

Planning for the 300mm Transition

Daniel Seligson, Technology and Manufacturing Engineering (TME), Intel Corp.

Index words: 300mm, manufacturing, cost, automation, standards, consortium

Abstract

Beginning in 1993 a small group of people at Intel began thinking seriously about a transition from 200mm manufacturing to the next wafer size. By early 1994, the industry reached consensus that the right size was 300mm. Late that year, we formed a cross-disciplinary team whose responsibilities included defining the 300mm goals for each of six functional domains: equipment, automation, factory and facilities design, EHS, manufacturing operations, and materials. In addition, we built a cost model that helped inform the above process while providing indicators of success that cut across all the domains.

The defined goals were widely disseminated in the industry through various channels including supplier management organizations and consortia. The equipment selection process was augmented to include the new requirements for tools, and the processes were extended to include all tools needed for a factory, not merely the most expensive and technically sophisticated. Task forces were established to answer urgent questions about lot size, carrier design, mini-environments, and any other issues that arose. In parallel, a new international consortium was built from the existing SEMATECH infrastructure, although with more clear cut goals and achievable objectives than SEMATECH had. This new organization, called I300I, became a critical part of the industry consensus building and supplier management. While we have gone through major industry transitions between program start and now, in areas such as market segmentation, transition timing, and cost focus, the goals developed in 1995 remain sound. Today, selection teams are in place for all equipment, from materials handling to clean parts, many of them selecting tools for new semiconductor processes never implemented at Intel. The teams are synchronized to produce a first set of tools in January 2000 and high volume 0.13um manufacturing capability in the second half of 2002. The productivity, as measured by reduction in die cost, is expected to exceed our targets of 30%. For comparison, only about a 15% productivity increase was achieved in moving from a 150mm to a 200mm wafer size.

Introduction

The goal of this paper is to provide a comprehensive summary of the critical issues and key events of the transition to 300mm. We begin with the decision to move to a new wafer size and how we arrived at 300mm. We then discuss the consensus that this would be an industry transition, rather than one led by a single company. The process by which Intel came up with the requirements we tried to drive in the industry is reviewed next, followed by a brief look at the cost model and productivity. Automation, in particular automated materials handling, received a lot of focus in our planning. Its importance and implications are explained in the next section, together with an explanation of a major improvement in factory design. Other equipment, safety, and factory requirements are then mentioned, where we also provide a one-page summary of all requirements. Getting the message distributed within Intel and the industry was a huge effort, so a special section is devoted to how that was done. The two consortia, I300I and SELETE, played large roles that cut across many aspects of planning, so we also review their contributions. We then discuss the many improvements and the entirely new business processes that have been put in place to help manage the transition as a whole. In the final two sections, we cover the high-level review processes Intel has used to make overall program decisions, and the outlook for the future from the vantage point of Q4 1998 is briefly discussed.

Selecting the Next Wafer Size

Intel began 200mm manufacturing in 1993, following a two-year development effort. The larger wafer size gave us nearly twice as many die for every wafer moved, but the growth in the microprocessor business and the growth in die size led us to conclude that before the end of the decade we would need to be adding factories at the rate of two per year. The complexities of construction, site selection, staffing and management development required to support that growth suggested that it was time to begin thinking seriously about the next wafer size.

Many of the costs of manufacturing are proportional to the number of wafers moved, and not to the area processed.

So, size does matter. Initially, Intel favored 400mm and Applied Materials, the world's largest semiconductor equipment manufacturer, put forward arguments for 350mm. But the silicon wafer manufacturers told us that if we were contemplating a wafer size change within a decade, the largest wafer we could have was 300mm.

Their arguments were based on the fact that the length of the boule should increase at least in proportion to the new wafer size, and its weight therefore by the cube of the dimension. The starting material for a 300mm boule will be between 300kg and 450kg. New inventions would be required to cost effectively manufacture wafers larger than 300mm. The correctness of this view can be seen in the existence, schedule and goals of the Japanese Super Silicon Crystal Institute Corp, a partnership of government and industry that focuses on developing wafers of 400mm and larger by 2004.

Process equipment was predicted to have no particular issues with any of these wafer sizes, although initially there were fears that the gravitational sag of hot wafers in a furnace would induce slip. The flat panel display industry was at that time using substrates in which a 350mm circle could be inscribed easily, so 300mm looked well within reach technically.

An Industry Transition

The transition from 100mm to 150mm in 1983 and 1984 was led by Intel's groups in Albuquerque, New Mexico. There is no historical data on its cost effectiveness. The transition to 200mm was led by IBM in Burlington, Vermont. First tools were delivered in January 1986, and 1 Mbit DRAMs were first qualified for production in early 1998. When Intel made this transition with production beginning in 1993, die costs were reduced by 10% to 15% when comparing new 200mm and new 150mm fabs. But when making the more realistic comparison of a new 200mm fab to a 150mm factory upgraded to meet new technology requirements, no die cost improvement was achieved. IBM may have achieved a greater reduction because they converted from 125mm wafers.¹ Both transitions were unpleasant experiences for the lead company in that it had to bear the burden of development costs, manufacturing delays, and poor equipment performance, all at little or no cost benefit.

¹ Die cost reductions are driven largely by the ratio of the number of die on the final and initial wafer sizes. For the three transitions 125->200, 150->200 and 200->300, these ratios are approximately 2.7, 1.9 and 2.4 respectively. The exact ratios depend on company-specific product issues.

These unfortunate precedents led us to conclude that no single company was smart enough or large enough to do a wafer size conversion by itself, and that the 300mm transition should be an industry one. The hoped for benefits were common performance objectives, shared learning, cost sharing, and more efficient and accelerated development facilitated by use of widely accepted standards.

Standards were expected to be important in such diverse areas as wafers and wafer carriers, data and material transfer protocols, safety, and others. Historically, standards had been defined *post facto* on the basis of successful implementations. Early implementers were almost guaranteed to be outside the standards. For this transition, the industry believed it would be better to define the standards at the outset and ensure that everyone adhere to them. SEMI, the international standards organization, was thus destined to play a key role.

Shared learning meant that we would do as much pre-competitive work as possible², particularly in the testing of equipment and modular components. The results of these tests would be fed back to the supplier, along with a roadmap for improvement. They would also be made available to the members of the industry group performing the testing, although at that time the group was not yet defined. The initial vision was of an Underwriter's LabTM, providing a stamp of approval instead of a *caveat emptor*.

