FabTime Newsletter

Volume 24, No. 5 November/December 2023

Information

Publisher: FabTime helps wafer fabs drive improvement via our factory insight engine and our cycle time management course. The FabTime factory insight engine empowers fab managers and production leads to evaluate their facility conditions in real-time. Unlike generic analytics tools, FabTime's dynamic interface includes adaptable pre-set charts and dashboards, ensuring decision makers have access to reports on-demand without burdening IT departments with reporting responsibilities. Contact us for a demo.

Editor: Jennifer Robinson

Date: Wednesday, November 15, 2023 – Vol. 24, No. 5

Contributors: Frank Michael Schulze (X-Fab); Marc Engel (Agileo Automation); Andy Beers (ABM Consulting)

Keywords: Factory Behavior; Wafer Size; Factory Size; Variability; Hot Lots; Material Handling

Table of Contents

Welcome

Welcome to Volume 24, Number 5 of the FabTime Cycle Time Management Newsletter. In this issue, we have a quick recap of the recent Fab Owners Alliance meeting, an announcement about FabTime's new chart engine, and a variety of articles from Jennifer's LinkedIn. Our software tip of the month is about using our Process Time Lot List and Tool State Gantt charts to analyze lot cascading, with the goal of driving improvements in tool utilization.

In our subscriber discussion forum, we share poll results on a question about top contributors to fab variability. We also share a subscriber question about managing the cycle time of external processes and a response to our recent software tip of the month about using smaller carts. We also share an example of a chart requested by a prospective customer for assessing the cycle time impact of hot lots.

In our relatively brief main article, we respond to a question from a subscriber who asked which of our recommendations are more applicable to 300mm fabs vs. less automated 200mm and smaller factories. We categorize some operational differences between these factories and share a few recommendations for 300mm (and highly automated 200mm) fabs.

Thanks for reading! This will be the final newsletter of 2023. Wishing you all a joyful holiday season and a happy and productive 2024 – Jennifer, Frank, Lara, and the FabTime Team

Community News/Announcements

FOA Golf Sponsorship

In mid-October, we attended the final Fab Owners Alliance meeting of the year in Lubbock, Texas. FabTime was pleased to sponsor the golf event on the first day, where our Director of Customer Success, Elaine Jacobson (pictured right) participated alongside other FOA members. FabTime Co-founder and COO, Jennifer Robinson (pictured left), joined Elaine for the full membership meeting on the second day. It was great seeing familiar faces and meeting some of our newsletter subscribers for the first time. We look forward to what next year has in store for the FOA!

New Chart Engine Releasing to Early Adopters Soon

FabTime's new chart engine continues our multi-year evolution to meet our end-users' request that FabTime be more powerful and more user-friendly. End-users wanted more advanced functionality in FabTime, to cut the time wasted exporting data to Excel or another package for analysis. But they also wanted FabTime to be more friendly to new users, so a wider audience could benefit from it. We've made major improvements each year in that direction:

- 2020: Replaced our static data table with AG Grid, bringing client-side data table sorting, filtering, and pivoting into the browser.
- 2021: Overhauled the user interface, charts menu, and search functionality, radically improving ease of use.
- 2022: Added free-form home page tabs and sparkline charts, for more flexible dashboarding.
- 2023: Migrated to a modern chart engine, enabling live charts (chart image responds automatically to data grid filtering and sorting), boxplots, histograms, and multi-pane charts that present both detailed data and summary statistics.

These changes improve our customers' ability to work uninterrupted in FabTime and lessen the need to export static datasets for analysis via different platforms. Having live charts linked to the data grid allows users to dig into the data and see graphical answers immediately.

FabTime customers interested in becoming early adopters for the new charting engine should contact support@FabTime.com. Prospective customers who would like to see a demo should visit our Request A Demo web page.

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 repost of an announcement by ATREG, Inc. about the purchase of the Nexperia-owned wafer fab in Newport, Wales by Vishay Intertechnology, Inc. Here's the full news release. FabTime joins many in the industry in expressing gratitude at seeing this site in line for ownership stability after a tumultuous few years. The hardworking Newport team deserves it, and Vishay will be lucky to have them in the family. [LinkedIn Post.]

