INFICON
Inspired by visions. Proven by success.

FabTime® Newsletter

Volume 26, No. 4 Q3 2025

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, September 9, 2025 – Vol. 26, No. 4

Keywords: Cost and cycle time; wafer fab management; financial trade-offs

Table of Contents

Welcome

Welcome to Vol. 26, No. 4 of the FabTime Cycle Time Newsletter. In this issue, we explore seven choices in fab management that trade higher cycle time for lower cost. We discuss why these decisions increase cycle time and suggest ways to mitigate and/or quantify each effect. We have subscriber discussion about WIP pareto charts and cycle time of conveyor systems, and announcements about the recent FOA meeting and new natural language chart generation in FabTime.

Thanks for reading! – Jennifer

Community News/Announcements

In this issue's community announcements section, we have highlights from the Q3 Fab Owners Alliance meeting and exciting news about automatic chart creation from natural language queries in FabTime. To suggest industry announcements, please reach out to Jennifer via this form.

Highlights of the Q3 FOA Meeting

The Q3 SEMI Fab Owners Alliance meeting was held at Analog Devices in Beaverton, OR on July 29-30. There were 175 attendees from around the world, a mix of device makers and suppliers (plus the fabulous SEMI FOA organizing team).

Jennifer Robinson and **Birender Kahlon** represented INFICON at the meeting. Jennifer is on the advisory council for the Women of the FOA (WFOA) and is also a member of the Supplier Partner Advisory Council (SPAC).

The WFOA held a breakfast meeting with more than 30 attendees, representing SEMI's progress in encouraging more women to attend this event. Jennifer led a session on recruiting and retaining female employees in the semiconductor industry. This was followed by a panel discussion with networking tips, and a getting-to-know-you bingo event. Shown in the photo above, left to right: Christie Baker (SEMI), Jeannie Jesson (WGNSTAR), Jennifer Robinson, Vrunda Bhagwat (camLine) and Raschell Hickmott (FABMATICS USA). Shown below: the full WFOA group.

With other SPAC members, Jennifer also presented ideas for improving support and communication between FOA suppliers and device makers in the closed session for device makers. A meeting for suppliers was held in parallel with the device maker meeting, with many good ideas shared for future FOA improvements.

The next FOA meeting (device maker meeting and speed networking event) will be held at SEMICON West, in Phoenix, AZ on October 9th. INFICON will be an exhibitor at SEMICON. Look for John Behnke, our Smart Manufacturing GM, and other team members.

Natural Language Chart Creation in FabTime

In previous newsletters, we introduced AskJen™, our AI-powered chat engine based on content from this newsletter. Our newest chat-driven innovation is the ability to use natural language queries to generate FabTime charts. For example:

- Show me yesterday's fab moves by hour.
- Do I have any P1 lots on hold?
- Which tools are currently in standby-WIP-waiting?
- Show me the average cycle time per visit by tool group over the past two days for P2 lots.
- And much more…

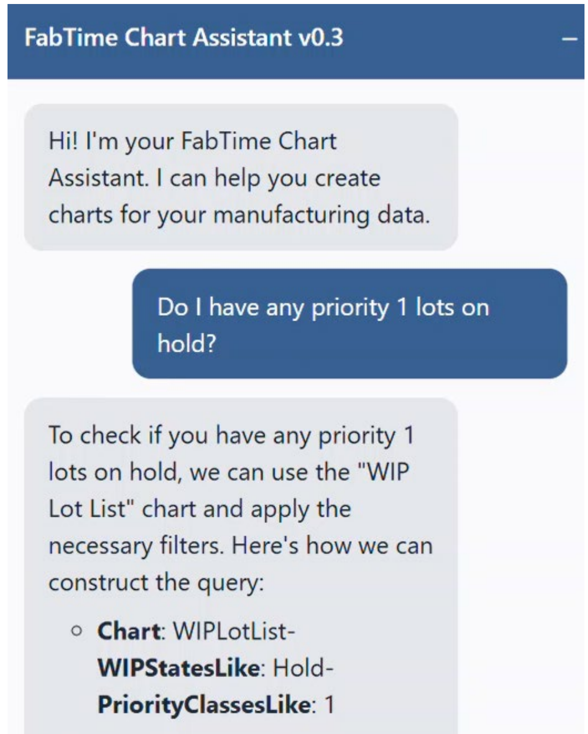

The Chart Assistant shows an overview of what chart it will use to satisfy the request, then generates the chart, already filtered as requested.

This feature is in early stages, but we believe that it's going to be incredibly useful. Reach out to Jennifer for more information or to see a demonstration.