Cost sharing really meant that the equipment suppliers would pay for the equipment development, and a manufacturers' consortium would pay for its testing. During the two previous transitions, when Intel and IBM paved the way, the suppliers were relatively small and weak. By the mid-90s, several of the suppliers had annual revenues exceeding \$1B, with R&D budgets capable of supporting development of a new generation of tools. The increased health of the supplier base can be attributed to prolonged growth in the semiconductor industry, consolidation in the equipment industry, and, in the US at least, the attentions of SEMATECH. Such a model could not have worked for the 150mm and 200mm transitions.

Finally, we recognized that having common performance objectives for tool performance would be a critical element of any industry transition. One of the lessons of SEMATECH was that everyone's secrets were the same, but they resided

² Pre-competitive work was not possible during the previous wafer size transitions. Not until the semiconductor crises of the mid 1980s were laws passed to allow the pre-competitive work such as that done at SEMATECH.

separately in each company. This compartmentalization of knowledge increased the workload on the supplier and decreased efficiency. Working through the difficult early years of SEMATECH, the US manufacturers learned to give direction to the supplier base without compromising their individual intellectual property. This greatly simplified their workload and guaranteed, on average at least, a higher quality product. Since the previous wafer size conversion had shown modest paper cost benefit and less real cost benefit, a higher quality product was badly needed.

An industry transition also meant that no one company would be first, or would need to be first. Nevertheless, by late 1994, Motorola had established itself as the industry leader, although they participated in the various standards' bodies and consortia. Their aggressive schedule and forcefully stated opinions helped accelerate the pace of decision making on important issues. Other companies chose to say that they wouldn't be first. Later, when Motorola's program was put on hold, enthusiasm still remained high. The pundits concluded that no company wanted to build the first 300mm fab, but neither did any company want to build the last 200mm fab.

Developing Intel's Requirements

In late 1994, an *ad hoc* 300mm discussion group, involving Process Equipment Development (PED) and California Technology and Manufacturing (CTM), realized that because 300mm fabs were still far off in our future, it was an opportune moment to begin developing a vision of a 300mm world. We invited representatives from all of Intel's stakeholding groups to the discussion. Since we had no authority to command, those who joined us were only the 'can do' types who wanted to be there, so progress was swift. By early 1995, we had built a group that covered major problem areas and so took on the task of writing a handbook defining our requirements in six different functional domains: equipment, automation, factory and facilities design, EHS, manufacturing operations, and materials. By mid-year, we had named ourselves the Cross Functional Working Group (CFWG) laying ourselves crosswise on a set of committees (SCS) whose job it was to select new tools and manage roadmaps for the evolution of process equipment, automation, factory design, chemicals and their use, manufacturing operations, and so on.

The requirements for 300mm automated materials handling, for example, would necessarily add requirements to the process equipment whose roadmaps were owned by an equipment SCS. This necessitated many rounds of negotiation among the various SCSs, complicated by the fact that the SCSs had neither the time nor inclination to listen carefully to

arcane discussions about a wafer size that existed on no Intel plans. Nevertheless, over the course of 1995, each domain developed its requirements' package, reviewing and refining the material at the CFWG, an audience increasingly attentive to the full spectrum of issues relevant to a 300mm transition.

Several problems emerged whose solution was outside the range of any one existing group. To solve these problems, we established task forces. Questions of lot size, lot buffering, lot carrier design, reticle carrier design, assembly components, and mini-environments were answered in this way. When relevant, our task forces would work closely with industry standards' organizations too. In the case of lot size, industry opinion ranged from 1 to 50 wafers per lot in mid-1994. Single wafer processing experiments at Texas Instruments (TI) suggested that the low latency and short throughput times could be advantageous. Motorola advocated single wafer transport initially. Meanwhile, the experience of Intel and the large DRAM manufacturers was that larger lot sizes were favored.

Trying to separate the emotion from the facts, we did a numerical factory simulation, varying lot size from 13 to 50 wafers, assessing factory capacity and several cost drivers. Small lot size implies more lot moves per unit time. We set the lower bound of our study to 13 because we knew that smaller lots would overwhelm any realizable automated materials handling system (AMHS). The lot size team looked at issues including ergonomics, metrology, floor space, supplier capability, capital equipment costs, labor, and total wafer cost. They drew three important conclusions. First, ergonomic considerations would preclude regular manual handling of a lot with 13 wafers or more; therefore, extensive AMHS would be needed for any lot size. Second, total processed wafer cost decreased with increasing lot size, so larger lots were favored. Third, AMHS on suppliers' drawing boards would be unable to move 50 wafer lots; they were too heavy. The answer, therefore, was 25 wafers per lot, and that became the Intel position.

A memorable event took place in February 1995 at the SEMI Standards meeting in New Orleans. In a crowded stuffy room, representatives of Motorola, TI, Intel, the Japanese DRAM manufacturers, and several dozen equipment suppliers debated the merits of different lot sizes. Motorola had retreated from its vision of single wafer transfer and had begun promoting 13. Someone asked, "is there anyone here that favors 13 over 25?" In response, only TI would even entertain the discussion. Motorola was defeated on the standards front; they could marshal the industry by dangling the hope of purchase orders, but not on the strength of their technical arguments or their negotiation skills. Gradually, Motorola

began to withdraw further and further from the industry center choosing to go it alone, convinced that the consortia and standards activities would only slow them down. When the dust settled, the industry standard for lot size was not one size, but two, 13 and 25. This added unnecessary development costs in both dollars and time for the suppliers of process tools, loadports, carriers, and material handling systems, to name a few.

The lot size debate is not as clear as portrayed here. Semiconductor manufacturers with business models different from an Intel or DRAM manufacturer, who make large volumes of a small number of products, might arrive at a different lot size. These business model issues were not factored into our analysis. In addition, the impact of the assumptions we used, particularly those having to do with lot-based rather than wafer-based metrology sampling plans, was never examined. Nevertheless, there is wide consensus today that 25 is preferred, although it introduces difficulties for companies manufacturing a large number of different short-run products.

The methodology used to make the lot size decision is typical of that employed for the other truly cross-functional decisions. It is not within the scope of this paper to include a detailed discussion of each, however.

The Cost Model³

As mentioned earlier, the move to 300mm was driven by the fact that it would provide an exit from the anticipated scenario of building two factories per year. Cost reduction was a factor, but more of an opportunity than an overriding force. To understand and then exploit the opportunity, the CFWG took on the task of developing a 300mm cost model. There was scant documentation of the cost targets for the 200mm transition and no documentation of performance against those targets, other than the overall null result. We believed that a more thorough modeling effort, followed by aggressive goal setting based on model parameters and detailed management of progress towards those goals, would yield a more felicitous outcome.

