




Gigascale Silicon Research Center Annual Report 1999

DISTRIBUTION STATEMENT A
Approved for Public Release
Distribution Unlimited





Table of Contents

1. Executive Summary	Page 2
2. Description of Current Research	Page 3
3. Activities and Interaction with Sponsors	Page 15
4. Research Project Summaries	Page 22
5. Student Information	Page 44
6. Financial Summary	Page 61



1. Executive Summary

It has been a very active year for the GSRC since we held our initial kickoff meeting in December 1998. The technical goals of the GSRC concern the development of algorithms, tools, and methodologies for increasing design productivity of reliable silicon-based electronic systems. However, of equal importance, the GSRC is an experiment in a new way of carrying out University/Industry/Government-based collaborative research. It concerns the development of a national resource--a community of researchers, from universities, industry, and government agencies, developing a compelling, long-range research vision and working together to see that vision transformed into a reality. To that end, our primary goal for this first year has been to establish such a shared understanding of the long-range issues facing the development of semiconductors and the systems they enable, to develop the infrastructure and behaviors needed to enable close collaboration, and to understand the fundamental issues that limit our abilities to improve designer productivity.

Our research focuses on four main Themes, and these themes form the basis for the research summary presented in this report. In addition to the significant progress the GSRC has made in achieving its organizational goals this year, substantial results have been achieved in each of the research themes. These results are described in the body of this report and in particular in the individual research project summaries. Highlights include the following:

- We have organized and hosted numerous quarterly reviews and research workshops throughout the year, with active attendance and participation by our industrial and government partners (see Section 3 below).
- We have developed four very challenging research Themes to focus the research and attract sponsor involvement in the research. We have developed an effective working relationship with the Berkeley Wireless Research Center for design examples.
- We have participated in many external meetings and promoted the common themes emerging from the GSRC activities within the research and development communities.
- Our annual review was held on December 9 & 10, 1999, and was attended by more than 150 industry, faculty, government, and student participants. It was widely regarded as an excellent review and certainly representative of a very productive year.

Financially, we have followed our initial budget quite closely. Due to a slow start in hiring, especially postdoctoral researchers, our spending rate indicated that we would exit the year significantly under budget. After discussion and approval from our Coordinating Council, by beefing up some of our themes and increasing our spending rate we exit our first year approximately \$0.5M under our original budget, but spending at a rate that will catch us up to plan by the end of Year 2. In our first year, we were significantly under spent on GSRC administration relative to budget. These funds were also redirected towards research projects. A detailed budget summary is included at the end of this report.

The major emphasis for next year will be on continuing to establish, and deepen, our working relationship with our sponsors. In the first part of the year, we will be visiting a number of sponsors to discuss more effective company-specific areas and mechanisms for collaboration. In addition, we will continue to strengthen and deepen our collaboration with the Interconnect FCRP. In fact, our first quarterly meeting of the GSRC will be joint with Interconnect faculty and students.



2. Description of Current Research

The MARCO/DARPA GSRC was primarily formed to meet the challenges presented by the growing design productivity gap, as described by SEMATECH some years ago. That is, while Moore's Law continues to grow at an average compounded annual growth rate, as measured in terms of our ability to manufacture logic transistors, of around 58%, the productivity of designers, as measured in terms of their ability to design and implement correct and testable transistors per staff-month seems to be growing at an under 25% compounded annual rate. This leads to a major gap in our ability to utilize effectively the potential of the silicon manufacturing process. The challenges presented by this growing design productivity gap involve all aspects of the design, verification, and test of silicon integrated circuits--from the *problems of the small*, motivated primarily by the decrease in minimum feature size and its implications, to *problems of the large*, motivated primarily by the rapid increase in the potential complexity of future integrated circuits, and including *problems of the diverse*, where the modern System-on-a-Chip (SOC) will include a wide variety of design styles (e.g. analog, RF) and technologies (e.g. FPGA, FRAM).

When we formed our initial team to consider the implications of these problems on design productivity and test, we quickly realized that by focusing our attention on any one aspect of the problem (e.g. physical design, verification, test, or system design), or even on every aspect but in isolation, we could see no way to solve the productivity problem. That is, to meet the challenges presented by the design productivity gap we needed to change the overall problem to one we could solve, rather than trying to simply improve the capacity or quality-of-results of existing algorithms and tools. Put another way, the challenge is largely one of methodology, not just algorithms and tools. Moreover, to address a methodology problem, it is clear we need a comprehensive approach, where researchers in many areas, from detailed process modeling and physical design, through to the system architectural level and even embedded software, work together on a thorough re-formulation of the overall design problem. In this sense, collaboration across many design technology disciplines, where the goal is a new and comprehensive approach to the design of integrated electronic systems, is a key aspect of the GSRC activity that we believe is one that differentiates this research effort from almost all others in the field.

To motivate, and appropriately constrain, such an ambitious effort, our team believes it is essential to carry out the design technology research in the context of one or more real design problems. Given the limited initial funding for the GSRC, and its broad scope as implied above, we were not able to fund the design of an advanced and complex circuit within the Center. However, we have formed a research partnership with another design center--the Berkeley Wireless Research Center (BWRC)[8], led by Professors Bob Brodersen and Jan Rabaey at Berkeley. The BWRC has selected three initial design targets, as described in detail on their web site, that emphasize complexity, low power, and very high performance. This emphasis is consistent with a number of the initial challenges of the GSRC effort and so we are collaborating in the development of new design technology to support their design efforts. As the GSRC evolves and tighter relationships are developed with our industrial sponsors, we expect to work more closely with them in this area as well. We also believe that the topic of circuit design and the ongoing support of advanced designs in the university environment would be an excellent theme for one of the new FCRPs to be established in the future, and we are very pleased to see that such a new FCRP has been proposed by the SIA and MARCO. To help organize and direct our initial research emphasis, in May 1998 the GSRC team developed an initial technical mission statement, as presented in Figure 1.



Empowering designers
to realize the potential of gigascale silicon by
rebuilding the RTL Foundation and by enabling
scaleable, heterogeneous, component-based design

Figure 1. GSRC Initial Mission Statement

Simply stated, two key aspects of the technical mission can be thought of as relating to the RTL-level-and-below, and the RTL-level-and-above. First, we must establish a reliable, predictable, and efficient foundation on which to build our systems, and then we must find ways of utilizing that foundation to build complex chips rapidly, correctly, and that can be yielded and tested.

To address this mission, in our initial proposal, we organized our team into a number of major Thrust areas. Each Thrust was classified as either a *technical thrust* or a thrust related to *methodology*. The Thrusts and their relationships are illustrated schematically in Figure 2.




Figure 2. GSRC Thrust Organization

We began by describing three 'horizontal' thrusts, related to *system design*, *circuit fabrics*, and *technology modeling*. The term *circuit fabrics* was selected to try to capture the idea that for a particular layout style (e.g. SRAM, standard cell, custom datapath, analog), there are an

associated set of design issues, timing, power and other metrics, and perhaps even fabric-specific yield and yield-enhancement approaches. Each one of these would be regarded as a different *circuit fabric* in our approach.



In addition to these Thrusts, we identified two key 'vertical', or cross-cutting technical thrusts: *verification* and *test*. The methodologies and techniques used in these areas must comprehend all aspects of the design problem, from system to physical design. Finally, we included a number of *methodology*-related thrusts: *collaborative design*, the *design driver* described earlier, and the topic of *design metrics*. In the metrics Thrust, the goal is to try to identify and quantify metrics appropriate to the overall chip design problem. For example, what might be the most useful way of measuring logic functional density in a design? Our goal here is to establish metrics for the objective comparison of different solutions to particular design technology-related problems.

Each Thrust has a thrust leader, and faculty choose to align their research interests with one or more thrusts. The faculty currently funded by the GSRC are listed on the web site at <http://www.gigascale.org>. More details of the specific research agendas proposed by each Thrust are also available on the GSRC web site.

We chose this organization with the principal goal of establishing a mechanism for collaboration among researchers in the same or different institutions and across the entire GSRC. We believe that this goal of collaboration is central to our effectiveness in identifying new and appropriate design methodologies for future chip designs.

The GSRC received its initial funding in December 1998, and since then we have met as an entire group for five two-day workshops, as well as many times as sub-groups, either physically or on-line. These meetings are described in detail on the web site, with full presentation and audio/video materials, as well as in the sponsor interaction summary later in this report. Over the past year, as well as carrying out the specific research related to each area, we have spent considerable time trying to really understand the methodological challenges associated with gigascale silicon. This has led to a current working structure within the GSRC that is organized around four groups that we refer to as Themes within the Center.

Over the past year, the GSRC members have organized themselves into four working groups, each one involving faculty from a number of Thrusts, or technical disciplines. The titles of these working groups are: *Constructive Fabrics* (led by Larry Pileggi), *Calibrating Achievable Design* (led by Andrew Kahng), *Component/Communication-Based Design* (led by Alberto Sangiovanni-Vincentelli), and *Fully Programmable Systems* (led by Kurt Keutzer). The first two Themes can be thought of as addressing issues raised in the first part of the GSRC mission statement: "Rebuild the RTL Foundation" while the second two Themes are associated with the idea of "Enabling Scaleable, Heterogeneous, Component-Based Design." More specific information regarding each Theme is available on the GSRC web site.

The emphasis and status of each Theme is outlined below. If funding for the GSRC is continued and grown beyond Year 2, we expect to add additional Themes, while we expect the Thrust organization, representing a management structure, to remain largely unchanged.

A. COMMUNICATION / COMPONENT-BASED DESIGNS

Theme Leader: Alberto Sangiovani-Vincentelli (U.C. Berkeley)

When considering the *problems-of-the-large* mentioned earlier, which are motivated primarily by increasing design complexity and time-to-market pressures, one of the few remaining productivity tools to exploit is the idea of design reuse. Clearly, this is not a new idea. At the printed circuit board level, each packaged component on the board is an example of design reuse. In many ways, in an ASIC design methodology design reuse is exploited regularly at the gate and medium-scale integration levels, where the ASIC library defines a collection of reusable, pre-characterized components. And of course in the software world, design reuse has been a goal for many years, most recently promoted in the form of object-oriented design with reuse of classes, as well as the component-based design approach used most often to assemble user interfaces or in products like Microsoft Visual Basic.

At the system level, in our research we make a clear distinction between the *functionality* to be implemented and the *architecture* (software and/or hardware) and *microarchitecture* upon which it is to be implemented. In a component-based approach to design, the most critical aspect of the implementation architecture is not the choice of the components themselves, but rather the *communication infrastructure*--the protocols and interfaces--used to compose the components into a system that implements the desired functionality. So while we advocate a component-based approach to design reuse, we also believe the research emphasis must be placed on the formalisms (both combinational and sequential) and implementation details associated with the composition of multiple components. Since chips are getting bigger and faster, and a single clocked-synchronous approach to managing on-chip concurrency will either be very wasteful of power or performance (or both!), we believe a major research challenge in the design of such an infrastructure is the efficient and reliable implementation of concurrency, including the interfaces to the external world. The optimal implementation and verification of (possibly concurrent), reliable communication among a collection of components on a complex chip, and their associated chip-level input-output requirements, is a major emphasis of this Theme. Here, optimality is measured in terms of some combination of performance, power, and cost, and the interpretation of 'reliable' will vary according to the situation as well. For example, some values may need to be transferred between components within a specifically bounded interval, while others may have much less stringent requirements for reliable communication. As a result, we do not believe a single, general approach (e.g. a scheme implemented on a single, standard bus) is a viable overall solution. Rather, we must develop a set of *general principles* that can be applied to and adapted for many different communication requirements and implementation strategies. These principles must be couched in a formal framework that allows for formal, or complete, verification of the critical properties of the communication medium, when such verification is necessary to guarantee reliable and correct functionality.

Most of the ideas described above were included in our initial proposal, now almost two years old. How has our thinking in this area evolved over the past year? One major change has been to recognize the importance of what we now refer to as a *platform-based* approach to design. In this regard, when thinking about communication among components it is useful to consider the successful (and unsuccessful!) examples from the past, both in hardware and in component-based software.

One very successful approach to the reliable composition of multiple components, particularly if they have been supplied from multiple sources, has been to use a very simple, clearly



unambiguous protocol (e.g. Unix pipes, RS232, TCP/IP, PCI bus). Such Simple Universal Protocols (SUP), while clearly far from optimal in terms of performance for any *particular* situation, gain their advantage from their relative simplicity and universality. SUPs have a clear and important role to play in system design, especially when it is desired to combine components designed in different organizations.

The other common situation where a component-based approach has been applied successfully is when a *single organization* is responsible for all of the components the designer might choose to use. In this case, even though the communication protocol(s) might be more complex and the requirements they impose on the implementation system might be harder to deliver, the 'single owner' of the components can verify that all of the relevant combinations of components do in fact function correctly. An example of such a Single Owner Protocol(s) (SOP) situation is Visual Basic. In this case, Microsoft is responsible for ensuring that the composition of various Visual Basic components has a predictable and meaningful behavior. Digital's MASSBUS would be considered a SOP, while its UNIBUS would be considered a SUP.

The situations in which component-based approaches have failed in the past are where the interfaces and protocols are complex and where components are developed by different groups in different locations. Usually, it is the verification issue that limits the utility of such systems.

One can think of a SOP as an example of a *platform-based* approach to systems design. That is, in a SOP world the set of components and capabilities from which the user is permitted to assemble an overall system is *restricted* in some way. Those restrictions are selected to maintain as large a degree of functional and performance flexibility as possible for a particular class of applications, while at the same time providing some bounds on the set of possible interactions and so reducing the chance (ideally, eliminating the chance) of both design and implementation errors. In most cases, the restriction is also consistent with an effective testing strategy as well.

We believe that platform-based approaches to the implementation of microelectronic systems, where the platforms may be implemented as a collection of hardware components and their interconnect, as a collection of software components and an underlying execution model, or as a combination of the two, represents a major methodology change for the semiconductor industry and is likely to be the most effective approach to the implementation of design reuse.

Major Results this Year: Professors Henzinger, Lee, and Sangiovanni-Vincentelli have played a key role in defining the underlying abstract syntax and semantics for the industry-wide system-level design language (SLDL) standards effort. As well as the general approaches outlined above, this Theme is where the majority of our verification effort (led by Professor Randy Bryant, CMU) has been focused. The verification team has focused their attention on interfaces among components, and initially on busses. Working with IBM and standards groups, the team has already found a number of errors in the PCI bus standard and has also provided significant feedback and assistance to IBM in their internal bus activity (described in more detail later in this report.)

Significant Changes in Research Directions: As mentioned above, the idea of a platform-based approach has emerged as a key aspect of our work. We have promoted this idea in the broader research community over the past eighteen months and it has now been adopted quite widely. This has also led to a significant change to our overall mission, as outlined below, and a significant emphasis on the separation of function and microarchitecture. This change has also influenced our approach to fully programmable systems, as outlined in the next section. Finally,



many of these new ideas and approaches have been promoted within the DARPA community--we believe a platform-based approach is key to providing military differentiation in a COTS-based world--and we believe this aspect of our work has had a significant influence on DARPA programs and the thinking of DARPA-funded investigators.

B. FULLY PROGRAMMABLE SYSTEMS

Theme Leader: Kurt Keutzer (U.C. Berkeley)

The other major GSRC Theme in the systems area, led by Kurt Keutzer, is concerned with exploring the limits of programmability in microelectronic systems. While the topic of design technology for embedded software and embedded run-time environments was certainly part of the original GSRC proposal, this is an area that has altered in its scope and increased both in its importance and focus within the GSRC over the past year. There are many reasons why improving the effectiveness of a programmable solution to the implementation of applications would be a better choice than developing a custom hardware solution.

The cost of a state-of-the-art microelectronics fabrication facility continues to rise and it is estimated that a new 0.18 μ m high-volume manufacturing plant costs approximately \$2-3B today. This cost is driving a continued consolidation of the back-end manufacturing process, with many existing vertically integrated manufacturing players either partnering with silicon manufacturing partners (mostly overseas, e.g. TSMC, UMC today) or moving to an entirely out-sourced manufacturing model. This increasing cost is also prejudicing the manufacturers towards parts that have guaranteed high-volume production from a single mask set (or that are likely to have high volume production, if successful.) This requirement translates to better response time and higher priorities at times when global manufacturing resources are in short supply.

In addition, the NRE costs associated with the design and tooling of complex chips are growing rapidly and the ITRS predicts that while manufacturing complex designs will be practical, at least down to 50nm minimum feature sizes, the production of practical masks and exposure systems will likely be a major bottleneck for the development of such chips. That is, the cost of masks will grow even more rapidly for these fine geometries, adding even more to the up-front NRE for a new design.

A single mask set and probe card for a state-of-the-art chip costs over \$1M today, up from less than \$100K a decade ago (note: this does not include the design cost). In addition, the cost of developing and implementing a comprehensive test for such complex designs will continue to represent an increasing fraction of a total design cost unless new approaches are developed. Once you have a clear idea and well-defined specification for a new chip, we estimate that the NRE requirements for the design and manufacture of samples for a new, complex CMOS chip in an existing state-of-the-art process today is \$7-10M. This is also a significant increase from a few years ago, and will continue to grow.

