Electronic Design Automation Industry Council

EDA Roadmap Taskforce Report

Design of Microprocessors

Date of issue: 3/28/99

Published by: Si2, Inc.

Table of Contents

LIST OF FIGURES

Task Force Recommendations

FOREWORD

This is the Report of the EDA Roadmap Taskforce on the Design of Microprocessors. This Roadmap is first in a series that attempts to project the future of technology for the design of electronic systems, particularly semiconductor integrated circuits. Starting with the design practices and tools in general use in 1998, it projects the changes required to meet the needs of the design community five years in the future, i.e. in 2003.

It is built on the National Technology Roadmap for Semiconductors 97 (NTRS97)[1]. We strongly recommend that readers of this report be familiar with NTRS97, it can be read or downloaded from the www.sematec.org[4] web site.

NTRS97 is the latest in a series that projects the evolution of semiconductor processing technology for five years into the future from date of issue. These Roadmaps have been remarkably accurate even while facing the formidable challenge of making projections in a technological field noted for its astounding rate of development. The Roadmaps are aggressive, making projections which at the time of issue seem futuristic — yet, in practice their forecasts are more often exceeded than not.

The NTRS97 authors approach their forecasting challenge in a unique manner, which futurists in other disciplines could well emulate. They set a stretch but attainable goal of an overall end result —an electronic system on a semiconductor chip which has defined capabilities well beyond the state of the art at the time. They then specify the evolution of each technology that must be attained in order to reach that goal. Some developments may take more effort than forecast, while others less. It is assumed that the worldwide semiconductor industrial, vendor, and university community will shift resources as needed to bring the lagging technologies through on time. Hence, the study is not a crystal-ball-gazing exercise, but a rigorous, schedulable plan of action to attain the goal.

NTRS97 identified *Design and Test* as a key technology whose evolution must be accelerated if the overall goals are to be reached. Readers are particularly directed to that chapter of the report.

This EDA Roadmap takes up the challenge of NTRS97 by digging deeper into Design and Test issues attempting to identify areas of critical R&D need or major design and test paradigm shifts.

This Roadmap is sponsored by the EDA (Electronic Design Automation) Industry Council[2]. The Industry Council is a committee that comprises EDA suppliers including EDAC[3], EDA users from both system companies, and semiconductor companies, as well as industry organizations including SEMATECH[4], SRC[5], and Si2[6]. This Taskforce was funded by DARPA[7] and Si2 administrated the work.

The EDA Industry Council chartered the EDA Roadmap Taskforce to select one electronic marketplace and examine changes that are required to design effectively five years hence. Implied by the Industry Council is the worry that the industry is unprepared for the future. Their hope is that an early prediction of needs and capabilities will allow enough lead-time for the development of the required advancements.

The EDA Roadmap Taskforce first met in February 1998. The Taskforce is comprised of 30 experts with semiconductor processing, design, or EDA backgrounds. They included representatives from both commercial vendors and customers of EDA. Many of the participants have a history in more than one of these fields. They examined several target markets to study (telecom, PC, consumer electronics, automotive, etc.) and chose High-end Microprocessor design as the most challenging. We expected that other taskforces would concentrate on such markets as telecom, PCs, consumer electronics, etc. The Taskforce then targeted designs that will begin in the year 2003, the time frame when semiconductor processes are forecasted to allow feature sizes smaller than one hundred nanometers.

The EDA Roadmap Taskforce met six times during 1998. It studied the NTRS97 report in detail. It obtained EDA information from many sources, including extrapolations from public plans that

are published by commercial EDA companies. Much of the design information that the Taskforce obtained is anecdotal and it is not publishable. All anecdotal data was verified by determining if similar sources gave equivalent responses. In fact, the “stories” are surprisingly similar.

Investigation began at points where design paradigm shifts will or must occur. This included forecasting directions of semiconductor processing, design approaches, and electronic design automation.

Next, the Taskforce examined alternative solutions for identified paradigm shifts. The solutions included modification of semiconductor processing and design methodology, and R&D on new EDA tools. The report to follow contains the results of this work.

ACKNOWLEDGEMENTS

The Taskforce was comprised of highly experienced technologists. They strove to advance our science. Each member's affiliation indicates the member's association at the time that the member joined the committee.

Paul Weil, Co-Chair, Intel
Arny Goldfein, Co-Chair, Cadence
Shakhar Borkar, Intel
Pat Capapano, Motorola
Jason Cong, UCLA
Don Cottrell, Si2
Carlos Dangelo, LSILogic
Steve Grout, SEMATECH
Bill Joyner, SRC
Andrew B. Kahng, UCLA
Sani Nassif, IBM
A. Richard Newton, UCB
Bryan Preas, Xerox
Steve Sapiro, Consultant
Chuck Shaw, Cadence
Peter Suaris, Mentor
Rob Smith, Actel

The Taskforce included the assistance from many other individuals. We would like to recognize the following for their special contributions to this effort. Each individual's affiliation is shown at the time of joining the study. We apologize for any names that we might have omitted in error.

Andy Charnas, Sun
John Cohn, IBM
Mark Birnbaum, Fujitsu
Phil Fisher, SEMATECH
Mark Flomenhoft, Intel
Randy Harr, Synopsys
Bernd Koenemann, LogicVision
Bob Lashley, Sun
Farid Najm, Univ of Ill
Phil Nigh, IBM
Hillel Ofek, Dynalogic
Saila Ponnapalli, IBM
Larry Rosenberg, VSI
Naresh Sehgal, Intel
Dick Smith, Cadence
Ron Walther, IBM

SECTION 1

The Challenge to EDA Design and Test

The ability to manufacture a semiconductor chip with certain capabilities is the driving force in electronic system technology. The evolution of that capability is the subject of the SIA National Technology Roadmap for Semiconductors '97 (NTRS97)[1] and it forecasts semiconductor fabrication capabilities in 2003 and beyond.

Integrated Circuit chips cannot be manufactured unless they can first be designed. NTRS97 identified Design and Test (D&T) as a critical technology whose evolution risked not being able to meet the challenges of designing the chips that could be fabricated in 2003. This EDA Roadmap Task took up the challenge of defining how D&T must evolve to prevent it from becoming the bottleneck, thus limiting the growth of the entire electronics industry.

The Taskforce is concerned that new EDA tools may not be ready for these new designs. Chip designs will start before the end of 2001. A gap may initially exist between EDA capabilities and the chip designers' needs. Among the goals of the Taskforce, is to focus EDA R&D efforts on the most critical problem areas.

It was concluded that without additional research in a number of areas, the required changes would probably never be made, compounding the problem. We hope that this will be thoroughly studied and debated throughout the industry, and that resources will be allocated to accelerate necessary activities to bring it to fruition.

To scope the problem, the Taskforce studied needs for high-end microprocessor design starts that will begin in 2003 — the time frame when semiconductor processes are forecasted to allow feature sizes smaller than one hundred nanometers.

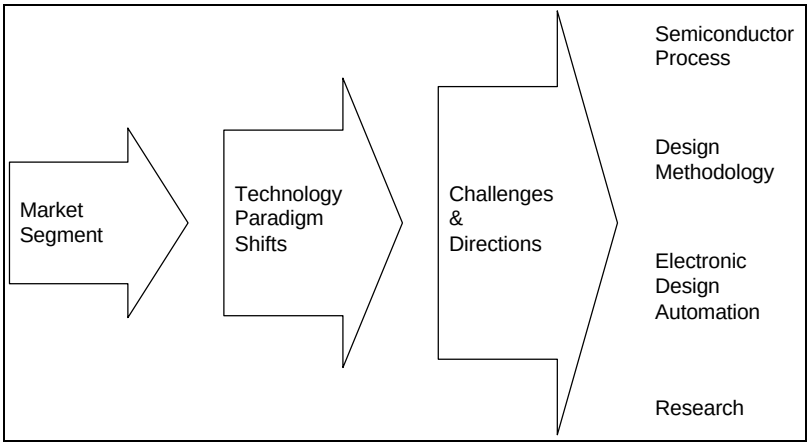

Figure 1: Taskforce Process - Focus on Change

To identify where paradigm shifts will occur, the Taskforce forecasting directions of the semiconductor processing, design approaches, and electronic design automation using the critical design points within these microprocessors.

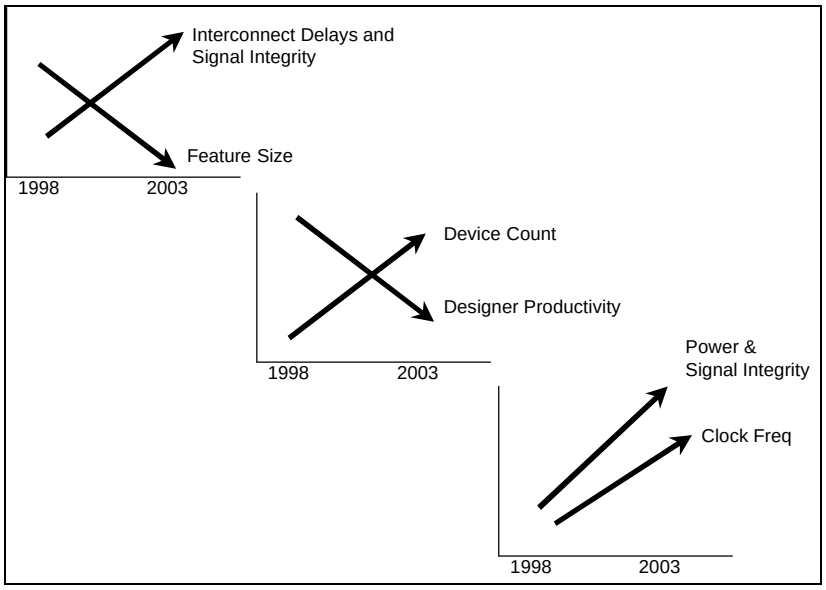

Figure 2: Technology Trends

The semiconductor process information is extracted from the NTRS97. The forecasts presumed that each technology will be enhanced in intervening years. The Taskforce built on these enhancements in order to identify design paradigm shifts and to emphasize other aspects where technology must be reshaped or created.

Next, alternative solutions for the paradigm shifts were examined. The solutions included modification of semiconductor processing, changing the design methodology, and creation of new EDA tools. Some of the technology changes impact design positively.

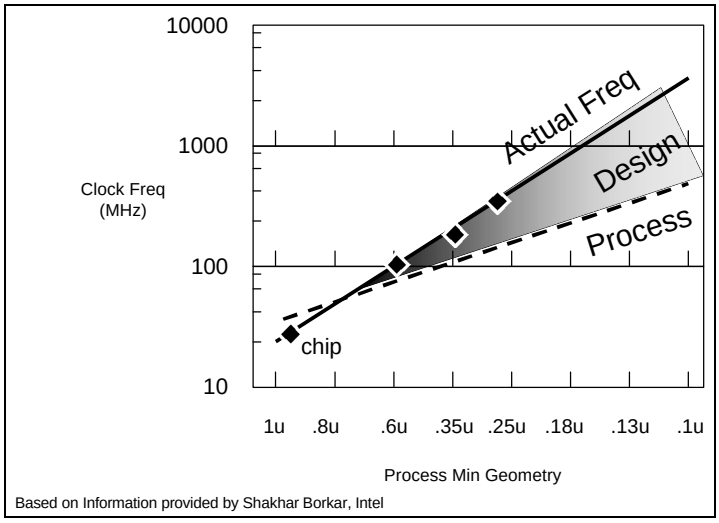

Figure 3: Pushing Frequency through Process and Design Changes

For example, small feature geometries make high-speed circuits possible as a result of the inverse relationship of gate length to RC time constants. But often there are second-order negative effects. For example, increasing the number of transistors operating at high speed, even at reduced supply voltages, will consume greater amounts of power.

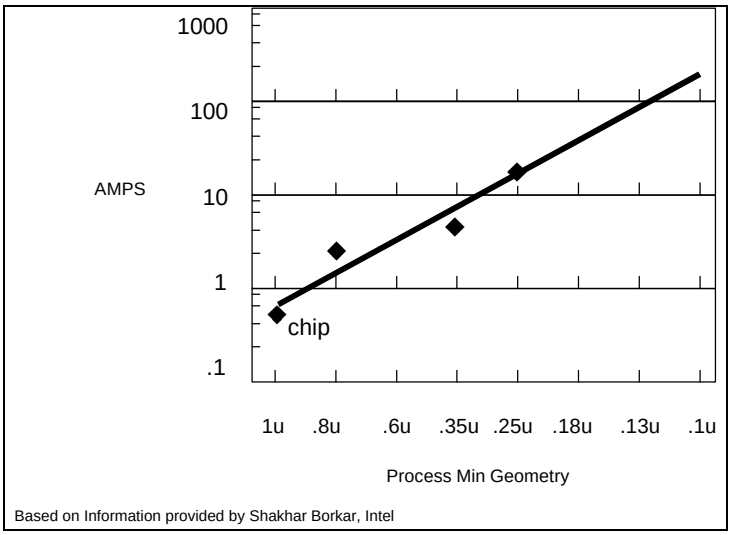

Figure 4: Power Rises Sharply

Target Design

Year 2003 - Microprocessor Data Sheet

The Taskforce felt that it was important to identify characteristics of the typical high-end microprocessor design in the 2003 timeframe to allow evaluation of difficulties and study of possible solutions. These microprocessors will be used in server computers to perform mathematical and data operations, and to perform multimedia and speech analysis / synthesis functions. Chip information, from the NTRS97, for microprocessors that will enter manufacturing in the 2003 timeframe was used as the basis of the evaluation. These products are expected to reach production in the 2004 timeframe. The Taskforce expects that these chips will be comprised of 200 million transistors of which at least 50 million transistors will constitute logic gates. The chip will be fabricated in a CMOS semiconductor process with 100 nanometer or smaller minimum geometry design rules. The overall die size will exceed 500 square millimeters. The chips will be connected into packages using four thousand "bumps", where over half the bumps will be used for power and ground. The chips will run at high speed with basic clock frequencies above 3.5 GHz, and with very fast slew rate requiring internal frequencies above 100 GHz.

