FabTime Cycle Time Management Newsletter

Volume 3, No. 3 March 2002

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.

Editor: Jennifer Robinson

Table of Contents

Contributors: Walt Trybula, Tim Stanley, and Kristin Rust (all from International SEMATECH); Jim Wright and Bill Gardner (both from Headway Technologies); Mike Hillis (AMD); Sanjay Jain (Virginia Tech); Philip Fontes (NEC); Stuart Carr (consultant); Paul Czarnocki (consultant); and Toby Patterson (PolarFab)

Welcome

Welcome to Volume 3, Number 3 of the FabTime Cycle Time Management Newsletter. This month, we are pleased to include a press release about our first FabTime customer, Headway Technologies. Headway recently chose to upgrade to FabTime Version 3.0, and also to renew their software maintenance contract. We look forward to continuing to work with them in the coming year.

I would also like to announce that I'll be presenting a paper at the Modeling and Analysis for Semiconductor Manufacturing Conference in Phoenix early next month (abstract below). The paper describes a hot lot expediting project performed by another FabTime customer, Mike Hillis at AMD. I hope to see some of you at the conference! And speaking of conferences, Frank Chance plans to be at SEMICON Europa this year - send him an email at Frank.Chance@FabTime.com if you'll be attending, and would like to meet.

This month we have lots of subscriber feedback on various topics, particularly on last month's hot lot discussion. My thanks to all who have contributed. I think you'll all find the discussions interesting. Our new topic for the month concerns quantifying the impact of tool dedication on cycle time. We present a formula for queue time as a function of traffic intensity, process time, and number of tools in the tool group, and show why, according to this formula, queue time tends to improve as tool dedication is lessened (for the same overall traffic intensity).

Thanks for reading! -- Jennifer

FabTime

325M Sharon Park Dr. #219
Menlo Park CA 94025
Tel: 408 549 9932
Fax: 408 549 9941
www.FabTime.com

Community News/Announcements

Headway Technologies Upgrades to Version 3.0 of FabTime Software, Renews Maintenance Contract

Menlo Park, CA. March 13, 2002 - FabTime Inc. today announced that Headway Technologies has upgraded to Version 3.0 of FabTime's cycle time management software, and also renewed its maintenance contract for the next year. Headway, a recording head wafer manufacturer located in Milpitas, CA, was the first customer for FabTime's software.

Version 3.0 is a significant software upgrade for Headway. The new functionality includes detailed visibility into the impact of tool downtime events. The average chart response time has also been improved in the new version, with response times now of less than one second.

"I use FabTime every day, and so do the supervisors who report to me," said Jim Wright, Headway's Production Manager. "The data that I need is right on my home page where I need it when I come in every morning."

"The FabTime installation has been very successful for us," said Bill Gardner, Headway's Director of Manufacturing. "With this new version, our equipment maintenance and engineering organizations have started using the software also, in addition to manufacturing, which made it easier for us to justify renewing the maintenance contract."

"We are very happy to be able to continue to support Headway," said Jennifer Robinson, chief operating officer of FabTime. "They have been a great customer for us, and have provided feedback that has helped us to improve the software dramatically over the past 18 months."

More information about FabTime's software is available at www.FabTime.com/software.htm.

About Headway Technologies

Headway Technologies designs and manufactures recording heads for high performance hard disk drives. Headway is a part of the TDK group of companies (NYSE: TDK), the largest independent recording head supplier to the hard disk drive industry. Headway's wafer fabrication facility is located in Milpitas, California. The company's website is located at www.Headway.com.

Press Release Distribution List

FabTime has recently created a press release distribution list. Most of our press releases will continue to be included in our monthly cycle time management newsletter. However, people on the press release distribution list will hear about them immediately upon publication, instead of waiting for the next newsletter. If you would like to receive email notification whenever FabTime publishes a new press release, simply email Press.Releases-@FabTime.com to be added to the list. The distribution list will only be used for sending you press release notifications, and your email address will not be shared with anyone outside of FabTime.