Interesting Reads

Semiconductor Insights and Productivity Tips

Recent articles shared on Jennifer's LinkedIn include:

- The results from **Thomas Beeg's** survey regarding wafer fab automation are now available. He received responses from 24 200mm front end fabs (far fewer for other types of facilities). I personally found most interesting the list of biggest challenges faced by the fabs (efficiency, integration, labor shortage, etc.) and the list of factors preventing fabs from adding automation (ROI calculations and lack of budget top that list). But do click through and look at the whole thing.
- A WSJ article about further complexity in the semiconductor industry as "A U.S. official wants to revoke waivers that allow global chip makers to access American technology in China, people familiar with the matter said. The move could strain relations with South Korea and Taiwan, whose companies would be most affected by the change."
- An interesting piece in Tom's Hardware about "zombie fabs" in China (fab shells that were never fully built out). The reasons cited for this at the end of the article are valid everywhere:

  "China's wave of semiconductor production companies' failures highlights a fundamental reality about the chip industry: large-scale manufacturing requires more than capital and ambition. Without sustained expertise, supply chain depth, and long-term planning, even the best-funded initiatives can quickly fall apart."

- A few interesting tidbits reported in the WSJ about TSMC. TSMC is accelerating their Arizona expansion by "several quarters," reports that yields at the first fab there are already comparable to those in Taiwan and is expecting significant growth in chip demand going forward (particularly due to growth in humanoid robots).
- On the other hand, per the WSJ, Intel "announced it would again delay the target date for the completion of a $28 billion chip plant it is building near Columbus, Ohio, 'to ensure that spending is aligned with demand' for the company's chips. Originally slated to be built by the end of 2025, the project likely won't be completed until after 2030." Intel will also lay off 15% of their workforce, mainly middle management.

For more industry news, connect with Jennifer on LinkedIn.

Subscriber Discussion Forum

We have subscriber discussion about use cases for average WIP in pareto charts and what to call time between departures in a conveyor system. If there is a topic you've been wondering about, please let us know.

When should I use the regular WIP Pareto chart vs. the Average WIP Pareto chart in FabTime?

One of our INFICON team members noticed that FabTime includes both a WIP Pareto chart and an Average WIP Pareto chart and wondered about the use cases for these variants.

Response from Jennifer: Originally, FabTime used starting WIP for the WIP line on moves charts as well as for WIP trend and pareto charts. We changed this in 2011. Now, the WIP lines on the moves charts in FabTime show average WIP. Our goal with this change was to make the WIP line displayed on the moves charts more accurate, and less subject to distortion from start-of-period effects. This change also made the moves charts more consistent with the turns charts, which had previously been calculated based on average WIP.

Our next step was thinking about which value to display for WIP charts. WIP lot list charts show a list of WIP at a point in time, so have no need of an average. Standard WIP pareto charts in FabTime are also based on a single point in time, so also have no need for averages. WIP trend charts have always shown WIP at the start of each period, with the option to drill down to the WIP lot list. It made sense to continue using starting WIP here, so that when drilling down, the results would be consistent. If the data table shows 322 lots in Photo at the start of the day, and I click to drill down, I expect to see 322 lots. If I tried to click down from an average WIP value, it wouldn't be clear which WIP list to display.

However, as mentioned above, starting WIP can be volatile between periods (especially at the tool group or operation level). We created average WIP trend charts to smooth out that variability and provide consistency with turns charts (and with the WIP lines on the moves charts). We added average WIP pareto charts for similar reasons. Unlike the standard WIP pareto charts, which are generated at a single point in time, average WIP pareto charts cover a time range. Currently average WIP pareto charts are not available in stacked form.

In all cases, average WIP can be computed using either a point-average or a time-average. A point-average is computed by dividing each period into sub-periods and computing the average of starting WIP for each sub-period within the period. People who wish to use starting WIP rather than average WIP can set the sub-period length equal to the period length, resulting in average WIP being equal to starting WIP (as the “average” of a single value, the starting WIP for the period). If the sub-period is left blank, time-average WIP is computed by summing the amount of time spent at each WIP value, then dividing by the total time.

Getting back to use cases:

- Use the WIP Pareto chart if a particular point in time is of interest. For example, if we want to see inactive WIP by area right now, we could generate a WIP pareto chart, sliced by area, filtered to only include WIP that has been at its current operation for more than 12 hours (or whatever threshold you use to designate inactive WIP). A quick look at this chart will tell us which areas are struggling more with inactive WIP right now, and we can drill down to get a list of the relevant lots in each area. An example, also stacked by product, is shown below. Note this chart is for a single point in time.

