FabTime Cycle Time Management Newsletter

Volume 12, No. 1 February 2011

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 "Excel (all rows)" link to support export of all rows of data, and increased speed for building large data tables.

Editor: Jennifer Robinson

Contributors: Shmulik Perez (Micron); Bob Kotcher (Philips Lumileds)

Table of Contents

Welcome

Welcome to Volume 12, Number 1 of the FabTime Cycle Time Management Newsletter! We hope that those of you in northern locations are weathering the season reasonably well, and that spring arrives soon. In this issue we have three announcements, one about a survey from WWK, another with a call for papers, and the third about staying in touch with FabTime via my LinkedIn profile. Our software tip of the month is about using the new lot line yield charts in FabTime. We only have one subscriber discussion question, but it is quite detailed (about fab management in a multi-constraint environment).

In our main article this month, we address the difference between confidence intervals and prediction intervals. Both can be applied to simulated or actual recorded data, anything where you have repeated, variable observations (cycle times, WIP, etc.). Confidence intervals are used to estimate an underlying value that can't be directly observed, like the "true" mean cycle time for a product line. Prediction intervals, instead, are used to establish a range in which it is likely that a future observation will occur, given a series of past observations. So, for example, you might use a prediction interval to predict the upper and lower bound of expected fab throughput next week. We hope that you'll find this discussion useful.

Thanks for reading – Jennifer

FabTime

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

Community News/Announcements

WWK Semiconductor Manufacturing Technology Survey

February 1, 2011 (Pleasanton, CA) – Wright Williams & Kelly, Inc. (WWK), a cost & productivity management software and consulting services company, announced today the start of its 2011 survey on equipment and process timing in the semiconductor industry. The survey results will be consolidated and provided to all participants free of charge. Participation in the survey is the only way to receive a full set of results. The survey form can be downloaded from the WWK web site at: http://www.wwk.com/2011survey.pdf .

Call for Papers: MASM Conference

The 7th International Conference on Modeling and Analysis of Semiconductor Manufacturing (MASM) will be held in Phoenix, Arizona, December 11-14, 2011.

The 2011 MASM aims to again be a forum for the exchange of ideas and best practices between researchers and practitioners from around the world involved in modeling and analysis of high-tech manufacturing systems. The MASM 2011 conference will be fully contained within the Winter Simulation Conference 2011 (WSC 2011), the leading conference in discrete event simulation. WSC 2011 features a comprehensive program ranging from introductory tutorials to state-of-the-art research and practice. WSC 2011 will take place in Phoenix, Arizona, USA. All attendees of the MASM conference will register for WSC 2011 at the same cost. All participants of WSC 2011 can attend MASM 2011 sessions.

While the conference is mostly focused on the current semiconductor industry state-of-the-art, neither presenters nor attendees need to be in the IC industry to participate. We are interested in any methodologies, research, and/or applications from other related industries such as TFT-LCD, flexible displays, bio-chip, solid state lighting (LED) and photovoltaic (PV) that might also share or want to share common and new practices.

The conference organizers are Stéphane Dauzère-Pérès, CMP, Ecole des Mines de Saint-Etienne, France and John W. Fowler, Arizona State University, USA. The deadline to submit papers is April 1st. See the full call for papers for more details: http://www.wintersim.org/wsc2011/MASM.htm

LinkedIn

We've been finding LinkedIn increasingly helpful in keeping up with colleagues as they change jobs (or just change company names, as seems to happen more and more frequently every day). If you use LinkedIn, and would like to use it to stay connected with FabTime through such transitions, you can connect with Jennifer here: http://www.linkedin.com/in/jenniferrobinsonfabtime.