It is also the case that design validation is still the limiting factor in both time-to-market for complex embedded systems, as well as in the predictability of time-to-market (often even more important). This is due to the increasingly significant effects of physics in modern fabrication processes (e.g. affecting on-chip communication, reliable power distribution--"Do I actually get what I designed?") as well as the impact of increasing design complexity ("Do I actually get what I want?"). As chips become more complex and take on more system functionality, one of the most difficult challenges is not modeling the behavior of the chip itself, but rather it is accurately modeling the behavior of the environment in which the chip is to be used. This increase in the

context complexity of modern chip design problems significantly increases the value of being able to prototype the design in its actual final application, operating at real system speeds, before widespread deployment. The specification of the actual design requirement is often incomplete or is evolving over time. The ability of the designer to correct and turn a design quickly is becoming increasingly important, both in the commercial sector as well as in complex military applications.

We believe that these factors will increase the value of pre-characterized, optimized, and pre-verified micro-architectural families (*platforms*, as described above), as well as a *strong preference for a programmable approach*, if energy, delay, cost and reliability objectives can be met by a programmable solution.

We believe that these factors, when considered together, are *likely to lead to the creation of parameterized "standard programmable platforms" for the implementation of embedded systems*, rather than the "unique assemblies of components" approach for each new design.

However, for such approaches to be viable, it is essential that we develop the methodologies, tools, and appropriate algorithms to support the efficient development of programmable, platform-based designs.

One of the major challenges in this regard is developing an understanding what it means to program a complex SOC efficiently. The central question to be addressed here is: "What is the programmers' model?" Or "How should the programmer view the underlying hardware and input/output systems?"




Figure 3. Major Aspects of a Programmable Solution

On the one hand we want to hide as many of the details of the underlying implementation as possible, while on the other we want to make visible a sufficient level of control that the application programmer can develop an efficient solution--in terms of performance, power, and reliable functionality. Based on the earlier discussion, it is also clear that any successful solution to this problem must comprehend the issue of concurrency, both in terms of instruction-level

parallelism as well as more general forms of concurrency. This is particularly important when viewed in the context of power consumed per unit of work done (e.g. mW/MIPS). Today, depending on the application domain and nature of the programmable processor, programmed implementations of many common algorithms can be as much as 10,000 times worse than a hardware-only single-chip solution in terms of this metric. Clearly, such a penalty is unacceptable in many application domains. Of course, this is not a new problem in the general sense. The design of efficient, instruction-level parallel processor architectures, microarchitectures, and associated compilation technology is a very active field of research today. However, much of this work has been targeted towards high-end processors rather than embedded systems, and the issue of optimal power/performance has not been addressed comprehensively.

The major aspects of this problem are illustrated above in Figure 3. To gain the maximum benefit from the final silicon, we believe our research agenda must break down the conventional hardware/software boundary and explore more thoroughly the possible advantages available by exploring architecture and microarchitectural tradeoffs in conjunction with programming models, software, and compilation. Due to our need to deal with concurrency explicitly and efficiently, this work must also transcend aspects of the conventional operating system boundary as well

Major Results: The overall software for an algorithmic architecture for the Mescal system has been established and the major components of the system (e.g. Trimaran compiler base) have been acquired and evaluated for use. After reviewing a number of possible design vehicles for this work, especially BWRC designs, the team has selected high-performance network processing and in particular Intel IXP1200-like systems as an initial emphasis for the work.

Significant Changes in Research Directions: The idea of a fully programmable systems theme in itself was a significant change in research emphasis. We initially considered a "mostly programmable" emphasis, but from the perspective of the long-range nature of the research it was decided to push the limit and explore just how far a programmable approach could be taken. In addition, the realization that general concurrency issues and parallelism would be a key challenge and should be considered as a major emphasis of the work, in compilation, architecture, and verification, was a significant change in overall thinking.

C. CONSTRUCTIVE FABRICS

Theme Leader: Larry Pileggi (Carnegie Mellon)

Even if there is a move toward fully programmable platforms and component-based design, it is still necessary to design and implement the silicon itself, reliably and predictably. How to do that effectively in the gigatransistor world, where interconnect delays are a dominant factor, is the topic of this theme headed by Larry Pileggi (CMU).

As mentioned earlier, the term *circuit fabrics* was selected to try to capture the idea that for a particular layout style (e.g. SRAM, standard cell, custom datapath, analog), there are an associated set of design issues, timing, power and other metrics, and perhaps even fabric-specific yield and yield-enhancement approaches. Each one of these would be regarded as a different *circuit fabric* in our approach. An IP block, a complete system-on-chip or a new type of circuit design could be a fabric.



The idea behind constructive fabrics is on-the-fly assembly and synthesis of soft IP blocks, while trading off delay, power, area and wiring concerns. The goal is a predictable and reliable implementation of a component from the microarchitectural level. In fact, this goal represents a substantial expansion of our charter from the original proposal to "rebuild the RTL foundation."

Over a 10-year horizon, we seek to provide an ability to construct a 50,000-gate component with the same ease as a 50-gate macrocell today, and to let designers assemble 20,000 or more of these 50,000-gate components with the same level of automation they now enjoy in putting together a 50,000-gate standard cell block. At the same time, we must preserve the performance of a microprocessor-class design because these components will form the core of a wide class of circuits in a reuse-based approach. In summary, our goal is *"Microprocessor-quality components (fabrics) on an ASIC schedule."*

The emphasis is on soft IP because it's more portable to multiple foundries than hard IP, and so the major challenge is to overcome the present performance advantage of hard IP. The term "constructive" is used to capture the idea that the system architect must be able to predict the likely performance, power, and cost attributes of a component at the architectural level before the component is actually implemented at the detailed level. In that sense, the component must be "constructed" in a very predictable manner. In much of the work in this Theme, such a construction involves the planning and prediction of the wiring aspects of the fabric first, before the transistor and gate locations are considered explicitly. We are nowhere near that level of automation today.

Then there's the sheer task of assembling thousands of complex blocks. We believe that all the nasty low-level details must be abstracted away if such an approach is to work. Designers don't worry too much about individual cells today. When we reach gigascale, we hope we can provide a methodology that extends this idea so that they won't think much about details inside a 50,000-gate block.

Research projects under the constructive-fabrics theme run from high-level design to transistor-level details. One effort is looking at the problem of simultaneously partitioning and retiming large numbers of blocks. Another is looking at new approaches to synthesis using "don't care" wires. Yet another is examining repetitive patterns in IC layout, so as to make verification possible on a gigatransistor chip. All of these activities are described in detail in the project summaries.

Entirely new types of circuit fabrics are part of the research as well. Marek-Sadowska's new "wave-steered" fabric, for instance, promises predictability and performance as clocks steer signals through binary decision diagram (BDD) nodes. The concept of "noise-immune fabrics", pioneered by Brayton, Sangiovanni-Vincentelli and their students, suggests that signal and ground shielding is routed to every wire in the design. While this is a fairly extreme view, evaluating such radical approaches is what the GSRC is about.

In addition, with the rapid advance of IC manufacture technology, it has been recognized that, as the clock rate and signal switching speed continue to increase, the fidelity of power and ground distribution systems has become one of the major bottlenecks for further advances of high-speed circuits and systems. For a gigascale integrated system on chip, a die may not be able to be packaged if the die and package are not co-designed and co-simulated. Area-IO opens up a new paradigm for trading-off on-chip interconnect versus in-package interconnect. In the past, we



have worked with Intel on a power distribution scheme based flip-chip and area-IO technology, which has been used recently in a product chip. Right now we are working with Intel on the modeling and analysis of a power and ground distribution system for a future microprocessor chip. Clearly, this work must be integrated with the Interconnect FCRP as well.

Compared with other analog circuits, RF wireless applications generally make more extensive use of passive components, especially inductors. However, many of the passive components are poor candidates for on-chip integration due to their large values and stringent requirements on their high quality factor, Q . Integrating passive components on a silicon substrate together with highly integrated CMOS die in a single package reduces undesirable parasitics and improves the predictability of circuit behavior. Bell Labs / Lucent Technologies has been using our modeling and analysis tools to design integrated passives for their wireless products.

Finally, Professor Kaushik Roy has also developed a dynamic noise model to analyze noise immunity of precharge-evaluate circuits. This year, he spent six weeks at Intel in the Circuits Research Laboratory, to validate and update the model.

Major Results: We have developed and evaluated a number of new circuit fabrics, including the wave-steering and noise-immune fabrics mentioned above. We have developed a number of basic techniques in the areas of clock design, testability, and block-based implementation and sequential time management. Most importantly, we have developed working relationships with a number of companies, including IBM, Intel, Synopsys and Cadence, in this important implementation-oriented area.

Significant Changes in Research Directions: The idea of built-in fully-program-testable systems (including analog) is a substantial change and new focus in our test effort. The realization of noise-immune fabrics and their characteristics is a radical new concept.

D. CALIBRATING ACHIEVABLE DESIGN

Theme Leader: Andrew Kahng (U.C. Los Angeles)

This theme is very different from the other three. It is not as much focused on technical aspects of how to do design, but more on the process of how we do CAD research, how we evaluate it and how we integrate that evaluation back into a feedback process for improving our efforts. If the key issue is keeping up with Moore's Law (the productivity gap), then shortening the time-to-implementation of a state-of-the-art EDA system can provide as much overall productivity gain as a major new algorithm or methodology.

This theme, led by Andrew Kahng (UCLA) attempts to answer three questions: What is the design problem going to look like in the future? How can we put together the right EDA solutions to address it quickly? And did a given solution really improve the overall design process?

We anticipate that the overall impact of this work is that we would have much more directed and focused research. We would not have graduate students wasting their time reinventing the wheel. We would also have better conduits among designers, EDA tool organizations and academic R&D.

The first aspect of the theme centers on the GTX (for GSRC Technology Extrapolation), a tool that provides a more formal way to develop a collaborative road map. We think of the tool as a



"Living Roadmap." It seeks to answer such questions as when inductance matters or whether 3-D extraction is required for timing convergence-questions that numerous companies and research groups are struggling to answer on their own.

The GTX tool takes in parameter files and rules, and outputs various kinds of reports, plots and graphs. It is a framework that allows one to capture one's knowledge, manipulate that knowledge via rule chains, and visualize the implications. It can be built by many people simultaneously in the true nature of our collaborative goals.

The second aspect of the calibration theme is a relatively new idea: IP for CAD developers. If one looks at modern EDA systems, one sees some very fundamental building blocks, like placers, Steiner tree constructors and delay calculators. The argument for "EDA reuse" in the context of this GSRC theme is very analogous to that for hardware systems — it is a time-to-market and quality issue.

From that concept has arisen the GSRC multi-university Bookshelf project that collects and stores reusable EDA algorithms. At the end of our first year of work, the Bookshelf consists of 10 active "slots," or categories, that contain data formats, benchmarks, binary utilities, source code and documentation. Two of the slots are from Professor Kahng's UCLA research on placement and hypergraph partitioning.

Finally, this theme seeks to measure the impact of EDA tools through a metrics-based approach. In this work, a library of metrics-oriented API's is used to build into a tool, so as to produce human-readable output for such metrics as wire length and CPU time. It transmits data via XML format over the Internet to an Oracle database server, where the results of an EDA tool or system can be evaluated and inefficiencies or pathologies found. The metrics transmitter in UCLA's current placer imposes less than one percent overhead on run time, so overall system performance is not degraded.

Major Results: The GTX system has been built and a number of industry-wide models have been implemented in the system. With its Java interface, GTX can be downloaded and used from the GSRC website. The Bookshelf framework has been implemented and ten slots established. Standards for tools have been determined in a number of areas. The prototype of the metrics system has been implemented on a university-based tool flow. The goal is to convince commercial developers to support the metrics interfaces in their systems.

Significant Changes in Research Directions: The evolution of GTX as a "Living Roadmap" is a new idea, which we have developed and are promoting to the general ITRS Roadmap community. The Bookshelf concept, as currently articulated and implemented, was not part of our original proposal but has become an important aspect of our work.

SUMMARY OF RESEARCH STATUS AND PROGRESS

The SIA/DARPA MARCO program, and specifically the Gigascale Silicon Research Center for Design and Test, is an experiment in the area of university-based collaborative research. It involves close collaboration among faculty, with industry, and with other Department of Defense-sponsored research groups. The initial team has evolved and refined its research emphasis over the past year and we believe we have already made significant progress in identifying and formulating a number of the key methodological and technical research problems we will face in

design technology for integrated systems in the years to come. Just as important, every member of the research team has a shared understanding of this comprehensive research vision.

As our funding increases, and as the mix of problems changes, we expect our major Themes to evolve as well. We will expand the research team and adjust the mix to match this evolution. One unifying aspect of our work to date is that we believe an important methodology shift, one that would result in a major improvement in design productivity, cost effectiveness, and time-to-market, is a shift to a programmable, platform-based approach to SOC design. To enable this transition, there are many challenging research issues the community must address. These include developing a better understanding of on-chip as well as off-chip communication requirements in this context, especially with respect to reliable and predictable concurrent communication, and how to fit them into a robust verification framework. We also believe that for a variety of both technical as well as business reasons, if we could develop effective programmable approaches to the customization of the behavior of such platforms, there would be an extremely high pay-off for the semiconductor industry. As a result of our work this year and for the reasons outlined in the sections above, we have changed our initial mission statement to the one below.

Empowering designers
to realize the potential of gigascale silicon by
enabling an efficient, single-pass route to an implementation from a microarchitecture and
by enabling **scaleable, heterogeneous, component-based design**

Figure 4. GSRC Current Mission Statement

We now firmly believe that the architecture/microarchitecture boundary will be the next major evolving pivot-point for the industry and for design methodology. We believe that programmability, in all its forms, and managing concurrency, in all its manifestations in hardware as well as software, will be two of the major challenges facing the semiconductor industry in the years ahead. Finally, we believe that a reliable and predictable, while still very efficient, implementation flow from the microarchitectural level will be essential to realize the potential of gigascale silicon.

REFERENCES

- [1] Semiconductor Industry Association(SIA) can be found at <http://www.semichips.org/>
- [2] SEMI/SEMATECH can be found with SEMATECH at <http://www.sematech.org/public/index.htm>
- [3] The MARCO web site is <http://marco.fcrp.org/>
- [4] The Defense Advanced Research Projects Agency (DARPA) web site is <http://www.darpa.mil>
- [5] The International Technology Roadmap for Semiconductors can be found at <http://www.itrs.net/ntrs/publntrs.nsf>
- [6] The web site for the Interconnect FCRP is located at <http://www.ifc.gatech.edu/>
- [7] The GSRC web site is presently located at <http://gigascale.eecs.berkeley.edu>
- [8] The BWRC web site is located at <http://bwrc.eecs.berkeley.edu>



3. Activities and Interaction with Sponsors

Over the past year, GSRC faculty, postdoctoral researchers, and students have engaged with industry and government sponsors in many different ways and in a wide variety of forums. While it is not possible to capture all of the interactions, especially the informal interactions and influence on industry and government via the press or via presentations or panel discussions at conferences, in this section we summarize many of the more formal interactions between the GSRC and its sponsors. In this first year of GSRC activity, the majority of these interactions have emphasized the development of a shared understanding of what joint projects are possible, practical, and of mutual interest. A number of the collaborations are under way and some have actually already delivered useful results to our partners. Overall, our faculty have been involved in over 120 GSRC-related meetings with industry and government agencies this year, where the current perspectives and plans of the GSRC have been discussed in the context of the needs and emphases of our sponsors. We have contributed as individuals to the International Technology Roadmap for Semiconductors, as well as to many DARPA requests for materials and planning workshops, in particular the ITO SOC workshop. While we count only four company researchers currently in residence at GSRC locations, there are a substantial number of ongoing collaborations being undertaken via remote interaction with the sponsors, many of which are presented below.

JOINT MEETINGS AND WORKSHOPS

GSRC Quarterly Review Meetings

During the past year, the GSRC sponsored five "quarterly" two-day meetings and these meetings are all available for review on the GSRC web site (<http://www.gigascale.org>). All presentations were captured with full audio, video, and slides, along with summaries of the meetings. Representatives of Sponsor companies and DARPA were invited to all of the quarterly reviews, with an average of eight sponsor participants per meeting. At each meeting, a sponsoring company representative was invited to give an evening presentation and to participate in follow-on discussion. During the year, representatives from Intel, IBM, Motorola, and Applied Materials made such presentations. At our first meeting, we involved a number of faculty from the Interconnect FCRP and Professors Daly and Wong (Stanford) made presentations to us regarding the design-related work in the interconnect FCRP.

IBM visit to CMU (January 27th, 1999, CMU)

This was a broad kickoff meeting to evaluate possible joint GSRC activities with CMU researchers. Ellen Yoffa, John Darringer, Steve Oakland, Mike Trick, Bill Lee, and Rich Colbourne of IBM, along with Professors Larry Pileggi, Wojtek Maly, Randy Bryant, Ed Clarke from CMU were in attendance. Overall, about 30 researchers, including students and postdocs, participated in the meeting. This meeting led to some important collaborative projects that have already produced useful research results on bus and protocol verification for IBM, as described later in this section.

Software Day (April 5th, Berkeley Wireless Research Center (BWRC))

Dr John Reekie, GSRC, led a software development skills and techniques all-day workshop at Berkeley. Using Code Complete as a basis for the workshop, along with notes developed in the Ptolemy project, the 42 attendees from four universities and including four industrial

representatives learned a variety of important lessons regarding the development of supportable and reusable code. Overall, it was an excellent workshop and we received rave reviews from the attendees. All materials from the Workshop are posted on the GSRC Website under [Introduction to Software Practice](#).