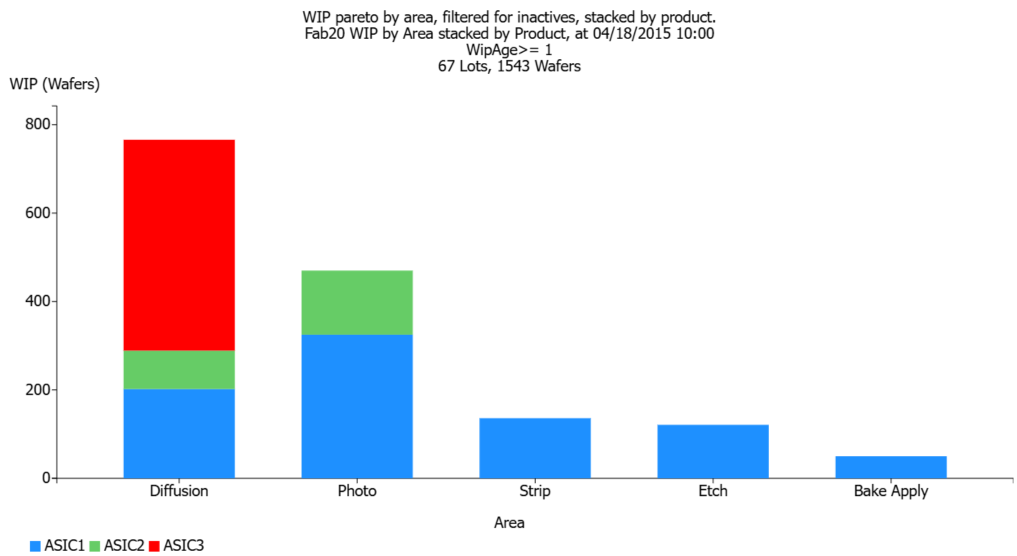

- Use the Average WIP Pareto chart to look at longer-term trends in the distribution of the WIP according to the pareto slice variable. For example, if we want to see the distribution of WIP across the line (for WIP balance), we might use the average WIP pareto chart for the current day, sliced by segment of the line. We’ll see the total WIP in each segment, averaged across the day. An example is shown below, with from and to dates for the chart covering a single workday.

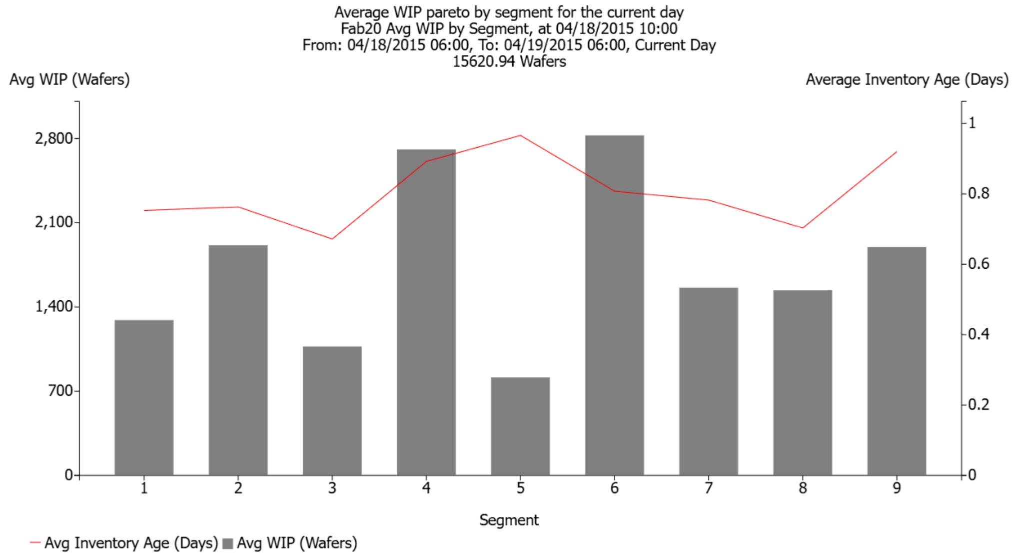

What do we call time between departures on a conveyor tool?

A **longtime subscriber wrote** to ask: "What do we call the interval between units out? For instance, let's say that parts take an hour to travel through a conveyor oven. They can be loaded a minimum of one minute apart. So, what would we call the one minute? I've always called it "takt time" even though that's technically incorrect because takt time is based on customer demand, not equipment capability. It seems that every source online says that the one minute is "cycle time." Yet in wafer fabs we've always defined "cycle time" as the time it takes for a part to travel the whole process (60 minutes in the above example). What's your take? Given wafer fabs' definition of "cycle time," what would we call this one minute?