▪ **Design Size:**	
- Total Transistors	200 X 10^6
- Total Logic Transistors	**50 X 10^6**
- Wiring levels	8
▪ **Scaling:**	
- Target Process for Microprocessors Ship)	**100 nm** (2003 Starts for 2005
- Chip Size	520 mm^2
▪ **Frequency:**	
- Local Clock Freq.	3.5 GHz
-	3rd Harmonic = **9 GHz**
-	Slew rate = **150 Ghz**
▪ **Chip Statistics**	
- Chip I/Os:	4000
- Wiring levels	8
- **Total Interconnect length**	**2840 m/chip**

Figure 5: Target Chip Data Sheet

Designing 0.2 billion transistors which will be switching at microwave frequencies will be a severe challenge. It will unavoidably cause several fundamental challenges as a result of the design magnitude, feature size, and clock frequencies.

Design Magnitude: Designing a 0.2 billion transistor chip, while presuming that a typical designer with present automation can design 1000 transistors per day, would take approximately 500 person-years. Even assuming most of the transistors are for memory and the existence of advanced techniques for design reuse, this will be a formidable task. The Taskforce identified designer productivity as a very important area to investigate.

Feature size: With 100 nm being sub-visible light (1000 Angstroms), process assumptions must change, and many current design assumptions must change as well. Design must be accomplished by exploiting a series of simplifications. Design factors previously considered second-order effects will become dominant effects, and it is likely that key simplifying assumptions will no longer be valid. The Taskforce identified the consideration of small geometry effects as one of the changes that needs to be investigated.

High frequency: The microprocessor clock frequency of 3.5 GHz is in the microwave range. If typical signals have sharp edges then at least three odd harmonics will be required to maintain the waveform shape. From a Fourier transform standpoint, the third harmonic is 10.5 GHz, and the slew rates are faster. Design teams must be prepared to deal with harmonic frequencies as high as 100 GHz. Further, at these frequencies, any interconnection approximately one millimeter long will need to be designed through non-traditional (transmission line) means. The Taskforce identified high frequency interconnect as an important area to investigate.

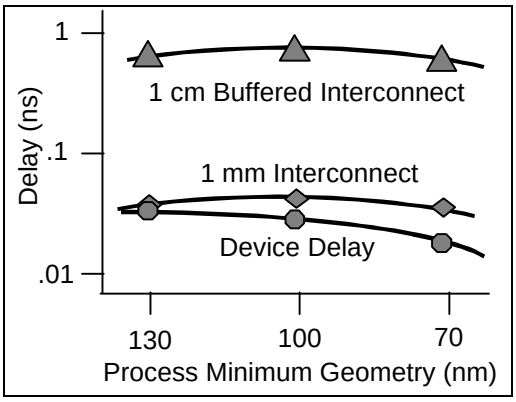

Figure 6: Global and Local Interconnect Delays vs. Gate Delays

The Taskforce objective is to define approaches to design larger chips with fewer engineers, while successfully dealing with new challenges introduced by the sub-100 nm technology. This report identifies enhancements and modifications to semiconductor processing, design methodology, and electronic design automation necessary to reach this objective.

The next sections of this report study these challenges in detail. It then spells out its problem areas and makes conclusive recommendations to deal with these problems.

SECTION 2

Power

In the Microprocessor market sector, speed is the ultimate differentiator. Therefore, the designer's highest priority is the maximization of clock frequency for the utmost processor speed. To optimize speed, other constraints need to be relaxed. One design constraint that can be adjusted is power. Microprocessor designers forecasted that, in the future, microprocessors would consume hundreds of watts per chip. Even though there are a number of mechanisms of removing the heat produced from this excessive amount of power, high currents, and thermal gradients are major concerns.

Power dissipation presents a number of design challenges. Inactivating portions of a chip can reduce power, but power-up can cause uncontrolled transients. Few tools are available today to guide designers in analyzing their design for worst-case power. The following discussion emphasizes the electrical effects of power dissipation. However, physical effects also must be predicted and analyzed. A very localized high transient power dissipation can create thermal gradients which initiate cracks in passivation layers and rupture solder bonds. There is a need to develop tools for these analyses.

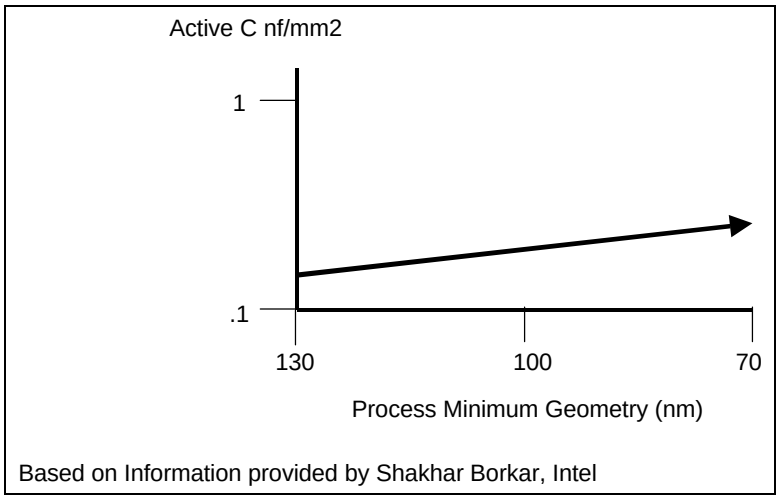

Figure 7: Growth in Active Capacitance Leading to Power Growth

A simplified approach to examine power (P) is to consider it as a function of capacitance (C), frequency (f), and the power supply voltage (V_{dd}). The total capacitance is composed of active and interconnect components that are being charged/discharged between the logic "zero" and "one" levels. As a technology scales down, more capacitance per unit area results, since the gate oxide thickness is decreased and the horizontal wire spacing is reduced (causing increased mutual capacitance between interconnects).

With smaller channel lengths however, transistor drive increases and results in a net increase in the clock frequency (f). Combined with a general trend to larger chips (and therefore yet more capacitance) these two phenomena overcome the gradual reduction in V_{dd} and result in a net increase in power dissipation as technology is scaled.

A rule of thumb for estimating power consumed by a digital circuit is:

$$P = K\ C\ f\ (V_{dd})^2$$

where K is a proportionality constant which accounts for nodes not switching at every clock cycle. To achieve ever-higher performance, circuit families (e.g. dynamic/domino CMOS) are being used that have a higher K constant than classic static CMOS.

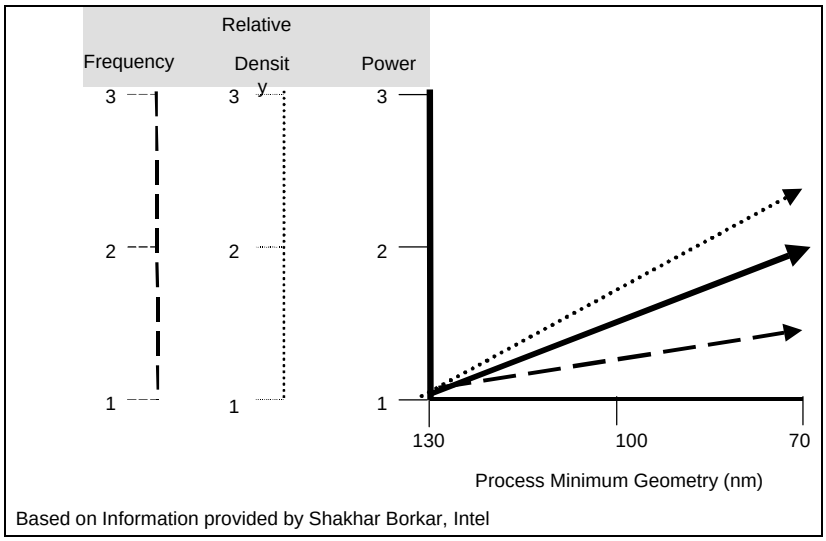

Figure 8: Supply Current

Microprocessors dissipating 100 to 150 watts are currently being designed, and power dissipation of 300 watts are anticipated (predictions of chips requiring 1000 watts were discussed.) This represents a dramatic increase in power from today. Further, since the supply voltage is expected to decline to about 1 volt, the current will increase faster than the power. This means that average currents of greater than 300 amps will be produced, and transient currents will double that value for durations up to 1 nanosecond.

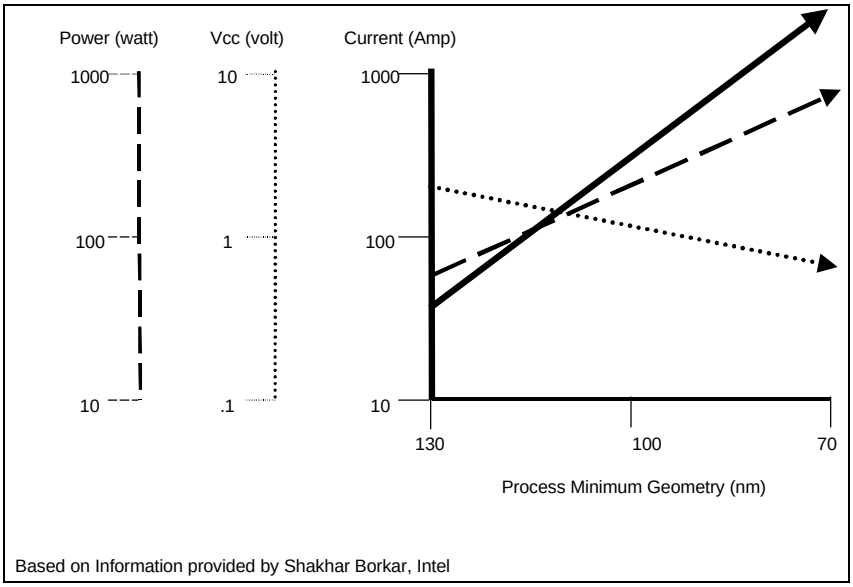

Figure 9: Power / Voltage = Current

A major challenge will be to maintain a near constant supply voltage when delivering (on the average) hundreds of amps, with a variation in supply current of up to hundreds of amps.

With lower supply voltages, the threshold voltage (V_T) of MOSFETs is scaled down in order to maintain performance. Most often this is accomplished by maintaining V_T as a fixed proportion of V_{dd}. Because the sub-threshold current is an exponential function of V_T, this reduction results in a very rapid increase in leakage currents. This leakage current acts to increase the background power dissipation and noise, and is very important for low power design.

With higher power supply currents and increased frequencies, and the use of low-power design techniques that involve clock gating and/or selective shutdown, the current change (di/dt) term increases, while the resistance in the power network is reduced in order to keep rail droop under control. This effectively lowers the threshold at which the dynamic voltage drop (L di/dt) is comparable to the static drop (R). This means that rail inductance becomes more important for power supply and noise performance. Even small mutual inductances between the power grid and signal interconnects will disrupt signals because of the large di/dt.

Since power is not expended uniformly throughout the chip, different areas of the chip will experience different temperatures, resulting in thermal gradients across the chip. These gradients may cause timing variability (hot circuits are slower), which in turn can require lower overall design performance in order to compensate for increased timing variability.

There is significant interaction between the power design, clock design, and other global controls. For most microprocessors, the clock network is the largest single consumer of power. The clock runs at full speed over the longest wires thus requiring appropriate ground return path design in order to keep inductance in check.

Due to otherwise excessive power consumption, large chips will turn off portions of their circuitry. However, simultaneously awakening large portions of a chip may cause power surges. Thus, power surges can be expected in common situations such as power up and reset. These are all dramatic effects that will require modifications in processing, design methodology, and in design automation.

For both power and ground distribution and for clock/return distribution, the chip package will play a larger role. The Taskforce foresees increased use of flip-chip technology with thousands of small bumps connecting chip to package with many of them devoted to power and clock distribution with the equipotential planes moving off the chip and into the multi-layer package.

For what remains on the chip, the need to keep electrical noise down demands that the clock network be appropriately shielded to isolate this circuitry. The Taskforce recommends that additional layers of metal be used for shielding as well as power and ground planes. In order to reduce power supply droop, it will be necessary to employ on-chip-decoupling capacitance.

