

# Advance Program

Sunday, 04 May 2003

**2:30 PM - 5:45 PM**

**Session TUT:** TUTORIAL SESSION

**Session Chair:** Tulin Mangir, *TM Assoc & Calif Univ, Santa Monica, CA, USA*

**TUT1 2:30 PM - 3:15 PM**

**The Challenges of Telematics: Where to From Here?** J. MacFarlane, *OnStar, Oakland, CA, USA*

The proliferation of mobile devices with location and communication capabilities continues to indicate a consumer desire for telematics services. Proliferation of in-vehicle electronics that provide specific mobile services, e.g. navigation and hands-free cellular, also indicates a consumer willingness to integrate telematics into their driving experience. However, the technology and services development for the telematics market has not grown at a fast pace. The purpose of this tutorial is to highlight the realities of developing vehicle-based services in terms of the business model, consumer demand and technology aspects, and, identify some of the development areas that still must be addressed before telematics can deliver to the high expectations that were originally created by the industry.

**TUT2 3:15 PM - 4:00 PM**

**Telematics and Its Place in Vehicle Infrastructure**, V. Rasin, *Ford Motor Company, Dearborn, MI, USA*

A few years ago the word 'telematics' was very popular. The future looked bright for car manufacturers with the promise of additional revenue streams from new technologies pouring into the company coffers. Today, reality has settled in, together with the understanding that the implementation and integration of telematics technologies into the vehicle is a complicated process, which must correlate with many other technical, business and marketing variables. The presentation will introduce the original meaning of telematics as a discipline and will show the transformation of this discipline over time, including various attempts to incorporate telematics technologies in the vehicle, successful or not. We will talk about where telematics technologies stand today and the approaches currently being taken by some car manufacturers. Finally, we will briefly describe the telematics-related developments at Ford Motor Company.

**4:00 PM - 4:15 PM**

**BREAK**

**TUT3 4:15 PM - 5:00 PM**

**Ad Hoc Networks**, L. Mitchell, *Sun Microsystems, Toronto, ON, Canada*

A widely accepted definition of an Ad Hoc network is generally taken as "Ad-Hoc networks are wireless, self-organizing systems formed by co-operating nodes within communication range of each other that form temporary networks. Their topology is dynamic, decentralized, ever changing and the nodes may move around arbitrarily."

This tutorial will focus on enough material to allow the participants to have some understanding of the architecture and issues surrounding Ad Hoc networks. Consideration of lower level mechanisms such as security, access control, discovery and QoS, self configuration and scalability are covered. However since such a discussion is not complete without an exploration of higher level software architectural frameworks, then this topic will be available if time permits.

**TUT4 5:00 PM - 5:45 PM**

**TBD**

**6:00 PM - 6:15 PM**

**Kickoff Meeting for Working Group Leaders**

Monday, 05 May 2003

**7:55 AM - 8:00 AM**

**Welcoming Remarks**

**Workshop Co-Chair:**

Alfred Mondelli, *SAIC, McLean, VA, USA*

**8:00 AM - 10:00 AM**

**Session MA:** OPTICAL INTERCONNECT TECHNOLOGIES

**Session Chair:** David T. Neilson, *Lucent Technologies, Holmdel, NJ, USA*

**MA1 8:00 AM - 8:30 AM****High Speed VCSEL and Photodiode Arrays**, M. Hibbs-Brenner, *Honeywell Inc., Plymouth, MN, USA*

The implementation of optical interconnects requires optoelectronic devices which are easy to produce in array form, are highly reliable, provide high modulation rates at low power dissipation, and can be driven at voltages compatible with common I.C. technologies. The practical realization of optical interconnects has been enabled by the development of Vertical Cavity Surface Emitting Lasers (VCSELs) and vertically detecting p-i-n photodiodes. Discrete devices and one-dimensional arrays are incorporated into optical subassemblies for use in transceiver modules which form the basis of enterprise level or cabinet to cabinet data links. The state-of-the-art for commercially available components is currently making the transition from data rates of 2-2.5Gbps to 10Gbps. However, further penetration of optical interconnects into systems, such as from board to board within a cabinet, will require higher levels of integration in order to achieve the required space, power dissipation and cost targets. This talk will cover a range of on-going technology developments, including the requirements, performance, and reliability of 10Gbps components, array packaging approaches, and chip scale integration of VCSELs, photodiodes, lenses and integrated circuits.

