FabTime Cycle Time Management Newsletter

Volume 10, No. 4 May 2009

Information

Mission: To discuss issues relating to proactive wafer fab cycle time management

Publisher: FabTime Inc. FabTime sells cycle time management software for wafer fab managers. New features in the software this month include support for factory non-scheduled time in predicting lot completion dates, and new Dispatch Precision trend, pareto, and list charts.

Editor: Jennifer Robinson

Contributors: Frans Brouwers (NXP Semiconductors); Ian Chizmar (DayStar Technologies); Eliot Parkinson (International Rectifier – Newport)

Table of Contents

Welcome

Welcome to Volume 10, Number 4 of the FabTime Cycle Time Management Newsletter! We hope that this issue finds you all well, and starting to see some improvements in your sector of the economy. In this issue, we have two conference announcements, one about the IMEC Technology Forum scheduled for Brussels in June, and another about the AEC/APC Symposium scheduled for Michigan in September. Our FabTime user tip of the month is about using new average WIP Trend and Pareto charts in FabTime (and using the same averaging functionality in the WIP Turns charts).

This month we have rolled the subscriber discussion section into the main article. We have four interesting and detailed discussions ongoing with subscribers related to: dispatch precision (a dispatch compliance metric); equipment uptime reporting (the main topic of the last issue); granularity of tool state reporting and modification of transactional data; and calculation of degree of lateness for in-process lots. In light of the substantive nature of these discussions (and with many thanks to the subscribers who have contributed), we've decided instead of a new main article to simply highlight these four topics. We welcome your feedback!

Thanks for reading!—Jennifer

FabTime

Tel: (408) 549-9932
Fax: (408) 549-9941
www.FabTime.com
Sales@FabTime.com

Community News/Announcements

IMEC Technology Forum 2009: June 2 to 4 in Brussels

We found this announcement on Future Fab Connect (an industry-focused social networking site that we recommend: http://futurefabconnect.ning.com), and thought that it looked interesting (though we will not be attending).

"What do the "Zen of Snow" application on an Apple iPhone® and satellites transmitting weather data and telephone signals have in common? To a scientist or technologist, the answer is easy: Nanoelectronics – those minute slips of technology known as integrated circuits, or semiconductors. After all, it is semiconductor technology that has enabled computing and communications to fit in the palm of your hand and let you enjoy turning your iPhone into a snow globe or carry on a business call at the airport. Consumers rarely think of all the ways that technology simplifies their lives, but executives managing the world's leading high-tech companies are constantly monitoring problems and trends and thinking of solutions for each one. After all, engineers and scientists are fix-it people.

The IMEC Technology Forum brings together executives from companies and institutes in Asia, Europe and the USA to discuss advances in science and technology and to present visions for future innovations. Formerly known as the Annual Research Review Meeting (ARRM), the IMEC Tech Forum is now in its 13th year. This year, the Forum will highlight the theme of science and technology innovation for the next 25 years to recognize IMEC's quarter-century anniversary." More details can be found at: http://www.itf2009.be.

Call for Papers: AEC/APC Symposium XXI: September 27-30, 2009 • Ann Arbor, Michigan

We received the following call for papers for the AEC/APC Symposium.

"This year's symposium will review recent technical advancements in order to assure alignment with the needs of IC manufacturers, semiconductor equipment suppliers, and software, sensor, and metrology suppliers. Advancements in related industries—such as solar, LCD, and memory devices—will also be discussed in order to assess how synergy between these industries can be better leveraged.

The symposium will be built around topics such as, but not limited to (abridged list, others are on the website):

- Factory-wide and enterprise-wide applications and deployment
- Real-time data collection and data management
- Benefits and justification (ROI, Coo, OEE)
- Tool productivity data collection/analysis
- E-diagnostics, E-manufacturing, and EEC
- APC and APC-related advancements in solar, LCD and memory devices industries
- APC applications to back-end semiconductor manufacturing

Abstracts should be one page of text (maximum of 1,000 words) and one page for figures. Full instructions for submitting abstracts can be found on the AEC/APC Symposium website (www.aecapcsymposium.org). Submission Deadline: May 22, 2009

FabTime welcomes the opportunity to publish community announcements. Send them to newsletter@FabTime.com.

