21st Century Semiconductor Manufacturing Capabilities

Eugene S. Meieran, Intel Fellow
Technology Manufacturing Engineering, Intel Corp.

Index words: manufacturing, operational cost, NGM program

Abstract

Semiconductor device manufacturers face many difficult challenges as we enter the 21st century. Some are direct consequences of adherence to Gordon Moore's Law, which states that device complexity doubles about every 18 months. Feature size reduction, increased wafer diameter, increased chip size, ultra-clean processing, and defect reduction among others are manifestations that have a direct bearing on the cost and quality of products, factory flexibility in responding to changing technology or business conditions, and on the timelines of product delivery to the ultimate customer.

Regardless of these tremendously complex problems, the industry is focused on meeting the predictions of Moore's Law, for which enormous resources are expended.

One of the great challenges ascribed to Moore's Law, that facility costs increase on a semi-log scale, is now known as Moore's Second Law. However, unlike his First Law, the industry would prefer to depart from Second Law predictions to avoid hugely expensive ($20 Billion) future fabs and attendant high chip costs. Logistics control, inventory management, better facility design, supplier management programs, etc. are all responses to Second Law predictions, to which many resources have been devoted.

Other pressures on factory management are emerging. In addition to cost considerations, reduction in feature size and increasingly complex devices, the Massachusetts Institute of Technology/Leaders For Manufacturing-led program, "Next Generation Manufacturing" (NGM) identifies the following issues as significant:

- globalization of supplier, customer, and factory base
- exponential growth of information and knowledge management capabilities that enable faster and better decisions
- development of new materials and processes at atomic scale dimensions
- faster delivery of higher quality products to an increasingly demanding customer
- rising awareness of environmental and energy concerns

This paper discusses the technological responses of industry management and university faculty to the predictions of Moore's Second Law. Special attention is given to knowledge management and operational modeling and simulation technology. These processes help us better understand the benefits of various alternatives used to affect factory performance as traditional methods such as yield improvement, automation, increased wafer size, equipment reliability, etc. lose their leverage.

Introduction

Gordon Moore first proposed the law that bears his name in the late '60's: chip complexity (as defined by the number of active elements on a single semiconductor chip) will double about every device generation, usually taken as about 18 calendar months [1]. This law has now been valid for more than three decades, and it appears likely to be valid for several more device generations, as shown in Figure 1.

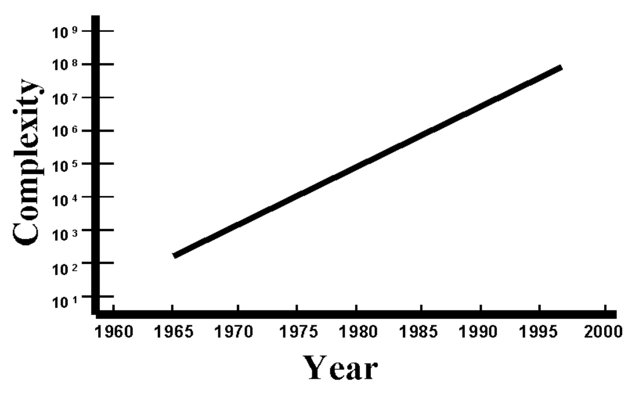

Figure 1: Moore's First Law

The compelling desire of the semiconductor industry to follow Moore's Law has affected high-volume device manufacturing, driving both the cost per bit of the devices and the overall cost of the fabrication and assembly facilities needed to build them. (Additional effects such as those on the ramp rate towards high-volume manufacturing are also experienced, but these are not discussed in this paper.)

For Moore's Law to remain valid, feature size must continue to be reduced, but since this reduction is insufficient in and of itself, chip size must continue to increase. Together, these two trends have not only maintained Moore's Law, but have accounted for the phenomenal success of our industry, since the cost per device element has now decreased by several orders of magnitude! Compared to *every* other commodity in the world, semiconductor chips are *cheap*, and continue to get cheaper (on a per element basis) every year.

The reduction in cost per active chip element is shown in Figure 2. Notice that while this cost continues to decrease, there appears to be a break in the curve: one section follows early predictions of Moore's Law, and the other departs from these predictions. This will be discussed later.

