S. K. Tewksbury, M. Uppuluri and L. A. Hornak, "Interconnections/micro-networks for integrated microelectronics," Global Telecommunications Conference, 1992. Conference Record., GLOBECOM '92. Communication for Global Users., IEEE, Orlando, FL, 1992, pp. 180-186 vol.1.
doi: 10.1109/GLOCOM.1992.276496

Interconnections/Micro-Networks for Integrated Microelectronics

S.K. Tewksbury, M. Uppuluri, and L.A. Hornak
Dept. of Electrical and Computer Engineering
West Virginia University
Morgantown, WV 26506

Abstract

As the scale of integration of monolithic ICs has increased, formerly multi-chip systems have evolved to single chip circuits and new system architectures and applications have been enabled. Assuming continuing increases in the scale of integration through the present decade, multi-chip systems will evolve to become distributed digital systems scaled to small areas and volumes. Traditional, passive interconnections will simultaneously evolve to "micro-networks", to some extent scaling present local area networks to the small dimensions of integrated microelectronic systems. The technologies of microelectronics provide new capabilities which can be exploited to develop novel communications networks on a small scale. At the same time, important constraints, quite different from traditional constraints on LANs, will impact the design of micro-networks. This paper looks ahead to the evolution of such interconnections into the era of ULSI.

1 Introduction

The introduction of successive generations of silicon integrated circuit technologies (familiar through well-known acronyms for those generations: *SSI*, *MSI*, *LSI*, and *VLSI* corresponding to small, medium, large and very large scales of integration, respectively) has launched major new directions during the evolution of digital electronics. The next generation of traditional silicon device technology (*ULSI* or ultra-large scale integration) will allow complex, high-performance *system on a chip* and perhaps even *distributed systems on a chip.*

As the integrated circuit (IC) technology evolves toward ULSI, significant changes (certainly evolutionary and perhaps revolutionary) in the system-level packaging technologies already appearing for high-end VLSI will also see significant technology changes. Together, the evolution of ICs to systems on a chip and the introduction of advanced packaging technologies to efficiently implement multi-chip systems will define a new communications environment for distributed systems, distributed over the small area of integrated microelectronics. These miniaturized distributed systems establish the probability that the local area networks in present distributed systems will evolve to miniaturized networks, here called "micro-networks" [1,2].

In particular, following previous historical trends in which previously multi-chip functions evolved into single chip functions (and large-scale systems evolved into smaller scale and lower cost systems) as the IC technology advanced, one possible model for future ULSI systems is a highly-integrated, small area realization of what we presently know as distributed systems. This model is compelling since considerable research and development is being directed at combining the resources of presently distributed computers, for example, into a high-performance, networked computer applied to a single user application. The knowledge and experience gained from these present distributed systems will provide the basis for useful, miniaturized high performance systems of the future. Although perhaps a compelling model for future ULSI systems, there are important new capabilities provided within the world of microelectronic interconnections not available for traditional LANs. At the same time, there are important constraints which may render traditional LAN networks as impractical. The result is an opportunity to explore new possibilities and constraints within the area of communications, ignoring the details of the application and focussing on the communications function as a fundamental, microelectronic system function in its own right.

This overview of interconnections within microelectronic systems is provided from the viewpoint of a communications-oriented audience, drawing on this expected evolution to ULSI systems. Section 2 discusses general issues and opportunities arising in the area of interconnections and packaging, along with some important characteristics suggesting the need for new communication network architectures and protocols for "micro-networks". Section 3 reviews several of the technology issues, with particular emphasis on advanced packaging evolving for present VLSI systems. Section 4 summarizes performance models for the "communications channels" provided by microelectronic interconnections. Finally, Section 4 considers high speed communications within a microelectronic system.

