Fab Cycle Time Improvement through Inventory Control:
A Wafer Starts Approach

by

Matthew J. Ward

Bachelor of Science in Mechanical Engineering, University of Illinois at Urbana-Champaign (1995)

Master of Science in Mechanical Engineering, University of Illinois at Urbana-Champaign (1997)

Submitted to the Engineering Systems Division
and the MIT Sloan School of Management
In Partial Fulfillment of the Requirements for the Degrees of

Master of Science in Engineering Systems
Master of Business Administration

In Conjunction with the Leaders for Manufacturing Program at the
Massachusetts Institute of Technology
June 2007.

Signature of Author_____
May 11, 2007
MIT Sloan School of Management
Engineering Systems Division

Certified by_____
Stephen C. Graves, Thesis Supervisor
Abraham Siegel Professor of Management

Certified by_____
Deborah J. Nightingale, Thesis Supervisor
Professor of Engineering Systems and Aeronautics Engineering

Accepted by_____
Debbie Berechman
Executive Director, MIT Sloan MBA Program

Accepted by_____
Richard de Neufville
Professor of Engineering Systems
Chair, Engineering Systems Division Education Committee

[THIS PAGE INTENTIONALLY LEFT BLANK]

Fab Cycle Time Improvement through Inventory Control:
A Wafer Starts Approach

By Matthew J. Ward

Submitted to the MIT Sloan School of Management and the Engineering Systems Division on May 11, 2007 in Partial Fulfillment of the Requirements for the Degrees of

Master of Business Administration
and
Master of Engineering Systems

Abstract

Intel's Colorado Springs wafer fabrication facility, known internally as F23, has undertaken several initiatives to reduce cycle time including High Precision Maintenance (HPM), content reduction through the application of Manufacturing Excellence (mX), effective utilization of production equipment, and aggressive inventory control. Each of these efforts has contributed to the marked improvement F23 achieved throughout 2006. F23's cycle time efficiency, the ratio of raw process cycle time to actual fab cycle time, improved from 12% (worst amongst Intel facilities) to greater than 35% (best amongst Intel sites), and overall cycle time was reduced by more than 61% in 2006. Inventory control was found to have a major impact on factory cycle time and performance.

F23 controls its factory work-in-process, WIP, inventory through the F23 Wafer Starts Protocol. F23 utilizes Little's Law (Cycle Time = Inventory / Output) to identify target WIP inventory levels required to achieve particular cycle time goals. The target inventory is then achieved by modulating wafer starts. To do this, the Wafer Starts Protocol monitors the inventory of the overall fab and the constraint operations and suggests the amount of wafers to start for each shift. Maintaining the target inventory level drives the overall factory cycle time towards the cycle time goal. Using the starts protocol, F23 has reduced its inventory by 44% while ramping factory output.

During the implementation of this wafer starts protocol, F23 began tracking a new inventory metric to determine factory performance. Critical WIP ratio was introduced to evaluate the factory's inventory relative to the theoretical minimum inventory based upon a given factory output level and raw process time. F23 also found that this metric provides an effective comparison of inventory level between fabs.

The Fab23 Wafer Starts Protocol is one of the ways in which F23 has applied Manufacturing Science tactics and principles to drive cycle time improvements. F23 has found that inventory control can have significant impacts on factory cycle time. This is one of the reasons why F23 was able to achieve dramatic cycle time improvement.

Thesis Supervisor: Stephen C. Graves
Title: Abraham Siegel Professor of Management, MIT Sloan School of Management

Thesis Supervisor: Deborah J. Nightingale
Title: Professor of Practice of Aeronautics and Astronautics and Engineering Systems, Department of Aerospace and Aeronautics and Engineering Systems Division

[THIS PAGE INTENTIONALLY LEFT BLANK]

Acknowledgements

I would like the Intel personnel at F23 for opportunity to work with them on their improvement initiatives. I count myself as very fortunate to have had the opportunity to work with them and learn from their collective efforts. In particular, I would like to thank Zane Rakes and Clay Carlson for their direction and advice during my internship. They were willing to invest the effort to ensure that my time at F23 was valuable for everyone involved. I would also like to thank Andrea Balmores and Eric Slyter. I appreciate their willingness to invite me to work closely with them as we developed the wafers starts protocol. Joe Lewis must also be acknowledged. Joe was an invaluable resource for my internship. Joe was always willing to answer my questions and provide further background on the detailed operations of F23. Finally I would also like to thank my family and my fiancé. It is through their support and love that I am able to accomplish all of my goals.

[THIS PAGE INTENTIONALLY LEFT BLANK]

Table of Contents

[THIS PAGE INTENTIONALLY LEFT BLANK]

Table of Figures

[THIS PAGE INTENTIONALLY LEFT BLANK]

1: Introduction

Cycle time has become a key competitive aspect of the semiconductor industry. Consumer demand and product lifecycles have become less certain as consumer electronics have proliferated into a wide range of devices. To better accommodate this competitive environment, Intel's manufacturing organization has made cycle time improvement a strategic objective. With this direction the various organizations within Intel's internal supply chain have pursued reductions in their specific cycle times. Within Intel, fab cycle time is defined as the time required to produce wafers of functioning microprocessors from raw silicon[a].

Intel's wafer fabrication facility in Colorado Springs (F23) is a 200mm wafer fabrication facility. It is the primary facility producing a wireless communication product that is important for Intel's wireless applications. In order to achieve the cycle time improvement objective, F23 has pursued a number of initiatives. These initiatives can be grouped into four general categories: tool maintenance, effective utilization of tools, waste elimination, and inventory control. All of these efforts have contributed to F23's cycle time success.

[a] Intel uses two specific measures of fab cycle time: static lot cycle time and dynamic lot cycle time. Static lot cycle time is measured for individual lots as the time from when the lot enters the fab for manufacturing to the time it exits the fab. Dynamic cycle time is calculated by determining the average time lots spend at each operation, and then summing those average times for all operations within the fab for the manufacturing process. A lot consists of 25 wafers, and each wafer can contain hundreds of microprocessors.

1.1: HPM:

F23 continues to improve manufacturing equipment (tool) performance through its High Precision Maintenance (HPM) efforts. Based upon the Total Productive Maintenance of the Toyota Production System, HPM is Intel's company wide initiative to improve tool performance. Much of Intel's HPM methodology is based upon the practices of Agilent Technologies in Fort Collins, Colorado, described by Jim Leflar[1]. Rebecca Fearing examined the scheduling of maintenance activities at Intel, and her thesis provides an overview of Intel's HPM practices[2].

1.2: Ue:

F23 has relentlessly pursued the effective utilization of its manufacturing tools. The goal of effective utilization (Ue) is to ensure that all manufacturing tools that are available and have WIP present are in operation. F23 uses a web-based software application to measure tool status. The system indicates when tools are in operation, idle with WIP, idle without WIP, and down. Manufacturing management reviews the factory performance daily and requires factory personnel to document the causes for all "idle with WIP" occurrences with durations of 30 minutes or longer. This focus on tool status has driven a systematic reduction in the number of occurrences where tools, that are operational, are idle while WIP is present.

1.3: mX:

Manufacturing Excellence (mX) is the term Intel has applied to its lean manufacturing initiative. The seminal document underpinning Intel's mX initiative is the Spear and Bowen article describing the key aspects of the Toyota Production System[3]. F23 has been practicing mX for approximately 3 years. A primary focus of F23s' mX efforts is the reduction of waste. F23's waste reduction has occurred in various forms; some examples include the elimination of entire manufacturing operations, reduction in distances for technician to travel while performing tasks, and removal of automation transactions that technicians must perform to process WIP. These efforts to reduce content have improved the time required for both value-added and non-value added activities.

1.4: Wafer Starts Protocol:

F23 has pursued specific cycle time targets by aggressively controlling the amount of inventory present within the fab. F23 controls its fab inventory by varying the introduction of new wafers into the factory. This approach differs from the typical Intel fab approach of starting wafers by rigidly adhering to a predetermined wafer starts schedule, except in extreme circumstances. F23's wafer starts protocol is the focus of this paper. For a general overview of semiconductor manufacturing, the work of Quirk and Serda[4] is recommended.

1.5: Thesis Structure

Chapter 2 provides an overview of WIP inventory control in the general literature and as it pertains to semiconductor manufacturing.

Chapter 3 describes the Wafer Starts Protocol, the inventory control methodology developed at F23, and compares it to Intel's traditional wafer starts process.

Chapter 4 presents the results of implementing the F23 Wafer Start Protocol in actual production.

Chapter 5 describes the Critical WIP ratio which was used to determine F23's performance in reducing inventory and to compare different Intel fabs.

Chapter 6 discusses the various challenges that F23 faced in changing the way wafer starts are managed within Intel.

Chapter 7 summarizes the lessons learned through the implementation of the F23 Wafer Starts Protocol.

Chapter 8 suggests next steps that could be taken to advance the Wafer Starts Protocol.

Chapter 9 highlights the factors that contributed to F23's ability to break some Intel norms and experiment with inventory control.

Chapter 10 summarizes the conclusions gained from implementing the F23 Wafer Starts Protocol in production and its impact on F23's operation.

2: Overview of WIP Inventory Control

WIP management policies can have significant impact on factory cycle time. The impact of inventory on factory cycle time is well established. In 1961, Little[5] related the queue length, the wait time and the arrival rate for a queuing system in a relationship that has since become known as Little's Law. For an operations management context, the application of Little's Law translates to a relationship between the work-in-process inventory L, the cycle time λ, and the output rate W of a manufacturing system where $L = \lambda W$. The reentrant nature of semiconductor manufacturing further complicates the decisions associated with WIP management. Graves et al[6] noted that in a reentrant work environment the decision on how to sequence production material can significantly affect cycle time. Therefore in semiconductor manufacturing, the WIP management policy must comprehend both lot release, the process of introducing new production material into the fab, and lot scheduling, the determination of what order to run the particular lots that are present at a given operation. Wein[7] conducted a comprehensive comparison of lot release and lot scheduling combinations. In his study, Wein concludes that control over the input of material into a fab, lot release, has the greatest impact on fab cycle time.

