FabTime Cycle Time Management Newsletter

Volume 9, No. 4 May 2008

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 WIP goals, WIP State Trend charts, and support for site-specific WIP attributes and filters.

Editor: Jennifer Robinson

Table of Contents

- Welcome
- Community News/Announcements
- FabTime User Tip of the Month – Alert based on Home Page Chart Data Rows
- Subscriber Discussion Forum
- **Main Topic – Dynamic X-Factor and Shipped Lot X-Factor**
- Current Subscribers

Welcome

Welcome to Volume 9, Number 4 of the FabTime Cycle Time Management Newsletter! Things remain fairly busy with FabTime, as we work on installation and training for various new FabTime customers (we're currently at 17 sites using our web-based dashboard software). We've also been adding lots of great new functionality to the product, most of it in direct response to suggestions from our customers. If you would like to see an updated demo, just email me at Jennifer.Robinson@FabTime.com, and I'll be happy to set something up with you.

In this issue we have a brief summary of upcoming industry conferences in our community announcements section. Our FabTime software user tip of the month is about using the home page chart alert functionality. We have one subscriber discussion questions, about the transition from paper to electronic travelers, for which we could use your input. Our main article this month is about the comparison between dynamic x-factor (a point estimate measured as total WIP divided by WIP running on tools) and shipped lot cycle time x-factor. We show that although in the long run, DXF can be used to predict x-factor, various issues sometimes make it difficult to draw exact comparisons between this week's DXF and some future week's shipped lot x-factor value. We hope that you find this of interest, and 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

Upcoming Industry Conferences

Here is a brief listing of upcoming industry conferences that readers of this newsletter may find of interest:

- Semicon West will take place July 15-17 at the Moscone Center in San Francisco. See www.semiconwest.org. The Fab Owners Association meeting (members only) will take place at Semicon on July 17th. See www.waferfabs.org.
- The International Sematech Manufacturing Initiative 5th Symposium on Manufacturing Effectiveness will take place October 22-23 in Austin, Texas. Details are available at http://ismi.-sematech.org/ismisymposium/index.htm.
- The International Symposium on Semiconductor Manufacturing (ISSM 2008) will take place in Tokyo, Japan October 27-29. Details are available at https://www.semiconportal.com/issm/.
- The Winter Simulation Conference will take place in Miami, FL December 7-10. See http://www.wintersim.org/. The Modeling and Analysis of Semiconductor Manufacturing (MASM 2008) will take place as part of the same conference. See http://www.wintersim.org/MASM.htm.

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

FabTime User Tip of the Month

Alert based on Home Page Chart Data Rows

There are quite a few alerts available in FabTime. You can have alerts based on areas, operations, toolgroups, lots, lot priority, and tool states. However, it's been our experience that sometimes people want to have alerts based on other, very specific pieces of information. To support this, we added the Home Page Chart alert function. This alert lets you set a trigger in FabTime so that you'll be notified if the number of rows on a particular chart reaches a specified value, or changes from one value to another. This alert is most useful for point in time charts, such as a WIP Lot List chart, where the chart displays some number of items.

For example, suppose you wish to be alerted whenever the WIP in queue for a bottleneck single-path tool falls below 5 lots. The home page chart to create is a WIP lot list filtered for the bottleneck tool, with queue choice = "in queue". Here is an example of how the alert works:

9:55am: The WIP lot list displays 6 data rows. (There are six lots in queue)

10:00am: A new "Number of Data Rows" <= 5 alert is defined for the WIP lot list chart, with a sleep-after time of 2 hours. (You'll find the row for "New Home Page Chart Alert" at the bottom of the Alert screen - your home page charts are listed in the drop-down by tab, and then chart name).

10:05am. When the alert is checked, the number of data rows (6) does not meet the alert condition, so the alert does not trigger.

10:07am: A lot in queue begins processing on the bottleneck tool.

10:10am: When the alert is checked, the chart displays only 5 data rows (5 lots in queue), so the alert triggers, and sends you an email. The alert goes to sleep for 2 hours.

12:10pm: The alert is checked for the first time after being asleep for 2 hours. If the number of data rows on the WIP lot list chart is <= 5, the alert triggers.

etc.

