Integrated cost and productivity learning in CMOS semiconductor manufacturing

by G. A. Leonovich
A. P. Franchino
W. J. Miller
U. E. Tsou

This paper describes a cost and productivity learning process that was carried out on a large-capacity CMOS manufacturing line at the IBM Burlington facility from 1991 to 1993. Major productivity gains were realized through process and tool improvements affecting yield, and through work-in-progress optimization and scrap reduction. Significant cost learning was also accomplished through tool cost management, capital depreciation and space cost reductions. and manpower optimization.

Introduction

This paper gives an overview of a cost- and productivity-learning process that occurred on an IBM Burlington CMOS manufacturing line from 1991 through 1993. The line is a large-capacity (more than 25000 wafers per month) fabricator producing CMOS chips using devices of 0.5-μm to 0.7-μm channel length. Called the CMOS IV line, its products include 4Mb DRAM chips, OEM-compatible and IBM PowerPC™ microprocessors, and a variety of other logic products for both internal and OEM markets. The majority of these products compete in the commodity marketplace, and in general their prices are set by that market. Memory prices have followed a 27% per year "learning" rate for the past 20 years or more, and competitive costs have been well understood. In 1991, our costs were not competitive in this environment, and unless we could begin to realize significant and sustained cost learning, the future was not encouraging. Beginning in 1992, therefore, we undertook a concerted set of product cost learning activities to correct the problem. The following paper describes many of the key activities and their outcomes.

Defined as a managed, well-understood, continuous rate of decline in the manufacturing cost per unit of product output, unit cost learning in the semiconductor industry takes two paths. The first, properly called *productivity* learning, seeks to increase output at a given fixed cost. The second, which is true *cost* learning, seeks to reduce the cost of a given, fixed output. Most plans work both paths simultaneously, but on a percentage basis, improving

0018-8646/95/$3.00 © 1995 IBM

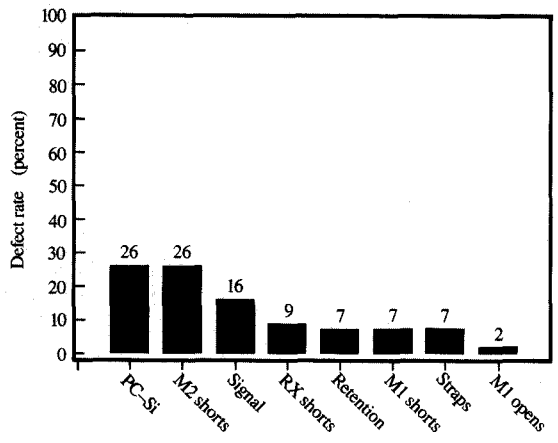

Figure 1

Pareto chart showing defect rates for eight defect categories in a production line for 0.5-μm CMOS.

output generates twice as much learning value as reducing costs. Therefore, 70 percent of the learning value generally results from activities that improve productivity, with the remaining 30 percent resulting from cost reduction. This paper discusses first the major productivity improvements our line has achieved [electrical yield, management of work in progress (WIP) to optimize output, and scrap reduction], followed by key cost-reduction activities (raw wafer, capital and space, and manpower). A discussion of cost measurement tools (cost-of-ownership analysis and activity-based management) is also included in the section on cost reduction.

Productivity increases

- *Electrical yield*

A key historical focus for reducing costs has been to improve the electrical test yield so that more good product leaves the production line per wafer processed. This was a primary concern for the CMOS IV line, but a difference in our approach to this goal in CMOS IV technology was the use of new diagnostic techniques and cross-discipline teams to work on problems. The teams were authorized to make changes to both the design and the base process in order to optimize yield and total productivity. Such changes were focused on improving the robustness or "process window" (resistance to down-time, scrap, and rework) of both design and process.

The traditional methodology for improving yield was for characterization groups to review the data from the functional test and defect test monitors, perform failure analysis, and report the results to the appropriate process engineer. Each process engineer would then work in a specific area, usually separate from the other engineers, to reduce defects by improving the tools or the process. In theory this sounds like a good way to work, but what often happens is that process engineers work in isolation from one another and from the characterization engineers. With the complexity of today's process and designs, a given problem is no longer confined to one area, but usually involves many process areas and requires the coordinated effort of both process and characterization engineers to identify and correct.

In the past, the characterization groups reported all problems, and the process engineers worked on all of them. With the increased use of computer data analysis, the characterization engineers have been able to give more value to reported data. Characterization engineers can now track data by tool, day, and hour through individual operations in the process. This allows fast and efficient auditing of a yield or defect anomaly through all operations and pinpointing the particular tool or operation that is causing the problem. A process engineer can now see the data on a graph in real time and can determine what changes occurred or may have occurred that caused the problem. In this way, reaction time is made much faster, and a tool or process can be taken off line until it is fixed. Characterization engineers can also track data through individual tools within an operation and determine which tools run better or worse than others, and the process engineers can use this information to improve the tools or processes. This tool-to-tool variation is used to weigh the impact on productivity of shutting down a tool, as opposed to continuing to run the tool with degraded yields.

One of the most powerful techniques used is the ability to correlate functional data to defect data or electrical monitor data. This correlation allows the engineers to build a pareto chart (**Figure 1**) of the defects or defect categories and their impact on yield. Using these data, teams of process and characterization engineers can focus on correcting the high-impact problems and expend less resource working on those for which the yield loss is relatively low. With the complexity of current processes, problems can have a multitude of potential sources and can interact with and compound one another. As a result, teamwork between engineers and technicians with a broad background of expertise is critical.

Of the many changes that have taken place in the CMOS IV fabricator business processes, the use of teamwork within the engineering community may be the most important. Teamwork was used from the inception of the line, with a joint effort between development engineers and line process engineers to introduce the CMOS IV

4Mb DRAM into volume production. Instead of the development engineers handing off a new product and process to the manufacturing line and leaving to develop the next product, they worked jointly with the process engineers for more than a year and are still in close communication. By working jointly on the early high-impact items, process engineers were able to learn the new technology more quickly, and development engineers were able to learn the manufacturing issues that process engineering and manufacturing deal with and to understand the need to design processes and products that are more manufacturable. The effectiveness of this team concept was demonstrated during the work on the most significant yield-limiting problem shown in Figure 1, PC-to-NW* (PC–Si) shorts.