- A WSJ deep dive into the impact of semiconductor equipment manufacturers on technical progress within the industry. Of course, it's not true (per the title of the article: The Most Important Tech Company You've Never Heard of Is a Major Reason Computers Keep Getting Faster) that most subscribers here "have never heard of" Applied Materials, but we still think you'll find it interesting. [LinkedIn Post.]
- A New York Times piece about an announcement by the Department of Defense that it was awarding $238 million to set up eight semiconductor manufacturing hubs around the US for promoting innovation in the semiconductor industry. The funds are one of the earliest releases of the Chips Act. "The awards announced Wednesday will go to research institutes, consortiums and universities located in New York, Arizona, Indiana, Ohio, California, North Carolina and Massachusetts, defense officials said." [LinkedIn Post.] In related news, the WSJ reports that Intel is in the lead to get billions in funding for secure defense-chip facilities. [LinkedIn Post.]
- A piece by the WSJ's Greg Ip positing that the problem with American manufacturing is that factories aren't productive enough. The semiconductor industry is one of the sectors he discusses (see late in the article), noting that: "In semiconductors, U.S. companies still dominate design, while steadily ceding production to Asia. Intel is the last major U.S. firm that both designs and makes chips, but its manufacturing capabilities have fallen far behind Taiwan Semiconductor Manufacturing Co... Even constructing a fab's clean room involves pouring the concrete and welding the pipes in just such a way to avoid tiny imprecisions that ultimately reduce yields." [LinkedIn Post.]
- A Bloomberg article about how "Japan's government-backed Rapidus Corporation is lobbying chipmakers, their suppliers and research institutions to set up shop on the island of Hokkaido in a bid to make the remote region a center of semiconductor innovation within a decade... Instead of competing against global chip giants that make large volumes of all-purpose chips, the Tokyo-based company will focus on pioneering specialized chips, such as low-power-consumption AI chips." [LinkedIn Post.]
- A blog post at Factory Physics and Automation by FabTime friend, Thomas Beeg of Fabmatics USA, about time links between process steps. Please do check it out. We learned some new things about this common "fab performance detractor" (the main topic in Issue 24.02 earlier this year), and we're sure you will, too. [LinkedIn Post.]
- A fun piece for Industrial Engineers in the WSJ exploring different methods for reducing the cycle time of the airplane boarding process. As anyone who reads this newsletter would expect, a key issue is the variability in how long different individuals take to stow their belongings. The Southwest system works well compared to traditional methods because passengers spread out themselves once they are on the plane, which sounds like something of a WIP smoothing policy. [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 Jennifer.Robinson@FabTime.com.

FabTime® Software Tip of the Month

Use the Process Time Lot List to Gain Insights into Cascading of Lots

Our Director of Customer Success, Elaine Jacobson, suggested this month's software tip based on work that she's doing with customers. You can use FabTime's new Process Time Lot List chart, together with the Tool State Gantt and WIP Lot List charts, to assess how well your fab is doing with cascading on key tools. In this context, cascading refers to loading lots onto a tool in the most efficient way to minimize any idle time between the lots. Cascade analysis can help identify inefficiencies in tool loading when there's a pile of WIP waiting and help drive improvements in tool utilization.

To do this type of analysis:

1. Enter "Process Time Lot List" in the "Search FabTime" box and select the chart.
2. Specify your time period of interest (usually "Current Day" or "Prior Day") and filter for a Tool of interest.
3. Add the resulting chart to a new home page tab (you could name it "Cascade Analysis").
4. Select the "Tool State Gantt" chart from the "Search FabTime" box and configure it to match your selections from the Process Time Lot List, then add that chart to the same Home Page.
5. Set the home page tab to two columns and click above each chart to select the "Left/Right" configuration. This will align the two charts so that you can see them one above the other.
    a. Because the y-axis labels may differ in width, you may find that the charts don't quite align. To tweak the alignment so that the time scales of the charts line up, change the "Columns" drop-down in the "Formatting" section of the home page to "Free Form." You can then re-size each chart individually and drag them to your desired location.
6. Look for:
    a. Overlap in pre-process and process times for subsequent lots for tools that allow multi-lot staging (good).
    b. Gaps in the loading of lots onto the tool while the tool was available and WIP was waiting (not good).
7. As needed, generate the WIP Lot List chart at a key time to see what WIP was in queue when the tool was sitting idle or came off a period of unavailable time.

In the example below, the Tool State Gantt chart shows a patch of dark gray (standby-wip-waiting) between lots 1949 and 2131 (see two-sided arrow below). We can also look at the WIP Lot List Chart, grouped by lot priority, at the time that the tool became available after a period of extended downtime. This shows us which lots were in queue and lets us confirm that the right lot was processed first (Lot 1998, the longest waiting of the higher priority lots). Finally, anywhere that we see the dark purple post-process time represents a potential opportunity for improvement, as this is time that lots have finished processing, but have not yet been moved out of the tool.