There are several methods for controlling inventory release. The Toyota Production System employs the concepts of Just-In-Time (JIT) and Kanban cards[8], where cards are used to maintain a constant inventory level between operations. Levitt and Abraham[9] considered the application of Kanban in wafer fabrication facility and determined a method for calculating the proper number of Kanban cards when multiple constraints exist. There have also been actual implementations of JIT within the

semiconductor industry. Martin-Vega et al[10] describe one such example where Kanban cards were employed within a particular functional area of a Harris Semiconductor fab.

Spearman et al[11] suggested the concept of CONWIP as an alternative to Kanban. In CONWIP a constant inventory level is maintained within the overall factory with the assumption that material will naturally queue in front of bottleneck operations. The constant level is achieved by not introducing another lot into the fab until one has exited. Another alternative is workload regulation; in workload regulation, an attempt is made to more accurately control the workload of constraints. The amount of work present at the constraint is monitored and material is released in order to maintain a queue of WIP in front of the constraint that has a certain amount of processing time. Rose's[12] CONLOAD extends the concept of workload regulation by considering the average cycle times of lots of various types present within the factory.

The concept of inventory control is also not new within Intel. McBride and Kempf[13] suggested that inventory control charts of all constraint and near-constraint equipment sets could be used to determine the release of wafer starts into the factory. In 2005, Jason Connally[14] suggested the examination of CONWIP as part of F23's future lean manufacturing efforts, and in 2006, Rebecca Fearing[2] posited that F23 might have excess inventory.

3: F23 Inventory Control

Based upon the status of the literature and earlier efforts at Intel, F23 embarked on a period of examination and experimentation to identify a specific inventory control mechanism during the second quarter of 2006. It is relevant to briefly review the various concepts in so far as it shows the extent to which F23 treated this effort as a true learning exercise, role modeling the basic tenets of mX.

3.1: Original Wafer Starts Process

To provide the proper context in which to understand the F23 Wafer Starts Protocol it is useful to examine how wafer starts are typically determined for Intel fabs. At a high level the process involves the planning and manufacturing organizations matching customer demand with manufacturing capacities. Manufacturing and Planning determine the desired output from the fab and then back calculate the wafer starts required to achieve the desired output using the expected yield and the assumed standard cycle time. After that point, Intel's manufacturing system concentrates on maintaining its wafer starts, yield, and standard cycle time per the agreement with Planning.

From their agreement with Planning, each wafer fabrication facility has a schedule for its required wafer starts per week. Each fab then determines their standard wafer starts per shift by equally distributing the wafer starts per week over the total number of shifts each week. In general practice, the standard wafers starts per shift are then started each shift regardless of the particular status of the factory. Some exceptions occur when extreme circumstances exist in the factory (e.g. the fab is experiencing dramatic tool down times or has incurred a major excursion event).

For completeness, it should be noted that some fabs employ a "carryover" process to accommodate manufacturing steps near the start of their manufacturing process that can only process wafers in certain increments (e.g. 100 or 200 wafers). The standard wafer starts per shift are then adjusted up (or down) to match the specific increments of the manufacturing step and the decremented (or additional) starts are carried over to future shifts, until an increment can be achieved that matches the downstream process capacity. On the surface this practice makes it appear as though some fabs adjust their starts, but they are only adjusting starts to accommodate capacity of an early constraint and on the whole, they start their predetermined amount.

3.2: F23's Experimentation

F23 examined a number of concepts to identify its wafer starts protocol, some of which were actually tested in production. This exploration included the idea of using the WIP inventory present at constraint operations as the primary factor for making overall factory inventory decisions. A concept to implement a "supermarket" between the front end and back end of the line was also considered[b]. Segment pace[c] and depletion/replenishment rates were considered as alternatives to inventory. Weighting the relative importance for inventory towards the front of the line as greater than inventory towards the end of the line was also contemplated. Through the examination of this range of concepts, it was determined that there was a myriad of ways to vary inventory, each of which with its own level of complexity. In addition, various levels of

[b] In semiconductor manufacturing, the term "front end" of the line is often used as a reference to the operations in fab that are before the addition of metal layers. "Back end" of the line refers to the operations associated with metal layers.

[c] Intel refers to portions of its manufacturing line by "segments". Each segment of the line is associated with a time period of 1 week of the standard cycle time for a lot through the fab. For example, the operations associated Segment 2 are the operations that a lot would see during its second week in the factory.

analysis can be used to evaluate each proposal. However, the results of such analysis are subject to the inherent challenges in modeling an entire working fab. After much discussion, it was determined that the best course of action was experimentation in actual production. In addition, as discussed in more detail in the Contributing Factors section of this paper, the time constraints F23 was facing warranted something that could be quickly implemented. Therefore a simple, straight-forward methodology that could be implemented in production and then modified as needed was identified as the goal for the effort. The result of this effort was the F23 Wafer Starts Protocol.

3.3: F23 Wafer Starts Protocol

At the core of F23's wafer starts protocol is an application of Little's Law. For a given target cycle time and output rate, F23 identifies a target WIP inventory. Little's Law states that

$$L = \lambda W$$

where L is inventory, λ is fab output rate, and W is cycle time. For a given fab output rate λ and a specific target cycle time W_{target}, the target inventory level is

$$L_{target} = \lambda W_{target}.$$

With a target inventory identified, the next step is to vary the standard wafer starts amount each shift to achieve the desired inventory level. The standard wafer starts amount is considered to be the steady state amount of wafer starts needed to maintain the

desired steady state fab output rate. The standard wafer starts per shift $WS_{standard}$ can be found by dividing the weekly fab output rate λ by the number of factory shifts each week:

$$WS_{standard} = \frac{\lambda}{\text{number of shifts}}$$

The variation of wafer starts is accomplished by applying two adjustments to the standard wafer starts amount. The first order response is based upon overall fab inventory level and, in general, is larger than the second order response which is based upon inventory levels at those operations identified as constraints.

3.3.1: First Order Response: Overall Fab Inventory

The first order response is determined by comparing the current WIP inventory level of the overall fab to several levels of control limits established around the target level for fab WIP inventory (the overall fab inventory needed to achieve the cycle time goal). The greater the deviation from the target level the greater the magnitude of the first order adjustment on the number of wafers to start.

3.3.2: Second Order Response: Constraint Inventory

The second order response is determined by comparing the current WIP inventory level of the operations identified as constraints to several levels of control limits established for the target WIP inventory for those operations. The greater the deviation from the target level the greater the magnitude of the second order adjustment on wafer starts.

The adjustment amounts for both the first order and second order responses are described in Table 1. The first order response adjustment involves a percentage adjustment, either up or down, of the standard wafer starts amount. The second order response applies an adjustment, up or down, of a lot or two of wafers.

Table 1: First and Second Order Response

Fab Inventory	Low more than 1% low	Medium Low 1% to 0.5% low	OK 0.5% low to 1% high	Medium High 1% to 3% high	High more than 3% high
First Order Response	30% increase	15% increase	0% change	15% reduction	50% reduction
Constraint Inventory	Low more than 10% low	Medium Low 10% to 3% low	OK 3% low to 15% high	Medium High 15% to 45% high	High more than 45% high
Second Order Response	2 lot increase	1 lot increase	No Change	1 lot reduction	2 lot reduction

Once both the first and second order adjustments are determined, they are applied to the standard wafer starts amount.

$$WS_{adjusted} = WS_{standard} + ADJ_{firstorder} + ADJ_{secondorder}$$

3.3.3: Carryover

With the first and second order adjustments applied, the next step is to apply carryover from the previous shift. Carryover arises because wafers can only be started in increments of 100 in order to optimize loadings for a diffusion operation that occurs very early in the production process. Therefore it is necessary to either round up or round

down the amount of wafers started to the nearest increment of 100. Once the carryover from the previous shift is applied, the wafer starts amount is adjusted to the nearest 100 with the difference "carried over" to the next shift.

$$WS_{rounded} = Round(WS_{adjusted} + Carryover_{previous})$$

where the amount $WS_{adjusted} + Carryover_{previous}$ is rounded to the nearest 100 and

$$Carryover_{next} = (WS_{adjusted} + Carryover_{previous}) - WS_{rounded}).$$

The process is repeated for every shift using the current fab inventory and constraint operations inventories, the carryover from the previous shift, and the carryover for the next shift.

3.3.4: Example Calculation

To better illustrate the Wafer Starts Protocol. Assume a wafer fabrication facility that has

- standard wafers starts per shift =150 wafers
- current fab inventory = 9950 wafers
- target fab inventory = 10000 wafers
- constraint operation target inventory = 300 wafers
- current constraint operations inventory = 500 wafers

- there is an early operation in the line that can only process 100 wafers at a time
- each lot contains 25 wafers

First Order Response

The current fab inventory is 0.5% below the target. From Table 1 the first order adjustment is to increase the standard wafer starts per shift by 15% or 23 wafers.

Second Order Response

The current constraint operations are 66% over their inventory target. From Table 1 the second order adjustment is to decrease the standard wafer starts amount by 2 lots (50 wafers).