Starting from a 0.35um 200mm wafer cost model, which was all that was available at the time, we examined the major line items. These line items included, among others, capital depreciation, direct and indirect materials, utilities consumed,

factory infrastructure, labor, and site overhead. We challenged ourselves to imagine a 0.35um 300mm cost model and to discover the mechanisms that underlie the scaling of each line item as we move from the smaller to the larger wafer size. By operating at the same technology generation, we were able to separate the wafer size issues from the process technology issues. For each line item we had therefore not merely a scaling number, but a concept.

This turned out to be quite powerful, for instance in modeling labor costs. Received wisdom had it that labor content per wafer increased 20% when we had increased wafer size in the past. So, we put into our model an equation looking something like this:

$$\text{labor}(300) = 1.2 * \text{labor}(200) \quad [1]$$

Over time we realized that this factor of 1.2 was a parameter,

$$\text{Relative Labor} = 1.2 \quad [2]$$

that we could control, not necessarily something predetermined. Achieving less than 1.2 would require that our operations' groups put plans in place to make it happen, possibly in conjunction with an AMHS group. Operations absorbed this and asserted that the right goal was parity or better as in,

$$\text{Relative Labor} \leq 1.0 \quad [3]$$

In this way, working with stakeholders, we set goals for all line items and for the groups that had a stake in them. The groups have gradually accepted more aggressive goals as they see the need, the opportunity, and the way.

Table 1 is an example of the methodology applied to a 0.25um DRAM process⁴. It illustrates that the scaling is substantially different for different line items. It is typical of what we see in the full-blown model Intel uses for planning purposes.

Line Item	200mm Scaling Factor Test wafers 45				
	300mm				
	\$	%		\$	%
Depreciation	793	41	1.50	1189	35
Labor	232	12	1.00	232	7
Maintenance	155	8	1.50	232	7
Consumables:					
Direct materials	90	5	4.50	405	12

³For a more complete analysis, see *Semiconductor International*, January 1998, Daniel Seligson, "The Economics of 300mm Processing."

⁴The underlying 200mm DRAM cost components are drawn from Jack Saltich's paper in the proceedings of ISSM '94. This table was published in *Semiconductor International* in January 1998.

Test Wafers	45	2	4.50	203	6
Indirect materials	445	23	2.00	890	26
Other	174	9	1.30	226	7
Totals	1934	100		3378	100
Equivalent				1407	73

Table 1: 0.25um DRAM cost model for 200mm and 300mm wafers

Finished wafer cost in a new factory is dominated by equipment capital cost. Meanwhile, the desire to minimize the number of new factories was the original driving force behind the move to 300mm wafers. For this reason, two parameters that have received special attention have been the Relative Capital Cost and the Relative Footprint of the toolset. With X referring to either Capital Cost or Footprint, these parameters are defined as

$$\text{Relative } X = \frac{X(300)}{X(200)} * \frac{\text{OutputCapacity}(200)}{\text{OutputCapacity}(300)} \quad [4]$$

Practically speaking, Relative Capital Cost is the ratio of the capital costs required to build a 300mm and a 200mm factory, each having the same number of wafer starts per unit time. Similarly, Relative Footprint is the ratio of the two factory areas. While defined for the factory in aggregate, these parameters can be measured for any individual tool. A customer can easily compare product offerings from multiple suppliers if consistent 200mm normalizations are used. The model parameters are easily visualized, as are their knobs.

One of the most powerful messages of the 300mm transition emerged from the realization that the value of these two parameters would determine whether the transition met its twin goals of increasing factory and capital productivity. The die cost model could thus be used to identify the parameter values required for success. The linear die cost model, which predicts costs given parameter values as input, could be inverted to produce parameter values given cost as input. Contours of constant die cost are lines in the space of Relative Capital Cost and Relative Footprint. Our senior management challenged us to deliver more than 30% die cost reduction, more than twice the reduction that had been expected in the 200mm transition. The inverted model told us that the Relative Capital Cost and Relative Footprint had to be below the red line in Figure 1. (In this figure, each point represents a different kind of process equipment. The area of plotted points is proportional to the product of individual tool cost and the number of tools required for a high-volume factory.)

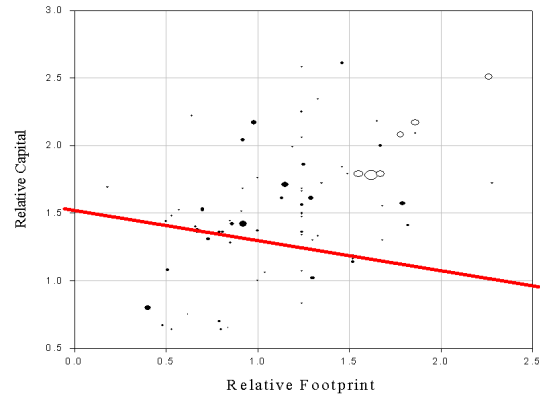


Figure 1: Productivity scaling factors for tools in a 0.18um logic process flow (data from early 1997)

This in turn gave us a simple tool to communicate with suppliers and the industry. On the whole, the new toolset needed to be below the line or we wouldn't get the sought after return, and we wouldn't make the transition. The supplier could quickly determine, by using equation 4, where they stood with respect to the line. The line became a high-level design target which provided guidance on the tradeoff between cost and footprint. Further simplifying the message for greater impact, we focused on a single point on the line, Relative Capital Cost = 1.3 and Relative Footprint = 1.0. Our vision of the future consisted of a 300mm factory that was the same size as our 200mm factories, producing the same number of wafers per week, and requiring a capital investment no more than 30% larger, as shown in Figure 2.




Figure 2: The macro view of the 300mm vision in which 200mm and 300mm factories are compared

When estimating die cost, two different scenarios are considered. In the first, you compare new different wafer size factories. This is the Greenfield Scenario, and our target for this is a 30% die cost reduction, obtained at a Relative Capital Cost of 1.3 and Relative Footprint of 1.0. In practice, the Reuse Scenario is more realistic. In this, you compare a new 300mm factory with an existing 200mm factory, one whose equipment set has been upgraded to meet the new technology requirements. Upgrading one of our 0.18 or 0.25um production lines to 0.13um technology results in a line whose capital basis is about half that of a new 200mm line. Figure 3 shows die cost vs. Relative Capital Cost for both the Greenfield and the Reuse Scenarios. It illustrates the point that the Relative Capital cost must be approximately 1.3 in order for the transition to achieve any real return. It also explains why the 200mm transition, with its 15% die cost improvement for the Greenfield Scenario, returned no benefit in practice.