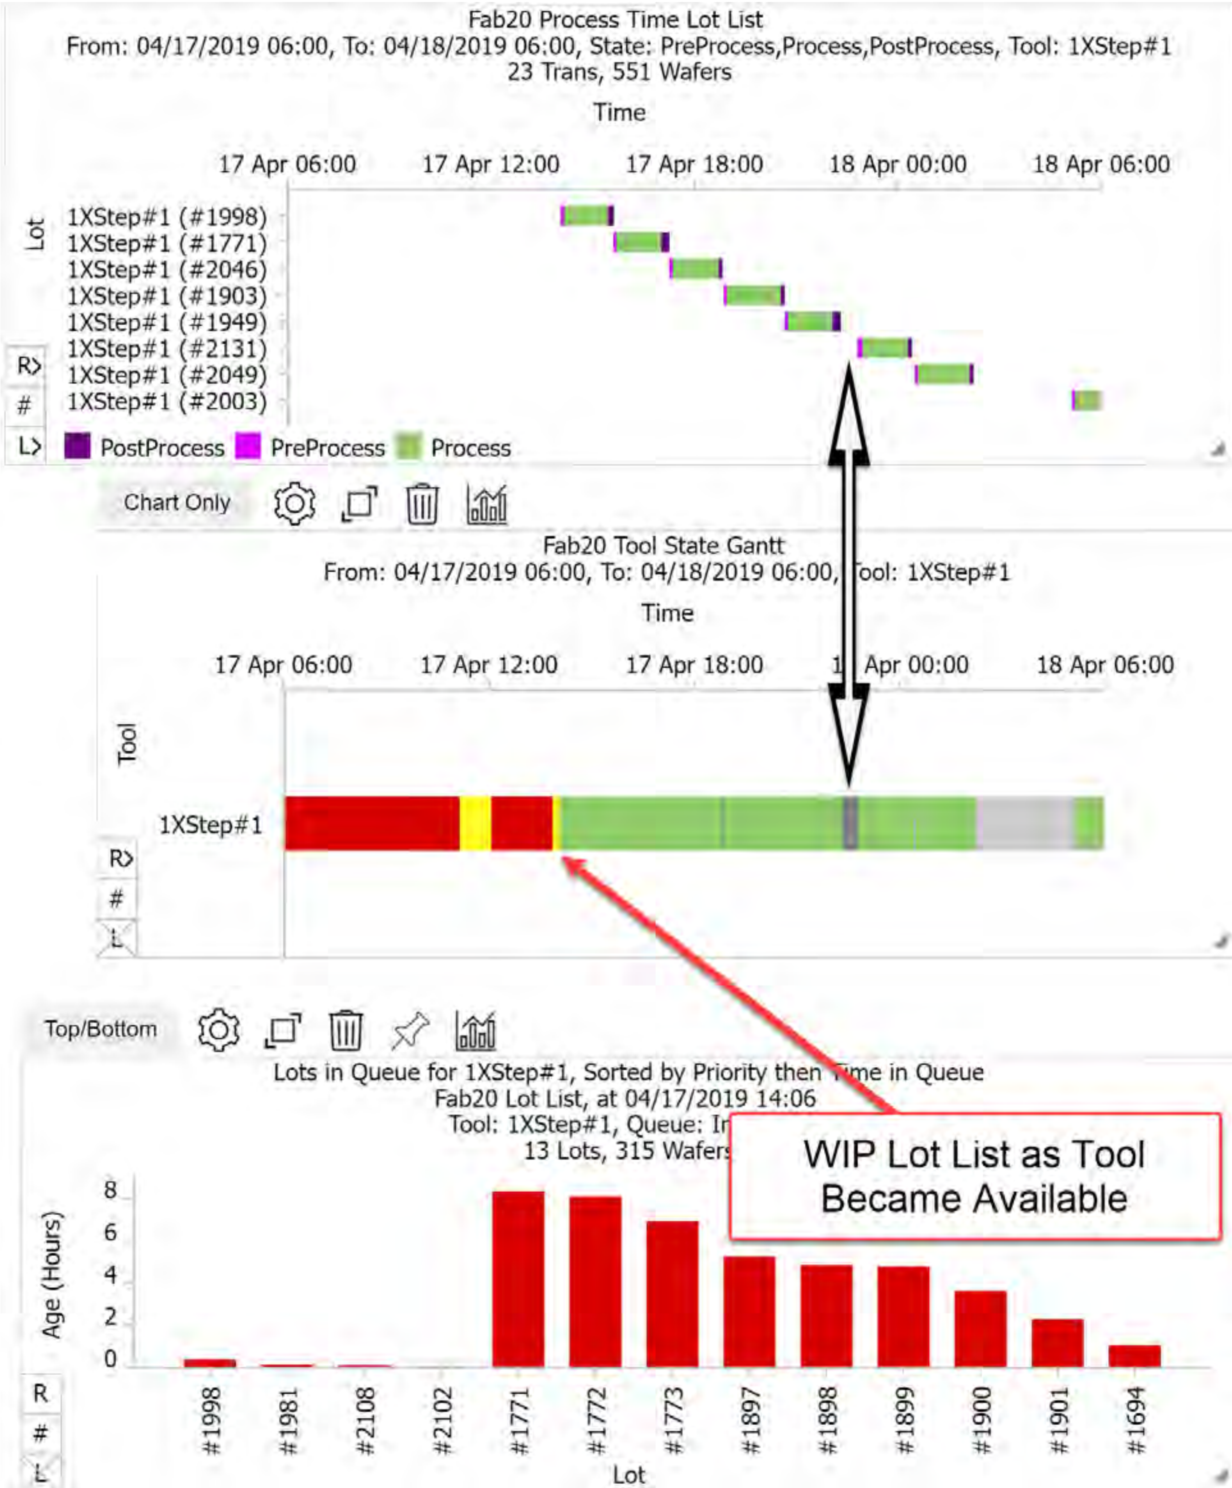

We hope you find this tip useful.

FabTime software customers can 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

Poll Results on Variability Contributors

Last month we shared a poll on LinkedIn asking people which of four fab complexities they thought contributed the most to variability in their fabs. We excluded downtime because we already knew that this was a major contributor, and asked about four other things: product mix, process time variation, holds, and batch processing. With 55 people voting, product mix edged out the second choice, holds. The poll is closed now, but if you have additional thoughts to share, please let us know.

Managing Cycle Time for External Processes

An anonymous subscriber wrote: "At my previous employer, I was responsible for overseeing the outside processing that was done for a process that we did not have the capability for. Additionally, at my current employer, which is not a semiconductor fab, I also oversee outside processes that are not done in house. I have noticed when it comes to cycle time it's hard to predict an accurate cycle time for a lot to complete as often the driving factor relies on the outside vendor and their turnaround time, which has a lot of variability. We have a base cycle time built into the route that is used to best predict the overall cycle time for a lot to complete. When the vendor misses the return date this adds to the cycle time and often ends with missed on time delivery. I was wondering if you have any previous articles covering this topic or any suggestions on how to improve the cycle time to make it become more stable.