Result

The protocol identifies a wafer starts amount of 155 wafers (150+23-50). Next the carryover must be considered. Since the downstream can only process lots in 100 wafer increments, the 155 amount must be rounded to the nearest 100 increment. In this case, the rounded amount would be 200 wafers. Therefore 200 wafers would actually be started into manufacturing and a negative carryover (-45 wafers) would be "carried over" to the wafer starts calculation for the next shift.

4: Results

In the second quarter of 2006, F23 began experimenting with different methods of inventory control. In WW33[d], the Wafer Starts Protocol as described above was implemented in production. The data shown in the following sections were collected using Intel's existing data systems for inventory and cycle time measurement.

4.1: WIP Inventory

The results of F23's effort to control WIP inventory are shown in Figure 1. The graph shows the overall fab inventory relative to the target level. The implementation of the Wafer Starts Protocol in WW33 resulted in specific inventory targets based upon fab cycle time goals. The four target level's that were pursued throughout the remainder of 2006 are labeled on the graph. The wafer starts protocol's efficacy in adjusting the fabs inventory is clearly shown. Overall, the fab inventory was reduced by 44% from the overall inventory high in WW22 to the end of year level. It must be noted that this inventory reduction was achieved while F23 was increasing factory output.

In Figure 2, the wafer starts and factory output are shown for Q3 through Q4 of 2006. The graph also shows the resulting inventory and planned fab capacity for comparison purposes. It can be seen where wafer starts reduced around WW33, WW39, WW43, and WW45 as protocol transitioned to new inventory targets. The graph also shows that F23 was successful in increasing factory output while reducing inventory and did not experience any negative effect in output. Also from the graph it can be seen that F23 was operating at its full capacity during this same period.

[d] Internally, Intel uses a work week (WW) calendar to indicate dates. The first week of the year is given the moniker WW01 while the last week of the year is described by WW52. WW33 is the 33rd week of the year.

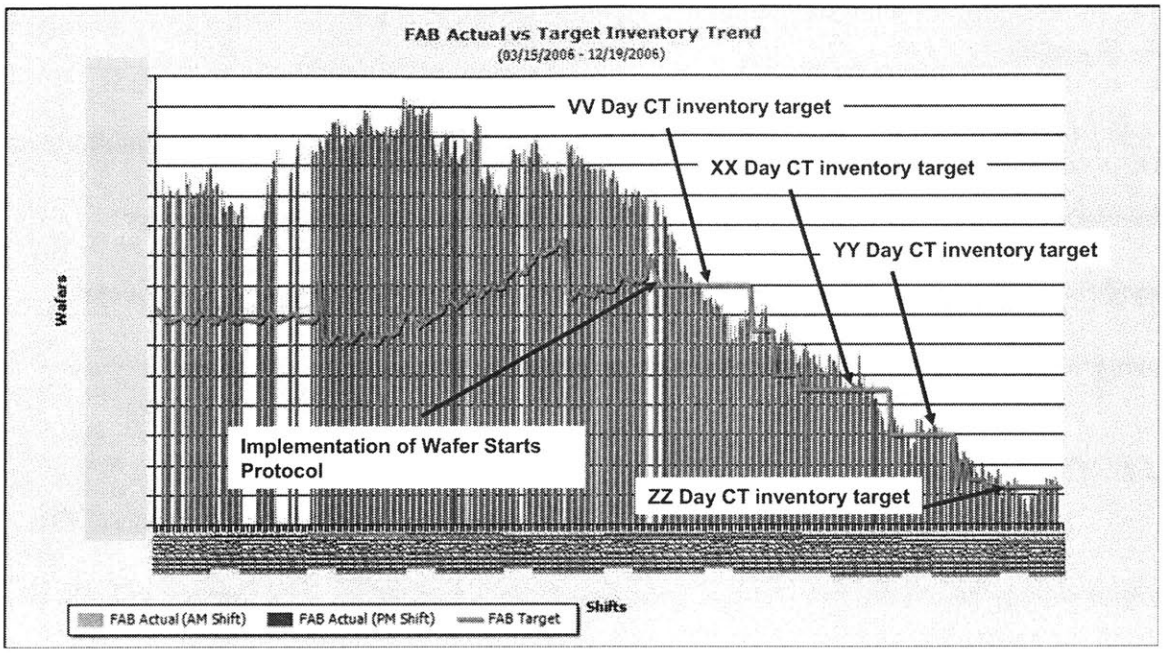

Figure 1: Fab Inventory

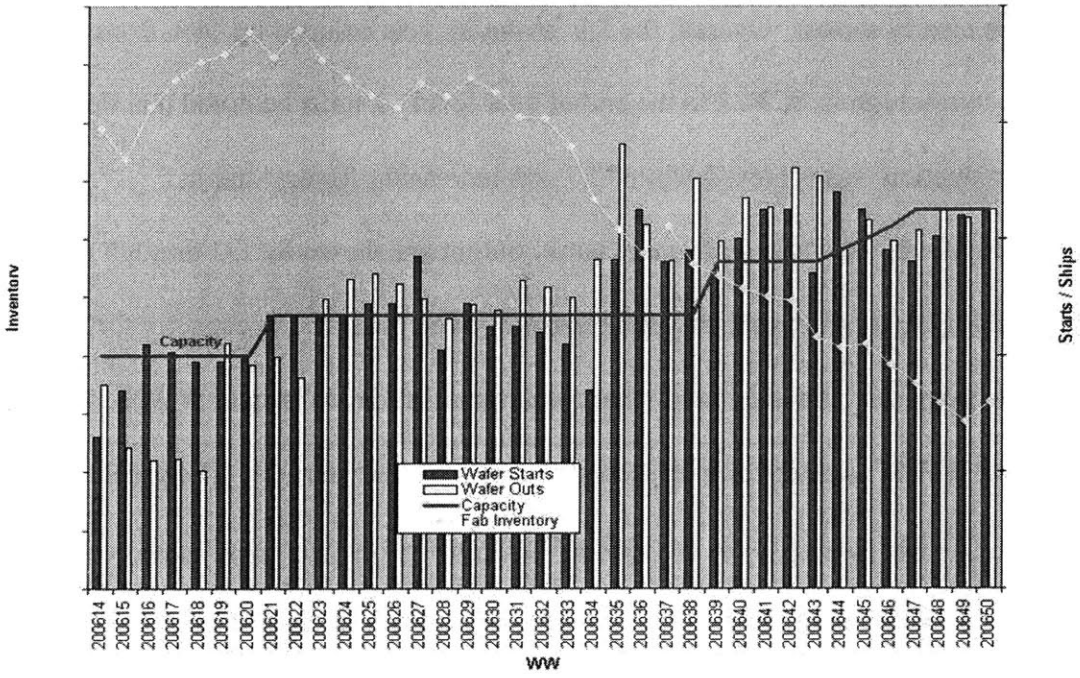

Figure 2: Fab Inventory vs. Wafer Starts / Factory Output

4.2: Cycle Time

Figure 3 displays the fab's dynamic cycle time performance from WW12 through WW50. At Intel dynamic cycle time is calculated by determining the average time lots spend at each operation for a given period of time, and then summing the average times at each operation for all operations in the overall manufacturing process. From WW23 to WW50, the cycle time was reduced by over 61%. The WW50 cycle time is consistent with a cycle time performance, in terms of days per mask layer, near semiconductor industry records. In Leachman's paper describing the SLIM WIP management policy developed at Samsung[15], a performance of 1.3 to 1.6 days per mask layer was reported. Such a performance is considered to be near the best in the industry.

Comparing Figure 1 and Figure 3, it can be seen that F23 successfully achieved each of the cycle time goals set by the wafer starts protocol. The fact that F23 set inventory targets based upon cycle time goals, reduced the inventory to those target levels, and achieved the cycle time goals, indicates that Little's Law holds in practical application despite the complexities of real world semiconductor fabs.

We note that anomaly conditions affected factory cycle time during WW37 and WW43. In WW37, the factory experienced an excursion requiring the factory to stop production, and factory had to cease production in WW43 because of blizzard which made it unsafe for worker to travel to the factory. However these anomaly data points show that the factory was able to resume its improvement even after being severely impacted by external events. Reduced factory inventory did not affect the factories ability to resume production and continue further cycle time reduction.

Figure 4 displays the cumulative distribution of lot static cycle times for WW33 through WW45. Static cycle time is the total time that a particular lot was in the fab. It

is measured from the time a lot is introduced into the factory to the time that it exits. The data for Figure 4 is grouped by work week. For each work week, the static cycle time for each lot exiting the fab that week is displayed. The median lot cycle time is displayed on the graph by the 50% demarcation. This graph clearly shows that static cycle times decreased from WW33 to WW45. In fact the median lot static cycle time dropped 49% during that time period.