0-7803-0608-2/92/$3.00 © 1992 IEEE

2 Microelectronic Interconnections as Communications Networks

The term "microelectronic interconnection" generally conveys the concept of a point-to-point interconnection between microelectronic devices on an integrated circuit. Here, a more general definition is used, namely a system-level interconnection based on thin film technologies. The generalization is important since the familiar thin film technologies used for integrated circuits are being applied (with significant differences in the technology details) to advanced packaging. Of particular importance are the so called *multi-chip modules* or *MCMs* [3,4], using a substrate compatible with thin film fabrication of electrical interconnections as a replacement for traditional printed circuit boards. Rather than individually packaging the integrated circuits within a generic plastic or ceramic package with a standard footprint of I/O pins suitable for mounting on a printed circuit board, unpackaged ICs are directly mounted on the multi-chip module. The much closer spacing of ICs and the much higher interconnection density supported by the multi-chip modules are expected to provide a high-performance, system-level microintegrated technology for present, high-end VLSI and for future ULSI circuits.

já existia!

SoC

With the arrival of the "system-on-a-chip", the advanced packaging will provide the physical environment for implementing communications between ICs comprising a miniaturized, distributed system. However, there may also be important cases in which a single integrated circuit itself will be a distributed system (e.g. an array of microprocessors or a collection of peripheral circuits). In this case, the traditional interconnections seen within an IC will also provide an environment for communications between systems within the chip. The electrical performance of such interconnections (both on-chip and MCM) has been studied extensively, including models developed both for microwave applications [5,6,7,8,9] and for high-speed IC environments [10,11,12,13,14,15] to evaluate their ability to handle high speed data transfers (e.g. Gbit/sec data rates on individual lines).

Even without the extension to MCMs, higher performance printed circuit boards and high pinout packaging for advanced VLSI components are rapidly being developed. These high-performance interconnection environments suggest that the evolving microelectronic system communications networks will differ substantially from present LANs for distributed systems. The analogy here is with distributed systems, rather than with massively parallel computing systems, since the latter are generally constrained by some view of regularity of structure. That regularity of structure is too strong an assumption for our purposes in this review. On the other hand, the massively parallel systems provide a useful view of communications as they are themselves scaled to more compact systems. To some extent, the scaled, massively parallel computing systems are an extension of presently well-known regular VLSI computation structures such as systolic arrays or wavefront processors.

Below, some important contrasts between micronetworks and conventional communication network technologies are briefly summarized to illustrate the new arena for research and development which is emerging.

2.1 Parallel vs Serial Communication Channels

It is assumed here that VLSI/ULSI circuits will evolve to ICs with "communication ports" rather than the traditional "I/O" ports. The transition of terminology from I/O port to communications port is already seen in VLSI circuits such as the Transputer and some high-end digital signal processors designed for easy implementation of array signal processors The microprocessor bus (with address, data, and control lines) is a primitive example of a "communications port", allowing a microprocessor to communicate with a variety of peripheral system ICs. At the same time, the single-master/multiple-slave structure of most microprocessor busses will become increasingly confining as miniaturized distributed systems evolve with several "masters" requiring simultaneous access to communication channels. It is this view of microelectronic systems evolving to distributed, communicating masters which suggests that traditional system busses will become less effective as VLSI evolves to ULSI.

!

Traditional communication networks (LANs, WANs, etc) are based on serial data transfers. Not only are the binary digits of the data word(s) time-division multiplexed onto a single line but also the destination address and generally the control signals are also time-division multiplexed onto that line. The dominance of serial data communications is partly a result of the economies necessary for long distance communication lines. Even if multiple lines per communication channel were used, difficulties maintaining equal end-to-end delays on the separate lines would complicate communication control. Such historical constraints are likely to be greatly relaxed for miniaturized distributed systems.

For example, given the edge dimension of an IC and the interconnection densities expected, communication channels for microelectronic systems are likely to be parallel data channels (i.e. space-division multiplexing of the binary digits of the information on separate wires), rather than the serial data channels often seen in traditional communications systems. The use of several, parallel "wires" to carry information over a communication channel pro-

vides an important means to achieve high net data rates within miniaturized distributed systems, acting as a multiplier of the data rate per wire to obtain the net communications rate.

