

FabTime® Newsletter

Volume 27, No. 1

March 2026

Information

Publisher: Acquired by INFICON in early 2024, FabTime has been helping fabs with cycle time and performance improvement since 1999. FabTime's [flexible reporting software](#), [cycle time management course](#), and this newsletter are now part of the INFICON [Intelligent Manufacturing Systems \(IMS\)](#) group.

Editor: Jennifer Robinson, Cycle Time Evangelist for INFICON

Date: Tuesday, March 17, 2026 – Vol. 27, No. 1

Keywords: Fab Cycle Time; Performance Improvement; Waddington Effect; Metrics and Goals

Table of Contents

- Welcome
- Community News/Announcements
- Interesting Reads
- Subscriber Discussion Forum – Expansion to recommended cycle time improvement process; the Waddington Effect in general aviation
- **Main Topic – Choices in Fab Management: Part 3: Trading Better Cycle Time for the Status Quo**
- Current Subscribers

Welcome

Welcome to Vol. 27, No. 1 of the FabTime Cycle Time Newsletter. For our main topic, we complete our three-part series on tradeoffs in fab cycle time management, highlighting five instances where the status quo can impede cycle time improvement efforts. We also share follow-up discussion for subscribers on two previous topics. We have announcements about the recent FOA Summit in Austin, TX, an upcoming user group meeting, and a new INFICON paper on process time prediction.

Thanks for reading! – Jennifer

Community News/Announcements

In this month's community announcements, we share a recap from the recent FOA Summit, held in Austin, TX, a User Group Meeting announcement, and a new INFICON article about using machine learning and advanced statistics for process time prediction. To suggest announcements, please [reach out to Jennifer via this form](#).

INFICON team attends the FOA summit in Austin, TX, sponsors poker night

Chris Bode, Emilio Martinez, Jennifer Robinson, Mario Faria, John Behnke, and Dave Wizelman (as shown left to right in the above photo) represented INFICON at the recent Fab Owners Alliance (FOA) Summit in Austin, TX.



Jennifer, who is on the Women of the FOA (WFOA) advisory council, as well as the overall FOA Supplier Partner Advisory Council, especially enjoyed the WFOA event, which included icebreakers, dinner, a panel-led group discussion, and a painting activity. The photo below shows the WFOA meeting attendees – a much larger group of women than could be found at any FOA event prior to three years ago.



At the main meeting, John Behnke presented a case study with STMicroelectronics about INFICON's new Impact Manager product. INFICON also sponsored the post-meeting Poker

Tournament event, where much fun was had by all. The next FOA meeting will be held in July at the Northrop Grumman facility in Maryland. Our team hopes to see you there!

INFICON User Group Meeting to be held in Southeast Asia in late 2026

INFICON is currently assessing options for an **in-person User Group Meeting** in late 2026. Our current plan is to host the event in Southeast Asia in September. All FPS, FabTime, and FabGuard customers are welcome to attend. Please contact [Mike Neel](#), Director of Marketing, for more information.

AI team uses machine learning and advanced statistical models for process time prediction and smarter scheduling

Readers may enjoy a recently published article by **Igor Kuvychko** and **Michael Neel** from INFICON in the November/December 2025 issue of Semiconductor Digest: [The Critical Role of Process Time Prediction in Semiconductor Manufacturing](#). The article discusses why process time prediction is important for optimal scheduling decisions and how machine learning and advanced statistical models can help drive smarter scheduling. The article includes lessons learned from INFICON's work in this area, as well as predictions about the road ahead for advanced prediction models.