System-Level Design Languages and Formal Methods (April 28th, UC Berkeley)

The purpose of the meeting was to explore ways in which the formal modeling and verification community, and the GSRC, might help the industry-wide SLDL effort to reach a useful outcome. The meeting was organized by Professors Edward A. Lee (Berkeley) and Carolyn Talcott (Stanford) and there were eight attendees, including representatives from VHDL International, Stanford, UC Berkeley (Newton and Lee), and U. of Cincinnati.

DSM Data Model Review (June 1st, BWRC)

We met with representatives from Intel, IBM, Cadence, the GSRC and the BWRC, to discuss the appropriate role for the GSRC in the development of an industry-standard data model. We reviewed the Intel Nike data model, as well as the IBM/Sematech CHDStd data model and associated tools. We agreed that provided funding could be found, the GSRC would be a potential candidate for the review and development of a new industrial data model. It was agreed that Si2 would be a likely supporter of such a project and would handle the overall management of the data. We agreed to hold a second review of the ideas in September timeframe. All of the material associated with this meeting can be found in the datamodel SIG on the GSRC web site.

GSRC Bookshelf and Data Modeling Workshop (2nd September, BWRC)

Led by Andrew Kahng, UCLA, we spent half of the day reviewing and discussing the software management requirements and other aspects of the implementation of the GSRC Bookshelf project. Christopher Hylands, Jon Forrest, and Aaron Walberg of the GSRC technical staff were active participants in the meeting, along with student and postdocs from Berkeley and UCLA. In the afternoon, we shifted gears to the data modeling effort, with representatives from Intel (3), IBM (2), Cadence, Sematech and Si2. Details of the meeting can be found in the GSRC datamodel SIG on the web site.

GSRC Test Thrust Meeting with Industry (Sep. 28th, Atlantic City during 1999 ITC)

This meeting was organized by Professor Tim Cheng to discuss and review GSRC test-related activities. There were about eighteen attendees, including the three GSRC test group faculty. Companies and universities represented included: Intel (Rob Roy), HP (Rob Aitken), Lucent (Vishwani Agrawal), Synopsys (Tom Williams and Rohit Kapur), Motorola (Magdy S. Abadir), IBM (Phil Nigh), Texas Instruments (Debashis Bhattacharya), LogicVision (Yervant Zorian), SUN (Scott Davidson), UCSB (Tim Cheng and Angela Krstic), Purdue (Kaushik Roy), UCSD (Sujit Dey).

GSRC Semantics Project and SLDL (Dec 3rd, BWRC)


This meeting was organized by Professors Edward A. Lee, Tom Henzinger, and Alberto Sangiovanni-Vincentelli. Following on to the April meeting on this topic, this group explored the role of the GSRC in helping industry develop a useful system-level design and description language and modeling systems. There were about five additional attendees, including representatives from VHDL International, Texas Instruments, and SRI



GSRC WEB SITE

One of the major collaborative tools we have created this year has been our web site at <http://www.gigascale.org>. Since the new site was up and running in late May, 1999, we have had continuous increase in web traffic, much of which has been direct interaction with our GSRC Special Interest Groups (SIGs) or the viewing and downloading of presentations and presentation materials. The figure below shows the web traffic at our site over the past six months.

GSRC Website Hits/Week, July-December 1999



During the last two months of 1999, the site had over 42,000 accesses, with almost 100,000 hits. In the same period, the site served half a gigabyte of page data and the average visitor visited 10 pages on the site. Of the recent visitors to the site, over 95% have come from outside the Berkeley domain and 42% of those have come from ".com" organizations while 20% came from ".edu"'s. While only permitted access to top-level information about the GSRC (the site is fully password protected), we have recorded visits from 23 different foreign countries to the site in the last two months as well.

We have recently begun a web-cast seminar series on the site as well, organized by Professor Robert Brayton, Berkeley, and plan to continue that weekly during 2000.

We believe that our investment in the Web and associated tools for collaboration is a key part of the GSRC experiment and are very pleased with the extent to which the site and its associated tools and services have helped our efforts.

ADVISORY GROUPS

The GSRC has formed two advisory groups: a Technical Advisory Group (TAG) and an Executive Advisory Group (EAG). On the afternoon of September 23rd, we held our first Executive Advisory Group meeting at the Stanford Park Hotel. As well as executives from sponsoring companies, the EAG includes representatives from SIA, SRC, DARPA, the BWRC and academia. We reviewed both the form and content of the GSRC and presented the group with

a set of questions. Following about 45 minutes of closed deliberation, Juri Matisoo, SIA, reported on their suggestions. In summary, the group felt we were doing a very good job, to the extent they could tell from the presentations and the discussions, but that convincing (educating) the MARCO member companies about the value of our work would probably be our biggest challenge.

The TAG consists of eight senior technical people from sponsoring companies and met three times during 1999, usually coincident with a quarterly review, and the GSRC Associate Director has been in regular contact with the TAG members between meetings as well and the TAG input and suggestions have been invaluable to us.

FEW-ON-FEW TECHNICAL INTERACTIONS

Throughout the year, GSRC faculty, staff and students were involved in numerous small meetings with sponsor companies and government agencies relating to GSRC projects. Some of the key interactions, as reported by the GSRC faculty, are summarized below. Due to the broad and multidisciplinary nature of the interactions, they are organized only very loosely and by general topic.

Professors Bryant and Clarke (CMU) have also developed a very effective working relationship with the IBM ASIC Division, specifically the group concerned with the IBM CoreConnect Bus. Last Summer, Bryant's GSRC-supported student Amit Goel worked with Bill Lee at IBM in Research Triangle Park. They performed interesting work on using model checking to verify their bus arbiter. In the process they found numerous places where the specification was not written precisely enough, and they also found some bugs in the arbiter design itself. This experience has convinced IBM that formal specification and verification is an important part of their bus architecture strategy. Amit & Bill have submitted a paper to DAC. Subsequently, the GSRC researchers at CMU have obtained a license to the CoreConnect bus documentation and VHDL models to do follow-on research. The CMU group is also in the process of licensing Rulebase, an IBM model checker based on SMV (developed by Ken McMillan at CMU). Professor Bryant visited the IBM facility in Burlington, Vermont in October for a daylong discussion on cooperative activities.

Professor Dill (Stanford) presented at the Intel MP Platform Validation Seminar on November 12th, hosted by Moenes Iskarous. He met with various people, including Ching-Tsun Chou, Moenes Iskarous, Fred Rastgar, Mani Azimi, and Steve Smith.

On November 19th, Professors Dill and Henzinger (Berkeley), along with GSRC graduate students Kanna Shimizu and Chris Wilson participated in a meeting at Intel hosted by Mani Azimi, with Moenes, Ching-Tsun, Fred Rastgar, and Mani Azimi.

Professor Tim Cheng's group had interactions with several MARCO member companies in the test area. AT ITC, he organized a GSRC test thrust meeting (see above) with a group of industrial test experts (from nine companies including Intel, IBM, Synopsys, Motorola, HP, TI, Lucent and SUN) to get industrial feedbacks on GSRC test activities and to establish effective technical interactions with MARCO member companies in the test area. He and his students have also been meeting regularly with Rob Aitken of HP, Janusz Rajski of Mentor Graphics, and Vishwani Agrawal of Lucent for in-depth technical discussions of GSRC test research directions and projects. They have also been interacting regularly with Greg Tollefson (Intel's SRC



representative in the test area), Rob Roy and Deb Mukherjee of Intel on coordinating and joint planning of both GSRC and SRC directions and projects in test.

Professor Cheng's group has been working closely with Synopsys (Ed Cheng's group) and HP (Norman Chang's group) in the area of DSM analysis and testing especially for the topic of delay testing/dynamic timing analysis considering power supply noise and crosstalk. One of our GSRC projects on deep submicron testing is in close collaboration with a Synopsys researcher (Dr. Jiang) which has already resulted in a joint paper that appeared in 1999 International Test Conference (also posted on the GSRC web site).

Professor Dey (UCSD) had close interactions throughout the year with several groups in Intel that are directly involved in automation and test for Intel processors. These include the Test Technology group in Intel, especially their manager R. Galivanche, and senior technical staff S. Chakravarty, Aditya Mukherjee of the IA-64 Processor Division, and P. Parvathala, Intel Arizona, who is responsible for defining self-test solutions for Intel processors in the future. Professor Dey also visited the Test Technology group in Intel in September, and presented the GSRC test work to obtain feedback.

Professor Dey has also met and discussed GSRC test plans with Mentor Graphics (J. Rajski, Chief Scientist), Lucent Bell Labs (T.J. Chakraborty, Senior Technical Staff), and Motorola (M. Abadir, Manager, Test and Logic Verification).

Professor Kaushik Roy has developed a dynamic noise model to analyze noise immunity of precharge-evaluate circuits. He spent six weeks at Intel (Circuits Research Laboratory, Shekhar Borkar, Portland, Oregon) to validate/update the model. He has been working closely with Vivek De, Shekhar Borkar, Desmond Kirkpatrick (the current Intel GSRC assignee to Berkeley) and others regarding this work.

Professor Cong (UCLA) has worked with Intel (through Prakash Arunachalam) about our effort to release the Interconnect Performance Estimation Models (IPEM) on the GSRC bookshelf during the October SRC review meeting at Georgia Tech (part of the early work on IPEM was developed under SRC support). Intel showed a great deal of interest in getting IPEM for their early full chip planning effort. Prakash's manager (Mohan, Mosur) contacted Cong to request a copy of the IPEM program during ICCAD. The IPEM set of executable models has been made available on the GSRC website. Cong and his students and postdoc are currently working with Intel to provide the source code to fit their APIs.

Professor Marek-Sadowska (UCSB) has discussed GSRC related research with Robi Dutta (Synopsys) and Ashok Kapoor (LSI Logic), as well as a number of representatives from Cadence design systems.

On March 3, 1999, Professor Dai (UCSC) presented "Achievable Designs: System-In-Package (SIP)" at Intel, Santa Clara. There were approximately 20 people attending the seminar. On April 2, 1999, Professor Dai presented "Integration of Large-Scale FPGA and DRAM Using Chip-on-Chip Technology" and "Modeling and Analysis of Integrated Passives for Single Package RF System" at Bell Labs, Lucent Technologies. There were approximately 20 people attending the seminar. On Sept. 9, 1999, he also presented "Soft Block Packing Based on Bounded Slicing Grid (BSG)" at Austin Research Lab, IBM, where there were approximately 10 people attending the seminar.



After attending SI2 workshop on CHDSstd at IBM in November, Professor Dai has begun interfacing his soft block packer to the IBM database. He is also working directly with IBM Austin Labs to address the critical timing issues in design planning for the GHz microprocessor. The first published paper on buffer insertion during floorplanning was the result of this collaboration with IBM Austin Labs.

Professor Larry Pileggi has worked with Sani Nassif and David LaPotin at IBM/ARL to collect data on the manufacturing variations associated with the clock distribution on the gigahertz processor. Future interactions are planned with students spending time at ARL. Aneesh Koorapaty, a GSRC-funded student, spent the summer at IBM-Austin working with Paul Villarubia on the combined physical-synthesis problem.

The Communication-based design theme, led by professor Sangiovanni-Vincentelli, has had a strong interaction with several members of the MARCO initiatives. Cadence has followed very closely the work carried out in his theme. As mentioned elsewhere, his work on standard interchange format for system level design has been inspired and closely followed by TI. Strong interaction with Intel with three ad hoc meetings: one in Berkeley and two in Santa Clara at Intel location to discuss how to use bus protocol formal verification in the design of IA64 and IA32 platforms and how to adopt a rigorous design methodology for communication in multi-processor systems. Lucent Microelectronics has been informed of this Theme's activity by Sangiovanni-Vincentelli in a recent meeting (December 2nd) and is particularly interested in platform-based design.

Professor Sharad Malik (Princeton) and his students have had continued interaction with Conexant (formerly Rockwell Semiconductor; main contact Alan Taylor) on a GSRC DSP compiler related project under the Fully Programmable Systems Theme. Conexant is looking for more ways to get involved in the Mescal project in GSRC, once they have resolved some IP related issues at their end. Conexant has been funding Malik outside of the MARCO project and plans to continue this funding over and above MARCO

PRESENTATIONS AND SEMINARS

SRC PTAB Presentation and Discussion (March 3rd, Irvine CA)

Richard Newton made a presentation to the SRC PTAB regarding the emphasis and plans of the GSRC. He emphasized the very different nature of the form of the research: emphasis on inter and intra university collaboration, as well as direct involvement with industry and government researchers. While the GSRC is emphasizing a long-range research agenda, he stressed that it was not the "range" of the research that was the fundamental differentiator between the SRC and the GSRC. He stressed the importance of the interaction between the two programs, and that having PI's involved in both programs was both a good idea and very healthy for the overall research agenda. He also volunteered to participate directly in the PTAB meetings, if requested, or that Kurt Keutzer was willing to participate as well. In that way, any questions regarding the GSRC and its interaction, or potential overlap with SRC research projects could be dealt with directly. Bill Joyner felt that was a good idea and agreed to discuss it with the PTAB and SRC management, and let Richard know.



SIA Presentation (March 16, Washington DC)

On March 16th, 1999 Richard Newton made a presentation to the SIA Technology Strategy Committee (Jim Meindl was there representing the Interconnect FRC). The emphasis of this presentation was on the organizational aspects of the GSRC, and can be found on the GSRC web site.

SIA Presentation (June 7th, San Francisco Airport)

At the request of Juri Matisoo, a second presentation was made to the SIA Technology Strategy Committee. In this presentation, the emphasis was on the research content of the GSRC and the directions the group was headed. Richard introduced the notions of the Themes and how they related to the Thrust organization.

ICCD Keynote Presentations (October 11th, Austin Texas)

Professors Newton and Sangiovanni-Vincentelli (UC Berkeley) presented keynote talks at the ICCD conference that were related to the GSRC activities. Newton's presentation (coauthored with Professor Keutzer, UC Berkeley) was an overall summary of the efforts in the GSRC, emphasizing the system-level aspects. The full paper and presentation can be found on the GSRC website.

Constructive Fabrics & Calibrating Achievable Design Workshop (Nov. 4th, BWRC)

Over twenty participants from UCLA, CMU, Berkeley, UCSB and the BWRC, including industrial visitors and GSRC assignees participated in a review of these Themes. This meeting was particularly useful in terms of focusing the goals and directions of the Theme activities, as well as for signing up contributions to the CAD Theme.

Bookshelf Workshop (November 20-21st, UCLA)

Organized by Professor Kahng (UCLA), this workshop involved over twenty participants (live and teleconference) with an emphasis on the organization and development of the EDA community bookshelf project. This was a major effort and resulted in significant progress in Bookshelf organization and direction.

4. Research Project Summaries

(Taken from Project Posters at the GSRC 1999 Annual Review)

Project Summaries for Theme 1: Calibrating Achievable Design

GTX: The MARCO GSRC Technology Extrapolation System

Faculty: Prof. Andrew Kahng (U.C. Los Angeles)

Researchers: Andrew Caldwell, Farinaz Koushanfar, Hua Lu, Igor Markov, Michael Oliver, Dirk Stroobandt

Technology extrapolation -- i.e., the calibration and prediction of achievable design in future technology generations -- drives the evolution of VLSI system architectures, design methodologies, and design tools. Via roadmapping efforts such as the International Technology Roadmap for Semiconductors (ITRS), technology extrapolation also influences levels of investment in various areas of academic research, private-sector entrepreneurial activity, and other facets of VLSI design automation.

This project involves GTX, the MARCO GSRC Technology Extrapolation system. GTX provides a robust, portable framework for the interactive specification and comparison of alternative modeling choices, e.g., for predicting system cycle time, die size, or power dissipation. Unlike previous "hard-coded" systems, GTX adopts a paradigm wherein parameters and rules allow users to flexibly capture attributes and relationships germane to VLSI technology and design. Serialized user-defined rules can be composed in numerous ways to define rule chains, which are then executed by a derivation engine to perform studies. Supporting grammars, parameter naming conventions, extension mechanisms, etc. enable GTX to incorporate -- and serve as a repository for -- literally unlimited forms of domain knowledge.

We detail the architecture of the GTX engine and how it achieves the various goals of this project. We specifically highlight ways in which GTX acts as an open-source infrastructure allowing added value from its users. Finally, we present various scenarios demonstrating use of the system.

Assessing Parameter and Model Sensitivities of Cycle-Time Predictions Using GTX

Faculty: Prof. Andrew Kahng (U.C. Los Angeles)

Researchers: Farinaz Koushanfar, Hua Lu, Dirk Stroobandt

The GTX (GSRC Technology Extrapolation) system serves as a flexible platform for integration and comparison of various studies that are aimed at calibrating and predicting achievable design in future technology generations. The flexibility of GTX makes it particularly useful for (i) development of new studies that model particular aspects of design and technology, and (ii) emulation, comparison, and evaluation of various technology extrapolation methods. In this poster, we highlight the ability of GTX to evaluate the sensitivity of existing (or newly developed) estimation methods to their input parameters and to their implicit modeling choices.