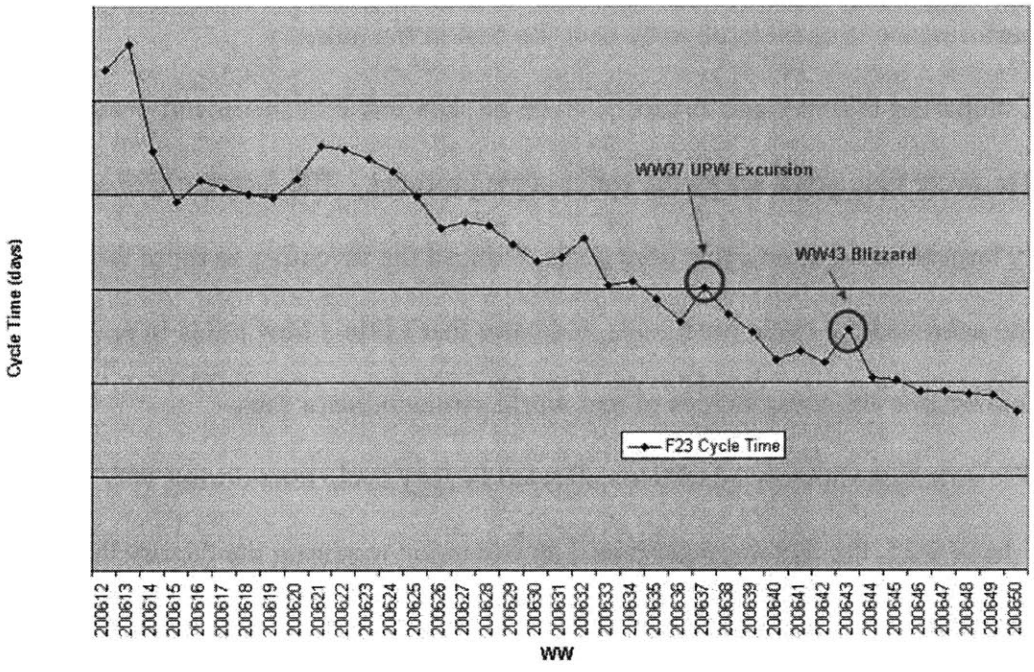

Figure 3: Fab Dynamic Cycle Time

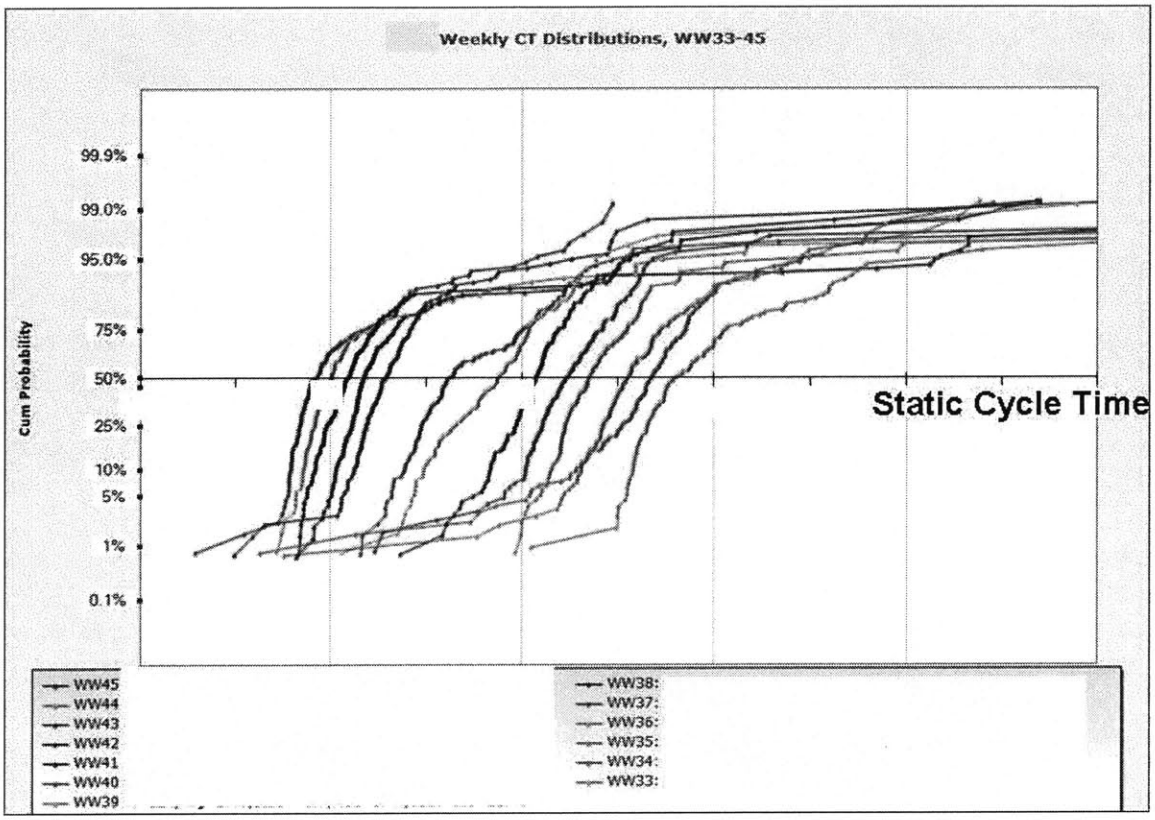

Figure 4: Static Cycle Time

4.3: WIP Turns

F23's WIP turn performance is shown in Figure 5. The WIP turn metric measures the relative turnover of inventory within the factory as a function of activities. An activity is defined as any operation that changes the physical characteristics of the wafer. WIP turns are often used as a relative comparison of factory speed at Intel rather than cycle time. This is because cycle times vary between process technologies due to the inherent differences in raw process times between process technologies and, therefore, cannot be compared directly. The WIP turn metric is defined as

$$WT = \frac{Activities}{INV}$$

where

Activities = the total number of activities performed in the factory for a given period of time

and

INV= the average factory inventory during that period of time.

The data indicates that F23 steadily increased its WIP turns from WW32 to WW50, consistent with the cycle time data of Figure 3. Of particular note, F23 was able to maintain its WIP turns in excess of 3.0 for 10 weeks in Q4. In addition, F23's WIP turns were approximately 3.5 or greater for 7 weeks. These WIP turn results are significant compared to other Intel fabs and show that an Intel facility can exceed the 3.0 WIP turn barrier that is often considered to be indicative of an "extremely fast" fab.

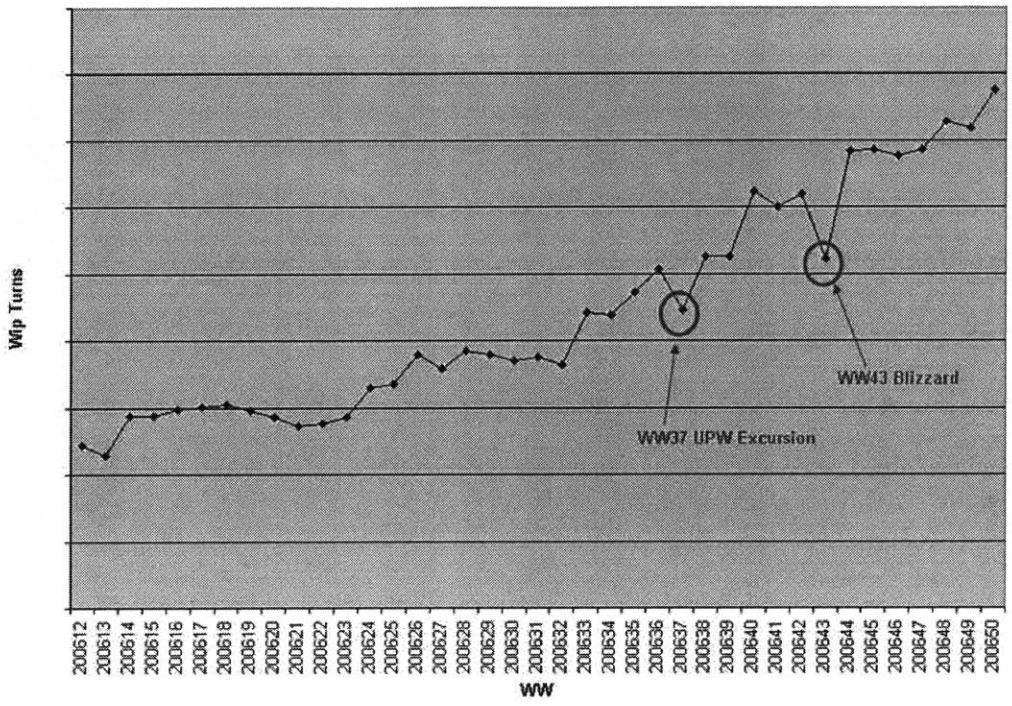

Figure 5: Fab WIP Turns

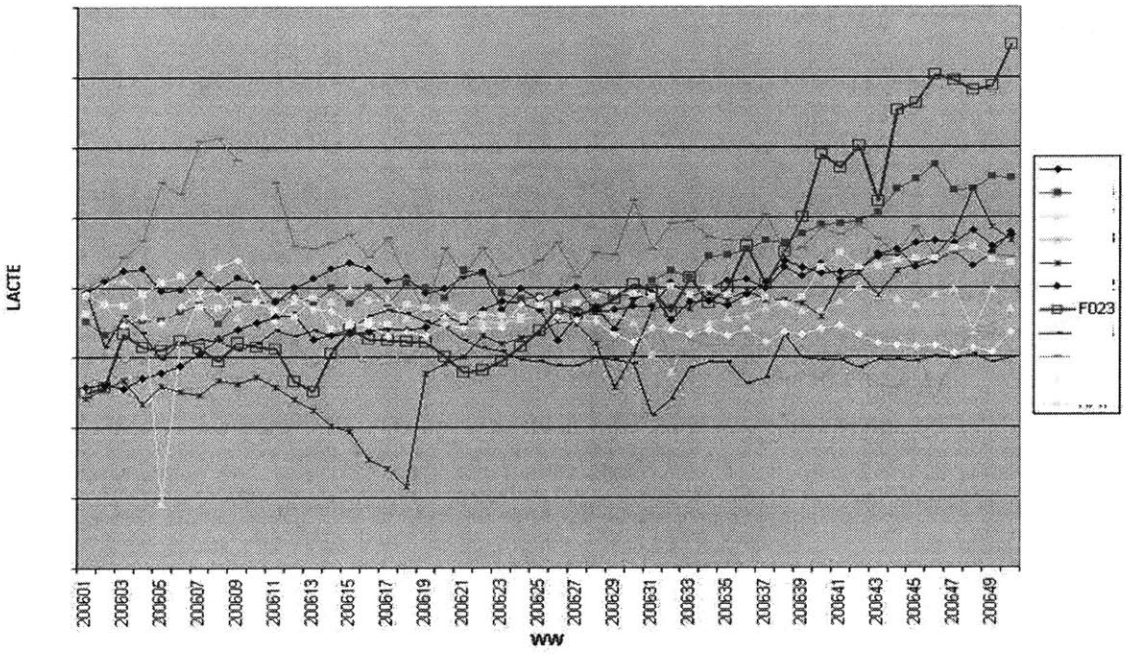

Figure 6: LACTE for all Intel Fabs

4.4: Load Adjust Cycle Time Efficiency

Figure 6 displays LACTE data for Intel fabs. LACTE, Load Adjusted Cycle Time Efficiency, provides a relative comparison of cycle time performance between different fabs since it incorporates the specific raw process times for each fabs process technologies as well as the fabs loading relative to its total capacity. LACTE is defined as

$$LACTE = \frac{L}{C} \cdot \frac{RPT}{CT}$$

where

L = Factory Loading as measured by actual wafer starts per week,

C = Factory capacity as measured by possible wafer starts per week,

RPT = Raw Processing Time for one lot,

and

CT = Dynamic Fab Cycle Time.