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

FabTime User Tip of the Month

Measure Lot Line Yields

As has been discussed in the FabTime newsletter in the past (Issue 9.06: Definitions for Short-Tem Line Yield Metrics), short-term calculations for line yield can be a bit tricky, particularly in the presence of split lots. In our recently added Lot Line Yield charts, we have attempted to make a complex calculation as simple as possible. Here's how the calculations work:

For each time period, we include any lots that have either a ship or unit scrap transaction for themselves or any descendent lots during that time period. We also include any lots that were completely closed during the time period (e.g. fully scrapped), with no sub-units remaining open. [Note that a lot can appear in the calculations for more than one time period, if there are child lots that ship during different time periods.]

Then, for each started lot, lot line yield = 100.0 * ShippedUnits / (StartedUnits - OpenUnitsAsOfEndTime).

Here ShippedUnits is the sum of all shipped units for StartedAsLot or any descendent, prior to EndTime (even if they shipped before the StartTime of the chart period).

OpenUnitsAsOfEndtime is the sum of all remaining open units for StartedAsLot or any descendent, as-of EndTime.

So, for example, if during this time period we ship 15 wafers of a 25 wafer lot, and have 5 wafers still open (as one or more child lots, with 5 other wafers scrapped in the past), and no other wafers that shipped earlier, then our lot line yield is 100*(15/(25-5)) = 100*(15/20) = 75%. If, however, we ship 15 wafers from a lot during this time period, but we previously shipped 5 wafers from that lot, and we have 5 wafers still open, then our lot line yield is 100*(20/20) = 100%.

If a lot is completely scrapped, so that no open units remain, then the lot line yield is, of course, zero.

You can see individual lot line yield results on the Lot Line Yield List chart (as shown below). The data is also rolled up in the Lot Line Yield Trend and Pareto charts. (Calculating total ShippedUnits / (StartedUnits - OpenUnitsAsOfEndTime) across all included lots.

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

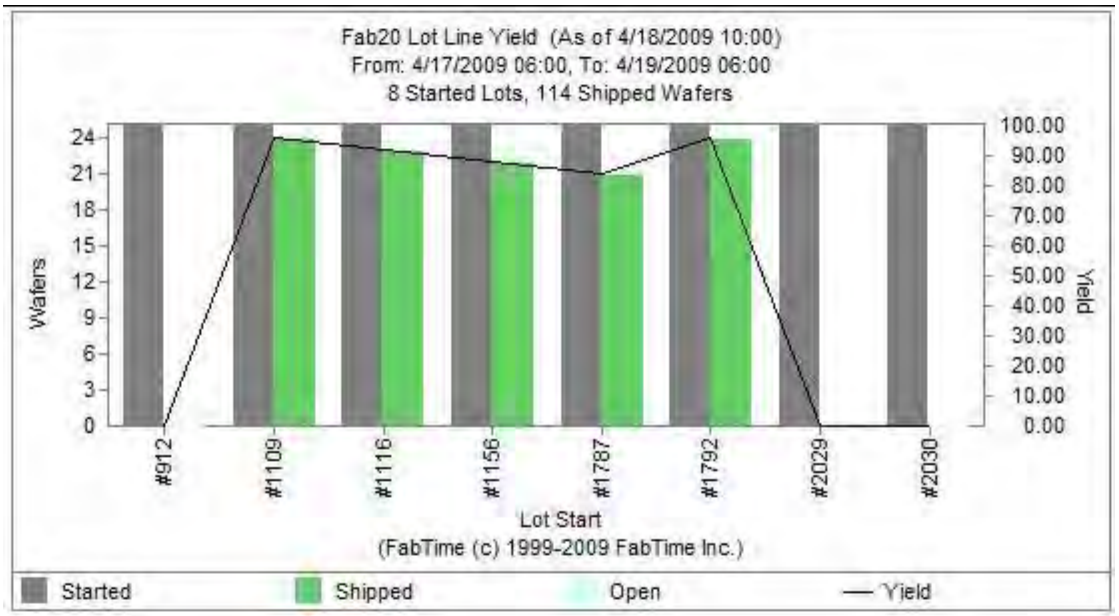

Subscriber Discussion Forum

Managing in a Multi-Constraint Environment

Shmulik Perez from Micron in Israel submitted several questions concerning multiple-constraint environments in wafer fabs.