**MA2 8:30 AM - 9:00 AM****Optical Data Links over Polymeric Backplanes**, R. Michalzik, *University of Ulm, Ulm, Germany*

Optical interconnects are firmly established as economical approach to inter-system data links, preferably on the rack-to-rack level. It is generally agreed that the future increase of data throughput within computer frames will necessitate the introduction of optical technology, where optical backplanes are the most promising candidate. At present, manufacturers of electrical backplanes are already requested by customers to include optical solutions into their product portfolio. This presentation will provide some overview over approaches to optical backplanes and closely related hybrid electrical/optical printed circuit board technology. We will address requirements on vertical-cavity surface-emitting lasers (VCSELs) as transmitter elements of choice and present their current performance data as well as discuss challenges on the receiving end of large-area multimode waveguides. A focus will be put on recent experiments demonstrating 10 Gbit/s data transmission over various polymer waveguide based prototype backplane samples with data paths of up to 1 meter length.

**MA3 9:00 AM - 9:30 AM****Polymer Waveguides as Optical Interconnect Technology**, B. L. Booth, *Optical Crosslinks Inc., Kennett Square, PA, USA*

Interest continues to grow for the application of polymer waveguide systems to board level and chip-to-chip interconnects. Applications include point-to-point and other functionalities, such as splitters or embedded filters, which are either bonded to substrates or as freestanding flexible films. Near term commercial opportunities are aimed at datacom and military defense applications using multimode 800 nm region interconnections. Considerable activity involves high-speed interconnection for computer equipment at the chip level. The talk will review approaches to serve these markets with focus on the approach currently being taken by Optical CrossLinks' exploiting GuideLink™ (formerly DuPont Polyguide™). This technology is a self-developing high volume pre-coated film system for creating optical waveguides limited only by the photomask size in dimensions. Critical issues which must be met for a cost effective high performance system will be reviewed with GuideLink™ performance examples including precision micromachining, packaging, connectorization, coupling with I/O mirrors, stacked multilayer circuits, processing and operational compatibility. The emphasis will be on examples of prototype applications and system performance. Application versatility and future potential will also be discussed.

**MA4 9:30 AM - 10:00 AM****Research, Development and Commercialization of 2-D Array Optical Interconnects**, A. Krishnamoorthy, *Aralight, Jamesburg, NJ, USA*

We will discuss the development of an optoelectronic-VLSI technology that permits active optoelectronic devices to be "intimately" integrated with Silicon VLSI electronics. We will detail recent progress in commercializing this technology and describe a high-density 120Gbit/s 2-D parallel optical interconnect product that we are currently shipping.

**10:00 AM – 10:15 AM****BREAK****10:15 AM - 12:15 PM****Session MB: THE CASE FOR AND AGAINST OPTICAL INTERCONNECTS****Session Chair:****MB1 10:15 AM - 10:45 AM****Scaling Methods for Electrical Interconnects to Meet the Performance Requirements of Microprocessor Platforms**, R. Mooney, *Intel Corporation, Hillsboro, OR, USA*

Performance demands of microprocessor platforms continue to scale according to a Moore's law trend. In order to meet these requirements, we must either scale the interconnect performance at an equal rate; increase the complexity and size of the cache memory system, or both. Electrical interconnects have traditionally been used to meet these needs. I'll discuss historical scaling of these interconnects along with potential barriers to future scaling. I'll then discuss methods of dealing with these issues; focusing on the opportunities for continued cost effective scaling of electrical interconnect technologies for microprocessor platforms. I'll give performance projections demonstrating that electrical interconnects can continue to meet the requirements of these systems for several generations.