PC-to-NW is a short from the gate to the source–drain region of a transistor. This defect has been the number one yield limiter of the 4Mb DRAM from the inception of the product. The causes of this defect are closely related to signal margin failures and interconnection failures (straps), which are also shown on the pareto chart (Figure 1). These items are related, since whatever fixes one often degrades the other two. The defect has many different sources, such as thick selective silicon, thick titanium, under-stripping of the silicide, over-etching of the spacers, and/or embedded foreign material. The technical issues and reasons for each are very complex, and are not discussed in this paper. What is important to understand here is that the sources of PC–NW cannot be solved by one engineer.

A team was formed that included representatives of all process and development engineering areas. The team met on a regular basis to address this failure mode, and its action plans cut across department boundaries. As a result of this effort, a significant reduction was made in the incidence of this defect (**Figure 2**). At the same time, designers made changes to the chip, making the product less susceptible to the defect and shrinking the design to make it more productive. This activity has been so successful that a new team of process and development engineers is working on a new design and process that essentially eliminates the PC-to-NW defect and shrinks the product by another 40%. This design will be more robust and also more productive (**Figure 3**).

This concept of cross-department and cross-functional teams has expanded throughout the CMOS IV line to address many key issues, including not only yield, but also productivity, tool availability, tool throughput, and certain controls that are necessary but do not add value. In addition, process steps have been eliminated (**Figure 4**), which reduces the cycle time of the product, improves yield learning, reduces capital expenditures, reduces manpower, and reduces material usage.

*PC-to-NW: polysilicon-to-n-well; the polysilicon is that used to form the word lines.

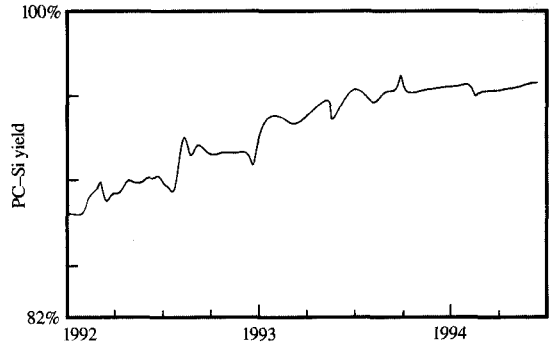

Figure 2

Yield improvement due to reduction of PC–Si defects (shorts from gate to source–drain).

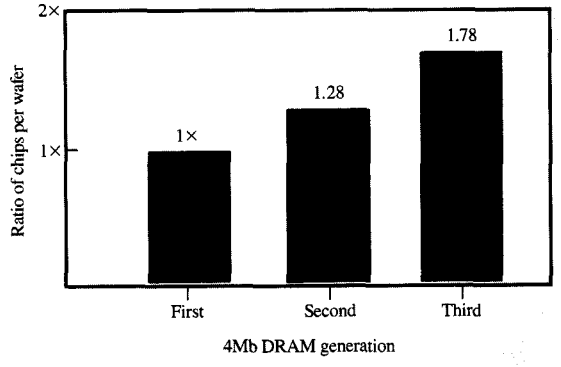

Figure 3

Chips-per-wafer yield improvement across three generations of 4Mb DRAM.

- *Work-in-progress optimization*

One of the least understood components of cost reduction involves the optimization of work in progress (WIP). Production requires three things—tools, workers, and WIP. Tool planning is quite sophisticated, as is manpower management (to be addressed later), which leaves optimizing WIP as a key management objective. The 970 fabricator is a multiple-product, multiple-turnaround-time environment within CMOS IV, with a product set that competes in the commodity market. Cost and serviceability

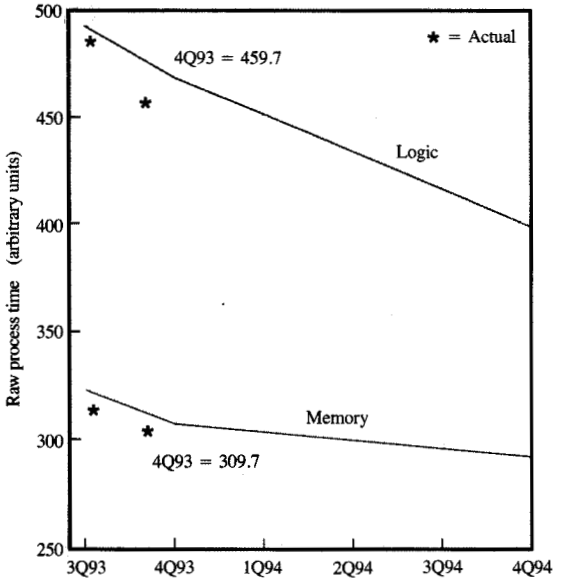

Figure 4

Reduction in process time due to elimination of process steps.

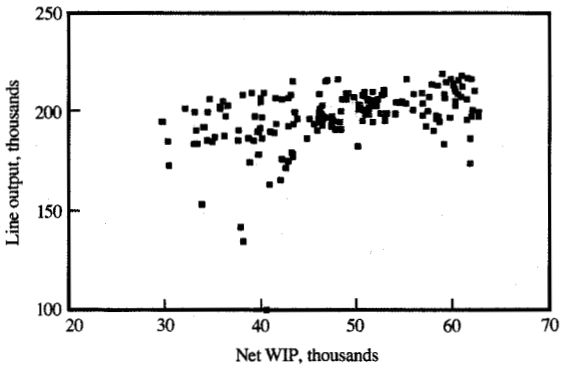

Net WIP = gross WIP − shortfall

Figure 5

"Hockey-stick" curve: Line output (combined for all stages) vs. net work in progress (WIP), over the period from October 1992 to March 1993 for 0.5-μm CMOS.

are the key line measurements, and achieving maximum line output is critical. (Line output is the daily sum of all wafer movements in the line, not just end-of-line output.

Both in queuing theory and in intuitive approaches to line loading, the relationship between WIP and output is known. If one begins with an empty line and a low level of starts, the product will move quickly and the level of output will rise with the number of starts; the line is WIP-starved, and line movement increases linearly with additional WIP. If the start rate continues to increase, the loading eventually saturates the line, which results in a static level of line output regardless of additional WIP or starts. Using historical data on WIP levels and output, we have identified the WIP level at which these two trends meet. We refer to the plot of WIP vs. output as the line "hockey stick," from the general shape of the data distribution. For WIP levels below the breakpoint, the line is under-utilized, and both movement and output fall off with lower WIP levels. Above the breakpoint, output is static, and cycle time increases with increasing WIP with no additional output. **Figure 5** shows the 970 line hockey stick for a six-month period centered on year-end 1992. At this point in time, we needed 48000 wafers for maximum output. The scatter within the hockey stick is a result of shortages of tools or people. At this WIP level and output level, the line average cycle time is about 3.8 times the direct raw process time for a wafer lot. Using certain tailored methods, we have been able to run specific lots at speeds approaching 1.5×, but the average of all the lots must equal 3.8×. If we chose to run a 3.0× line cycle time, it would lead to very high levels of lost output (revenue). Similarly, if we chose to run the line at WIP levels above saturation, we would suffer higher inventory costs and exposure to yield or engineering change issues for no useful purpose.