FabTime Response: That is a challenging question. We have some recommendations for improving due date performance in Issue 23.03: Delivery Performance in Overloaded Wafer Fabs. We also discussed Understanding Why Lots Miss Their Due Dates in Issue 18.03. Both of those articles were focused on operations with the factory.

Your situation is different. It's difficult to improve something that you don't have control over. Obviously, you can reduce the chance of missing delivery targets by including a larger cycle time buffer for the external processing. You'd use something like the 95^{th} percentile of their historic cycle time performance, if you wanted to be very safe. Alternatively, you could provide incentives for the outside vendor to meet delivery targets, or provide delivery performance improvement recommendations to the outside vendor.

Another idea would be to gather data for the supplier cycle times and present it back to the supplier in the form of a boxplot, with individual data points plotted on the boxplot, as shown below. This would highlight the spread of the distribution. If you label the outliers with the highest cycle time, you could ask about individual lots. ("What happened with lot J47? Why did it take two weeks when most lots were much faster?") Alternatively, you could present the data as a histogram, but this would not be as useful as you wouldn't have the individual examples to point to and ask for accountability. The discussion is likely more productive when it is specific. "Why did lots X, Y, and Z take so long?"

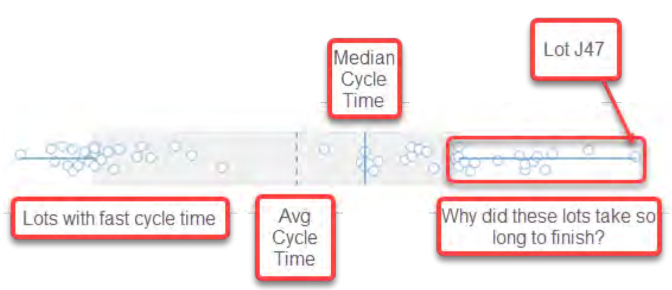

Those are the only ideas that we have, but we're sure that other subscribers have much more experience with this situation than we do. Subscribers, do you have any tips to share?