**MB2 10:45 AM - 11:15 AM****Recent Advances in Modeling and Simulation of Multimode Optical Links**, B. K. Whitlock, *RSoft Design Group, Inc., Milpitas, CA, USA*, J. Morikuni, P. V. Mena, *RSoft Design Group, Inc., Des Plaines, IL, USA* and R. Scarmozzino, *RSoft Design Group, Inc., Ossining, NY, USA*

Recently, there has been a strong interest in multimode optical communication systems for such applications as 1 Gigabit and 10 Gigabit Ethernet due to their performance benefits vs. copper solutions and potential cost savings over single-mode optical systems. This has led to

recent advances in modeling, simulation, and analysis of multimode optical data links. Various techniques have been used over the years, including purely analytical methods, system-level waveform simulations using linear and nonlinear models, and detailed device-level simulations of specific aspects of the physical links. Most recently, approaches have been developed that combine the complexity and accuracy of device-level simulation with the efficiency and utility of system-level simulation. We will describe a new link-level simulation platform specifically targeted toward the design, analysis, and simulation of multimode systems. This new tool enables the accurate analysis of an electromagnetic field starting with its generation by a laser, continuing with its propagation through spatially dependent components such as multimode fibers, and ending with its detection at a photoreceiver. Sample applications include multimode laser simulation, laser/fiber coupling, fiber/detector offset, fiber differential mode delay, transceiver/fiber encircled flux, and the like. This multimode simulation capability is housed within a broader link-level simulation infrastructure; thus, the extremely detailed spatio-temporal electromagnetic field calculations are implemented in a manner that allows straightforward system-level analysis.

### **MB3 11:15 AM - 11:45 AM**

#### **Short-Haul Optical Interconnect and Its Route to Effective Deployment, J. M. Van Campenhout, Ghent University, Gent, Belgium**

For years, meter and centimeter-range optical parallel interconnects have been advocated as a possible escape route from the imminent interconnect bottleneck predicted by several authors, including the successive issues of the SIA road map. The fundamental beneficial properties of optical interconnects have been studied extensively, and are by now well known to most researchers. Yet, the expected great breakthrough of the use of these interconnects in real systems other than long-haul telecommunications has not happened so far, and on-going research projects do not seem to be geared to make this quantum leap to happen either. A deeper insight into the underlying reasons for this would be highly welcome.

In previous work we have addressed this issue by studying the potentialities of short-range optical interconnect from a low-level systems point of view. This was done through a variety of techniques, ranging from analytical system and circuit modeling techniques over detailed simulations of realistic system contexts, to actual design and implementation of a system-level demonstrator. Part of this work was done in the context of the large-scale European Community sponsored research project OIIC. At this moment, a second major demonstrator system is being designed in the context of the IO project, homing in onto a real-life application in high-performance IP-router equipment.

This work has enabled us to draw some conclusions identifying the underlying reasons for the slow adoption of optical interconnects in real systems. In the presentation we shall discuss these issues, and indicate in what way we are addressing them in the ongoing IO project.

### **MB4 11:45 AM - 12:15 PM**

#### **Metal vs Optical Interconnects, K. Saraswat and P. Kapur, Stanford University, Stanford, CA, USA**

The continuous miniaturization of devices, while conducive to improvement of transistors, has left interconnects as performance bottleneck. On-chip copper (Cu) interconnects, at best, improve in absolute terms but cannot keep up with the gate delay improvements. While, in the worst case of global wires, their performance deteriorates even in absolute terms, in spite of introduction of low dielectric constant materials for capacitance reduction. On the other hand, off-chip metal interconnects such as those required for board communications present a second bottleneck owing to demands of keeping up with enormous increase in aggregate chip bandwidths due to both higher integration levels and faster clock speeds. This presents an opportunity for optical interconnect, albeit with scale and application specific requirements. We examine the role optics can potentially play in freeing this short distance interconnects bottleneck. We start by evaluating the on-chip metal interconnect performance from both power and delay perspective. Although, it is recognized that these interconnects will deteriorate, our modeling in the light of realistic technology effects on metal resistance, predicts a situation worse than previously anticipated. Above constraints have been insignificant in the past, but are becoming increasingly important with the aggressive scaling suggested in the ITRS. Scaling induced increase in electron surface scattering, fractional cross section area occupied by the high resistivity barrier and realistic interconnect operation temperature will lead to a significant rise in the effective resistivity of Cu. We model these effects and their impact on delay and indirectly on power dissipation and find it to be significantly deteriorating. Optics could be a promising alternative with superior isolation properties (cross talk) and potential delay/power advantage. We next model the optical interconnect system delay and power dissipation and compare it with the best available Cu interconnect on a chip in the form of repeated wires. There exists a critical length well within the chip size, beyond which the optical system is faster than the fastest copper system with repeaters. The delay-power scaling trends for both systems are also examined. For off-chip interconnects, although, high speed serial links on Cu over tens of centimeter have been reported, signal integrity and timing precision requirements at these high bit rates can drive up the power dissipation. Optics could provide these speeds potentially at a lower power. Further, it can provide much higher bandwidths at longer distances. Off-chip metal vs. optical interconnect comparison is in progress and has so far involved optical link power optimization and comparisons with a state-of-the-art electrical signaling system. There is quite a bit of doubt as to the adequacy of conventional metal based clock distribution schemes to meet performance specifications in the future. Therefore, one of the first application of optical interconnects on-chip is believed to be in clock distribution because of reduction in skew and jitter and in some cases even power advantages.