Once the WIP target has been set, it must be deployed to the various areas and departments. Using our history data base, we generate a scatter plot of WIP vs. output for each department. Most departments show a "classical" triangular plot (**Figure 6**): For low levels of WIP, the department is unable to achieve maximum output, but once the WIP level rises above some optimum point, no additional output is achieved. This optimum point becomes the target WIP for that department. Normally, the sum of the department WIP targets will equal the optimum WIP for the line as a whole, since the shape of the center-line for the department scatter plots is the statistical cause of the hockey-stick distribution. Understanding the relationship of WIP vs. output at the department level, we realized that exceeding the WIP target for one department does not make up for falling short of the target in another. The "over-WIP" department may achieve some additional output on average, but the "under-WIP" department will lose enough to more than offset any gain. We developed the term *net WIP*, which is total line WIP minus the sum of the under-WIP levels of those departments below their

optimum WIP. For example, if the line had 50000 wafers, but three departments were each under their WIP targets by 1000 wafers, the net WIP would be 47000: [50000 − 3(1000)]. It is this net or effective WIP which is plotted in Figure 5. The keys to maximum line output are, therefore, WIP level and balance. **Figure 7** shows the running hourly deviation of output from target, and net WIP deviation from target for the same interval. When the net WIP "delta" increases, we see an almost instantaneous loss of output.

When we choose to maximize output (and therefore revenue), we must understand that the resulting WIP levels (inventory) and cycle time are a simple calculation. Over time, our maximum line output has increased through a variety of means (improvements in process and tools, additional tools, and reduction in unproductive tool operation), and the optimum WIP level has been reduced by more repeatable product flow ("tack" time) and improved tool and process reliability. Over the past two years, we have reduced WIP levels by 20%, improved output by 20%, and reduced the optimum WIP by 15%, achieving not only more output with less WIP, but also faster cycle time through WIP management.

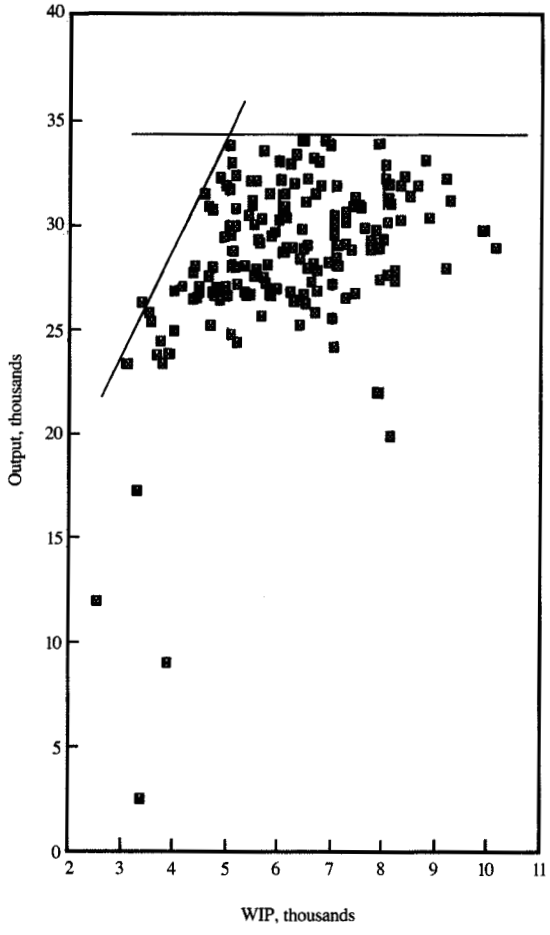

Figure 6

Example plot of output vs. WIP for one department.

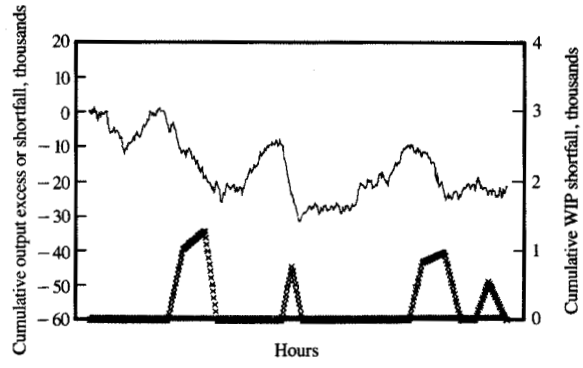

Figure 7

Cumulative hourly output excess or shortfall compared with WIP for a baseline output of 230000 wafers per hour.

Productivity increases

• *Scrap reduction*

Improving the mechanical yield is critical for reducing the cost of products, since spreading the relatively fixed costs of materials, space, depreciation, and manpower over more output reduces the unit cost. In starting up a new product in a new process, misprocessing or miscentered processing will significantly affect the line's wafer output (scrap) as well as the electrical yield. As product complexity has increased (64Kb to 1Mb to 4Mb DRAMs), the opportunity for errors to occur has also increased, resulting in lower mechanical yield. With concurrent increases in internal cost emphasis and OEM competition, these higher scrap levels became unacceptable. A team approach was used to improve CMOS IV scrap performance.

The team analysis showed that the keys to mechanical yield improvement were improved controls and root cause analysis that resulted in the implementation of real fixes. A control review team was formed to investigate and improve the controls that were used to run the process and the tools in CMOS IV fabrication. A scrap review team was formed to review all scrap and identify the root cause. Both teams comprised a mixture of operators, technicians, engineers, and managers.