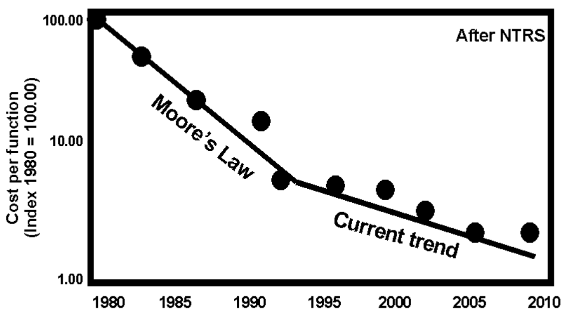

Figure 2: Cost per chip element

Many programs are associated with following Moore's Law and each has consequences for the cost per chip element, as shown in Table 1.

	1975	1997	2003
Chip complexity (index to 1)	1	10	100
Feature size reduction, µ m	2	0.25	0.08
Chip size increase, mm^2	30	150	600
Wafer diameter, mm	50	200	300
Facility automation, %	5	60	80
Die yield, % good	40	85	95
Line yield, % good	40	90	95
Assembly/test yield, %	90	99	99
Defect levels, DPM	2%	500	50

Table 1: Programs to maintain Moore's Law

Most of these programs tend to contribute to a reduction in chip element cost, but some of them, especially those dealing directly with increased chip and process complexity, tend to increase that cost. Fortunately, scaling, reduced feature size, improved yield, and increased wafer diameter more than make up for the added costs of more expensive packages and more complex processing.

Figure 3 shows the other major consequence of following Moore's Law. The reduction in cost per chip element is just offset by the increase in element density, leading to an essentially constant cost per individual chip. However, as a result, overall factory costs increase almost exponentially as we struggle to meet the ever increasing demand for more and larger high-performance chips. In order to meet cost per chip goals, cost per factory has increased to the point where high-volume factories now cost several billion dollars! So being successful in reducing chip costs brings its own share of additional problems. Building, equipping, and maintaining billion dollar factories tax even the most successful companies. This explosion of factory cost has come to be known as Moore's Second Law—one we do NOT wish to follow with such great zeal!

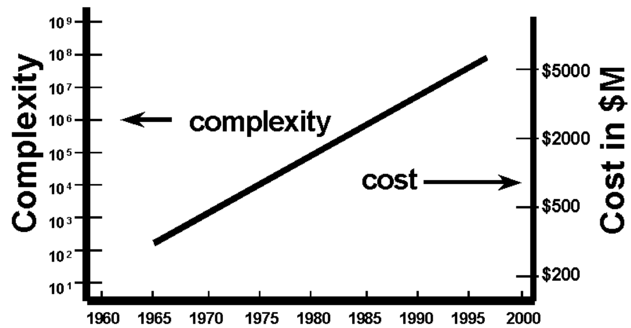

Figure 3: Moore's Second Law

Many of the same programs that have driven cost per chip element down are also responsible for the trend shown in Figure 3. In addition, some operational programs that have had little direct influence on cost per chip element have significant influence on factory cost. These additional programs are shown in Table 2.

	1975	1997	2003
Chip complexity (index to 1)	1	10	100
Feature size reduction, µ m	2	0.25	0.08
Chip size increase, mm^2	30	150	600
Wafer diameter, mm	50	200	300
Facility automation, %	5	60	80
Die yield, % good	40	85	95
Line yield, % good	40	90	95
Assembly/test yield, %	90	99	99
Operational efficiency	1	10	100
Equipment cost	1	10	50
Defect levels, DPM	2%	500	50

Table 2: Factory cost control programs

Tables 1 and 2 are combined, below, in Table 3, which shows two emerging problems with regard to both cost per chip element and factory cost containment:

1. Some goals of the programs are in conflict: lowering the cost per element actually adds to factory cost.
2. The leverage of some of the programs is diminishing: for example, we will not exceed 100% yield or 100% automation.

Hence, other means are necessary to meet cost projection goals for factories and chip elements.