FabTime User Tip of the Month

Display Average WIP Trend and Pareto Charts

The WIP Trend chart in FabTime displays, for each time period, the WIP that matches the chart's filters at the start of the time period. So, for instance, a WIP trend chart with a period length of 24 hours will show the WIP and inventory age at the start of each day. This is generally fine when looking at WIP for the fab as a whole, because total fab WIP tends not to change dramatically from hour to hour. However, when looking at more granular data, such as WIP by Area or ToolGroup, reporting the WIP at the start of each day may not be detailed enough, especially for areas that feature large WIP fluctuations from hour to hour. To address this, we've recently added (in Patch 96) two new charts: Average WIP Trend and Average WIP Pareto.

To use either of these charts, simply generate the chart from the Chart list (WIP Charts category). The default for the chart will be to show data by 24-hour periods, averaged by 12 hour sub-periods. This means that instead of displaying the WIP at the start of each day, the chart will divide each day into two, 12-hour sub-periods, and display the average starting WIP across those two sub-periods. If you would like to see a more granular average, you can change the sub-period to one hour. In this case, FabTime will take snapshots of the WIP at the start of each hour, average those, and report the resulting value.

The examples below (and on the next page) show the difference that this can make at the ToolGroup level. The first chart sets the sub-period length equal to

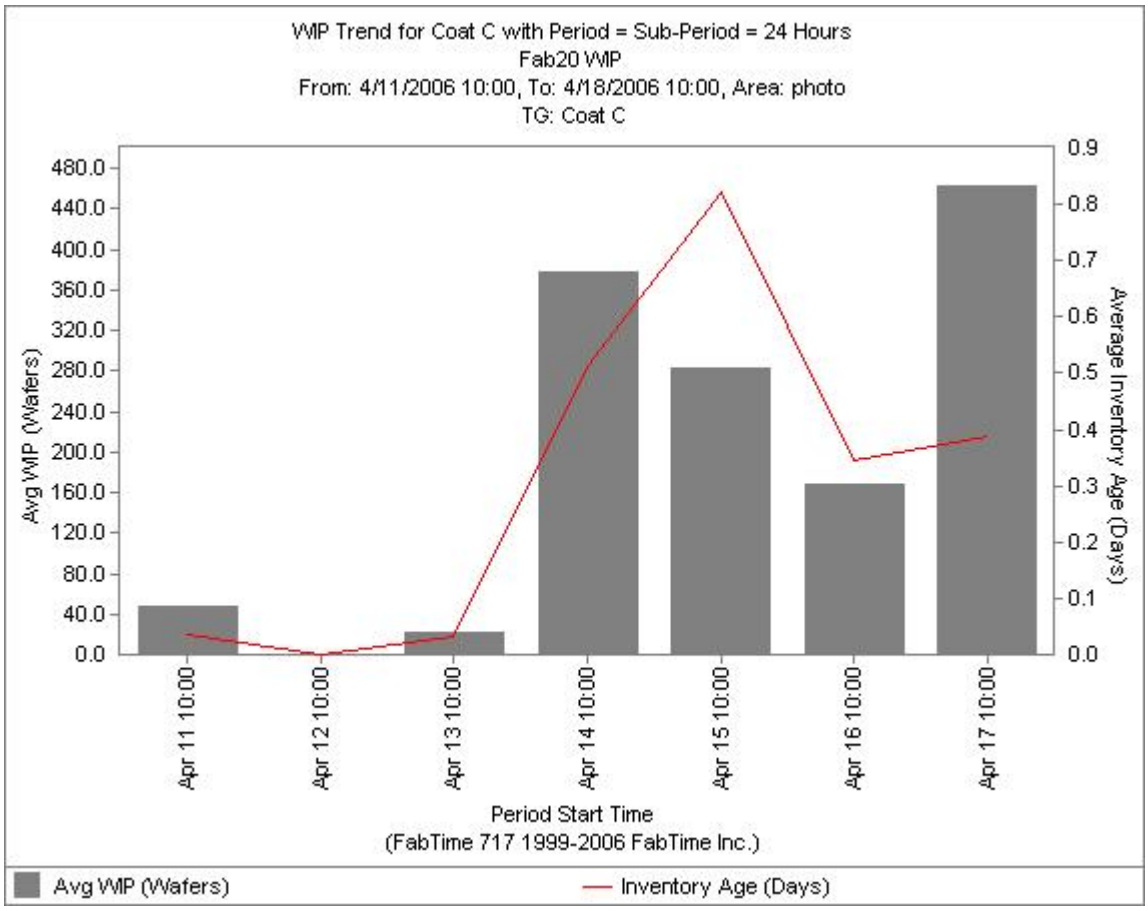

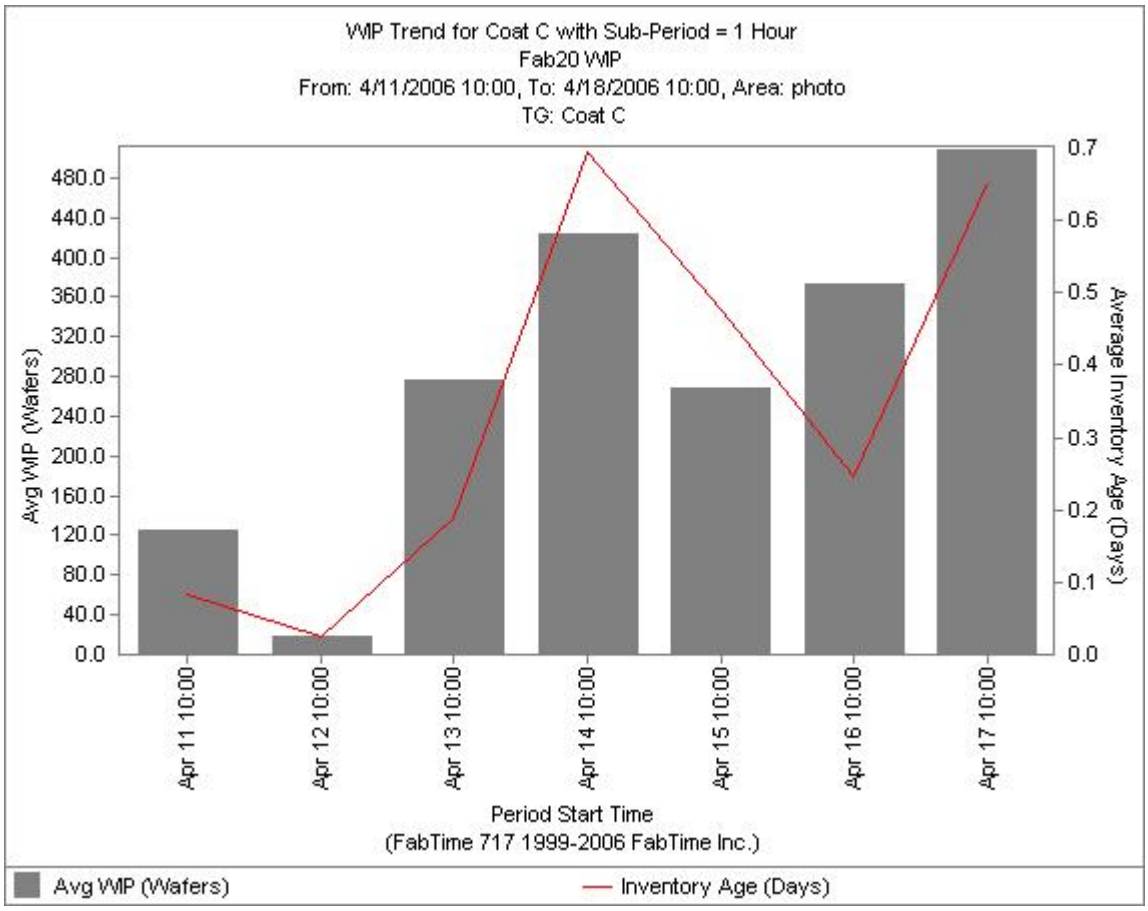

the period length (essentially, no averaging), while the second chart uses a one-hour sub-period length.