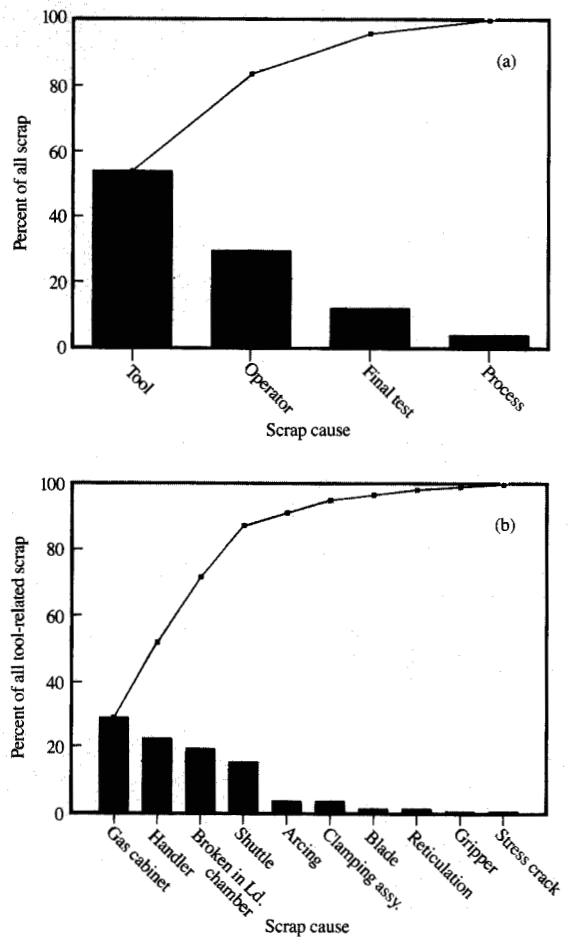

Figure 8

Pareto charts of scrap tendencies during reactive ion etching (RIE): (a) all causes; (b) tool-related causes.

The process control review (PCR) team was charged to work with the individual tool groups, such as the photolithography stepper team or the reactive ion etch (RIE) team, to review each group's tool and process controls. There are 30 such tool groups, and each group's controls were reviewed to ensure that they were sufficient to prevent scrap, maintain the process in control, and put action plans in place to address the top scrap items on the pareto. An example is shown in the RIE team paretos [**Figures 8(a)** and **8(b)**], which list all of the historical causes of scrap. A team composed of operators, technicians, engineers, and management worked on reducing or eliminating several of the top detractors. The PCR team not only ensured that there were proper controls, but also looked for unproductive or unnecessary controls to eliminate. An example of this was a monitor that was run with every lot of product during the RIE operation. During the review of the RIE team controls, the PCR team noticed that the monitor was never out of control, and recommended that the monitor be eliminated. This reduced the cost of the process by eliminating both the monitor and some related work that was adding no value to the operation.

The scrap review team was chartered to investigate all scrap. The team put in place a business process to define the steps that were to occur from the time at which product was first put on hold until the point of final disposition. (The term *disposition* can mean to ship as is, to rework, or to scrap.) This business process was key (**Figure 9**) and was well defined so that all groups would follow it closely. The weekly review meetings became a team effort between manufacturing and engineering, allowing an open forum in which each group could learn from the others' activities in scrap reduction. As required by the business process, after wafers are put on hold, the individual who took the action is responsible for disposition, and can enlist the help of engineering or other groups to achieve disposition as quickly as possible. If wafers are scrapped or need rework because of misprocessing or tool malfunction, the area responsible for that scrap is required to present to the scrap review team a root cause analysis and a root cause fix. It is critical for the scrap review team to ensure that a root cause fix is in place in order to prevent further product loss. If the root cause results from a lack of control, the scrap review team asks the process control review team to do a special review of that area to help improve the controls and to ensure that the proper ones are in place. A major side benefit of this business process is that the amount of product on hold and the amount of time a lot is on hold have both been reduced by an average factor of 3 (**Figures 10** and **11**).

Neither the scrap reduction team nor the PCR team would be successful without the full cooperation and support of management, engineering, and operators. To promote this, an incentive plan was established to support a mechanical process yield of 95% for the line as a whole. The incentive program consisted of three levels of achievement and reward. At the start of the incentive program, a scrap baseline was defined for each tool group to use in setting goals for its area.

Level 1 required that a team successfully complete a process control review, ensuring that its area had a closed business process in place to identify owners, processes, dependencies, customer set, and controls. Key areas that were examined during these reviews were training methods, certifications, communication across teams, and measurements. As a reward for meeting these requirements, team members each received a small IBM

