winter Simulation Conference paper. He noted that "determining appropriate due dates for this rule is a critical task." He performed full fab simulation experiments that looked at "whether there is a due date setting that minimizes the average cycle times of the lots". He found that when commit dates are set too low (especially at high utilization rates), average cycle time increases significantly (becoming much higher than it would have been under first-in-first-out dispatching) and on-time delivery performance is also poor. CR only did better than FIFO in terms of OTD in his study when due dates were set correctly (that is, with a long enough planned cycle time).

What seems to happen is that when the due dates are set too low, as the fab fills up, lots later in the process flow start exceeding their due dates. At highly loaded tools, these late lots are prioritized ahead of more recently arrivals waiting for layers closer to the start of the flow. Because the tools are so busy, the more recently released lots that are not yet late can end up stuck in queue until they are late. This increases variability in the fab, especially at batch tools, which of course increases average cycle time. Long downtimes on key tools exacerbates the situation (as it does most situations). For more detailed explanations of this behavior, see the Rose paper as well as FabTime Newsletter 15.01.

What Can You Do to Mitigate the Due Date / CR Problem?

One obvious recommendation is to set more realistic commit dates if you are going to use CR (or any other due-date-based dispatch rule). Rose recommends using a simulation model to estimate average cycle time under FIFO, and using that plus a small buffer to set due dates used in CR. In practice, most fabs today don't have a detailed enough, updated simulation model to do this. But fabs do have data on actual average cycle times. They should also have data on what x-factor they are getting for lots right now. Your fab will get better results from CR (in terms of OTD and average cycle time) if you use the actual x-factor to set commit dates (by major flow, if you have it) than you'll get using out-of-date, unrealistic historical values.

This option may not be within the control of the people running the fab, however, depending on where due dates are set and what customer expectations are.

Another possibility is to modify the dispatch rule to avoid using critical ratio or other pure due-date focused rules. If your commit dates are so tight relative to your current cycle time that all the lots end up late anyway, a first-pass approach could be to use priority-FIFO. This will reduce variability relative to CR and thus should reduce average cycle times in the fab. This means that at least some of the lots will likely be out on time (especially the higher priority lots).

A better long-term approach may be to use a dispatch rule that reduces the variance of cycle time. When the variance of the cycle time distribution is reduced, attaining OTD is easier. A 2013 Winter Simulation Conference paper by Zhou and Rose looks at dispatch rules for reducing the variance of fab cycle time. They use a two-layer approach in which the first layer selects lots to fulfill a WIP balance requirement and the second layer selects lots from within the first layer based on other performance measures. When due date performance is used in that second layer (either explicitly or implicitly via x-factors), the combined rule is successful (in simulation studies) in improving cycle time variance and delivery performance.

What Else Can You Do to Improve Delivery Performance?

If you want better delivery performance in general, you need to tighten the distribution of lot cycle times. You tighten distributions by reducing variability. In addition to the points made above about dispatch rules, to reduce the variability of lot-to-lot cycle times you should:

- Cut down on the number of hot lots in the fab, especially hand-carry lots. Hot lots add variability that drives up cycle time. Resist pressure to make ever-increasing quantities of lots hot. If everything is a hot lot then nothing is a hot lot.
- Focus your equipment maintenance team on doing everything possible to reduce the duration of the longest downtime events. Use Green-to-Green (see Issue 20.02) charts or similar metrics that focus on the variability of downtime, with emphasis on eliminating the longest periods of unavailable time (including Engineering time). This will help reduce WIP bubbles in the line and make lot-to-lot cycle times more consistent.
- Set caps on setup minimization policies. In heavily loaded fabs that have very strict setup avoidance policies, you can have low-volume lots that sit… forever. You need a cap on how long a lot can wait before you force a setup. Similarly, for heavily loaded batch tools. Yes, it's generally good to run full batches when they are available (without holding tools idle to wait for full batches). But you do need

to make sure that your low volume lots get processed sometime, even if you occasionally don't run the largest possible batch.

- Plan starts to minimize setups and maximize batch sizes. For example, start lots of the same product/product family together to reduce setups. Start lots together that share the same recipe at the first visit to a furnace, to reduce time spent waiting to form early batches.
- Increase tool qualification. Review operations that have a single qualified tool, and even operations that only have two or three qualified tools. Work to qualify additional tools where possible (to increase flexibility). This will help to keep lots at those single or dual path steps from being disproportionately impacted by tool downtime events. By keeping those lots moving, you help keep their cycle times closer to that of other lots, tightening the cycle time distribution. (See Issue 20.05: The Impact of Tool Qualification on Cycle Time).
- Use Dynamic X-Factor (DXF) to identify and monitor short-term bottlenecks affecting specific products or specific tools. Responding to these problems quickly helps keep lot-to-lot cycle times closer together. For example, if DXF for a product spikes, even though overall fab DXF remains within expected limits, this indicates that the product in question is not getting processed and will end up with higher cycle time. DXF is a point estimate that measures total WIP divided by WIP that is currently running on tools. See Issues 4.08. 5.03, and 9.04.