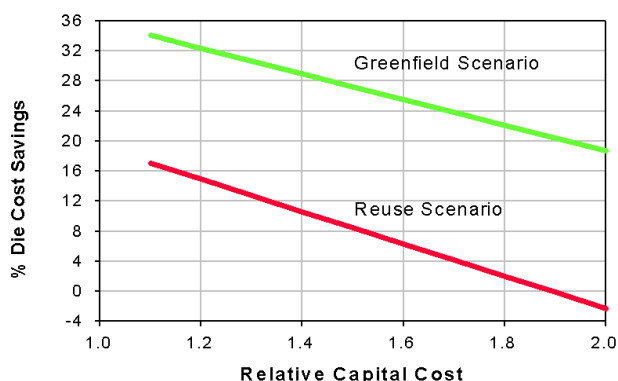


Figure 3: Die cost reduction for two conversion scenarios. Cost is important, but other measures are important too. What we are really talking about in this transition is increased productivity, where productivity is defined as follows:

$$\text{Productivity} \equiv \left(\frac{\text{something}}{\text{die out}} \right)^{-1} \quad [5]$$

Capital productivity increases when the Relative Capital Cost decreases. Factory space productivity increases when the Relative Footprint decreases. Some of the other parameters we have measured and tracked include the productivity of chemicals, natural resources, utilities, and labor, as mentioned earlier. Just as we set values for the capital and footprint productivity and put plans in place to achieve them, e.g., through our supplier management organizations, so also have we done for these other measures of productivity. The 300mm transition is a watershed of productivity improve-

ments for the industry. To date, all of the productivity targets we set in 1995 appear realizable or very nearly so, suggesting perhaps that we didn't reach far enough.

Over the years of 300mm planning, responsibility for managing the cost model has moved from engineering to finance, which is where it resides today. Its evolution has depended on maintaining a close relationship with engineering.

AMHS and a New Factory Design

In 200mm, we saw the first widespread use of AMHS. There were industry standards for such simple but important features as the height at which wafer carriers, or pods, rest on tools and the orientation of those pods with respect to the tools. But these standards were not implemented uniformly, making intrabay AMHS solutions custom, and therefore expensive, in every case.

Based on favorable outcomes with 200mm intrabay delivery to diffusion furnaces, with improvements in labor and capital productivity, we worked hard to define a set of AMHS standards that would facilitate 100% intrabay delivery. This was also desirable from an ergonomics point of view, as discovered in the lot size analysis⁵. It was important to get the industry to agree on what was required and to communicate those requirements to the suppliers before the equipment arrived. Such a set of standards, focussing on the tool loadports, would decouple the AMHS from the equipment, but it would also be essential that the standards were implemented by all suppliers. One example of how AMHS standards lower cost is that they enable a mechanically simple 1-axis robotic transfer from AMHS to tool to be used, rather than the 6-axis robots required in our present day 200mm implementations. For nearly two years after the standards were defined, managers at Intel would ask, "How many of the tools will arrive with the standards in place?" The answer, "They're requirements, so 100%," would be followed by challenges based on prior history of meeting such commitments, particularly in the area of automation. Today we all accept that the answer is 100%. That result is the product of extensive internal and external education forums and supplier management programs. Figure 4 illustrates the main points of the AMHS and loadport interface standards.

⁵ We struggled to find the right justification for 100% intrabay delivery. A return on investment analysis depended on too many unknowns. Ultimately, Manufacturing Operations' policy of striving towards an incident and injury free workplace provided the justification we have come to accept.

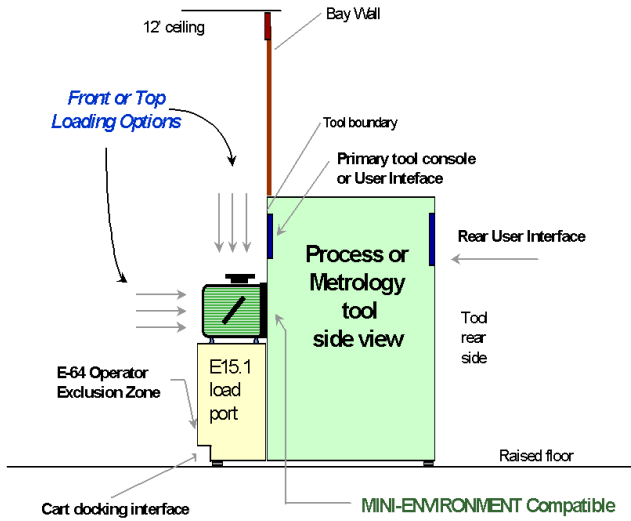


Figure 4: Features of the AMHS and Loadport interface standards (picture courtesy of D. Pillai)

Independent of 300mm planning, Intel has established an ongoing benchmarking program whose objective is to gather

information about the relative performance of Intel against other semiconductor manufacturers. In the mid-1990s, while 300mm planning was blooming, this group learned that our fab construction costs were quite high compared to the best manufacturers. A task force was established to develop a new factory design for 300mm, setting cost targets 30% lower than our most recent 200mm factory. Prior to this goal being set, the initial value for the Relative Factory Costs in the cost model was 1.08, much worse than the task force's target of 0.70. The value 1.08 was based on historical precedent. This serves to illustrate how costs had spiraled upwards in the past and how 300mm successfully reversed the trend.

Similar to the CFWG, the task force drew on all stakeholders, with a surprisingly large amount of input from AMHS. The result (see Figure 5) was a unique, and for Intel, radical design. It disposed of many features that had previously been considered essential. Further, it included novel solutions to problems brought about by 300mm, in particular problems related to the fact that lot storage space requirements would increase significantly. The final cost, although still only an estimation, was 0.62, beating the task force target.