	Cost per function	**Factory Cost**
Complexity increase	Up	Up
Feature size reduction	Down	Up
Chip size increase	Down	Up
Wafer size increase	Down	Slowing
Facility automation	Down	Slowing
Die yield	Down	Slowing
Line yield	Down	Slowing
Assembly/test yield	Down	Even
Operational efficiency	Down	Down

Table 3: Comparison of programs

The major program that does **not** suffer from topping out or from conflict is improving operational efficiency. However, before we discuss this, some additional forces acting on the manufacturing environment are discussed.

Emerging Trends

The additional forces acting on the manufacturing environment have little to do with Moore's Law. These forces are discussed in the National Science Foundation sponsored program, "Next Generation Manufacturing" conducted by the Leaders For Manufacturing program at MIT, the Agility Forum and the Technology to Enable Lean Manufacturing [2]. The major issues are listed in Table 4.

Globalization refers to the fact that for a number of reasons, industries are locating manufacturing facilities in many geographical locations, utilizing a supply of skilled workers at reasonable wages and servicing a widely dispersed customer base. As a result, suppliers of parts, materials, and equipment for these factories have had to become globalized, since operating conditions for manufacturers dictate that short time to delivery to the local customer is a matter of competitive necessity.

•	**Manufacturing globalization:**
	- Factories
	- Suppliers
	- Customers
•	**Increased global competition**
•	**Increased customer expectations**
•	**New technologies and processes**
•	**Environmentally aware manufacturing**
•	**Human factors:**
	- Training and retraining
	- Redeployment
	- Organizational structure
	- Wages and reward structure
	- Globally dispersed collaboration
•	**Pervasive information technology:**
	- Computation
	- Communication

Table 4: Emerging manufacturing needs

Due to the pervasive and timely availability of information and knowledge, global competition is more aggressive: new products are developed and brought to market quickly to globally distributed customers. Consequently, there is an erosion of what had been known as customer loyalty. Just as industrial jobs are no longer secure for life, brand-name loyalty on the part of a customer is not likely to survive; customers shop around for the most convenient or persuasive supplier.

Customers' expectations are increasing: they expect on-time delivery of high-quality customized products at prices reflective of high-volume manufacturing costs, and great service; otherwise, they will find other suppliers without hesitation. Quality is a given, not a differentiator; if one producer's product does not exhibit high quality, the customer will quickly find someone else.

Environmental concerns are also becoming more important in response to government regulations and societal concerns. Industry is recognizing that environmentally sound manufacturing is more rewarding than environmentally insensitive manufacturing.

Firms expecting to compete in the next millennium will have to play this ball game, on this playing field, with these new rules, encumbered as well by the needs and requirements listed in Table 3. These are the challenges the semiconductor industry faces as markets change, customer requirements change, and political and socio-economic forces affect how business is carried out.

Information Technology Responses

Two items listed in Table 4 were not discussed above: new technologies, materials, and processes; and greater access to global information and knowledge. The first refers to the fact that we can now create materials and structures on an atomic scale, with properties hitherto not only unavailable, but undreamed of. New products such as micro-motors, micro-refrigerators, micro-turbines, device analysis tools, and packaging will probably generate significant business in the not-too-distant future. However, since these do not concern semiconductor chip costs at the moment, they are not discussed further here. For details on these opportunities, see the NGM report [2].

The explosion of information technology (IT) is however another story. Indeed, information technology—the pervasive generation, storage, distribution and use of information and knowledge—seems to be the technology that may help resolve ALL the dilemmas of cost and competitiveness. IT can help with the declining rate of cost per chip element and increased cost per factory, as well as those emerging manufacturing needs identified in Table 4. In the remainder of this paper, we discuss how specific elements of information technology can be used to significantly impact all these issues.

Two applications of information technology that appear to have the greatest leverage are operational modeling and simulation, and management of knowledge assets and intellectual capital. In addition, these programs also affect the third way of decreasing these costs, the ramp speed to high-volume manufacturing. (For example, Intel's Copy EXACTLY! policy is one way of managing our corporate knowledge and wisdom to increase ramp speed.) However, ramp rate improvement is not discussed further in this paper. Please refer to "The Evolution of Intel's Copy EXACTLY! Technology Transfer Method" in this issue of the *Intel Technology Journal* for a fuller description of this important program.