As show in Figure 6, F23 LACTE performance was below average for the first half of 2006, when compared to other Intel fabs. Then F23's improved from worst in the Virtual Factory in WW21 to best-in-class by WW50. In WW50, F23 exceeded the nearest fab by approximately 12 percentage points. This fact is significant since several

fabs have comparable tool performance to F23 and have similar HPM programs, yet only F23 is practicing active, aggressive inventory management.

4.5: Activities

In conversations with Intel personnel who were skeptical of F23's improvement, it was posited that perhaps F23 was focusing on short-term success with cycle time while sacrificing longer term factory performance. Their thought process was that at some point F23 would "starve" its manufacturing line because of its decreasing inventory and eventually see a drop in output. Figure 2 indicates that F23 was able to maintain its factory output for more than two quarters while lowering its inventory. It could be argued that sustained factory output over two quarters suggests effective overall line performance for the entire fab, but critics might argue that factory output alone does not provide confirmation that the entire factory is maintaining its performance.

Within Intel, activities are used as an indicator of pace of the entire manufacturing line. Figure 7 displays the weekly activities totals with the corresponding fab WIP inventories for WW27 through WW50. The graph clearly shows that during this time period F23 maintained its activity performance while its inventory decreased more than 40%; the only anomalies were in WW37 and WW43 due to the corresponding excursion and blizzard. In combination, Figure 2 and Figure 7 indicate that F23 was able to lower its inventory and not "starve" its line. This suggests that F23 results are sustainable. In fact Figure 7 shows that the activities per week actually increased slightly during this period. This suggests that the combination of F23's HPM, Ue, and mX initiatives have

allowed F23 to increase its capacity; the combination of initiatives has increased the "operational cadence" of the factory.

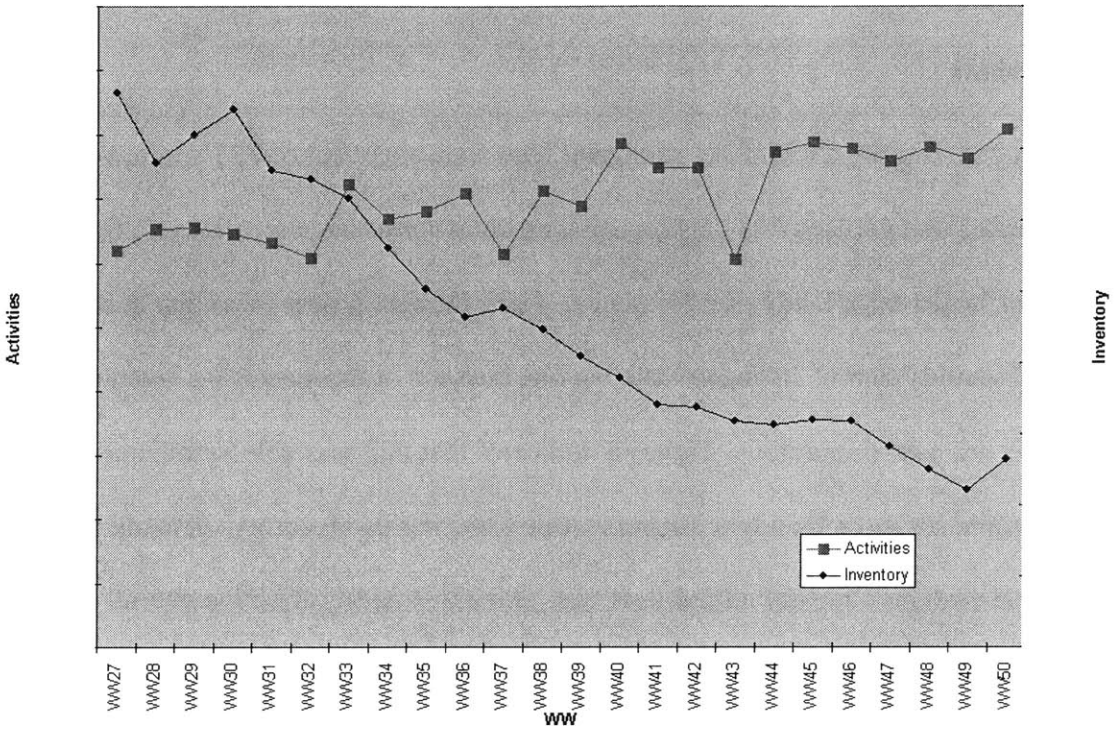

Figure 7: F23 Weekly Activities and Inventory

4.6: Cost and Information Turns

Another result of F23's cycle time improvement was a reduction in cost. F23's cost per wafer decreased dramatically from the beginning of Q3 to the end of Q4. This improvement could be attributed to tool performance improvement and accounting adjustments for reduced inventory. However, F23's process engineering group leaders attribute the reduction to another aspect of the improvement effort: the speed with which they were able to collect data. With F23's cycle time improvement, process engineers were able to conduct experiments and collect the data even faster. This faster

information turn allowed engineers to more quickly improve process capability. The improved process capability had a direct impact on scrap and defect rates, and therefore, reduced F23's costs.

At the outset of the effort to control inventory, the explicit goal was to reduce factory WIP inventory and factory cycle time. Such improvements obviously improve a factory's ability to respond to market changes due to the reduced cycle time and lower the factory costs through lower inventory. The more subtle result of increasing information turns also directly impacts the factory's bottom line. In addition, F23's increased rate of gathering information also provides F23 with enhanced flexibility to accommodate future changes no matter what their impetus. The impact of faster information turns further supports the premise that faster factory cycle time is better for business, even if standard accounting procedures struggle to highlight the importance of fast cycle time.

5: Critical WIP Ratio

During the course of the F23's inventory reduction efforts, the fab inventory was monitored closely, and the actual fab cycle time was compared with predictions. Throughout this comparison and data tracking, we developed the concept of comparing the current inventory to a theoretical minimum inventory. This comparison provided a measure of the factory's performance in reducing waste (inventory is considered a waste within lean manufacturing methodologies[16]).

The theoretical minimum fab inventory is found using Little's Law with the technology's raw process time and the factory's output level:

$$Inv_{theoretical} = RPT \cdot Output$$

Where

RPT = raw process time for one lot (each lot typically contains 25 wafers)

and

Output = the target fab output.

The comparison of actual inventory to a theoretical minimum suggests a natural metric: the ratio of the two values. This ratio is called the critical WIP ratio:

$$CWR = \frac{Inv_{actual}}{Inv_{theoretical}}$$

where

Inv_{actual} = the current fab inventory

and

$Inv_{theoretical}$ = the theoretical minimum fab inventory.

Figure 8 shows the F23's CWR compared with the actual fab cycle time. We see that as CWR was reduced so too was cycle time.

Any value of CWR greater than 1.0 suggests that the factory has excess inventory. Some amount of excess inventory is unavoidable no matter how much waste reduction is achieved. Inventory is needed to buffer against variability, and additional cycle time, and therefore inventory, must be expected due to internal transportation, set-up times, and other non-value-added activities. The measurement of CWR also provides a clear indication of how the factory is performing relative to the ideal state. The communication of an ideal state is a key aspect of any lean implementation according to Flinchbaugh and Carlino.[17]

What is considered to be world class? Since lean manufacturing is somewhat new to the semiconductor industry, it may be necessary to look at another industry to identify what is truly best-in-class for manufacturers. In a conversation with consultants from the Lean Learning Center in Novi, Michigan, it was suggested that a Toyota manufacturing plants may operate with CWR in the range of 1.5 to 2.0. Though no specific data was obtained to validate this estimate, it does provide a framework in which to judge the success of waste elimination (inventory reduction) at Intel fabs. However, it must be noted that such comparisons between Intel and external companies can be significantly affected by how the ideal state is defined. We have elected to identify wafers as our unit of measure for CWR. However Intel still "batches" wafers in groups of twenty-five to form lots. In addition, each wafer is essentially a batch itself since it contains group of individual products. Toyota eschews batching where possible and identifies individual vehicles, i.e. individual products, as their unit of measure. This must be taken into account if quantitative comparisons are desired.