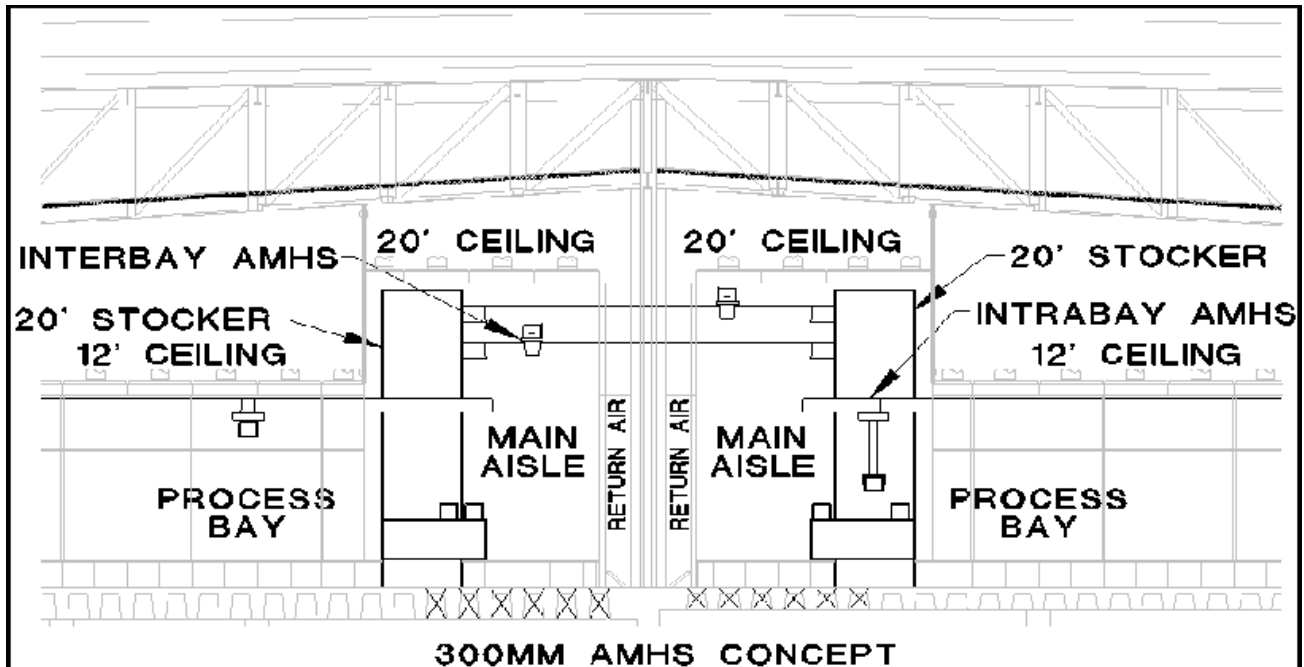


Figure 5: Cross section (transverse to the main aisle) of the new 300mm factory concept. Note the tall center section providing efficient storage of WIP. Figure courtesy of Noel Acker.

Other Equipment Performance Requirements

Chemical usage and environmental stress had become an increasingly large issue, accompanied often by negative publicity. Analyses showed that approximately 30% of our factory building costs were tied up in facilities that managed the ebb and flow of chemicals and utilities. We extended the notion of productivity to these items, establishing targets for the use of hazardous air pollutants, electrical power, scrubbed exhaust, and others, all measured relative to 200mm and normalized to capacity. For some of these, we didn't know whether the targets were achievable, and the measured data is not yet in, but the projections are that we will come close, resulting in an overall reduction of more than 50% per die.

In order to reduce the costs associated with the installation and qualification of tools, we developed a set of guidelines for standard practices, as well as setting relative targets calling for a 30% reduction in time and cost. We have since been working closely with suppliers to help them understand the

guidelines, forecast their costs, and address issues when the forecasts fail to meet the goals.

The handbook of these and other requirements formed a thick presentation that took more than four hours to deliver when first rolled out to some suppliers in November 1995. This was useful but tedious. We realized that we needed a one-page document understood by everyone and capable of being carried around in an executive's shirt pocket. Table 2 is the one-page summary containing all the key requirements from all the participating groups. We identify the relevant industry specs, and we include a column identifying the I300I position on each requirement. (I300I is an international consortium that is explained later.) The I300I column emphasizes the fact that our requirements are more than Intel's requirements. Ultimately, they became requirements for the entire industry.

300mm Equipment Performance Requirements				
Domain	Topic	Specification	References	I300I Position
EHS	Safety and Ergonomics	100% Compliance	SEMI S2-93A and S8-95	same
	EU Machinery, EMI and Low Voltage Directives	100% Compliance	EU/CE, E33-94	same
	Relative Recordable Injury Rate	<1.0		
	Relative Hazardous Air Pollutants Emissions per wspw	< or =0.5		same
	Relative Perfluorocarbon Emissions per wspw	< or =0.5		same
	Relative Volatile Organic Compounds Emissions per wspw	< or =0.5		same
	Relative Emissions and Utilities Consumption per wspw (e.g., water, process electrical, scrubbed exhaust, bulk gases)	< =1.0		same ¹
Equipment	Relative Factory Footprint per wspw	<1.0		same
	Relative Capital Cost per wspw	<1.3		same
	Capital Equipment Availability	>90%		same ¹
	Tool Install/Qual Cost as % of tool cost	<6%		same
	Relative Tool Installation Duration	<0.7		same
	Relative Tool Qualification Duration	<0.7		same
	Relative Spares and Maintenance \$/wspw	<1.0		same
	Relative Monitor Wafer Usage per wspw	<0.25		same
	Edge Exclusion	3mm		same
	Electrical Voltage Drop-out Immunity	100% Compliance	IEEE-446 (CEEMA) Curve	same
Mini-Environment Compatibility	100% Compliance	SEMI E44, E52 (FIMS)	same ¹	
Automation	Load/Unload Port Interface	100% Compliance	SEMI E15.1-0697	same
	Integrated Box Opener/Loader	100% Compliance	SEMI E63 (BOLTS)	same ⁴
	Parallel I/O Interface for Automatic Load/Unload	100% Compliance	SEMI E23-96	same
	Cart Docking Interface	100% Compliance	SEMI E64	same ⁵
	Lot Size	25 wafers	SEMI E1.9-0697	13/25 ⁶
	Lot Buffering	2 Unicassette Ports minimum with capability of cascading. (Note: High runrate tools will need special attention.)		same
	Carrier Architecture	Horizontal transport, Front loading, Cassetteless, 10mm pitch, Kinematic couplings	SEMI E62, E57-1296, E47.1-0697, E1.9-0697	same ⁷
	Embedded Controller	SECS II, GEM, HSMS	SEMI E5-96, E30-95, E37-95, E37.1-96	same ⁸
Network Connectivity	RS-232 & Ethernet with TCP/IP			

Table 2: 300mm equipment performance requirements

Communicating the Requirements

The SCS-ratified CFWG requirements were communicated through an extensive infrastructure of existing supplier management channels, primarily through the department that is now called Capital Equipment Development (CED). A small group of 300mm content experts teamed with engineers and managers responsible for the performance of individual suppliers to deliver the 4-hour requirements package to each of more than a dozen suppliers by the middle of 1996. These meetings served to educate our own teams as well, since 300mm was new and generally of low priority compared to more pressing 200mm programs. We used these meetings to understand the suppliers' 300mm program status and their reactions to our requirements as well. Management of supplier performance to the requirements was, at least in principle, done not by the 300mm content experts, but by the individual supplier team. We produced a video to facilitate delivery of the handbook content both within Intel and to the supplier base. We have also held topic-specific meetings with a large community of suppliers, the purpose being to clarify particularly subtle aspects of the requirements.