Operational modeling and simulation (OM&S) and management of knowledge assets and intellectual capital applications have different purposes. OM&S is used to lower the cost and speed up the process of trying alternative solutions to different operational scenarios. It can provide quicker and more accurate answers to questions such as how much equipment or how many people are needed to perform a given number of activities; how can a factory be laid out for improved efficiency; how can equipment be best located to provide high throughput and still be easily accessible for maintenance; or how equipment operation can be best scheduled to improve overall capital utilization. In order to answer these questions, different alternatives can be tried out on the computer, saving months or years of physical experimentation time, and millions or even tens of millions of dollars of experimental materials and equipment time.

In Knowledge Management (KM), ever more transient users can access vast sources of data, information and knowledge in real time to enable them to make more informed and higher quality decisions. This information is wide in scope and sufficiently deep to enable one versed in the use of such technology to make and execute decisions with unparalleled ability. Considering that the value of a corporation is more and more dependent on intellectual assets (patents, know-how, trade secrets, processing and product knowledge, best-known methods, etc.) than on capital assets (equipment, buildings, rights of way, etc.) it is not surprising that significant attention is now being paid to knowledge management.

Both OM&S and KM can be beneficially applied to the many domains of manufacturing including scheduling, using the theory of constraints tied to operational models; enterprise integration tied to enterprise models; electronic commerce; capacity planning and factory layout improvements, tied to comprehensive factory models; improved equipment utilization and performance, tied to equipment and material handling models. All these domains can benefit significantly from OM&S and KM. Using these methods, we can now start to overcome some of the limitations we face as yields approach 100%, as factory automation approaches an economical limit, and as increased wafer diameter and increased package complexity continue to add to the cost of running a large factory.

Some examples of how these two information and knowledge capabilities can be used to help improve operational efficiency are illustrated below.

Operational Modeling

OM&S is used widely in process development, wafer fabrication, assembly test, manufacturing support, and other parts of the manufacturing enterprise. Savings accrued through the use of OM&S can be substantial, in the hundreds of millions of dollars.

Generally speaking, OM&S capabilities are directly linked to improvement of major factory performance metrics: cost reduction, delivery improvement, quality improvement or product performance improvement. Factory improvement issues are often stated thus: "If I change this and that, how does the result affect my bottom line performance?" or "What if I did this instead of that (if I added or removed people from the line; if I laid out the equipment differently; if I used this strategy vs. that one to schedule downtime, and so on), how would factory performance be changed?" Consequently, OM&S programs are often called "what-if" scenarios. They are used to save time and money. Running a physical experiment, i.e., re-laying out a product line, can take months or years compared to running a simulation, which can take minutes or hours. Or, running a physical experiment can cost too much. Running a single experiment in an operating fab could cost hundreds of thousands of dollars.

Application of OM&S Technology

The following are examples of how OM&S technology can be applied:

- Comparison of Continuous Flow Manufacturing (CFM) to current Functional Flow Manufacturing practices in the production of Single Edge Connector Cartridge (SECc) modules may be applicable to other manufacturing facilities.
- Dedication of particular stepper lenses to particular lots in fabs to improve overall factory performance.
- Increase in WIP turns using full factory simulation to enhance use of information to improve performance.
- Evaluation of the effects of lot size on factory performance to determine optimum lot size.
- Evaluation of the effects of modifying operational policies on scheduling use of factory equipment to increase utilization without adding more equipment.

More detailed discussions of applications of operational modeling may be found in Court Hilton's paper entitled "Manufacturing Operations System Design and Analysis" and Karl Kempf's paper "Improving Throughput Across the Factory Life-Cycle" also appearing in this Q4'98 issue of the *Intel Technology Journal.*

Note that all of the above examples are specific applications; they do something for someone who has a specific issue to resolve. As such, they are highly beneficial. But the real payoff comes when all these applications are linked through some integrated, hierarchical model. The benefits of such a model can be imagined by comparing it to Microsoft Windows*. In Microsoft Windows, each application (Word*, Excel*, PowerPoint*, etc.) is individually very useful, but the ability to share textual and image objects between applications greatly enhances the whole. The total Windows environment is more than just the sum of its parts.