Interesting Reads

Semiconductor Insights and Productivity Tips

Recent articles shared on Jennifer's LinkedIn include:

- Newsletter friend **Thomas Beeg** ran a survey on his blog, Factory Physics and Automation, asking: When will we see the first humanoid robots in wafer FABs in regular production? His initial post briefly [outlined the rationale for bringing in these robots](#), as well as the many challenges in the way. The survey found that "about 50% of the voters expect deployment within the next 3 years, but there is also a large fraction which points to a much later date." [See more details here](#). Meanwhile, an article that I read in my local San Jose paper shared a more [skeptical view on humanoid robots from Roomba co-creator Rodney Brooks](#).
- Readers won't want to miss **Liz Allan's** [Global IC Fabs and Facilities report](#) in Semiconductor Engineering. "Semiconductor companies announced a significant number of facilities in 2025 as global onshoring efforts continued across manufacturing, materials, packaging, design, and R&D." The article includes a table of "more than 170 notable chip industry fab or facility investments and updates in 2025." A highlights section near the top provides a quicker overview.
- See also one sign of good news in the semiconductor wafer fab ecosystem from STMicroelectronics, as [reported by Mauro Orru in the Wall Street Journal](#), and [another from Infineon Technologies](#), which is "increas(ing) investments in the fiscal year through September to boost manufacturing as semiconductor demand for artificial intelligence keeps booming."
- **Robert Quinn** has some [interesting thoughts about developments in the semiconductor industry in 2026](#). There's thought-provoking discussion in the comments, too. You may not agree with all (or any) of his predictions, but the article is well worth reading.

- [According to the WSJ](#), TSMC “is planning a significant expansion in Arizona as part of a potential trade agreement between the U.S. and Taiwan... With the new pledge, the company would build several new fabs in Arizona, which would bring its total footprint there to roughly a dozen plants.” Of note, TSMC “has also committed to building two new factories to make semiconductors known as packaging chips, which provide supporting functions.” See also [confirmation of the signed agreement](#).

For more industry news, [connect with Jennifer on LinkedIn](#).

Subscriber Discussion Forum

We have subscriber discussion about an optional fourth step in our recommended process for improving fab cycle time, as well as a new study showing the Waddington Effect. If there is a topic you’ve been wondering about, [please let us know](#).

Extension to our three-step process for improving wafer fab cycle time: an optional fourth step about explaining what you’ve learned to others

I was reading [an article recently by Arthur Brooks](#). A prescription for happiness that Brooks included struck me. He said (emphasis mine): “To get enduringly happier, you need to **understand the science, change your habits and behavior** in tangible ways, and **explain what you have learned to others**.”

This reminded me of the newsletter article that I published in Issue 26.01: [Why and How Fabs Should Focus on CT Improvement](#). I proposed that to improve fab cycle time, people need to **understand the fundamentals of factory behavior, collect actual data** on where the problems are, and **change behaviors in the factory**. The first and third of these match Brooks’ prescription (understanding the science and making changes in behavior). I added a recommendation for data collection that Brooks didn’t mention. My three-step process for cycle time improvement didn’t include Brooks’ final point about being able to explain what you have learned to others, but I think that could be a useful optional fourth step.

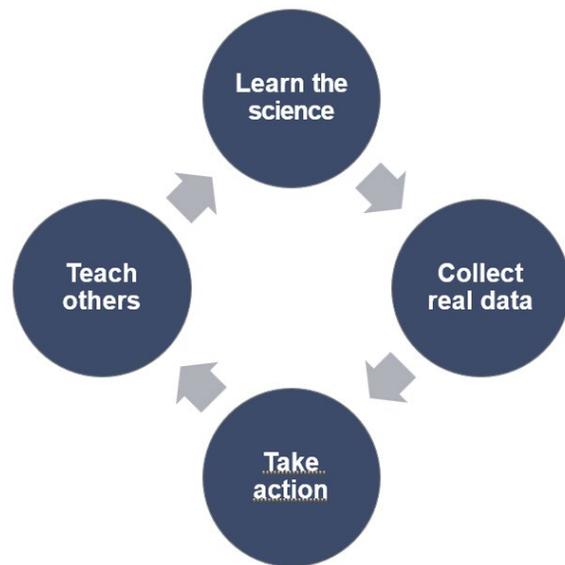
There are two ways to think about sharing your learning with others in the context of fab cycle time improvement. First, participating in a fab cycle time improvement class or reading these newsletters or other resources can increase understanding. But the best way to know that you really understand something is to be able to teach it to someone else. For me, the process of writing the newsletter and teaching cycle time classes helps me to better understand each topic that I cover. It’s easy to gloss over in my own mind when I kind of grasp something. But to share it with someone else, I have to really dig into it and understand it at another level. This is true for factory physics or anything else one learns about. If you try teaching someone else about what you’ve learned, you’ll help them and enhance your own understanding.

