FabTime Newsletter

Volume 23, No. 1 February 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 Beeg (Fabmatics)

Date: Tuesday, February 15, 2022 – Vol. 23, No. 1

Keywords: Queueing Models; Operating Curves; Holds; Operators

Table of Contents

Welcome

Welcome to Volume 23, Number 1 of the FabTime Cycle Time Management Newsletter. We hope that the new year is treating you all well. In this issue we have a question assessing potential interest in a multi-company session of our cycle time management course, an announcement about the FOA Collaborative Forum (in person as of this writing), and some highlights from Jennifer's LinkedIn. We have a FabTime software tip about the use of our new search bar. We have subscriber discussion about various topics, including a recommendation for a new blog on Factory Physics that we believe subscribers to this newsletter will enjoy.

We've talked many times in this newsletter about using queueing model-based operating curves to illustrate the impact of various factors on wafer fab cycle time. In our main article this month, we discuss the idea of populating these operating curves with data from real fabs. Reasons to do this include estimating the impact on cycle time from changes in utilization or other parameters and identifying places in the fab where the cycle time is worse than might be expected given the known characteristics of a tool group. We discuss techniques and pitfalls of collecting data for this effort. We show a detailed example from our demonstration server and then highlight possible uses of the operating curves. As always, we welcome your feedback.

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

Community News/Announcements

Multi-Company Session of Cycle Time Management Course

We have held many sessions of our remote half-day cycle time management course for companies over the past few months. We are currently assessing whether there is sufficient demand to justify holding a multi-company session of the course that people could register for individually. If this is something you would be interested in, please contact Jennifer.Robinson@FabTime.com.

The cycle time management course is a four-hour class delivered via Teams in two two-hour sessions, usually during the same week. The first portion covers the fundamental drivers of cycle time (utilization, variability, and number of qualified tools), while the second portion covers metrics, equipment downtime, and the impact of various tactical/operational factors (staffing, etc.) on cycle time. Request more information about the course from our website.

FOA Collaborative Forum

After two years of virtual meetings, the SEMI Fab Owners Alliance (FOA) is schedule to be back in person with the 10th Annual Collaborative Forum next month. The meeting will be held March 2-3 in Scottsdale, Arizona. The forum features joint success presentations by Solution Providers and Device Makers along with talks from industry analysts and invited presenters. Jennifer will be representing FabTime in person at this year's forum and hopes to see some of you there.

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 posts have included:

- Micron Technology is considering building a $40B wafer fab in Austin, TX. The Austonia article says that CA, NC, and AZ are being considered. I have to say that I will be flat out astonished if they choose California. Thank you to Robert Quinn for the link. There is a robust discussion of this news, including speculation on where this fab will be built, in the comments of the LinkedIn post.
- See also this LI post about Intel's plans to build two new wafer fabs near Columbus, Ohio. According to the WSJ story, "The company also pledged $100 million toward partnerships with educational institutions to build a pipeline of talent and bolster research programs in the region."
- What do people think? Should the post-pandemic office become more of a clubhouse where people go to network with colleagues, while the solo work takes place at home? This WSJ article proposes that it should. [Link to post.]
- This WSJ article provides further evidence that our recent newsletter on running a fab during a staffing shortage was timely. "In the U.S. alone, around 70,000 to 90,000 workers or more will have to be added by 2025 from 2020 levels to meet the most critical workforce needs for anticipated fab expansion, according to a report by Eightfold.ai, a talent-management company." [Link to post.]
- The WSJ also carried a piece about increased capital equipment spending by wafer fabs at the "trailing edge" (Analog Devices, onsemi, NXP, Microchip, etc.). The article includes this warning: "The risk, of course, is if demand cools before new capacity is fully utilized. Analysts for UBS predicted in a report earlier this month that trailing-edge chip makers could be in a state of "foundry oversupply" next year, hurting gross margins." Time will tell! [Link to post.]

For more industry news, connect with Jennifer on LinkedIn: http://www.linkedin.com/in/jenniferrobinsonfabtime

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 FabTime's New Search Bar to Find Information

We've been gradually rolling out FabTime's new user interface (UI) to customer sites. For those of you who are using the new version (Patch 114 or higher), we wanted to share some tips about using the new search bar. The search bar is available from anywhere in FabTime and will autofill with suggestions as you start typing. As always, you can search for a type of chart by typing in the chart name. But you can also:

- Type in a lot number to go directly to the Lot History chart for that lot.
- Type in a tool name to go directly to either the Moves Lot List or the WIP Lot List for that tool.
- Type in a person's name to select from a list of that person's shared home page tabs.
- Type in a home page tab name to select from a list of your own and other shared tabs of that name.
- Type in a topic to see a list of charts, FabTime newsletters, tips, and help articles that have that term in the title. For example, when we search FabTime's demo server for information about bottlenecks we see a home page tab that Jennifer has created, six past newsletters and three past tips.