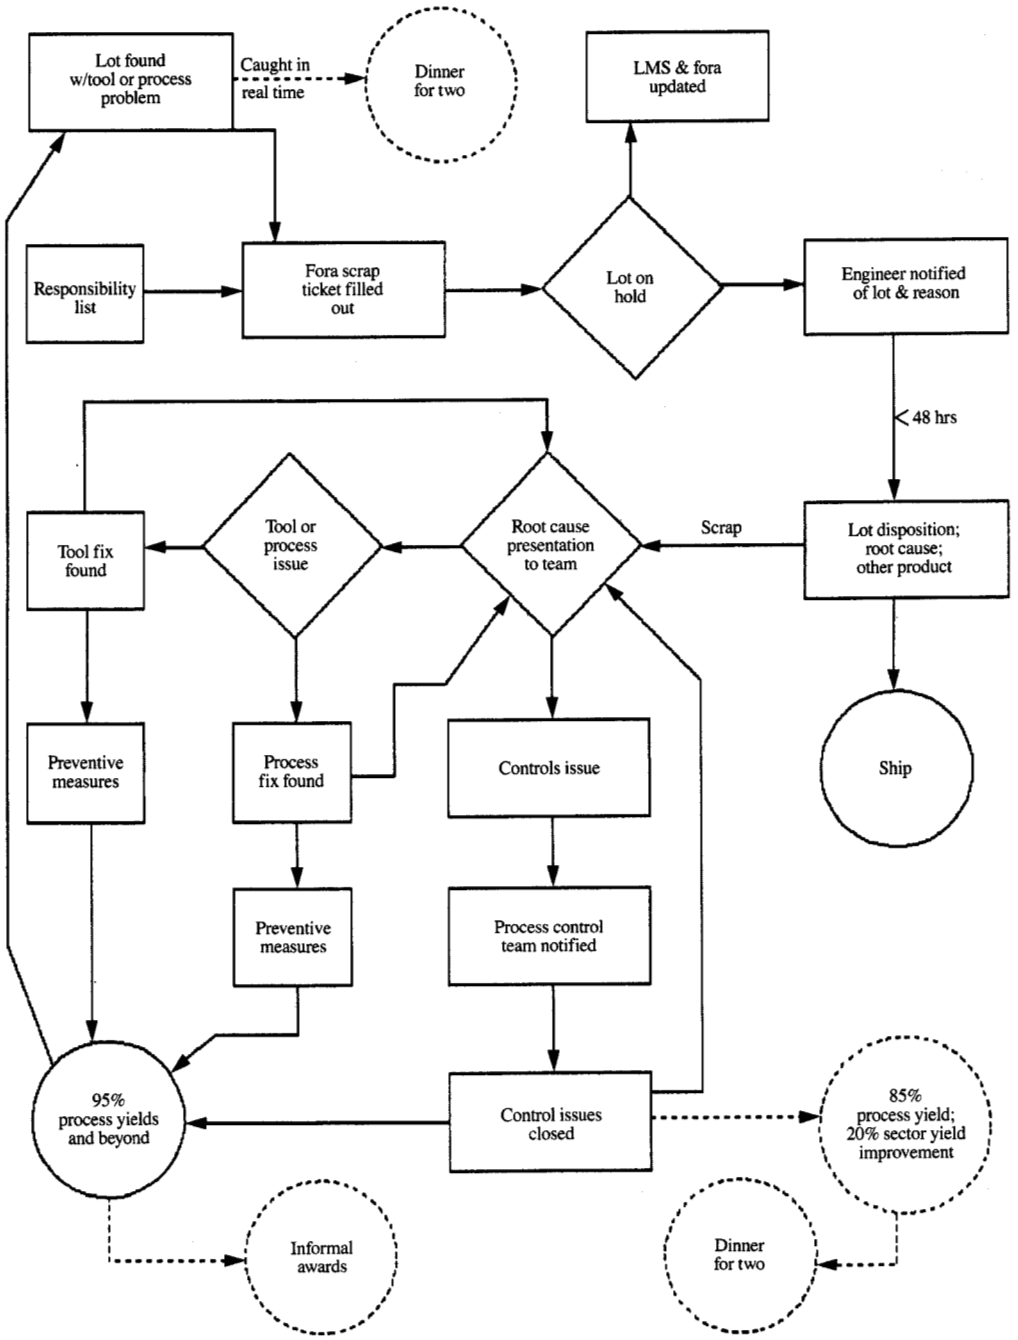

Figure 9

Flowchart of scrap disposition process.

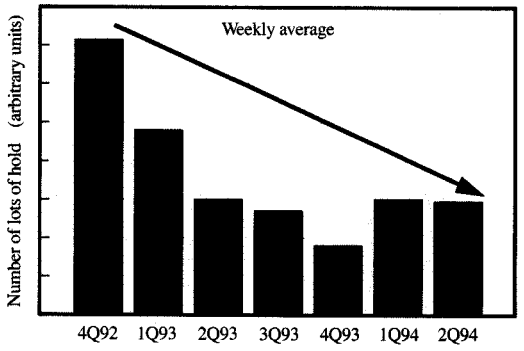

Figure 10

Trend chart for WIP quantity-on-hold.

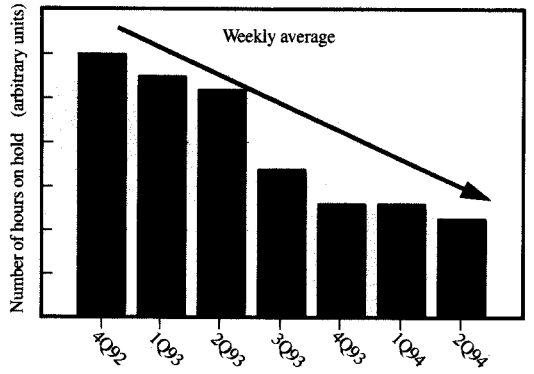

Figure 11

Trend chart for WIP time-on-hold.

gift. Level 2 required that each tool group reduce its scrap level by 20% (for a 13-week average), and that the total scrap for the line improve approximately 14% with no loss in device test yield. Level 3 was similar to level 2, requiring a further 10% reduction in scrap for each tool group and another 14% improvement in total line scrap.

The results of the program have been very positive. In less than one year, mechanical process yield is running at record levels. Although level-3 yields have not yet been achieved, the line is running consistently at the 93–94% level, which has never been done before. To achieve level 3, the PCR team is beginning control reviews again, and

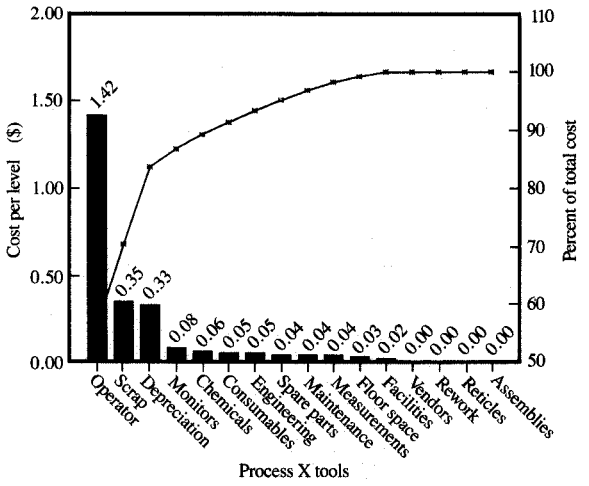

Figure 12

Average cost of ownership by element for an example process.

the scrap team is reviewing the top ten scrap areas. The full cost of this program will eventually reach $1.3 million, but the increase in revenue for the first year has been more than $65 million (a 50× return). At this time we are planning to implement a fourth level in the program, which will take into account not only mechanical yield but test yield and costs.

Cost reductions

- *Approach and tools*

The key cost drivers in semiconductor manufacturing, accounting for 75 to 80% of total costs, are depreciation, space, raw wafers, and manpower. In this section we discuss the important approaches we used in reducing our manufacturing costs, as well as the tools used to do this job. Cost of ownership (COO) and activity-based costing (ABC) were the key tools used to understand costs and provide a focus for cost learning. The CMOS IV fabricator has developed a cost-of-ownership model that allows cost calculations at the process level. The key cost elements included in the COO model are shown in **Figure 12**.

With activity-based management and activity-based costing (ABM/ABC) [1], we attempt to go one step further and define the causes of cost within a category. ABM uses the COO model as well as additional input such as manpower drivers, product drivers, and quality drivers. Using this system, we can understand the impacts of various product programs. In a multiple-part-number