FabTime Conference Presentation Announcement

Jennifer Robinson will be presenting a paper at the Modeling and Analysis for Semiconductor Manufacturing (MASM) Conference, which will take place in Phoenix April 10th-12th. The list of papers being presented in the same track can be found at http://www.eas.asu.edu/~masmlab/masm2002/omsti.htm, along with links to more information about the conference. Jennifer's paper describes an AMD project on super-expediting lots in a wafer fab. The primary author of the paper is Mike Hillis, cycle time and line yield improvement manager from AMD. The title is "Extremely Hot Lots: Super-Expediting in a 0.18 Micron Wafer Fab". The abstract follows:

"AMD's Fab 25 has a high rate of technological change that sometimes requires "super-expediting" of extremely high priority lots. AMD developed a seven-step procedure for managing such lots: 1) ensure production management buy-in for lot priorities; 2) identify a "key expediter" responsible for getting the lot out quickly; 3) clearly state expectations for all involved in processing the lots; 4) define tactical communication methods; 5) establish a system for tracking and forecasting lot performance; 6) execute the tracking and forecasting system; and 7) report progress and get assistance from management as needed. AMD applied this procedure to the first lot of a critical new product. This lot went from being 14 days behind schedule to shipping early."

Job Change Announcement: Sanjay Jain

Sanjay Jain has recently started as Research Associate Professor with the Grado Department of Industrial and Systems Engineering at Virginia Tech. Sanjay is based at Virginia Tech's Alexandria Research Institute as the Northern Virginia representative of the Center for High Performance Manufacturing. The center is devoted to improving the performance of manufacturers, particularly in the state of Virginia. In the past, Sanjay has worked with the semiconductor industry in Singapore in the area of supply chain management and planning and scheduling, and he is looking forward to working with partners from the industry here. For more information, please visit the center website at www.eng.vt.edu/chpm.

FabTime welcomes the opportunity to publish community news and announcements. Simply send them to Jennifer.-Robinson@FabTime.com.

Subscriber Discussion Forum

Hot Lots - Issue 3.2

Last month's main article on hot lots sparked quite a bit of reader discussion, as follows:

Walt Trybula (International SEMATECH) wrote: "Regarding your hot lot discussion. If you continue to increase the number of hot lots the cycle time increases as you indicated; however, you reach a point where thee cycle time starts to decrease. I won't provide a detailed analysis, but only indicate that the cycle time for 100% hot lots == 0% hot lots in a fifo system. I have actually seen a situation with the shortest cycle time was with almost 10% hot lots. In general, the average total cycle time for 0% and 100% should be equal. The price you pay for expediting is the delay in non-expedited. We are in agreement."

Tim Stanley (International SEMATECH) wrote: "With regard to your article on Hot Lots, the real problem with Hot Lots as implemented in fabs I have worked in, is that they do several additional things beyond just prioritize queues. They hold tools, impose additional setups, smaller batches and smaller lots. These things not only impact non-hot lot cycle time, they reduce the overall capacity of the fab. One of the people in my group at SEMATECH, Kristin Rust, did her Masters Thesis on this subject. If one of your readers would like additional information on her work, she can be contacted at Kristin.Rust@SEMATECH.org."

Kristin Rust added: "I am sending you a copy of a paper presented at SMOMS 2001. It summarizes the work presented in my Masters Report. I hope this proves helpful/interesting to you and/or your readers. Certainly, you are correct in your observation that hot lots have not been well covered in literature. Sematech has plans to further investigate hot lots in new simulation scenarios. Specifically, we are working with Brooks Automated Planning and Logistics to develop an extension to their simulation software that will model more accurately managerial decisions which move hot lots through the factory faster. Also, one of our current interns is planning in his Masters Report to model single wafer processing of hot lots." If anyone would like a copy of Kristin's paper, you can request it from her directly, or by writing to Jennifer.Robinson@FabTime.com.

Paul Czarnocki (consultant) wrote: "The issue of quantifying the number of Hot lots allowed in a manufacturing area has been an issue facing every Fab that I've worked in. Most Fabs have different types of Hot lot designations. For instance, the highest priority, a Rocket lot, would require the operator to be notified of its arrival time in order for the tool to be set up specifically for this lot ahead of time. This decreases the utilization of this tool but may be a small price to pay for the quick processing of this lot to get the data or in support of a customer with a line down situation. Rocket lots are relatively few in number. An attempt is made to have only one such lot in a given area of the line.