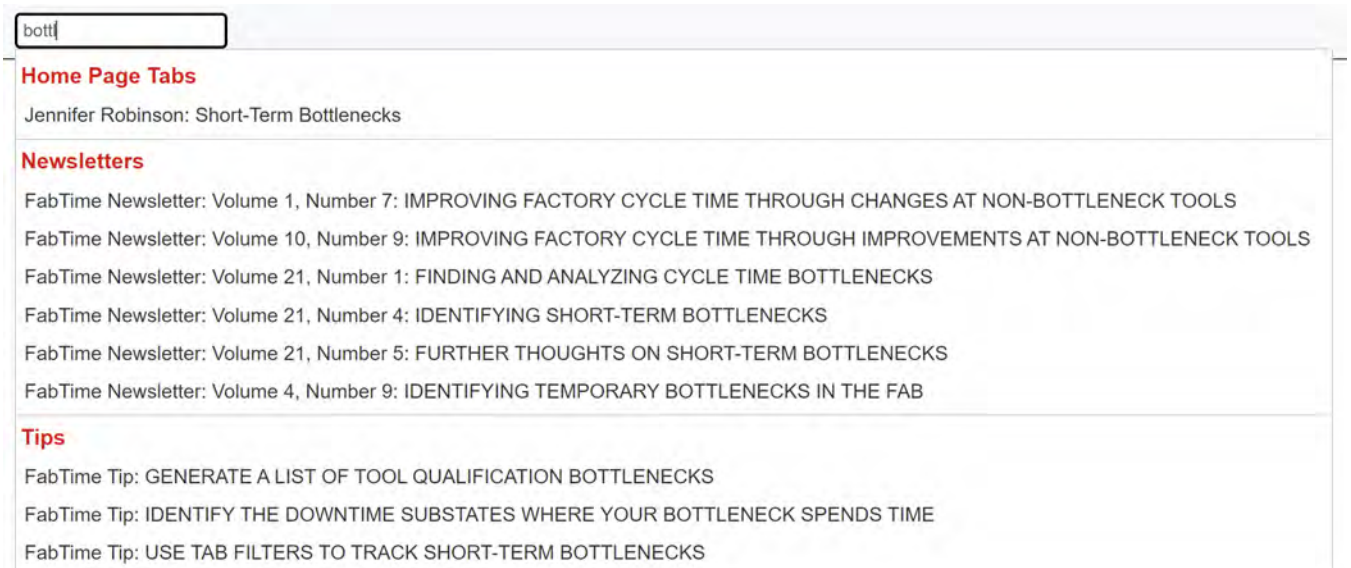

We hope you find this tip useful. Subscribe to the separate Tip of the Month email list (with additional images for customers only) here: http://www.fabtime.com/tip-of-the-month.php. Thanks!

Subscriber Discussion Forum

Setting Limits for Holds

An **anonymous subscriber** wrote: "I am working to set limits for Hold lots in my production facility. I am wondering if you have any guide/newsletter related to HOLD limits that I can refer to as a reference?"

FabTime Response: We last discussed this in Issue 6.06 (available to subscribers for download from the FabTime newsletter archive). We also discuss recommendations for managing holds in our cycle time management course. However, neither of these references talks about hold limits. Our general feeling is that you should always work to have fewer holds, because holds add to cycle time for the lots on hold and add to variability that affects all lots. However, we've always discussed this in terms of how you should manage the holds that you have, rather than thinking about a specific limit on the number of holds. Do any other subscribers have feedback on setting limits for the number of holds in a fab?

Open Questions from Our Cycle Time Classes

Several interesting questions have arisen during our cycle time management classes. We are sharing them here, in the hope that some will spark subscriber discussion.

- In a high mix fab that has lots with different numbers of wafers, how do people keep operators from prioritizing larger lots at the expense of smaller ones (if operators are motivated by moves)?
- How do people manage the variability from metrology sampling? Sampling adds variability, but you don't want to have to add testing for everything.
- Are any other fabs dedicating weekends to PMs or analyzing failure events for time-of-day patterns?
- Has anyone made other findings from the use of Dynamic X-Factor in a fab? We know that DXF can help with a forward look at cycle time and can help identify issues with not keeping tools busy during shift change. Have there been other findings as people use DXF in practice?

New Blog About Factory Physics from a Long-Time Subscriber

A long-time subscriber to this newsletter, **Thomas Beeg of Fabmatics**, has started a new blog about Factory Physics. He has an interesting early four-post series asking how in practice fabs should determine which tool groups are the bottlenecks. He runs through various candidates and includes a reader poll on fab metrics, as well as an interesting discussion about fab utilization profiles. We highly recommend that readers of this newsletter give this series a look and subscribe to receive Thomas' future posts by email. You can download the full bottleneck series as a PDF file here. Alternatively, go to the landing page for the blog and scroll down to find the oldest post at the bottom of the page, and then scroll upward to find parts 2-4 on the blog. A newer ongoing series on the blog is about the chip shortage and fab performance and looks at the dollar costs and benefits of cycle time reduction based on operating curves.

Issue 22.05: Operators

Labor Shortages (including operators): In the comments on a recent LinkedIn post by Jennifer about where Micron may be building a new fab, several people expressed concerns about the labor shortage in the industry. These are medium-to-long-term concerns. If all the companies discussing building fabs in the US go through with their plans, there is a high likelihood that these fabs will struggle to find qualified employees. This is something that SEMI is working on, but efforts by individual companies (apprenticeship programs, etc.) will be necessary, too.