2.2 Small Area Network Latencies

Management of information flow through communication networks typically requires some form of end-to-end control. This can be a particularly difficult problem if the delays (latencies) seen in moving information from one end of the network to the other considerably exceed (as in the case of long distance lightwave communications) the period at which successive data bits are inserted onto a communications channel. A natural result of the latency conditions (particularly their impact on network access and control) has been to transfer a large number of data bits upon each setup of an end-to-end connection, leading to communication networks based on transfer of packets containing blocks of data

Data transfers on traditional, point-to-point microelectronic interconnections need no end-to-end control protocol since the line is dedicated to transfer signals from a given source to the far end destinations. The standard protocols used to control a microprocessor bus protocols are based on a simple handshake procedure ("ready for data, data ready, data taken, etc.). Such simple handshake procedures seem rather primitive compared to their communications counterparts, particularly since one must complete several exchanges of information for each data transfer. The end-to-end latencies need not be excessive compared to the clock period (in contrast to traditional communications networks) and if each handshake triggers the parallel transfer of several data and address digits during one bus cycle (though this would require a large number of parallel data lines. A memory to processor data transfer can be also organized to move a block of data on successive high speed clock cycles, once the connection has been established, the analog of packet-based communications. In this sense, the simplicity of the bus protocol does not intrinsically restrict the communications rates except at the level of very fine grain communication (i.e. one data transfer event for each set of handshake control events).

However, the simple handshake protocol above tends to presume a single master active at any time (otherwise many clock cycles can be consumed in selecting and enabling a new master). As microelectronic systems evolve to miniaturized distributed systems with multiple masters communicating among the ICs, traditional communication problems such as conflicts, congestion, etc. will arise and the simple handshake protocol of the microprocessor bus will become an increasingly serious performance limit. For such reasons, it can be expected that protocols for the miniaturized communication networks within microelectronic distributed systems will undoubtedly borrow strongly from the general protocol approaches developed in traditional communications systems. However, major opportunities for efficiency arise since (1) end-to-end latency is small, (2) data, address, and control can propagate simultaneously on spatially separated wires in a communications channel, and (3) the topography and detailed electrical behavior of the micro-network can be fully engineered by the microelectronic systems designer (in contrast to the a-priori unknown lengths and distribution of lines in a traditional communications network as the network evolves).

2.3 Non-Uniform Microelectronic Communications Channels

Although, as noted above, the designer of a microelectronics communications network need not anticipate a vast range of possible communication channel installations by different users of the network, that designer does confront a significant problem not seen in traditional communication networks, as discussed here.

The communication lines of a traditional communication network are generally highly uniform structures. For example, the entire link between source and receiver for an optical fiber communication channel is a highly uniform physical structure with well-defined end-to-end signal propagation characteristics largely independent of the local details of the optical fiber. Given the source-end connection, the receiver-end connection and the overall length of the optical fiber, the performance of the overall optical fiber link can be rather accurately predicted.

Microelectronic interconnections impose numerous localized structures leading to end-to-end characteristics which depend on the local details of the interconnection. For example, most microelectronic interconnections use multi-level metalization, with the overall communication channel constructed from a number of short links within given metalization levels and via connections between those levels [10]. Microelectronic interconnections generally also pass through different packaging levels, encountering wire bonds, solder bumps, package interconnections, etc. [11,12,13,14,8]. Although this distributed structure has only a limited impact on the end-to-end performance of the interconnection at low data rates, the localized structure increasingly impacts signals traversing the interconnection as the data rate increases.

A microelectronic interconnection confronts not only the above problem of electrically important physical inhomogenieties along its length but also couplings to neighboring interconnections running along side or crossing over that interconnection [9,15]. Such couplings com-

promise to some extent the convenience of parallel data links for individual communication channels and certainly limit the achievable density of interconnections. However, so long as the couplings and inhomogenieties of an interconnection can be adequately modeled and their effects controlled, the communication channel(s) can be designed to provide the desired level of performance. In particular, the permanent structure of the microelectronics communications channel (in contrast to the unknown routing of cables and fibers in conventional communication networks) provides the necessary control to handle the electrical complications of the microelectronic interconnections. Electrical models for microelectronic system interconnections are summarized later.