The "Change in Data Rows" alert can be used to notify you of things like a change in the number of down tools in your area, an increase in the number of wafers scrapped so far today, an increase or decrease in the number of hot lots in your area, etc. You can find detailed examples in the Help for the Alert page.

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

Subscriber Discussion Forum

Paper vs. Electronic Travelers

An **anonymous subscriber** asked: "I'm looking for a way to eliminate our clean room paper travelers, and wondering if you could share some of your thoughts and ideas, knowing what we do here...Our hang-up is being so used to having people see from the traveler where to deliver the box to...and then to see from the traveler that the lot is ready to load, etc..... To eliminate the paper traveler, we are thinking to go with hand held devices, or something to allow people to "know" where in the process flow the wafer is, versus physically where the wafer is... We wondered if you had any thoughts on this, from all of your visits to different fabs."

FabTime Response: We're afraid that we don't have any experience with this question, so we are opening it up to the subscriber community. Anyone made this transition? Any suggestions for things to do, or things to watch out for, in making the change? We, and the subscriber asking the question, would much appreciate your input.

FabTime welcomes the opportunity to publish subscriber discussion questions and responses.

Dynamic X-Factor and Shipped Lot X-Factor

Introduction

One of our subscribers wrote in with a question this month that we thought was significant enough to warrant a full newsletter article. This subscriber explained that his fab was tracking dynamic x-factor, which is a metric reported at each observation as Total WIP in the Fab / WIP in Process (or, WIP running on tools). FabTime described this metric back in Issue 4.08, after learning about it through a paper written by Johnishi et. al. (full reference below). We went into some additional detail about it in Issue 5.03. We showed, using Little's Law, that when averaged across a number of observations, dynamic x-factor should be an early indicator of a fab's shipped lot x-factor.

We like dynamic x-factor as a metric because it is intuitive and easy to calculate (much easier than calculating theoretical cycle time), and gives an early warning of when future cycle times are likely to increase. We know of a number of fabs that are using dynamic x-factor in a control chart-like fashion, by which they take note if it drifts upward, outside of normal fluctuations. These fabs also use dynamic x-factor to highlight short-term, periodic effects in the fab, such as shift change.

But the fab that wrote to us last week reported that they were averaging their hourly dynamic x-factor values over a two week period, and then comparing that to the weighted average x-factor for shipped lots. And they found that the reported day-to-day dynamic x-factor was significantly higher than the shipped lot cycle time x-factor. They didn't think that it was a short-term effect, because they said that both values (DXF and x-factor) had been stable for quite some time. They asked us if we had any insight into what might be causing this difference. In this issue, we will discuss, in general terms, factors that could cause the dynamic x-factor prediction to be significantly different from a fab's reported shipped lot cycle time x-factor.

The Relationship between Dynamic X-Factor and Cycle Time X-Factor

To begin, let us revisit the mathematical relationship between dynamic x-factor and shipped lot cycle time x-factor. In this section, we repeat the mathematical justification from Issue 4.8. We show mathematically what already makes sense intuitively: that in the long term, dynamic x-factor will work out to be equal to the traditional cycle time x-factor. Dynamic x-factor says: of the WIP we have in the line, how much are we working on at any given point, and how much do we have sitting?

Say the dynamic x-factor works out to be four, for example. This means that for every lot in process in the fab, there are three lots in queue (or in transit to the queue). Every time a lot gets processed, it first has to wait in queue (on average) for those three other lots to be processed. So, we would expect that its average cycle time by operation would be about four times the average process time (consisting of three intervals of queue time while other lots are processed, plus one interval of actual process time).