Using a small sub-period on a chart over a very long overall time period may result in slower chart generation. You can experiment with the sub-periods, and see where tighter sub-periods make a difference for your reporting. This WIP averaging capability is also now available on the WIP Turns charts, so that the resulting charts display moves for the time period, divided by Average WIP during that time period. To revert to the original versions of the Turns charts, which used starting WIP, simply set the sub-period length equal to the period length (so that no averaging will occur). You can set a default sub-period value in the Defaults section on the Chart List page. The same default will apply to Average WIP charts and WIP Turns charts.

One other comment: When using WIP Turns charts for very granular slice-by variables (operations or individual tools, for example), please note that the WIP used in the denominator of the Turns calculation is still the average WIP present at that operation or tool. It is not some sort of cumulative WIP number, showing how much WIP arrived to the tool during the period (you can use the Arrivals Charts to get that arrival data). High-speed tools that generate many moves, but never have much WIP waiting, will display very high WIP Turns values at the tool and operation level. For this reason, WIP Turns is generally more useful as a metric at a higher level (such as Area or Fab).

If you have any questions about this feature (or any other software-related issues), just use the Feedback form in the software.

Responses to Four Recent Discussion Topics

Introduction

We have four interesting and detailed discussions ongoing with subscribers related to dispatch precision (a dispatch compliance metric), equipment uptime reporting, granularity of tool state reporting, and identification of lateness for in-process lots. In light of the substantive nature of these discussions (and with many thanks to the subscribers who have contributed), we've decided instead of a new main article to simply highlight these four topics. We welcome your feedback!

Dispatch Precision

Two people wrote in response to last month's subscriber discussion topic of the dispatch compliance metric "Dispatch Precision", where we defined:

Dispatch Precision% for a Lot = 100% * (1.0 - ((Lot's order on dispatch list) - 1)*(1/(#lots on list)))

Frans Brouwers from **NXP Semiconductors** wrote: "Your proposal is a good first estimate. Most probably, when you first start measuring by this method, the figures will show a lot of room for improvement. So that is adequate to the goal. But still, I would like suggest to some more details:

- In some cases the difference between lot number 1 and lot number 2 is not important for the total fab and/or order performance. So, if, for the operator on the shopfloor, it is handy to pick lot number 2, there is no objection. But how to put this in software? In my time as production control manager I never found a solution for that. But maybe somebody else will.

- One other example. Suppose lot number 1 is a speed lot. Suppose lot number 2 is a normal lot. In this case taking lot number 2 is very much undesirable for the logistic performance. But for capacity reasons an operator might choice for lot number 2. This case might be easily put in software. For instance by giving speed lots 10 points and normal lots 1 point. Suppose there are 20 lots in the list. Taking lot number 2 would lead to a performance of 95%. I would suggest to increase the number of point subtracted by a factor of 10. This would lead to a performance of 50%.

FabTime Response: Frans' first point is, of course, quite complex. Trading off operator efficiency vs. the goals of the dispatch system can be quite tricky. Our view is that dispatch compliance scores should not be used punitively, but rather should be used for learning. For example, if the operator is consistently having trouble running the first lot, there may be a lot delivery issue, outside of the control of that operator. Dispatch compliance scores, analyzed carefully, can help to identify these situations. Regarding Frans' second point; we agree that taking lot priority into account in some fashion could be a useful extension of this method.