Also, in the context of fab operations, it’s difficult, perhaps impossible, to explain what you’ve learned to others unless you have some sort of images to show. This is probably true for anything complex. I recommend leaning on graphs or other images when you do teach people things. I think this is a real benefit of FabTime and Factory Dashboard: you don’t just get a set of numbers telling you what’s happening. You get graphs that can help you to convey as much information to others as possible.

So, to improve your fab operations, I now recommend that you:

1. Learn about the fundamentals of fab behavior.
2. Collect actual data to see where the current problems lie.
3. Take action to change behavior in the fab.
4. Teach steps 1 to 3 to others in your organization.

Let me know how that works out for you!



The Waddington Effect on small general aviation airplanes

Back in [Issue 25.06](#), we wrote about the Waddington Effect, which found that (according to research on bombers in World War II) unscheduled downtime events increase for a time after scheduled maintenance, rather than decreasing. We thought that readers would be interested to learn about a more recent study that observed the same effect in small general aviation (GA) airplanes. A November 2025 [article on the Savvy Aviation website](#) describes and links to a pre-print of the study by **Daniele Paolo Scarpazza** and **Joseph A. Hutter**. (Thank you to **Frank Chance** for the link.)

Per Savvy Aviation, “(t)he analysis found that during the first hour after a piston GA airplane flies after its annual or 100-hour inspection, it is 33.8 percent more likely to have an aircraft-caused accident or serious incident than the human-caused event baseline. This declines to 30.0 percent during the second hour, 27.5 percent during the third hour—sure sounds like the Waddington effect, doesn’t it? —and continues to decline hour-by-hour until the risk of aircraft-caused and human-caused events becomes equal after 31 hours TSLI.” The author concludes that if the analysis is correct, the FAA’s mandated 100-hour inspections for GA airplanes may well be causing more harm than good.

We welcome the opportunity to publish subscriber discussion questions and responses. [Submit your responses here](#).

Main Article: Choices in Fab Management: Part 3: Trading Better Cycle Time for the Status Quo

By Jennifer Robinson

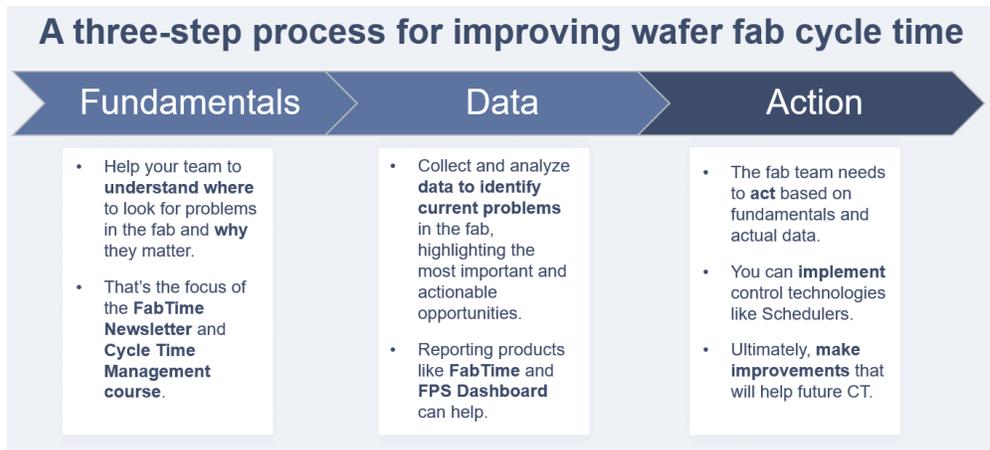
A goal in most wafer fabs is low, predictable cycle time. While there are structural conditions in fabs that make achieving this goal difficult (high product mix, reentrant flow, time constraints), there are also **choices** that fab teams make that drive up cycle times. It's been my experience that fabs regularly trade better cycle time for three things: lower costs/higher revenue, higher yields, and what we'll call the status quo (resistance to change).