So part of the evolving OM&S effort is aimed at defining a modeling hierarchy, and establishing the links and infrastructure between modeling elements, to make the entire modeling environment much more than the sum of the individual components. This is schematically illustrated in Figure 4, where the NOW environment shows individual models, distributed through the manufacturing enterprise, and the FUTURE scenario shows an evenly distributed, linked hierarchy of models.

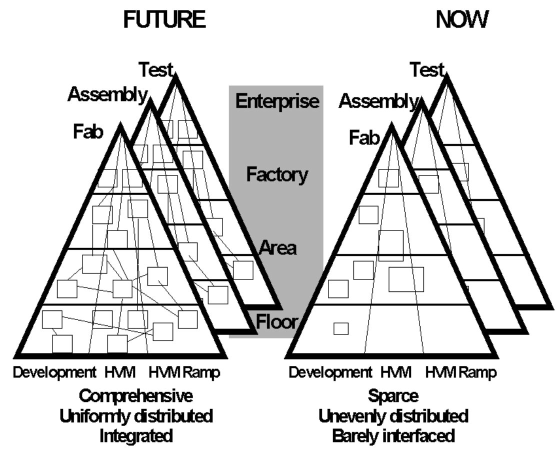

Figure 4: Modeling hierarchy

The scope of operational modeling is very broad, as illustrated in Figure 5. For convenience, the operational environment has been divided into three roughly equal domains: those dealing *directly* with product (the PHYSICAL DOMAIN), those dealing with the data and information associated both with the product and with the factory itself (the INFORMATION DOMAIN), and those dealing with background and support issues (the INFRASTRUCTURE DOMAIN). Each of these domains is itself sub-categorized, as shown in Figure 5.

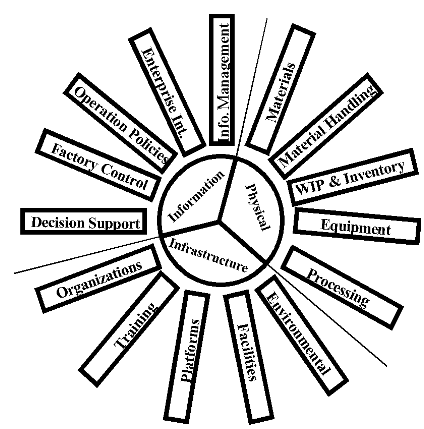

Figure 5: Model scope

Each sub-category is made up of sub-sub-categories, and so on, until one reaches the lowest level of the model hierarchy. Hence, each topic can have applications, roadmaps, goals, interfaces, etc.; the question is, how many of these topics

* Other brands and names are the property of their respective owners.

have common elements and should actually be integrated with one another. This integration is both lateral, meaning across equivalent levels of hierarchy, as well as being up and down the chain of model hierarchy. It raises interesting philosophical questions about model integration, as well as deep practical questions of how one may make modeling capabilities more cost-effective and efficient.

Knowledge Management

Whereas OM&S technology provides a fairly direct link between the capability of a technology and factory performance, knowledge management (KM) technology is one step removed from such a direct link. Indeed, KM is a logical counterpart of physical asset management, the leveraging of our physical capital (land, factories, computers, equipment, etc.) to improve profitability. KM leverages "knowledge capital" (patents, trademarks, know-how, competencies, skills, tacit or unwritten knowledge, relationships, etc.). Since, at the present time, the value of these intellectual assets is not really understood, the first goal of KM is to define a set of metrics that allows one to know even if there is any leverage to intellectual capital.

One rough estimate may be made by comparing the value of a company in the eyes of its stockholders to the paper value of the company's physical assets. In the case of Intel, the stock value (shares outstanding times price) is about $120 Billion, while the physical assets have a value of about $25 Billion. The difference, about $95 Billion, or four times the physical asset value, may be ascribed to non-physical assets!

KM capabilities may be defined using the following model. KM is divided into four large domains: the creation of knowledge, the capture and structure of knowledge, the dissemination of knowledge, and the application of knowledge. Some attributes of each of these four categories are shown below in Table 5.