A second priority is for Hot lots. Generally, they will not require a change to the tool but will be the next lot processed on the appropriate tool. The quantity of these type of lots are relatively high. Their impact on overall Fab cycle time is difficult for management to see at any given time because they most often rely on printed reports showing activities and product cycle time. Visibility is much clearer with the use of the FabTime Cycle Time software. It will show that these type of lots will make an impact on all other lots running in the line. Monitoring the total number of hot lots can then be better managed." Note: Paul worked until recently for AMD's Fab 25, and was instrumental in AMD's decision to install FabTime's cycle time management software.

Stuart Carr (consultant) wrote: "The discussion about hot lots and 2 priority classes reminded me of my thesis! In the unlikely event that anyone is interested in some gory math details around queueing models of a single workcenter with 2 priority classes, you could refer them to my thesis. In it, I give a recursive formula to compute the steady-state distribution of the size of the low and high priority queues (and, hence, the distribution of the wait time for each queue). The thesis title is "A Partial Make-To-Order Production-Inventory Strategy for Industrial Manufacturers", from Jan. 1995, School of Operations Research and Industrial Engineering, Cornell University. If someone's interested, I could email them the postscript file."

Another subscriber wrote: "Our system has 4 separate priority levels - regular and type 1-3 with 1 being the highest priority. We limit the total number of lots that have priority since the impact is not limited to capacity but is actually much more significant in people resources needed to manage and track these lots. The rule is that you can have up to X total priority lots, up to Y total priority 1 and 2 lots and up to Z priority 1 lots (X>Y>Z of course). We also have a "pain" indicator attached to some of the lots (mainly new products) and limit this level. We have a capacity impact calculation that captures when a tool will lose capacity (and gain CT over the baseline) to check we do not lose too much capacity. All in all - we still feel that we do not understand all that is needed on this subject and a lot of the management is guess work!"

Production Equipment Efficiency (PEE) - Issue 3.1

Philip Fontes (NEC) wrote: "I am considering proposing to our Line Maintenance group your view-point on using PEE as a superior metric over OEE since it does not work against Cycle Time improvement goals. I was just wondering how much greater an effort it will be to segregate down time and idle time into "good" and "bad" based on whether or not product was waiting. We have tool condition codes that have been created in the MES such as "Idle-no operator," "Idle-no product," and "Down for repair-lots waiting," but their use is not regular or consistent. Do others who use PEE associate this "good" and "bad" time through algorithms/clever analysis of reported tool status, and known lot status (with appropriate recipe requirement and tool restrictions,) or do they rely on operators to designate the "tool/product" condition? Which do you feel would be more successful in implementing?"

Phil also added this comment, which we particularly liked, in regards to the newsletter as a whole:

" . . . . When the reading stack grows - it actually never seams to shrink - the Fab Time newsletter always rises to the top!"

FabTime Response:

We appreciate the positive feedback - it definitely helps to keep us motivated. Regarding your question, I don't actually know how people implement this metric. I wrote about it because it seems so obvious as a solution for people who care about cycle time. But I think that you make an excellent point about the difficulty of separating out the "good" vs. "bad" downtime in practice, especially when people are accustomed to thinking of all downtime as "bad".

I do know of one company that just groups productive time and standby time together in their tool-level reporting, because their goal is to maximize the time that each tool is available for processing WIP, whether there is WIP in front of the tool at the time or not. They then use move goals to track whether each tool is actually spending enough time processing WIP. This seems to serve their purpose fairly well.

Anyway, what the PEE calculation is saying is that if you can arrange for your downtime to occur, as much as possible, when there is no WIP waiting at a tool, then you should not be penalized for that downtime when making improvement efforts. Clearly, it is better if don't have lots waiting while you are performing maintenance events. But I don't know how people break out this "good" vs. "bad" downtime in practice, or even if they do. I will have to raise this as a question for the next newsletter issue.