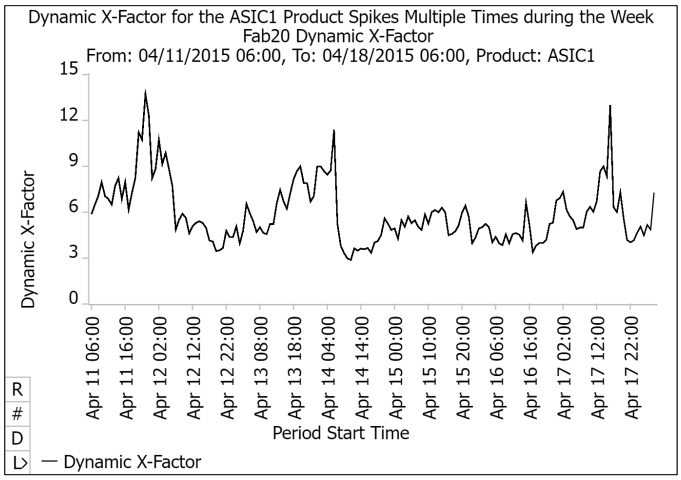

- Where operators have a choice about what lot they select from the dispatch list, avoid measuring their performance solely on number of moves. Otherwise, you can run into situations where the lots with longer process times, or that require setups, are quietly not selected.
- Get lots off hold as quickly as possible. Where you have future holds, make sure there's a backup available if the person who entered the future hold is unavailable.

You'll improve delivery performance if you can help bring in the outliers, those lots that end up with disproportionately high cycle times.

Conclusions

In today's capacity-constrained environment, many fabs are overloaded. They have squeezed capacity buffers to try to get more wafers out and as a result seen average cycle times increase. Their customers are also under pressure from further along in the supply chain. Automotive factories were particularly hard hit initially by this, but many categories of products have since been affected. These customers are pushing fabs to maintain the same commit dates that they offered previously, even though it is much harder to meet these dates at higher fab utilizations.

Fabs are operating with a high percentage of lots in process that are already behind schedule. When those fabs use due-date-based dispatch rules, performance can degrade rapidly. Using Critical Ratio in a situation where the due dates are set unrealistically low has been shown to increase average cycle time and harm on-time delivery. We recommend that fabs that are using CR set more realistic due dates. If this is not possible, then the safer course is to turn to dispatch rules that either don't rely on due dates or incorporate them in a cycle time variance reduction framework.

In addition to dispatch rule selection, fabs should focus on operating practices that help reduce the distribution of lot-to-lot cycle time. These include minimizing the number of hot lots and tracking and reducing the longest periods of unavailable time on tools. They should also ensure that lot volume lots keep up with other lots (at least to some degree) by paying attention to batching and setup minimization policies.

This is a challenging time to be responsible for meeting customer expectations regarding cycle time. There's no quick and painless solution. But we do hope that these suggestions help.

Closing Questions for Newsletter Subscribers

Have you run into problems lately with customer commit dates that don't reflect the current cycle time performance of the fab? What have you done to improve the situation? Have you tried dispatch rules to help with line balance or cycle time variance reduction? What other approaches have worked in your fab for reducing the variability of the cycle time distribution?

Acknowledgements

This article was inspired by a discussion with **Hani Ofeck**. Hani also suggested some of the above recommendations. We are grateful for her input.

Further Reading

- O. Rose, "Some Issues of the Critical Ration Dispatch Rule in Semiconductor Manufacturing," *Proceedings of the 2002 Winter Simulation Conference*, E. Yücesan, C.-H. Chen, J. L. Snowdon, and J. M. Charnes, editors, 2002. Available for download from the 2002 WinterSim archive.
- Zhou, Z. and O. Rose, "Cycle Time Variance Minimization for WIP Balance Approaches in Wafer Fabs," *Proceedings of the 2013 Winter Simulation Conference*, R. Pasupathy, S.-H. Kim, A. Tolk, R. Hill, and M. E. Kuhl, editors, 2013. Available for download from the 2013 WinterSim archive.
- All FabTime newsletter issues referenced above are available for subscriber download from the FabTime newsletter archive. The current password is "FabTimeCommunity".

Subscriber List

Total number of subscribers: 2904

Top 20 subscribing companies:

- onsemi (192)
- Infineon Technologies (147)
- Analog Devices (136)
- Intel Corporation (126)
- Micron Technology, Inc. (119)
- GlobalFoundries (103)
- Microchip Technology (102)
- NXP Semiconductors (83)
- Skyworks Solutions, Inc. (70)
- Carsem M Sdn Bhd (69)
- STMicroelectronics (69)
- Western Digital Corporation Inc. (63)
- Seagate Technology (57)
- Texas Instruments (56)
- X-FAB Inc. (52)
- Wolfspeed, Inc. (43)
- Qualcomm (38)
- Tower Semiconductor (31)
- Honeywell (30)

Top 3 subscribing universities:

- Ecole des Mines de Saint-Etienne (EMSE) (9)
- Arizona State University (7)
- Virginia Tech (7)

New companies and universities this month:

- Boeing
- Booz Allen Hamilton
- Lund University
- Photronics
- QuantumScape

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 to the newsletter. Past issues of the newsletter are now available in PDF for download by newsletter subscribers from FabTime's website. To request the current password, email your request to Jennifer.Robinson@FabTime.com.

To subscribe to the newsletter, send email to newsletter@FabTime.com, or visit our website. 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: If you would like more information about our web-based dashboard for improving fab cycle times, please visit our website. A sample home page and a sample page from FabTime's new Charts menu are shown below.