More formally (though note that this is a justification, not a formal proof), let

W_{tot} = Total cycle time (average)

TPT_{tot} = Total process time (average)

L_{tot} = Total WIP in the fab

λ_{Start} = Arrival rate into the fab

TPT_{avg} = Average process time for a single step

N_{Steps} = Number of steps

L_{proc} = Total WIP currently being processed

λ_{Proc} = Total arrival rate to all individual steps (summed across all steps)

We want to estimate x-factor = cycle time / theoretical process time

$$\text{X-Factor} = W_{tot} / TPT_{tot} \quad (1)$$

We know that, on average, total cycle time = total WIP / total system arrival rate (from Little's Law) and so we have:

$$Wtot = L_{tot}/\lambda_{Start} \quad (2)$$

Now, the total process time is the sum of the process times for the individual steps, and we have:

$$TPT_{tot} = TPT_{avg} * N_{Steps} \quad (3)$$

The average process time for a step is equal to the wafers in process at the step divided by the arrival rate to the step. This is also from Little's Law (it applies to the whole system, or the queue, or the process time). And we have:

$$TPT_{avg} = L_{proc} / \lambda_{Proc} \quad (4)$$

So, substituting equation (2) and equation (3) into equation (1), our estimate for cycle time x-factor is:

$$\text{X-Factor} = W_{tot} / TPT_{tot} = (L_{tot}/\lambda_{Start}) / (TPT_{avg} * N_{Steps})$$

Substituting in equation (4) for TPTavg, we have:

$$= (L_{tot}/\lambda_{Start}) / ((L_{proc} / \lambda_{Proc}) * N_{Steps})$$

Rearranging terms (inverting and bringing up the denominator), we get:

$$= (L_{tot} / L_{proc}) * (\lambda_{Proc}) / (\lambda_{Start} * N_{Steps})$$

Now, the total arrival rate to the process steps is equal to the arrival rate into the system multiplied by the number of steps, and so we have:

$$\lambda_{Proc} = \lambda_{Start} * N_{Steps}$$

This means that the right-hand factor above goes to 1, leaving

$$\text{X-Factor} = (L_{tot} / L_{proc}),$$

which is total WIP divided by WIP in process. This is the definition of dynamic x-factor.

Of course in practice we need to take quite a number of samples of DXF to get an average value, because we cannot be sure of the distribution of DXF (as was pointed out by Alexander Schoemig in Issue 5.03). And we must take care not to allow sampling issues in our DXF values (for instance, if you were to only record DXF at shift change).

Issues Potentially Affecting DXF Values

The above mathematical justification, based on Little's Law, suggests that in the long run, if averaged over enough point estimates, DXF should be a predictor for x-factor. In practice, however, a number of systematic issues can cause a short-term divergence between DXF and x-factor values.

Time Lag: DXF is a predictor of what the future cycle time is going to look like. This mean that the DXF value for this week may not resemble the average x-factor for the lots shipped this week, particularly for fabs with rapidly changing cycle times. In general, today's DXF reflects the approximate x-factor ratio that the lots currently in the line will have when they ship. Assuming that, on average, the lots currently in process are half-finished (which should be true in steady state, but will not necessarily be true during ramp-up or ramp-down period), today's DXF should (when averaged across various estimates) predict the shipped lot x-factor in about 1/2 a cycle of average cycle time. That is, if the average cycle time is 60 days, then this week's DXF should give an estimate about 30 days ahead.

However, this doesn't mean that you can simply take this week's shipped lot cycle time x-factor, look back to a time approximately half the average cycle time ago, and expect that DXF to exactly reflect the current shipped lot cycle time value. The reason is that most fabs consist of a mix of longer and shorter cycle time products (either because of number of layers or because of different lot priorities). If this week you happen to have shipped a

large number of hot lots, for example, the resulting x-factor will be lower than that predicted by the overall DXF. In the long term, these things will average out, but you will certainly see short-term fluctuations based on your product mix. Looking at DXF by either product family or priority class can help to smooth out this type of mix-related fluctuation.

An example, from a simulation model (NOT actual data) is shown below and at the top of the next page. For this fab, the historical dynamic x-factor 2 months ago averaged approximately 3. For the two most recent weeks of shipped lot cycle times, the average x-factor was 2.8 one week and 4.4 the next. Filtering to only display the main product runner, however (not shown), smoothed things out considerably, and made the charts line up much more closely.