My personal feeling is that increasing the reporting burden on operators should be avoided if possible, and that some sort of algorithms or estimates should be possible. As a first pass, you could just neglect the "down no product" time altogether. This gives something between the traditional OEE and the newer PEE. You don't have a way to encourage people to perform maintenance events when no lots are waiting through this approach, but you do account for the "idle no product" time, and you in general continue to encourage people to reduce the time spent on downtime and maintenance events. But I'll have to ask the other newsletter subscribers if anyone has feedback on the use of PEE in practice.

Single Wafer Lots

Walt Trybula of International SEMATECH would like to exchange views on single wafer lots. He writes: "The question has been raised about single wafer lot cycle time. I am looking for opinions on how people are modeling this. One case is to run the wafer and evaluate the results before permitting the next wafer to start -- but a throughput of 6 wafers an hour is not real good. However, early development wafers could employ that approach. Once the process is proven and the design stable, then starting multiple designs through the stepper without permitting any evaluations is fine. I'm curious on what algorithms people are employing to handle the percentages of product types and reprocessing of failures. Another question that can be raised is do single wafer lots have higher rework rates?" If anyone has any thoughts on this subject, we, and Walt, would be interested to hear them.

Characteristic Curves

Philip Fontes also asked: "When attempting to calculate/plot characteristics curves for a fab, do you typically plot "complete average-cycle time" against "overall-average fab % tool utilization," or does one just focus in on the time through a select few "bottleneck" tools against the respective % utilizations for those bottleneck tools?"

FabTime Response:

The way I have seen characteristic curves most commonly used is to plot "complete average cycle time" (actually weighted average cycle time for shipped lots, across all products, sometimes normalized by dividing by the theoretical process time) vs. overall fab loading. Rather than being an average of the tool utilizations, the fab utilization is equal to the loading of the bottleneck tool group. The reason for this is that once you load the bottleneck to 100%, you can't increase throughput in the fab, so it provides the best overall guide for how the fab is loaded.

The reason that most people use the factory, or bottleneck, loading, rather than generating different curves for different bottleneck tools has to do, I think, with local vs. global optima. A fab is a complex environment, with reentrant flow making the performance of the tool groups particularly inter-related. This makes it hard to look at the impact of performance improvements without looking at their impact on the entire fab. Naturally, if you improve the performance of the bottleneck, you will almost always improve the performance of the overall fab, so there is some value to looking at characteristic curves of the bottleneck (or bottlenecks) alone. How much value depends on the type of improvement that you are evaluating (e.g. for AMHS, I would always look at factory models, but if I just wanted to know the impact of a change in setup policy on the bottleneck, I might only look at that toolgroup).

Since people usually care most about improving the overall cycle time of lots shipped, most studies that I have seen use simulation to generate the factory-level curve under each of several different alternatives. Then they compare the performance of the different alternatives by either comparing cycle times at some factory loading value, or comparing the factory loading values that correspond to a given average cycle time. An excellent paper that discusses this methodology is:

J. W. Fowler, S. Brown, H. Gold, and A. Schoemig, "Measurable Improvements in Cycle-Time-Constrained Capacity," *Proceedings of the 6th IEEE/UCS/SEMI International Symposium on Semiconductor Manufacturing (ISSM)*, October 6-8, 1997, San Francisco, A21-A24. (The full paper can be requested from http://www.fabtime.com/abs_SiemFab.htm).

As a side note, you might notice that I having been using the term "loading" instead of "utilization" in this discussion. There is a slight difference in these measures. Tool group loading is the actual input rate divided by the maximum process rate that the tool can sustain, after accounting for downtime, etc. So if you have a tool that is down 25% of the time, busy 25% of the time, and idle 50% of the time, that tool is utilized (busy) 50% of the time, but its loading is 33.3%, because it is using 25% out of the 75% time available for processing (or 1/3). This example assumes that the downtime percentage is independent of loading, otherwise the calculations are a bit more involved.

Foundry Performance Data

Toby Patterson of PolarFab asked: "I would like to know if FabTime has ever put together a report comparing the cycle times and yields of world-wide foundries. If so, could you tell me how I may obtain this report? If not, do you know if any reports exist??"