An **anonymous subscriber** wrote in response to the last issue to share the actions ongoing at his company to combat operator shortages.

> "Thanks for sharing. I appreciate the recommendations. Some sharing from our company: Manpower shortages are the exact problem I am having at my site right now as well. Multiple actions we have in place/ongoing are:
>
> 1. Getting operators cross trained/certified for more operations
> 2. Getting other employees (not currently operators) certified to perform certain operations
> 3. Increasing the operator population….
> 4. Focusing on critical/bottleneck operations and letting non-bottleneck tools spend time idle. However, this really impacts overall factory cycle time and efficiency. I have also noticed that operators are getting burned out as we push for maximum capacity usage and moves.

My team is continuously working to improve the situation here…once again thanks for picking up the right issues and proposing some good recommendations."

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

Using Operating Curves in Real Fabs

Introduction

During a chip shortage like the one we're in, many fabs are running at a higher utilization rate than usual, and consequently have WIP piling up at various tool groups. We talked in Issue 21.01 about how to identify cycle time bottlenecks. In the Subscriber Discussion Forum above we linked to a series by Thomas Beeg that looks at a similar question. But here's another question that people have been asking Jennifer of late:

> How do you tell where in the fab the cycle time x-factor is worse than you would expect, based on current utilization and variability levels?

This is the kind of question that one can answer with a full-fab simulation model. However, keeping such a model up to date in the ever-changing environment of a fab has proven unsustainable for many fabs.

This is a question that can be answered, to at least some degree of accuracy, using queueing models. FabTime's Operating Curve Spreadsheet (available for download by subscribers from our newsletter archive) uses a queueing model to generate the operating curve for a tool group. You can compare up to three input scenarios. We use this spreadsheet extensively in our cycle time management course to help students build intuition about how things like arrival and process time variability, number of qualified tools, and downtime distribution affect cycle time through a tool group.

It is possible to also use the spreadsheet to build operating curves for actual tool groups from your fab. In this article we discuss why this might be a worthwhile effort, how to go about collecting the right input data, and what to do with the operating curves once you have them.

Reasons to Build Operating Curves for Real Tools

Here are three reasons to consider building operating curves for actual tool groups:

1. To estimate the impact on cycle time from utilization changes due to changes in product mix.
2. To assess the potential improvement in cycle time if we can make other changes like increasing the level of tool qualification, buying another tool, or decreasing variability.
3. To identify tool groups where the cycle time is significantly higher than the queueing models predict. These tools are potentially low-hanging fruit. The cycle time isn't only high because of fundamentals like utilization, variability, or number of tools (if it's higher than the operating curve predicts). It's high because of something else, something in how the tool is being operated. Maybe we have standby-WIP-waiting time or soft dedication that we're not capturing, for example.

It is difficult to build operating curves using historical data, because we typically don't have enough observations of cycle time at different levels of utilization. Even if we do have this data, it's likely that other things in the fab have changed in between observations.

Again, one could build a detailed, full-fab simulation model to address these questions. But those models require considerable maintenance, as well as a specialized skillset to run them and interpret the results. For fabs that don't already have such models in place, creating one is a daunting task. By building queueing model-based operating curves for individual tool groups, we can get results more quickly (we don't have to design long-running experiments), and we can focus first on the tool groups that are of the most interest. What tools are those?

- The top capacity bottlenecks (the tool groups with the highest utilization)
- The top cycle time bottlenecks (see Issue 21.01)
- One-of-a-kind tools (or tools that operate like one-of-a-kind tools because of process restrictions)
- Tools where we expect significant utilization changes due to changes in product mix

Of course, there are shortcomings to the use of queuing models of the individual tool groups:

- Queueing models are only approximations. They won't capture the full level of detail of the fab, particularly the impact of operators. They don't readily capture complexities like variations in lot size due to rework and scrap, cluster tools and other linked tools, and time constraints between process steps. Even batch processing is not modeled in our Operating Curve Spreadsheet.
- By looking at operating curves for individual tool groups, we fail to directly capture the interaction between tool groups (though we can approximate that by measuring arrival variability).

But we're not necessarily looking for a detailed model that tells us exactly what the cycle time is going to be through the fab as a whole tomorrow, or next week. We're looking for something reasonably quick and easy to implement that can identify low-hanging fruit for improvement efforts and let us do a bit of what-if analysis on which changes might be fruitful. The trick is to find some of these, and then go off and work on the improvement efforts in the real fab.

Collecting Data to Populate Queueing Models

The primary inputs of FabTime's Operating Curve Spreadsheet for a tool group are:

- Average process time
- Number of tools
- Coefficient of variation (CV) of time between arrivals
- CV of process times
- Percentage downtime
- CV of repair time
- Mean time between failures

(There are also options for modeling batch arrivals and the impact of hot lots on regular lot cycle time, but we'll leave those to be discussed in future articles. The spreadsheet does not include batch processing.)