Ian Chizmar of **DayStar Technologies** also wrote in response to the Dispatch Precision article, asking for clarification. He said: "Is the precision cumulative, and if so, is it cumulative per process step, or per Lot? For example, at step A I have 9 Lots on my list and I choose the 3rd, so my precision for that incident is 77.8%. After that Lot processes, the next dispatch list shows 8 Lots and I choose the 2nd one. For that incident, my precision has improved to 87.5%. So is the overall precision for step A over the course of a day 77.8% * 87.5% * X… (for as many lots that I process in step A for a specified duration)? Could this be applied to the Lot itself as it moves through the line? If the Lot shows up at step A and is listed at #1, but 2 other Lots get processed before it, it's own precision could be calculated as 3rd out of however many Lots were on the List when it first arrived at that step. Then

for each subsequent step, the cumulative precision could be calculated… ? If at every move out the Lot was evaluated as to which position it held on the next dispatch list, then this could show how often fab technicians are following the dispatch list. This might get a little wacky because we all know that Lot priorities and their position in a dispatch rule evaluation can change quickly. Not to mention if the Fab tried to enforce the 1st Lot always getting selected, or if the technicians are able to override the list. Just curious if you and your team have thought about that."

FabTime Response: If we understand the question properly, the calculations are not cumulative. If numbers are rolled up they are simply averaged, not multiplied. The precision is always calculated based on the ordering of the lots at the time the dispatch list is requested by the operator, not when the lot arrives, because many of the calculations are dynamic and must be evaluated when the list is presented to the operator (e.g. a higher priority lot could appear, or downstream WIP conditions could change, due dates could be adjusted, etc). For looking at average order (rather than precision) there is a separate chart that shows the ordering, so it's easy to tell if operators are consistently choosing one of the first 2 lots, etc."

So, to clarify further, we've only envisioned this as a tool-focused metric. Whenever an operator requests a dispatch list, and chooses a lot, a dispatch precision value is generated. These values can be averaged over time, and across tools. Our other dispatch performance chart reports, again by tool, the rank in the dispatch list for the selected lot. This is then averaged over time (and can be averaged across tools, etc.). So a value of 3 means that, on average, for the tools currently included in the chart, the operators selected the third lot from the dispatch list.

We haven't looked at measuring dispatch compliance on a lot-by-lot basis. We agree that this could be a bit tricky, given changing lot priorities, etc. We do have charts that look at the lot's performance to schedule, and cycle time performance (including breakdown of process time vs. non-process time), and we think those are better metrics right now for tracking lot progress.

What do the rest of you think about this?

Issue 10.03: Equipment Availability versus Equipment Uptime and Manufacturing Time

Ian Chizmar also responded to last month's main article, about the metrics Equipment Uptime and Manufacturing Time. He said: "As for the other point of interest in your article, equipment state modeling happens to be one of my favorite topics to discuss about fab operations! I have gone through the E10/E79/OEE discussions with Industrial and Manufacturing Engineers alike, many times. I have even had the opportunity to create a 13-state equipment model derivative of E10 (the goal there was to support a 70+ "bucket" retroactive OEE reporting model). I strongly believe that tool efficiency metrics MUST be calculated from the Tool's perspective. When it comes to Engineering Time or Productive Time, the two should only be differentiated by state because the Tool is actually being used for different purposes, but Engineering Time should not "ding" the Tool's Availability Efficiency, or Operational Efficiency. As far as the Tool is concerned, it's doing exactly what it's supposed to. Even in the case of an Engineering event putting the tool in a state where Lots cannot be processed because parameter variables are being changed, the Tool is still "available" to process Lots from its perspective, it just can't stop the engineer from using itself to create recipes etc. So I guess I disagree with your separation of Equipment Uptime = (Productive + Standby + Engineering) and Manufacturing Time = (Productive +

Standby). I think they are equivalent. Rather than try to classify these differences by Up Time/Manufacturing Time, I think it's more valuable to report on how much of that Up/Manufacturing time was spent processing experiment/R&D lots, or how much time an engineer spent working on the tool without actually causing a physical "down" event. This can be captured with simple Lot attributes and special event logging. With this kind of data, decisions about WIP balancing and prioritization can be made without putting the Tool in the middle… when that happens, someone always gets misinterpreted, whether it be manufacturing, R&D, or maintenance."