Critical WIP Ratio also provides a relative comparison amongst different fabs since it incorporates the specific raw process times and output capacity of each factory. The findings of a preliminary comparison of CWR amongst Intel fabs are consistent with the results observed at F23: the lower the CWR, the better the cycle time performance. This initial comparison also showed that other Intel fabs are currently operating at a higher CWR than F23, a result that may suggest that other Intel fabs have opportunity for cycle time improvement through inventory control. The comparison of CWR's also showed that F23's inventory level in WW22 relative to its theoretical minimum level is

not an abnormal state for Intel factories. Several factories had CWR's between 4.0 and 5.5.

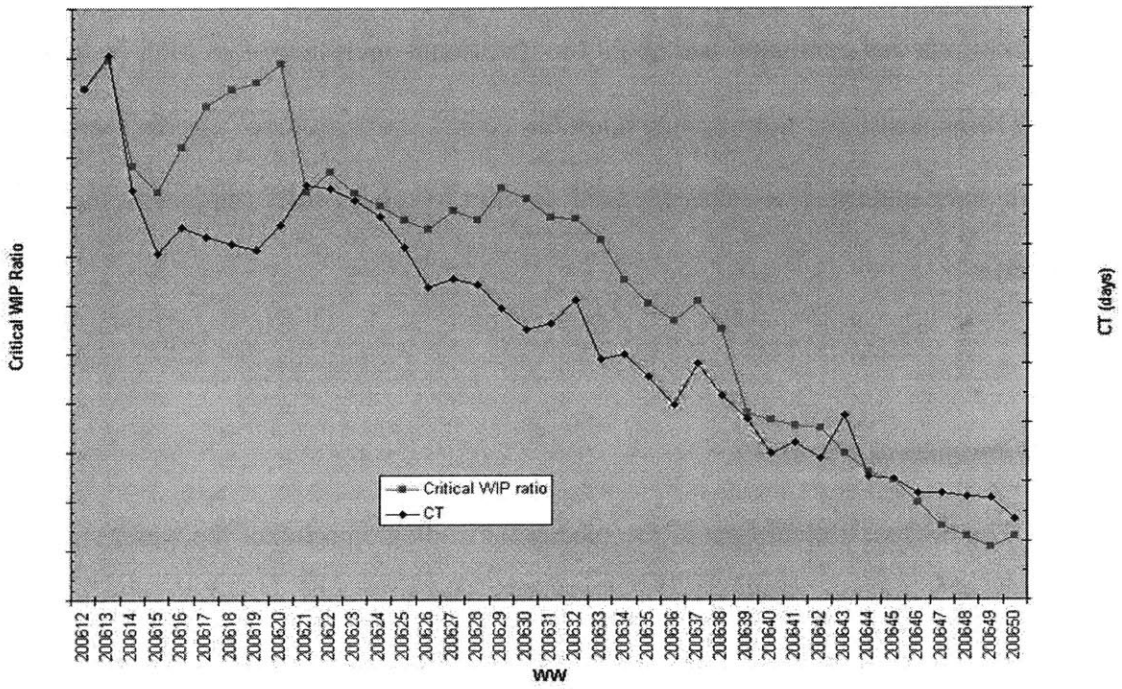

Figure 8: Critical WIP Ratio vs. Fab Cycle Time

6: Challenges

Several challenges were encountered during the implementation and use of the wafers starts protocol. These challenges range from technical issues with the model to preconceptions that personnel had about the traditional operation of an Intel factory. Though these issues did not impinge upon the wafers starts protocol's performance, these issues do exist and must be acknowledged in order to successfully implement the protocol.

6.1: Philosophical Hurdles

The biggest impediment to the successful implementation of the wafer starts was the mental hurdle of modifying how a typical Intel factory operates. The controlled variation of wafer starts represented a paradigm shift in Intel's operational thinking. Intel's internal supply chain has been traditionally starts-oriented. Divisional product groups and manufacturing organizations agree to production schedules based upon wafer starts.

In using a starts-based methodology, the assumption by all parties during this planning exercise is that the material that was started on a certain date will exit the factory in a standard number of weeks. This standard time is referred to as the Plan of Record (POR)[e] fab cycle time. Such an assumption may or may not be correct since it does not explicitly incorporate the variability of actual production. For example, deviations in tool performance could cause yield or cycle time issues that would cause

[e] Within Intel Plan of Record refers to the established baseline often with proven historical results or formal approval.

fewer wafers to exit the factory in the POR cycle time. Nevertheless, Intel's internal supply is based upon wafer starts, with the assumption that a certain number of wafer starts will result in a certain number of wafer outputs in a set period time.

In addition, Divisional product organizations are able to track the performance of their agreement with Manufacturing. They are able to observe exactly how many wafers were started by Manufacturing, and they will inquire with Manufacturing if they notice any deviations from the planned amount of wafer starts. This direct visibility into the manufacturing pipeline does not allow the manufacturing organization much flexibility in responding to variability in their production operations.

To implement an active inventory control like wafer starts protocol, the focus must shift from a wafer starts perspective to factory outs (i.e. wafers that exit the factory) perspective. Manufacturing and divisional groups must commit to a certain output since manufacturing may vary starts to control inventory. Manufacturing must then endeavor to hit the committed output to meet the divisional product group's demand. This change in perspective may seem rather straight forward, but changing the long-standing perspectives of a large, successful, rather complex organization is not without challenges.

In addition to this factory performance perspective, there were internal perspectives that had to be modified to provide the proper support. Several factory personnel and group leaders initially expressed concern about potential issues or pitfalls associated with reducing factory inventory. The most common concern was the belief that if the factory reduced its inventory it may not be able to deliver its committed output. Perhaps the general assumption underlying this concern was that the current factory inventory was set at a specific level in order to sufficiently buffer against the current tool

variability. Such an assumption is consistent with the mindset that the current way of doing things is the best or the result of some explicit design or forethought to make possible the current level of performance. Hammer[18] specifies that identifying and defying a key organizational assumption is an essential aspect of achieving true operational innovation. In particular, he states that an ideal assumption to challenge is one that interferes with achieving a strategic objective. High inventory does exactly that; the assumption that the current inventory levels are needed to buffer against tool variability obviates Intel's cycle time improvement efforts. To overcome this perspective, F23's factory management had to demonstrate firm support for the effort and, essentially, role-model the "experiment to find improvement" rule of their mX philosophy.

Additionally, factory personnel expressed concern over the increased frequency of set-ups or changeovers that would result due to lower inventory. This is a concern that would result from any sort of cycle time improvement. However since these personnel were unaccustomed to cycle time improvement, this was yet another mental hurdle that had to be overcome in order to gain everyone's support. F23's management accomplished this by pursuing a temporary solution of dedicating tools and by acknowledging that an increased number of set-ups would be required but communicating this would be a natural result of improving factory cycle time. This concern also highlights a future opportunity for improvement efforts since the traditional Intel perspective has been to avoid set-ups or changeovers. Because of that perspective, there exists potential improvement opportunities to be had by beginning the process of addressing set-up times.

6.2: Technical Issues

A key aspect of the current F23 wafer starts protocol is the identification of the constraint operations or collection of constraint operations. This is a standard concept in most operations philosophies but in actual practice, at least in an Intel fab, this activity is less tractable. In general, the general impression is that the lithography tools represent the constraint operations in the fab. However, due to a range of issues, this is not necessarily the case. The issue is complicated further by the various manners in which constraints are typically identified. Whether using planned tool capacity (based upon anticipated tool performance and availability), actual tool performance (based upon historical run rates and availability), or tracking problematic tools (based on manufacturing's tracking of problematic tools each shift), the actual identification of specific constraints is non-trivial. Often due to the variability associated with tool performance, it is not possible to identify a single constraint tool set for the factory. There can sometimes be a range of toolsets that are "constraint" or "near-constraint" tool sets.

F23 addressed this problem by identifying a range of operations as the constraint. At times this might include an entire segment week as the constraint (a segment week represents the operations associated with a specific week of the standard fab cycle time). Occasionally other operations would be added to this collection of constraints, typically operations associated with the previous or following segment weeks. Also, the range of operations had to be reviewed as F23 adjusted to new inventory levels and target fab cycle times, thus changing the operations associated with specific segment weeks. There were also situations where atypical events occurred, and, rather than modify the starts

protocol for unusual events, the wafer starts protocol was briefly manually overridden to modify the starts response. An example of such an event is when a tool set, that typically has very high availability and therefore there is a small number of tools in this particular fleet, goes down for an extended period of time.

Another technical issue that occurred and was difficult for the model to comprehend was removal of WI P from operations in the middle of the line. This event occurred when WIP was put on hold due to changes in Divisional demand. In such situations, the Division would prefer to halt the production of specific products until they had more clarification from their customers. This would result in material being temporarily removed from the line and placed in stores until some time, identified by Division, when this material could be re-released into the line. To the wafer starts protocol, this would appear like a sudden reduction in fab inventory when the material was placed into storage, or a sudden increase in the fab inventory when the material was re-released into the line. To avoid sudden fluctuations in the wafer starts response, F23 would try to anticipate Division's direction and manually override wafer starts amounts in order to smooth the transitions between these conditions. The placement of WIP on hold by Division was a rather rare event and is nearly impossible to anticipate. Therefore, permanent modifications to the starts protocol were not made to respond to such events.

7: Lessons Learned

Several lessons learned were identified as a result of the initial implementation and experimentation with the wafer starts protocol.

7.1: Impact of Inventory Control

Active and aggressive work-in-process inventory control can have a significant impact on fab cycle time. F23 achieved large cycle time improvements after implementing the wafer starts protocol. The standard belief within Intel is that in order to improve cycle time significant improvements in tool performance have to be made. In general that concept is true. However, a more detailed exploration begs the question of under what operating regime does that statement hold true? In F23's case, the factory was significantly over inventory, and, as a result, significant cycle time improvement could be achieved without substantial improvement in tool performance.