3 Emerging Interconnection Technologies for Microelectronics

Interconnections within microelectronic systems traditionally are embedded within a hierarchy of packaging. On-chip interconnections provide very high density and, by minimizing length, high performance. Chip-to-chip interconnections must traverse the chip package structures and the interconnection environment of the circuit board (conventional printed circuit board and/or MCM). For complex systems, higher levels of packaging are seen in the backplane interconnections (and their driver/receiver circuitry). The microelectronic systems considered here are bounded by the backplane level of interconnection (e.g. consist of at most a single backplane with associated printed circuit boards, MCMs, and conventionally packaged ICs). Communications in traditional microelectronic systems generally appears at the higher levels of the packaging hierarchy (e.g. at the backplane level or higher). However, as VLSI evolves to ULSI and the communicating system components move deeper within the packaging hierarchy, it will be necessary to move the nodes of the communications environment into lower levels of the packaging hierarchy.

The traditional technologies for microelectronic systems are likely to evolve towards significantly "higher tech" supporting technologies. The multi-chip modules have already been briefly mentioned. By eliminating the first level of packaging (i.e. the packages used for individual ICs), significantly higher performance chip-to-chip interconnections are expected. The extent to which MCMs will replace conventionally packaged ICs is presently a subject of serious debate. The issue is not merely cost, though the present high cost of MCMs is certainly a concern. Additional issues include the so called "good bare die" issue, reflecting the difficulty in obtaining fully tested but unpackaged ICs from major IC manufacturers. More subtle but also practical issues arise from the loss of a package with a standard footprint able to handle, through the versatility of wire bonds, changes in the precise layout of IC bonding pads through successive generations of technologies and among different suppliers of a given IC function. Overall, the short term potential for insertion of MCMs as routine parts of microelectronic systems is unsettled due to these and other practical issues. However, the long term probability that MCMs will become a mainstream packaging approach is high, given the pressures imposed by very high pinout and very high speed ICs. A good example of the type of highly integrated IC which the future must easily support is the *Alpha* microprocessor recently announced by Digital Equipment Corporation. With several hundred I/O pins and clock rates of 200 MHz and higher, conventional packaging will be severely pressed to provide adequate performance.

Multi-chip modules, however, represent merely one important direction for improved, system-level technologies for microelectronics. The advantages of 3-D system structures are well known but have been previously merely the dreams of a few. New technologies are aggressively being applied to demonstrate three dimensional system level structures, based on IC technologies and on MCM technologies [16,17,19]. An example of a 3-D integrated circuit technology (using a stack of IC die) is seen in the recent work by TI for memory [16]. By producing a "pancake" stack of unpackaged memory ICs and providing a general purpose set of I/O pins to the stack, very high density memory systems are achieved. Hughes continues work on their three-dimensional signal processor [17], using "pancake" stacks of wafer-scale integrated signal processor circuits. 3-D system structures using stacked MCMs are also being aggressively investigated. The traditional limitation of very compact, 3-D structures has been removal of the heat generated by electronics near the center of the stack. This remains an important limitation but integration of heat removal elements within the stack is expected to relax this limitation.

Another important, though controversial, direction for microelectronics technologies is the integration of electronics to support communications directly within the "circuit board". Although this is clearly not realistic for conventional circuit boards, such integrated electronics is clearly possible (though practicality remains in dispute) for MCMs using silicon substrates. Rather than merely placing passive interconnections on a virgin silicon wafer, the "active circuit board" would include drivers, switches, and other circuit functions necessary to support communications among VLSI and ULSI circuits with minimum I/O drive power dissipation. The alternative is to place the electronics directly in the VLSI/ULSI circuits, rather than in the silicon circuit board. However, the combination of a large number of I/O pins and very high data

rates imposes considerable power dissipation and I/O port area penalties on the VLSI/ULSI circuit. These penalties are relaxed if line drivers are integrated directly on the circuit board. Performance limitations which lead to any active circuitry on the silicon circuit board will enable a vast amount of active circuitry for other functions (i.e. it is as easy to fabricate a silicon based MCM with a large number of transistors as with a few transistors).

The active silicon circuit board represents a rudimentary form of wafer-scale integration (WSI) [18]. Traditional applications targeted for WSI have involved massively parallel, regularly organized arrays of processors to provide high performance computation. Array structures have ranged from systolic arrays for matrix computations and other applications through more general purpose computing arrays with programmable interconnections to image processing arrays. Providing reorganization of interconnections to bypass faulty portions of the arrays such that the rate of data flow between point-to-point connected computational nodes is not severely compromised has been both a focus and a practical limitation of WSI. However, the communications network function that could arise through use of active silicon circuit boards (with computing nodes added as attached ICs) provides a rather different set of tradeoffs and deserves separate consideration.