The two areas that require most attention are items 2 and 3 in Table 5: the collection, structuring, and indexing of knowledge, and the secure, rapid dissemination of knowledge to potential users. Of primary interest are metrics: understanding how to value the intellectual assets of the enterprise, and indexing: the categorization of knowledge for rapid and ubiquitous application. Also of great significance is the knowledge tool environment. Much like the information tools of prior generations, knowledge tools are rapidly emerging and evolving. We expect that a knowledge tool environment similar in concept to the Windows* information environment will emerge, thereby allowing us to exchange knowledge objects in much the same way as we already exchange information objects.

1. Knowledge Creation
- Research
- Brainstorming
- Strategizing
- Synthesizing
2. Knowledge Structure
- Data and knowledge databases
- Indexing
- Training development
- Report generation
- Knowledge management tools
3. Knowledge Dissemination
- Inter- and Intranet
- Education and training
- Electronic mail
- Reading
- Browsers and interfaces
- Security precautions
4. Knowledge Application
- Problem solving
- Strategizing
- Decision making
- Managing and metrics

Table 5: Knowledge management domain

Some potential areas where knowledge management can be applied are as follows:

- Understanding and matching of core competencies of individuals with attributes of job needs.
- Providing a "Knowledge Atlas," a visual environment in which employees can guide themselves to find knowledge items, for example, "how do I do this?"; "who do I see to do that?"; or "who is the expert on this?".
- Developing tools that leverage an employee's job skills, allowing people to take on more responsible jobs using knowledge assistants for help.
- Better problem solving by providing access to vast and comprehensive knowledge bases of past occurrences, tied to the nature of a problem rather than to simple keyword searches.

Knowledge management tools will help make us a more efficient company by providing access to knowledge to people who need it, wherever they are and whatever the problem set. We should then be able to make faster and wiser decisions, resulting in significant improvements in factory and even enterprise performance.

Organizational Issues

Pursuit of information and knowledge technology, as given in the examples above, is not free. In particular, in addition to the obvious need for technical skills, there is a need to understand and respond to the managerial and organizational skills required for success.

At one time, the resources required to operate a factory consisted almost universally of people who had their hands on the product: moving it, processing it, assembling it, storing it. Currently, the trend is towards having a greater percentage of the workforce spending time on the processing of data and information. They gather data, analyze data, and convert these data to information. This information is then stored, transmitted, and disseminated, so that decisions can be made and our knowledge increased. Meanwhile, the total workforce is decreasing through physical and logical productivity improvement.

The result of these two trends is schematically illustrated in Figure 5 below. The total workforce is decreasing, while the percentage of IT and software personnel is increasing.

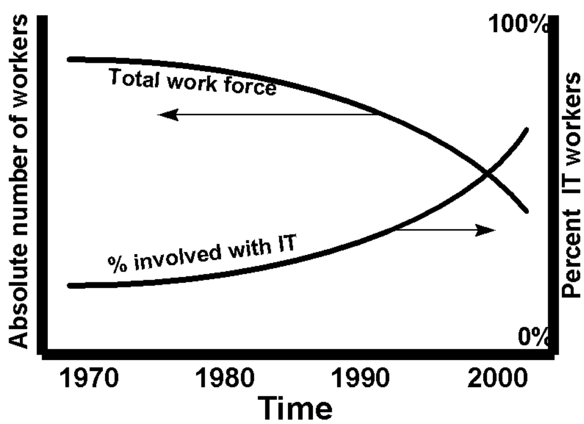

Figure 6: IT headcount projections

There are two personnel issues to confront as a result of these trends: the first is the evolution of the factory workforce from process-centric to one that is more information-centric. The processing domain is equipment dominated, where our equipment suppliers own the core competencies. As more and more information processing is incorporated into the factory, more technologists will be necessary in the IT processing field. However, this problem is fairly manageable; Intel is an expert at managing technology.

The real issues are those of organization and management. Managing process is straightforward: align the management organizations functionally, for example, with cross-cutting metrics such as yield, cost, delivery, etc. Managing the information organization is different, however. The cross-cutting disciplines such as platforms, software, and databases are not conducive to factory management, but the information technology does not map well to the traditional metrics of yield, delivery, etc. Furthermore, the skills of management need to be different. Management needs to be more proficient in IT skills; their current skill set is technologically oriented towards processing technology.

These management and organizational issues need to be dealt with concurrent with the growth of IT technology.