"Background:

Over the last couple of years we have increased significantly our capacity. Obviously, in order to minimize capital spending, we have added, at each one of the ramp-up phases, additional tools, just for the top constraints. Of course, the decision to add tools was done only after exploiting all other opportunities to increase the tool's run rate by performing engineering improvement projects.

Now, towards the completion of this capacity increase effort, we are at the position that about 15 toolsets, are considered as constraints (with capacity within 2-3% of the stated capacity), and these toolsets are running about 30% of our Fab process steps. This is a situation that we are not used to … The immediate implications that we are facing are nonlinear output, resulting from variability at the constraints, and increased WIP and cycle time. As we all know from the Theory of Constraints - capacity lost at the constraints is unrecoverable …

My questions:

1. Was there any study performed on running a multi constraint environment?

2. What is the number of constraint toolsets which is considered to be 'reasonable' in a wafer fab?

3. What are the proposed methods to maintain the cadence of the line in such an environment?

4. What are the best practices to synchronize scheduled down time along the chain of constraints?

5. What other topics we should be aware of and take care of proactively?"

FabTime Response: We have not seen any research that addresses this topic to this level of detail. It has been our experience that most fabs run with a small number of near-constraints, such that the honor of being "the" bottleneck shifts frequently between them. Having several near-bottleneck tools is perhaps inevitable, given the high cost of semiconductor equipment. However, the more constraints you have, the more difficult the fab is to operate.

As Shmulik mentioned, we can look back to the production management classic, Eli Goldratt's *The Goal*. In Chapter 11, factory management guru Jonah says: "the closer you come to a balanced plant, the closer you are to bankruptcy". (See our full review of *The Goal* here: http://www.fabtime.com/goal.shtml). This is because variability leads to intermittent starvation of the constraints, causing lost time that cannot be recovered. Goldratt says that there should only be one bottleneck in a factory, and that everything else should be sublimated to that bottleneck. Of course that's easier said than done, particularly in an environment with reentrant flow and frequent product mix changes. (See FabTime Newsletter 1.04 for a discussion of applying Theory of Constraints to wafer fabs - email newsletter@FabTime.com to request a copy.)

As far as maintaining the cadence of the line in the presence of multiple constraints, we're starting to see more and more fabs making dispatching decisions upstream that favor lots headed to the bottleneck. How you do this when you have a tool that feeds several different near-bottlenecks, we can't say for certain.

In general, we recommend doing everything possible to keep WIP spread evenly through the fab, and keep the bottlenecks from starving. You might want to set target queues for all of the bottlenecks, and have some sort of flexible dispatch system that prioritizes lots going to whichever of your key tools is closest to going below those targets. These target queues would then help protect you in the presence of scheduled downtime events, because your downstream constraints would each have a buffer to work off in the event of decreased flow from an upstream maintenance event. (Of course it can be difficult to even say what is upstream vs. downstream, in the presence of reentrant flow.)

We would be interested to hear from our subscribers on these issues. What do you think is a "reasonable" number of constraint toolsets to have in a fab at one time? Three to five sounds reasonable to us, particularly if they are a bit spread out through the line. Fifteen sounds too high to manage (as evidenced by the issues at Micron that prompted this question in the first place). What do you all think? If you have any references on running fabs in the presence of multiple constraints, we would love to share them in the next newsletter issue. Perhaps we'll have enough feedback to expand this into next month's main article.