We integrate three highly influential cycle time models within GTX -- SUSPENS (Stanford University System Performance Simulator), BACPAC (Berkeley Advanced Chip Performance

Calculator) and the model of Fisher and Nesbitt which provides cycle time values for the current SIA ITRS roadmap. We first compare the clock frequencies that result when primary input parameters are common to all models, and evaluate the models' sensitivities to input parameter changes (parameter sensitivity) as well as to changes in the components of the estimation model (model sensitivity). Our results reveal a surprisingly high level of uncertainty inherent in predictions of future CPU timing. In particular, existing cycle time models are extremely sensitive to both modeling choices and to changes in device parameters. We find that SUSPENS is very sensitive to parameter changes, while BACPAC and Fisher/Nesbitt are less sensitive to most parameters (but differ significantly in the parameters to which they are sensitive). A number of interesting modeling sensitivities -- i.e., sensitivity of the technology extrapolation to (changes in) part of a given computation chain -- are also observed.

"Chip-package Co-design", integration of top level design of IC with bottom level design of package

Faculty: Prof. Wayne Dai (UC-Santa Cruz)

Researchers: Huaizhi Wu, Paul Morton, Shuo Zhang

For a gigascale integrated system on chip, a die may not be able to be packaged if the die and package are not co-designed and co-simulated. IC designers, not the package house, will design the package as a part of the design planning for the die. By integrating our soft wire planning, soft block packing and noise-avoidance routing with package design, we aim at developing a new system design methodology in which top level design of IC and bottom level design of package are performed concurrently.

Device Driver Automation

Faculty: Prof. Kurt Keutzer (U.C. Berkeley)

Researchers: Sherry Xu and Minxi Go

Writing device driver has become the bottleneck in embedded system design. Manually implementation is a time consuming and error prone task. We are currently carrying out a survey on academic researches and commercial tools for device driver automation. We will also discuss problems with those tools and future work to solve some of those problems.

An Analysis and Classification of System Level Description Languages

Faculty: Prof. Kurt Keutzer (U.C. Berkeley)

Researchers: Andrew Mihal

The goal of the Mescal project is to provide a development environment for embedded systems that encompasses the entire design flow, from high-level specification to implementation on a programmable architecture. A fundamental aspect of this research is the formulation of a design model, which is abstract enough to allow for high-level modeling yet expressive enough to allow designers to take full advantage of the intricacies of the architecture. The balance between these requirements has been the key issue at the center of numerous system-level description language (SLDL) research projects. In this work, we survey existing design languages and analyze their strengths and weaknesses against the criteria set forth in the Systems Level Design section of the Industry Standards Roadmap. This analysis uncovers the key aspects of a successful design model. By applying the lessons learned from previous SLDL projects to Mescal's own design model, we hope to create a powerful and functional embedded systems design environment.

Enabling CAD-IP Reuse: The GSRC Bookshelf

Faculty: Prof. Andrew B. Kahng (U.C. Los Angeles)

Researchers: Andrew Caldwell, Igor Markov

This work proposes a new infrastructure that is aimed at enabling scalable CAD-IP reuse. Our infrastructure, which we call the GSRC Bookshelf, combines features of (i) a public-domain, plug-and-play repository and (ii) a publication medium for leading-edge implementations, problem formulations, testcases and solutions.

The GSRC Bookshelf and its overarching goal of CAD-IP reuse seek to facilitate methodology and metaheuristic innovation, as well as algorithm adoption by arbitrary organizations in either industry or academia. This repository also provides a means of convergence to appropriate data models, APIs and in-memory representations. Specific objectives include the following:

- (i) Creating incentive structure to motivate research/effort toward leading-edge implementations.
- (ii) Providing a scalable and user-friendly software repository with multiple referred and unreferred mechanisms to encourage the widest possible participation.
- (iii) Facilitating valid comparisons with leading-edge algorithms and implementations (and invalidating excuses for not doing so) by making them available to all.
- (iv) Providing a widely accessible forum for industry to present relevant problem formulations, use models and testcases.
- (v) Allowing industry access to leading implementations, enabling in-house comparisons before committing to a new approach or algorithm.
- (vi) Supporting interoperability of solvers between problem domains.
- (vii) Encouraging reuse of CAD-IP and fundamental algorithms research by serving as a "one stop shop" for both historically significant and leading-edge implementations.

Design and development of the framework for CAD-IP reuse, and the Bookshelf repository itself, naturally draw on experiences with implementation repositories, benchmarking, and algorithm evaluation methodologies. Where previous efforts fail to ensure domain coverage/applicability, comparability, and/or reusability of entries, our proposed implementation repository specifically addresses these issues, drawing on experiences in electronic publication and the body of software engineering, open source and free software.

METRICS: A System Architecture for Design Process Optimization

Faculty: Prof. Andrew Kahng (U.C. Los Angeles)

Researchers: Stefanus Mantik

We describe the architecture and prototype implementation of METRICS, a system aimed at improving design productivity through new infrastructure for design process optimization. A key precept for METRICS is that measuring a design process is a prerequisite to learning how to optimize that design process and continuously achieve maximum productivity. METRICS, therefore, (i) gathers characteristics of design artifacts, design process, and communication during system development effort, and (ii) analyzes and compares that data to analogous data from prior efforts. METRICS infrastructure consists of three parts: (i) a standard metrics schema, along with metrics transmittal capabilities embedded directly into EDA tools or into wrappers around tools; (ii) a metrics data warehouse and API for metrics retrieval; and (iii) data mining and visualization capabilities for project prediction, tracking, and diagnosis.

Salient aspects of METRICS include the following. First, a standard metrics schema, along with standard naming and semantics, allows a metric from one tool to have the same meaning as the same metric from another tool from a different vendor. Second, transmittal APIs that are easily embeddable within tools allow freedom from the "log files" that currently provide only limited visibility into EDA tools. With appropriate security and access restrictions, these APIs can prevent loss of proprietary information while yet enabling detailed tracking of the design process. Third, at the heart of METRICS is a centralized data warehouse that stores metrics information. Several means of data retrieval and visualization (e.g., web-based project tracking and prediction) afford user flexibility. Finally, industry-standard components and protocols (http, XML, Java, Oracle8i, etc.) are used to create a robust, reliable system prototype.

Modeling of Giga-Scale IC Implementation Strategy Trade-Offs, Step I: Cost Forecasting

Faculty: Prof. Wojciech Maly (Carnegie Mellon University)

Researchers: Pranab K. Nag

The goal of the research described in this abstract is the development of cost modeling capabilities that could be used in IC design flows, with the stress on cost forecasting techniques that are capable of predicting cost for yet non-existing technologies and IC designs. We have begun our research, using as a foundation a-priori cost modeling research conducted at CMU for last ten years. To date our cost modeling efforts had been focussed on three main areas: cost of manufacturing, impact of yield and yield learning on cost, and cost of testing. Thus, we developed methodologies and models to assess impact of manufacturability, testability and diagnosability of a given design. One important outcome of this effort has been the implementation of a discrete event software system for yield and cost forecasting. Through extensive simulation experiments we have been able to demonstrate that much of the cost impact is directly or indirectly related to many strategic design decisions.

Based on this foundation of development and understanding, we have directed our focus towards modeling system design related attributes and their interdependences with manufacturing, testing and diagnosis. Presently, our emphasis is on developing models to relate chip area and yield to number of available metal layers. Specifically, we are focusing on cost estimation for designs intended for 0.25 micron technologies and subsequently project the cost for 0.18 micron technology. Our experimental results show that the optimum number of metal layers can be very sensitive to both fabrication cost and nature of the IC design. Nevertheless, this approach is just one thread of inter-dependence. Another thread will be to explore net-length reduction and increasing number of contacts.

From this study we conclude that such cost estimation techniques are certainly useful and should be built right into the design process flow. In order to eventually achieve this goal we have chosen a medium-sized system-on-chip design for further analysis and modeling.

Modeling of Giga-Scale IC Implementation Strategy Trade-Offs, Step II: Case Study of Limitations of Traditional Monolithic Implementation of Modern ICs

Faculty: Prof. Wojciech Maly (Carnegie Mellon University)

Researchers: Steven Y. Deng

It is well understood that the "interconnect crisis" can be postponed by only a couple of years via development of new (e.g. low-k) interconnect materials. Therefore, one should seek the solution to the interconnect crisis by exploiting "design dimensions" of interconnect delay problem. One



component of this dimension is the length of the worst case connection, which will have to increase with the predicted increase of the die size. It will have to grow; as long as, integrated circuits are placed and routed in two-dimensional space. The obvious solution to this problem would be, therefore, a 3-D circuit integration. But it is not quite technically feasible yet. Consequently, an intermediate solution, which, for instance, uses a 3-D stack of traditional monolithic ICs, should be considered as one of the avenues of addressing the interconnect crisis. Also new system design methodologies should be developed to fully take advantages of such a new system integration strategy. It is not done in reality.

We have begun our investigations by focusing on a case study that assesses potential of 2.5-D system integration strategy (i.e. a strategy in which IC chips are bonded together one on the top of another). Our goal is to find out when and under what circumstances it is more cost effective to implement a given design on two (or more) chips connected by vertical connectors build in any place of chips overlapping areas, the hypothesis being that this strategy can lead to a substantial reduction of the total silicon area (and thus gain yield) and substantial reduction in interconnect length (and thus increase in performance).

In order to assess validity of the above hypothesis, we developed two "2.5D" CAD tools – simulated annealing based standard cell placer and a general floorplanner. The simulated annealing based placer explores fine-grained characteristics by placing a gate-level netlist into two folded dies that are interconnected by "2.5D vias". Our experiments, performed on several industrial circuits, show roughly 33% and 40% reduction in area and total wire length, respectively. Such gains are only possible if the "2.5D via" can be placed on the top of the standard cells or it has to be small (comparable to an inverter) in size. The BSG based floorplanner works on the top level of physical design. Two well known benchmark circuits are used for this purpose. Since the number of modules available in the netlist are small, but individually large in size, the area reduction is trivial (2% and 3%). However, 20% and 30% reductions in total wire length (global wires at the top level of design hierarchy) are observed. From this study, we conclude that 2.5D integration scheme may offer substantial advantage over traditional monolithic approach.

Noise-avoidance in detail routing and global routing

Faculty: Prof. Wayne Dai (UC-Santa Cruz)

Researchers: Huaizhi Wu, Paul Morton, Shuo Zhang

As technology scales down, lateral coupling becomes increasingly more significant than coupling to ground. Currently, most coupling analysis is done with the precise geometry information of wires, which can only be obtained after the detail routing. But at same time, the whole layout is so determined that it is difficult to fix for the noise problem. To resolve this confliction, we have a new approach in which the coupling noise can be identified and diminished the immediately following topological routing, which is the point in the routing process that gives the best trade off between the ability to detect noise problems and the ability to correct the problem. Furthermore, in the earlier stage, global routing, we can also evaluate the degree of coupling conflict between critical nets and decrease it. Such pre-detail-routing method provides the most freedom to avoid the coupling noise.

The Real Impact of Vias on Available Routing

Faculty: Prof. Andrew Kahng (U.C. Los Angeles)

Researchers: Stefanus Mantik and Dirk Stroobandt

Vias used to connect wire segments between two adjacent (H- and V-) layers are typically not considered harmful. However, vias that provide a connection between two layers further apart (e.g., between the bottom and top layers) also take up resources on all intermediate layers. Other wires on the intermediate layers are prevented from using the space taken up by the vias. Traditionally, the via impact has been estimated by only taking into account the actual area occupied by the vias. We show that this greatly underestimates the via impact: routing wires in the neighborhood of vias cannot be achieved without using detours and extra resources. We therefore propose a more realistic via impact model that takes wiring criteria into account. Our model probabilistically captures the number of routing segments that effectively can no longer be used due to the presence of vias.

We have implemented our new via impact model within GTX, the GSRC Technology Extrapolation system. GTX allows us to quickly analyze the results and the impact of changes in several parameters of the model. With this new model, we are able to predict more accurately the number of wiring layers needed for accommodating signal wiring. The model's assumptions and predictions are validated through an experimental framework in which real designs are placed and routed and the effects of additional vias measured.

Soft block packing with soft wire planning

Faculty: Prof. Wayne Dai (UC-Santa Cruz)

Researchers: Huaizhi Wu, Paul Morton, Shuo Zhang

In consistency with the theme of constructive fabrics, which aims at providing the flexibility required for system level fabric assembly by synthesizing soft fabrics, we have incorporated soft blocks and soft wires into our floorplanning and wireplanning stages of physical design. With the capability of reshaping soft blocks, we can greatly reduce the total packing area and inter-block wire length to meet the area and timing constraints. In addition, by incorporating soft wires whose sizes and shapes have not been fixed yet, we can further reduce interconnect delay. Such geometry flexibility of soft blocks and soft wires can considerably increase the likelihood of first-pass silicon success.

Project Summaries for Theme 2: Constructive Fabrics

Analysis, Modeling and Test Generation for New DSM Failure Modes

Faculty: Prof. Tim Cheng (U.C. Santa Barbara)

Researchers: Angela Krstic, Yi-Min Jiang, Jing-Jia Liou

Noise effects such as power supply noise and crosstalk can significantly affect the performance of deep submicron designs. Existing analysis and test generation techniques cannot capture the effects of noise on cell/interconnect delays. This is because these delay effects are highly input pattern dependent. We propose new techniques for modeling, abstraction and analysis for emerging deep submicron defects to support efficient test development, test evaluation, test generation and BIST.



Binary and Multi-valued SPFD-based Wire Removal in PLA Networks

Faculty: Profs. Robert Brayton, Alberto Sangiovanni-Vincentelli (UC-Berkeley)

Researchers: Sunil Khatri, Sinha

This paper describes the application of binary and multi-valued SPFD-based wire removal techniques for circuit implementations utilizing networks of PLAs. It has been shown that a design style based on a multi-level network of approximately equal-sized PLAs results in a dense, fast, and crosstalk-resistant layout. Wire removal is a technique where the total number of wires between individual circuit nodes is reduced, either by removing wires, or replacing them with other existing wires. Three separate wire removal experiments are performed. Either wire removal is invoked before clustering the original netlist into a network of PLAs, or after clustering, or both before and after clustering. For wire removal before clustering, binary SPFD-based wire removal is used. For wire removal after clustering, multi-valued SPFD-based wire removal is used since the multi-output PLAs can be viewed as multi-valued single output nodes. We demonstrate that these techniques are effective. The most effective approach is to perform wire removal both before and after clustering. Using these techniques, we obtain a reduction in placed and routed circuit area of about 11%. This reduction is significantly higher (about 20%) for the larger circuits we used in our experiments.

Design and Test of High Speed Scaled CMOS Circuits

Faculty: Prof. Kaushik Roy (Purdue University)

Researchers: Seung Hoon Choi

A dynamic noise model is developed and used to analyze the noise immunity of high-speed circuits. Considering that the primary source of noise-injection in the circuit is cross-talk, a simple metric represented as voltage-time product can be used to quantify the dynamic noise-margin. Based on the information on dynamic noise-margin of the circuit, we develop test strategies for cross-talk faults. Future work also will include design of noise-tolerant circuits.

Don't Care Wires in Logical/Physical Design

Faculty: Prof. Robert K. Brayton (UC-Berkeley)

Researchers: Philip Chong, Yunjian Jiang, Sunil Khatri, Subarna Sinha

A layout is obtained before the complete functionalities of individual cells are set. In particular, given an initial decomposition of a logic module, we determine, for each input to the nodes of the multi-level network, a set of "compatible" alternate wires, i.e. for every pin of the network, there is a set of sources. These sets are compatible, like don't cares in logic; any combination can be used (as long as there is exactly one source for each pin) and each combination determines a different set of functionalities within the cells. SPFDs are used to find compatible sets of alternate wires. We use this wiring flexibility to construct a wireplan; the final assignment of sources to pins depends on the ease with which the resulting netlist can be placed and wired. Once an assignment is chosen, the logic within each cell is determined. This partially reverses the classical approach of doing logic synthesis first, followed by physical layout and leads to an "information" driven placement methodology. We show results demonstrating total wire length improvements using this new flexibility.

FSM performance optimization through Wave Steering, Part I: Decomposition Theory

Faculty: Prof. Malgorzata Marek-Sadowska

Researchers: Luca Macchiarulo, Shih-Ming Shu

Wave steering is a systematic technique that allows to substantially increase throughput of combinational circuits. The presence of latency-dominated constraints (also known as iteration bounds), prevents a straightforward extension of any pipelining scheme, including wave-steering, to sequential circuits. However, wave steering permits to choose the order in which inputs are introduced. Placing the state variables at the top most part of the circuit allows to reduce the length of the iteration loop to the number of state variables. This can be further improved by using classical Hartmanis-Stearns decomposition theory in order to obtain cascade decompositions of machines with a smaller number of variables in the feedback loop. The results on MCNC benchmarks are encouraging.

FSM performance optimization through Wave Steering, Part II: One-Hot Encoding

Faculty: Prof. Malgorzata Marek-Sadowska

Researcher: Luca Macchiarulo

Another way of relaxing the feedback constraint takes advantage of the properties of one-hot encoding. When an FSM is one-hot encoded, each state can be implemented by a different machine whose inputs are the PIs and other state bits. The iteration bound, in this case, is limited by the maximum number of incoming edges in a STG representation of the machine that can be reduced by an appropriate state-splitting technique. Other circuit modifications permit further improvements. The single state-variable machine can be arranged in a standard cell fashion, thus simplifying the placement and routing issues (all PIs can be routed simultaneously by aligning the PI-part of the machines). This also allows a simple output function realization.