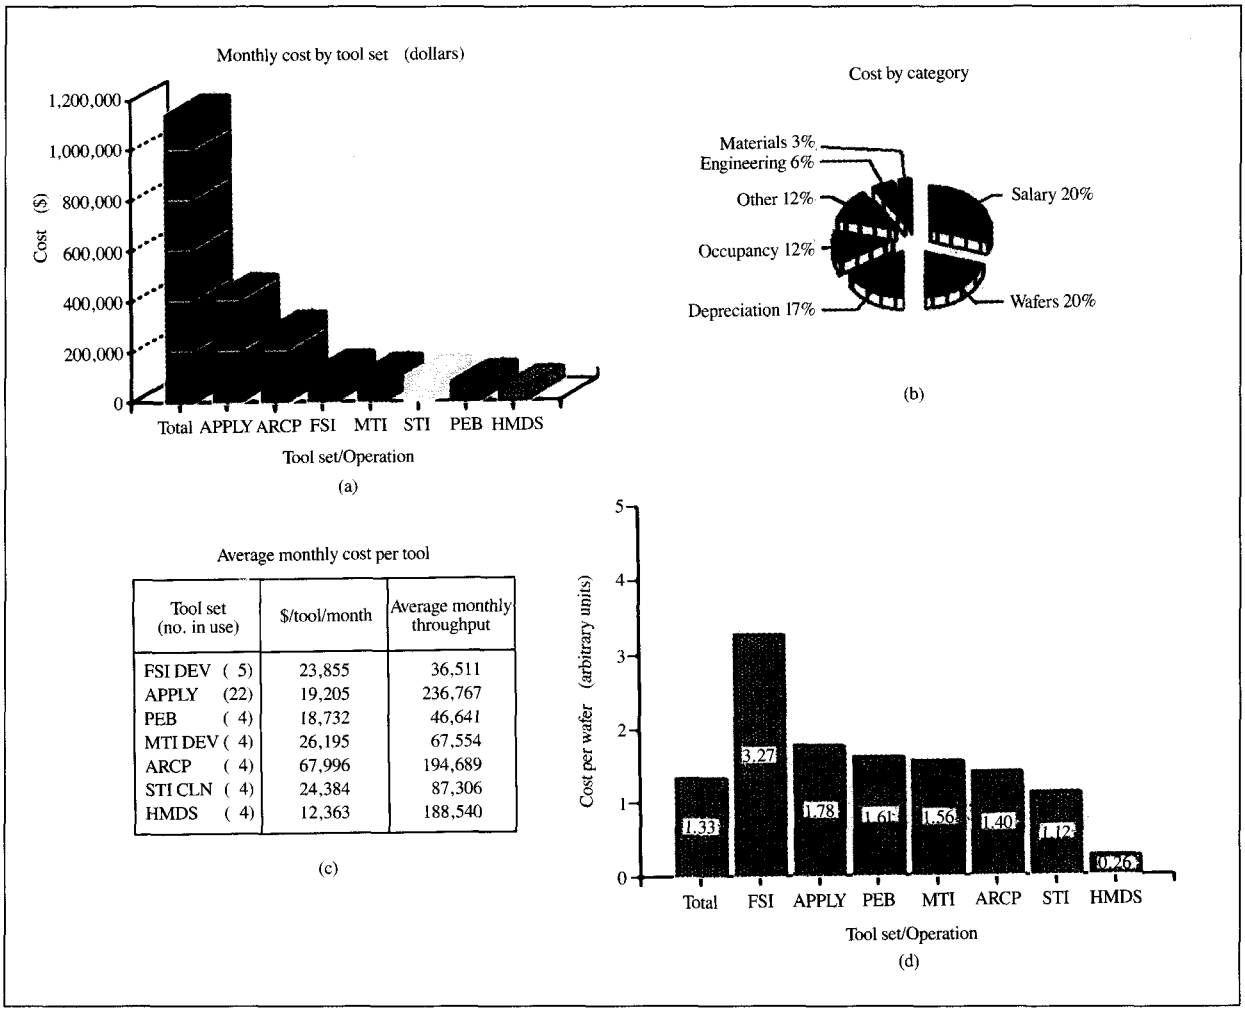

Figure 13

Activity-based cost breakdown for an example department: (a) total cost (all categories) per tool set or operation; (b) cost by category; (c) average monthly cost per tool; (d) average processing cost per wafer.

environment, the definition and separation of product families is an integral piece of information for "designing in" a cost-effective process and making strategic menu-planning decisions. An example of a department's ABM breakout is shown in **Figure 13**. Part (a) describes the monthly cost for the whole department, broken down by tool set. Part (b) breaks down the department's cost by cost element. Part (c) is a tabular chart listing the number of tools and their monthly cost and throughput. These data in turn generate part (d), a bar graph of wafer processing cost by tool set. The most expensive tool set to run is the ARCPs, but it does not have the highest wafer processing cost, since its throughput is high. While the cost to run the FSI is moderate, it has very high unit cost because of its low throughput.

- *Raw wafers*

Raw wafers, costing more than $50 million a year, have been a prime target of our cost-reduction activities. The reduction of raw wafer usage comes from improvements in wafer mechanical yield (scrap) and reductions in virgin monitor usage. Improvements in scrap have already been addressed. Reductions in monitor wafer usage have been achieved through development of an improved wafer-reclamation process, and by improving the reliability and capability of processes and tools, thus reducing the need for monitor measurements.

Statistical techniques were used to balance the risk between building defective products and the added cost of monitoring. Decreasing the monitoring frequency by making the process more robust or by improving process

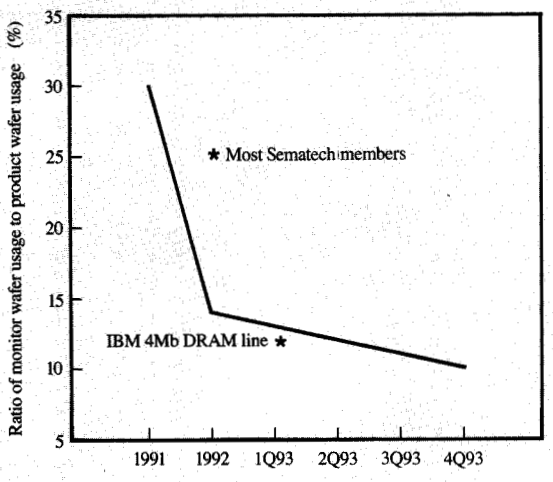

Figure 14

Trend in usage of raw wafers for monitoring.

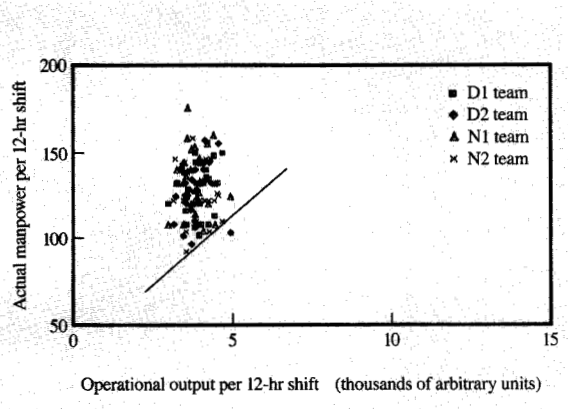

Figure 15

Manpower vs. value-adding output for a department whose staffing level is near the minimum required.

tool stability are the most desirable solutions. However, since it is not possible to completely eliminate the need for monitoring, improving the reclamation and reuse efficiency of the monitor wafers is also very important. Recleaning between uses and repolishing used wafers has enabled us to increase the reuse of monitor wafers from three or four passes to more than twenty. By focusing on both reduction of usage and reclamation efficiency, the CMOS IV fabricator has been able to reduce its monitor usage by more than 80% in three years (**Figure 14**).