We covered choices regarding cost and cycle time in [Issue 26.04](#) and choices regarding cycle time and yield in [Issue 26.05](#). In this article, we'll discuss the third of these: changes that fabs could make that would likely reduce cycle time, but that require overcoming the powerful force of inertia.

A sign of this resistance to change is hearing statements like:

- "But we've always done it this way."
- "There's no reason to make that change *today*."
- "It's what management wants. What can we do?"

Back in Issue 26.01: [Why and How Fabs Should Focus on CT Improvement](#), I proposed a three-step process for improving fab cycle time (shown below). After learning about the fundamentals and collecting data, I highlighted the need for the fab team to act. (You can read about a potential fourth step in this month's [Subscriber Discussion Forum](#).) The fact is, in fabs and in life, it's not unusual to know that a change would be beneficial and still choose not to act. Or, at least, to choose not to act just yet. Below, we share some specific examples drawn from fabs. Our hope is that reviewing these may trigger a re-assessment of the status quo for some fabs.

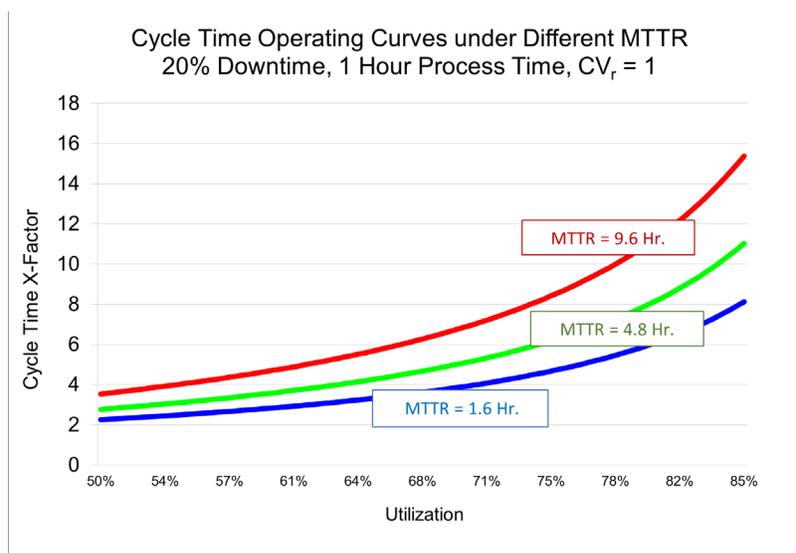


Five habits to break to improve wafer fab cycle time

The easiest way to do something, on any given day, is to do it the same way you did it yesterday. Below, we review five operating practices that could, if changed, result in better wafer fab cycle time. We discuss the likely outcome of each vote for the status quo and suggest ideas for mitigating those impacts.

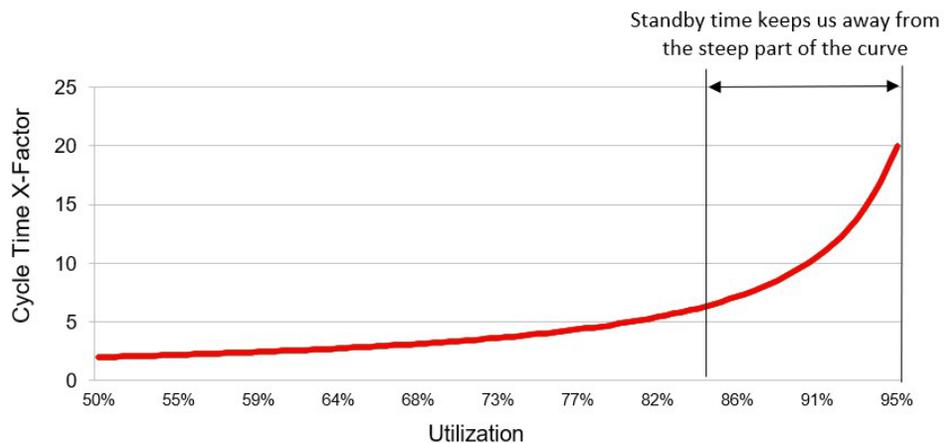
Of course, not all fabs do all these things. But I think most readers will find at least a few of these familiar. These are all behaviors that I've heard about over the years from people who work in fabs.