FabTime Response: We agree with Ian that engineering time should be counted as "good" when tracking the performance of the maintenance team, and when looking at how well the tool is performing relative to, say, what the vendor promised. However, we still think that there is value, from an operations perspective, in knowing the total time that a tool was available for use by the manufacturing organization, and how much of that time was actually used to process wafers. This is the ratio (productive time / productive + standby) that drives cycle time performance of the tool. From a cycle time perspective, it doesn't matter at all if the tool was unavailable because it was down or because it was being used by an engineer - it still wasn't available for running production wafers, or as standby time to buffer against variability. So, while we agree that there's a place for tool-centered metrics, and that we should include engineering time in equipment uptime for that reason, we still believe that there is also a place for a Manufacturing Uptime metric.

But of course we can agree to disagree – that's why there are so many metrics out there. We'll be interested to see if this discussion generates further feedback from other subscribers.

Granularity of Sub-State Reporting

Ian then followed up with an additional question about tool states, asking: "One thing we have debated here is the proper use of sub-states in an equipment model. For example, we all agree that "Standby" and "Standby No WIP" should be separated, but Mark proposed the use of several other sub-states, such as "Standby No Operator", or "Standby No Durables". I have seen this type of breakdown in real-time models really clog up the smooth execution of the state model because 1) it gets more confusing for the operators, and 2) it makes the transition model get more complicated which further reduces flexibility when engineers/technicians are just trying to do their job and quickly log the data. With multiple sub-states, accurate history is more and more dependent on real-time manual logging, and we all know that real-time manual logging is never as accurate as we'd like it to be. As I mentioned in my original feedback, I have had success with capturing the same kind of information as you get from sub states, with retroactive reporting that can see subsequent events and make a better determination of which type of Standby was occurring (i.e. no operator, no durable). Basically, you're looking at the deltas between certain events. Of course, there is always a factor of interpretation, so the accuracy will never be 100%, but I feel that since it frees up the real-time state model from unnecessary transition restrictions, this is a better approach. My manager disagrees because he doesn't think there is anything you can learn after the fact to make a better determination than what you can already know in real time. My question to you is: what is your experience/preference/thoughts on how to handle granular sub-states (or "time buckets" as I like to call them in a retroactive report)?"

FabTime Response: Regarding the sub-states, our feeling, like yours, is that it's

best to refrain from having operators do more logging than is necessary. When reporting Standby-WIP Waiting vs. Standby - No WIP, we make that distinction at the time, in FabTime, by looking at the WIP transactions. That is, when a lot is logged out of the tool, taking the WIP in tool down to zero, we transition the tool to standby. Then we check whether or not there is WIP in queue at the tool, according to whether or not WIP has left the previous operation, and transition to Standby-WIP Waiting vs. Standby - No WIP accordingly. This isn't perfect, either, because we generally rely on the move out transaction from the prior step, to say that there is WIP at the tool. Sometimes, the WIP is in transition. But we find it's still a useful breakdown of time to have, and one that doesn't rely on additional operator logging. Any time you see a significant portion of Standby-WIP Waiting, especially when accompanied by high per-visit cycle times, this suggests that some operational improvement could be made going forward (changing staffing levels, changing how lots are transferred between steps, etc.). Standby-No WIP, of course, usually requires no particular action (with the exception of attempts to better balance WIP between tools, but that's a larger-scale project).

We have done other work occasionally with retrospective changes to equipment states, mainly for fabs in which minimal logging is done. For example, we have created virtual begin run transactions, based on planned process times, for fabs that only log end run/move out transactions. This allows the fabs to get some idea of queue time vs. process time, though this is of course less accurate than information obtained by logging both begin run and end run transactions.

Our general view on this is that these types of transaction manipulations are most worthwhile if done in a seamless, automated fashion (and one that people agree with), as near to real-time as possible. We're not big fans of making manual changes to the data after that fact - we would tend to agree with your manager that this is less useful. But we try to offer as much flexibility as we can to our software customers, and leave them in the driver's seat on these types of decisions. In truth, over the years we've learned that some fabs just don't have the resources to get more granular logging from their operators, and we'll do whatever we can to help fill those gaps, in a systematic manner.