Response from Jennifer: I have not thought about the time between lots out on a conveyor system, because I don't see that much in wafer fabs. It seems to me that your one minute is something like a load time that is required between lots. I certainly wouldn't call it cycle time. I also don't think of the one hour as cycle time, because when I think of cycle time, I think of it as including queue time. Cycle time is the time from when the lot arrives to the tool until it finishes processing. In your case, the lots have a process time of an hour, with something like a one-minute load or setup time between each lot. Because of the conveyor system, you do need to consider the total number of lots processed over time in assessing the capacity of the system. But I still think that the cycle time for each lot is one hour plus queue time.

If lots are processed continuously through the tool, then the time between departures will be one minute. But if the tool is sometimes loaded less frequently (e.g. it's not run at capacity), then you'll have longer times between departures. I've never heard time between departures called anything but that. Maybe you could say it's the pace of the tool – the rate at which you can load the tool – or perhaps the cadence of the tool.

Do any other subscribers have thoughts on this question?

We welcome the opportunity to publish subscriber discussion questions and responses. Submit your responses here.

Main Article: Fab Trade-Offs: Higher Cycle Time for Lower Cost

By Jennifer Robinson

A goal in most wafer fabs is low, predictable cycle time. While there are some 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).

In this article, we'll discuss the first of these: choices fabs make that drive cycle times higher in the interest of reducing cost or increasing revenue. We will cover choices regarding yield and the status quo in later issues. I'm sure those of you who have worked in wafer fabs will have suggestions to supplement this discussion. Feel free to reach out to me with your contributions.

Seven things that fabs do to reduce costs and/or increase revenue that can be counterproductive for cycle time

Seven examples are included below of choices fabs make in the interest of cost reduction that may have a detrimental impact on cycle time. We'll discuss the likely outcome of each choice and suggest ideas for mitigating those impacts.

Of course, not all fabs do all of 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. **Eliminate Buffer Capacity:** To cut costs, management pushes for the **elimination of standby time** on as many tool groups as possible or directs capacity planners to **plan for 100% utilization** on the various tools.
    - **Likely outcome:** Higher utilization on tools leads to higher cycle time for lots processed on those tools. See The Three Fundamental Drivers of Wafer Fab Cycle Time (Issue 6.05 or 22.04), Resolving the Cycle Time vs. Utilization Conflict (Issue 7.06) or our webinar on this subject (scroll to bottom of page) for more details. When tools are pushed near 100% utilization over the longer term, cycle time and WIP can build to unsustainable levels.
    - **What you can do:** Share the above resources to help your management team understand the reasons why a capacity buffer helps keep cycle times from getting out of control. If general articles are not sufficient, pull some per-visit cycle time data for your highest utilization tools and show how high the queue times are relative to other, lower utilization tools. Trend charts for key tools can show how short-term cycle time increases with utilization.

      As an example, the chart below shows E10 tool state data by tool group, sorted in descending order by standby-no-WIP time (the light gray bars) for our FabTime demo server. Note the gray marked line showing WIP (which we know from Little's Law correlates with cycle time). WIP is highest for the tool groups with limited standby time, and drops significantly for the other tool groups, which all have more of a standby time buffer.

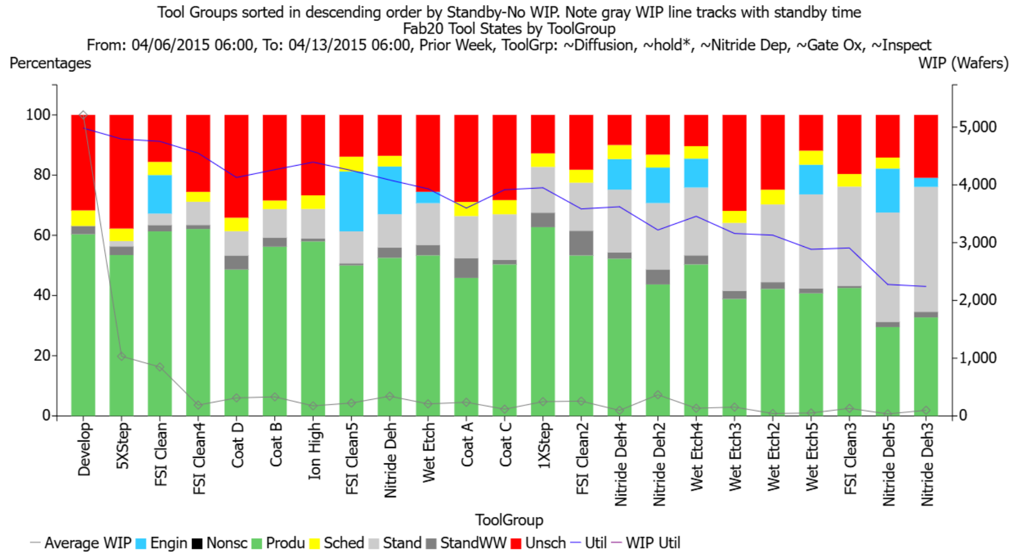

There are also things that you can do to increase the standby time buffer on tools by reducing standby-WIP-waiting time, improving equipment uptime, and scheduling engineering time more efficiently. See A New Metric for the Functional Utilization that Drives Cycle Time (Issue 25.05) for more details.