WIP that Never Ships: By default, your dynamic x-factor estimate includes all of the WIP in the line in the numerator. This can include lots on extended hold, lots with yield problems, etc. Sometimes these lots never end up being shipped - they sit around in extended hold for years, or are scrapped. This means that they drive up the DXF estimate, but are not included at all in the shipped lot x-factor (by which you take, for each lot, actual cycle time divided by theoretical cycle time). This will tend to bias the DXF value high relative to the shipped lot x-factor. One way to cope with this in reporting DXF is to exclude lots on extended hold (e.g. in storage) from the hourly DXF estimates. However, it is difficult to guard against other types of scrapped lots, since they cannot be systematically identified until they are scrapped. Note that you do generally want to include regular holds in your DXF calculation (in the numerator), since this hold time is part of the cycle time for the lots.

Inconsistent Estimates of Theoretical Cycle Time: One of the primary benefits of using DXF instead of regular x-factor is that you don't need to know each lot's theoretical cycle time. You just need to be able to tell whether each lot is in process or not. However, if you're going to compare

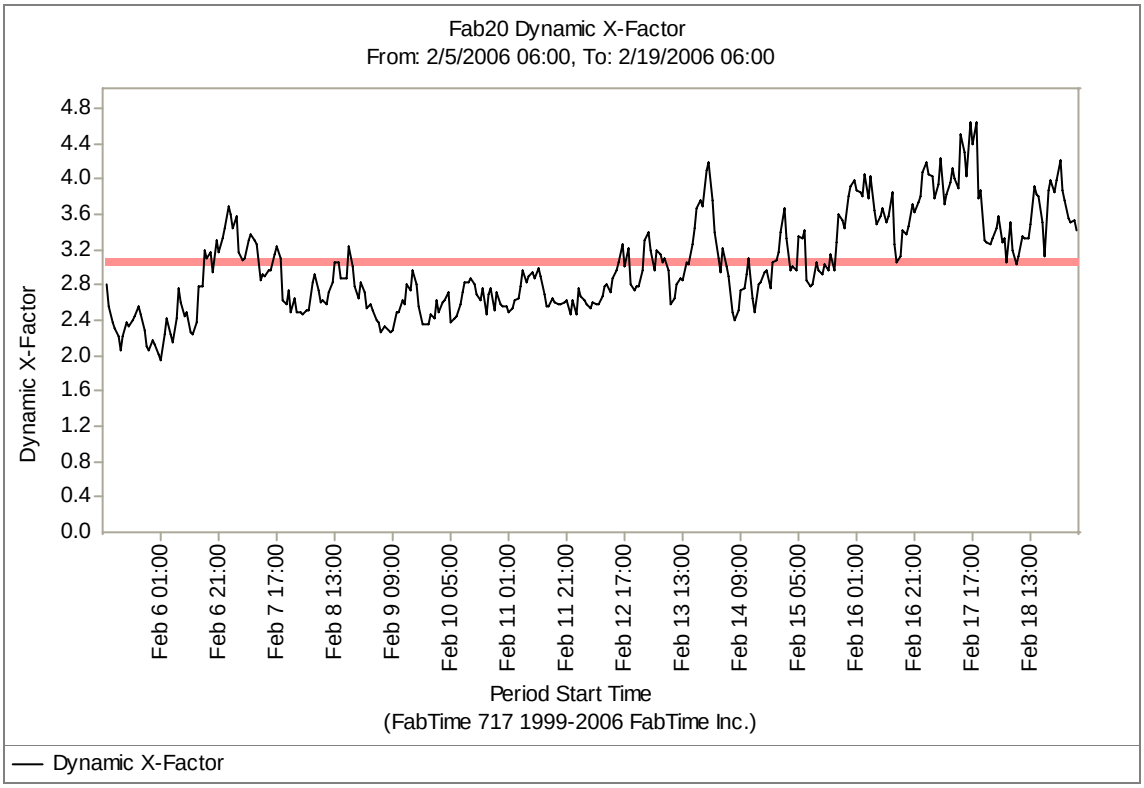

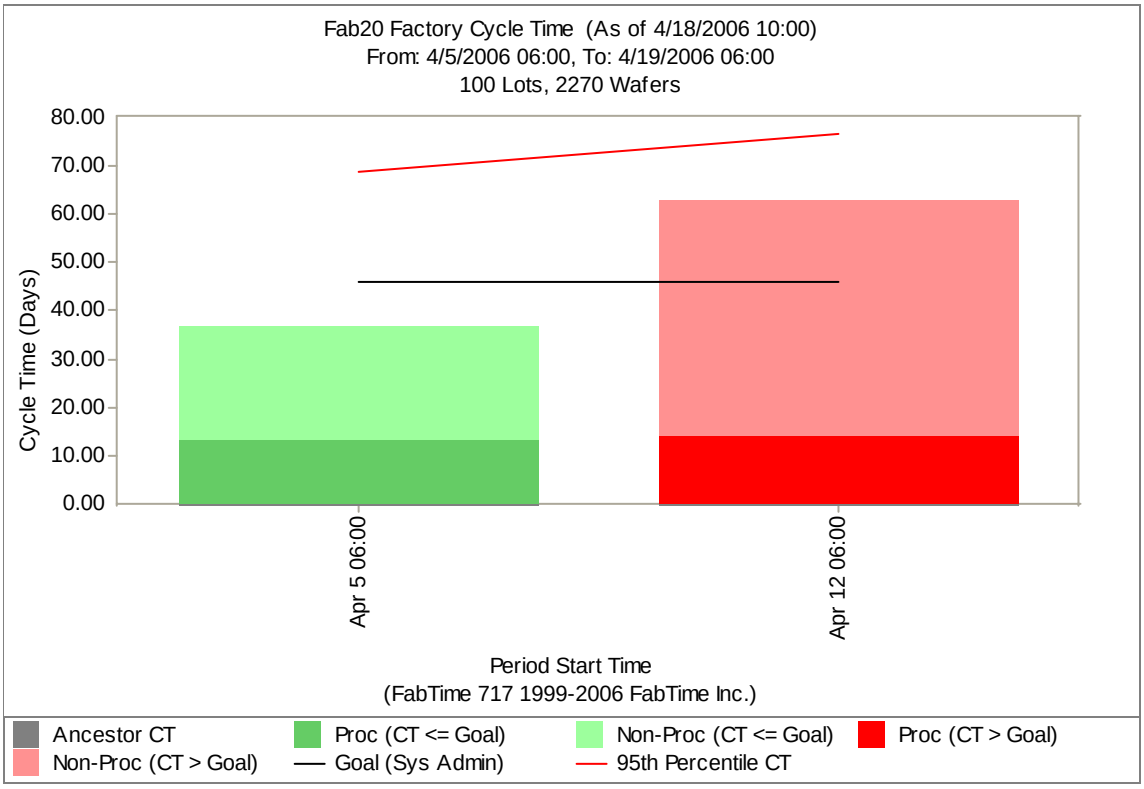

DXF and x-factor, you need to make sure that the theoretical cycle time used was calculated in a manner consistent with that used in DXF. You can't, for example, compare to a theoretical CT value that has a bunch of travel time built into it, since DXF counts that travel time as queue time. Similarly, if you got your theoretical value by running a hand-carry through the fab, the number is probably inflated by about 20% relative to the true, process-time-only theoretical cycle time. In this case, your DXF will implicitly be based on a lower number than your theoretical, and will look high relative to the overall x-factor.