7.2: Perception of Factory Loading

Another conclusion of the wafer starts team was that starts-based inventory goaling methods can provide incorrect perceptions of factor loadings. F23 had previously examined their loading based upon whether or not they had started their predefined starts amounts. As a result, F23 was operating in an operating regime similar to other Intel fabs, and therefore the assumption was that the factory was properly loaded. Occasionally, there would be concerns over inventory if unusual conditions arose but, for the most part, the inventory level of the factory was simply accepted as an outcome that resulted from the factory's conditions.

As the team began to a focus on output rather than starts, this perception changed. The team began tracking the fab's inventory level using the Critical WIP Ratio, and this measure showed how much the fab was over its theoretical minimum inventory. When F23 implemented the wafer starts protocol, its Critical WIP ratio was 5 or greater, similar to other Intel factories during the same time period. F23 transitioned to a critical WIP ratio of less than 3 over the course of Q3 and Q4 of 2006. The wafer starts team was struck by how "easy" it was for F23 to reduce the inventory over this range of critical WIP ratio while maintaining factory output. This ease in reduction further suggests that the perception that the original inventory levels were appropriate was incorrect, and that, in fact, the factory was overloaded with inventory.

The perception that the original inventory was appropriate may have been based on the assumption that it was needed to buffer against tool variability. Though inventory is needed to buffer against tool variability, F23's results suggest that its buffer inventory was excessive and was not at a level predefined by any calculation based upon tool variability. These findings have changed the perspectives of F23's personnel. Now the perception exists that the fab was greatly over inventory at the outset of wafer starts protocol implementation.

7.3: Little's Law Holds True

At the outset of the implementation of the wafers starts protocol, team members discussed how the factory may respond to the reduction of inventory. Though the concept of Little's Law was fully understood, the team was not certain if it would hold true in actual practice. In fact, in conversations with various personnel, it was suggested that Little's Law, though theoretically correct, would not accurately apply to a real

system such as a fab with all of its complexity and variability. Counter to this supposition, the team found that Little's Law did hold true. The team was able to successfully set inventory targets based upon cycle time goals, reduce the inventory to that level, and achieve the desire cycle time goal. These results were achieved at each of the four major inventory targets set during Q3 and Q4. In retrospect, this finding may seem self-evident, but at the moment of implementation, these results were not so certain.

8: Next Steps

Further work with inventory control using the Wafer Starts Protocol is planned. This work involves refinement of the mechanisms employed by the protocol. In particular, an examination of the use of segment pace to trigger variation in starts is a potential area of exploration. In addition, a continued review of factory constraints and constraint identification is warranted. The former being necessary as the factory transitions to new operating regimes potentially resulting in different tools becoming problematic. Another item of focus as the operating conditions change is the control limits of the protocol's responses. As the factory transitions to new operating regimes with significantly different inventory levels, it is necessary to ensure that the first order and second order responses are properly set to the new inventory levels in order to ensure the proper protocol response. The proliferation of the wafer starts protocol to other Intel sites may also require additional standardization and documentation.

In addition to the technical aspects of the wafer starts protocol, F23's improvement effort has shown that some changes in Intel's way of conducting business may be required. In some cases the changes may be significant enough to represent a paradigm shift. One example is the extent to which Intel's internal supply chain uses wafer starts as a metric, a potential issue when varying wafer starts is used as the mechanism for inventory control. The transition from using wafer starts to using wafer outs as the metric does not seem terribly difficult, but in actual practice such a change is not very easy for a complex arrangement of systems and long-standing practices.

In general, factory personnel must continue to "learn" how to conduct day to day activities in the new operating regime associated with very fast cycle time. For example,

the expectations and norms for responsiveness to issues must adapt to the new environment. With very fast cycle time and low factory inventory, it is imperative for issue resolution to be very fast. Solving problems quickly is nothing new but as factory cycle time drops significantly, the need for an increased "cadence" of activities is critical as there is not sufficient inventory to support extended down time. In fact, there will likely be other activities that also need to adapt to the new pace and increase their cadence in order to properly support fast factory operations.

We should note that F23's continued implementation of cycle time improvement and inventory reduction will likely uncover additional issues going forward. In lean manufacturing methodology, this is often described by the analogy of inventory to water. As inventory is reduced, similar to the lowering of the water line, new issues or "rocks" are uncovered. The central concept of lean manufacturing is to then eliminate those issues, removing the rocks, and then continuing the inventory reduction, lowering the water line, ever looking to remove future problems and support future cycle time improvements. This concept was reinforced with F23 personnel in order to properly manage expectations regarding the challenges that would inevitably be incurred as a result of lowering inventory. Figure 9 shows the graphic that was shared with operators, managers, and senior management to convey this concept. Going forward F23 will need to continue pursuing all of its initiatives for cycle time improvement. Further improvements will likely be even more difficult to achieve and will require the contributions of all efforts.

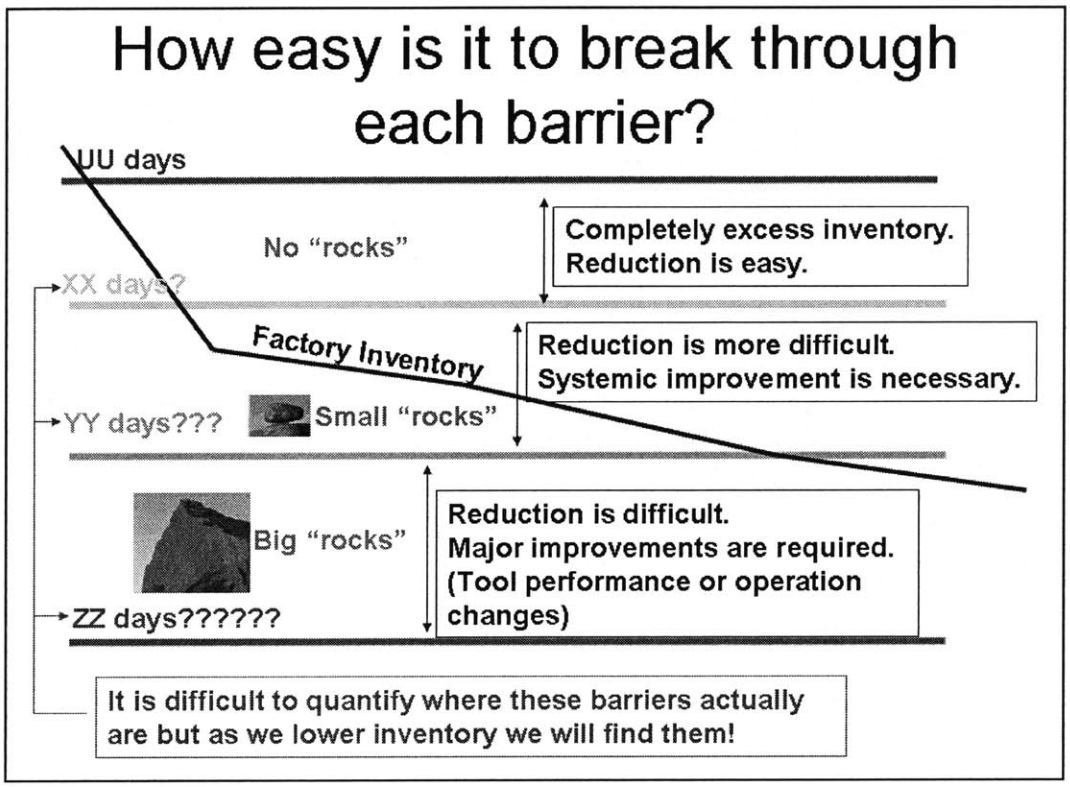

Figure 9: Uncovering the Rocks

Finally, as mentioned above, the critical WIP ratio comparisons of various fabs suggest that some other Intel factories may benefit from inventory control. Therefore the proliferation of the wafer starts protocol may be beneficial. Additional examination of this hypothesis should be performed, and if warranted, the wafer starts protocol should be introduced at other factories. We must recognize that such an endeavor will need to be done carefully as it is often difficult to transfer improvements or lessons learned between organizations, even within the same overall group. Hayes et al[19] discuss this common challenge. They attribute these difficulties to a number of issues: the costs of transferring the knowledge, the incentives systems, the degree of differences in operating

environment, and pride. Therefore any efforts to proliferate the F23 wafer starts protocol to other sites would need to be conducted with sensitivity to such potential issues.

9: Contributing Factors

So why was F23 so successful in making a dramatic change in their performance? To answer that question, one must consider the broader environment in which F23 existed during this time period. Several conditions existed at F23 that contributed to establishing an environment amenable to significant change. These conditions allowed F23 to overcome the social and political challenges that typically thwart significant change.

9.1: Momentum

First, F23 has been implementing mX for approximately 3 years. A majority of factory personnel have been exposed to lean concepts during actual implementations or through the internal mX course that introduces the concepts of lean manufacturing. Additionally, a number of the managers have attended week-long, lean manufacturing courses at the Lean Learning Center in Novi, Michigan. F23 has conducted a number of Kaizen events and the successes of these efforts have been documented and shared with all factory personnel. With this level of exposure to lean concepts, there is not a great deal of resistance to new ideas or new approaches. The factory personnel are fairly accustomed to the concepts of continuous improvement. It is quite possible that F23 has finally achieved a critical mass of change agents and level of understanding of implementing process improvements. Therefore, the conditions were finally right for the dramatic changes and improvements that occurred through aggressive inventory control.

9.2: Management Support

F23's continuous improvement efforts have also enjoyed the support of factory management. F23 has two manufacturing managers in charge of production at the facility. These co-managers are not only supportive of the F23's improvement efforts, but they have also served as change agents themselves. One of these managers has been one of the major proponents of mX within Intel, and the other manager has an extensive understanding of manufacturing science. This combination of backgrounds and knowledge of the potential benefits provided a willingness to explore inventory control.