Lateness Reporting for In-Process Lots

Eliot Parkinson from International Rectifier (a newsletter subscriber and FabTime customer from the UK) recently opened up a discussion with us regarding the ongoing identification of late lots. The question is, what calculation do people use to determine if a lot mid-way through processing is early or late. What we have done in FabTime in the past is compute the elapsed time between the lot's start date and its due date, and divide this elapsed time by the sum of all planned cycle times for completed steps and future steps. The result of this division is the required ratio of actual to planned cycle times that the lot must meet in order to ship on time. Then we use this required ratio to compute required out times for each future operation. This required out time is compared to the expected time that the lot will finish each operation (based on the sum of planned cycle times for future steps, plus any expected factory nonscheduled time). At each point in time, a delta exists between the required time and the expected time, and the lot is early or late at that point by the difference between these two values. The lot is early or late overall depending on the difference between the expected shipment date and the due date.

Eliot, however, suggested that this method wasn't quite serving his company's needs. He wrote: "We have a very diverse product

mix here and we run a high % of Engineering and development material. Our products have very diverse processing times, test times and some even go off site for part processing. There has been a need to be able to see exactly the position with any material late to its schedule during its life in the process, as our priority lots change often. This report would then enable us to see detail on lots that are behind to the CT plan, so we could then go and reprioritize or set rules through FabTime dispatch so we can get them back on plan. Using the normal method of extrapolation telling us "you are late now and if you extrapolate you will be late by xxx in 10 days" is OK but most of the time not accurate as there are so many changes/testime etc that affect the lots."

The bottom line is that we found our existing method of computing required out times computationally inefficient to perform for every lot every few minutes, because it depends on the full past history of each lot. And upon reflection, we concluded that it was non-intuitive for two lots with the same flow, started on the same date, and with the same due dates, to have different required out times for the remainder of their flows (which can happen, e.g. if one of the two lots revisits steps for rework, that affects the required cycle time ratio discussed above). So we decided to simplify things a bit, and just calculate, after each extract from the MES, each lot's future planned cycle time. That is, sum up planned cycle time (or some multiple of theoretical cycle time, for lots planned with specific x-factors) from the end of the flow backwards to our current step (plus any nonscheduled time), subtract that from the due-date, and the result is the required out time for the current step. This way, what the fab has done in the past to this lot is irrelevant. Only the remaining steps in the flow and the due date are relevant. And this calculation can be completed very quickly for all lots, every time we receive new data from the MES.

What do you all think? Do you have a scheduled time that you expect to complete each operation when the lot begins, or do you just work backwards, on the fly, off of the due date and the remaining process time? Does time spent off the normal flow (e.g. in rework) affect the required out times for future operations? Or do you have some other method for identifying how early or late each lot is, at any given moment in time.

We welcome your feedback on any or all of these four topics. Send your messages to Jennifer.Robinson@FabTime.com, and we will include them (attributed or not, as you prefer) in the next issue.

Subscriber List

Total number of subscribers: 2793, from 468 companies and universities.

Top 21 subscribing companies:

- Maxim Integrated Products, Inc. (217)
- Intel Corporation (150)
- Chartered Semiconductor Mfg (85)
- Micron Technology, Inc. (83)
- X-FAB Inc. (73)
- Western Digital Corporation (69)
- Texas Instruments (62)
- Freescale Semiconductor (59)
- Analog Devices (57)
- ON Semiconductor (56)
- TECH Semiconductor Singapore (56)
- International Rectifier (55)
- NEC Electronics (53)
- IBM (47)
- STMicroelectronics (46)
- Infineon Technologies (44)
- NXP Semiconductors (43)
- Cypress Semiconductor (39)
- Seagate Technology (36)
- BAE Systems (30)
- National Semiconductor (30)

Top 3 subscribing universities:

- Virginia Tech (11)
- Arizona State University (8)
- Ben Gurion Univ. of the Negev (8)

New companies and universities this month:

- DayStar Technologies
- Nimble Consulting Services

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 and receive the current issue of the newsletter each month. Past issues of the newsletter are currently only available to customers of FabTime's web-based digital dashboard software or cycle time management course.