On the other hand, many fabs use a theoretical cycle time value based on some optimal processing rate of the tool. DXF is based on the time lots actually spend in process. This means that if your actual process times tend to be longer than your theoretical cycle time estimates, your DXF estimate may be low relative to what you report for shipped lot x-factor. For example, suppose that your planned theoretical cycle time is 10 days, and your shipped lot cycle time is 40 days. Your shipped lot x-factor will be 4. However, now suppose that your actual process time is 12 days. Your DXF estimate will implicitly be based on that 12 days. All else being equal, you'll get an estimate of 40/12 = 3.33X using DXF and actual process times, instead of an estimate of 4X using the theoretical cycle time.

Logging Issues: A related point to the above is that logging issues can cause the dynamic x-factor values to be inaccurate. We've seen tool groups, for example, where the operators routinely wait to log the BeginRun until immediately prior to the EndRun. This means that your DXF estimate won't correctly show the amount of WIP process, and will tend look artificially high (because the denominator will be artificially low). Alternatively, you might have issues with WIP that has finished processing sitting in the tool, waiting for the move out transaction to be recorded. This WIP looks like WIP in Process to the DXF estimate, and tends to artificially lower the value. In some cases, there could be other systematic logging issues that are affecting DXF, such as tools

where BeginRun and EndRun aren't logged at all.