In addition to the support of the manufacturing managers, F23 improvement efforts also have the support of the factory manager. F23's factory manager is a graduate of the Leaders For Manufacturing program at MIT. As a result, he has a firm understanding of lean manufacturing concepts and manufacturing science. But beyond just an understanding of the concepts and potential benefits, F23's factory manager also has a proven record of success within Intel. This reputation for delivering results provides a level of credibility that allows the factory manager to challenge long-standing Intel practices. As a result, F23's factory manager is able to provide a political umbrella under which F23's change agents can operate free from resistance from other high-level stakeholders. For example, the support of both the manufacturing managers and the factory manager was sufficient to overcome concerns raised by corporate planning organizations regarding deviation from planned wafer starts, a critical aspect of the wafer starts protocol.

9.3: Personal Ownership

Another condition that has allowed F23 to implement such dramatic changes is the fact that F23 is the only Intel factory producing its particular products or process technology. As a result, F23 has no other facilities that must approve changes. Typically Intel factories function as members of a network of remote facilities, known as the Virtual Factory (VF). The VF enforces Copy Exactly!, Intel's methodology for ensuring all factories that produce a certain product or technology do so in exactly the same manner. The VF accomplishes this through a series of approval forums and change control processes. Being the only factory producing its products, F23 is the only member of its particular VF, and therefore only needs the approval of its internal stakeholders to make changes. This situation has perhaps made F23's change efforts easier to implement since they did not require gaining the approval of external stakeholders, some of whom have had little exposure to lean manufacturing concepts.

9.4: Sense of Urgency

Finally, a major factor that has contributed to F23's willingness to implement significant changes in its operations is F23's status in Intel's overall manufacturing roadmap. Currently, F23 is not scheduled for production of products beyond the current process technology in production at the site. As a result, F23 is currently facing the risk of closure or possible sale. This situation creates a dramatic "burning platform" to motivate personnel to implement change. Since Intel is a merit-based company, the possibility exists to gain a position on the production roadmap by demonstrating better performance than other Intel facilities. F23 has decided to accomplish this by demonstrating the best cycle time performance within Intel, and therefore they have been

willing to accept potential risky proposals such as aggressive inventory control in order to achieve that performance and reserve themselves a position in Intel's long-term plans. They have named this initiative "Lead with Speed", a phrase now used in most of the site's employee communications and update meetings.

9.5: Right Combination

These factors address a number of the issues that Kotter[20] identifies as the most common causes for major change or improvement initiatives to be unsuccessful. According to Kotter, the eight common errors implementing change efforts are

1. Not establishing a great enough sense of urgency
2. Not creating a powerful enough guiding coalition
3. Lacking a vision
4. Under-communicating the vision by a factor of ten
5. Not removing obstacles to the new vision
6. Not systematically planning for and creating short-term wins
7. Declaring victory too soon
8. Not anchoring changes in the corporation culture.

F23's sense of urgency addresses error #1 while its management support of cycle time improvement addresses error #2. F23's development of "Lead with Speed" vision circumvents error #3. With F23's celebration of exceeding the 3.0 WIP turn mark and the constant communication describing how various efforts support the "Lead with Speed" initiative, F23 has effectively avoided errors #4 and #6. F23 Management stopping the planning organizations from interfering with the wafer starts protocol is an example of how F23 has avoided error #5. So far F23 has successfully addressed six of Kotter's

eight errors typically made during continuous improvement efforts. Because they are continuing to pursue lower and lower inventory, F23 is on their way to resolving the problems associated with errors #7 and #8.

The combination of these conditions conspired to provide F23 with an environment conducive to experimenting with new operational methodologies and challenging long-standing Intel operational assumptions. The recognition of these conditions further highlights the challenge associated with replicating the success of F23 at other Intel sites. It will be possible to transplant certain technical mechanisms, such as a wafer starts protocol, but the question remains whether it will be possible to transfer the "alchemy" associated with the entire set of F23 cycle time improvements and environmental factors.

10: Conclusions

F23 has several cycle time improvement initiatives underway. All of these efforts have contributed to the marked improvements that F23 demonstrated throughout 2006. F23 successfully reduced its cycle time by more than 60% and it consistently ran at WIP turns levels of greater than 3.0, a level often thought to be extremely difficult to achieve even on rare occasions. Going forward all of these efforts will be needed to achieve F23's cycle time improvement goals.

F23 found that active inventory control had a significant impact on cycle time. F23 developed the Wafer Start Protocol to provide a systematic method for varying wafer starts in order to control inventory. In doing so, F23 challenged a number of assumptions held within Intel. In particular, F23 found that significant cycle time improvement can be achieved even without significant improvement in equipment performance. F23 reduced its inventory by approximately 44%. This result suggests that F23's original condition was most likely a state of excessive inventory even though it was similar to the condition of a number of other Intel factories. This conclusion is supported by initial comparisons of the critical WIP ratios for other Intel fabs.

F23's continued pursuit of lower inventory levels has uncovered a number of issues, and future inventory reductions are expected to do the same. This provides continuing opportunities to eliminate more problems and foster additional improvement. F23 has seized upon these opportunities and eagerly awaits the identification of more issues knowing that such a course of action allows them to achieve ever greater cycle time improvements.

Seeking out additional problems is not typical of Intel factories. Such a mindset change is indicative of the state of improvement F23 has achieved. A number of factors have contributed to F23 success. It may not be possible to emulate all of these factors at other Intel factories but hopefully the lessons learned at F23 will provide other Intel facilities with improvement ideas, and the other facilities will be able to gain from these concepts. F23's example demonstrates that a typical Intel facility can achieve cycle time results on par with the best in the semiconductor industry if it is willing to make real changes for the sake of improvement.

11: References

[1] Leflar, Jim A. *Practical TPM: Successful Equipment Management at Agilent Technologies.* Portland: Productivity Press. 2001.

[2] Fearing, Rebecca. "Managing Preventative Maintenance Activities at Intel Corporation". Thesis, Massachusetts Institute of Technology, 2006.

[3] Spears, Steven and Kent H. Bowen. "Decoding the DNA of the Toyota Production System." *Harvard Business Review.* September-October, 1999. pp. 96-106.

[4] Quirk, Michael and Julian Serda. *Semiconductor Manufacturing Technology.* Upper Saddle River, NJ: Prentice Hall, 2001.

[5] Little, J. D. C. "A proof of the queuing formula L=λW," *Operations Research*, vol. 9, pp. 383-387, 1961.

[6] Graves, Stephen C., Harlan C. Meal, Daniel Stefak, and Abdel Hamid Zeghmi. "Scheduling of Re-Entrant Flow Shops." *Journal of Operations Management.* Vol. 3. No. 4. August, 1983, pp. 197-207.

[7] Wein, L.M. "Scheduling semiconductor wafer fabrication," *IEEE Transactions on Semiconductor Manufacturing*, Volume 1, Issue 3, Aug 1988 Page(s):115 – 130.

[8] Hopp, Wallace J., Mark L. Spearman. *Factory Physics: The Foundation of Manufacturing Management.* Chicago: Irwin/Mcgraw-Hill, 2001. pp. 162-166.

[9] Levitt, M.E.; Abraham, J.A. "Just-in-time methods for semiconductor manufacturing" *Advanced Semiconductor Manufacturing Conference and Workshop*, 1990. ASMC 90 Proceedings. IEEE/SEMI 1990; 11-12 Sept. 1990 Page(s):3 – 9

[10] Martin-Vega, L.A.; Pippin, M.; Gerdon, E.; Burcham, R. "Applying just-in-time in a wafer fab: a case study." *IEEE Transactions on Semiconductor Manufacturing*, Volume 2, Issue 1, Feb. 1989 Page(s):16 – 22.

[11] Spearman, Mark L.; Woodruff, David L.; Hopp, Wallace J. "CONWIP: a pull alternative to kanban" *International Journal of Production Research*, May90, Vol. 28 Issue 5, p879-894.

[12] Rose, Oliver. "CONLOAD – A New Lot Release Rule for Semiconductor Wafer Fabs." *Proceedings of the 1999 Winter Simulation Conference.*

[13] McBride, Eric and Karl Kempf. "Metering Fab Starts by Application of Control Charts." *Intel Manufacturing for Excellence Conference.* 1997.

[14] Connally, Jason. "Introducing Pull Methodologies in a Semiconductor Fab." Thesis, Massachusetts Institute of Technology, 2005.

[15] Leachman, Robert C.; Jeenyoung Kang; Lin, Vincent. "SLIM: Short Cycle Time and Low Inventory in Manufacturing at Samsung Electronics". *Interfaces*, Jan/Feb2002, Vol. 32 Issue 1, p61-77.

[16] Dennis, Pascal. *Lean Production Simplified.* New York: Productivity Press, 2002, pp. 20-24.

[17] Flinchbaugh, Jaime and Andy Carlino. *The Hitchhiker's Guide to Lean: Lessons from the Road.* Dearborn: Society of Manufacturing Engineers, 2006, pp. 33-36.

[18] Hammer, Michael. "Deep Change". *Harvard Business Review*, April, 2004. pp. 84-93.

[19] Hayes, Robert, Gary Pisano, David Upton, and Steven Wheelwright. *Operations Strategy, and Technology: Pursuing the Competitive Edge.* Danvers: John Wiley & Sons, 2005, pp. 154-156, 284-288.

[20] Kotter, John P. "Leading Change: Why Transformation Efforts Fail." *Harvard Business Review*, March-April, 1995. pp. 59-67.

S111-54