FabTime response: We haven't heard of a report like this, but think that it would be very useful. Many of the foundries have been known to make this type of information publicly available, so it's certainly possible for there to be a report out there, but we don't know of one. If any other subscribers do, please let us know. Thanks!

How Much Does Tool Dedication Inflate Cycle Time?

Introduction

We talked back in issue 1.08 about the fact that single path tools tend to drive up cycle times. The question is, how much does tool dedication inflate cycle times? The are sometimes important reasons to have dedicated tools. What's needed is a way to explore trade-offs. In this article, we present an approximation for queue time as a function of number of machines in a tool group. This approximation clearly shows that queue time decreases as the number of tools in the group increases (for the same total traffic intensity of the tool group).

Queue Time as a Function of Number of Servers

Sakasegawa proposed an approximation for average queue time as a function of traffic intensity, processing rate, and number of tools. This approximation is appropriate for the case where the time between arrivals is exponentially distributed (highly variable) and the process times also follow an exponential distribution.

Let

rho = traffic intensity of the group (arrival rate divided by total processing rate) (dimensionless)

mu = processing rate of each tool in the tool group (e.g. lots/hour)

s = number of tools in the tool group

Then

Queue Time ~= [rho ^ (SQRT{ 2 * (s + 1)} - 1)] / [s (1 - rho) * mu].

When s = 1, this equation simplifies to

Queue Time = [rho] / [(1 - rho) * mu], which is the queue time for an M/M/1 queue (as discussed in the previous issue).

When s is greater than 1, the numerator contains rho raised to some power greater than 1. Since rho is always a fraction between 0 and 1 (required to be less than 1 for stability), the numerator will decrease as it is raised to a larger and larger power. And so the numerator will decrease as s increases. The denominator will increase linearly as s increases (since we have s multiplied by (1 - rho)). Decreasing the numerator and increasing the denominator both serve to decrease the total value of the queue time.

Example

Let rho = 0.90 (the tools in the group are 90% loaded) and mu = 4 lots/hour for each tool (15 minute average process time).

When s = 1, QT = 0.90 / ((1 - 0.90) * 4) = 0.9 / 0.4 = 2.25 hours

When s = 2, QT ~= [0.90^(sqrt(6) - 1)] / [2* (1 - 0.90) * 4] = [0.90^1.449] / [2 * 0.4] = 0.858 / 0.8 = 1.07 hours

Note that as we go from 1 to 2 tools, rho in the numerator is raised to a power of 1.449. Since rho (0.9) is less than 1, raising it to a power greater than one yields a lower result (0.858). The denominator, by contrast, is doubled in size (since it is multiplied by s). The resulting queue time decreases by slightly more than 50%.

This decrease continues as s is increased further, though the steepest drop occurs from one tool to two tools, as shown below:

s = 1, QT ~= 2.25
s = 2, QT ~= 1.07
s = 3, QT ~= 0.69
s = 4, QT ~= 0.50
s = 5, QT ~= 0.39
s = 6, QT ~= 0.31
s = 7, QT ~= 0.26
s = 8, QT ~= 0.22
s = 9, QT ~= .19
s = 10, QT ~= .17

Note that in all 10 cases, the traffic intensity on each tool is the same. Suppose that you have ten layers and ten steppers. If you dedicate one stepper for each layer, you have the s=1 case (complete tool dedication). If you can cross-qualify each stepper for two layers, so that every layer has a choice of two steppers, you have the s=2 case. If you can cross-qualify each stepper to run all ten layers, you have the s=10 case. This method gives an approximation for the cycle time penalty of various degrees of tool dedication.

Intuitive Explanation

In the example above, the decrease in

queue time is quite dramatic as the number of tools increases. Once there are 10 tools in the tool group, the queue time is practically insignificant, even though the tools are all loaded to 90% of capacity. Think about it this way. When you have a group with 10 tools, each loaded to 90% of capacity, each tool is going to be idle 10% of the time (neglecting downtime). If arrival times and process times are highly random, then much of the time, when a lot arrives, at least one tool will be idle and ready to process that lot. By contrast, if you only have one tool in a group, and lots can arrive any time, 90% of the time, when a lot arrives, the tool will be busy, and the lot will have to wait.