1. **Group PMs and other unavailable periods:** When a tool is down, there's a natural tendency for the maintenance team to say: "Since this tool is down anyway, let's also take care of this other required maintenance."
 - **Likely outcome:** Grouping PMs can be efficient in reducing the amount of required qualification time for the tool, and in the allocation of the maintenance team's time. However, this practice is generally terrible for cycle time, especially for one-of-a-kind tools. The issue is that longer periods of unavailable time allow WIP bubbles to build up, resulting in higher queue times. This is described in detail in Issue 22.01 [On Breaking Up PMs and Other Unavailable Periods](#).
 - As an illustration of this behavior, consider the operating curves below. In all cases, a one-of-a-kind tool is unavailable 20% of the time. On the blue curve the tool is taken down, on average, every 8 hours for 1.6 hours (20%). On the green curve, the tool is down once every 24 hours for 4.8 hours (still 20%). On the red curve, the tool is down once every 72 hours for 14.4 hours (still 20%). We can see that the x-factor is about twice as high for the red curve as it is for the blue curve. Even if, on the blue curve, the tool needs to spend some extra time doing quals, such that the utilization is a bit higher, it's still much better to be on the blue curve than the red one. If you were to ask people from the manufacturing team with long experience whether they would prefer short, frequent interruptions or long, infrequent interruptions, they will almost always pick the short, frequent interruptions. This is because they've experienced the pain of fab disruption that comes from long downtime events.



- **What you can do:** The best way to keep this behavior from occurring is to move from tracking mean time between failures (MTBF) to tracking either mean time to repair (MTTR) or green-to-green time. (Green-to-green time is the total time elapsed between when the tool first becomes unavailable until it becomes available again, even if several scheduled and/or unscheduled downtime instances were logged in the interim. See [Issue 25.04](#) or Issue 20.02 for details). Keeping the average duration of unavailable time low while continuing to monitor overall availability will help keep cycle times reasonable. Because this dynamic tends to be counter-intuitive for many people, this is an area in which cycle time training can be beneficial. See our [cycle time class](#) for more information.
2. **Let operator preferences lead to one-of-a-kind tools:** In fabs that give operators some degree of flexibility regarding which lot to run on which tool, we sometimes observe operator preferences for certain tools. This is called “soft dedication.” It can occur, for example, if one tool in a tool group is inconveniently located, or if a newer tool is preferred by the operators to the older version. (Or, I suppose, if the older tool is preferred to the newer version.)
- **Likely outcome:** Soft dedication leads to lower utilization levels on some tools, and higher utilization levels on others. This is typically a mis-match from what was expected in the fab’s capacity planning system. The tools with higher-than-expected utilization will have less buffer to protect against variability and hence will have higher cycle times.
 - The graph below shows the relationship between cycle time and utilization for a tool, and how standby time keeps the tool away from the steep part of the curve. Soft dedication reduces standby time for the tools that the operators prefer, resulting in higher cycle time per visit. The increase is non-linear, so this is particularly an issue when it occurs on a tool that already has a relatively high utilization rate. For tool groups with considerable excess capacity, soft dedication is less of an issue.