Early in the design cycle, tools are needed to estimate power dissipation accurately for each major section of these very large microprocessors. These early predictions need to estimate power dissipation, thermal behavior, and current requirements for different areas of a chip. Power planning tools must be fully aware of the current distribution and power dissipation of underlying technology. That entails knowledge of wire resistance, capacitance, cross capacitance, mutual-inductance and self-inductance parameters, and layout physical design rules. The prediction must anticipate layout issues such as blockages i.e. portions of the chip where the power grid may be only partial.

In order to reduce the power surges that occur at clock edges, it may be important to adopt self-timed and asynchronous design techniques for major portions of the chip designs. Some designs, which are entirely asynchronous, are already reaching the market (although not at the high end). Current design tools do not adequately synthesize or analyze self-timed and asynchronous designs. The Taskforce suggests that a new category of asynchronous and self-timed design automation tools may need to be invented.

Recommendation I: Power

- **Semiconductor Process Changes Required for Increased Power**
  - Additional Metal Layers for Power Planes
  - Additional Metal Layers for Shielding
  - On Chip Decoupling Capacitors
- **Power Management Design Methodology**
  - Increase Usage of Gated Clocks
  - Staggered Clock
  - Self Timed and Asynchronous Design
- **Design Automation Required for Increased Power**
  - Early Prediction of Power
  - Self-Inductive and Mutual-Inductive Effects to Signal Line Avoidance Software.
  - Power Dependent Timing Verification

SECTION 3

Signal Integrity and Delay Variation

Microprocessors are designed for high-speed operation. The Taskforce investigated factors that limit speeds (GHz) in microprocessors that will be fabricated in CMOS processes with 100 nm minimum geometries. Much of the data that has been used is from the NTRS97 but the data has been enhanced from other sources.

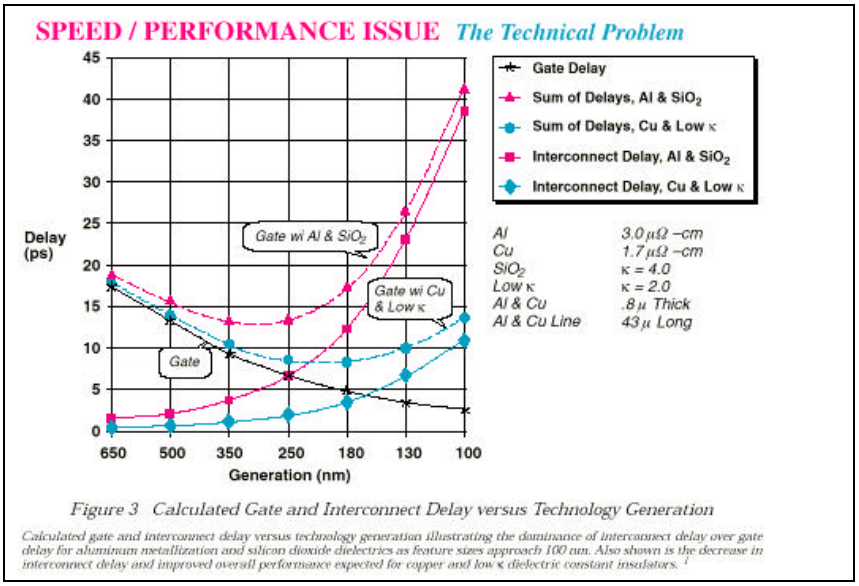

Figure 10: Gate / Interconnect Delay

The gate delay noted in Figure 10 is with no load. The effect of interconnect delay ranges from negligible to dominant. For short distances such as within a cell interconnect delay and signal distortion will be negligible. The Taskforce assumed the use of copper and low-k dielectrics. However, even with the use of copper and low-k materials, the effect of interconnect on delay will dominate over other delay characteristics for major elements on the chip.

Gate Delay	**3 ps**
On Chip Parameter Variability	+/-10%
Average Interconnect Delay	12 ps
Time of Flight	5 ps/mm
Interconnect Resistance	100 ohms/mm
Self Inductance Signal Lines	**0.5 nh/**mm
Mutual Inductance signal to signal	**0.3 nh**/mm
Crosstalk	0.2 pf/mm
Crosstalk Ratio	**.6** $C_{interconnect}$ / C_{total}
Reflections	Non--terminated long routes above 9 Ghz
RF antenna	**2.5 mm**

Figure 11: Signal Integrity and Delay Variation

At GHz frequencies, logic signals will not be clean digital waveforms with slew rates. They will be similar to analog signals that cross switching points. These signals will be distorted through interaction with other signals. This will upset the signal's integrity and alter the time that the signal will take to reach a switching point. This variable time will manifest itself as variable "delay" time.

Physical phenomena that were heretofore ignored will become predominant. Most notable among these will be self-inductance (particularly in the wide global interconnections), mutual capacitance, and mutual inductance. Mutual capacitance will be noticeable in higher impedance circuits with low current nets, and self and mutual inductance will be dominant in high current nets and especially high di/dt nets. Each will cause the signal's integrity to be diminished.

The mutual capacitance and mutual inductance will couple one signal to another. The couplings (or crosstalk) will be proportional to wire dimensions and pitch, and the relative slew rates of the two signals, so that the amount of coupling will vary based on functional operation. One effect of this may be that signals transition before, during, or after the time they need to switch. Pre-charged networks may present incorrect information to the following stage switches. Waveform glitches may impact delay even more significantly if they occur at the switching point where they can cause an actual change of state.

Long word-length microprocessor architectures with multi-clocks were examined. These synchronous microprocessors switch many nets simultaneously. The advantage of long word architectures is that many parallel operations can be accomplished concurrently. However, this leads to multiple signal integrity problems. In many cases, combinations of mutual capacitors and inductors will affect many signals.

A simple pair of interconnects is used to illustrate the mutual interference problems. One signal is considered of primary interest. The primary signal's response is affected by another signal, which is an aggressor. If a primary signal and aggressor transition close to the same time, then the timing of both signals will be modified from nominal. For example, if both signals transition at the same time with the same polarity, then there is no change in charge across the mutual capacitance between them. The signal will change faster on both the primary and aggressor nets. If the signals are opposite in polarity, then the mutual capacitance will require double the normal charge and will lengthen the delay on both primary and aggressor nets.

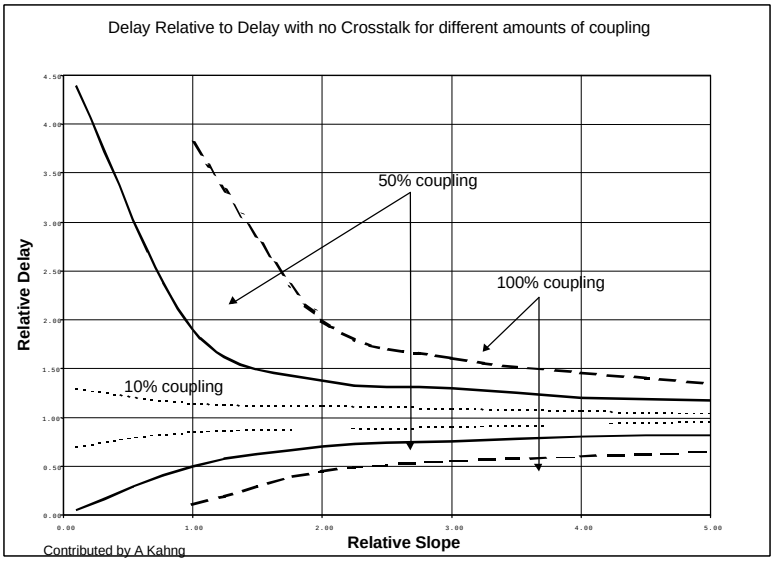

Figure 12: Delay Variation

During multiple signal transitions, some signals will rise faster and some slower. Additionally, the signals will be seriously distorted. Because 100 nm geometry transistors are low gain, some of these distortions will be transmitted to subsequent stages. These stages will have other aggressors that add other distortions. The manifestation of distortion in a digital system is delay variation. Even pre-charging a net will result in delay variation. Each of these distorted signals will reach its final value if given enough time. However, in high-speed designs, precautions will need to be taken to assure that the variations will not cause a system malfunction.

Interconnections of approximately 1000 nm (10x minimum geometry) or less may remain unaffected by other signals. For interconnections that are approximately 100u in length (1000x minimum geometry),

the effects of other signals will limit performance and precautions will be needed. Long interconnects at the chip level, like buses, clocks, reset, and power distribution, must be designed to control the effects of mutual signal distortion. Tools may be needed to support several classes of interconnection resources, depending on length and delay/noise criticality. In an interconnect-centric world, the available types of interconnections between blocks become a critical class of design options.

At some point, buffers will be added to reach a practical maximum delay per mm. However, when using buffers/inverters, time of flight increases. Inverter insertion will be most effective if inserted in a physically staggered pattern. The staggering will have the effect of balancing the aggressor signals by inverting the aggressor's polarity. This will limit the effect in some challenging situations, while it does not help in other situations like power distribution and low skew clocks. The primary benefits of adding buffering are the reshaping of the signal and reducing the effects of cross capacitance. The use of Schmidt Triggers or similar circuits with hysteresis may be more effective in minimizing the impact of delay variations.

Designers will also need to consider the use of shielding to limit the effects of long parallel interconnections or high current. Running signals over a functional block will often induce undesirable effects on the block's timing. Shielding will be needed to greatly reduce signal interference, but will need to be carefully designed in order not to introduce undesired side effects.

Delay uncertainty will be a function of the direct and indirect relationships between signals. If multiple signals are always spaced temporally, then their mutual effect will be diminished. If these signals often switch at the same time and at random polarities then many will impair the timing of the others.

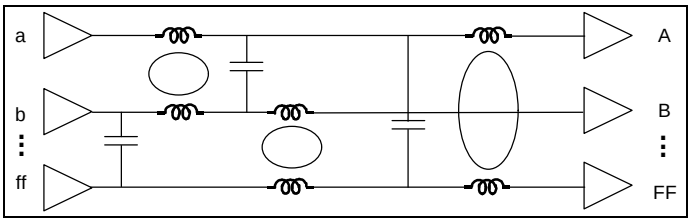

Figure 13: New Signal Interrelationships

Microprocessors are synchronous and many signals transition concurrently. To avert problems, designers and / or tools need to look for long runs of parallel wires and analyze them for coupling problems.

There are several essential EDA technologies. One such technology, parasitic extraction, derives the mutual inductance and capacitance as well as other parameters that affect delay. Because signal integrity analysis must capture relationships between the signals, delay models will be multi-port and statistical. Additionally, parameter and temperature variations across a chip will also contribute significantly to changes in delays. Therefore, additional ports will be needed for these variables. These models will be evaluated in a multi-path logic analysis engine.

The Taskforce recommends that the multi-path logic analysis engine be built into physical design. This multi-path analyzer will need to carry the relationships between signals so that mutual relations will be properly modeled. Physical design must be "signal integrity" aware. Tools must support transparent maintenance of timing models and avoid timing problems whenever possible by using buffering, shielding, variable and asymmetric width / spacing of interconnections, and interlacing quiet signal lines next to active signal lines. The higher frequencies will require the equivalent of twisted pairs, coax cable, transmission lines, and even wave-guides. Sensitive circuits, e.g. those which cannot tolerate the added capacitance and/or inductance of parallel shielding lines, may need to employ driven shields using techniques from analog circuit design. EDA tool enhancements will be needed to include such circuitry.

The Taskforce recommends that physical design be "interconnect-centric". The Taskforce further recommends that physical design be completed in hierarchical stages and that longer, more challenging connections be analyzed first, and then either completed or redesigned before dealing with less challenging details. The higher levels of physical hierarchy with power delivery, global signals such as clocks, and busses and signals between blocks will be designed first. During the design of the chip level, the lower level functions will be specified and budgets will be set for functional delays. Functional design

at the high level will control signal integrity by minimizing the number of signal transitions that implicitly interfere, and by minimizing the number of global signals that are required. Thereby, elements at lower levels of the design hierarchy will be designed aware of chip level signals that will dominate delay. This lower level of physical design will concentrate on achieving the intended delay budgets.

Recommendation II: Signal Integrity and Delay Uncertainty

- **Semiconductor Process Changes**
  - Additional Metal Layers for Shielding
  - Low mutual capacitance and low mutual inductance between signals including power
- **Design Methodology**
  - Hierarchical Design that is Interconnect-centric
  - Staggered Signals
- **Design Automation**
  - Physical Design that is Signal Integrity Aware
  - Multi-Port Delay Models
  - Multi-Path Timing Analyzer
  - Interconnect-centric Design Tools which emphasize High Level Physical Design

SECTION 4

Design Productivity

A major Taskforce objective was to find approaches to designing microprocessors using fewer engineers and with shorter cycle times, while being concerned with more complex electrical/physical details.

Details dominate the effort to design a chip. Many of the simplifications applied to the chip design process today will no longer be valid. The good news is that designers are a resilient and resourceful bunch, and will keep designing high-end microprocessor chips in spite of the roadblocks thrown in their path. However, billions of transistors, switching at microwave frequencies, and complex cross-signal interactions will exceed any engineer's ability to analyze the entire chip design and uncover timing problems. In previous generations, designs were limited in speed by the transistors, and interconnect was usually considered contributing second order effects. In future generations, speed will be limited by interconnect and transistor performance becomes a second order consideration Solving interconnect difficulties will dominate design cycle. The design system and methodology must evolve to support such changes.