The Consortia

In 1987, a group of American semiconductor manufacturers and the US government formed a consortium aimed at reversing the losses the US industry had suffered at the hands of the Japanese. The consortium was called SEMATECH (Semiconductor Manufacturing Technology), which celebrated its move into its Austin headquarters in November 1988. SEMATECH focussed on rebuilding the infrastructure of the American equipment suppliers. There is little doubt among the faithful that it was successful. Some argue, however, that the cost was excessive and that a set of conference rooms, a coffee machine, and a legal umbrella would have accomplished as much and would have saved us building a high-overhead factory.

By the mid-1990s, two developments were shaping some major changes within SEMATECH. Firstly, 300mm was identified as the next wafer size, so SEMATECH initiated a 300mm program. As mentioned at the outset, there was agreement amongst most players that 300mm should be different, and in particular that the players should cooperate as much as legally possible. Along those lines, in late 1994, a SEMATECH task force ran a series of meetings organized by process type. Suppliers and member company (MC) representatives were invited, and the agendas included identifying critical problems that needed resolution, prospective MC schedules, and performance specifications. In parallel, SEMATECH was preparing cost analyses, based largely on estimates from the

suppliers of what they wanted to charge. A summary meeting was held at year's end, during which, among other things, suppliers and customers resolutely rejected the notion of bridge tools⁶.

The second development was that the SEMATECH MCs and the US government decided to end their partnership. This naturally led to the idea that a SEMATECH 300mm program should have international membership. Combining that notion with the desire to have a consistent testing methodology, the international 300mm initiative or I300I was spun off from SEMATECH. It would use SEMATECH infrastructure, but would be international and would secure most of its funding from its own MCs and not from SEMATECH resources. Ultimately, it grew to 13 MCs, 6 outside the US, but none from Japan. Dues were approximately \$2M per MC per year. Initially, the scope of I300I was limited to delivering test results on a set of 0.25um 300mm tools by the end of 1997.

As SEMATECH began to deliver results in 1990, the Japanese bubble economy of the 1980s began to deflate. By the mid-1990s this was compounded, for our industry, by the amalgam of advancing semiconductor companies from Korea and Taiwan. Rather than join I300I, with its distinctly American flavor and slightly jingoistic aftertaste, the Japanese formed their own SEMATECH-like organization called SELETE, the Semiconductor Leading Edge Corporation. As SEMATECH did for the Americans, so was SELETE meant to do for the Japanese, i.e., give the (Japanese) MCs and (Japanese) suppliers a competitive advantage over their counterparts in other countries. Annual dues were approximately five times higher than those of I300I, but expected outcomes were greater: they included improvement and development of new tooling. SELETE was following the SEMATECH model while SEMATECH itself was moving away from it. Initially, SEMATECH had tried to rally the Americans with wartime fervor, citing the Manhattan Project and other heroic feats. SELETE perhaps was trying to recreate the heroic deeds of Japan's VLSI program in the late 1970s and early 1980s, which developed the 64k DRAM and which was responsible for vaulting the Japanese into a leadership position in semiconductors.

⁶ Bridge tools are used across a wafer size transition; they had been commonplace at the 200mm transition, with the result that 200mm tools offered little economic advantage over 150mm tools. Bridge tools are making a bit of comeback today, but with a few exceptions, a single tool cannot meet our performance expectations at two different wafer sizes.

The stated purpose of I300I was to distribute, over the 13 MCs, the costs of evaluating tools, as well as to provide those evaluations in a consistent manner. To that end, an extensive Demonstration Test Methodology (DTM) was developed that provided MCs and suppliers with a common language to discuss and measure tools. The scope of the DTM, and the number of tools tested determined the I300I budget. Testing of tools was the prime deliverable, and the DTM has been applied to 60 demonstrations, of varying degrees of extensiveness, as of this writing.

There were two additional somewhat unexpected results that I300I delivered that have been very valuable. First, the schedule they set for the demonstrations became the *de facto* schedule for the industry conversion. Motorola was trying to move the industry forward with unilateral pronouncements, supplier meetings, and promises of purchase orders to come. Other manufacturers were announcing their own schedules with less fanfare. The result was that there was no industry schedule for this industry transition until I300I published a schedule for the 1997 demonstrations. This served to align the demands of the manufacturers and the readiness of the suppliers. Some have said that it introduced delays because leading customers backed off until after the demo results became available, but these original schedules were marketing schedules, not manufacturing schedules. In 1998, now that the time line has been pushed out twice, many suppliers are irate with I300I for publishing such unreliable schedules. However, I300I has merely repeated what its MCs have told it, and the forces driving the delay are powerful indeed, although analysis of those is beyond the scope of this paper.

The second additional result was that I300I has become the *de facto* standards body for 300mm worldwide. As part of its effort to develop the DTM, it had to identify *what* would be tested. For each different tool type, appropriate performance measures and target values, such as, etch rates, process uniformity, or capital costs per wafer per unit time, were developed⁷. In addition, I300I worked to develop a set of requirements (AMHS, safety, and other areas) common to all tools. A high degree of compliance to the common requirements, as determined by I300I, became the key to further testing. These requirements, agreed to by the 13 MCs, became the defining standard of 300mm. SEMI is the official international standards-setting organization of our industry. The SEMI ballot-

ing process was used to ratify what the customers had agreed upon at I300I.⁸

The fact that we had two consortia, I300I and SELETE, served to complicate matters because the I300I requirements could not truly be the industry requirements until they were negotiated with the Japanese. Additionally, SELETE explicitly stayed out of the game of setting standards, leaving that to another organization called J300. Teams from I300I and J300 began meeting by mid-1996, actively trying to avoid a built-in divergence in requirements. At SEMICON West in 1997, the two organizations announced and distributed their single set of requirements for the 300mm generation. These were called the Global Joint Guidance on 300mm Semiconductor Factories.⁹ They are barely distinguishable from the set of requirements developed by our CFWG. At SEMICON Japan 1998, we expect a similar document to be published covering the requirements of 300mm assembly equipment.