Summary of DXF vs. X-Factor Issues:

Remember, the equation in question is:

DXF =? X-Factor

or

Total WIP / WIP in Tools =? Shipped Lot CT / Theoretical CT.

If DXF is systematically higher than shipped lot x-factor, one or more of these things may be happening:

1. The estimate being used for theoretical cycle time is high, relative to the time that lots are logged in to tools in the MES. This could be because of lots not having their BeginRun recorded properly, or because the theoretical cycle time value is based on an inflated number.

2. The average recorded shipped lot cycle time is too low, because of lots that were part of the WIP, but never shipped (scrapped lots, or lots on permanent hold), that had higher than average cycle time.

3. WIP in Tools is artificially low because of logging issues.

4. Total WIP is high, relative to the WIP that's actually shipping from the line, due to product mix changes.

If DXF is systematically lower than shipped lot x-factor, then one or more of these things is probably happening:

1. The estimate being used for theoretical cycle time is low, relative to the time that lots are logged in to tools in the MES. This could be because lots spend extra time "in process" due to not being moved out promptly, or because the theoretical cycle time is based on unrealistically low cycle time performance.

2. The average recorded shipped lot cycle time is too high on a short-term basis, because of product mix changes, relative to the DXF being compared.

3. WIP in Tools is artificially high because of logging issues, and lots not getting moved out of tools.

4. Total WIP is low, relative to the WIP that's actually shipping from the line. This shouldn't happen in the long-term, but could happen on a short-term basis, if the fab is pulling WIP from the end of the line to meet shipment targets.

In addition to the above, time lag, rework, and product mix change issues could be playing a part.

Conclusions

Dynamic x-factor is a predictor for shipped lot cycle time x-factor. If your DXF value starts drifting upward, and you do not do something to correct the situation, you will see increases in shipped lot cycle time at some future date. This must be true, because a rising DXF means that you either have more WIP in the line (which translates to longer cycle times down the road), or it means that you aren't doing as good a job of keeping WIP running on tools (which also translates to longer cycle times). However, on a short-term basis, you can't necessarily look at the graph of your fab-wide DXF and expect it to exactly track with future shipped lot cycle time values on a particular date. The two primary reasons for this are 1) that product mix changes combined with the time lag between DXF and x-factor make short-term comparisons difficult; and 2) systematic issues in how you measure and report theoretical cycle time, and how you log transactions in your MES, can lead to differences in the reported values for the two metrics. You can manage the first issue, somewhat, by looking at DXF on a family and/or priority class basis. You can also investigate your logging issues, and your theoretical cycle time estimates, and make changes to bring the DXF and x-factor values closer together, if necessary. However, even if you do not see a one-to-one absolute correspondence between

these two values, we think that reporting DXF on a short-term basis is still useful. DXF gives you an early warning of worsening cycle time, and it gives you valuable information about short-term, periodic behavior, such as shift change effects. It is also very easy to measure, and to explain to people. Therefore, we still highly recommend it as a fab-wide metric, and a barometer of the fab's overall health.

Closing Questions for FabTime Subscribers

Do you report DXF in your fab? Have you compared it to actual x-factor? Have you found it to be biased high, or low, relative to shipped lot cycle time values for your fab?

Further Reading

S. Johnishi, K. Ozawa and N. Satoh, "Dynamic X-Factor Application for Optimizing Lot Control for Agile Manufacturing," *Proceedings of the 2002 International Symposium on Semiconductor Manufacturing (ISSM2002)*, Tokyo, Japan, 2002.

J. Robinson and F. Chance, "Dynamic X-Factor", *FabTime Newsletter*, Vol. 4, No. 8, 2003.

J. Robinson and F. Chance, "Dynamic X-Factor Revisited", *FabTime Newsletter*, Vol. 5, No. 3, 2004.

If any subscriber to the newsletter would like a copy of issue 4.8 or 5.3 between now and when Issue 9.05 is published, simply email your request to newsletter@-FabTime.com. We regret that we cannot distribute the Johnishi paper (where we first learned of DXF), because of copyright issues.