Because most of the timing effects will be layout related, simple functional analysis will not be effective for physical design verifications. Current analysis techniques such as logic and circuit simulation have already exceeded their limits. Static analysis that is based on gate or transistor verification will overwhelm the designer with exceeding large numbers of false-path errors.

However, the time it takes to deliver a good chip design depends on many other critical issues such as, a wrong or incomplete specification, or poor design practices. Only by attacking all critical issues, will improved design productivity result.

PRODUCTIVITY FACTORS

Design productivity is a by-product of two interrelated activities. The first is the efficiency of the team producing the design. Microprocessor design teams of today have a wide variation in the team sizes. Some teams are as much as an order of magnitude larger than others. Factoring in design differences do not account for this wide variance. The factor that appears common to the most efficient design teams is that they are populated with "tall-thin" designers who had previously worked together as a team.

The second is the number of design iterations (or spins) required to achieve a correct implementation. The most effective way to shorten design cycle time is to reduce design iterations. One of the major causes of multiple iterations is inconsistent views of the design . Often, one team does not know about the latest changes made by another team or individual. Similarly, because the original specification may not be clear, one team may think that a bus address is 0:32 while the another thinks that it is 32:0. Often a designer will make a change to the design in his personal design responsibility scope and verify it, but forget to update the master design database.

EDA systems must assist both of these interrelated activities and designers' capability must grow so that they are capable of working across hierarchy. EDA needs to optimize design processes and focus design activities and methodology so as to allow the engineers' efforts to be focused on the most critical tasks.

GUIDELINES

A set of guiding principles is offered for creating High-end Microprocessors using 100 nm processes. These guiding principles may seem simple but their impact can be enormous.

- Avoid problems rather than identifying and correcting the errors later

  Verification is the biggest bottleneck in the design process. Tools must be provided which do not require repetitive verification. Development of verification management methodologies, and supporting tools, is required so that automation simplifies the verification process. For example, commonly, a cell is designed and then placed and routed. This is followed by an analysis of the interactions between that cell with the surrounding interconnections and other

circuitry. This analysis may then uncover a resulting design error, which causes the design loop to be reentered. Design teams can prevent problems by defining and following strict design rules. A typical rule may be to include buffering and shielding during the design of a cell to make it resistant to signal interference.

- Verify once and use the results forever

Designers are encouraged to thoroughly evaluate each hierarchy level and encapsulate that design entity as a known-good design. The encapsulation will be used within the next higher design level. Each encapsulation must encompass a broad range of design parameters. Certainly among these parameters are function, timing, signal integrity factors, power, instantaneous current, size, pinout, and design assumptions or constraints.

- Change design focus to be interconnect-centric

The timing of 100 nm designs depends upon the interaction between signals. Characteristics of all but local interconnections dominate over active devices. Physical and electrical design must consider interconnect effects first, then the active devices. What is suggesting is to forecast the interconnection properties then later, precisely define interconnect details. Use the forecasts and a precise layout to refine the acceptable range of each property throughout the hierarchy.

- Tether design changes to the old versions to control the magnitude of any change

The Taskforce suggests that controlling the effect of a change is as important as the change itself. The tools must track changes and their hierarchical effect so that there are no surprises during implementation.

Recommendation III: Guiding Principles

- **Avoid problems**
- **Verify once**
- **Interconnect-centric design**
- **Tether design changes**

In the light of the preceding guiding principles, several strawman design approaches were evaluated. Some were not capable of achieving high-speed 100 nm designs. Others would be burdened with immense details at high levels of hierarchy. The Taskforce feels that the following design approach will meet the challenges of 100 nm microprocessor designs.

MEET-IN-THE-MIDDLE

It is felt that a formalized meet-in-the-middle design approach will be extremely advantageous. Design will be hierarchical. As usual, the design process begins with architectural definition and ends with a physical realization and testing method. The intersection of these domains is the middle. For microprocessor design, the middle is the full chip level. The full chip level is where the lowest level of architecture design and the highest level of the physical design meet. In the middle, the process of resolving ambiguities begins by recognizing risks and clearing away conflicting requirements.

Recommendation IV: Meet in the Middle Design Approach

- **Semiconductor Process Changes**
  - N/A
- **Design Methodology**
  - Actively Avoid Problems
    - *Add Shielding*
    - *Add Buffers*
    - *Interconnect Models*
    - *Verify Block*
  - Full Chip Level First
    - *Full Chip Layout*
    - *Forecast Block Specifications*
- **Design Automation**
  - Verify Chip Using Models

AVOID PROBLEMS

Applying the four principles (see Recommendation III) leads to some conclusions. Avoiding problems implies that bottom-up design cannot be performed without thoroughly specified goals. Goals need to encompass the full spectrum of design parameters from delay to signal integrity, and from power/current requirements to area/pinout. The goals must be based on credible estimates and be as accurate as practical. Timing information, which heretofore was a back-annotated afterthought, needs to be accurately predicted and budgeted. For critical parameters, the degree of uncertainty must be forecasted.

ESTIMATION

Budgets are a means of passing constraints and ranges of ambiguity down hierarchical levels. The budgets must, at each hierarchical level, enable realizable designs. One characteristic of efficient design teams is the ability to estimate/forecast well. These teams have learned to estimate efficiently because they have experienced tall-thin designers who can focus effectively across design levels. Their high level decisions are a natural outcome of low level consideration. It is suggested that estimation become institutionalized and embedded within the EDA infrastructure. Special methodologies and EDA tools need to be developed to estimate the key properties of each block in the design hierarchy as accurately as possible.

Estimation is foreseen as a joining of top-down specification and bottom-up prediction. Prediction is never exact, but the quality of the design in the middle depends on being close. For the full chip level, it must be close enough to complete the level designed. The design will include power distribution, built -in test capabilities, dominant signal (such as clock and reset) distribution, inter-block timing, and soft error control.

Estimating is a conjunction of historical information and tools. Nearly every microprocessor design team has individuals who have designed microprocessors previously. Once captured, the historical experience can be used as a starting point for the next microprocessor design. The same is true for the large blocks (intellectual property). It is expected that automatic generators will create data paths, register stacks, and memory (these generators can also build the models.) It is expected that synthesizers will create random logic. These synthesizers will be advanced in order to produce increasingly accurate models for high-level verification. Designers, at times, may resort to prototype experiments using new blocks or

models. However, model building should be automated with special tools that confirm accuracy of the generated models.

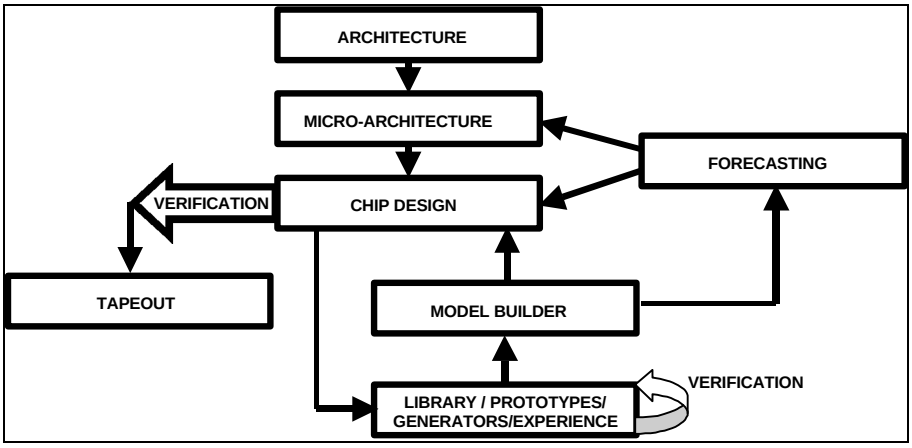

Figure 14: Interconnect-centric Design System

Some design methodology changes will improve the quality of estimation and of the chip design - whereas, software can prevent problems by active avoidance. One example is adding shielding between interconnects that are likely to interfere. This shielding stabilizes the circuit operation and can minimize the effect of an interconnect on surrounding entities. It also will make it possible to predict a block's operation more accurately. Problem avoidance reaps a side benefit of predictability and although such avoidance methodologies may appear at first to require additional silicon area, correct operation forces that result. It is suggested that these schemas be increasingly leveraged.

SIMPLIFY THE DESIGN PROCESS

Some methodology changes can simplify the entire design task. Power and signal integrity are both affected by the number of signal transitions. Minimizing the transition counts and (especially) controlling transitions on longer length interconnects will reduce power and noise problems.

With meet-in-the-middle design techniques physical design starts at the full chip level. As was stated earlier, design needs to be interconnect-centric, since interconnect has the largest effect on design with longer interconnections having the most effect. This includes all of the long signal lines, busses, global clock signals, reset signals, test signals, and power distribution lines. The longest interconnections are at the chip level, then at the block level interconnect, and finally at the local silicon and cell level. To design these interconnections, the high-level blocks of a design must be estimated. The blocks must be placed in the chip structure and interconnects between them will dominate the performance. These blocks are thereby specified in-place where they will be used. Care can be taken to reuse blocks repetitively to save design and verification time.

After global chip level interconnect is completed, each higher level block specification is enhanced with the dominant interconnect effects. These include loading effects, signal integrity effects due to inter-layer crossings or nearby interconnect on the same layer, clock/sequence skew, and power fluctuations. The block's pinout and dimensions are also delineated. This enhanced block specification allows a team to have, at all times, an intact hierarchy that is divisible into sub-activities that can be reassembled. Since each block always coexists with its environment, the design is successively refined and a modification within one block can be reflected to the rest of the design through the chip level hierarchy. Thus, the new design is tethered to the old design and degree of change can be controlled.

Another advantage of meet-in-the-middle is early verification. The estimated blocks are complete models, where lower level details (behavior, delays, signal integrity, power/current, area, and pinout) are abstracted into high-level block representations. Often the high-level blocks will have built-in checkers to assure that their functions are being used appropriately. As a design proceeds from estimation to realization, the model/verification accuracy will improve. As a design matures, the verification is completed hierarchically. As blocks are refined, only the modified blocks will need new parameterization

and verification will be maintained. To make this possible, the details of the each design entity must be captured using physical data (layout) whenever available to increase its accuracy. In this way, engineers can concentrate on the interconnect between the modeled functions (entities).

Controlling change is also important to assure convergence and to limit the amount of (re)verification. Whenever a design modification is proposed, it must be tethered to the preceding design by determining that it meets block and chip level constraints.

The process is one of successive design refinement cycles, converging to a well-formed complete block design. During the refinement process, advanced tool capabilities should transparently (at least until a design error is detected) maintain models defining block properties, and verify that successive refinements are converging upon a well-formed result. Some tools may also need to accomplish trend analysis and be capable of issuing warnings, providing diagnostics, and escalating issues.

MODELING

The lower levels of detail are abstracted into models that are used at different levels of the hierarchy, up to the blocks at the full chip level. This is reverse synthesis. Handling detail at a high level requires elaborate modeling techniques. These techniques trade the difficulty of analyzing a billion individual elements for the difficulty of building abstracted models. Verification is an NP complete problem, but model building time and accuracy can be traded off against one another to minimize the effort. For example, power can be modeled as turn-on current, reset current, clock edge current, nominal operation current, worst-case operating current as well as final-test current. Each element of the model requires an evaluation and these evaluations must be automated so any entity can be rapidly evaluated so that the higher levels of modeling can utilize the lower level parameters

In addition to the four guiding principles of Recommendation III, some additional guidelines are recommended:

Recommendation V: Additional Guidelines

- **Semiconductor Process Changes**
  - N/A
- **Design Methodology**
  - Reduce design verification by making designs regular.
  - Simplify design by adding microcode to reduce hardware complexity.
  - Plan power, signals, soft error handling, testing approaches, area and pinout at the chip level to make creating detail immensely simpler.
- **Design Automation**
  - N/A

FORECASTING

Design methodology must change to enable 100 nm microprocessor design. New design automation software is needed to support this new methodology. A primary requirement is an enhanced ability to forecast and estimate block characteristics before they exist. Another requirement is for compilers that build models that the forecaster and high-level verification can use.

- **Budgets / Specifications**
  - Power Distribution
  - Built -in-Test
  - Dominant Signal (Such As Clock and Reset)
  - Block Delay Distribution
  - Signal Integrity
  - Soft Error Control
  - Area / Pinout
  - Function
- **Audit Design vs. Budget**

Figure 15: Forecasting

Forecasting is intelligent estimation based upon libraries and captured knowledge. It helps model early stages of a design and audit the later stages of that design. The objective is to continuously build a reservoir of design knowledge and detect design problems by analyzing differences within the information . The estimation process functions by extrapolating from current available data to a design proposal. When the extrapolation exceeds limits, more data is requested. This data may be supplied through experience, libraries, generators, or prototyping. As a design matures, the Forecaster audits implementations to assure that specifications are met. It interrogates a wide range of characteristics such as soft errors, signal integrity, power, testing, and area compliance. If a design exceeds any bound, the designer needs to be modified as early as possible. The approach must allow for successive refinement of a design and build a multi-team concurrent environment. For example, if the chip current exceeds the power capacity, then either the power rails must be redesigned or the chip must be re-architected to reduce power requirements.