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

Prediction Intervals vs. Confidence Intervals

Introduction

In the semiconductor industry, as in most manufacturing industries, discrete event simulation is sometimes used to predict dynamic metrics such as WIP and cycle time. In putting this data into use, the question arises: given my assumptions, what range of outputs can I reasonably expect to see? And how do I calculate this? Is this the same as computing a confidence interval for the data? The answer to that last question is no. Calculating a range in which we expect the actual performance to fall (a prediction interval) is not the same as estimating a confidence interval. We recently had to take some time to clarify this in our own minds, and thought that a brief explanation of the difference between confidence intervals and prediction intervals would be of use to some of our newsletter subscribers.

Confidence Intervals

A confidence interval is what you use when you believe that your underlying data follows a distribution that has a true mean value, and you would like to make a statement, within a certain degree of confidence, about what that true mean value is. That degree of confidence is called a confidence level.

Confidence intervals are commonly used to estimate a mean value in cases where the distribution of the underlying population of data is not known, and we have repeated samples of data (from simulation experiments, or from other recorded data points). In these cases, we calculate (from the sample data):

s = estimated standard deviation of the data

xbar = estimated mean of the data

n = the number of samples (e.g. the number of simulation experiments)

C = the desired confidence level

The sample mean can usually be assumed to follow something called the t distribution (an introductory statistics textbook will outline what you need to check in order to make this assumption). The t distribution is used for testing when the sample size is small, and is generally bell-shaped like the normal distribution. The t distribution follows a mean, mu (approximated by xbar), and a standard deviation $s/\sqrt{n}$. It is described as having n-1 degrees of freedom. As the sample size n increases, the t distribution becomes closer to the normal distribution.

For a population with unknown mean mu and unknown standard deviation, a confidence interval for the mean, based on a random sample of size n, is

$$\text{xbar} \pm t^* \times (s/\sqrt{n})$$

where t* is the upper (1-C)/2 critical value for the t distribution with n-1 degrees of freedom. There are tables where you can look up the values for the t distribution for different degrees of freedom and confidence level. (Or, if you know the distribution of your underlying data, there are other tables you can use, but the general method is the same.)

Confidence intervals are useful for understanding certain types of experimental data. With a large enough sample size, one can calculate quite narrow confidence intervals, and have a pretty good idea that the true mean of the underlying distribution of the data lies within the window.

For example, assume there is a true unknown expected value for the average WIP, call it W. If we repeated the above procedure an infinite number of times and generated an infinite number of confidence intervals, and the assumptions we have made about normality etc are true, then 95% of these infinite number of confidence intervals will include W, e.g. only 5% of the time will we generate intervals that do not include W.

However, confidence intervals are not what we want if we want to predict, with a certain degree of assurance, the upper and lower bounds for, say, the cycle time of an individual lot. It doesn't matter if we can say with 95% confidence that the true mean cycle time for this product line is 47.3 days. That doesn't mean that this particular lot will have a cycle time of 47.3 days (in fact, given the degree of variability in a fab, that would be extremely unlikely). It matters if we can say with 95% assurance that the cycle time of this lot will be between 43 and 50 days. This is a prediction interval, as discussed below.

Prediction Intervals

A prediction interval is an estimate of an interval in which future observations will fall, with a certain probability, given what has already been observed (Wikipedia).

We think that the use of and calculation of prediction intervals are best illustrated using an example. (Frank came up with this example, based on his experience driving to his wife's school.)

Suppose we want to know how long it takes to drive from home to school. Assume there is a true unknown expected value for this driving time, and it is D.

If we drive to school 100 times and record the driving times d(n) where n=1 to 100, then we can construct a 95% confidence interval for D, the true unknown expected driving time. It will probably come out to something like (14,16) minutes.

If we drive to school 500 times and record the driving times and construct a new 95% confidence interval for D, it will most likely be a tighter interval because we'll have more data, so it will come out to something like (14.75, 15.25) minutes. If we have 5,000 observations, then the interval will be even shorter, e.g. (14.99, 15.03) minutes.

However, if we want to tell someone "I bet you $50 that today I can make it to school between 14.99 and 15.03 minutes", that is not a good bet. What we are looking for here is a *prediction interval* for the individual d(n), the time it takes to drive to school on any given day.