**12:15 PM – 12:30 PM**

#### **Workshop Problem Statement**

Michael McFadden, *University of Delaware, Newark, DE, USA*

**12:30 PM – 2:00 PM**

#### **Luncheon & Working Group Session I**

**2:00 PM – 4:00 PM**

#### **Working Group Session II**

**6:00 PM – 7:00 PM**

#### **Welcome Reception**

## Tuesday, 06 May 2003

### TuA1 8:00 AM - 8:30 AM

**Optical Interconnects, Outside the Computer Forever**, D. Huang, *Sun Microsystems, Palo Alto, CA, USA*, T. Sze, *Sun Microsystems, San Diego, CA, USA*, A. Landin, *Sun Microsystems, Menlo Park, CA, USA*, R. Lytel, *Sun Microsystems, Palo Alto, CA, USA* and H. Davidson, *Sun Microsystems, Menlo Park, CA, USA*

At lengths appropriate for chips, boards, and cabinets, it can be shown that optical interconnects outperform wires on almost all performance metrics. This is true at any scale length that supports guided waves, and for some free space cases. Despite this modeled advantage they are not common in computers. We will give a broad overview of current computer design, and compare optical to electrical methods at different levels in the interconnect hierarchy.

### TuA2 8:30 AM - 9:00 AM

**Perspectives on Future Interconnect Technologies**, E. M. Koontz and G.A. Frantz, *Texas Instruments Inc., Dallas, TX, USA*

Innovative product concepts for additional network peripherals, creative and new differentiated services, and more bandwidth-demanding applications under development will influence the need for increased bandwidth within the communication network infrastructure. The increase in peripherals, applications, and services will thus have a significant impact on the future highspeed interconnect requirements and architectures within a range of network infrastructure equipment. Furthermore, the addition of distributed services, such as storage applications, will likewise affect the selection of future interconnect technologies.

Binary electrical, multi-level electrical, single wavelength optical, and multi-wavelength optical solutions are examples of the technologies under investigation for the highspeed interconnections of the future. Although the optimum solution is not yet identified, in order to better down select, a number of topics are in need of further investigation; packaging, backward compatibility, and power and thermal management, to name a few. From the perspective of an electronic component supplier, the discussion will touch upon the bandwidth driving applications as well as a few of the areas in need of additional investigation for realization of highspeed interconnect solutions.

### TuA3 9:00 AM - 9:30 AM

**Optical Interconnect Requirements/Solutions for High Speed Switches**, R. P. Luijten, *IBM Research, Ruschlikon, Switzerland*

PRIZMA is a successful switch architecture that entered the switch market in 1995 with a 6.4 Gbps chip. Currently we have a 512 Gbps solution on a single card. This card requires 1024 connector pins at 2 GHz at the card edge to transfer the switched data packets. Half of the electrical power of the card is used to transport bits, while the other half is only used for the switching function. An important system metric is Gbps/ Watt/cubic inch for the aggregate effective switch throughput. To achieve market acceptance, this metric must not get worse with each new system generation. Working on our future Prizma evolution, we believe that at signaling rates around 10 Gbps the transition to intrarack optical interconnects must be made to improve the above power metric. We are working on an opto-electrical high-speed/high-density link and packaging technology that is expected to meet our system needs, and represents an evolutionary step from current FR4 card manufacturing technology.