- **Estimation Basis**
  - Experience
  - History
  - Intellectual Property
  - Generators
  - Prototyping
- **Model Building**
  - In-place Models Including Interconnection
  - Backannotate Physical Design Characteristics
  - Full Range of Design Parameters

Figure 16: Automated Model Builder

The Forecaster is a key element in an interconnect-centric design environment where interconnect is designed first with estimates of the functional blocks. This code must extract signal interference, block timing uncertainties, and power consumption from the physical design.

Another key element is to allow estimation and knowledge collection. The design automation process that is proposed requires that many models be generated rapidly. Therefore, an Automated Model Builder is paramount for a productive design environment. Without this capability designers may be forced to trade-off design efficiency against modeling expense. This would lead to no net gain in productivity. The Automated Model Builder must absorb the physical design characteristics for all the levels of design including the chip. It abstracts lower hierarchical level information in a form that higher levels in the hierarchy can use. This extraction must cover the full range of design parameters including signal interference, block timing with uncertainties, and power consumption. Further, the block level model needs to support enhanced delay variability due to the uncertainties of circuits and logic paths.

The Taskforce proposes significant changes to design methodology and design automation. These changes are directly forced by 100 nm processing. These changes are also forced by microprocessor design's continuance to grow in complexity, which magnifies the need for productivity. The design methodology proposed is built on the principles of problem avoidance, verification only once, interconnect-centric design, and the tethering of design changes. The Taskforce recommends that efforts focus on these new methodologies and design automation as early as possible so industry will be ready for the challenges ahead.

Recommendation VI: Productivity

- **Semiconductor Process Changes**
  - N/A
- **Design Methodology**
  - Meet in the Middle Design Approaches
  - Guidelines
  - - Verify Using Hierarchical Models
- **Design Automation**
  - Forecaster
  - Auditor
  - Model Builder

SECTION 5

Types of Testing

Testing is done many times in the design and manufacture of Integrated Circuits. Design engineers test to confirm their design and to characterize it. Manufacturing tests to reject bad devices, and in sample quantities to confirm reliability and maintain processes within specifications. Where incipient early failures cannot be detected in final test, devices must be burnt-in and then re-tested, a costly process step to be avoided if at all possible. Each will have its problems in 2003, but because the manufacturing tests contribute by far the largest set of problems, the Taskforce concentrated on them only. Manufacturing's screening tests are not only costly in themselves, but are costly in yields when they err on either side – rejecting good devices or allowing bad devices to continue processing or be shipped.

Electrical Testing Challenges

Designs are moving ahead at a faster rate than current tester technology. The amount of circuitry is exploding, requiring huge numbers of test vectors. Architecture is increasingly complex – many diverse functions on the same chip demanding a great variety of test methods.

With chip clock frequencies moving to 3.5 GHz, and with edge-rate frequencies moving well beyond that, target testing frequencies and complexity of test head circuitry become severe problems. Variation in interconnect parasitics (complex distributed R-L-C-M due to coupling and over-the-cell routing) can cause timing variations particularly difficult to test since those effects are often signal-combination-specific minor changes in timing and signal levels. This often manifests itself as clock jitter.

Some blocks on the chip may have been acquired as Intellectual Property (IP) from outside sources, and the details of their design not known even by the design team itself. The tests required by these blocks may not be known in sufficient detail.

Physical and Thermal Challenges

Making contact to chips with 4000 I/O bumps switching hundreds of amps of current, some with power sources less than 1 volt will present real challenges at the wafer probe stage. Electrically, high currents must flow through very small pressure contacts with tolerable voltage drop and with acceptably low noise. Thermally, while packaged devices at final test can be clamped into test fixtures approximating their final assembled configuration, this is not possible at wafer probe. As chip sizes increase, the contact heights become inherently less planar, increasing the problems of making solid, reliable temporary pressure contacts to them.

Heat cannot flow through the pressure contacts in the same way that it can through bumps when bonded into packages – means must be found to provide reliable, adequate heat flow and to reduce heat generation in the wafer probe stage.

Traditional test methods cannot continue to be employed which depend upon accessing internal nodes. Geometries become ultra-fine, multi-level interconnects make much circuitry inaccessible, and when packaged by flip-chip techniques, back-access is nearly prohibited.

Some proposed test methods, such as using Electron-Beam technology, impose environmental requirements such as vacuum, which will only compound these problems.

Economic Challenges

With feature size rapidly reducing and with wafer and die sizes remaining the same, manufacturing cost per transistor is declining rapidly. Unfortunately, if traditional approaches to manufacturing test continue to be pursued, equipment will continue to become more complex and therefore more costly, and testing

times will lengthen, requiring more of the expensive machines. Test equipment depreciation per transistor is forecasted to at best remain flat. Figure 17 shows these two relationships both historically and as forecast by the NTRS97 and other industry information.

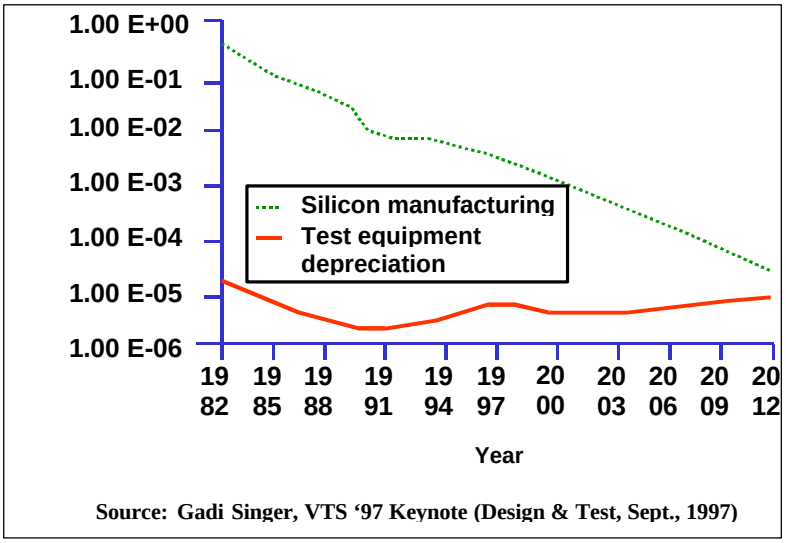

Figure 17: Test Cost Impact on Product Pricing

Test Yield Challenges

Product requirements and feature size scaling will result in silicon speed increases that will increase faster than the required increase in overall timing accuracy (OTA) of traditional manufacturing test equipment.

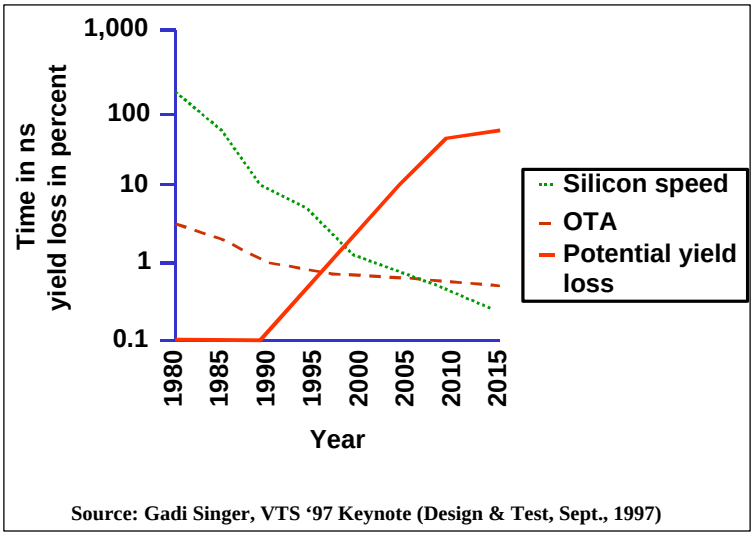

Figure 18: Yield Loss due to Guard Banding

Figure 18 shows these two relationships both historically and as forecast by the NTRS97 and other industry sources. Manufacturing test use of increased guard-banding against timing defects to compensate for the loss of timing accuracy would result in a finite and significant manufacturing YIELD LOSS. The Taskforce feels that this would unacceptably increase unit costs, so that increased guard-banding cannot be used, and other solutions must be found.

Burn-In and Reliability Challenges

With the increased number of chip input/output pins, increased supply-current, increased power dissipation, and rising chip-to-substrate interconnection complexity, the cost per burn-in socket is expected to rise sharply. In addition, alternatives to uncover potential chip manufacturing failures, e.g., use of I_{ddq} testing, are at risk for use in the future. The figure 5.3 shows these two relationships both historically and as forecast by the NTRS97 and other industry information. Alternatives to burn-in must be developed.

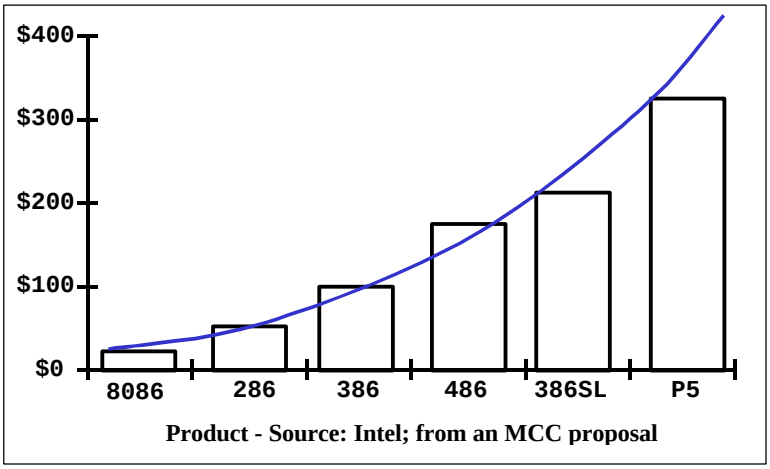

Figure 19: Cost Per Burn-in Socket Position

But the early failures against which burn-in is targeted are not the only reliability consideration. Some fabrication processes are being dramatically changed as we move to 100 nm devices, particularly the replacement of aluminum with copper in interconnect layers.

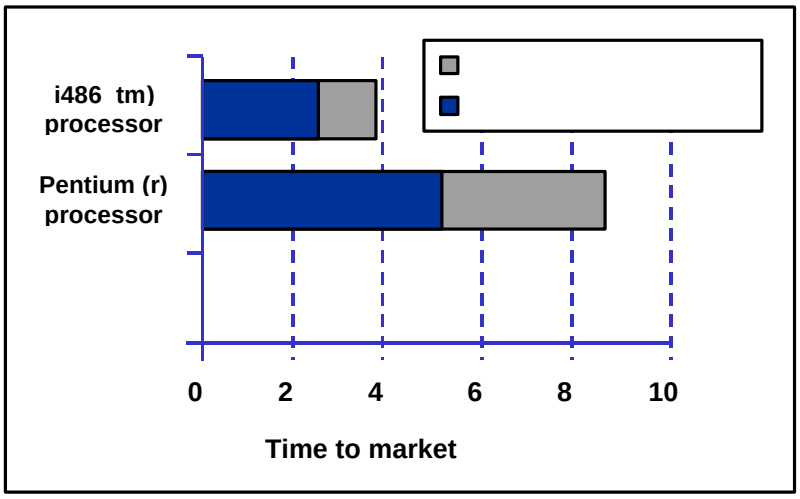

Source: Carbine and Feltman (Intel) at ITC

Figure 20: Time to Market and Volume Increasing

These dramatically change the long-term reliability picture. Failure models since the invention of the integrated circuit have been built around the properties of aluminum—electromigration in particular. Many design “rules of thumb” have these models at their base. Just as designers must not continue using obsolete rules, test engineers must not continue using reliability projections and make reliability tests based on incorrect models. A new set of fault models is needed.

Time-To-Volume Challenges

Time-to-Volume is as important as initial product release and overall Time-to-Market. This is especially true for products targeted at low end, high volume markets. The below figure shows these two relationships both historically and as forecast by the NTRS97 and other industry information. Increased test equipment complexity and therefore procurement time could delay manufacturing ramp-up unacceptably in 2003 – other solutions are imperative.