Response to Cycle Time Tip 007: If you use carts to transport lots between steps, consider smaller carts

A commenter on LinkedIn, **Frank Michael Schulze from X-Fab**, responded to a post about the October cycle time tip about minimizing variability from carts, writing: "Alternative: decouple the task of tool load/unloading and transport by using dedicated personnel for each. The transport job is done continuously in a "bus line" like way. By this, the maximum transport time of a finished lot only depends on the bus frequency and route length. While I am not the originator of this idea, my colleagues and I have found this to be an excellent approach to solve the variability issue in arrival times. It also provided enough motivation to look at the next step (hardware automation: even more constant/predictable, not dependent on staff availability - and well accepted by the people on the shopfloor as nobody wants to move many Kilometer per day each day)."

FabTime Response: This approach makes a lot of sense to us. It's something of a way to add the benefits of an automated material handling system, without the reliability risk (or cost!). Thank you for sharing!

Cycle Time Impact of Hot Lots

A prospective customer asked us this summer if we had a chart in FabTime that could show the impact of front-of-the-line hot lots on the cycle time of lots already in queue. We didn't have a standard chart for this, but we thought it was a useful insight to capture, so we went ahead and implemented a new Hot Lot Cycle Time Impact List chart. To calculate hot lot cycle time impact, FabTime does the following:

1. Generate a list of Arrivals within the chart date range, for hot lots with the specified PriorityClass filter.
    - Any lots with ToolGroup=Undefined are excluded from this list.
    - Lots not matching the specified PriorityClass filter are assumed to be non-hot lots.
2. For each hot lot Arrival, search for the corresponding BeginRun with timestamp on or after the Arrival at the same Operation/Step.
    - It is assumed that BeginRun transactions will always be preceded by an Arrival transaction.
3. For each hot lot BeginRun, find the time of the subsequent BeginRun for any lot at the same ToolGroup.
4. For each hot lot, find any non-hot lots that were in queue for the same ToolGroup at the time of the hot lot Arrival, and that have not had a BeginRun prior to the hot lot BeginRun.
5. Approximate the delay of non-hot lots as the delta between the time of the hot lot BeginRun and the time of the subsequent BeginRun (for any lot at the same ToolGroup).
    - Argument: The hot lot jumped ahead of the non-hot lots, thus delaying each of them. Each lot that arrived prior to the hot lot could possibly have started processing when the hot lot started processing, but was delayed at a minimum to the subsequent BeginRun time.

For example, assume in the chart below that Lot1 and Lot2 are non-hot lots, and Lot3 is a hot lot.

	8:00	9:00	10:00	11:00	12:00	13:00	14:00	Cycle Time
Lot 1	Queue	Process						3 hours
Lot 2		Queue				Process		6 hours
*Lot 3 **(HOT)***			Queue	Process				3 hours
								Total: 12 hours

- At 08:00, Lot1 arrives at Step1 (ToolGroup1).
- At 09:00, Lot1 begins running on Tool1. Lot2 arrives at Step1 (ToolGroup1).
- At 10:00, Lot3 (hot) arrives at Step1 (ToolGroup1).
- At 11:00, Lot1 completes running on Tool1. Lot3 (hot) begins running on Tool1.
- At 13:00, Lot3 (hot) completes running on Tool1. Lot2 begins running on Tool1.
- At 15:00, Lot2 completes running on Tool1.

In this example, Lot2 was delayed by 2 hours due to Lot3 (hot). This delay is computed as the delta between Lot3's BeginRun (11:00) and the time of the subsequent BeginRun (13:00).