The path from CFWG to Global Joint Guidance is our spec or standards pipeline. Following problem identification, we form a multidisciplinary group within Intel to develop a need-driven schedule to solve these problems. Some members of this group need to be well informed about and participating in relevant external activities and industry groups. If there were no such activity, they would initiate it. The group will meet regularly and review progress at the monthly CFWG meetings and if necessary with other stakeholders within Intel. Once we reach consensus within Intel, we try to influence the industry groups. The CFWG has imposed the requirement on the Intel group that their solution must ultimately be the industry standard, so there is give and take until we're all in agreement. It is the case today that for every topic of interest to us there is a parallel group at I300I, and the I300I group then takes the topic to their counterpart in Japan, the J300. In parallel, if there are SEMI standards to be voted on, a few individuals at Intel will begin an exhaustive campaign to ensure high SEMI voter turnout. This is important because a low turnout invalidates the vote, slowing down the

⁷ www.semiatech.org/public/division/300/metrics.htm

⁸ Early on, manufacturers and suppliers alike had expressed the view that the SEMI process was too slow to be effective for defining standards at the beginning of the transition, rather than after the fact as had happened before. To date, the SEMI process has worked well enough, and all parties have accepted the fact that I300I requirements may precede formal SEMI acceptance, but that the goal remains to obtain SEMI standards status.

⁹ www.semiatech.org/public/division/300/guide.htm

setting of standards. Everything is in order and successfully exercised for the standards defining AMHS interfaces, lot carriers, mini-environments, some assembly issues, and others. More are still being worked on.

New and Improved Business Processes

Within Intel, a number of business processes have been improved or invented to ensure that the overall goals for the 300mm transition are met. A few deserve special mention, although most details are beyond the scope of this paper.

Intel has a time-tested equipment-selection methodology that is managed by the strategic committees mentioned earlier. Because at 300mm there were a number of new requirements that do not pertain to 200mm equipment, we developed a new set of selection-training materials and delivered them to key people from every selection team. The new selection methodology and requirements augment the existing system; they do not replace it.

Intel uses a purchase spec to define the requirements of our equipment and to manage supplier performance to those goals. The generic form of this document was updated to be consistent with the requirements that eventually became the Global Joint Guidance. Some other changes were included too, where those changes would drive increased productivity of the toolset. For instance, we changed the specification for individual tool scrap rates to be consistent with results already achieved within Intel, thereby increasing the overall line yield targets by approximately 2%. The achieved results had not yet been, and likely never would be, incorporated into the 200mm purchase spec. The writing of the new spec served as yet another scrubbing of our requirements, which ensured that we asked for industry standards, not some unique ones favored by a special interest group within Intel.

Several sources of data are used to set targets for tool performance, to provide real time measurements of tool performance, and to order tools based on some combination of the above. The sources tend to be in conflict, sully the decision making process. For 300mm, we put in place a single database called the Selection Database (SDB), which serves as the sole source for decision-making data. The stakeholders have defined and agreed on a process for updating the database and for using it for various applications. The system has some shortcomings, e.g., it is not particularly user friendly, and we find ourselves needing to remind people that this is the sole source. But, it does work, and it enables engineers and managers to quickly answer questions and develop summaries that would otherwise be obtainable only by extensive scurrying.

Each selection team has the responsibility to develop cost models to manage suppliers to productivity expectations and AMHS interface requirements, to measure utilities and natural resource consumption, and to do several other tasks. While we had trained the teams on the requirements, we needed to oversee them on their performance to a degree never done before. If we failed to do that, we could be pretty sure that the full benefits of 300mm would not be realized. We put in place the Selection Synergy Working Group, chartered by the CFWG, whose initial function was to provide a single forum for answering questions about requirements and 300mm business processes. Somewhat later its function changed to ensuring that the teams were functioning properly, that they had membership from all the right groups at Intel, and that they had schedules for testing. Finally, today, the Synergy WG's role is to monitor progress on key deliverables of the selection team, to delve as deeply as possible into technical issues related to meeting 300mm targets, to challenge the teams to exceed the targets, and to provide clear summaries to management of performance across the toolset. For instance, at 200mm, it would be an exhausting effort to put together data on a single topic, such as, compliance to emissions' standards, across all the tools being selected within a given timeframe, even though in 200mm only a few tools are being selected at once. Through the action of the Synergy WG, we get summaries across the entire toolset on each of approximately ten different topics. In principle, the functions of the Synergy WG could be completed by the committees managing selection, but they have asked the CFWG to manage these details for them, while retaining responsibility for making the selection itself.

For 200mm selections, each team is individually responsible for securing the test wafers it needs to properly exercise tools. The test wafer flows are typically run in our development fabs. Since Intel has no 300mm tools, running test wafer flows is impossible so we replaced the distributed model of test wafer acquisition with a centralized model called the Silicon Clearing House (SiCH). The SiCH had three primary functions: (1) to determine the flows needed for all selections, (2) to secure processing of these flows outside Intel, and (3) to manage the logistics of pushing wafers from shipping and receiving to sites where processing could occur (primarily Austin, Santa Clara, and Japan), and then redistributing the wafers to the teams as needed. In this fashion, approximately 5000 wafer passes have been processed to date (Figure 6), with another 5000 anticipated prior to completion of all the selections. This is the most complex and underappreciated task of the entire 300mm program.

300mm Si CLEARING HOUSE WAFER MOVES




Figure 6: 300mm silicon clearing house wafer moves. Figure courtesy of Melton Bost.

The strategic committees managing selections are limited to examining the most expensive tools. The number of tools they manage is counted in tens, whereas the number of line items in the full catalog of everything needed for a factory is counted in thousands. Intel's change control policy, called Copy **EXACTLY!**, demands that every item be the same in all factories running the same process technology. The business processes for ensuring compliance on the less expensive tools were *ad hoc* at best; many decisions were not reviewed and were left up to individuals. For 300mm, the CFWG was asked to manage the selection of the thousands of tools falling in the cracks between the strategic committees. Many of these items are not wafer-size dependent, so they default to today's choices. Many others fail to meet criteria related to process sensitivity or total cost, so teams are asked to follow a formal selection process, but the results are not reviewed. The remainder, again numbering in the tens, will be reviewed by the CFWG as we near selection.

Getting to GO!

While the CFWG and the various other committees staffed by middle management served to set most direction and manage most details, for Intel to make the multi-billion dollar decision to go to 300mm, a higher level committee was needed.