The prediction interval can be generated using a variety of methods; the one that we find easiest to understand is the bootstrap method. Here we are not making any assumptions about distributions, just using the data we've collected:

- Suppose we have 1000 observations of the driving time, d(1) to d(1000).
- To construct at 95% *prediction interval* for the time it takes to drive to school on any given day, we find the 2.5th percentile and 97.5th percentile and then our prediction interval is everything between these two numbers.
- So we would sort our d(n) from lowest to highest to give us dsort(1) to dsort(1000).
- The 2.5th percentile is dsort(25) (the 25th value). Suppose this is 11 minutes.
- The 97.5th percentile is dsort(975) (the 975th value). Suppose this is 18 minutes.
- Then our 95% PREDICTION INTERVAL for the time it takes to drive on any given day is (11,18) minutes.
- So then we would feel comfortable saying "I bet you $50 that today I can make it to school with total driving time between 11 and 18 minutes."

Of course there are other methods for estimating prediction intervals, ones that do take the distribution of the underlying data into account. Consult the references below for more details.

Conclusions

Whenever we have variability in data, whether that data is obtained via simulation experiments or direct observation of the fab, the distribution of the data includes important information. When we're trying to make predictions, it's often not enough to look at an expected average value, no matter how confident we are that this average value is the "true" average. What we need is a range of values in which we can have a high degree of confidence that future behavior will fall.

Cycle times are a good example of this. While the average expected cycle time value may be important for general planning (and certainly for benchmarking), if we're going to make shipment commitments, we need a more precise idea of when specific lots can be completed. This is where prediction intervals come in. A prediction interval is an estimated range in which future values are likely to occur (with a certain probability), based on actual observations already taken.

The data-rich environment of a fab lends itself to the calculation of prediction intervals (at least in cases where things like product mix haven't changed significantly). Discrete event simulation is also an excellent tool for generating prediction intervals, since it is usually easy to generate many observations of data. The trick is not to confuse prediction intervals (estimates of likely future values) with confidence intervals (estimates for a true underlying value that can't be directly observed).

Closing Questions for FabTime Subscribers

Do you use prediction intervals in any of your fab planning (cycle times, WIP levels, throughputs)? If so, what methods do you use for calculation?

Further Reading

■ NIST/SEMATECH Engineering Statistics Handbook, Section 7.1.4. What are confidence intervals?, http://www.itl.nist.gov/div898/handbook/prc/section1/prc14.htm , 2003.

■ NIST/SEMATECH Engineering Statistics Handbook, Section 4.5.1.2. How can I predict the value and estimate the uncertainty of a single response? http://www.itl.nist.gov/div898/handbook/pmd/section5/pmd512.htm , 2003.

■ George E. P. Box, William G. Hunter, and J. Stuart Hunter, *Statistics for Experiments: An Introduction to Design, Data Analysis, and Model Building*, Wiley, 1978.

■ Allan J. Rossman and Beth Chance, *Workshop Statistics: Discovery with Data, Third Edition*, Wiley, 2008.

Yale University Statistics Department, "Confidence Intervals", http://www.stat.yale.edu/Courses/1997-98/101/confint.htm .

Acknowledgements

This article originated from an email discussion between our friend Bob Kotcher (Philips Lumileds) and FabTime's Frank Chance. Frank's wife, Statistics Professor Beth Chance from CalPoly in San Luis Obispo, consulted with us on the technical details. We are grateful to Bob for raising a question that got us thinking, and to Beth for her statistical guidance.

Subscriber List

Total number of subscribers: 2679, from 453 companies and universities.