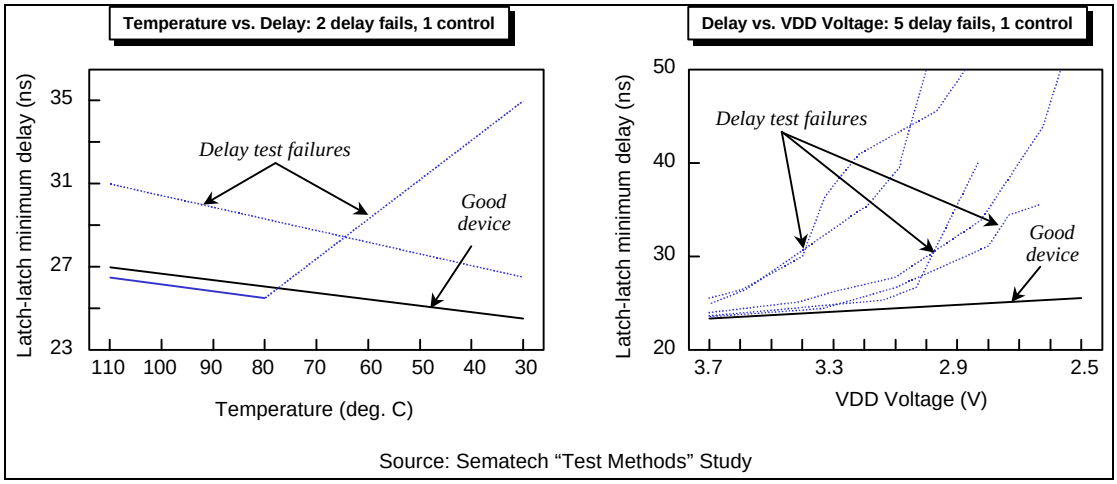

Figure 21: Temperature and Voltage Stress Identifies Faulty Behavior

Defect Identification through Accelerated Stress Testing

An important 100 nm test challenge to resolve will be to establish if parametric stress conditions can help to distinguish defect behavior from normal behavior, for example, the use of temperature stress or voltage stress in finding defective devices. It was noted above that I_{ddq} may no longer be effective, yet elimination of burn-in is crucial. Figure 21 shows example data illustrating the possible roles of parametric stress testing in manufacturing testing.

SHIFT TO BUILT-IN-SELF-TEST (BIST)

The Taskforce recommends full implementation of BIST (Built-In-Self-Test - the chip should test itself.) The challenges above will only be met by taking a fundamentally different approach to manufacturing test. The function of the test hardware should be reduced to providing power, to triggering test initiation, and to collecting pass/fail data from the chip itself as it performs its own tests.

Fortunately, high-end microprocessors are uniquely suited to implement BIST. Inherently on the chip are the controls, clocks, memories, registers, timers, and level-differentiating circuitry needed to implement full-chip BIST. Some increase in these facilities may be required, but its cost is miniscule compared to the alternatives discussed above. A side benefit of BIST is that it remains available for use throughout the life of the microprocessor, and can be made a part of the end system's self test as memory test is used today.

Using BIST eliminates all the timing accuracy problems, the contact parasitic effects and the test head complexities. A timing test can be performed totally on-chip, by initiating a sequence of actions through a significantly long path and looping the result back to a timed gate – if the signal arrives within the time window that the gate is clocked open, the device passes, otherwise a failure is registered. We are not aware of any specification which could not be tested using fully on-chip BIST.

The reduction in complexity and therefore in cost of test equipment would be dramatic. Test time itself could become less important, since many chips could be testing themselves simultaneously. Only power supply and heat removal problems would need consideration.

Implementing BIST requires that it be an integral part of the chip design, hence our next recommendation:

DESIGN TESTS DURING CHIP DESIGN

Design for Test (DFT), if it is to be based on BIST, must be carried on concurrently with the R&D effort to design the chip itself and at as high a priority. No longer will it work to fully design a chip to meet its specifications and only then turn it over to Test Engineers to figure out how to test it.

No fundamentally different capabilities or architectures will need to be designed into the chip to implement BIST, but the required number, location and properties of the various elements must be included. Hence, Design Reviews should also be Test Reviews, and Test Engineers be an integral part of the design team.

A new set of CAD software will be needed to support this function. It should monitor coverage, reduce the "bookkeeping", insert the tests themselves and generally be an integral part of the design flow, interfaced well and seamlessly into that flow.

DEVELOP NEW FAULT MODELS & PLAN TESTS TO THEM

Manufacturing tests are performed solely to detect faults. Every source of potential fault must be identified, its probability of occurrence and severity of effect known and then a test strategy developed and implemented to detect it. At the root of this effort is the fault model.

New, reliable fault models will need to be developed for all of the types of net and signal interference, including complex R-L-C-M parasitic coupling and fault modes for both over and through the cell routing. The longer-term mechanisms of failure must be re-evaluated. Where they have changed, for example through the use of copper instead of aluminum, changes in screening, sampling and environmental testing must be implemented.

Universities have taken the lead in the past in much of this modeling work. This is a fertile area for them to contribute again.

Recommendation VII: Test

- **Semiconductor Process Changes**
  - N/A
- **Design Methodology**
  - Design Tests during Chip Design
    Test Engineer must be integral member of design team
    New CAD software needed for DFT
- **Design Automation**
  - Develop New Fault Models & Plan Tests to Them
  - Copper and other processing changes will invalidate some current models
    Effective testing strategies must be based on good fault models
    Fertile area for University participation

EDA System Structure

As the number of tools and the kinds of data that have to be passed between them increase, these tools must become reusable software components that make their data and services available to other tools via standardized Application Procedural Interfaces (APIs). Two facts are coming together that make this important. The first is that the number of tools will increase along with the growing number of design issues that must be considered. Thus design flows will require the use of more tools and, coupled with a growing level of design abstraction, there will be steady increase in the necessity for shared data between tools. Therefore, design systems will need to be more than mere collections of tools, but rather tightly integrated suites of tools that explicitly support homogeneous design flows.

Data needs to be managed consistently throughout design flows, without forcing an explosion in the number of file formats or undue redundancy in persistent data. Net, Pin, Cell and Instance names tend to be passed from the top of the design hierarchy to the bottom, passing through all tools. Parasitic information is passed from the extractors to timing analyzers and simulators, bottom-up. Constraints are propagated and enhanced in detail as they pass from floor-planner to synthesis tools, to placers, to global routers, and on to detail routers.

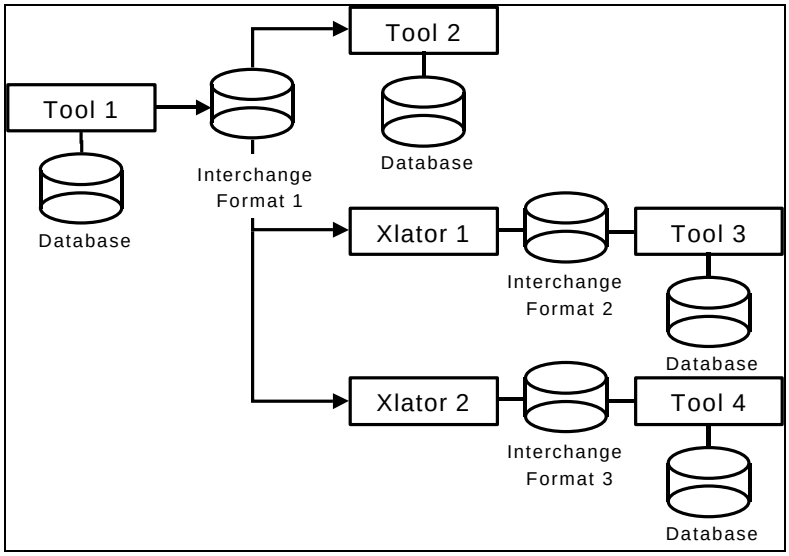

Figure 22: 6.1 File Centric EDA

As the number of tools grows and the size of the design grows, three problems arise:

- The number of tools that need to access any particular subset of the design data increases
- The size of the design databases explodes as a function of (design-complexity x number of representations x number of formats generated by translators)

Replacing the entire data by a tool as the result of only incremental changes becomes unduly costly

Using Component Object Interfaces (e.g. CORBA, COM, and Java Beans) ameliorates some of these problems because:

- There is one set of access code for each interface rather than for each tool x each format
- The size of the design database is proportional to (design-complexity x number of representations) and there only needs to be one format for each kind of data's external storage format. This saves processing time because once a tool creates a datum, any tool can access it by calling the correct API without processing any unneeded data.

- Any incrementally updated data is immediately available to any downstream tools

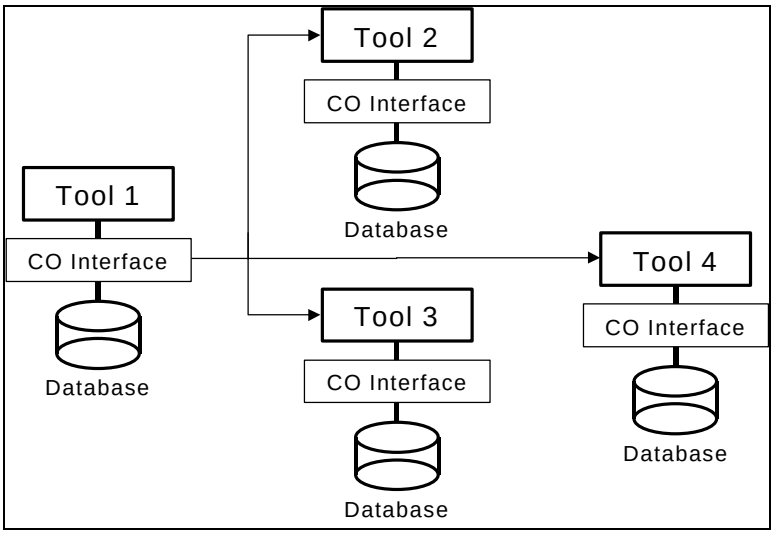

Figure 23: API Centric EDA

The second trend is based upon the tendency of CAD tool developers to use generic operating system services to the full extent possible, but not to allow the lack of such a service to limit the power of their tools. For example, early EDA layout editing tools had tool specific code for rendering the shapes in a layout, while most current tools use X-Windows or MFC code to perform these services. Access to design data has benefited from use of operating system services, but the access model will change as tools shift to the use of component objects.

Currently, EDA tools commonly communicate their data by exchanging files. Here, each tool records its resulting data into a file (in a de facto standard format, a public standard format or a proprietary format). This causes four problems:

- Each tool that needs to access the design data is required to parse the entire design file, and process all the data in the file (even if it only needs a small subset)
- If two tools that do not process the same file format need to exchange some design data, they need a translator program to map the data from one tool's format to the other's, after which the second tool will read it all again
- When a new version of a tool is released, it commonly has added new information or structure to its file format requiring either: (1) find a new revision of the receptor tools that accept the new format; (2), revise the translator to accept the new format; or, (3) write a new translator to map the new format
- Providing a multi-tool accessible file format imposes a heavy extra burden on its developers because: (1) they must provide documentation of the public format that is sufficient for people to write readers and writers in other tools; and, (2) they need to try to minimize changes in public file formats so as not to impose too heavy a maintenance burden on their users

Object Interfaces eliminate these problems:

- A tool that binds the interface can read only relevant data that it needs and not irrelevant data
- The provider of the interface maps the file data to the interface so each tool that binds the interface uses the same code to access the file data
- Interfaces are themselves versioned so if a tool changes a file format, it may support both the old version of the interface and the new version, which means that tools that depend upon the old format don't break when the new tool comes on stream

There are already a large number of support tools that help an interface supplier make it accessible to potential users and in addition, the versioning helps to reduce the impediment to evolving tools

The Taskforce feels that the expansion of EDA tools needs the kind of services that this component object architecture structure affords. The result will provide users suites of flexible tools, involving few translation steps, and evolve rapidly.

Recommendation VIII: EDA System Structure

- **Semiconductor Process Changes**
  - N/A
- **Design Methodology**
  - N/A
- **Design Automation**
  - Develop with Object Interfaces
    - *Read only Relevant data*
    - *Interface Maps from files to Objects*
    - *Versioned Interfaces*
    - *Multi-tool accessible formats*

SECTION 7

A Vision of 100 nm Microprocessor Design

The Taskforce has identified future transitions required to design microprocessors. These transitions are not limited to microprocessors. Many other chip designs, like graphics controllers, share the same characteristics with other custom chips like graphic controllers. These chips will be built in silicon processes with geometries less than 100 nm and which will operate at GHz frequencies. Even though the Taskforce concentrated on microprocessor design, it is expected that each of the paradigm shifts will affect a wide range of designs.

It is apparent that there are multiple solutions to the paradigm shifts that the Taskforce has investigated. These solutions span processing, design methodology, and design automation. Some solutions will be preferable compared to others. Some challenges will be so difficult to control that multiple techniques will be required to alleviate their effects. Each individual paradigm shift is only an element of the whole. This section of the Roadmap examines the interrelationships between design issues, and how solutions relate and complement each other.

PROCESS ENHANCEMENTS

The Taskforce has identified signal integrity as a major challenge. Signal distortion will increase as a result of power transients and signal interference and many details arise concurrently to upset signals thus requiring they be re-timed. Further, to deliver hundreds of amps to sections of the circuit instantly, the signals on the bus lines need to be stable in a determined amount of time and not have bus signals with wide variations in delay.

It is not a simple matter of designing around signal integrity, as there are too many problem sources for that approach. An engineer cannot concentrate on details when there are billions of details therefore, the challenge is in the handling of details while efficiently doing design. A significant methodology change is needed to remedy the problem. The characteristic of the problem needs to change.

SIGNAL INTEGRITY

The most effective technique of controlling inter-signal interference is by shielding and the use of power planes to supply primary current to the chip. These techniques should be expected since they have been used in PWBs and MCMs for many years.