2. **Soft Idle Tools:** To reduce variable costs (electricity, consumables, maintenance) during a slow period, management directs the team to **"soft idle" certain tools** and thus run with higher utilization on the ones that remain. Typically, these tools are placed into a non-scheduled state (though there is some debate about that), where they may be re-activated following qualification if needed.

   - **Likely outcome:** In addition to the impacts from higher utilization (see above) on the remaining tools, idling tools leads to smaller tool groups. Where this results in groups with fewer than three qualified tools (especially where it results in only one qualified tool), expect cycle time per visit to be 25-100% higher than it would be otherwise. See Managing One-of-a-Kind Tools (Issue 23.05) or The Impact of Tool Qualification on Cycle Time (Issue 20.05).

     The graphic below shows estimated cycle time per visit for tool groups with different levels of redundancy (90% utilization in all cases, and moderate variability). Soft idling one of 10 tools to go to nine won't have much impact, but soft idling one of two tools to go down to one will double the per-visit cycle time.

   - **What you can do:** Push back where "soft idling" results in one-of-a-kind tools or tool groups with very little standby time. Where possible, keep the soft idled tools up to date on major PMs, so that they can be brought back online relatively quickly if an issue arises with one of the remaining tools.

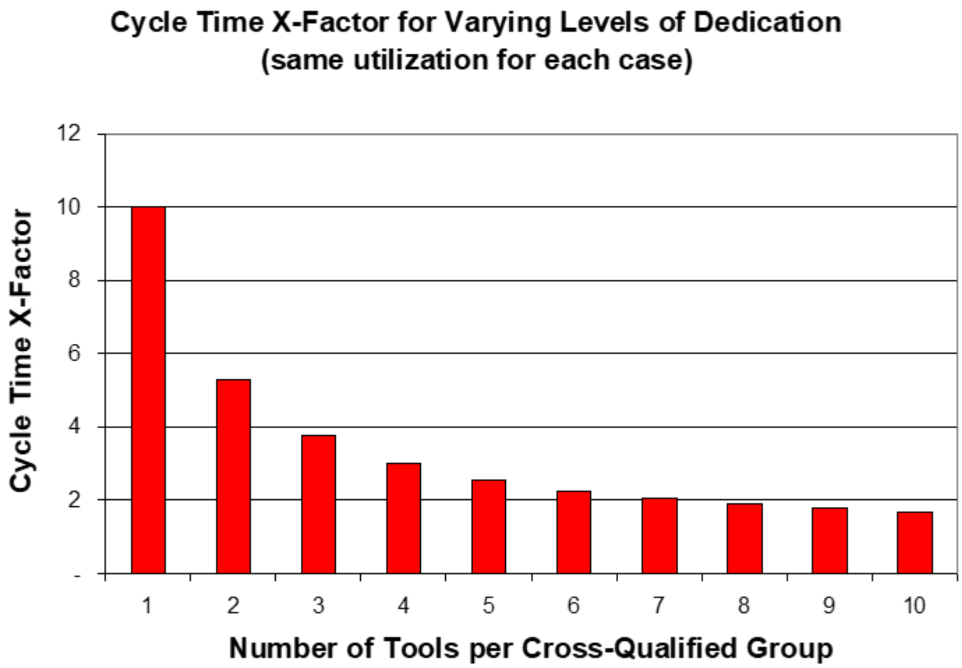

3. **Run Too Many Hot Lots:** To meet demands for key customers in the face of high overall cycle time, the production control team requests **ever-increasing numbers of hot lots**. This behavior tends to be aimed more at increasing revenue than decreasing costs.
   - **Likely outcome:** Higher variability from the hot lots (especially when they are hand carry lots) may drive average cycle time higher, resulting in the need for even more hot lots to satisfy customers. A vicious cycle will ensue.
   - **What you can do:** Push back by measuring the impact of the hot lots on the cycle time of regular lots and (if you are in a vicious cycle) show how this is increasing. See Issue 25.01 for a description of a metric to capture this, which is now available in our FabTime reporting module. See also Cycle Time and Hot Lots: Updated (Issue 19.03), available in PDF from our FabTime Newsletter Archive.
4. **Reduce Spare Parts and Service Contracts:** To reduce costs, the equipment engineering department **cuts back on spare parts and lets service contracts lapse,** and/or **cuts back on the number of in-house maintenance techs**.
   - **Likely outcome:** Longer periods of unscheduled downtime will be observed while waiting for parts or outside service technicians. Similar delays may come from waiting while the in-house maintenance tech is busy somewhere else. Longer downtimes, especially on one-of-a-kind tools, are well-known to contribute to longer cycle times. See Improve Fab Cycle Time by Tracking the Right Equipment Reliability Metrics (Issue 25.04) for more details.
   - **What you can do:** Track your unscheduled downtimes by detailed sub-states (like INFICON's Enhanced Tool Performance states) and look for time spent waiting for parts or waiting for a service technician. Observe where this correlates with high cycle time per visit. Hint: this will be most obvious on one-of-a-kind tools. Use the worst instances of long unscheduled downtimes on smaller tool groups to guide potential