Average process time for all lots going through a tool group is reasonably straightforward to calculate from recent fab data. If you have begin-run and end-run (or start process and move out) transactions, you should be able to record the set of process times for all lots through the tool group for the past two weeks or month and take an average of that. One complication here arises in the case of cascading tools, where a second lot can start before the first one finishes. Fortunately, the queueing model results are not particularly sensitive to this variable. Our advice is to get whatever actual data you can get with minimal effort, and not worry too much about the data for this variable being perfect.

Number of qualified tools in the tool group sounds straightforward but can be tricky in practice. This is because, unless we are looking at a one-of-a-kind tool, we typically have recipes that can only be performed on some sub-set of the tools in a tool group. Sometimes fabs break large tool groups into smaller sub-groups and dedicate certain recipes to each. If your dedication strategy results in stable sub-groups, you may want to model the operating curves of the sub-groups separately.

The right number to use here for number of tools is not the total number of tools in the full group, but rather, a representative number that captures the number of tools available for the average lot that passes through that tool group (or sub-group). We realize that this number might vary considerably for different recipes. But that's the beauty of the Operating Curve Spreadsheet, which allows you to enter data for up to three different scenarios. Enter three different values for number of qualified tools and see what the impact is for operations that have more vs. fewer qualified tools. A side benefit may be understanding how much cycle time improvement is available from qualifying additional tools.

CV of times between arrivals to a tool group is also straightforward, provided you have some sort of arrival transaction recorded in your MES. In the simplest case, you can use the move out of the prior operation and treat that as the arrival transaction (though this will underestimate variability due to lot transport). Each time an arrival transaction is recorded you start the clock for the interarrival time until the next lot. Once you have a sequence of interarrival times, you can calculate the CV of the series (standard deviation divided by average – both are standard functions in Excel). For batch arrivals, your sequence of interarrival times may look like: 1 hour (until the batch arrives and the first lot is scanned), then 0 minutes, 0 minutes, 0 minutes as the other lots in the batch are scanned. It's important to include the 0 (or very low) values. Batch arrivals are a significant source of arrival variability in wafer fabs. An example is shown above. (See Issues 4.05, 7.08, and 16.05 for more details on arrival variability.)

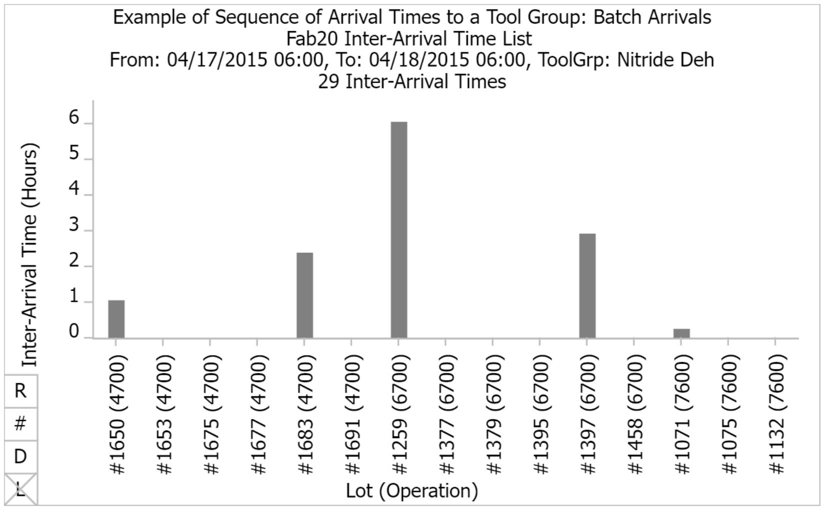

A caveat is in order regarding arrival CVs for bottleneck tools in real fabs. The arrival CV measures the variability in arrivals to the back of the queue. For tools that are sometimes starved, this is a valid measure. However, if you have a tool group that nearly always has a queue in front of it, the effect of the arrival variability on the operating curve is dampened. If you validate your queueing model against actual fab data for a bottleneck tool and the operating curve spreadsheet predictions are significantly higher than the data observed in the fab, it's probably worth checking to see if the tool group always had a queue in front of it during the time of your data collection. If so, there is an argument for setting the arrival CV to something closer to zero. In this case (for a very heavily loaded bottleneck), if there is always someone there to load the tool, then the process time variability and downtime variability will dominate.

A second caveat concerns tools where the WIP is held in queue at a prior step instead of the current step (e.g., due to time constraints). In this case, the arrivals are not independent (something required by the queueing models), but rather correlated between the two steps. The queueing models may not be sufficiently valid to use in this situation.

CV of process times can be calculated from the same dataset that was used for calculating average process time. Take that set of data and compute standard deviation divided by average. The problem here is that this is a lower bound for the process time variability. What we should really use here to be most accurate is the CV of the effective process time, where effective process time is the time from when a lot arrives at the front of the queue, ready to be processed, until an end-run transaction is logged. This could include setup time and downtime. From the lot's perspective, this time is effectively part of the process time.