### TuA4 9:30 AM - 10:00 AM

**Connecting to the Optical Network: Evolution Over the Next 10 Years**, R. J. Runser, *Telcordia Technologies, Inc., Columbia, MD, USA*

The accessibility of vast amounts of data for computing applications is critical since databases, data collection, and information sources are geographically separated. Real-time processing of large amounts of data at multiple computing sites is also becoming important for solving large-scale computing problems. Although the technology to support 10 Tb/s over fiber using DWDM technology has been demonstrated, the economics needed to leverage these advances in the design of future interconnect architectures and access network interfaces are vastly different. Local area networks have yet to take advantage of WDM technologies where the bandwidth requirements today are far less than the core network and the dominant costs of deployment are the optical transceivers required at each access node. The design of this interface for high performance computers is likely to undergo significant change over the next decade as bandwidth demands approach the level supplied in core network links. Extending backplanes and interconnects over significant distances may also become necessary to take advantage of distributed computing facilities. This presentation will discuss the trends over the next 10 years that may develop to provide the access to large bandwidth pipes needed to link high performance computing platforms over the external network.

10:00 AM – 10:15 AM

BREAK

10:15 AM - 12:15 PM

Session TuB: FUTURE INTERCONNECT TECHNOLOGY

Session Chair:

### TuB1 10:15 AM - 10:45 AM

**Nanophotonics in the Future for Optical Interconnects?** R. E. Slusher, *Lucent Technologies, Murray Hill, NJ, USA*

New optical materials and planar lightwave (PLW) structures, where nanometer features are important, promise chip level integration of thousands of optical devices along with high-speed electronics. This next generation of opto-electronic chips may play a major role in integrating the advantages of optical transport with the processing and storage capabilities of electronics for applications in the optical interconnect field. Using the example application of an optical packet router chip with a throughput of 2 Tb/s, several materials and PLW structures will be compared.

Recent material advances include silicon quantum dots and electro-optic polymers. Silicon quantum dots efficiently emit radiation near 900nm wavelengths and may provide enough gain for optical amplifiers and laser sources that could be fabricated using CMOS

compatible processing. Electro-optic polymers have potential for fabrication on silicon-based chips where they could serve in electro-optic modulators or switches.

High index contrast photonic crystal and total internal reflection structures will be compared. Are there any advantages for 2D or 3D photonic crystals in PLW circuits compared to more conventional total internal reflection structures? What is the optimum index contrast in order to minimize losses and optimize the density of optical devices? How advantageous are resonator structures in enhancing optical modulators and switches?

**TuB2 10:45 AM - 11:15 AM**

**Future OE Devices**, P. D. Dapkus, University of Southern California, Los Angeles, CA, USA

ABSTRACT NOT AVAILABLE

**TuB3 11:15 AM - 11:45 AM**

**Meso- and Nano-Scopic Technologies for Chip-Scale Optical Interconnects**, D. W. Prather, University of Delaware, Newark, DE, USA

In the course of this talk we will discuss the design, fabrication, characterization, and integration of meso- and nano-photonic devices for integrated chip-scale optical interconnects. In particular, two emerging technologies for optical interconnect applications will be addressed: the first is a mesoscale free space optical system based on diffractive optics and the second is based on a nanoscale planar approach that uses photonic crystals. To this end, the talk will introduce the development and use computational electromagnetic analysis tools, such as the plane wave method and the finite-difference time-domain method, for the design of optical devices on the meso- and nano-scopic scale. For fabrication we will present our newly developed methods for high aspect ratio and multi-layer photonic crystals in silicon, patterned with direct write electron beam and interferometric UV lithographies. In addition, we will also discuss our custom grayscale lithography and etching processes for both guided wave and free space optical devices. For characterization we will present our work on high efficiency coupling to photonic crystal devices and integrated free space optical networks. Lastly, we will also discuss proposed methods for their integration into actual chip-scale optoelectronic circuits.

**TuB4 11:45 AM - 12:15 PM (Invited)**

**Design Rules for Very Short Reach (VSR) Optics**, D. V. Plant, McGill University, Montreal, PQ, Canada

We will present work on very short reach (VSR) optical interconnects for chip and board applications. We will discuss how spatial parallelism can offer high density interconnects, and we will present design rules for achieving these interconnects.