The directions above retain electrical interconnections as the means of chip-to-chip interconnections. However, there has been considerable study of optical interconnections between ICs and between circuit boards (e.g. optical backplanes). The use of optical interconnections is, like the other significant advances in microelectronic system-level interconnections, a controversial topic. The potential advantages have been addressed well in several publications [19]-[29]. The practical constraints are also well known. Semiconductor lasers may have significantly lower reliability than conventional electronic drivers. Optical components remain expensive. The opto-electronic interfaces must have high performance, be very inexpensive, and provide easy connections between the optics and the electronics. Although these and other concerns are serious questions, considerable progress can be seen in the research towards optical interconnections. For example, high density arrays of vertically emitting lasers demonstrate the potential for very high connectivity using free space optics [29]. Holographic interconnection techniques have been extensively explored (see for example [20]), and extended to recent techniques using diffractive optics. Free space optics have been reported [19] as a means of achieving a 3-D stacked wafer system which does not require extensive, mechanical contacts as would be necessary for electrical interconnections. Although much of the emphasis on optical interconnection research has been toward very high speed and/or high density interconnections (avoiding the limits of electrical interconnections between conventionally packaged ICs), optical interconnections provide some important though less glamorous practical advantages. As mentioned above, they may relax the need for rigid mechanically connected components using advanced packaging (the rigid connections making disassembly and repair/replacement of defective components difficult). In addition, optical interconnections provide the important function of electrical isolation, with regard to signals, between components. Direct electrical switching of optical signals to achieve a microelectronic communications network presents problems since traditional couplers do not scale to sufficiently small dimensions. However, optical modulators may provide the necessary switching function.

4 High Speed Miniaturized Distributed Systems

For those accustomed to Gbit/sec data rates on very long distance optical fiber communication networks, the relatively modest data rates in the 20-50 Mbit/sec range seen in microelectronic systems with short distance interconnects is perhaps a surprise. However, such relatively low data rates reflect the cost, power dissipation limits, and physical interconnections seen in contemporary electronic systems.

The architecture of a present microelectronic system is left unchanged while the underlying VLSI technology is scaled to higher densities, the size of that microelectronic system decreases and the lengths of its interconnections also decrease. The result is the potential ability to scale up the data rates on interconnections as the system is scaled to smaller sizes. The scaling of the VLSI technology to higher densities carried the potential of higher power dissipation unless the power dissipation of the individual transistors and the interconnection drivers also decreases during the scaling. Under assumptions of an unchanged power dissipation per integrated circuit (e.g. < 5 Watt/chip), the speed of the individual transistors and line drivers must remain approximately unchanged as the VLSI technology is scaled. Under such constraints and for several other reasons, there is little opportunity to exploit the potential for well controlled behavior on the shorter system-level, chip-to-chip lines at higher frequencies. For a considerable time, characteristic clock rates for data transfers in microelectronic systems have been stuck below 100 MHz as the IC technologies evolved to higher density components.

Recently, there has been a rapid growth in the clock rate of microprocessors, seeking to maximize the performance of microprocessor-based computers and worksta-

tions. Rather than scaling the speed of the entire system, the microprocessor clock rate increases while much of the remainder of the system remains at the same speed (e.g. memory speed, peripheral speeds, etc.). The transition between the constant speed and increasing speed portions of the system are, in this case, handled architecturally (e.g. through use of cache memories to provide a data rate transition between the processor and the traditional user memory). Such architectural adaptations will increasingly be sought to avoid excessive requirements for high speed communications as processors reach toward higher speeds.

Despite the possibilities of architectural approaches to relax communication rate requirements, pressures to develop significantly higher data rates are likely to increase as the technology passes the level enabling a "system on a chip" and advances to even higher performance capabilities (particularly for applications which are not easily implemented using a distributed architecture). Obtaining significant increases in system speed will favor advances not only in high speed ICs but also in high speed packaging and interconnection technologies for multi-chip systems using those ICs. Confronting the objective of high speed communications in miniaturized distributed systems is the limit discussed earlier. In particular, an end-to-end microelectronic system interconnection is a highly heterogeneous structure, reflecting the multi-level metalization of ICs, of MCMs, and of printed circuit boards and reflecting the connections between the various levels of packaging (e.g. traversal of the IC's package, passage through the drivers and receivers of the backplane connections, etc). There will therefore be a considerable possible range of performance differences among interconnections within a microelectronic system. Management of the microelectronic system design will require accurate and usable models of interconnection behavior, including their local structures.