The upside benefit of adding layers of metalization is design efficiency. The downside is a significant increase in processing cost. However, other solutions to signal integrity require spreading the elements of a design further apart, which requires larger chip area and at substantial costs. Also, the *di/dt* on power lines, clock, bus and other long high fanout interconnects will be substantial, and the inductive as well as the mutual inductive effect will be dominating.

The Taskforce recommends that the one power distribution layer be expanded to as many as four layers of metal. Two of these layers will serve as power and ground planes and two layers will serve as interconnect shielding and block shielding. The power and ground planes should be adjacent to each other with a maximum capacitance between the layers. This will reduce power spikes but additional suppression will most likely be required.

The Taskforce recommends that process technology either build a large integrated capacitor into the chip, or use MCM technologies to attach an array of capacitors to the chip. The capacitance between power and ground will need to be capable of sustaining voltages for as long as a nanosecond while the power surges reach hundreds of amps. That implies that a large capacitance is required.

Recommendation IX: Process Modifications

- **Semiconductor Process Changes**
  - Power Delivery
    - *Power and Ground Planes*
    - *On chip and/or On MCM Bypass Capacitors*
  - Signal Integrity Assurance
    - *Shielding*
    - *Low mutual capacitance and mutual inductance materials*
- **Design Methodology**
  - N/A
- **Design Automation**
  - N/A

Soft Errors

Soft-errors are random nondestructive events resulting in recoverable circuit error, which are caused by the induced noise from subatomic particles or alpha particles resulting from radioactive decay of materials. When the prevalence of signal upset is high enough, processing must find a means to harden circuits against soft errors. Silicon-on-insulator (SOI) is one method for achieving hardening. SOI has, also, the potential of being a superior high frequency technology. The Taskforce recognizes the risks that conversion to SOI may entail, but suggests that it may be necessary and the technology must be ready.

METHODOLOGY ENHANCEMENTS

Chip design methodology will require even greater changes. The objective of design is the development of a chip or chip/system that is functionally working, and that meets performance requirements in a predictable amount of design time. 100 nm designs are more complex than previous generations because many details that will overwhelm design, because there will be more transistors to design and because designs will not converge without a strategic plan.

- The Taskforce recommends that designers use a meet-in-the-middle approach

This approach uses some rigid hierarchies as the basis for design. At the top of the hierarchy is architectural definition and at the bottom is physical layout and test. The middle is the full chip level where architectural design meets physical design.

- The Taskforce recommends more use of staggered clocks or asynchronous logic

The Taskforce believes that an increasing amount of design will be accomplished at the full chip level. This may include global optimizations such as use of staggered clocks or asynchronous logic or minimizing the number of signal transitions.

- The Taskforce recommends use of rule based design

A typical rule may be to include buffering and shielding in a cell to make it signal interference resistant. Verification at the full chip level is reaching the end of its useful life. At the full chip level, there are just too many things to verify. Prevention and techniques to avoid predicaments must become the approach. From a timing standpoint, timing delays and the delay variability will be determined at low levels of hierarchy and abstracted into forms that can be applied at higher levels. The same approach is necessary for power and transient power, for test and built-in means of testing at speed, and for physical characteristics such as area and pinout.

Global interconnect will dominate over transistor delays and design must become interconnect-centric. Because the longest interconnections are between blocks, this part of the physical design

must be completed first. The rest of a design will flow from the decisions made at this level. To make these decisions, a design engineer will estimate the characteristics of functional blocks in an analogous manner to the way an engineer estimates interconnect today.

- The Taskforce recommends that low level design be considered a physical design task with constraints

In order to generate a block, its function is specified. Commonly these functions are synthesized or compiled into low level blocks that are then placed and routed. In GHz chips, however, physical design predominates over function. The function has degrees of freedom to allow the other constraints to be met. It is common to see similar system architectures be designed with different functional blocks. The methodology involves one-time verification at this level, then repeated use known-good pre-validated abstractions at other levels.

One of these constraints is a high level functional specification. At the completion of the low level design, a designer needs to build a higher level block-model that abstracts the detail but maintains the essential fidelity required to assemble blocks. It is at the block level where the full chip level will be assembled and verified. The Taskforce believes that design convergence is a function of attention paid to initial specification, attention to detail at the lowest level of hierarchy, and attention to abstracting the low level of detail (including physical back annotation) to a chip level block. As blocks require iteration, then the change should be tethered to the old circuit to reduce the magnitude of the verification effort.

Recommendation X: New Design Methodologies

- **Semiconductor Process Changes**
  - N/A
- **Design Methodology**
  - Signal Integrity Design Methodology
    - *Meet-at-the-full-chip level Design Approach*
    - *Hierarchical Design that is Interconnect-centric*
    - *Staggered Signals and Asynchronous Logic*
    - *Built-in Test*
  - Rules Based Design
    - *Constraints*
    - *Top Down Forecasting*
    - *Bottom up Model Building*
- **Design Automation**
  - N/A

High level decisions are a natural outcome of low level considerations. Low level development is driven by chip level budgets. These budgets will set the realizability of a design. So to begin a design process, an engineer will need to estimate or forecast the characteristics of the block being considered. The characteristics should include power requirements, built-in test, inter-block timing, and signal integrity factors. These factors will be assembled from archived designs and from modeling or prototyping efforts.

ELECTRONIC DESIGN AUTOMATION

New categories of design tools will be needed. These tools can be grouped into two categories. The first category is tools that improve a particular level of design. The second category is tools that improve inter-level design.

Captive design support activities and commercial EDA software companies have vigorously developed design automation in the first category. These EDA activities are numerous and they often invent new solutions for particular design problems. A few examples of needed tools in this first category are:.

- Power and supply current analysis must take into account: average as well as instantaneous surges. This analysis must determine how high power will affect signal integrity and thermal variation in delay.
- Multi-path with multi-signals timing analysis. Because delay is a function of layout as well as relationships between signal directions, a new class of timing analyzers is required to verify circuit operation.
- Placement/Route is one way to control signal integrity and power, and to minimizes the number of unwanted signal interactions. Placement/Route automation needs to support noise immune technologies such as exploitation of shielding, repeater/receiver hysteresis and transition slew rate control.
- Physical design generators for arithmetic as well as other functions need to combine synthesis/placement/routing for inter-block connections that are signal-integrity sensitive that exploits other techniques to improve noise margins. They provide rapid regeneration of pre-designed blocks (cores) using constraints for aspect ratio, timing, power, noise, etc.
- Microprocessor design will increasingly become a large software project coupled with a very regular and quite disciplined hardware design project. New code generators are needed to enable the hardware to make better use of regular structures. The design approach needs to be switched from local repair of problems to global avoidance.

The Taskforce feels that the industry will naturally find solutions for these real challenges. EDA has been capable of generating this class of solutions in the past this trend will undoubtedly continue. However, it is the second category, improving inter-level design, where EDA has not been generally successful. Inter-level implies multi-discipline, which often necessitates the integration of software that was never intended to co-exist.

As part of preparatory tasks, the Taskforce reviewed results from previous taskforces on EDA. Each one identified tool integration as the most important task ahead and yet, EDA tool integration has been met with very limited success.

- This Taskforce recommends that the picture of integration be changed. As long as each EDA sphere orbits at its own rate, alignment will be infrequent. Because there are billions of elements of data, EDA software structure must be built on a strong base of information flow with strict alignment between the tools in the flow. Thus, alignment must somehow be designed into the system. Alignment between design tools and thus EDA companies, for chip design requires a concerted methodology. The Taskforce is pessimistic that EDA tools can be interfaced into alignment and we are pessimistic that EDA be assembled into alignment. EDA has created spheres of excellence, which the Taskforce feels can be molded into alignment only through the use of Object Interface design.
- The Taskforce recommends that all EDA become interconnect-centric. The Taskforce further recommends that EDA software architectures incorporate a new set of tools. These tools aim at the same interconnect-centric goal from another view. The Interconnections that are most dominant are the longer lines that interconnect blocks or carry global signal such as clocks and busses. These interconnections are so important that they need to be known, minimized, and analyzed before defining lower levels of detail. They determine delay and skew between events and are the largest single power consumers. Further, their length and proximity to other interconnect and active devices causes the greatest level of signal integrity problems.

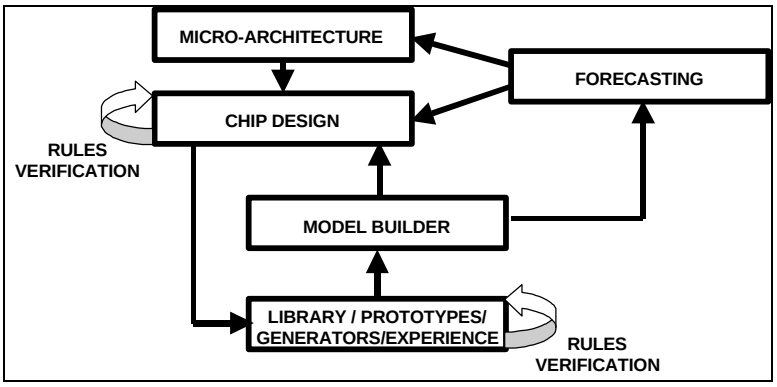

Figure 24: Interconnect-Centric EDA

- For interconnect to be designed, there has to be a set of functional blocks to connect. Therefore, the design process will initiate with estimation of the blocks that will later undergo detail design. The Taskforce suggests that EDA create an ability to forecast block design. This capability will be built upon libraries, archived design data, synthesizers and generators. This is a special case of rules driven design where each class of blocks will adhere to rules. For example:

  IF *blocktype* = MEMORY, THEN DO ___;

  IF *blocktype* = REGISTER STACK, THEN DO ___;

  IF *blocktype* = DECODER, THEN DO___.

  The forecaster value will depend on the design activity's ability to collect knowledge and build models. The models will map lower level detail into high level design. The high level design must deal will models for signal interference, block timing with uncertainties, and power consumption. The model builders will need to extrapolate from libraries, prototyping, or generators.

- The Taskforce recommends that verification be built into the design process. Rule base audits are very effective. An auditor queries implementations as they are created to avoid problems. It will interrogate the full range design space: signal integrity, power, testing, timing, and area compliance. Difficult situations will be eliminated before they cause major design iterations.

In summary, this Taskforce dares to ask software to begin again. For those who question if this is needed, let us examine EDA history. What happened to: Tegas, Calma, Daisy, Valid, Crosscheck, Scientific Calculations, and Silvar-Lisco. EDA is change. The only question is who will fill the void.

Recommendation XI: New EDA

- **Semiconductor Process Changes**
  - N/A
- **Design Methodology**
  - N/A
- **Design Automation**
  - New Design Tools
    - *Power Analyzer Considering Signal Integrity and Thermal variation in Delay*
    - *Multi-path/signal Timing Analysis*
    - *Integrated Synthesis, Placement and Routing that Controls Signal Interactions and Power Problems*
    - *BIST*
  - New Design System
    - *Object Interfaces*
    - *Forecaster and Estimator*
    - *Rules Verifier*
    - *Model Builders*

SECTION 8

Future Research

The Taskforce has encountered a number of challenges that require further research.

SOFT-ERRORS

Soft-errors are random non-destructive events resulting in recoverable circuit error, which are caused by the induced noise from subatomic particles. The perturbation caused by the particles interacting with silicon atoms can cause the value of a storage cell to change if the induced charge goes above or below a fixed level ($Q_{critical}$). Further, transient output pulses in logic circuits may indirectly produce changes in the state of other circuits if they occur at critical time periods, such as during clock or data transitions. There are two sources of these subatomic particles: alpha particles and cosmic particles.

Alpha particles are high-mass particles that result from the radioactive decay of materials within the IC or its contacts. Alpha particles can produce a great amount of ionization along their track, thus liberating large numbers of electron-hole pairs from atoms over a very short distance. If free electrons are captured on a capacitor charge plate, a single event upset (SEU) may occur resulting, for example, in the memory-state may be flipped. Additionally, transient output pulses may result on signal lines that indirectly produce changes in the state of circuits if they occur at critical time periods, such as during clock or data transitions.

Cosmic particles are high-energy particles that emanate from outer space. High energy cosmic particles can fracture the nucleus of silicon atoms thus causing a stream of secondary particles that can produce a great amount of ionization, thus liberating large numbers of electron-hole pairs from atoms over a very short distance. If free electrons are captured on capacitor charge plate, the memory-state may be flipped. Again, transient output pulses may result on signal lines that indirectly produce changes in the state of circuits if they occur at critical time periods, such as during clock or data transitions.

The probability of a soft-error occurring on an IC is a function of several factors making it difficult to quantify:

- The materials' properties, the altitude at which the IC must operate, and shielding each affect the probability of a particle strike.
- The feature size and feature packing, as well as the design itself affect the probability that a strike will cause an event upset (or soft-error).

Whether or not an occurring soft-error impacts the execution of the IC is a function of whether or not the bad state is used before it is refreshed to a correct state.

To understand the impact of feature size on the probability of soft-errors, the following should be considered. The charge of a memory device is approximated by Q=CV, where C is the gate capacitance and V is the voltage differential. Gate capacitance is approximated by $C=(\varepsilon_{GOX}A_{gate/})/T_{GOX}$, where ε_{GOX} is the dielectric constant for the gate, A_{gate} is the gate area (which scales as the square of the feature size), and T_{OX} is the oxide thickness (which scales with feature size). Thus, the gate capacitance (given no change in materials properties) decreases proportionally with feature size.