Subscriber List

Total number of subscribers: 2811, from 473 companies and universities. 21 consultants.

Top 20 subscribing companies:

- Maxim Integrated Products, Inc. (230)
- Intel Corporation (157)
- Micron Technology, Inc. (83)
- Analog Devices (66)
- Infineon Technologies (63)
- X-FAB Inc. (63)
- Freescale Semiconductor (62)
- NEC Electronics (61)
- Texas Instruments (59)
- STMicroelectronics (56)
- International Rectifier (55)
- ATMEL (54)
- Cypress Semiconductor (54)
- TECH Semiconductor Singapore (54)
- Chartered Semiconductor Mfg (53)
- ON Semiconductor (52)
- NXP Semiconductors (47)
- IBM (44)
- Spansion (36)
- Seagate Technology (32)

Top 3 subscribing universities:

- Virginia Tech (10)
- Ben Gurion Univ. of the Negev (8)
- Nanyang Technological University (7)

New companies and universities this month:

- SUMCO

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® Software Capacity Planning Module

CP Configuration

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

- Identification of the source of any additional data needed for the planning module.
- Automation of the process of importing the additional data into FabTime.
- Validation against client data.

Interested?

Contact FabTime for more information, or for a quote.

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

Do you need to answer questions like:

- Given a target product mix, do we need any new tools?
- Given the tools that we have, and the products that we are running, how many wafers can we expect to produce?
- Given our existing set of products and tools, what happens if the product mix changes? Where can we expect bottlenecks?

Are you tired of maintaining a standalone capacity planning spreadsheet?

FabTime's capacity planning module leverages the data already stored in the FabTime digital dashboard software, to make it easier to build capacity planning scenarios. The only required manual inputs are:

- Weekly ships per product.
- Product line yield percentages.

FabTime uses route information from the fab MES and calculates UPH data (tool speed) based on actual performance. FabTime also uses tool uptime performance to estimate availability (though this can be overridden). These inputs are used to generate predicted utilization percentages for each capacity type. Detailed intermediate calculations (UPH, tool productive time, tool rework percentage, etc.) are also available (an example for one tool is shown below). All outputs can be easily exported to Excel.

Capacity Planning Module Benefits

- Eliminate the need to maintain offline capacity planning models.
- Automatically update capacity planning data to reflect new conditions (process flows, tool uptime characteristics).
- Quickly run scenarios to anticipate (and avoid) bottlenecks caused by product mix changes.

C Type	Output	Value	Notes
1XStep	Rework Moves/Week	21	2004-09-06 10:00:00 to 2004-11-15 10:00:00
1XStep	Total Moves/Week	12310	2004-09-06 10:00:00 to 2004-11-15 10:00:00
1XStep	Rework Ratio	0	Rework Ratio = Rework Moves / Total Moves.
1XStep	Productive%	61	2004-09-06 10:00:00 to 2004-11-15 10:00:00
1XStep	Availability%	76.26	Availability = Productive% + Standby%.
1XStep	Historic Utilization%	79.99	Utilization (Mfg efficiency) = Productive% / Availability%.
1XStep	Productive(Rework)%	0.1	Productive(Rework)=Productive% * ReworkRatio.
1XStep	Net Availability%	76.15	Net availability% = Availability% - Productive(Rework)%.
1XStep	Arrivals (Units/Hour)	79.36	Based on total plan WGR=2025
1XStep	Tool Quantity	8	1XStep#1 ... 1XStep#8
1XStep	UPH	15.02	UPH = (TotalMoves/ToolQty) / (Productive% * 168)
1XStep	Required Hours/Day	126.84	Required hours = 24 * HourlyArrivalRate / UPH
1XStep	Predicted Utilization%	86.75	Util = 100 * ReqdHours / (24 * NetAvail * ToolQty / 100)
1XStep	Max WGR	2334.22	MaxWGR = PlanWGR / PredictedUtilization
1XStep	Historic WGR	2457.8	(Non Rework Moves) / (OperationCount / ProductCount).