In July 1995, the first 300mm Steering Committee was formed, with a charter to make a recommendation for action based on an analysis of the full set of relevant issues. Comprised of vice presidents and senior managers representing engineering, materials, development, strategic planning, manufacturing, and finance, we made our first recommendation in early 1996: Intel should proceed towards 300mm, intercepting it at our 0.18 μ m generation in Santa Clara, to be followed quickly by the 0.13 μ m generation in Portland.

A second similarly stacked committee was formed to steer the project into existence. A date for first deliveries was set, October 1, 1998, and planning commenced. The new steering committee defined a process of reviewing overall status (or risk), and the First Risk Assessment was done in October 1997. We determined that Intel's need for additional capacity and the industry's ability to deliver it at 300mm no longer coincided. The committee pushed out our intercept to 0.13 μ m, with first tools scheduled for an April 1, 1999 delivery. At the time, it appeared that more than six companies would put together pilot lines in 1998 and 1999.

Because the intercept no longer called for development at two generations and two sites, the committee task was greatly simplified, so it made itself once again. The Second Risk

Assessment was done in May 1998 and showed that the international slowdown signaled by the Asian financial crisis of 1997 had decimated the industry's plans for 300mm in 1998 and 1999. Only Siemens' pilot line, with partnership from Motorola and money from the German government, was likely to get off the ground in that period. Nevertheless, overall tool readiness appeared somewhat better, but real knowledge of it was too thin to make a startup decision. Furthermore, the schedule for the 0.13um intercept did not really demand first tools in April 1999; we could slip the first dock dates without compromising the process certification dates. Combining this with the slowdown in business and the reexamination of all major capital programs, the recommendations of the Second Risk Assessment group were (1) that we make no commitment to a startup date, (2) that we take the next six months to make detailed inquiries into tool readiness, and (3) that we do a Third Risk Assessment towards the end of 1998 to determine what date, if any, would work.

Fourth Quarter of 1998 Outlook

The overall industry slowdown has reached historic proportions, and the lack of a startup date has made management of the suppliers more difficult than usual, but not impossible. We have weathered the dog days of summer 1998 when layoffs in the equipment industry became routine, when the world's largest supplier withdrew from I300I undermining I300I's importance and seeming to threaten it with extinction, when at SEMICON West senior managers from various sides met but seemed near fisticuffs, when banner headlines in the San Francisco Bay area papers announced the death of 300mm, when every pundit, no matter how small, offered their opinion on the demerits of the industry transition, and when our own selection teams declined to fulfill program commitments because they felt management was itself not committed to continuing the program. Ignoring all the negativity, we are focussing on the critical task of getting the data on tool production worthiness, by working extensively with I300I, with suppliers, and with key groups at Intel.

As we enter the last quarter of 1998, we are vigorously exercising a process to review the teams' equipment data tool-by-tool, anticipating the Third Risk Assessment before year's end. As of this writing, we expect that approximately 80% of the required tools will be ready for selection and a January 2000 startup. As for the remaining 20% of the tools, there are no known showstoppers; they will just need to be managed carefully.

Overcapacity is driving manufacturers to extend the useful life of 200mm equipment where ever possible. The 200mm era, in the sense of new construction of 200mm lines by the major manufacturers, is effectively over. The next big opportunity for the suppliers is 300mm. The forecasted productivity of the 300mm equipment looks very good, appearing to come close to capital productivity targets, as shown with historical perspective in Figure 7. If the forecasts are borne out in practice, then 300mm will seem a much more attractive alternative to using 200mm fabs, even those where the equipment set is largely depreciated. When the realization of this counterintuitive result sinks in, as it has here, the industry will make the transition much more rapidly than it did when it went to 200mm¹⁰. The first year in which this will have a big impact on suppliers' revenue is 2000, and the first year it will have an impact on manufacturers' cost is 2002, possibly as late as 2003. This impact, as measured by die cost reduction, will be twice as large as that of any wafer size conversion in our short history. The two lessons we should take away from this are that such transitions should be planned as an industry and that wafer size should increase by at least 50%.

¹⁰ The transition to 200mm took place over many years. IBM began production in 1988, but the transition was still gaining momentum in 1993 when Intel began 200mm production.

Forecasting Productivity Improvement for new 300mm vs 200mm Fabs




Figure 7: Forecasting productivity improvements for new 300mm vs. 200mm fabs

Acknowledgments

First, I would like to thank the early members of the Cross Functional Working Group, particularly those who articulated the automated materials handling roadmaps and gave us all a model to follow. These individuals include Sandra Viarengo, Dev Pillai, Srinu Srinivasan, Bala Subramaniam, Jeff Pettinato, Fred Voltmer, Richard Parker, Noel Acker, and Yolanda Leung. Next, I would like to thank my other colleagues who have over time helped build the infrastructure that is now in place and working today. These include Winston Saunders, Carolin Seward, Dave Krick, Melton Bost, Larry Norman, Tom Garrett, Ken Moyle, Gordon McMillan, Rich Poliak, William Yeh, Roy deGroot, K.V. Ravi, Diana Harris, Brian Hunter, Stefan Radloff, Stephen Sumner, John Morrissey, Derek Youngson, Dan Enloe, Tom Abell, John Souza, and Kwang-Woo Choi. Also, Ann Marie Kenitzer and Peter Silverman have supported the supplier management and tool selection effort, providing staffing and invaluable critique throughout. Frank Robertson, the General Manager of I300I deserves special mention. There are many others, and I apologize to those I have left out unintentionally. Finally, I would like to acknowledge the contributions of Intel senior managers who provided inspiration, time, support, guidance, and patience over what has now become a long development effort. These include Gulsher Grewal, Don Turner, Don Rose, Paolo Gargini, Ken Thompson, Bob Jecmen, and Youssef El-Mansy.

Author's Biography

Dan Seligson spent his first ten years at Intel as a researcher in Bubble Memories, X-Ray Lithography, and Neural Networks. During this period, he did a two-year assignment at SEMATECH in Austin and was a researcher in residence at the Hebrew University in Jerusalem for six months in 1990. Subsequently, in 1994 he began managing equipment suppliers in the Thermal Processing Area and building the 300mm program. Today he is the Front End Synchronization and 300mm Manager in TME. Prior to joining Intel, he received a B.Sc. from MIT in 1977 and a Ph.D. from Berkeley in 1983, both in physics. He holds five patents and has a long publication list. Outside family and work, his interests include windsurfing, music, and photography. His e-mail is daniel.seligson@intel.com.