Examples in Everyday Life

The obvious example that comes to mind here is the grocery store. If instead of having 10 separate checkout lines your grocery store was configured to have a single line for all 10 cash registers, your average waiting time would decrease dramatically. This is because you would never be waiting in line for one cash register, while another one down the line became free. You would never be stuck immediately behind the person with 87 different coupons, because you would get the next available register. However, the day of having a single line at the grocery store is not likely to come any time soon. Grocery stores probably use separate lines because of space constraints (the carts are very large).

On the other hand, banks and airport check-in often share a single waiting line among multiple servers, e.g. each server is cross-qualified to handle a range of transactions (although the airlines will often provide a dedicated server for hot lots -- first class customers!).

Conclusions

We're not saying that you should completely eliminate tool dedication in your fab There are process reasons for it, just as there are customer service reasons why airlines have dedicated check-in counter for frequent travelers, and political reasons why transportation systems have car-pool lanes. What we are presenting is a formula that you can use to estimate the impact of tool dedication decisions on lot queue times. Understanding the consequences of tool dedication policies puts you in a better position to make decisions. For cycle time, it is better to have larger tool groups, on which more operations are processed, than to have smaller tool groups that are dedicated to fewer operations. This formula may help you to estimate "how much better?"

Closing Questions for FabTime Subscribers

- How do you make decisions about tool-dedication in your fab?
- Is there pressure to dedicate tools to particular operations for process / yield reasons?
- Do you have a procedure for reducing dedication when processes become more stable?

If you have any observations from your own experiences about cycle time and tool dedication trade-offs, please send them to Jennifer.Robinson@FabTime.com. They will be published in the next issue (with or without your name included, as you prefer).

References

- The Sakasegawa approximation can be found in the book Factory Physics: Second Edition, by W. J. Hopp and M. L. Spearman, pg. 272.

- For a discussion of tool-dedication and cross-qualification of steppers in an Infineon wafer fab, see the article by J. W. Fowler, S. Brown, H. Gold, and A.

Schoemig, "Measurable Improvements in Cycle-Time-Constrained Capacity," *Proceedings of the 6th ISSM Conference*, October 6-8, 1997, San Francisco, A21-A24. You can request a PDF copy of the full paper from our website (see abstract at www.fabtime.com/abs_SiemFab.htm).

■ Another article that explores the cycle time impact of tool dedication is: M. Mittler, "Two-Moment Analysis Of Alternative Tool Models With Random Breakdowns," *Proceedings of the 1996 IEEE Conference on Emerging Technologies and Factory Automation*, Kauai, HI, 546-552, 1996. Email Jennifer.Robinson@FabTime.com to request an abstract of this paper.

FabTime Recommendations

Universal Currency Converter

We recently needed to know the conversion rate from dollars to yen. A short Internet search yielded the Universal Currency Converter. This convenient web utility converts between every type of currency that we could think of, using an easy drop-down list interface. You simply enter the amount to be converted, select from and to values from the lists, and press go. The site says that it uses live currency rates. The converter was created by a company called XE. Their mission is "to facilitate the globalization of commerce." They also have a table on their homepage that gets refreshed every five minutes, showing how much a U.S. dollar and a Euro are each worth in 10 other countries. You can find the converter at http://www.xe.com. We think that the site is worth bookmarking.

The Elements of Style

FabTime's President, Frank Chance, recommends "The Elements of Style" by Strunk and White. Frank said "I like it because it's concise, and not the slightest bit wishy-washy. They say exactly what you should and shouldn't do. I think it's a big help in writing." You should be able to find the "The Elements of Style" at any bookstore, or at Amazon.com.