changes. This can be an instance where the financial benefits of keeping key tools running clearly outweighs cost savings on the maintenance side, but you'll need to gather sufficient data to show this.

As an example, the chart below shows green-to-green instances longer than 8 hours for a collection of photo tools. The medium blue color is "waiting for parts." The red color is "waiting for tech." If these waiting states correspond to poor cycle time for the associated tools, changes might be needed. In this case (from our demo server), the 1x Steppers have extra redundancy, so the delays are probably not critical.

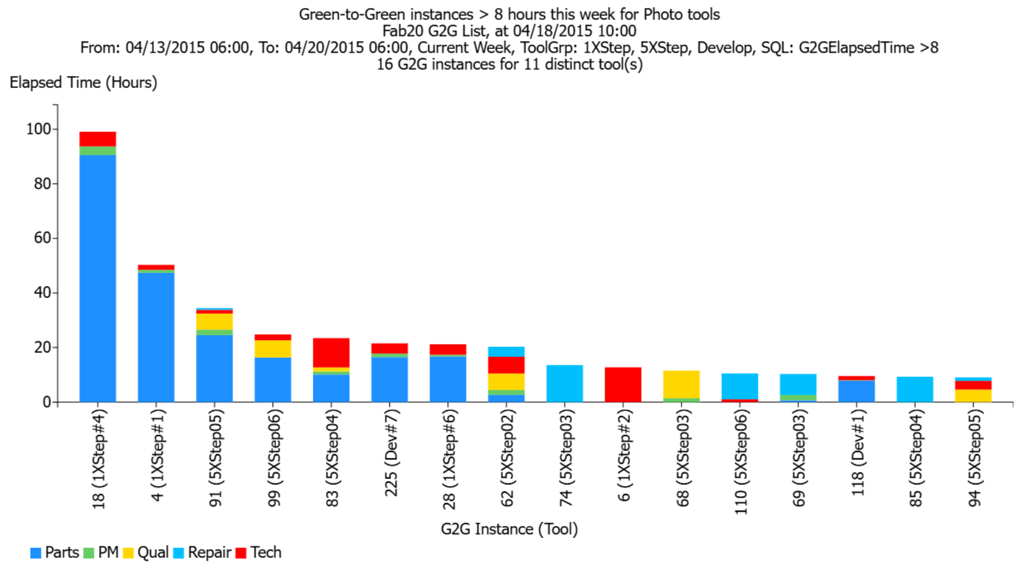

5. **Maintain an Insufficient Number of Operators:** To control costs, the production team either **fails to cross-train operators** or **doesn't hire enough operators**, such that (especially when anyone is out sick) operators end up constraining key tools in the fab.

   - **Likely outcome:** Because operators can't be in more than one place at one time, even bottleneck tools may spend time in a standby-WIP-waiting state. This is lost effective capacity that results in lower fab throughput and/or higher cycle times. Also (for manual fabs), because operators are overloaded, they may wait to fill carts before transferring lots between steps. This adds arrival variability to downstream steps, further increasing cycle time.

   - **What you can do:** Collect data on tool groups where standby-WIP-waiting time corresponds to higher per-visit cycle time or missed throughput. Reassign or recommend cross-training for operators where possible. If arrival variability due to full carts is an issue, consider a one-time investment in smaller carts, or implement policies to restrict the number of lots loaded onto any cart. An example of high standby-WIP-waiting corresponding to high cycle time per visit for a tool group is shown below. See A New Metric for the Functional Utilization that Drives Cycle Time (Issue 25.05) for more details.