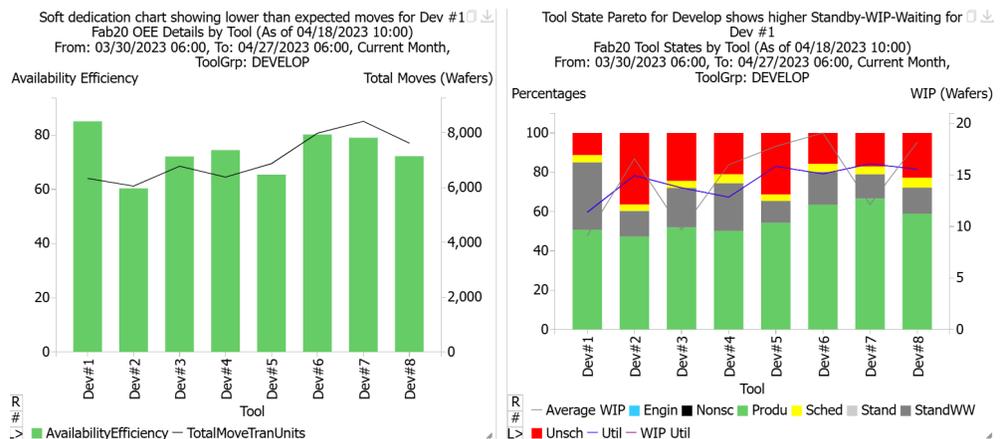
Utilization = Productive Time / (Productive + Standby Time)
X-Factor = 1 / (1 - Utilization)
As Standby Time approaches zero, Utilization approaches 100%



(c) 2024 by INFICON Inc. All rights reserved.

- **What you can do:** Implementing a scheduler and requiring schedule compliance will drive down soft dedication. Highly automated and scheduled fabs don't tend to see this issue.
- For fabs where operators have more leeway, we recommend checking to see if there are tool groups where soft dedication is contributing to high per-visit cycle times. You can do this by investigating tool groups that have high per-visit cycle times and comparing the moves per tool with availability efficiency. You're looking for mis-matches, where one or more tools has a lower move rate than would be predicted by the availability. High standby-WIP-waiting time on some tools in the tool group but not others can also be an indication of soft dedication, as the less favored tools sat idle even as WIP waited for the more favored tools to be available.

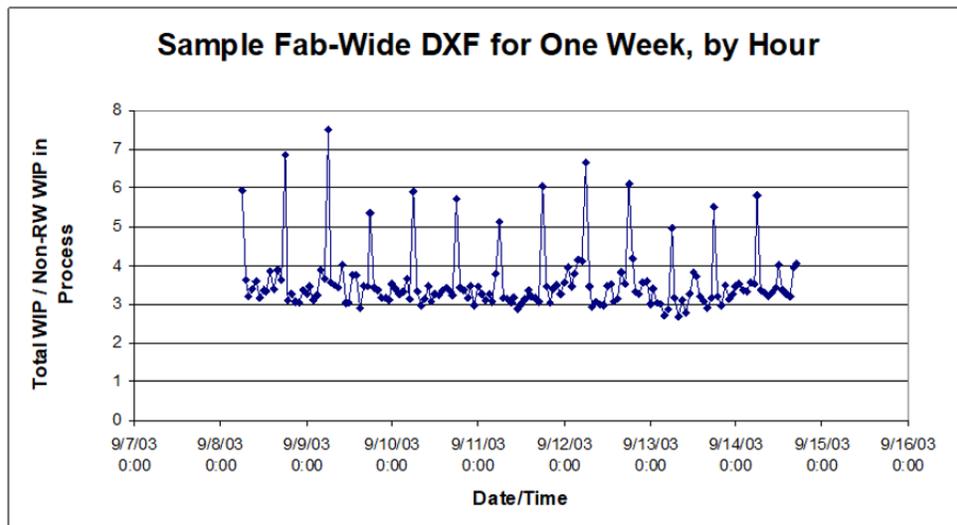
In the example below from the FabTime demo server, the chart on the left shows lower than expected moves (the black line) for Dev #1 compared to the other tools. The chart on the right shows higher than expected standby-WIP-waiting time for Dev #1.



- If you do find that soft dedication is increasing per visit cycle times, it makes sense to crack down on dispatch compliance in those areas. A mixture of education (about why we want similar utilization levels on all tools in a tool group) and data (showing how cycle time has increased because of the soft dedication in a particular area) can help drive change.
3. **Put lots on hold for down tools:** When a tool is expected to be unavailable for an extended period, some fabs will put lots requiring that tool on hold. The idea here is to let the operators know that there is nothing that needs to be done with those lots right now.
- **Likely outcome:** There are two problems with putting lots on hold for a down tool. First, the impact of the downtime on the lots' cycle time is masked. This can make it harder for maintenance techs to know which tools are the top priority for repairs and make it harder to justify tool improvement projects or spare parts contracts down the road. Second, the fab might not have systems to automatically take the lots back off hold when the tool becomes available. This unnecessarily inflates lot cycle time.
 - **What you can do:** The best thing is not to do this. Our recommendation is to not put lots on hold when the hold masks the reason for the lot's delay. If you are