Subscriber List

Total Subscribers: 771

1st Silicon (1)
3M Company (2)
Abbie Gregg Inc. (2)
ABB Semiconductors (4)
Adexa Corporation (1)
Advanced Micro Devices (38)
Affymetrix (1)
Agere Systems (5)
Agilent Technologies (4)
Aisin Indonesia (1)
Alfalight Canada (1)
Alpha Industries (1)
Alpha-Sang (1)
AMI Semiconductor (2)
Amkor (3)
AMR Research (1)
Analog Devices (5)
Andes University (1)
Applied Materials Corporation (9)
Aralight Corporation (2)
Arch Wireless (1)
Arizona State University (5)
Arkansas Tech University (1)
Asia Management Group (1)
ASM International NV (1)
ASML (3)
ATMEL (4)
Axcelis Technologies (1)
Axsun Technologies (1)
Bookham Technology Plc (1)
Boston Scientific (1)
Bovis Lend Lease Microelectronics (1)
BP Solar (3)
Brooks Automation (2)
Byelorussian State Economic Univ. (1)
Cabot Microelectronics Ltd. (1)
California Micro Devices (1)
California Polytechnic State University (1)
C&D Aerospace (1)
Cannon Precision (1)
Canon USA (1)
Carsem M Sdn Bhd (3)
Chartered Semiconductor Mfg (24)
CMC Electronics (1)
Compugraphics International Ltd. (1)
Conexant Systems, Inc. (4)
Continental Device India Ltd. (1)
Cornell University (1)
Corning (1)
Cox High Speed Internet (1)
C-Port Corporation (1)
Cree, Inc. (1)
Cronos Integrated Microsystems (1)
Cummins S. de R.L. de C.V. (1)
Cyberfab (1)
CyTerra Corporation (1)
Dallas Semiconductor (2)
Datacon Semiconductor Equipment (1)
Delphi Automotive Systems (1)
Delta Design (1)
Dick Williams and Associates (1)
DomainLogix Corporation (1)
Dominion Semiconductor (5)
Dow Corning Corporation (1)
Durham ATS Group (4)
Dwarkadas Associates (1)
Eastman Kodak Company (3)
Electroglas, Inc. - Statware Division (2)
e-METS Co, Ltd (1)
Ernst & Young (1)
eSilicon Corporation (1)
Eskay Corporation (1)
FabOptima GmbH (1)
FabTime (3)
Fairchild Imaging (1)
Fairchild Semiconductor (3)
Fraunhofer IPA (1)
Fujitsu Microelectronics, Inc. (1)
General Semiconductor (2)
Headway Technologies (2)
Hewlett-Packard Company (6)
Hitachi, Ltd. (1)
Hitachi Nippon Steel Semiconductor (4)
Huck Fasteners (1)
Hynix Semiconductor Mfg America Inc. (1)
IBM (10)
ICG / Semiconductor FabTech (1)
IDC (7)
IMEC (2)
Infineon Technologies (32)
Infosim Networking Solutions (1)
INSEAD (3)
Institut National Polytech. de Grenoble (2)