• *Capital and space*

Capital depreciation cost, together with the space required for tools, is a $100 million item for the CMOS IV fabricator. Control of this depreciation cost comes from 1) improving tool utilization and productivity to reduce the number of tools and avoid new purchases; 2) improving the process capability of the existing tool set, prolonging its useful lifetime and avoiding new purchases; and 3) ensuring the future extendibility of new tool purchases (a key criterion in the selection of new tools). For the department shown in Figure 13, the FSI tool set was totally eliminated and the workload shifted to the ARCPs tool. This transfer required an upgrade of $100 thousand, but decreased wafer processing costs by 20%. By focusing on these three areas, the Burlington CMOS IV fabricator has kept its capital depreciation cost level in spite of dramatic product menu shifts, tightening of product ground rules, and increased line throughput.

Space charges normally amount to 50% of a tool's depreciation cost, so cost-effective line layout without compromising worker safety should be the major guideline here. The focus should be not only on the physical-dimension requirements, but also on the facility service required and the cost of the space. Overstatement of the clean-room class or temperature/humidity requirements can increase cost significantly. We have reduced our total space requirement by tool productivity increases as well as timely removal of excess or obsolete tools from the line and reuse of the space. The cost per square foot has also been reduced by reducing air supply rates, with line cleanliness actually improving.

• *Manpower*

A great deal of effort has been expended by the industrial engineering community in support of manpower modeling, to improve the calculation of manpower requirements in support of production. As with most models, the output is only as valid as the input. Owing to human nature, most inputs concerning manpower err on the high side. For the past two years in the 970 fabricator, however, we have taken another approach. Instead of modeling, we have constructed a manufacturing data base containing actual product throughputs and manpower working in the fabricator at various levels of detail (line, department, tool). These data are analyzed to improve our understanding of the relationship of manpower in a given area to the output of that area. An example of this analysis for one of our departments is shown in **Figure 15**. Each point on this scatter plot is one twelve-hour shift for a department, and the time period covered is eight weeks

(or 112 shifts). The plot shows that for this particular department, temporary additional manpower did not lead to any increase in output, and we can conclude that output in this department will not increase with additional base manpower. Second, by examining the scatter plot in Figure 15 again, we can tell whether the department needs the level of manpower it currently has. If in fact a department's staffing is near the required level, as in Figure 15, we can expect that on days when manpower levels fall *below* this minimum level, output will suffer. (In contrast, **Figure 16** shows a department for which output sensitivity to manpower is not apparent at the current level of daily staffing. Knowing the minimal level allows us to set organizational targets, know what levels of absence we can afford, and react on a daily basis to overtime and vacation requests. One side effect of this method of manpower allocation is to encourage continuous learning within the organization. (The need for this is evident, for example, in the department represented in Figure 16.) Human nature perceives "unknowns" as problems, and methods will be discovered to allow continued normal operation at lower and lower levels of staffing until the minimum required level is known for every department.

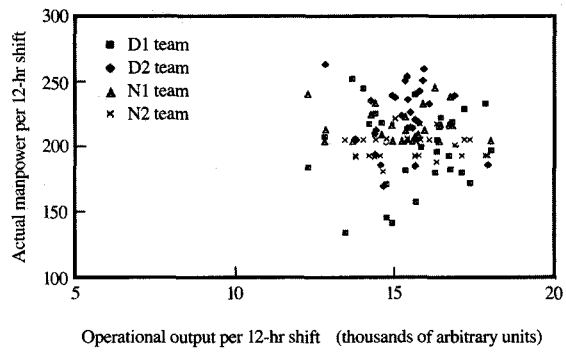

Figure 16

Manpower vs. value-adding output for a department whose staffing level appears to be above the minimum required.

The second technique used to manage manpower concerns efficiency. Using the normal definitions of value-adding and non-value-adding work, we benchmarked each organizational area as to its effective use of manpower. On the basis of previous work, we came to the conclusion that it takes approximately the same amount of time for an operator to "run" any tool in our line. Although our tools range widely in complexity, value in proportion to the operator's time was relatively constant. For all intents and purposes, each tool looked like a black box with a load/unload port. The labor required to fetch, verify, load, unload, and transport to the next operation was very similar for all of our tools (nine minutes). Industrial engineering has calculated a conversion factor to estimate the time an operator is at work but unable to "run the tool" (breaks, department meetings, etc.). This factor is approximately 30%, so our base rate for running a box of wafers through a tool is 12 minutes. Knowing how much product has been run over a period of time, together with how much labor has been expended, allows us to measure the efficiency (amount of labor spent on value-adding vs. non-value-adding tasks). **Figure 17** shows how the ten production departments in the 970 fabricator compare in terms of efficiency. Workers in a department with a 4× measurement spend three hours on non-value-adding activity for every hour they spend running the tool. Beginning in 1992, we were challenged to reduce our manpower costs by 30%. This figure was converted to a 40% non-value-adding (NVA) labor cost learning target, which was applied to each department. Those with higher levels of NVA work had to learn more in terms of absolute

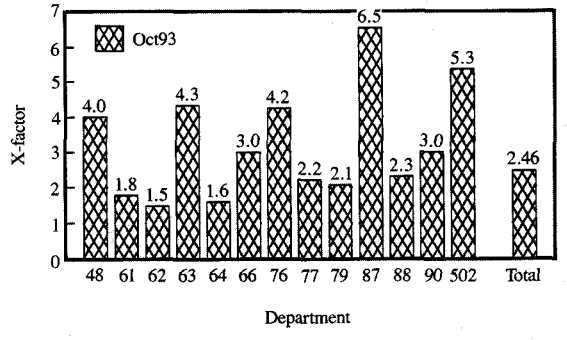

Figure 17

Manpower efficiency (ratio of total labor time to value-adding labor time) by department.

hours, while those with lower levels of NVA had lower learning targets. This methodology recognizes that requiring equal levels of absolute learning is not fair or possible. **Figure 18** shows efficiency learning by department over a 14-month time period, and **Figure 19** shows the savings in manpower.