Hierarchical Abstraction of RLC Interconnect Circuits

Faculty: Prof. Larry Pileggi (Carnegie Mellon University)

Researchers: Satrajit Gupta, Michael Beattie

New models for RLC interconnect that allow magnetic and electrostatic coupling between groups of interconnect objects to be macro-modeled via individual hierarchical interconnect components. Such models are necessary for parasitic extraction to keep pace with component-based gigascale system design.

Modeling and Analysis of Design Fluctuations to Facilitate Robust Design

Faculty: Prof. Larry Pileggi (Carnegie Mellon University)

Researchers: Emrah Acar

New methodologies for robust circuit design require new models and simulation algorithms that capture manufacturing and environmental fluctuations in the design. Modeling such phenomena becomes increasingly difficult since manufacturing variations are becoming more pronounced with larger systems and smaller feature sizes. This project takes an interconnect-centric view of the modeling and analysis problem to establish new foundations for capturing parameter fluctuations for worst- and best-case simulation.



NexSIS: Retiming and Architectural Floorplanning

Faculty: Profs. Robert K. Brayton and A. Richard Newton (U.C. Berkeley)

Researchers: Abdallah Tabbara

The conventional VLSI design flow consists of an integration of various steps and tools in order to synthesize a circuit. Typically the flow introduces a separation between the logic synthesis step and the physical design step. For designs with aggressive performance goals, this division entails hundreds of iterations between synthesis and physical design before converging to the desired implementation and achieving closure on design constraints, especially timing. In this work we address this issue and lay the groundwork for a new design flow that exploits (a) the recently proposed idea of planning for performance at the early stages of the flow in order to minimize design iteration, (b) recent work in the area of retiming, and (c) results of a recent study of delays in DSM. The retiming work promises to permit area and delay trade-off decisions at the early stages of floor-planning.

NexSIS: An IP Modeling and Integration Framework

Faculty: Profs. Robert K. Brayton, A. Richard Newton, and Alberto Sangiovanni-Vincentelli (U.C. Berkeley)

Researchers: NexSIS Team

NexSIS is a research effort intended to address design productivity, and timing closure issues in Deep Sub-Micron (DSM). Our goal is to develop a design flow and tool support for assembly, synthesis, and physical design of IP's in the context of System-on-Chip (SoC). We are investigating several research directions including design management and data modeling, as well as new architectural floorplanning, synthesis, and layout algorithms and techniques.

Nonlinear Programming Techniques for Floorplanning

Faculty: Prof. R.K. Brayton (U.C. Berkeley)

Researchers: Philip Chong

We have developed a technique for floorplanning based on general nonlinear programming techniques. The solution space for the floorplanning problem is made "smooth" by taking a differentiable problem that approximates the actual problem at hand. The approximate solution is then converted to a solution to the original problem while trying to minimize the perturbation of the modules being placed.

The primary advantage of our technique is that complex cost functions and constraints may be easily incorporated into the solution. For example, we have introduced timing constraints into the floorplanning process, a consideration that has yet to be addressed in the literature.

We have begun to try to compare our technique with existing floorplanning algorithms (i.e. sequence pair) in terms of quality of results and run times. We are also exploring extending this technique to placement of large sets of modules through clustering.

Pattern Repeatability in Deep Sub-micron VLSI Designs

Faculty: Prof. Wojciech Maly and Andrzej Strojwas (Carnegie Mellon University)

Researchers: Mariusz Niewczas and Michal Palusinski

The objective of the research described in this abstract is: first to reach an adequate level of understanding of what comprises a minimum number of geometrical patterns needed to describe any IC layout and then to apply the above knowledge in designing highly regular and well characterized components of gigascale circuits.

We have begun our research by proposing a technique for decomposing IC layout masks into repetitive geometrical patterns. A software package implementing this technique was designed to be capable of dealing with a very large industrial IC design. Recently, we have conducted first experiments using the above software to extract patterns from various IC designs, ranging from several thousand transistors (student's projects) up to state-of-the-art multi-million-transistor industrial design.

The obtained results clearly show the nature of pattern repeatability obtained with the existing design methodologies. We have found the total number of distinct patterns within a single layout mask is rather large. But typically there is a small group of patterns that covers most of the layout. On the other hand large numbers of the patterns occur only once. In addition we have found that there seems to be a lot of patterns with a "similar shape". This implies that perhaps some pattern differences are unnecessary and there is a potential to "equate them" without imposing unwanted limitations on the circuit's performance. The above finding will be used as a guidance for our future investigations.

Performance-Driven Multi-Level and Multi-Way Partitioning with Retiming

Faculty: Prof. Jason Cong (UC-Los Angeles)

Researchers: Sung Kyu Lim

Under the interconnect-centric design paradigm, partitioning is seen as the crucial step that defines the interconnect. To meet the performance requirement of today's complex design, performance driven partitioners must consider the amount as well as performance-related quality of the interconnect induced by partitioning. In this poster, we formulate the performance-driven multi-way circuit partitioning problem with consideration of the local and global interconnect delay. We develop an efficient algorithm HPM (Hierarchical Performance-driven Multi-level partitioning) that simultaneously considers cutsizes and delay minimization with retiming. HPM builds a multi-level cluster hierarchy and performs various refinement while gradually decomposing the clusters for simultaneous cutsizes and delay minimization. During the clustering phase of HPM, a delay oriented clustering builds the initial cluster structure to ensure the best possible subsequent retiming, which is then extended by a cutsizes oriented multi-level clustering. During the partitioning phase of HPM, a cutsizes oriented refinement is performed, followed by a simultaneous cutsizes and delay oriented refinement. The delay result is further improved by retiming. We provide comprehensive experimental justification for each step involved in HPM and in-depth analysis of cutsizes and delay tradeoff existing in the performance-driven partitioning problem. HPM obtains (i) 7% to 23% better delay compared to the state-of-the-art cutsizes minimization based hMetis at the expense of 19% increase in cutsizes, and (ii) 81% better cutsizes compared to the state-of-the-art delay minimization based PRIME at the expense of 6% increase in delay. This is a joint work with Chang Wu at Aplus Design Technologies.



Performance Estimation Models for Optimized Interconnects (IPEM)

Faculty: Prof. Jason Cong (UCLA)

Researchers: David Zhigang Pan and Wangning Long

The UCLA IPEM package provides a set of procedures to estimate VLSI interconnect performance under various interconnect optimization techniques for designing high-performance, deep submicron IC circuits. As a companion to the interconnect optimization engines in the UCLA TRIO (Tree-Repeater-Interconnect Optimization) package, IPEM is intended to be used for synthesis and design planning tools to fast predict the optimized interconnect behavior, and to ensure the overall design convergence. Our extensive experiments show that IPEM is reasonably accurate (90% on average) and extremely fast (constant time in practice, with a factor of 10,000x faster than TRIO). Moreover, it is very easy to be linked to user's application programs. Currently, we provide interconnect delay estimation models for two-pin nets under OWS (Optimal Wire Sizing), SDWS (Simultaneously Driver and Wire Sizing), BIWS (Buffer Insertion and Wire Sizing), and BISWS (Buffer Insertion, Sizing and Wire Sizing) optimizations. In addition, we provide an API to compute the critical length for buffer insertion with OWS optimization. For convenience to users, IPEM includes some default technology parameters based on the 1997 National Technology Roadmap for Semiconductors (NTRS) and UC Berkeley's Strawman technology. It also allows users to set up their own technology parameters.

A VLSI Design Methodology using a Network of PLAs Embedded in a Regular Layout Fabric

Faculty: Prof. Robert Brayton, Alberto Sangiovanni-Vincentelli (UC-Berkeley)

Researchers: Sunil Khatri

We present a VLSI design methodology to address the cross-talk problem, which is becoming increasingly important in Deep Sub-Micron (DSM) IC design. In our approach, we implement the logic netlist in the form of a network of medium sized PLAs. We utilize two regular layout "fabrics" in our methodology, one for areas where PLA logic is implemented, and another for routing regions between such logic blocks. We show that a single PLA implemented in the first fabric style is not only cross-talk immune, but also about 2X smaller and faster than a traditional standard-cell based implementation of the same logic. The second fabric, utilized in the routing region between individual PLAs, is also highly cross-talk immune. Additionally, in this fabric, power and ground signals are essentially "pre-routed" all over the die.

Our synthesis flow involves decomposing the design into a network of PLAs, each of which has a bounded width and height. The number of inputs and outputs of each PLA are flexible as long as the resulting PLA width is bounded. We perform folding of PLAs to achieve better logic density.

We have implemented the entire design flow using public domain software. Our scheme results in a reduction in the cross-talk between signal wires of between one and two orders of magnitude. As a result, for a 0.1 micron process, the delay variation due to cross-talk dramatically drops from 2.47:1 to 1.02:1. Additionally, our methodology results in circuits that are extremely fast and dense, with a timing improvement of about 15% and an overall area penalty of 2.4% compared to standard cells. The regular arrangement of metal conductors in our scheme results in low and highly predictable inductive and capacitive parasitics, resulting in highly predictable designs. The crosstalk immunity, high speed, low area overhead and high predictability of our methodology indicate that it is a strong candidate as the preferred design methodology in the DSM era.

Wire Sizing and Placement

Faculty: Prof. R.K. Brayton (UC-Berkeley)

Researcher: Philip Chong

A technique for estimating optimum wire sizes for each metal layer in a design was presented in [1]. Higher layers of metal tend to have wider wires, both reducing the delay along those wires and reducing the quantity and length of those wires as well. Thus there is a tradeoff here; a design may require only a few very fast wires, or it may require many wires, in which case they will necessarily be slower.

We consider an estimate for the wire length distribution of a design using a Rent's Rule-based technique. From this we can estimate an optimum size for wires for each layer of metal which will provide enough wiring capacity as well as satisfy the timing requirements for the design.

We hope to use this technique to aid placement. By considering some wires as "fast", we allow some flexibility in the placement process, since then certain pairs of cells may be placed further apart while still maintaining their timing relationship. To do so, we require knowledge of which wires will be fast and how fast they will be, and so we may be able to extend the work in [1] to provide such an estimate.

[1] P. Chong, R.K. Brayton, "Estimating and Optimizing Routing Utilization in DSM Design", in Workshop on System-Level Interconnect Prediction, 1999, pp. 97-102.

Wireload-Model Independent Technology Mapping

Faculty: Prof. Larry Pileggi

Researchers: Aneesh Koorapaty

Existing synthesis and technology mapping algorithms rely on statistical wireload models to estimate gate loading. As interconnect becomes more dominant, this modeling becomes more inaccurate and the corresponding synthesis becomes more error prone. For this reason there have been recent attempts at combining synthesis and physical design in an iterative fashion to continuously update the loading estimates. In contrast, this project attempts to entirely eliminate the need for wireload models by incorporating physical information into synthesis at the technology mapping level.

Project Summaries for Theme 3: Fully Programmable Systems

Adaptive and Retargetable Compiler Optimization

Faculty: Prof. Sharad Malik (Princeton)

Researchers: Manish Vachharajani

Most easily retargetable optimizing compilers perform optimizations that perform well for a set of applications, and will yield modest performance improvements on most microarchitectures. Retargetable research compilers, such as IMPACT and Trimaran, can optimize code reasonably well for a variety of Instruction Level Parallel (ILP) micro-architectures, but the optimization heuristics are still selected by hand and work well only on these ILP architectures. Since the compilers do not consider program properties, except for edge and block profiles, programs may



perform poorly if they are significantly different from the benchmark suite that is used to select optimizations and heuristics. A better compiler, for example, could tune register allocation heuristics used during scheduling to allow for more aggressive scheduling for program regions which cannot cause register spills. We are developing a retargetable compiler that can automatically select optimizations and optimization parameters based on the target microarchitecture as well as adapt the optimizations for the specific program being compiled to yield an aggressive, yet easily retargetable optimizing compiler.

Application Specific Customization of Embedded Operating Systems

Faculty: Prof. Sharad Malik (Princeton)

Researchers: Shaojie Wang

Embedded systems require small and efficient system software. General purpose operating systems (even commercial Real Time Operating Systems) cannot satisfy this requirement because they have to work on different applications in various environments. Consequently, embedded system developers may spend significant effort developing the operating system manually. This project aims to analyze the application and customize the operating system automatically. The customization includes tailoring the OS and customizing the implementation of OS primitives.

Dynamically Reconfigurable Computing for Systems on a Chip

Faculty: Prof. Sharad Malik (Princeton)

Researchers: Zhining Huang

This research examines the role of reconfigurable logic in systems on a chip design. Specifically we will study the overhead of storing and downloading the configuration code bits for different parts of an application in a dynamically reconfigurable coprocessor environment. This will provide insight into the granularity of the reconfigurable logic that is appropriate for this application. Our initial study will involve the examination of multimedia and communication systems. These will be profiled within the MESCAL compiler infrastructure (this is built on the Trimaran infrastructure). The ultimate goal of this project is to derive an algorithm that optimally selects the logic macros for use in the reconfigurable logic and compiles the reconfiguration state as part of the MESCAL compiler.

Integrating DSP Compiler Optimization Techniques in a VLIW Compiler

Faculty: Prof. Sharad Malik (Princeton University)

Researchers: Subramanian Rajagopalan

Existing compilation techniques for General Purpose Processors (GPPs) lack the ability to take advantage of the specialized features of Application Specific Processors (ASPs). Digital Signal Processors (DSPs), a certain type of ASP, provide low cost and low power solutions for computationally intensive applications. In this project, we are currently developing a compiler for DSPs by modifying an existing GPP compiler, namely the Trimaran compiler, which includes a vast library of optimizations and supports most of the advanced architectural features of a modern GPP. The Trimaran compiler also satisfies one of the key requirements of a DSP compiler, namely retargetability. The ultimate goal of this project is to unify the GPP and ASP compilation techniques into one common retargetable framework. This framework can then be used to



develop compilers for future GPPs that may include some of the features of ASPs to enhance the performance of some of the common applications.

MESCAL Architecture

Faculty: Prof. Kurt Keutzer (UC-Berkeley), Sharad Malik (Princeton)

Researcher: Niraj Shah, Michael Shilman, Ashok Sudarsanam, Scott Weber

We are developing an architecture template for the exploration and development of MESCAL architectures. Initial and future architecture and micro-architecture features are presented. Based on the attributes of a given application, an architect configures the template to create an application specific programmable solution.

Mescal Compiler Project

Faculty: Prof. Kurt Keutzer, Richard Newton (UC-Berkeley); Sharad Malik (Princeton)

Researchers: Ashok Sudarsanam, Niraj Shah, Scott Weber, Michael Shilman, Manish Vachharajani, Subbu Rajagopalan, Zhining Huang

The objective of the Mescal compiler group is to develop a high-quality retargetable compiler infrastructure that enables a set of interesting source applications to be efficiently mapped onto a family of target architectures/microarchitectures. The target architectures will be composed of multiple processing elements (PEs) and specialized hardware that are integrated onto a single IC. Furthermore, each PE may exhibit a high degree of instruction-level parallelism (ILP) via multiple functional units (FUs) that can execute in parallel. After a thorough evaluation of existing compiler infrastructures, we have decided to use the Trimaran 2.0 compiler as our base infrastructure since it features, among other things, strong VLIW data structure and algorithm support that enables it to efficiently exploit the multiple FUs within each PE.

One of the primary compiler research issues that we are tackling is how to effectively exploit the multiple PEs and specialized FUs of the target architecture by using hints provided by the programmer via the programmer's model. For instance, our hope is that by having the programmer specify the various concurrent tasks in the application via the programmer's model, the compiler can be much more effective at exploiting process-level parallelism and task-level parallelism across the multiple PEs. In our poster session, we will describe the work that we have been doing towards developing this high-quality compiler.

Mescal Design Driver

Faculty: Profs. Kurt Keutzer (UC-Berkeley), Sharad Malik (Princeton)

Researchers: Niraj Shah, Michael Shilman, Ashok Sudarsanam, Scott Weber

The goal of the Mescal project is to provide a programmer's model and software development environment that allows for efficient embedded system implementation onto a family of fully-programmable architectures/microarchitectures. In order to validate our approach, we are developing real-world design driver applications that represent interesting classes of next-generation embedded systems. This poster outlines some of the systems that we have evaluated, including music synthesis, wireless radios, and network routing/security. We compare the computational characteristics of these applications relative to our design system capabilities and motivate our first application focus, Virtual Private Networks.



MESCAL Programmer's Model

Faculty: Profs. Kurt Keutzer, Richard Newton (UC-Berkeley); Sharad Malik (Princeton)

Researchers: Niraj Shah, Michael Shilman, Ashok Sudarsanam, Scott Weber

We are developing an abstraction of the underlying MESCAL architecture that allows the programmer to express his/her knowledge about the system, including parallelism at different levels (process, thread, instruction, and bit), appropriate bindings of processes, and opportunities for optimizations.

Network Security Technologies and their Hardware/Software Implications

Faculty: Prof. Kurt Keutzer (UC-Berkeley)

Researchers: Mel Tsai, Scott Weber, Niraj Shah, Michael Shilman

Today's network implementations have an increasing demand for secure transactions and connections. To facilitate the high bandwidth requirements and compatibility issues of these secure connections, network hardware and software design methodologies may have to be changed.