To subscribe to the newsletter, send email to newsletter@FabTime.com, or use the form at www.FabTime.com/newsletter.htm. 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® Dispatching Module

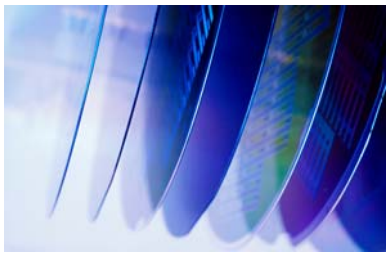

Dispatch Configuration and Support

We offer our dispatching and planning modules together for a single, fixed monthly fee (on top of your regular FabTime subscription). This includes:

- Dispatch rule and factor configuration via user-friendly web-based interface.
- Training.
- Dispatch list feed to the MES (if applicable).
- Support and upgrades.

Dispatch Factors

- Batch code at the current tool.
- Lot priority.
- Downstream tool priority.
- Current tool FIFO.
- Current tool idle time.
- Downstream batch efficiency.
- Critical ratio.
- Earliest-due-date.
- Current step processing time.
- Remaining processing time.
- Current step qualified tool count
- WIP level or staging time at downstream tools.
- Up to 20 other site-specific factors.

Interested?

Contact FabTime for technical details.

FabTime Inc.
Phone: +1 (408) 549-9932
Fax: +1 (408) 549-9941
Email: Sales@FabTime.com
Web: www.FabTime.com

Do your operators make the best possible dispatching decisions?

- Do you struggle to balance lot priorities and due dates with tool utilization and moves goals?
- Do your critical bottleneck tools ever starve?
- Do you use standard dispatch rules, but feel that your fab's situation is more complex, requiring custom blended rules?How well does your fab execute your dispatching strategy?

FabTime's dispatching module is an add-on to our **web-based digital dashboard software**. At any point, for any tool in your fab, FabTime will show you the list of all lots qualified to run on that tool. This list will be ordered by the dispatching logic that your site has selected for that tool. This logic can use standard dispatch rules such as Priority-FIFO and Critical Ratio. However, you can also create custom dispatching logic using any combination of dispatch factors (shown to the left).

You can display dispatch lists in FabTime, and/or export them back to your MES. FabTime also includes a dispatch reservation system to hold downstream tools when a lot is started on an upstream tool, as well as dispatch performance reporting.

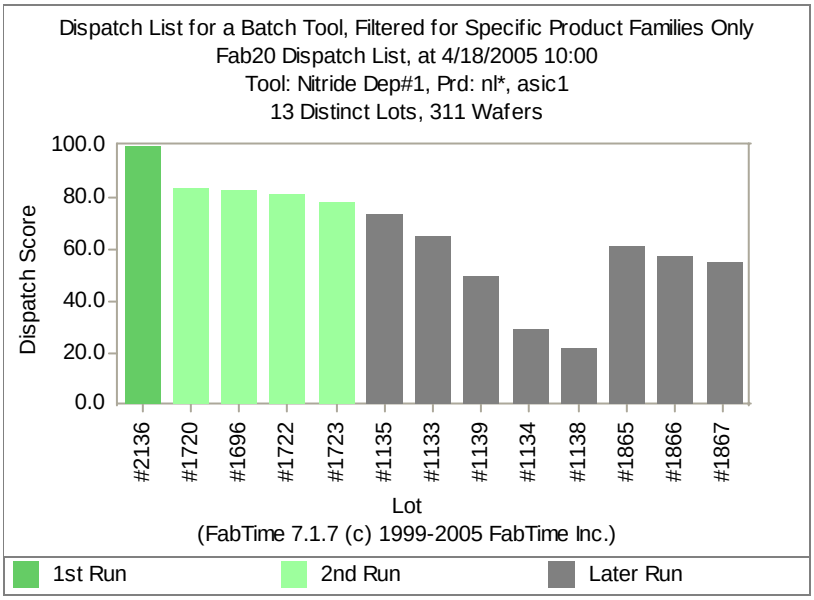

FabTime Dispatching Module Benefits

- Ensure that wafers needed by management are in fact the wafers that are run, while requiring less manual intervention on the part of management.
- Improve delivery to schedule, and the display of performance to schedule.
- Document the dispatching logic used by the best operators and make this available to all shifts.