Effective process times are a useful concept (see Issue 12.06 and the subscriber discussion in 13.01, plus the references listed below), but they are very difficult to compute in practice. This is because even if we start the clock when a lot reaches the front of the queue for a tool, things can change. While the tool is, say, down for a PM, another lot might arrive that has a higher priority, jumping in front of the first lot. Or the first lot might end up processed on a different tool altogether.

The conclusion that we have reached in terms of effective process times is that it's useful in theory to think of downtime and setup time as part of process time variability. But in practice, calculating the CV of effective process times isn't practical. Our recommendation is to use the CV of the sequence of actual process times and understand that this is a lower bound on the process time variability.

Downtime distribution is the trickiest thing to model in the Operating Curve Spreadsheet from actual fab data. The reason for this is that the spreadsheet only allows you to enter data for one downtime distribution. In practice, the distributions of scheduled and unscheduled downtime are probably quite different. We may also want to capture the impact of engineering time, as this is time that the tools are not available to manufacturing. If we think about what we're trying to capture with this downtime distribution (the impact of the variability from the tool being unavailable), a couple of possible solutions comes to mind.

One approach is to generate a list of all scheduled downtime, unscheduled downtime, and engineering time transactions for the tool group over a representative time period (two weeks or a month). Treat the average of these transaction times as a generalized mean time to repair (MTTR). Also calculate the CV of the transaction times to use as the CV of the repair time. Sum up the percentage of time spent in scheduled downtime, unscheduled downtime, and engineering time over the same period and use that as the percent downtime. Use the percent downtime and the MTTR to compute the MTBF to input into the spreadsheet.

We know that MTTR/MTBF = % Down. Therefore, MTBF = MTTR / % Down.

For example, if the MTTR is 5 and the % Down is 20, then the MTBF = 5 / (.2) = 25.

Another approach is to use the Green-to-Green instance data for the tool group instead of the scheduled downtime and unscheduled downtime transactions (see Issue 20.02), combining those with the engineering time transactions. Green-to-Green (G2G) instances measure the total time the tool is down between two up states (whether productive or standby). A G2G instance may contain multiple scheduled and/or unscheduled downtime transactions in a single instance. As such, the G2G transactions will have a higher CV than the individual downtime transactions. It is this total time that the tool is unavailable to manufacturing that has the most impact on cycle time, so this is probably the most accurate aggregated downtime measurement to use, if you have it available.

It should be noted that neither of these approaches is perfect, since we are combining non-like distributions (especially when we include the engineering time), but we do think it will give some sense of the downtime variability to include.

Actual Average Tool Utilization and **Actual Average Cycle Time X-Factor per Visit** are the final pieces of data to collect for validation of this effort. You should measure these over the same time period that you used for collecting the other data, especially the percent downtime. If using the Operating Curve Spreadsheet, you can validate your inputs by seeing how close the actual cycle time x-factor is to the prediction of the spreadsheet at the actual utilization value.

One potential complication here in measuring actual cycle time per visit x-factors is that, as mentioned above, for some tools the WIP (and hence queue time) is held at an earlier step (e.g. a clean step, or the step before some other time link). The queueing models will not be as applicable in this situation, though you could try combining the queue time across the two steps.

Note for FabTime Software Users: If you use FabTime's dashboard software, the Operating Curve chart can be used instead of the Operating Curve Spreadsheet. However, the full multi-tool approximation with downtime is not currently modeled in FabTime's software. We are working to expand the calculation used in the software and will update when that is available. The current version of the chart uses the simpler approximation for a one-of-a-kind tool with a variability factor. You can, however, use the data table for the Operating Curve chart to obtain the average utilization and the CV of arrival and process times for a tool group (shown in the data table at the "Operating Point."), and then use the spreadsheet version for a more detailed analysis. This is what we've done in the examples below.

Examples from Our Demo Server Data

To test out this methodology, we looked at three examples from our software demonstration server. This server is populated with data that was simulated using a dataset from an actual factory that was scrubbed and made publicly available many years ago. We looked at a 5x Stepper tool group with six tools, a Nitride Deh tool group with two tools, and a Wet Etch tool group with two tools. We recorded the following values over a two-week period.

Description	5x Stepper	Nitride Deh	Wet Etch
Average process time (hours)	1.476	0.95	0.92
Hot lot percentage	0%	0%	0%
Number of tools	6	2	2
Process time variability	0.17	0.1	0.22
Inter-batch arrival variability	2.94	3.13	2.46
Average arriving batch size (lots)	1	1	1
Arriving batch size variability	0	0	0
Repair time variability	1.377	2.09	2.36
Mean time between failures (hours)	12.54	6.95	6.49
Percentage downtime	40%	31%	30%

The resulting Operating Curves generated from our Operating Curve Spreadsheet are shown below:

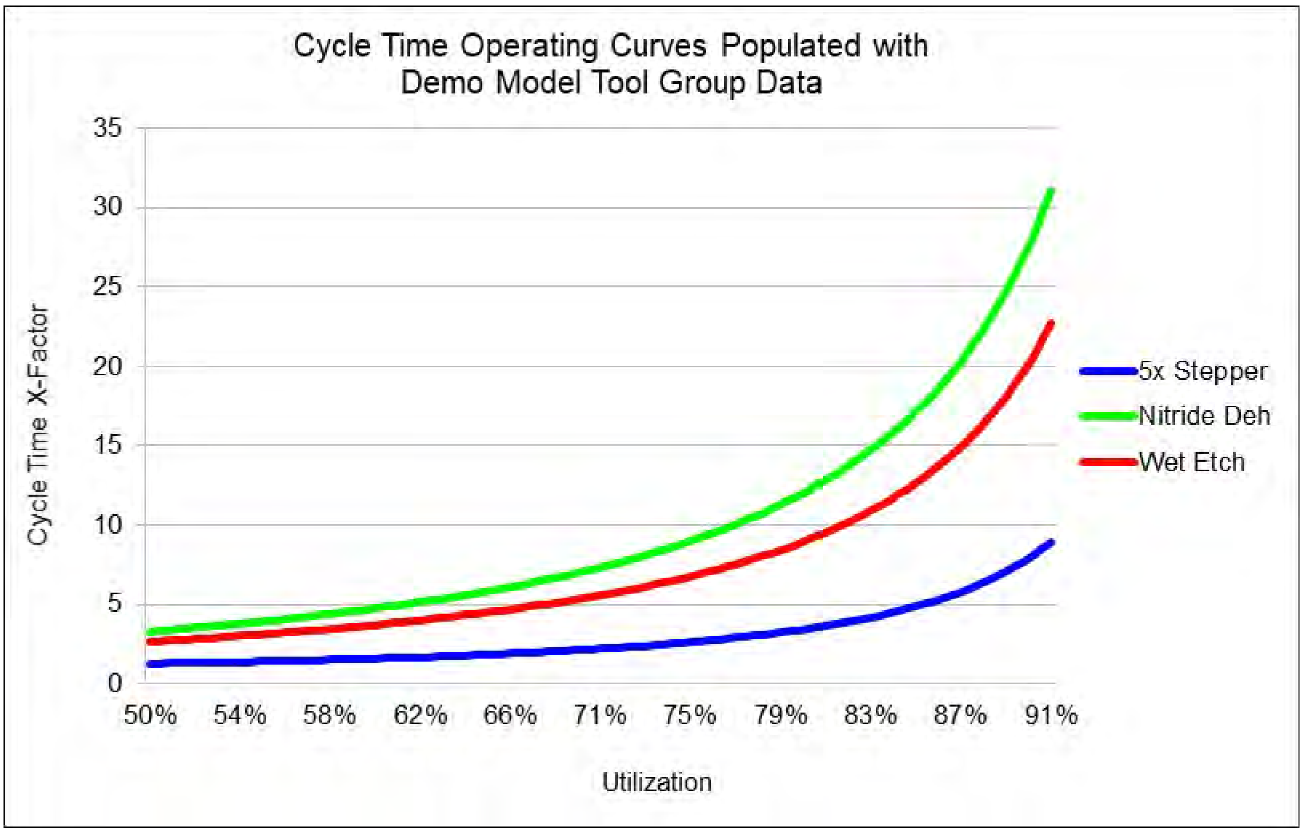

When we compared the prediction from the Operating Curve Spreadsheet to the values observed in our (simulated) demo, the predicted cycle time x-factor for the 5x Stepper group at 86% utilization was only about half of the (simulated) per visit x-factor. For the Nitride Deh group, the predicted x-factor from the Operating Curve Spreadsheet was approximately 10X, compared with 13.9X for the simulated value. The Wet Etch predicted value at 82% utilization was 10X, compared with 12X for the simulated value.

One loss factor not included in the analysis so far is capacity loss due to lack of operators. The 5X Stepper tool group spent 2.78% of time in a state of Standby-WIP-Waiting. This is time when the tool was available and lots were waiting, but there was no operator to load the tool. If we compute an adjusted utilization value in which we remove that 2.78% from the standby time, we get an adjusted utilization value of about 91%.

This doesn't sound like much of a change, but because the operating curve is nonlinear, the predicted x-factor at 91% is about 8.9X, much closer to the simulated value of 11.3X.

Adjusting the utilization values for the other two tool groups by treating the Standby-WIP-Waiting time as a loss factor, we see the following:

	Avg. Utilization	X-Factor from Demo	X-Factor from OpCurve	% Time Standby-WIP-Waiting	Adjusted Utilization	Adjusted OpCurve X-Factor
5X Stepper	86.4	11.3	5.2	2.78	91	8.9
Nitride Deh	76.2	13.9	9.8	3.4	80	12.2
Wet Etch	82.1	12	10.6	4.64	88	16.1

The Nitride Deh estimate is quite close when using the adjusted utilization value, while the Wet Etch estimate is too high when using the adjusted value.

One possible reason why the Operating Curve Spreadsheet value is usually lower than the simulated value is that we're using the CV of the observed process times rather than of the effective process times. These tool groups experience quite a bit of downtime, which means that the effective process time variability would likely be significantly higher if we could calculate it.