Self-Test for Analog and Mixed-Signal Components in SOC

Faculty: Prof. Kwang-Ting (Tim) Cheng (UC-Santa Barbara)

Researchers: Jiun-Lang Huang

In this project, we address the issues of reusing on-chip digital programmable components and analog-to-digital (AD) and/or digital-to-analog (DA) converters for on-chip test signal generation, measurement and response analysis to support self-testing analog/mixed signal components. The basic ideas are (1) utilizing simple yet relatively high-tolerance analog circuitry for the A/D and/or D/A conversions for on-chip stimuli generation and response acquisition, and (2) using on-chip processor cores for test synthesis and response analysis employing digital signal processing (DSP) techniques. In this work, we propose (1) a full-chip self-test architecture, (2) a BIST scheme for testing on-chip AD and DA converters, and (3) techniques for the characterization and calibration of a 1-bit first-order delta-sigma modulator for accurate on-chip measurement. The techniques are verified with numerical simulations, and we are conducting hardware experiments for further validation.

Self-Testing of Embedded Processors and System-on-Chips

Faculty: Prof. Sujit Dey

Researchers: Li Chen, Krishna Sekar

Conventional test techniques, involving the use of external testers, will be increasingly difficult as the speed of processors and system chips approaches the GHz range. An alternative to external testing is enabling the processor (or system chip) to test itself. Current hardware-based BIST techniques usually have unacceptably low fault coverage or prohibitively high costs. In this project, we are developing novel embedded software-hardware solutions to enable self-test and self-diagnosis at minimal costs.

A processor self-test technique that we have developed is based on embedding a self-test program that tests each component of the processor. By targeting the structural test needs of manageable components with the aid of processor functionality (instructions), this technique has the high fault

coverage advantage of deterministic structural testing and the at-speed advantage of functional testing. Most importantly, by relieving testers from test application, it enables at-speed testing of GHz processors with low speed testers. Applying the new self-test methodology on a processor core, we have demonstrated the feasibility of the technique to apply high-quality at-speed tests with no test overhead.

In the future, we will explore other hardware-software self-test and self-diagnosis techniques for more complex processor architectures, as well as system-on-chip self-testing using the programmability of the embedded processor cores.

Self-testing the Path Delay Faults in Embedded Processor Cores

Faculty: Prof. Kwang-Ting (Tim) Cheng (UCSB)

Researchers: Wei-Cheng Lai and Angela Krstic (UCSB)

We address the problem of testing path delay faults in a microprocessor using instructions. The focuses of the project are on identifying paths that are testable by using the instructions and on automatically generating the instruction sequence to detect testable paths. Such a test program can then be loaded into the memory in the system-on-chip for self-testing the processor core at speed.

It is observed that many structurally testable paths (i.e., paths testable through at-speed scan) in a microprocessor are not testable by its instructions. These paths do not need to be targeted for path delay testing. We propose an efficient method to identify those untestable paths. Our method uses the microprocessor's RTL description and extracts the spatial and temporal correlations among registers and flip-flops in the microprocessor. We use the extracted constraints to identify the untestable paths. Identification of such paths helps determine the achievable path delay fault coverage and reduce the subsequent test generation effort. The experimental results for two microprocessors indicate that a significant percentage of structurally paths need not be tested.

We further propose an instruction-level, constrained ATPG process for generating tests for path delay faults. During the test generation process, our ATPG attempts to sensitize the path as well as to meet extracted constraints. Therefore, the vector pairs generated are mostly realizable by instructions.

Project Summaries for Theme 4: Communication/Component-Based Design

Approximate Model Checking

Faculty: Prof. David Dill

Researchers: Shankar Govindaraju

Our goal is to increase the size of designs that can be model checked by trading off accuracy for size. Sacrificing accuracy is often acceptable because most correctness properties do not depend on all the details of the reachable state space of a system. The technique uses an over-approximation of a state set based on projections onto a collection of subsets of the state variables.



Bus Interface specification

Faculty: Prof. David Dill

Researchers: Kanna Shimizu

Bus specifications are currently informal, resulting in ambiguities and inconsistencies. We've been working on specifying busses using monitors, which could be written in synthesizable RTL. We have found a good style for writing specifications and debugging them using a CTL model checker. We wrote and debugged a specification for a subset of PCI using this method.

Checking Specifications for Consistency

Faculty: Profs. Randal E. Bryant and Edmund Clarke

Researchers: Amit Goel, Dong Wang, Pankajkumar Chauhan, Yuan Lu

Complicated specifications, such as those for bus architectures, are currently written in natural language such as English. These specifications are often error-prone and ambiguous leading to inconsistent specifications and varying interpretations.

We present a methodology to check temporal logic specifications for consistency. We require:

- * Satisfiability of the temporal formulae
- * Forward Progress possible in all reachable states (no dead states)
- * Modules should be receptive to actions controlled by their environment.

The first two requirements are easily checked by using CTL model checking on a maximal model for the properties. However, to check for receptivity, we need to realize that there are dependencies between events within a single synchronous step and that these dependencies affect the evaluation for receptiveness of modules. These dependencies lead to a natural division of the synchronous step into 'micro-steps'; each micro-step involves a set of variables that can be evaluated taking into account the dependencies. Receptivity requires that at each micro-step, the system should be able to make forward progress no-matter what the evaluation for other variables in previous micro-steps. Our algorithm is based on simple manipulations of BDD's for checking consistency.

Distributed Data Management for EDA

Faculty: Prof. Richard Newton (UC-Berkeley)

Researchers: Mark Spiller

The increasing sizes of modern chip designs as well as their design flow complexities have severely impacted the efficient management of design data. Incremental design techniques attempt to improve efficiency by limiting re-computation to only the areas of the design effected by the changes, thus saving the processing time across design iterations. However, the success of incremental design depends upon tight, integrated APIs that allow efficient data handling between the tools in the design flow. Today, most companies pull together third-party tool suites using file exchange for the hand-off points, for which incremental design is much more difficult.

While such improvisation and lack of clear standards among design tools cause heterogeneity in data interfaces, incremental design requires smooth, efficient data flow. We are working to identify the needs and requirements for an architecture that will provide integration between heterogeneous tools and data repositories while providing efficient support for collaborative

incremental design. Areas of interest include relaxed transaction models suited towards collaborative work [1] and intelligent, distributed caching of design data.

To this end, and as part of the GSRC collaborative design activity, we are reviewing existing industry standards for data representation (e.g. CHDStd[2]) as a vehicle for data management in a heterogeneous and widely distributed environment. We are gathering tool usage patterns and design set data throughout the design process, to gain an accurate model of design data size and usage trends. Results are being incorporated into prototyping efforts that use a platform-independent architecture (Windchill [3]) to analyze the performance trade-offs of different data management techniques. Our research goal is to find reliable design data management and transactional semantics (i.e. an architecture) for large data sets in a distributed, heterogeneous, and potentially unreliable network environment.

[1] A. Fox, et al., "Extensible Cluster-Based Scalable Network Services,"
Proceedings of the 16th ACM Symposium on Operating Systems Principles
(SOSP-16), St. Malo, France, October 1997.

[2] The Chip Hierarchical Design System Technical Data Standard (CHDStd),
<http://www.si2.org/CHDStd/>

[3] Windchill, from Parametric Technology, Inc. <http://www.ptc.com/windchill>

The Intercom Design Case Study

Faculty: Prof. Jan Rabaey, Alberto Sangiovanni-Vincentelli (UC-Berkeley)

Researchers: Josie Ammer, Per Bjureus, Fred Burghardt, Fernando De Bernardinis, Suet Fei Li, Sue Mellers, Vandana Prabhu, Marco Sgroi, Mike Sheets, Julio da Silva, and Arvind Thirunarayanan

A case study for the application of the design methodology for protocols based on CFSMs is presented in this poster. The application is an intercom system capable of connecting up to 32 users using a TDMA access scheme. The protocol stack has been appropriately layered and each layer has been specified using CFSMs. The tool used for the implementation of the CFSM machine is VCC from Cadence. VCC design flow is very suitable to the proposed design methodology and allows a clear functional simulation and a performance simulation. The results of both phases are shown and the trade-offs involved with different mappings/architectures are estimated. A final architecture has finally been selected for the actual implementation of the design.

JavaTime Embedded Systems Development Environment

Faculty: Prof. Richard Newton, Paul Hilfinger (UC-Berkeley)

Researchers: James Shin Young, Josh MacDonald, Michael Shilman

JavaTime is a software environment for embedded systems design. It is based on the use of a common model for systems representation, called the Abstract Reactive model, into which heterogeneous domains may be embedded. The domains represent different models of computation, such as finite state machines, discrete event, dataflow, and synchronous models, and allow designers to specify systems in a domain-specific manner. The domains are implemented as packages in the Java language, with API's similar to the 'native' syntax of the models.



Latency Insensitive Design Methodology

Faculty: Prof. Alberto Sangiovanni-Vincentelli (UC-Berkeley)

Researchers: Luca Carloni, Kenneth McMillan and Alexander Saldanha (Cadence Berkeley Laboratories)

In Deep Sub-Micron (DSM) designs, performance will depend critically on the latency of long wires. We propose a new synthesis methodology for synchronous systems that makes the design functionally insensitive to the latency of long wires. Given a synchronous specification of a design, we generate a functionally equivalent synchronous implementation that can tolerate arbitrary communication latency between latches. By using latches we can break a long wire in short segments which can be traversed while meeting a single clock cycle constraint. The overall goal is to obtain a design that is robust with respect to delays of long wires, in a shorter time by reducing the multiple iterations between logical and physical design, and with performance that is optimized with respect to the speed of the single components of the design. As a case study, we present the latency insensitive design of PDLX, an out-of-order microprocessor with speculative-execution.

A Methodology for Protocol Design

Faculty: Prof. Jan Rabaey, Alberto Sangiovanni-Vincentelli (UC-Berkeley)

Researchers: Fernando De Bernardinis, Marco Sgroi, Julio Silva

The goal is to define a methodology to design protocols that are functionally correct and efficient, while satisfying performance requirements and using a minimal amount of physical resources. The methodology we are developing is based on the notion of communication refinement [Rowson-Sangiovanni, DAC97]. In our approach the communication is refined from a high-level description to the final implementation by applying a sequence of refinement steps, such that each transformation preserves the original behavior and communication constraints are propagated in a top-down fashion to lower levels of abstraction. To make our methodology of practical use, we are developing a design environment to support protocol designers throughout the different phases of the design process and we are evaluating the methodology on a real case study (Intercom wireless system). This will require, among the other things, to introduce and experiment ECFSMs (Extended CFSMs), a formal model of computation for mixed control/dataflow systems, and develop tools for validation and architecture evaluation.

Partial model checking of embedded systems

Faculty: Prof. David Dill (Stanford)

Researchers: David Park and Ulrich Stern

We're building a checker for anomalies in multi-threaded Java programs. The system includes a translator from Java to an intermediate representation ("SAL"), and a model checker for SAL that can model program states that have dynamically allocated memory, data structures that vary in size, dynamically varying numbers of threads, and other characteristics of software that are not well supported by existing model checkers.

The Ptolemy Project

Faculty: Prof. Edward A. Lee (UC-Berkeley)

Researchers: John Davis, II, Christopher Hylands, Bart Kienhuis, Kees Vissers, Bilung Lee, Jie Liu, Xiaojun Liu, Steve Neuendorffer, John Reekie, Mary P. Stewart, Jeff Tsay, Yuhong Xiong

Ptolemy II is a set of Java packages supporting heterogeneous, concurrent modeling, simulation, and design of component-based systems. The emphasis is on a clean, modular software architecture, divided into a set of coherent, comprehensible packages. The kernel package supports definition and manipulation of clustered hierarchical graphs, which are collections of entities and relations between those entities. The actor package extends the kernel so that entities have functionality and can communicate via the relations. The domains extend the actor package by imposing models of computation on the interaction between entities. Various domains are already implemented like, for example, Communicating Sequential Processes, Continuous Time, Discrete Events, Finite State Machines, Process Networks, and Synchronous Dataflow.

Semi-formal verification using quasi-symbolic simulation

Faculty: Prof. David Dill (Stanford)

Researchers: Chris Wilson

Our goal is to provide verification that is much more efficient than conventional simulation, but more scalable than formal verification. Our approach is based on a form of approximate symbolic simulation, coupled with a SAT algorithm that can selectively do more detailed simulation.

Specification and Verification of Bus Protocols

Faculty: Prof. Kareem Sakallah (U. Michigan)

Researchers: Fadi Aloul

We are exploring a taxonomy of busses based on a formal specification of bus protocols. To identify the essence of what makes a communication conduit a "bus", we plan to develop a generic template for specifying busses that is suitable for describing commonly-used bus architectures. Availability of such template will also enable the development of efficient bus-specific verification algorithms. We are currently casting bus protocol verification as a Boolean SAT problem. We also plan to explore the use of mathematical induction to verify bus systems consisting of an arbitrary number of devices.

System-Level Power Management

Faculty: Prof. Giovanni De Micheli (Stanford)

Researchers: Yung-Hsiang Lu

Power management shuts down unused devices to save power. In this study, we build a Windows-NT based environment to implement, evaluate, and compare power management algorithms and point out directions for future improvement.



Temporal Logic Specifications for PCI Bus Protocol

Faculty: Prof. Edmund Clarke (Carnegie Mellon U.)

Researchers: Pankaj Chauhan, Yuan Lu, Dong Wang

Bus protocols are currently being specified in the English language, which is informal, often ambiguous, error-prone and non-executable. We propose a specification language based on temporal logic, and a methodology to visualize and model check the specifications. We translate the English protocol specifications to temporal logic properties. For each temporal logic formula, a state machine is built using only non-data signals. Simulation is done on the whole system composed of individual state machines, which generates signal timing diagrams for visualization. Then we abstract the data part in the protocol to obtain abstracted FSM models and analyze high level properties such as deadlock, livelock, receptivity etc. using model checking. We also investigate ATPG techniques to achieve more efficient checking methods. We demonstrate the above methodology on the PCI bus protocol. We give the specifications in our language for a subset of the PCI bus protocol and show some initial results.

Verification of an arbiter for the IBM CoreConnect Architecture

Faculty: Prof. Randal E. Bryant (Carnegie Mellon University)

Researchers: Amit Goel

Advances in silicon densities now allow many functions to be integrated onto a complex system-on-a-chip (SOC) design, which typically consists of a collection of cores interconnected by an internal bus architecture. The verification and specification of these buses poses a special challenge. We seek to better understand the difficulties associated with this task and what is required to make it easier.

We present the model checking effort for an arbiter core for the IBM CoreConnect Architecture and discuss our verification methodology, describing how it was influenced by the CoreConnect Architecture. We also present the bugs found and discuss the difficulties associated with verifying complex on-chip buses. Our findings highlight the need for better tools and methodologies for the specification of complex architectures and their verification. The insight gained from this work has motivated our work on Checking Specifications for Consistency.

Verifying IP Core based system-on-chip designs

Faculty: Prof. Edmund Clarke (Carnegie Mellon University)

Researchers: Pankaj Chauhan, Yuan Lu, Dong Wang

Verification of system-on-chip designs is becoming more and more challenging, especially in the light of heterogeneity of components involved in terms of protocol, functional and timing requirements. We describe a methodology for verifying such designs. The verification is decomposed into three tasks. First, we verify, once and for all, the standard bus interconnecting IP Cores in the system. The next task is to verify the glue logic, which connects the IP Cores to the buses. Finally, using the verified bus protocols and the IP Core designs, the complete system is verified. Various techniques including abstraction, assume-guarantee reasoning, induction and symmetry are required in order to contain the size of the problem. To illustrate our methodology, we verify the PCI Local Bus, a widely used bus protocol in system-on-chip designs. We demonstrate various modeling and verification techniques for buses by modeling the PCI bus with SMV. We have found two potential bugs in the PCI bus specifications.



Verifying Sequential Consistency in Cache Coherence Protocols

Faculty: Prof. Thomas A. Henzinger (UC-Berkeley)

Researchers: Ranjit Jhala, Rupak Majumdar, Shaz Qadeer

Shared memory multiprocessors are an important class of high performance computing systems. The design of a correct and efficient shared memory is one of the most difficult tasks in the design of such systems. Shared memory systems are typically studied as abstract models, parameterized by the number of processors, the number of memory locations, the number of data values and some criterion for correctness. Such a correctness criterion may be serial memory where the memory system behaves as if there is a centralized memory that services read and write requests atomically, such that a read to a location returns the latest value written to it. As this model may be too stringent, several other models like sequential consistency [Lamport79] have been proposed, where the global temporal order of memory requests are ignored and it is only required that some interleaving of the local temporal order of events at different processors be a trace of serial memory. Our goal is to verify sequential consistency for an arbitrary number of processors, memory locations and data values for cache coherence protocols in commercial multiprocessor systems by automated model checking techniques. Since model checking cannot be directly applied to such a parameterized system, we hope to develop compositional techniques. In particular, we shall model the cache coherence protocol used in the HP Runway bus and check if it is sequentially consistent.