Top 21 subscribing companies:

- Maxim Integrated Products, Inc (175)
- Intel Corporation (146)
- Micron Technology, Inc. (104)
- GLOBALFOUNDRIES (93)
- Western Digital Corporation (69)
- X-FAB Inc. (67)
- Texas Instruments (66)
- International Rectifier (63)
- TECH Semiconductor Singapore (61)
- STMicroelectronics (57)
- ON Semiconductor (56)
- Analog Devices (53)
- Freescale Semiconductor (53)
- IBM (48)
- NEC Electronics (46)
- Infineon Technologies (40)
- Cypress Semiconductor (38)
- Skyworks Solutions, Inc. (35)
- ATMEL (32)
- National Semiconductor (32)
- Seagate Technology (32)

Top 3 subscribing universities:

- Arizona State University (8)
- Ben Gurion Univ. of the Negev (8)
- Nanyang Technological University (8)

New companies and universities this month:

- BI Technologies
- Shanghai Huali Microelectronics (SHLMC)
- Smoltek AB
- AmFor Electronics
- Renewable Energy Corporation (REC)
- Oclaro Inc.
- Lam Research
- Wincup

Sampler Set of Other Subscribing Companies and Universities:

- Axcelis Technologies (1)
- Digital Optics Corporation (2)
- Dvir.ie Ltd. (1)
- Fort Dodge Animal Health (1)
- Gemalto (1)
- ICG / Semiconductor FabTech (1)
- Jazz Semiconductor (6)
- JDS Uniphase (2)
- KAIST (2)
- LSI Logic (2)
- MIT Lincoln Laboratory (4)
- Oklahoma State University (2)
- Propex (1)
- Rose-Hulman / GE Healthcare (2)
- Sandia National Labs (3)
- SUMCO USA (1)
- Titan Ind Limited (1)
- Tokyo Electron Ltd (2)
- University of Central Florida (1)
- University of Shanghai for Science and
- Technology (2)

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® Cycle Time Management Training

"It was helpful to see best-in-class methods for wafer fab cycle time management. Discussing these matters in-depth with you was quite valuable, as we could ask questions specific to our fab and processes."

Shinya Morishita
Manager, Wafer Engineering
TDK Corporation

Course Code: FT105

This course provides production personnel with the tools needed to manage cycle times. It covers:

- Cycle time relationships
- Metrics and goals
- Cycle time intuition

Price

$7500 plus travel expenses for delivery at your U.S. site for up to 20 participants, each additional participant $300. Discounts are available for multiple sessions.

Interested?

Contact FabTime 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 make the best possible decisions?

- Do your supervisors possess good cycle time intuition?
- Are you using metrics that identify cycle time problems early?
- Can you make operational changes to improve cycle time?

FabTime's Cycle Time Management Training is a one-day course designed to provide production personnel with an in-depth understanding of the issues that cause cycle time problems in a fab, and to suggest approaches for improving cycle times. A two-day version and a half-day executive management version are also available upon request. As of January 1, 2011, the course is only available for delivery at customer sites within the United States.

Prerequisites

Basic Excel skills for samples and exercises.

Who Can Benefit

This course is designed for production personnel such as production managers, module managers, shift supervisors, hot lot coordinators, and production control.

Skills Gained

Upon completion of this course, you will be able to:

- Identify appropriate cycle time management styles.
- Teach others about utilization and cycle time relationships.
- Define and calculate relevant metrics for cycle time.
- Teach others about Little's law and variability.
- Quantify the impact of single-path tools and hot lots.
- Apply cycle time intuition to operational decisions.

Sample Course Tools

Excel Cycle Time Simulator

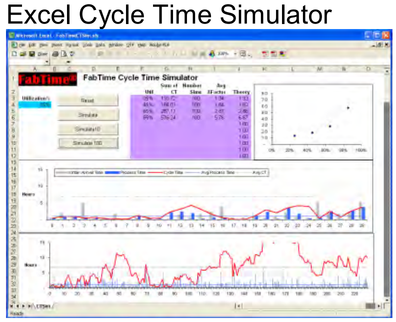

Staffing Delay Simulator