To build a causal cost model for manpower, we went a step further and generated a survey designed to break down our manpower usage by driver. **Figure 20** shows how the manpower for one of our departments broke out by use or activity. The remaining time is the amount that has not been applied to a work-related activity. The majority of

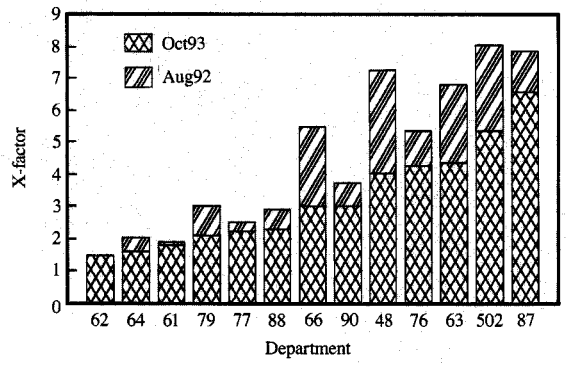

Figure 18

Manpower efficiency improvement by department for a 14-month period.

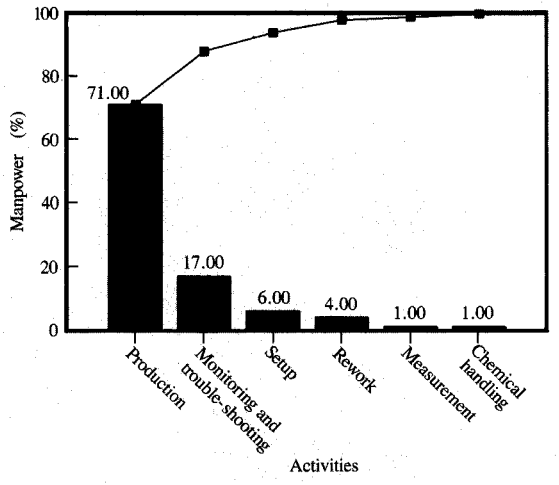

Figure 20

Activity distribution for the example department shown in Figure 16.

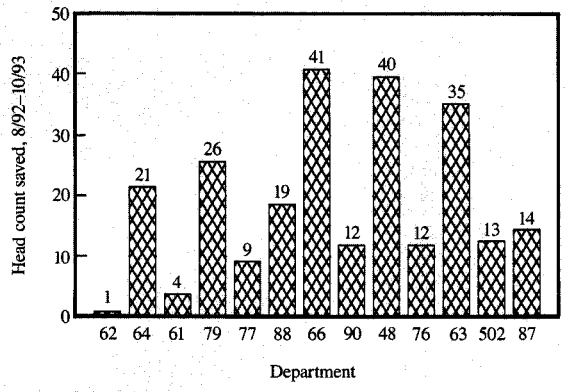

Figure 19

Manpower saving by department, achieved by the efficiency learning shown in Figure 18.

this remaining time is idle time of one form or another, including breaks, department meetings, and idle time due to lack of WIP or insufficient tool workload. We could then calculate how much each activity was costing, and found that the department shown in Figure 20 was spending 71% of its time on value-adding activities. This figure is directly comparable to the department efficiencies shown in Figure 18, and the two measurements were used to verify each other. Once this verification was complete, we were able to use the results of the manpower surveys to characterize each tool set and identify those areas with the most leverage.

Efficiency can be improved by eliminating unnecessary inspections, measurements, and setups, as long as the time made available by removing these actions results in higher throughput. Analysis of individual areas creates an objective understanding of the relationship between manpower and output, and permits much more accurate allocation of resources and learning plans across an organization.

Summary and conclusions

As mentioned in the Introduction, the cost and productivity learning activities just described have been ongoing in the IBM Burlington CMOS IV line for the past two years. Over that time we have realized a 50% per year reduction in unit costs and have achieved an overall level of competitiveness in our major product areas. This has been accomplished through an integrated, methodical, and team-oriented approach to line improvements. Everyone in our line contributes to the improvement process, and everyone benefits from the effort and achievement. The process must continue, however, essentially forever. We expect the competition to further reduce costs by 27% over the next year, and so must we. The key to remaining competitive is not a single invention or fix, but a business process that has the capacity for perpetual learning, and a team motivated to continue using (and improving) the process year after year. In the IBM Burlington CMOS IV line, we believe we have both.

Acknowledgments

The authors would like to acknowledge all of the operators, technicians, engineers, and management on the 970 fabricator team, without whose efforts none of the progress defined in this paper would have been possible.

PowerPC is a trademark of International Business Machines Corporation.

References

1. P. B. B. Turney, *Common Cents: The ABC Performance Breakthrough*, Cost Technology, Inc., Hillsboro, OR, 1991.

Received June 6, 1994; accepted for publication January 4, 1995

George A. Leonovich *IBM Microelectronics Division, Burlington facility, Essex Junction, Vermont 05452 (GLEONOVI at BTVVMOFS).* Mr. Leonovich received his B.S. in physics from the Rensselaer Polytechnic Institute in 1966. With significant experience in information systems, manufacturing engineering, and manufacturing, Mr. Leonovich is currently the manager responsible for line floor control systems and strategy in the B/970 manufacturing organization at IBM Burlington.

Anthony P. Franchino *IBM Microelectronics Division, Burlington facility, Essex Junction, Vermont 05452 (AFRANCHI at BTVVMOFS).* Mr. Franchino received the A.S. degree in electrical engineering from Rex Electronic Technology Institute, Nutley, New Jersey, in 1979. He is currently involved with worldwide supply-and-demand process reengineering in the IBM Microelectronics Division Manufacturing Program Office in Burlington. Mr. Franchino's experience includes quality engineering, manufacturing characterization engineering, manufacturing management, and Activity-Based Costing (ABC) coordination and development for both engineering and manufacturing models.

William J. Miller *IBM Microelectronics Division, Burlington facility, Essex Junction, Vermont 05452 (WJMILLER at BTVVMOFS).* Mr. Miller received a B.A. from Ithaca College, Ithaca, New York, and joined IBM at Burlington in 1979. With extensive experience in semiconductor engineering, including photolithography, RIE, and thin-film technologies, he is currently the manager of a photolithography and RIE engineering program in Burlington, involved in the introduction, fabrication, and cost reduction of advanced semiconductor products such as PowerPC and OEM-compatible logic chips and submicron DRAMs.

Uh-Po Eric Tsou *IBM Corporation, Route 100, Somers, New York 10589 (ETSOU at RHQVM07).* Dr. Tsou received his Ph.D. from the University of California at Los Angeles in 1985. With extensive experience in semiconductor manufacturing technology and production cost management, he has held several management positions in both electronic packaging and semiconductor processing. Dr. Tsou is currently program manager of advanced engineering for manufacturing on the IBM Corporate Manufacturing and Logistics staff.