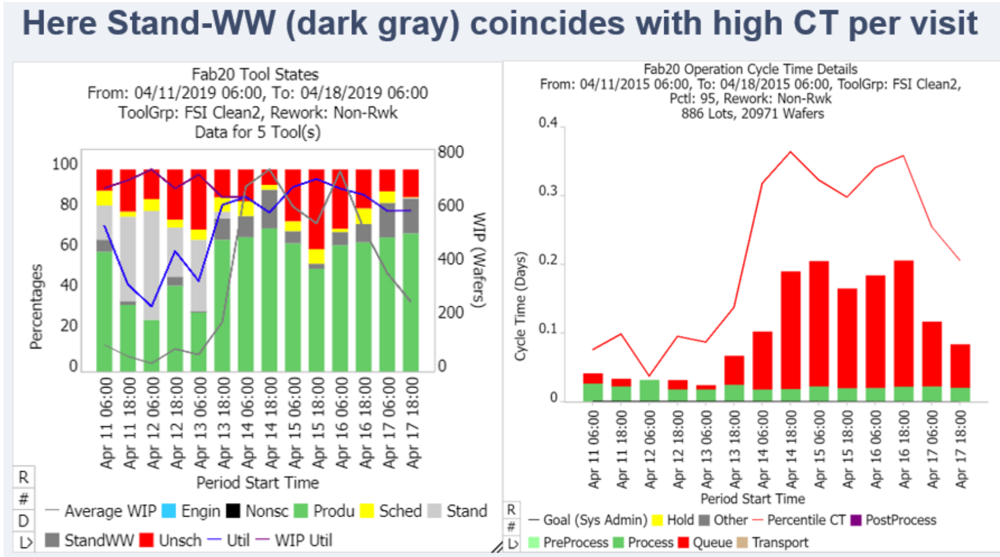

6. **Run Full Batches on Lightly Loaded Tools:** To reduce consumables costs, production management directs operators to **run (nearly) full batches on all batch tools**, including lightly loaded tools**.**

   - **Likely outcome:** For tools that are not heavily loaded, average queue time per visit can increase significantly when lots wait to form batches, as shown in the left-hand side of the image below. Arrival variability downstream of the batch tools also increases, spreading cycle time problems out through the fab. See Batch Loading Policies for Wafer Fabs (Issue 9.03) or 10 Recommendations for Fab Cycle Time Improvement (Issue 22.02), item four.

   - **What you can do:** Where consumables costs are a significant factor (e.g. a gold process), it may be necessary to run full batches. In this case, it's best to start lots that can be batched together in groups and keep them together as much as possible throughout the process flow. Fabs should bear in mind, however, that running full batches in a low utilization environment can dramatically increase fab cycle time. Where possible, the most lightly loaded tools should be run using a "greedy" policy, whereby the tool is started whenever any lots are waiting. Alternatively, fabs should try scaling back the required batch size and observe the result. This may be another instance where the cost saving from running fewer batches is not worth the cost in higher cycle time.