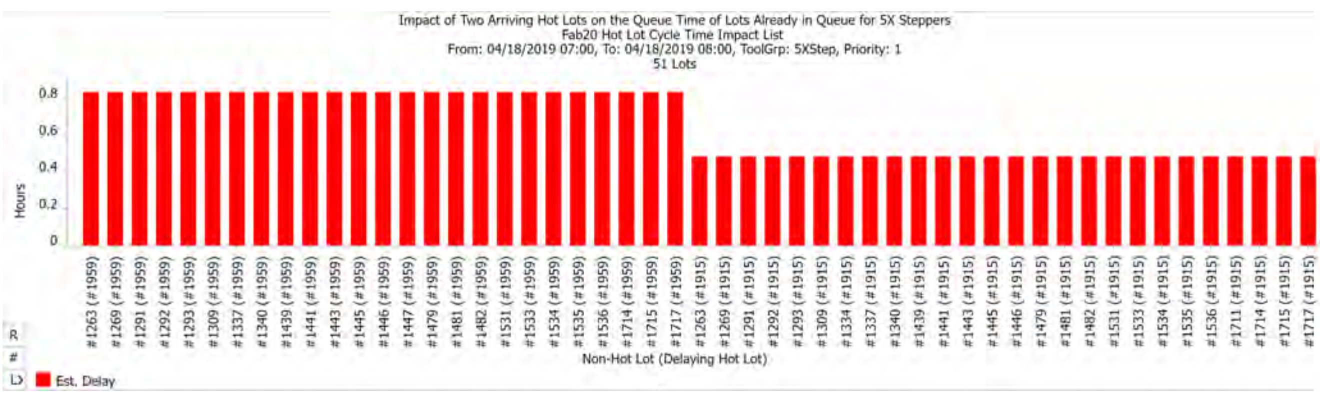

The above chart from FabTime's software shows another example. In this case, two hot lots arrive at the 5x Stepper tool group during the one-hour time period of the chart. Each of these lots jumps the queue and delays the roughly 25 lots already in queue by the hot lot process time (~0.8 hours for the first lot and ~0.5 hours for the second lot).

It's tempting to look at this example, with so many lots already in queue, and think that the hot lot delay should be averaged across the other lots, rather than incurred for each lot. But consider: if each of these two hot lots had not jumped the queue, every lot already in queue when they arrived would have been processed 0.8 + 0.5 hours sooner. The extra queue time applies to each lot, not to the lots on average.

The chart on the next page shows the average cycle time per visit through the 5x Stepper tool group, broken out by Priority 1 (hot) lots and Priority 2 (non-hot) lots. The Priority 1 lots incur only a tiny bit of queue time (the light red), primarily due to waiting for other hot lots, while the Priority 2 lots incur considerable queue time (much of it spent waiting for those Priority 1 lots).

Do other subscribers capture this data on the cycle time impact of hot lots? What do you do with it?

For more about the cycle time impact of front of the line hot lots on other lots, see Issue 19.03 of the newsletter, Cycle Time and Hot Lots: Updated, available for download by subscribers (along with all past newsletter issues) from the FabTime Newsletter Archive (current password: FabTimeCommunity).

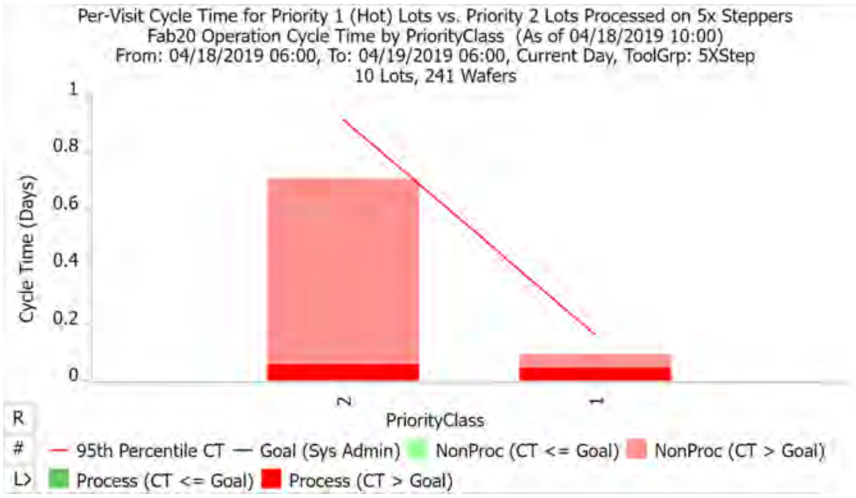

FabTime welcomes the opportunity to publish subscriber discussion questions and responses. Simply send your contributions to Jennifer.Robinson@FabTime.com.