One other observation worth noting in looking at these results is how much lower the cycle time values are for the stepper tool group with six tools than for the other two tool groups, which each have only two tools. This is consistent with our previous discussions about the impact of number of qualified tools on cycle time. See Issues 20.05 and 22.04. It's more important to stay away from the steep part of the operating curve for the Nitride Deh and Wet Etch tool groups than for the steppers (in this example).

Another Note for FabTime Software Users: Extracting these values from our demo server required using a set of several different charts and exporting some of them to Excel for additional calculations. One thing we did to streamline this process was add the charts to a home page tab and use tab filters to modify the tool group. We will work on software improvements to make this type of analysis easier in the future. For more details, including which charts to use, contact Jennifer.Robinson@FabTime.com.

Ideas for Using Operating Curves from Fab Data

If you build the operating curve for a tool group with data from your fab, you should first validate the operating curve against actual data, as we did in the above example. There are complexities in real fabs that may make this validation difficult, such as:

- Tool groups that always have a pile of WIP in front of them during the time period of interest, which can affect the right values to use for arrival CV.
- Tool groups where the WIP is held at a prior step and only released when the current tool is ready (e.g., due to time constraints). Here the queue time will be incurred at the prior tool, and not the current tool.
- Tool groups with significant standby-WIP-waiting time (as observed in the example).

But it's worth taking some time to make sure that the operating curves you observe make sense for the tool group you are looking at prior to drawing any conclusions. In general, the operating curves will probably be more representative for tool groups that are not operating right up close to 100% utilization. The queueing model assumes infinite WIP. In practice, the WIP in the fab is capped (though it might feel infinite some days).

Once you have built and validated the operating curve for a tool group, what are you going to do with it? As a starting point, you can create a separate instance of the spreadsheet for each tool group, and then vary the inputs to look at different scenarios, like:

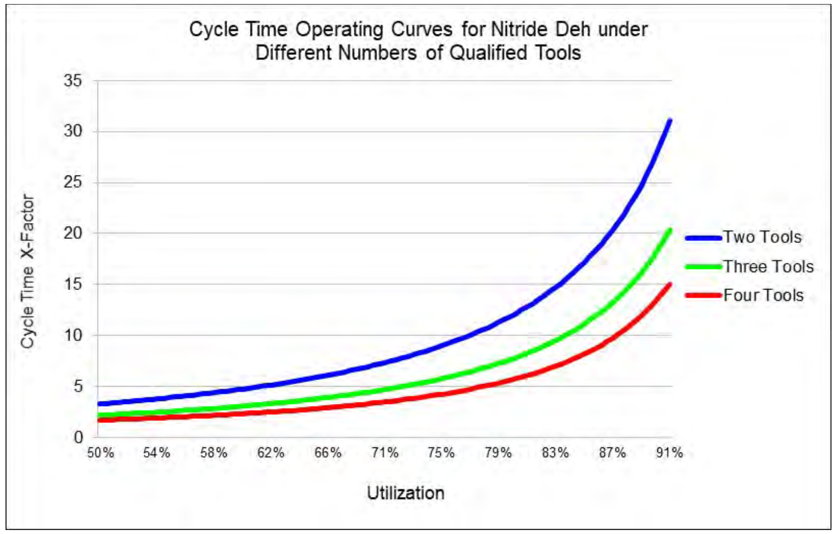

- What happens if we increase tool qualification? Or buy another tool? (Note the utilization decreases in the latter case). In the Nitride Deh operating curve from the example above, if we can increase the number of qualified tools from two to three, the average cycle time x-factor drops from 12X to 7.8X at the adjusted utilization value. Even if some extra setup time pushes up the utilization by another percent or two, the cycle time is still lower with the additional qualified tool. The resulting curves for two, three, and four qualified tools are shown to the right.
- What is the impact on the cycle time if we reduce the CV of the downtime distribution? Or reduce the MTTR? Note that we can't as easily look at different downtime percentages for the same tool group under different scenarios, since changing the percentage of downtime will affect the utilization value, putting us at a different location on the operating curve. To use the Operating Curve Spreadsheet to look at changes in the percentage of downtime, it's best to create separate versions of the spreadsheet. In each case, the utilization on the x-axis will be Productive Time / [Productive + Standby Time] for that percent downtime.
- What happens if we can reduce or eliminate the standby-WIP-waiting time, moving us to a different spot on the operating curve? Here you would leave the inputs the same for all three scenarios and look at being at different spots on the utilization curve.
- What will happen if our start rate increases, and our utilization increases? What must happen with downtime distribution or arrival variability to mitigate this impact?