5. Student Information

Carnegie Mellon University

Last Name: Goel	First Name: Amit
Work Address: 2132 Hamerschlag Hall, 5000 Forbes Avenue	
City: Pittsburgh	State: PA Zip: 15323
Work Phone: (412) 268-1112	Home Phone: (412) 361-3389
Email: agoel@ece.cmu.edu	URL: http://www.ece.cmu.edu/~agoel
Thesis Title: <i>Verifying Consistency of Protocol Specifications</i>	
Expected Graduation: 5/1/2000	Advisor: Bryant

Last Name: Chauhan	First Name: Pankajkumar
Work Address: Wean Hall, 5000 Forbes Avenue	
City: Pittsburgh	State: PA Zip: 15323
Work Phone: (412) 268-3564	Home Phone:
Email: pchauhan@cs.cmu.edu	URL: http://www.cs.cmu.edu/~pchauhan
Thesis Title: <i>Undecided</i>	
Expected Graduation: 5/1/2002	Advisor: Clarke

Last Name: Wang	First Name: Dong
Work Address: Wean Hall, 5000 Forbes Avenue	
City: Pittsburgh	State: PA Zip: 15323
Work Phone: (412) 268-3625	Home Phone: (412) 681-5130
Email: dongw@cs.cmu.edu	URL: http://www.cs.cmu.edu/~dongw
Thesis Title: <i>Formal Specification and Verification of Bus Protocols</i>	
Expected Graduation: 5/1/2002	Advisor: Clarke

Last Name: Deng	First Name: Yangdon
Work Address: 2131 Hamerschlag Hall, 5000 Forbes Avenue	
City: Pittsburgh	State: PA Zip: 15323
Work Phone: (412) 268-6648	Home Phone: (412) 605-0225
Email: yangdon@ece.cmu.edu	URL: http://www.ece.cmu.edu/~yangdon
Thesis Title: <i>undecided</i>	
Expected Graduation: 5/1/2002	Advisor: Maly

Last Name: Palusinski	First Name: Michal
Work Address: 2138 Hamerschlag Hall, 5000 Forbes Avenue	
City: Pittsburgh	State: Pennsylvania Zip: 15213
Work Phone: (412) 268-6641	Home Phone:
Email: michalp@cyrus.andrew.cmu.edu	URL:
Thesis Title: <i>Undecided</i>	
Expected Graduation: 5/1/2003	Advisor: Maly



Carnegie Mellon University

Last Name: Acar
Work Address: 5803 Fifth Avenue,
City: Pittsburgh
Work Phone: (412) 268-4276
Email: emrah@ece.cmu.edu
Thesis Title: *Undecided*
Expected Graduation: 5/1/2002

First Name: Emrah
State: Pennsylvania **Zip:** 15232
Home Phone: (412) 361-1359
URL:
Advisor: Pileggi

Last Name: Arunachalam
Work Address: 2119 Hamerschlag Hall, 5000 Forbes Avenue
City: Pittsburgh
Work Phone: (412) 268-4276
Email: ravia@ece.cmu.edu
Thesis Title: *Undecided*
Expected Graduation: 5/1/2001

First Name: Ravi
State: Pennsylvania **Zip:** 15232
Home Phone: (412) 687-7508
URL: <http://www.ece.cmu.edu/~ravia>
Advisor: Pileggi

Last Name: Gupta
Work Address: 2136 Hamerschlag Hall, 5000 Forbes Avenue
City: Pittsburgh
Work Phone: (412) 268-3096
Email: satrajit@andrew.cmu.edu
Thesis Title: *Hierarchical Abstraction of RLC Interconnect Circuits*
Expected Graduation: 5/1/2003

First Name: Satrajit
State: Pennsylvania **Zip:** 15232
Home Phone:
URL:
Advisor: Pileggi

Last Name: Koorapaty
Work Address: 2137 Hamerschlag Hall, 5000 Forbes Avenue
City: Pittsburgh
Work Phone: (412) 268-6645
Email: aneeshk@ece.cmu.edu
Thesis Title: *Undecided*
Expected Graduation: 5/1/2003

First Name: Aneesh
State: Pennsylvania **Zip:** 15232
Home Phone: (412) 681-1016
URL: <http://www.ece.cmu.edu/~aneesh>
Advisor: Pileggi



Princeton University

Last Name: Huang	First Name: Zhining
Work Address: Electrical Engineering, CE Group, E-Quad B217	
City: Princeton	State: New Jersey Zip: 08554
Work Phone: (609) 258-1354	Home Phone:
Email: znhuang@EE.Princeton.EDU	URL: http://www.ee.princeton.edu/~znhuang
Thesis Title: <i>Dynamically Reconfigurable Computing for Systems on a Chip</i>	
Expected Graduation: 6/1/2002	Advisor: Malik

Last Name: Vachharajani	First Name: Manish
Work Address: Department of Electrical Engineering, E-Quad	
City: Princeton	State: New Jersey Zip: 08544-5263
Work Phone: (609) 258-7379	Home Phone:
Email: manishv@ee.princeton.edu	URL:
Thesis Title: <i>Undecided</i>	
Expected Graduation: 6/1/2002	Advisor: Malik

Last Name: Wang	First Name: Shaojie
Work Address: Department of Electrical Engineering, CE Group	
City: Princeton	State: New Jersey Zip: 08544
Work Phone: (609) 258-1354	Home Phone:
Email: wsj@ee.princeton.edu	URL: http://www.ee.princeton.edu/~wsj
Thesis Title: <i>Application Specific Embedded Software Synthesis</i>	
Expected Graduation: 6/1/2002	Advisor: Malik



Purdue University

Last Name: Choi	First Name: Seung Hoon
Work Address: 1285 Electrical Engineering Building West, Box 393	
City: W. Lafayette	State: Indiana Zip: 47907-1285
Work Phone: (765) 494-0759	Home Phone: (765) 495-6522
Email: choi5@ecn.purdue.edu	URL: http://www.min.ecn.purdue.edu/~choi5
Thesis Title: <i>Undecided</i>	
Expected Graduation: 8/1/2002	Advisor: Roy

Stanford University

Last Name: Govindaraju	First Name: Shankar
Work Address: Gates Computer Science Building, Room 322,	
City: Stanford	State: California Zip: 94305
Work Phone: (650) 725-1775	Home Phone: (650) 491-4048
Email: shankar@leland.stanford.edu	URL: http://www.stanford.edu/~shankar
Thesis Title: <i>Approximate Symbolic Model Checking Using Overlapping Projections</i>	
Expected Graduation: 6/1/2000	Advisor: Dill

Last Name: Shimizu	First Name: Kanna
Work Address: Department of Computer Science, Gates Computer Science Building, Room 352	
City: Stanford	State: California Zip: 94305
Work Phone: (650) 725-1683	Home Phone: (650) 497-7571
Email: kannas@stanford.edu	URL: http://www.stanford.edu/~kannas
Thesis Title: <i>Undecided</i>	
Expected Graduation: 6/1/2002	Advisor: Dill

Last Name: Su	First Name: Xiao-Wu
Work Address: Gates Building, Room 226,	
City: Stanford	State: California Zip: 94305
Work Phone: (650) 725-3656	Home Phone: (408) 733-5727
Email: xsu@cs.stanford.edu	URL: http://www.sprout.stanford.edu
Thesis Title: <i>Undecided</i>	
Expected Graduation: 6/1/2000	Advisor: Dill

Last Name: Wilson	First Name: James
Work Address: Gates Computer Science Building, Room 312,	
City: Stanford	State: California Zip: 94305
Work Phone: (650) 725-3656	Home Phone: (408) 353-5058
Email: chriswi@leland.stanford.edu	URL:
Thesis Title: <i>Semi-Formal Verification Methods</i>	
Expected Graduation: 6/1/2001	Advisor: Dill

Last Name: Lu	First Name: Yung-Hsiang
Work Address: Department of Computer Science, Gates Rm. 3A-502, 353 Serra Mall	
City: Stanford	State: California Zip: 94305-9030
Work Phone: (650) 725-3924	Home Phone: (650) 566-9122
Email: luyung@stanford.edu	URL: http://www.stanford.edu/~luyung
Thesis Title: <i>Dynamic Power Management</i>	
Expected Graduation: 12/1/2001	Advisor: De Micheli



Stanford University

Last Name: Peymandoust	First Name: Armita
Work Address: Dept. of Computer Science, Gates Computer Science Rm. 3A-330, 353 Serra Mall	
City: Stanford	State: California Zip: 94305-9030
Work Phone: (650) 723-9035	Home Phone: (408) 531-9099
Email: armita@leland.stanford.edu	URL:
Thesis Title: <i>Undecided</i>	
Expected Graduation: 6/1/2002	Advisor: De Micheli

Last Name: Smith	First Name: James
Work Address: Gates Building, Room 326,	
City: Stanford	State: California Zip: 94035
Work Phone: (650) 725-3647	Home Phone: (650) 965-2039
Email: redskins@aglaia.stanford.edu	URL:
Thesis Title: <i>Allocation and Infoface Synthesis Algorithms for Design Reuse</i>	
Expected Graduation: 10/1/1999	Advisor: De Micheli

University of California, Berkeley

Last Name: Shilman	First Name: Michael
Work Address: Electrical Engineering and Computer Sciences, 211-137 Cory Hall	
City: Berkeley	State: California Zip: 94720-1772
Work Phone: (510) 642-5048	Home Phone: (510) 527-8480
Email: michael@eecs.berkeley.edu	URL: http://www.eecs.berkeley.edu/~michaels
Thesis Title: <i>Communications Visualization for Embedded Systems Design</i>	
Expected Graduation: 6/1/2001	Advisor: Newton

Last Name: Spiller	First Name: Mark
Work Address: Electrical Engineering and Computer Sciences, 146 Cory Hall	
City: Berkeley	State: California Zip: 94720-1770
Work Phone: (510) 642-3433	Home Phone:
Email: mds@eecs.berkeley.edu	URL: http://www-cad.eecs.berkeley.edu/~mds
Thesis Title: <i>Distributed Design Data Management for EDA</i>	
Expected Graduation: 12/1/2000	Advisor: Newton

Last Name: Tabbara	First Name: Abdallah
Work Address: Electrical Engineering and Computer Sciences, 211-150 Cory Hall	
City: Berkeley	State: California Zip: 94720-1770
Work Phone: (510) 643-5187	Home Phone: (510) 843-7987
Email: atabbara@eecs.berkeley.edu	URL: http://www-cad.eecs.berkeley.edu/~atabbara
Thesis Title: <i>Synthesis and Physical Design Methodologies for DSM</i>	
Expected Graduation: 6/ 1/2001	Advisor: Newton

Last Name: Young	First Name: James
Work Address: Electrical Engineering and Computer Sciences, 550-G Cory Hall	
City: Berkeley	State: California Zip: 94720
Work Phone: (510) 642-3652	Home Phone:
Email: jimy@eecs.berkeley.edu	URL: http://www-cad.eecs.berkeley.edu/~jimy
Thesis Title: <i>A Common Model Approach to Heterogeneous Specification of Embedded Systems</i>	
Expected Graduation: 12/ 1/2000	Advisor: Newton

Last Name: Chinnery	First Name: David
Work Address: Electrical Engineering and Computer Sciences, 211-23 Cory Hall	
City: Berkeley	State: California Zip: 94720-1770
Work Phone: (510) 642-5322	Home Phone: (510) 654-2776
Email: chinnery@eecs.berkeley.edu	URL: http://wwwcad.eecs.berkeley.edu/~chinnery
Thesis Title: <i>Investigation of Improving Standard Cell Design Towards Custom</i>	
Expected Graduation: 12/1/2003	Advisor: Keutzer

University of California, Berkeley

Last Name: Shah First Name: Niraj Work Address: 211-80 Cory Hall, #1770, City: Berkeley Work Phone: (510) 642-1459 Email: niraj@eecs.berkeley.edu Thesis Title: <i>Undecided</i> Expected Graduation: 5/ 1/2003		State: California Zip: 94720-1770 Home Phone: (510) 665-4533 URL: http://www-cad.eecs.berkeley.edu/~niraj/ Advisor: Keutzer
Last Name: Weber Work Address: 211-120 Cory Hall, #1772, City: Berkeley Work Phone: (510) 642-5322 Email: sjweber@eecs.berkeley.edu Thesis Title: <i>Undecided</i> Expected Graduation: 12/1/2003		First Name: Scott State: California Zip: 94720 Home Phone: (510) 704-0656 URL: Advisor: Keutzer
Last Name: Chong Work Address: U. C. Berkeley, 211-42 Cory Hall City: Berkeley Work Phone: (510) 643-5135 Email: pchong@eecs.berkeley.edu Thesis Title: <i>Undecided</i> Expected Graduation: 5/1/2002		First Name: Philip State: California Zip: 94720-1770 Home Phone: (510) 664-2802 URL: http://www/cad.eecs.berkeley.edu/~pchong Advisor: Brayton
Last Name: Jiang Work Address: Electrical Engineering and Computer Sciences, 211-176 Cory Hall City: Berkeley Work Phone: (510) 642-4332 Email: wjiang@eecs.berkeley.edu Thesis Title: <i>Undecided</i> Expected Graduation: 5/1/2002		First Name: Yunjian State: California Zip: 94720-1772 Home Phone: (510) 540-0324 URL: http://www.cad.eecs.berkeley.edu/~wjiang Advisor: Brayton
Last Name: Khatri Work Address: Electrical Engineering and Computer Sciences, 550-B 1772 Cory Hall City: Berkeley Work Phone: (510) 642-1662 Email: linus@eecs.berkeley.edu Thesis Title: <i>Cross-talk Noise Immune VLSI Design Using Regular Layout Fabrics</i> Expected Graduation: 12/ 1/1999		First Name: Sunil State: California Zip: 94720 Home Phone: (510) 486-1009 URL: http://www-cad.eecs.berkeley.edu/~linus Advisor: Brayton

University of California, Berkeley

Last Name: Tsai
Work Address: 2299 Piedmont Avenue, #608
City: Berkeley
Work Phone: (510) 642-5322
Email: tsaimelv@pilot.msu.edu
Thesis Title: *Undecided*
Expected Graduation: 12/1/2003

First Name: Mel
State: California **Zip:** 94720
Home Phone: (510) 664-2905
URL:
Advisor: Brayton

Last Name: Guo
Work Address: 1609 Spruce Street
City: Berkeley
Work Phone: (510) 666-3137
Email: along@eecs.berkeley.edu
Thesis Title: *Undecided*
Expected Graduation: 5/1/2002

First Name: Chynlong
State: California **Zip:** 94709
Home Phone: (510) 660-0393
URL: <http://infopad.eecs.berkeley.edu/~along>
Advisor: Rabaey

Last Name: Patel
Work Address: 2108 Allston Way
City: Berkeley
Work Phone: (510) 666-3140
Email: dpatel@eecs.berkeley.edu
Thesis Title: *Low Energy Ad-hoc Patrols*
Expected Graduation: 5/1/2000

First Name: Danny
State: California **Zip:** 94704
Home Phone: (510) 486-9848
URL:
Advisor: Rabaey

Last Name: De Bernardinis
Work Address: 231 Cory Hall #1772,
City: Berkeley
Work Phone: (510) 643-7039
Email: fdb@eecs.berkeley.edu
Thesis Title: *Undecided*
Expected Graduation: 5/1/2003

First Name: Fernando
State: California **Zip:** 94720
Home Phone:
URL: N/A
Advisor: Sangiovanni-Vincentelli

Last Name: Martin
Work Address: 550 Cory Hall, Mail Box 64 - #1770,
City: Berkeley
Work Phone: (510) 642-3652
Email: jmartin@ic.eecs.berkeley.edu
Thesis Title: *Undecided*
Expected Graduation: 12/1/2002

First Name: Jonathan
State: California **Zip:** 94720
Home Phone: (925) 458-8875
URL: <http://www-ad.eecs.berkeley.edu/~jmartin>
Advisor: Sangiovanni-Vincentelli



University of California, Berkeley

Last Name: Passerone	First Name: Roberto
Work Address: Electrical Engineering and Computer Sciences, 211-66 Cory Hall	
City: Berkeley	State: California Zip: 94720-1770
Work Phone: (510) 643-7039	Home Phone:
Email: roby@eecs.berkeley.edu	URL: http://www.eecs.berkeley.edu/~roby
Thesis Title: <i>Undecided</i>	
Expected Graduation: 5/1/2001	Advisor: Sangiovanni-Vincentelli

Last Name: Pinello	First Name: Claudio
Work Address: 550 Cory Hall, Mail Box 67 - 1770,	
City: Berkeley	State: California Zip: 94720
Work Phone: (510) 642-3433	Home Phone:
Email: pinello@eecs.berkeley.edu	URL: http://www.eecs.berkeley.edu/~roby
Thesis Title: <i>Hybrid Systems in Automotive Control Applications</i>	
Expected Graduation: 5/1/2003	Advisor: Sangiovanni-Vincentelli

Last Name: Chen	First Name: Rong
Work Address: 550 Cory Hall 1770	
City: Berkeley	State: California Zip: 94720
Work Phone: (510) 642-5322	Home Phone: (510) 649-9879
Email: rongchen@eecs.berkeley.edu	URL: http://www.eecs.berkeley.edu/~rongchen
Thesis Title: <i>Undecided</i>	
Expected Graduation: 5/1/2003	Advisor: Sangiovanni-Vincentelli