Main Article: Cycle Time Improvement for 300mm Fabs

By Jennifer Robinson

Introduction

We've published many recommendations for wafer fab cycle time improvement in this newsletter over the past 24 (!) years. A subscriber recently asked which issues and topics are more relevant for 300mm fabs and which are more useful for 200mm and smaller fabs. The fabs FabTime works most closely with tend to fall more into the latter category than the former. We don't have a lot of first-hand experience with 300mm-specific challenges, so I reached out to my LinkedIn community to help me frame the discussion. Based on that feedback, I outline some characteristics of 300 mm (and highly automated 200mm) fabs, and then share some specific recommendations. What I would like is for this relatively brief article to be a catalyst for further discussion within the newsletter community on this topic.

Community Input to Frame the Discussion

I asked my LinkedIn community the following:

> "Looking for Input on 300mm vs. 200mm Manufacturing: What would you say are the differences between the types of fabs that would impact operational improvement strategies? Generally speaking, I'd expect more automation, more advanced scheduling/dispatching, and fewer one-of-a-kind tools in 300mm fabs. What else am I missing?"

Marc Engel, CEO at Agileo Automation, wrote: "automation between 200mm and 300mm wafer fabs is no longer necessarily a differentiating criterion, since the so-called GEM300 standards, designed for 300mm wafer fabs, are now being used in some new 200mm wafer fabs that are state-of-the-art in terms of automation. Indeed, more and more new 200mm wafer fabs are using OHT/AMHS or AGV-type carrier delivery systems. 200mm and 300mm fabs also use standard EFEM with atmospheric and vacuum robots to load/unload wafers within process modules. This distinction between 200 mm and 300 mm remains valid for plants built several years ago from an automation standpoint. Indeed, these 200mm fabs were using SECSGEM standard only without any standardization of job definition and carrier delivery, which was a clear disadvantage for cycle time optimization throughout the fab. So, I'd say that the MES generation and carrier delivery automation used are more important in differentiating cycle time than the wafer size!

Of course, 300mm is clearly an advantage for the quantity of chips produced per wafer, and therefore has an impact on overall cycle time."

Andy Beers, President at ABM Consulting, added: "All true except; the complexity of the 300mm Factory Automation design and reliability means…. If the overhead track gets backed up even for a few minutes it can take away hours of productivity to clear the stack-up. Defectivity requirements also mean more critical wet cleans and recovery requirements (chamber seasoning)."

FabTime Response: Marc's main point is that it's not so much a matter of 300mm vs. 200mm, but a matter of more automated fabs, using newer MES systems, vs. less automated fabs. Andy highlighted in particular the additional risk carried by fabs with automated material handling of the overhead system tying up, well, everything. We'll wrap each of these points into the discussion below, with thanks to Andy and Marc for taking the time to respond to my question.

Characteristics of Larger, More Automated Fabs (300mm or newer 200mm)

Distinguishing characteristics of 300mm (or newer, highly automated 200mm) fabs may include:

- Wafer level tracking (vs. only lot level tracking) of move transactions.
- More automated logging of transactions (e.g., the tool sends an end run transaction to the MES when the lot is done processing, vs. less automated fabs where there is no transaction recorded until an operator is there to load the tool).
- Automated material handling, via overhead transit system or via AGV-type carrier systems (robots!).
- More advanced scheduling / dispatching, necessary to run the most automated systems, in contrast to more manual fabs where operators have some leeway in selecting the next lot to process.
- Fewer one-of-a-kind tools (simply as a matter of scale, because the fabs are larger).
- More chips per wafer for 300mm wafers (meaning that each wafer is more valuable, such that scrap is more expensive).
- Greater cost of the fab overall, meaning more pressure to maintain high utilization of the tools.
- More stringent processing requirements in some cases (wet cleans, etc.). Possibly more time links between process steps(?).

Of course, there is a spectrum in terms of level of automation for fabs, independent of wafer size. An older fab might undergo a project to add robots for material transfer, or upgrade to a more advanced MES. But this list is sufficient to give us some ideas about specific cycle time improvement recommendations that are relevant for the larger, more automated fabs, many of which are 300mm fabs.

Recommendations for Cycle Time Improvement in 300 Fabs

Here are a few recommendations that are applicable for 300mm and highly automated 200mm fabs.