going to do this, it's important to use a specific hold code that allows for later analysis (e.g., "Hold for Down Tool"). The more detailed the better. It's also very important to have some automated procedure for taking lots back off hold when the tool becomes available. At a minimum, you'll want to make sure that the list of lots on hold is reviewed every shift, to identify any lots that may have mistakenly slipped through the cracks. Even with these precautions, you will still have the issue of not being able to see how much WIP is in queue for each down tool.

4. **Leave tools idle during shift change:** Many fabs assess operator performance based on number of moves completed per shift. When this is the case, there can be little incentive for operators to begin processing lots that will not finish by the end of the shift.
 - **Likely outcome:** If operators leave tools idle with WIP waiting near the end of the shift, fabs will lose capacity on those tools. This lost capacity shows up as standby-WIP-waiting time on tool performance charts. Standby-WIP-waiting time pushes tools to a steeper place on the operating curve and drives up cycle time. See [Issue 25.05](#) for more details.
 - This behavior can sometimes be observed when looking at hourly trends in dynamic x-factor (DXF). DXF shows, at each point in time, the total WIP in the fab divided by non-rework WIP currently running on tools. Fabs that leave tools idle during shift change will see spikes in the DXF. The total WIP doesn't change much from hour to hour, but the WIP running on tools can decrease sharply, spiking DXF. An example from a real fab is shown below. The fab's hourly DXF normally floats between three and four, but spikes by more than 50% every 12 hours, at shift change.



- **What you can do:** There are several possible approaches to dampen this shift change effect. You can:
 - Implement a scheduler, and measure operator performance relative to schedule compliance, rather than moves.
 - Stagger shift change, so that not all operators in each area leave the fab at the same time. This will, if operators are sufficiently cross-trained, dampen the impact of tools left idle.

- Supplement moves with a different metric that rewards starting lots even if they won't finish by the end of the shift. FabTime implemented a metric like this called Earned Plan Hours back in 2013, at the request of a customer. See Issue 14.01: Overcoming Productivity Losses During Shift Change for more details.
 - Incorporate tracking and minimizing standby-WIP-waiting time on tools into your fab culture.
 - These ideas boil down to the idea that it is unreasonable to expect operators to keep tools running during shift change unless your metrics and shift schedules align with that behavior.
5. **Leave cycle time out of fab metrics:** Historically, fabs focused on throughput and yields, because these were key drivers of profitability. This resulted in the prevalence of metrics like moves, shipments, and scrap, and the absence of metrics focused on cycle time.
- **Likely outcome:** As illustrated with idle time during shift change, if people aren't measured on cycle time performance you can't expect them to make behavior changes that will drive cycle time improvement. Similarly, if you don't have metrics that tell you what to tackle now to improve future cycle time performance, you may find it challenging to know what actions to implement.
 - **What you can do:** To improve cycle time, it makes sense to have metrics that account for cycle time and/or WIP, and that tell us what to change to improve future performance.

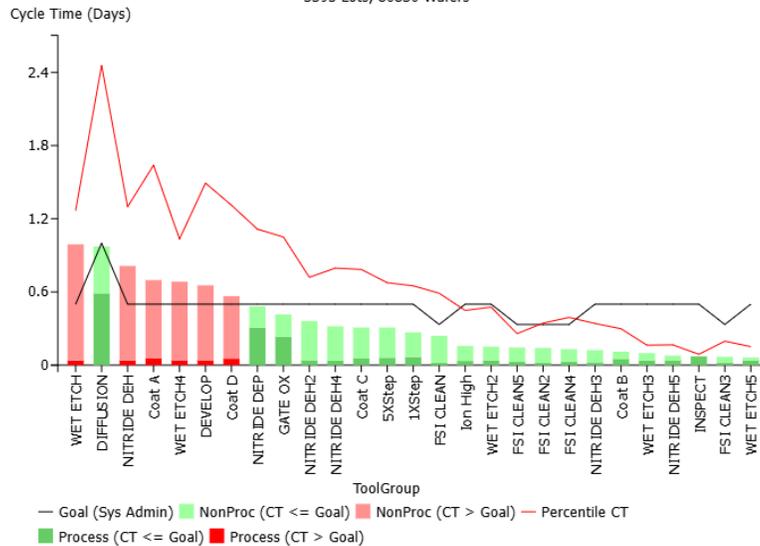
These metrics reflect **past cycle time performance:**