Integrated Device Technologies (2)
Integrated Technologies Company (2)
Intel Corporation (35)
Intelligent Quality Systems (1)
International Rectifier / HEXAM (3)
Intersil (3)
Interstar Technology (1)
Jacobs Consultancy (1)
James Nagel Associates (1)
JDS Uniphase (3)
Johnstech International Corp. (1)
Kansas State University (1)
Ken Rich Associates (1)
KLA-Tencor (1)
Kulicke & Soffa Industries, Inc. - K&S (2)
Kymata - Alcatel (1)
Lexmark International, Inc. (1)
Linear Technology (1)
Litel Instruments (2)
LSI Logic (9)
Lynx Photonic Networks (1)
Macronix International Co. (5)
Managed Outsourcing, Inc. (2)
Manufacturing Integration Technology (1)
MASA Group (1)
Maxim Integrated Products, Inc. (3)
Maxtor (1)
MECA Electronics, Inc. (1)
Medtronic (5)
MEMS Optical (1)
Methode Electronics, Inc, (1)
Metrology Perspectives Group (1)
Micrel Semiconductor (2)
Microchip Technology (1)
Micron Technology, Inc. (1)
Micro Photonix Int. (1)
MicroVision-Engineering GmbH (1)
Motorola Corporation (44)
MTE Associates (1)
Nanometrics (2)
Nanyang Technological University (3)
National Chiao Tung University (1)
National Semiconductor (10)
National Univ. of Ireland - Galway (1)
National University of Singapore (2)
NEC Electronics (7)
Nortel Networks (6)
Ohio State University (1)
Oklahoma State University (1)
ON Semiconductor (8)
Onix Microsystems (1)
Palmborg Associates, Inc. (2)
Pelita Harapan University (1)
Penn State University (1)
Peter Wolters CMP Systeme (1)
Philips (18)
Piezo Technology Inc. (1)
Planar Systems (2)
PolarFab (3)
Politecnico of Milano (1)
Powerex, Inc. (3)
PRI Automation (2)
Productivity Partners Ltd (1)
ProMOS Tech. (1)
Propsys Brightriver (1)
PSI Technologies, Inc. (1)
Quanta Display Inc. (1)
Ramsey Associates (1)
Raytheon (1)
Read-Rite Corporation (3)
Redicon Metal (1)
Rexam (1)
Rockwell Automation (1)
RTRON Corporation (2)
SAMES (1)
Samsung (14)
Saint-Gobain Company (1)
Seagate Technology (18)
SEMATECH (17)
Semiconductor Research Corp. (1)
SemiTorr NorthWest, Inc. (1)
Serus Corporation (1)
SEZ America, Inc. (1)
Shanghai Grace Semiconductor Mfg. (1)
SiGen Corporation (1)
Silicon Integrated Systems Corp (1)
Silicon Manufacturing Partners (4)
Silterra Malaysia Sdn. Bhd. (5)
Sipex Corporation (1)
Sony Semiconductor (1)
SoundView Technology (3)
SSMC (1)
STMicroelectronics (27)
Stonelake Ltd. (1)
Storage Technology de Puerto Rico (1)
Superconductor Technologies, Inc. (1)

Süss MicroTec AG (2)
Synergistic Applications, Inc. (1)
Synquest (2)
Takvorian Consulting (1)
TDK (1)
TECH Semiconductor Singapore (21)
Terosil, a.s. (1)
Texas A&M University (1)
Texas Instruments (13)
Tokyo Electron Deutschland (1)
Tower Semiconductor Ltd. (1)
Triniti Corporation (1)
TriQuint Semiconductor (7)
Tru-Si Technologies (1)
TRW (1)
TSMC (4)
UMC (7)
Unisem (1)
United Monolithic Semiconductors (1)
Unitopia Taiwan Corporation (2)
University of Arkansas (1)
University of California - Berkeley (4)
University of Mining and Metalurgy (1)
University Porto (1)
University of Southern California (1)
University of Texas at Austin (1)
University of Virginia (1)
University of Wuerzburg - Germany (1)
Velocium (1)
Virginia Tech (5)
Vitesse Semiconductor (1)
Wacker Siltronic (4)
WaferTech (10)
Win Semiconductor (1)
Wright Williams & Kelly (8)
Xerox Brazil (1)
X-FAB Texas, Inc. (3)
Yonsei University (1)
Zarlink Semiconductor (4)
Zetek PLC (1)
Unlisted Companies (15)

Consultants:
Carrie Beam
Vinay Binjrajka (PWC)
Javier Bonal
Steven Brown
Stuart Carr
Alison Cohen
Paul Czarnocki
Scott Erjavic
Greg Fernandez
Ted Forsman
Navi Grewal
Cory Hanosh
Norbie Lavigne
Michael Ray
Bill Parr
Nagaraja Jagannadha Rao
Lyle Rusanowski
Mark Spearman (Factory Physics, Inc.)
Dan Theodore
Craig Volonoski

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. To protect the privacy of our subscribers, email addresses are not printed in the newsletter. If you wish contact the subscribers from a particular company directly, simply email your request to the editor at Jennifer.Robinson@FabTime.com. To subscribe to the newsletter, send email to the same address. You can also subscribe online at www.FabTime.com/newsletter.htm. We will not, under any circumstances, give your personal information to anyone outside of FabTime.

Copyright (c) 2002 by FabTime Inc. All rights reserved. www.FabTime.com