1. Minimize process restrictions to ensure that each operation can be run on at least 2 qualified tools (ideally 3-4). Just because a fab has a high level of tool redundancy does not necessarily mean that this fab has at least two qualified tools for every recipe, particularly for fabs doing foundry work. But it's not helpful to have plenty of redundant tools if individual recipes are not cross-qualified. See Issue 20.05 for details.
2. Subject to not violating 1, consider some creation of sub-groups within the tool groups to reduce process time variability (different process times from lot to lot). It's also important when doing this to keep an eye on the relative utilization of the sub-groups in the presence of product mix changes. See Issue 22.04 for more on the impact of utilization and process time variability.
3. Look for and reduce long periods of unavailable time on tools (whether scheduled or unscheduled), especially in toolsets with less redundancy. That is, go beyond availability to look at variability of

availability and green-to-green time. See Issues 22.01 and 20.02 for more on the impact of long downtimes, and green-to-green time, and Issue 18.04 on measuring availability variability in general.

4. Use dispatch rules or scheduling algorithms that look downstream to avoid starving bottleneck tools and allowing too much WIP to pile up in one place (line balance). See Issue 15.01 on dispatching and line balance.
5. Analyze the efficiency of the AMHS to ensure that lots aren't being grouped unnecessarily for transit or spending excessive time in transit.
6. Put heavy focus on the reliability of the overhead transit system. It is something of a mega one-of-a-kind tool in the fab and can delay every lot if it fails.
7. Take advantage of more automated logging of transactions to report time that WIP spends in states like travel and post-processing. The former may indicate a need for optimization or expansion of the material handling system. The latter can be used to identify places where additional operators are needed (post-process time is time from when the tool finished processing until an operator logged a move out – it can be an indicator of insufficient operators).

Conclusions

The fundamental things that drive up cycle time in a discrete manufacturing environment are common to all factories. At the tool level, they boil down to utilization, variability, and number of qualified tools. The many complexities of wafer fabs (see Issue 24.04 for a recap) ultimately impact cycle time because they impact one or more of those three variables. Because we take a factory physics approach that focuses on these fundamentals, we've never differentiated much in this newsletter between the different types of fabs in our recommendations.

Our experience base is more with people working in 200mm and smaller fabs, so our recommendations have probably slanted somewhat in that direction (like last month's cycle time tip: if you use carts to transport lots between steps, consider smaller carts, or Issue 23.05 about managing true one-of-a-kind tools). In this issue, we have focused more specifically on 300mm fabs (or, as pointed out by one of our colleagues on LinkedIn, highly automated 200mm fabs). After thinking about characteristics of those fabs, we've shared some specific recommendations. We would love to have feedback from other subscribers on what we're missing.

Closing Questions for Newsletter Subscribers

What are we missing here in terms of distinguishing characteristics of 300mm and other highly automated wafer fabs? What are other recommendations for cycle time improvement that are more relevant for these automated fabs?

Further Reading

The following FabTime newsletter issues were mentioned above. All are available from the FabTime Newsletter Archive (current password: FabTimeCommunity).

- Dispatching and Line Balance (Issue 15.01)
- Measuring Variability of Availability (Issue 18.04)
- A Metric for Green-to-Green (G2G) Analysis (Issue 20.02)
- The Impact of Tool Qualification on Cycle Time (Issue 20.05)
- On Breaking Up PMs and Other Unavailable Periods (Issue 22.01)
- Fundamental Drivers of Wafer Fab Cycle Time (Issue 22.04)
- Managing One-of-a-Kind Tools (Issue 23.05)
- Commonly Reported Wafer Fab Cycle Time Contributors (Issue 24.04)

Subscriber List

Total number of subscribers: 2,857

Top 20 subscribing companies:

- Onsemi (146)
- Intel (136)
- Infineon (124)
- Micron Technology (116)
- Analog Devices (111)
- Microchip Technology (96)
- GlobalFoundries (86)
- NXP (78)
- STMicroelectronics (70)
- Skyworks Solutions (67)
- Texas Instruments (64)
- X-FAB (54)
- Western Digital (53)
- Seagate Technology (52)
- Wolfspeed (52)
- Carsem M Sdn Bhd (41)
- Qualcomm (36)
- Tower Semiconductor (34)
- Applied Materials (34)
- ASML (33)

Top 3 subscribing universities:

- Ecole des Mines de Saint-Etienne (EMSE) (7)
- Arizona State University (5)
- Ben Gurion University of the Negev (5)

New companies and universities this month:

- TU Dresden (1)

Note: Inclusion in the subscriber profile for this newsletter indicates an interest, on the part of individual subscribers, in cycle time management. It does not imply any endorsement of FabTime or its products by any individual or his or her company.

There is no charge to subscribe 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, please visit our website. To unsubscribe, reply to your newsletter email 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 factory insight engine, please visit our website. A sample home page and a sample page from FabTime's new Charts menu are shown below.