University of California, Los Angeles

Last Name: Lu	First Name: Hua
Work Address: Computer Science Department, EE	Department: 5th Floor, IV
City: Los Angeles	State: California Zip: 90024
Work Phone:	Home Phone: (213) 481-2035
Email: hua@ee.ucla.edu	URL:
Thesis Title: <i>YLIN Code Compression for TICBX</i>	
Expected Graduation: 7/1/2000	Advisor: Mangioni-Smith

Last Name: Lim	First Name: Sung
Work Address: Computer Science Department, 4651 Boelter Hall	
City: Los Angeles	State: California Zip: 90095
Work Phone: (310) 206-2279	Home Phone: (213) 739-1055
Email: limsk@cs.ucla.edu	URL: http://www.cs.ucla.edu/~limsk
Thesis Title: <i>Large Scale Circuit Clustering and Partitioning</i>	
Expected Graduation: 6/1/2000	Advisor: Cong

Last Name: Caldwell	First Name: Andrew
Work Address: Computer Science Department, 3731 Boelter Hall	
City: Los Angeles	State: California Zip: 90095
Work Phone: (310) 825-1802	Home Phone: (310) 839-0456
Email: caldwell@cs.ucla.edu	URL: http://www.visicad.cs.ucla.edu/~caldwell
Thesis Title: <i>Undecided</i>	
Expected Graduation: 7/1/2002	Advisor: Kahng

Last Name: Chen	First Name: Yu
Work Address: Computer Science Department, 3563 Boelter Hall	
City: Los Angeles	State: California Zip: 90095-1596
Work Phone: (310) 825-1872	Home Phone: (310) 208-4392
Email: ychen@cs.ucla.edu	URL: http://www.cs.ucla.edu/~yuchen
Thesis Title: <i>Undecided</i>	
Expected Graduation: 6/1/2002	Advisor: Kahng

Last Name: Koushanfar	First Name: Farinaz
Work Address: Computer Science Department, 3732 Boelter Hall	
City: Los Angeles	State: California Zip: 90095
Work Phone: (310) 825-1872	Home Phone: (310) 209-2891
Email: farinaz@cs.ucla.edu	URL: http://www.visicad.cs.ucla.edu
Thesis Title: <i>CAD Implementation for Substrate Noise in Analog/Digital Circuit</i>	
Expected Graduation: 6/1/2003	Advisor: Kahng



University of California, Los Angeles

Last Name: Mantik	First Name: Stefanus
Work Address: Computer Science Department, 3563C Boelter Hall	
City: Los Angeles	State: California Zip: 90095-1596
Work Phone: (310) 825-1872	Home Phone: (626) 291-5555
Email: stefanus@cs.ucla.edu	URL: http://www.cs.ucla.edu/~stefanus
Thesis Title: <i>METRICS: A System for Continuous Improvement of Design Technology</i>	
Expected Graduation: 6/1/2001	Advisor: Kahng

Last Name: Markov	First Name: Igor
Work Address: Computer Science Department, VLSI CAD Lab, 3732 Boelter Hall	
City: Los Angeles	State: California Zip: 90095-1596
Work Phone: (310) 206-0179	Home Phone: (310) 288-0820
Email: imarkov@cs.ucla.edu	URL: http://visicad.cs.ucla.edu/~imarkov
Thesis Title: <i>Global Objectives for Timing-Driven Placement</i>	
Expected Graduation: 6/1/2000	Advisor: Kahng

Last Name: Shallesh	First Name: Vaya
Work Address: Computer Science Department, 3732 Boelter Hall	
City: Los Angeles	State: California Zip: 90095
Work Phone: (310) 825-1872	Home Phone: (310) 625-0578
Email: vaya@cs.ucla.edu	URL:
Thesis Title: <i>Undecided</i>	
Expected Graduation: 4/1/2003	Advisor: Kahng



University of California, San Diego

Last Name: Chen	First Name: Li
Work Address: ECE-EBU-I/2704, 9500 Gilman Drive	
City: La Jolla	State: California Zip: 92093-0407
Work Phone: (858) 534-7883	Home Phone: (858) 552-1304
Email: lichen@ece.ucsd.edu	URL: http://www.esdat/~lichen
Thesis Title: <i>Determine Functional Self-test for Embedded Processor Cones and System-on Chip</i>	
Expected Graduation: 9/1/2002	Advisor: Dey

Last Name: Sekar	First Name: Krishna
Work Address: ECE-EBU-I/2704, 9500 Gilman Drive	
City: La Jolla	State: California Zip: 92093-0407
Work Phone: (858) 534-7883	Home Phone: (858) 587-1142
Email: ksekar@ece.ucsd.edu	URL: http://aspire.ucsd.edu/~ksekar
Thesis Title: <i>Self Testing of System-on Chips</i>	
Expected Graduation: 5/1/2003	Advisor: Dey

University of California, Santa Barbara

Last Name: Gunar Work Address: ECE Department, UCSB City: Santa Barbara Work Phone: (805) 893-5678 Email: guner@bigbend.ece.ucsb.edu Thesis Title: <i>VLSI Testing</i> Expected Graduation: 12/1/2004	First Name: Ahmet State: California Zip: 93106-9560 Home Phone: URL: Advisor: Cheng
Last Name: Huang Work Address: ECE Department, City: Santa Barbara Work Phone: (805) 893-5678 Email: lang@redwood.ece.ucsb.edu Thesis Title: <i>Fault Diagnosis and Testing for Analog/Mixed-signal Circuits</i> Expected Graduation: 11/1/1999	First Name: Jiun-Lang State: California Zip: 93106 Home Phone: (805) 403-6694 URL: Advisor: Cheng
Last Name: Iyer Work Address: ECE Department, UCSB City: Santa Barbara Work Phone: (805) 893-5678 Email: madiyer@yellowstone.ece.ucsb.edu Thesis Title: <i>Undecided</i> Expected Graduation: 6/1/2003	First Name: Madhu State: California Zip: 93106-9560 Home Phone: (805) 685-6356 URL: Advisor: Cheng
Last Name: Lai Work Address: ECE Department, City: Santa Barbara Work Phone: (805) 893-5678 Email: wlai@windcave.ece.ucsb.edu Thesis Title: <i>System on Chip Testing</i> Expected Graduation: 12/1/2002	First Name: Wei-Cheng State: California Zip: 93106-9560 Home Phone: URL: Advisor: Cheng
Last Name: Hu Work Address: ECE Mailbox #146, UCSB City: Santa Barbara Work Phone: (805) 893-5678 Email: hu@guitar.ece.ucsb.edu Thesis Title: <i>Undecided</i> Expected Graduation: 9/1/2004	First Name: Bo State: California Zip: 93117 Home Phone: (805) 685-1542 URL: Advisor: Marek-Sadowska



University of California, Santa Barbara

Last Name: Wang

Work Address: ECE Mailbox #1156, UCSB

City: Santa Barbara

Work Phone: (805) 893-5678

Email: wzy72@engineering.ucsb.edu

Thesis Title: *Undecided*

Expected Graduation: 7/1/2004

First Name: Zhiyuan

State: California **Zip:** 93117

Home Phone: (805) 571-7506

URL:

Advisor: Marek-Sadowska



University of California, Santa Cruz

Last Name: Morton	First Name: Paul
Work Address: CSE Department, 1156 High Street	
City: Santa Cruz	State: California Zip: 95064
Work Phone: (831) 459-4954	Home Phone: (831) 502-9309
Email: morton@cse.ucsc.edu	URL: http://www.cse.ucsc.edu/~morton
Thesis Title: <i>Undecided</i>	
Expected Graduation: 6/1/2001	Advisor: Dai

Last Name: Wu	First Name: Huaizhi
Work Address: CSE Department, 1156 High Street	
City: Santa Cruz	State: California Zip: 95064
Work Phone: (831) 459-4954	Home Phone: (831) 427-3058
Email: huaizhi@cse.ucsc.edu	URL: http://www.cse.ucsc.edu/~huaizhi
Thesis Title: <i>System-Level Interconnect Planning for Deep Submicron IC Design</i>	
Expected Graduation: 7/1/2002	Advisor: Dai

Last Name: Zhang	First Name: Shuo
Work Address: CSE Department, 1156 High Street	
City: Santa Cruz	State: California Zip: 95064
Work Phone: (831) 459-4954	Home Phone: (831) 423-3199
Email: shzhang@cse.ucsc.edu	URL:
Thesis Title: <i>Global Interconnection Optimization for System-in-Package</i>	
Expected Graduation: 7/1/2002	Advisor: Dai

University of Michigan

Last Name: Aloul	First Name: Fadi
Work Address: Electrical Engineering & Computer Science, EECS Building, 1301 Beal Avenue	
City: Ann Arbor	State: Michigan Zip: 48109-2122
Work Phone: (734) 936-2828	Home Phone: (313) 584-8875
Email: faloul@umich.edu	URL: http://www.eecs.umich.edu/~faloul
Thesis Title: <i>Undecided</i>	
Expected Graduation: 5/1/2001	Advisor: Sakallah

Last Name: Bell	First Name: Jeffrey
Work Address: 1301 Beale Avenue,	
City: Ann Arbor	State: Michigan Zip: 48109-2122
Work Phone: (734) 936-2828	Home Phone: (734) 761-9499
Email: jlbell@eecs.umich.edu	URL: http://www.eecs.umich.edu/~jlbell
Thesis Title: <i>Sequential Implications of Delay False Path Analysis</i>	
Expected Graduation: 5/1/2000	Advisor: Sakallah

Last Name: Kim	First Name: JoonYoung
Work Address: 1301 Beale Avenue,	
City: Ann Arbor	State: Michigan Zip: 48109-2122
Work Phone: (734) 936-2828	Home Phone:
Email: jykim@eecs.umich.edu	URL:
Thesis Title: <i>Testability and Performance Relationships of Logic Circuits</i>	
Expected Graduation: 12/1/2000	Advisor: Sakallah

Last Name: Nam	First Name: Gi-Joan
Work Address: 1301 Beale Avenue,	
City: Ann Arbor	State: Michigan Zip: 48109-2122
Work Phone: (734) 936-2828	Home Phone: (734) 763-1523
Email: criteria@eecs.umich.edu	URL: http://www.eecs.umich.edu/~criteria
Thesis Title: <i>Undecided</i>	
Expected Graduation: 5/1/2001	Advisor: Sakallah

6. Financial Summary

The financial summary is presented in the three tables contained in this section. Our guide for expenses, shown in Table 1, derives from a revised proposed budget sent to MARCO and DARPA on 5 November, 1998. In that document, two plans for spending were outlined: an approved budget (referred to as Plan B in the 1998 document) and a proposed budget (Plan A). The requested or proposed budget represents a ramp-up in research in Year 3 (2001) of the program, consistent with plans that have been proposed by MARCO for all of the Focused Research Centers.

**Table 1: Total Program Cost
Breakdown and Cost Share**

	Year 1 1999	Year 2 2000	Year 3 2001 Proposed	Year 3 2001 Approved	Total Proposed	Total Approved
Focus Research Center Total Cost	\$5,016,120	\$4,986,624	\$9,497,256	\$2,497,199	\$19,500,000	\$12,499,943
MARCO Funds	\$2,636,527	\$2,603,118	\$8,260,355	\$1,256,520	\$13,500,000	\$6,496,165
DARPA Funds	\$2,379,593	\$2,383,506	\$1,236,901	\$1,236,901	\$6,000,000	\$6,000,000
Total Cost Share*	\$288,744	\$247,100	\$307,855	\$132,670	\$843,699	\$668,514

* Not included in total cost

Table 2 shows in detail how we administered funds in 1999 and our projection for expenses in 2000, our second year of operation. As outlined in the Executive Summary of this report, we exit our first year underspent by \$591,219. The principal variation in expenses from the first year's budget came in salaries, principally in the research staff. Some of the variance was redirected toward research in the subcontracting universities.

As shown in our projects for the year 2000, we expect to catch up to our budget, exiting the year with a small cumulative variance of approximately \$40 k. We expect our research effort at Berkeley, as well as the subcontracting universities, to be at full staff, resulting in a slight negative variance in salaries.

Table 2: Actual (1999) and Projected (2000) Expenses vs. Budget

	1999			2000		
	Approved Budget (11/5/98)	Project Costs thru Dec	Variance (to 2000)	Approved Budget (11/5/98)	Projected Costs	Variance (to 2001)
UC Berkeley						
Salaries						
Research	1,080,259	659,830	420,430	1,086,285	1,000,776	85,509
Staff	305,557	245,890	59,667	312,177	543,869	(231,692)
Benefits/Tuition	367,423	200,420	167,003	377,995	284,271	93,724
S/T Salaries	1,753,239	1,106,140	647,099	1,776,457	1,828,917	(52,460)
Equipment & ODC's						
Inventory Equip.	187,500	219,201	(31,701)	190,600	190,600	0
Other Direct Costs	122,200	164,686	(42,486)	108,800	197,200	(88,400)
Lease/Utilities	192,500	130,000	62,500	175,000	237,500	(62,500)
S/T Equipment	502,200	513,887	(11,687)	474,400	625,300	(150,900)
Indirect Costs 19.5%	346,306	224,201	122,105	345,766	247,349	98,417
TOTAL - UCB	2,601,745	1,844,228	757,517	2,596,623	2,701,566	(104,943)
Subcontracts						
Stanford	420,000	445,304	(25,304)	420,000	479,784	(59,784)
Carnegie-Mellon	760,000	879,205	(119,205)	760,000	917,940	(157,940)
Princeton	150,000	164,030	(14,030)	150,000	222,753	(72,753)
Michigan	150,000	150,000	0	150,000	150,000	0
Purdue	0	45,000	(45,000)	0	111,066	(111,066)
UCLA	300,000	412,759	(112,759)	300,000	490,741	(190,741)
UCSB	300,000	300,000	0	300,000	300,000	0
UCSD	80,000	80,000	0	80,000	80,000	0
UCSC	79,995	79,995	0	79,995	79,995	0
New Projects ¹	150,000	0	150,000	150,000	0	150,000
S/T Subcontracts	2,389,995	2,556,293	(166,298)	2,389,995	2,832,279	(442,284)
Indirects Costs 19.5% ²	24,375	24,375	0	0	0	0
TOTAL - Subcontracts	2,414,370	2,580,668	(166,298)	2,389,995	2,832,279	(442,284)
TOTAL - GSRC	5,016,115	4,424,896	591,219	4,986,618	5,533,844	(547,226)
Cummulative Total	5,016,115	4,424,896	591,219	10,002,733	9,958,740	43,993
Total Variance end of Year 2						43,993

1. Work Planned for MIT, now Performed by Princeton

2. Indirect Costs Applied to First \$25,000 of Each Subcontract

The spending plans going forward to 2001, our third year of operation, must be considered under two scenarios: the approved and requested or proposed budgets referred to above in Table 1. Plans for both of those scenarios are shown in Table 3.

**Table 3: Approved and Requested Budgets
& Projected Expenses for Year 3**

Year 3 - 2001						
	Approved Budget	Requested Budget	Projected Costs (Approved Budget)	Projected Costs (Requested Budget)	Projected Variance (from Approved')	Projected Variance (from Requested')
UC Berkeley						
Sub-total Salaries	733,634	1,709,030	629,267	1,873,656	104,367	(164,626)
Sub-total Equipment	80,140	441,400	80,140	529,800	0	(88,400)
Indirect Costs 19.5%	119,062	325,052	102,732	353,551	16,330	(28,499)
TOTAL - UCB	932,836	2,475,482	812,139	2,757,007	120,697	(281,525)
Subcontracts						
Other Universities	1,448,335	2,240,000	1,448,335	2,840,978	0	(600,978)
New Projects	116,028	4,562,645	280,718	3,724,151	(164,690)	838,494
Sub-total Subcontracts	1,564,363	6,802,645	1,729,053	6,565,129	(164,690)	237,516
Indirect Costs 19.5%	0	219,140	0	219,128	0	12
TOTAL Subcontracts	1,564,363	7,021,785	1,729,053	6,784,257	(164,690)	237,528
TOTAL GSRC	2,497,199	9,497,267	2,541,192	9,541,264	(43,993)	(43,997)
Cummulative Total	12,499,932	19,500,000	12,499,932	19,500,000	0	0
Total Variance end of Year 3:					0	0

This table should be viewed as having three major columns: Budget, Costs, and Variance. Each of the major columns contains an "Approved" sub-column and a "Requested" sub-column. This arrangement readily identifies the projected spending plans under the two scenarios. In either case, the net variance at the end of Year 3 is zero.

A significant distinction between the two spending plans, and the principal reason for our request for ramped-up funding in Year 3, concerns the support of New Projects that extend the scope of



our research in directions, and on projects, that simply cannot be supported on the Approved budget. Our plan in 2000 is to identify those programs in the Center that are both under-funded and show promise for extraordinarily high payoff. We also plan a continuous evaluation of those programs already under way that require additional funding, either with new resources or redirection of approved resources. The combination of those "New Projects" and augmentation of deserving existing programs show up as additional budget requests and projected costs. The "New Projects" line in Table 3 is particularly relevant to the increased funding request.

If we are to receive approval to ramp up our efforts in Year 3, we need to begin the process of writing proposals early in Year 2. We eagerly seek that approval. Based on the excellent beginning to GSRC that we have witnessed in 1999, we firmly believe that the ramp-up is warranted.