- **Manufacturing cycle time, x-factor, and days per mask layer for shipped lots.** These are all standard metrics that have value in comparing across fabs and benchmarking overall fab performance. Because of the typically long process flows in fabs, these metrics are less useful in identifying improvement opportunities.

These metrics help identify **current cycle time bottlenecks:**

- **WIP hours.** This is a newer metric developed by the FabTime user community to identify cycle time bottlenecks. WIP hours expands on tracking WIP by tool group by computing the estimated hours required to process the WIP that is in queue. See Issue 20.03 for more details about WIP hours, and Issues 21.01, 21.03 and 21.04 for further discussion on short-term and cycle time bottlenecks.
- **X-Factor (or queue time) by tool group** for operations completed in the past day or two. Tool groups with a high queue time per visit are excellent places to start improvement efforts. In the example below, all tool groups colored in red are short-term cycle time bottlenecks that bear attention.

Cycle Time Bottlenecks: Past 2 Days
 Fab20 Operation Cycle Time by ToolGroup
 From: 04/16/2023 06:00, To: 04/18/2023 06:00, Relative Day -2 to -1, Pct: 95
 3395 Lots, 80830 Wafers



These metrics can give a **forward look at cycle time**:

- **Dynamic x-factor**, as defined above. When DXF is measured hourly and then averaged over time, it is equivalent to shipped lot x-factor. If it starts to drift upward, we get an early indication that shipped lot cycle time will (if the situation continues) increase in the future. See Issues 9.04 and 15.05 for more information about DXF.
- **WIP turns and dynamic cycle time**. WIP turns (moves per day divided by average WIP) give us an idea of the pace of the fab. If the turns rate is declining, we know that future cycle time will increase. The dynamic cycle time metric in FabTime adds the required number of steps for each lot to convert this pace into a future cycle time estimate in days. See [Issue 24.03: Forward-Looking Cycle Time Metrics](#) for more details.
 - Adding these cycle time-focused metrics can help your fab to identify improvement opportunities on a short-term and long-term basis and inspire people on your team to make necessary changes.

Conclusions

Trade-offs can be found in all aspects of life, and certainly in the complex environment of a wafer fab. Trade-offs between cycle time and cost and between cycle time and yield are actively managed. People decide about whether to spend money on tools, staff, or software in order to reduce cycle time. (See [part one of this series](#).) They decide how much inspection and tool qualification is necessary for yield improvement, in some cases accepting higher cycle time in exchange for better yields. (See [part two of this series](#).)

There's a third type of trade-off that people make in managing fab cycle time that is a more passive choice. There are practices that could improve cycle time if they were changed, but that require overcoming the powerful force of the status quo. In this article we've highlighted five of these: grouping PMs, allowing soft dedication, putting lots on hold for down tools, allowing tools

to be idle during shift change, and using metrics that do not account for cycle time. All of these, particularly the last two, can be influenced by careful selection of fab goals.

Closing Questions for Subscribers

Does your fab do any of the things mentioned above? What else would you add to this list? What should we talk about next quarter? (Any responses shared will be shared anonymously.)

Further Resources

All past FabTime newsletters are available in PDF format from the FabTime Newsletter Archive. Please [reach out to me](#) for the link or look in the most recent email issue of the newsletter. You can download individual issues or download a zip file containing all past issues. Some articles have been re-published on the INFICON website. Those are linked above where mentioned.

For a more in-depth discussion of how these choices apply to your site, consider hosting a session of our [four-hour web-based cycle time management course](#).