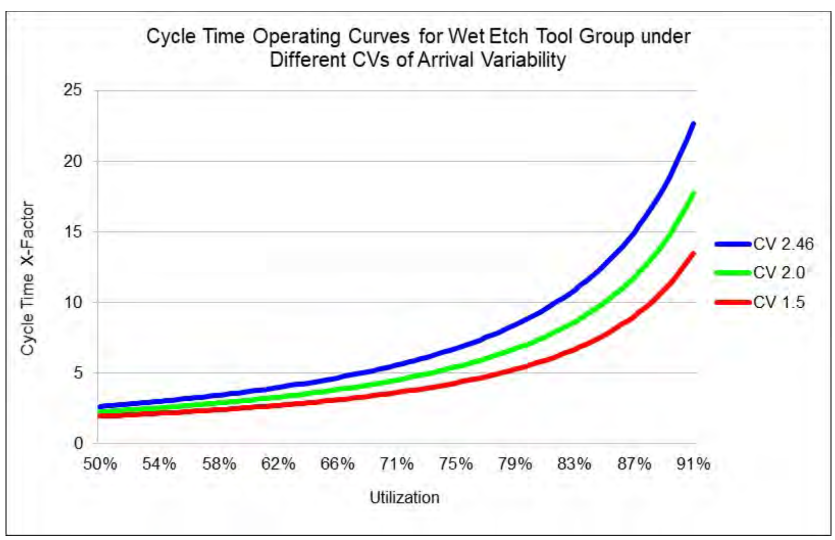

- What happens if we reduce the arrival variability? In the examples above, because the arrival variability is quite high, the cycle times drop significantly with reduction in arrival variability. Looking at an example of this potential improvement might be motivational for driving variability reduction efforts in the fab. An example varying the CV of arrivals for the Wet Etch tool group is shown to the right.

The other target that we had for building operating curves of real tool groups was to identify places where the cycle time in the fab is significantly higher than that predicted by the model. Where you see that happening, it's worth digging into the fab data to see what you might be missing:

- Is there a significant amount of standby-WIP-waiting for this tool group?
- Is there soft dedication, causing there to be a smaller number of qualified tools in practice than you expect?
- Is this a case where the effective process time variability is particularly high due to downtime, or how the tool is being operated?
- Is there more arrival variability than you are capturing because you're using move outs instead of arrival transactions? (This would probably be reflected broadly across different tool groups.)
- Is there something else that isn't being captured that is significant for this type of tool? Lot size variability? Something unusual in the dispatching? Holds?

There is no shortage of real-world complexities in wafer fabs to consider.

Conclusions

In this newsletter and our cycle time class, we've turned frequently to queueing model-based operating curves to illustrate the impact of various factors on cycle time. In this article, we've proposed taking the next step with these operating curves and using them to explore the actual and potential behavior of tool groups from real fabs.

The idea is not to exactly model what's going on in your fab via these operating curves. Rather, the idea is to see if you can build representative models of some key tool groups, and then use those models to identify opportunities for improvement. That's what really matters – making changes in the fab that reduce cycle time, so that you can get orders out to customers more quickly.

If you try this, please let us know how it turns out. We would love to hear (and share) your feedback.

Closing Questions for Newsletter Subscribers

Have you ever used operating curves to identify improvement opportunities for your fab? How did you generate them? From historical data, simulation models, or queueing models?

Acknowledgements

FabTime is grateful to participants in our recent Cycle Time Management Course sessions for discussions about the Operating Curve Spreadsheet. Thanks also to one of our customers for discussions about the use of the Operating Curve Spreadsheet in practice, and to **Thomas Beeg from Fabmatics** for discussions about operating curves on his blog.

Further Reading

All the past issues mentioned above, as well as the Operating Curve Spreadsheet, are available for download by subscribers from the FabTime newsletter archive. The current password is "FabTimeCommunity". See also:

- L.F.P. Etman, C.P.L. Veeger, E. Lefeber, I.J.B.F. Adan and J.E. Rooda, "Aggregate Modelling of Semiconductor Equipment Using Effective Process Times," Proceedings of the 2011 Winter Simulation Conference, 2011.
- J. H. Jacobs, L. F. P. Etman, E. J. J. van Campen, J. E. Rooda, "Characterization of Operational Time Variability using Effective Process Times," *IEEE Transactions on Semiconductor Manufacturing*, Vol. 16, No. 3, 511-520, 2003.

Subscriber List

Total number of subscribers: 2885

Top 20 subscribing companies:

- onsemi (198)
- Infineon Technologies (150)
- Analog Devices (includes Maxim) (138)
- Intel Corporation (125)
- Micron Technology, Inc. (120)
- GlobalFoundries (104)
- NXP Semiconductors (82)
- Skyworks Solutions, Inc. (81)
- Microchip Technology (72)
- Carsem M Sdn Bhd (69)
- STMicroelectronics (69)
- Western Digital Corporation Inc. (63)
- Seagate Technology (57)
- Texas Instruments (55)
- X-FAB Inc. (52)
- Wolfspeed, Inc. (44)
- Qualcomm (38)
- Tower Semiconductor (32)
- Hitachi Energy Ltd. (30)
- Honeywell (30)

Top 3 subscribing universities:

- Ecole des Mines de Saint-Etienne (EMSE) (9)
- Arizona State University (8)
- Virginia Tech (7)

New companies and universities this month:

- Forge Nano
- IMT Atlantique
- Industry Consulting Strobl
- Innolume
- Merit Sensor Systems, Inc.
- NexGen Power Systems
- Onto Innovation
- Silicon Mountain Contract Services
- Spectricity

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, or use the Contact form.

To subscribe to the newsletter, send email to newsletter@FabTime.com, or use the form at www.FabTime.com/newsletter-subscribe.php. 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.