**12:15 PM – 1:45 PM**

**Luncheon & Working Group Session III**

**1:45 PM – 3:45 PM**

**Working Group Session IV**

**3:45 PM – 7:30 PM**

**Free Afternoon**

**7:30 PM – 8:30 PM**

**Presentation of Workshop Problem Solutions**

**8:30 PM – 9:00 PM**

**Student Poster Session**

## Wednesday, 07 May 2003

**9:00 AM - 11:45 AM**

**Session PLE:** PLENARY SESSION

**Session Chair:** Marc Simpson, Oak Ridge National Laboratory, Oak Ridge, TN, USA

**PLE1 9:00 AM - 9:45 AM**

**High Speed Interconnects: Where We Have Been, Where We Are, Challenges Ahead**, H. Stone, NEC Research Institute, Inc., Princeton, NJ, USA



**Bio:**

Harold Stone is known for his work on the perfect-shuffle interconnection network that stimulated architectural advances in parallel computing. His textbooks in computer architecture include High-Performance Computer Architecture, Microprocessor Interfacing, and Computer Organization and Data Structures.

Dr. Stone received a BSEE degree from Princeton in 1960, and Ph. D. in Electrical Engineering in 1963 from the University of California at Berkeley. He has been active in the IEEE Computer Society and Communications Society as well as the ACM. Dr. Stone received a BSEE degree from Princeton in 1960, and Ph. D. in Electrical Engineering in 1963 from the University of California at Berkeley. Dr. Stone is a Fellow of the IEEE and ACM.

He received the IEEE Emanuel Piore Award in 1992 for contributions to parallel computation and education. In 1999 he received the Taylor Booth Award for contributions to education. He retired from NEC Research Institute in 2001.

**Abstract:** The key issue for high-speed interconnections is to grow the technology to match Moore's Law so that interconnect data rates can match the needs of processor technology. The issue started to become clear in the early 1990s when processor clocks were running twice the speed of the memory bus. Memory could not keep up then. By 2003 the ratio has reached 20 to 1, and the trend is for the ratio to get larger. Similar trends hold for the interconnection technology as a whole. This talk describes the reasons for the trends, various changes in interconnect technology that have helped reduce the problem, and what might prevail going forward. If we maintain the current course, future processors will starve for data. We need to see major breakthroughs in cost reduction that will enable high speed technologies to be deployed within systems.

**PLE2 10:00 AM - 10:45 AM**

**Optical Interconnects Today**, J. D. Crow, *Crowsquill Consulting, Claremont, CA, USA*



**BIO:**

Dr. Crow, a UC Berkeley graduate, has been a major contributor to optical interconnect technology in computer systems for the past few decades. At IBM, during the 90's, he led the company and government sponsored programs Nexus, AARTs, OETC, and Jitney, which researched parallel optical interconnect and developed flip chip mounted array OE and OEIC chips, Si carriers for parallel interconnects, and plastic molded parallel optical modules and cables. He interspersed this with the development of high-speed serial optical data links, for IBM storage network products, and 1-10 Gb/s LANs. He has contributed to optical technology roadmap studies for NAS, OIDA, and participated in IEEE and ANSI data link standards bodies, and helped organize IEEE conferences in the field. He is an IEEE Fellow.

**ABSTRACT:** Although the performance-cost roadmap trajectory for intrasystem interconnects has been lowered by the recent IT market pause, both computer systems and communications systems still project requirements for tens of terabits/sec of data flow within the system interconnect, within the decade. Electrical solutions require increasingly complex, clumsy, and expensive technology; yet the case for an optical solution is all but clear. This talk compares the parallel optical modules and wiring coming to the market with the system and electrical metrics; and discusses the speed, power, reliability, and manufacturability (cost) challenges for an optical interconnect technology. Strategies for both performance and cost enhancements will be discussed, emphasizing the need for merging rather than competing the optical and electrical interconnect options.

**PLE3 10:45 AM - 11:30 AM**

**TBD**

ABSTRACT NOT AVAILABLE

**11:30 PM – 11:45 PM**

**Closing Remarks**

Matthew Goodman, *Telcordia Technologies, Red Bank, NJ, USA*