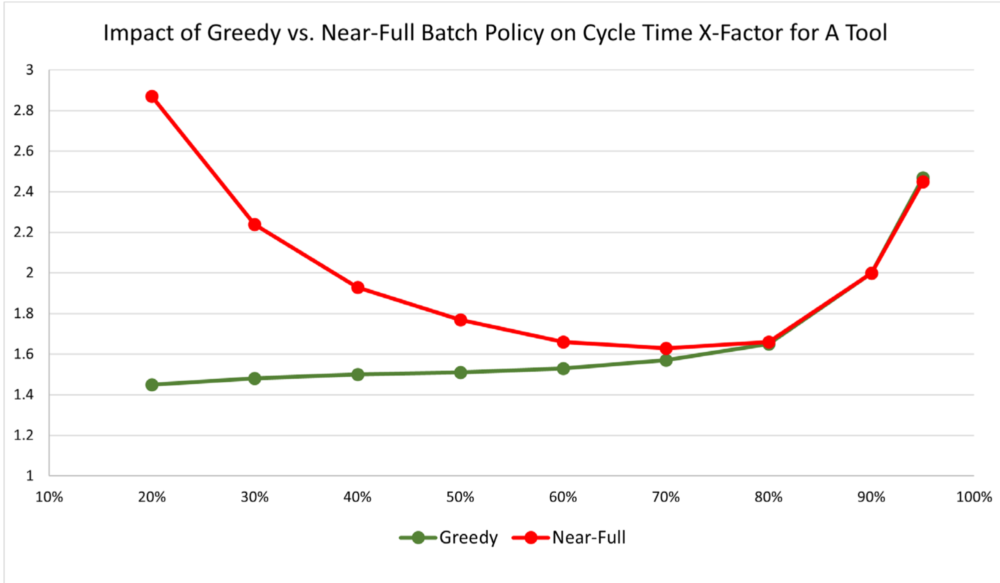

7. **Defer Investments to Make the Fab Smarter:** To keep costs down, fab management resists software investments that might make the fab "smarter." This could be anything from not having a formal MES to doing scheduling, dispatching and starts planning manually (or via spreadsheets). The fab likely relies on the experience of the production staff to make good decisions.
   - **Likely outcome:** Over time, the fab may fall behind compared with competitors. As key personnel retire or change jobs, passing on knowledge will become difficult. Customers may be unhappy about poor visibility into expected shipment dates. Industrial engineers may spend time manually recording events in the fab, rather than analyzing the data and suggesting improvements.
   - **What you can do:** Your options here depend, of course, on where you are starting from. Some of the transition nodes (like adding an MES) are big jumps, requiring a real commitment of time and money. But in many cases, there are smaller steps that you can take. You can add some sensors or FabGuard® systems to automate data collection. Even if you aren't ready for full fab scheduling, you may be able to improve your dispatch system. Getting a Digital Twin for your factory opens up all sorts of other options by helping you to store data in a structured, accessible way. Wherever you are starting from, the team at INFICON would be more than happy to talk with you about options to make your fab a little bit smarter (or a lot smarter).

     The picture below shows the SEMI Smart Manufacturing Roadmap. Wherever your fab is on the pyramid, you can benefit from moving up, and adding more integrated, autonomous solutions.

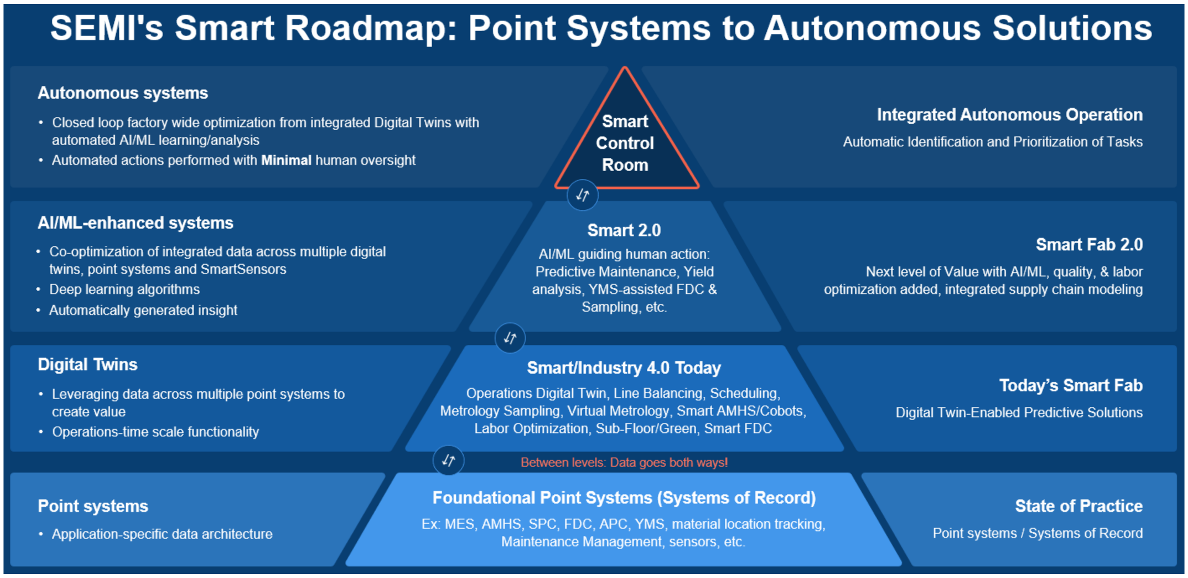

Conclusions

We all face trade-offs every day. In a wafer fab, some of the biggest trade-offs are between cycle time and cost (particularly in times when markets are uncertain). We want good cycle time, but we also want to increase throughput, to maximize return on expensive capital investment. We want good cycle time, but we also want to reduce variable costs, including labor, consumables, and spare parts. Understanding the cost side of these decisions is straightforward. Understanding the cycle time side of the equation is more challenging.

In this article, we've discussed some of the impacts that these cost-driven choices have on cycle time, and suggested ways to minimize those effects. In some cases, we recommend the use of data to better quantify the trade-offs. The goal, after all, is to optimize **overall** company performance, balancing cost and cycle time.

Closing Questions for Subscribers

Does your fab do any of the things mentioned above? Do you have other suggestions for mitigating the detrimental effect of these choices on cycle time? What else would you add to this list? What should we discuss in the next two issues (choices, or lack thereof, regarding yield and the status quo)? Your response will be kept confidential – your name and company name will not be shared publicly.

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. See also:

- 10 Recommendations for Fab Cycle Time Improvement (Issue 22.02)
- 10 More Recommendations for Improving Fab Cycle Time (Issue 24.01)

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.