As the data rate increases within a miniaturized microelectronic distributed system, we eventually enter the regime of distributed RC lines, lossy transmission lines, and dispersive transmission lines, along with signal distortions due to discontinuities such as vias and crossovers (see, for example, [5,6,7]). In addition, as the data rate increases, there are likely to be significant crosstalk problems as one seeks to combine parallel lines per communication channel, minimum separation between those parallel lines of a given channel, and fast risetime signals. For randomly routed interconnections in a highly disorganized system, the engineering problems are significant. However, such high data rates may be more manageable regular network structures.

5 Summary

Several trends in microelectronic interconnections have been reviewed. To a considerable extent, the interconnection technologies and their performance characterization and modeling will evolve smoothly as VLSI evolves toward ULSI. However, it has been suggested that these trends are particularly important when one considers the role that interconnections will play in future microelectronic systems. With each IC a system-on-a-chip (or containing several systems-on-the-chip), the interconnections between ICs perform the traditional function of a communications network forming the backbone of a distributed system. However, the physical environment for constructing such micro-networks is quite different from traditional communication network (e.g. LANs). For example, each communication channel of a micro-network can contain several "wires", rather than the single, serial data channel of a traditional communications network. In addition, end-to-end delays across the micro-networks are much less (relative to clock period) that those of traditional communication networks. A simple mapping of traditional network strategies and structures into the microelectronic system communications function may not adequately exploit the opportunities for performance provided by the microelectronic system technologies. Instead, new strategies and structures are likely to be needed.

References

[1] S.K. Tewksbury and L.A. Hornak, *Communication network issues and high density interconnects in large-scale distributed computing systems,* IEEE J. Selected Areas in Communications, vol. 6, p. 587, 1988.

[2] S.K. Tewksbury and L.A. Hornak, *Future physical environments and concurrent computations,* in *Concurrent Computations: Algorithms, Architecture and Technology,* S.K. Tewksbury, B.W. Dickinson and S.C. Schwartz (Eds), Plenum Press (1988).

[3] R.W. Johnson, R.K.F. Teng and J.W. Balde (eds), *Multichip Modules,* New York: IEEE Press (1991).

[4] G.F. Taylor, B.J. Donlan, J.F. McDonald, A.S. Bergendahl and R.H. Steinvorth, *The wafer-transmission module – wafer scale integration packaging,* Proc. IEEE Custom Int. Circuits Conf., pp. 55-58 (1985).

[5] K.W. Goosen and R.B. Hammond, *Modeling of picosecond pulse propagation in microstrip interconnections on integrated circuits,* IEEE Trans. Microwave Theory and Techniques, vol. MTT-37, pp. 469-478 (1989).

[6] J.R. Brews, *Transmission line models for lossy waveguide interconnections in VLSI,* IEEE Trans.

Electron Dev., vol. ED-33, pp. 1356-1365 (1986).

[7] A. Deutsch et al, *High speed signal propagation on lossy transmission lines,* IBM J. Res. Develop, vol. 34, pp. 601-615 (1990).

[8] O. Palusinski, J.C. Liao, P.E. Teschan, J.L. Prince and F. Quintero, *Electrical modeling of interconnections in multilayer packaging structures,* IEEE Trans. Components, Hybrids and Manuf. Technol., vol. CHMT-10, pp. 217-223 (1987).

[9] V.K. Tripathi and R.J. Bucolo, *Analysis and modeling of multilevel parallel and crossing interconnection lines,* IEEE Trans. Electron Dev., vol. ED-34, pp. 650-658 (1987).

[10] W.M. Loh, S.E. Swirhun, T.A. Schreyer, R.M. Swanson and K.C. Saraswat, *Modeling and measurement of contact resistances,* IEEE Trans. Electron Dev., vol. ED-34, pp. 512-524 (1987).