The relationship of charge to the Technology Roadmap is given in Figure 25. As shown, the state charge for latches and RAM decreases by 62.6% by the year 2003, as a result of decreasing feature size and V_{dd}. Since the charge on an electron is constant, this implies that the relative number of free electrons resulting from ionization decreases by the same percentage. Therefore, as an approximation (given no changes as a result of materials properties), the probability that a particle strike will cause an event upset might be expected to increase by the same proportion.

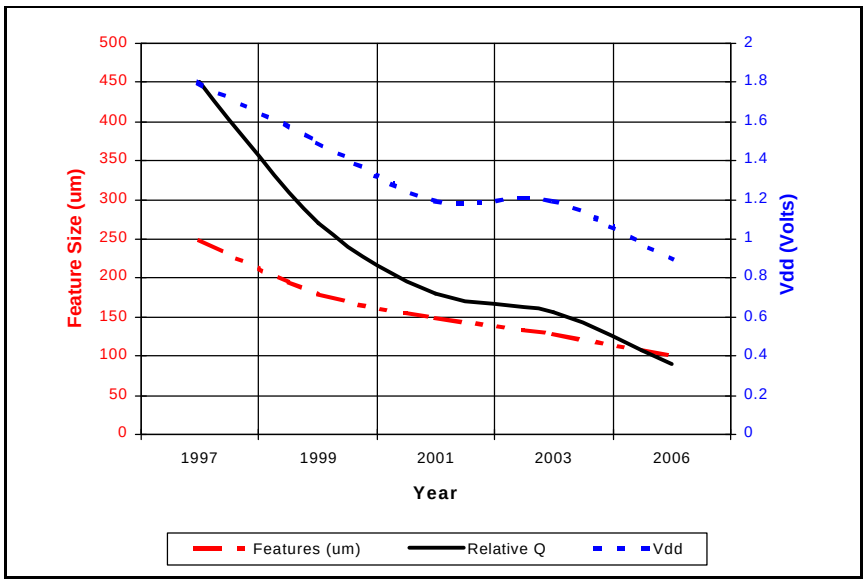

Figure 25: CMOS Charge is Decreasing

Prevention of particle strikes may be improved by the use of materials within the IC that result in less radioactive decay and by shielding the IC. Shielding the IC from cosmic particles is not practical given the extremely high energy levels. Shielding transistor gates from the free electrons that result from particle strikes may be improved by process technology such as silicon-on-insulator (SOI). With SOI, the oxide insulator can reduce the number of free electrons that are in close enough proximity to the gate to be swept up.

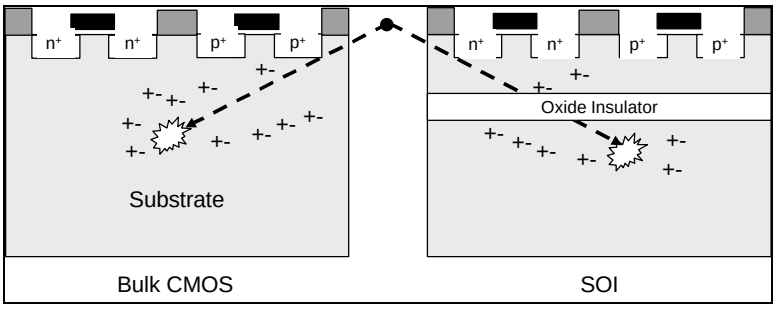

Figure 26: Soft Errors in CMOS and SOI

Detection and Correction must necessarily be designed in since soft-errors, by definition, are not hard physical defects. Therefore, it is expected that there will be increased use of parity checking and modular redundancy in microprocessor designs. Additionally, there will be more use of error correcting schemes that can correct for both single and multiple-event upsets within a functional unit —such as a memory word or register.

The probability of a soft-error that affects machine operation is extremely small, but it is increasing as ICs shrink. Even with a low probability, it is an area that deserves attention as a single error occurrence could cause the machine (using the IC) to crash or produce incorrect results. The Taskforce has evaluated several studies and has concluded that the exposure to soft errors could not be quantified, therefore further research is recommended. The following list of EDA research and development may be necessary based on the results of such research:

Detection/Correction Logic Insertion

Similar to the productivity advantages to supplying scan-register insertion tools in today's EDA systems, in the near future, automatic insertion of soft-error detection logic will be desirous. Quite possibly, this

support would allow the designer to identify susceptible functions (or regions) within a design and the detection technique (parity, modular redundancy, etc.) to be used, and allow computer automation to insert the correct detection circuitry.

Similarly, research into computer automation to insert correctly functional logic to process error correction techniques (with standard techniques such as Hamming codes, etc. available to the user) is desirable.

Physical Design

100 nm physical design tools will maintain a complex model of the rule-constrained and budget-limited physical design hierarchy elements. For each property managed in service of known-correct overall chip design, the physical tools may (for efficiency's sake) be required to compute the incremental differences in properties that are consequences of proposed local design changes. The specific requirements for such tools remain to be determined.

It is also likely that physical design considerations may affect the probability of multiple soft-errors within a design. Physical separation of circuits from their redundant copy or bits in a register may be an important consideration for effective placement algorithms. The density of latches and memory in regions of the IC will be important to understand such diverse factors as peak pulse power needs and the overall susceptibility to multiple-event upsets.

Modeling Tools

Research has shown that the critical charge level ($Q_{critical}$) is a function of the variances in a number of variables such as doping levels, spacing, and current levels. Further research and development of accurate and efficient modeling and simulation tools that can predict the probability distribution of soft error susceptibility across regions of the chip will be required. These tools must consider the manufacturing process variations, feature size and spacings, and frequency, etc.

ASYNCHRONOUS DESIGN

With the rise in power and especially with the instantaneous current surges that occur at clock edges, it will be increasingly important to disperse circuit switching times. One approach to reduce these periodic surges is asynchronous design. Asynchronous design has fallen into disuse because synchronous design has been greatly automated, while asynchronous design remains both more challenging and less well supported by tools. Research is needed to automate asynchronous design, making this design style both low-risk and relatively painless. A different approach to timing analyzers will be required because most presently deployed timing tools presume clocked synchronous design practices. Asynchronous synthesis will require enhancement and asynchronous testing approaches may need revision.

MEASURE DESIGN EFFECTIVENESS

As the Taskforce evaluated design data, it became apparent that design efficiency and productivity are not measured, and that management of productivity has been in most cases ad hoc. Those who report productivity experience a broad range of results. Often, one organization's productive environment will be less effective than another organization's unproductive design. It appears that the underlying issues go far beyond differences in expectations and perceptions. Methods are needed to measure 100 nm design team productivity. Likewise, measures of quality that determine how a design compares to criteria that are created for a process and machine architecture are needed.

Measures such as the number of transistors designed per engineering day have been used. However, when such factors are low, the value of the measure is questioned. Design complexity or other external factors often excuse poor performance. If a design team is rated by a factor, many managers will optimize the outcome for that factor. It should be an objective is to use these factors is to enhance design effectiveness and not manage an ongoing process. For that reason, a controlled experiment at the university level followed by industrial appraisal is preferable.

Multiple universities can be funded to design the same microprocessor. The microprocessor will be specified as a behavior and the design team's task will be to complete a chip design when investigating methodologies and automation techniques. In parallel with the design teams, a team of industrial

psychologists who will observe the process and factors such as task times, individual times, group times, strategies, successes, failures, iterations, along with EDA ability to perform effectively during each phase of a design. The design quality can be measured through chip size and speed along with time to design (equivalent to time to market). The factor will be drive future enhancements that will advance design effectiveness.

Recommendation XII: Further Research

- **Semiconductor Process Changes**
  - Soft-error Prevention
- **Design Methodology**
  - Measures of Design Effectiveness
- **Design Automation**
  - Physical Design for Soft-error Reduction
  - Asynchronous Design Automation

SECTION 9

FUTURE TASKFORCE ACTIVITIES

The EDA Roadmap Taskforce has investigated high-end microprocessor design. It considered other market segments as well, but concluded that attempting to analyze multiple market segments together may lead to different conclusions than focus on only one would yield. Two markets that the Taskforce feels are important to study, however, are High-End Core Based ASICs and Battery Powered Communications Products. It is recommend that two additional committees be assembled simultaneously to study both.

When staffing these committees the most important factor is identifying individuals to lead the committees with industrial backgrounds that are heavily weighted toward theoretical work. The leaders should have previously demonstrated an ability to lead volunteer organizations. These individuals must be personally committed to the technology, the Taskforce objectives, and to the other volunteers. They should build very active committees that meet physically together monthly. This Taskforce found the greatest attendance at meetings that were held in the Silicon Valley. Even the attendees that resided outside this area were more prone to attend meetings held in Silicon Valley. The sponsors must be willing to cover expenses of the committee's operation.

Recommendation XIII: Future Taskforce Activity

- **Future Taskforce Activities**
  - Battery Powered Communications
  - High-End Core Based ASICs

Appendix A: Acronyms

Source: NTRS97

A

AC, alternating current

API, application program interface

ASIC, application-specific integrated circuit

ATE, automatic test equipment

B

BiCMOS, bipolar complementary metal-oxide semiconductor

BIST, built-in self test

BIT, built-in test

BW, bandwidth

C

CAD, computer-aided design

CMOS, complementary metal-oxide semiconductor

COM, Common Object Model

CORBA, Common Object Request Broker Architecture

CPU, central processing unit

CTE, coefficient of thermal expansion

D

D&T, Design and Test

DARPA, Defense Advanced Research Programs Agency

DC, direct current

DFM, design for manufacturability

DFT, design for test

DLT, device level test

DoD, Department of Defense

DRAM, dynamic random access memory

DSM, deep sub-micron

DSP, digital signal processing

DUT, device under test

E

ECAD, engineering computer-aided design

ECC, error-correcting circuitry

ECL, emitter coupled logic

EDA, electronic design automation

EDAC, Electronic Design Automation Consortium

EDI, electronic data interchange

EDIF, electronic design interchange format

EM, electromigration

EMC, electromagnetic compatibility

EMI, electromagnetic interference

ESD, electrostatic discharge

F

FPGA, field programmable gate array

FSM, finite state machine

G, H

GHz, GigaHertz

GUI, graphical user interface

HDL, hardware description language

I

IC, integrated circuit

IDDQ, direct drain quiescent current

I/O, input/output

IP, ion projection OR intellectual property

IR, infrared

J, K, L

LSI, large-scale integration

M

MC, Monte Carlo model

MCM, multichip module

MLM, multi level metal

MOS, metal-oxide semiconductor

MOSCAP, metal-oxide semiconductor capacitor

MOSFET, metal-oxide semiconductor field effect transistor

MPU, microprocessor

N

Nm, nanometers

NMOS, negative channel metal-oxide semiconductor

Noise, introduction of unwanted voltage on a signal or power line

NSF, National Science Foundation

NTRS97, National Technology Roadmap for Semiconductors, 1997 Edition

O

OPC, optical proximity correction

OTA, overall timing accuracy

P

PMOS, positive channel metal-oxide semiconductor

PPC, process proximity correction

PSM, phase shift mask

PWB, printed wiring board

Q, R

R&D, research and development

RAM, random access memory

RF, radio frequency

RMS, root main square

RTL, Resistor Transistor Logic OR resistor transistor level

R-L-C-M, resistance-inductance-capacitance-mutual (inductance)

S

S/D, source/drain

SEU, Single Event Upset

SEMATECH, SEmiconductor Manufacturing TECHnology

SEMI, Semiconductor Equipment and Materials International

SI2, Semiconductor Integration Initiative

SIA, Semiconductor Industry Association

SMT, surface mount technology

SOC, system on a chip

SOI, silicon on insulator

SRAM, static random access memory

SRC, Semiconductor Research Corporation

T

TCAD, technology computer-aided design

TPG, test pattern generation

U

ULSI, ultra large scale integration

UV, ultraviolet

V

VHDL, VHSIC hardware descriptive language

VLSI, very large-scale integration

VSI, Virtual Sockets Interface

V_{dd} , Power Source Voltage

V_T , Threshold Voltage

W, X, Y, Z

Appendix B: References

[1] Semiconductor Industry Associates, *National Technology Roadmap for Semiconductor*s, 1997, www.sematech.org

[2] EDA Industry Council, www.si2.org/ic

[3] EDAC, Electronic Design Automation Consortium, 111 West Saint John Street, Suite 200, San Jose, California 95113, Phone: (408) 287-3322, www.edac.org

[4] SEMATECH, 2706 Montopolis, Austin, TX, Phone (512) 356-3500, www.sematech.org

[5] SRC, Semiconductor Research Corp, P.O. Box 12053, Research Triangle Park, NC 27709-2053, Phone: (919) 941-9400, www.src.org

[6] Silicon Integration Initiative, Inc., 4030 West Braker lane suite 550, Austin, TX 78759, Phone (512) 342-2244, www.si2.org

[7] DARPA, Defense Advanced Research Projects Agency, www.darpa.mil