Conclusions

It seems clear that our industry is departing from at least some of Gordon Moore's earlier quantitative predictions. One of these is illustrated in Figure 6. Gordon's 1974 tongue-in-cheek but genuine extrapolation of wafer sizes suggested that by the year 2000, we would have 57-inch diameter wafers! Clearly, this is off by about an order of magnitude. Yet simple extrapolations of Gordon's trends does lead to qualitatively correct predictions.

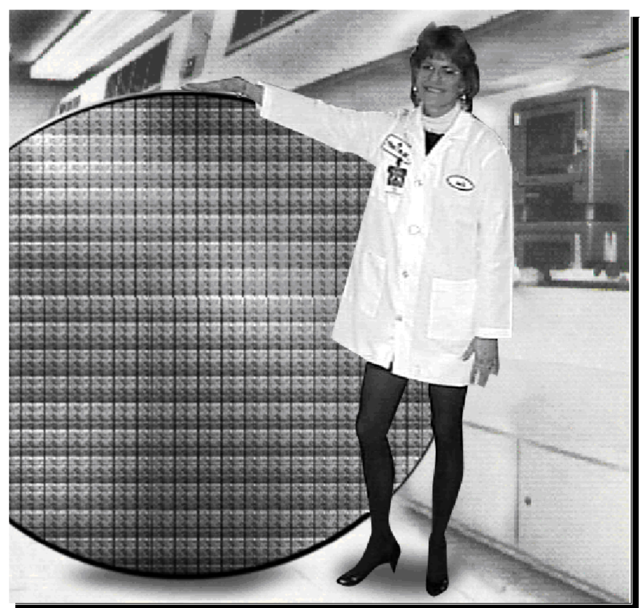

Figure 7: "Extrapolated" Year 1999 wafer size[1]

Regardless, two trends seem inescapable: everything in the production of semiconductor devices is moving toward more expensive factories, and there is swiftly expanding use of information and knowledge to reduce costs, improve delivery, and improve quality. These two trends need to be linked to try to alleviate the effects of the former by using the latter. At the same time, one must also recognize the emergence of other forces: the need for cleaner, safer, and less energy-consuming manufacturing enterprises, the evolution and indeed revolution of materials and materials' processing, and the change from local politics and culture to global politics and culture. All these trends will result in a significantly greater emphasis being put on manufacturing as a competitive weapon in the 21st century.

Acknowledgments

I would like to thank Gordon Moore, Karl Kempf, Court Hilton, Sri Sridharan, Scot Ruska, Bruce Sohn and John Birchak for valuable discussions that helped formulate the concepts discussed in this paper.

References

1. Gordon Moore, *Electronic Materials Symposium* Santa Clara, CA, March 1998.
2. "Next Generation Manufacturing: A Framework for Action," *Agility Forum*, Bethlehem, PA, 1997.

Author's Biography

Gene Meieran received his B.S. degree from Purdue University in 1959 and his Sc.D degree in materials science from MIT in 1963. He joined Fairchild R&D in 1963, where he specialized in the analysis and characterization of semiconductor device materials. He joined Intel in 1973 as manager of Package Development, responsible for developing new lines of packages for the emerging memory and microprocessor products. In 1977, he joined the Quality and Reliability staff, with responsibility for all Intel materials, the Materials Analysis Laboratory, and for manufacturing reliability functions. He has worked in Statistical Process Control (SPC) and advanced manufacturing strategy development in Intel's Technology Manufacturing Engineering group for the past 12 years.

Dr. Meieran taught technical courses in leading US universities and has given seminars and invited talks to many international universities. He has about 50 technical awards and has received three international awards based on technical talks.

He served on the Scientific/Education Advisory Board for Lawrence Berkeley Labs and on advisory boards for several university departments. He has been Director for Research for the MIT Leaders For Manufacturing Program since 1993 and has served on numerous government and industry panels dealing with manufacturing technology and policy issues.

In 1985, Gene was appointed an Intel Fellow, Intel's most senior technical position. In 1987, Purdue University elected him a Distinguished Engineering Alumni, and in 1998, he was elected to the National Academy of Engineering.

His e-mail is gene.s.meieran@intel.com.