[11] A.J. Rainal, *Performance limits of electrical interconnections to a high-speed chip,* IEEE Trans. Components, Hybrids and Manuf. Technol., vol. CHMT-11, pp. 260-266 (1988).

[12] R. H. Caverly, *Characteristic impedance of integrated circuit wire bonds,* IEEE Trans. Microwave Theory and Techniques, vol. MTT-34, pp. 982-984 (1986).

[13] D.Herrell and D.Carey, *High-frequency performance of TAB,* IEEE Trans. Components, Hybrids and Manuf. Technol., vol. CHMT-10, pp. 199-203 (1987).

[14] S.M. Wentworth, D.P. Neikirk and C.R. Brahce, *The high-frequency characteristics of tape automated bonding (TAB) interconnects,* IEEE Trans. Components, Hybrids and Manuf. Technol., vol. 12, pp. 340-347 (1987).

[15] K.C. Saraswat and F. Mohammadi, *Effect of scaling of interconnections on the time delay of VLSI circuits,* IEEE J. Solid-State Circuits, vol. SC-17, pp. 275-280 (1982).

[16] S.N. Shanken, R. Some, and M. Suer, *3-D processor packaging and interconnect,* Proc. GOMAC '91, pp. 151-154 (1991).

[17] J. Grinberg, G.R. Nudd and R.D. Etchells, *A cellular VLSI architecture,* IEEE Computer, pp. 69-81 (1984).

[18] S.K. Tewksbury, *Wafer Level System Integration: Implementation Issues,* Kluwer Academic Press (1989).

[19] L.A. Hornak and S.K. Tewksbury, *On the feasibility of through-wafer optical interconnects for hybrid wafer-scale integrated architectures,* IEEE Trans. Electron Dev., vol. ED-34, pp. 1557-1563 (1987).

[20] M.R. Feldman, S.C. Esener, C.C. Guest and S.H. Lee, *Comparison between optical and electrical interconnects based on power and speed considerations,* Applied Optics, vol. 27, pp. 1742-1751 (1988).

[21] W.H. Wu, L.A. Bergman, A.R. Johnson, C.C. Guest, S.C. Esener, P.K L. Yu, M.R. Feldman and S.H. Lee, *Implementation of optical interconnections for VLSI,* IEEE Trans. Electron Dev., vol. ED-34, pp. 706-714 (1987).

[22] A. Guha, J. Bristow, C. Sullivan and A. Husain, *Optical interconnections for massively parallel architectures,* Applied Optics, vol. 29, pp. 1077-1093 (1990).

[23] E.E. Frietman, W.V. Nifterick, L. Kekker and T.J.M. Jongeling, *Parallel optical interconnects: Implementation of optoelectronics in multiprocessor architectures,* Applied Optics, vol. 29, pp. 1161-1177 (1990).

[24] R. Kostuk, J.W. Goodman and L. Hesselink, *Optical imaging applied to microelectronic chip-to-chip interconnections,* Applied Optics, vol. 24, pp. 2851-2858 (1985).

[25] D.H. Hartman, *Digital high-speed interconnects: A study of the optical alternative,* Optical Engineering, vol. 25, pp. 1086-1102 (1986).

[26] R. Selvaraj, H.T. Lin and J.F. McDonald, *Integrated optical waveguides in polyimide for wafer scale integration,* IEEE J. Lightwave Technol., vol. LT-6, pp. 1034-1044 (1988).

[27] L.A. Hornak et al, *The impact of Polymer Integrated Optics on Silicon Wafer Area Networks,* in SPIE vol. 1337, *Nonlinear Optical Properties of Organic Materials III,* pp. 12-22 (1990).

[28] P.J. Ayliffe, J.W. Parker, A. Robinson, *Comparison of optical and electrical data interconnections at the board and backplane levels,* SPIE vol. 1281: Optical Interconnections and Networks, pp. 2-15 (1990).

[29] J.L. Jewell, J.P. Harbison, A. Sherer, Y.H. Lee and L.T. Florez, *Vertical-cavity surface-emitting lasers: Design, growth, fabrication, characterization,* IEEE J. Quantum Electronics, vol. QE-27, p. 1332 (1991).