

Incorporating Multi-Chip Module Packaging Constraints into System Design[†]

Vivek Garg, Steve Lacy, David E. Schimmel, Darrell Stogner,
Craig Ulmer, D. Scott Wills, and Sudhakar Yalamanchili

Packaging Research Center
School of Electrical and Computer Engineering,
Georgia Institute of Technology, Atlanta, Georgia 30332-0250
Phone: (404) 894-2940, Fax: (404) 853-9959, e-mail: sudha@ece.gatech.edu

Abstract

Computer system design addresses the optimization of metrics such as cost, performance, power, and reliability in the presence of physical constraints. The advent of large area, low cost Multi-Chip Modules (MCM) will lead to a new class of optimal system designs. This paper explores the early analysis of the impact of packaging technology on this design process. Our goal is to develop a suite of tools to evaluate computing system architectures under the constraints of various technologies. The design of the memory hierarchy in high speed microprocessors is used to explore the nature and type of trade-offs that can be made during the *conceptual design* [15] of computing systems.

1 Introduction

The optimization of metrics such as cost, performance, power, and reliability subject to physical constraints, is central to computer system design. An accurate analysis of the effect of packaging technology and options early in the system design process will enable the implementation of systems that *effectively* utilize available packaging technology. In particular we are interested in the impact of Multi-chip Module packaging technology on the architecture of computing systems. A first step towards this goal is the development of analysis tools to assist the systems designer in understanding the effect of packaging constraints on the computing system architecture.

This paper describes initial efforts towards assessing the impact of the next generation of MCM packaging technologies on the architecture of computing systems. The challenge is in the construction of an evaluation environ-

ment that will i) enable rapid prototyping of high level models of computing systems early in the design cycle, ii) facilitate the analysis of the relationships between packaging technology and system metrics such as cost, performance, power and reliability, and iii) enable an interactive exploration of trade-offs between packaging level options and architectural techniques. This paper illustrates our initial efforts towards these goals with an example of architectural trade-offs that can be made to optimize the use of MCM packaging technology.

The following section briefly reviews the major changes in physical constraints that are provided by MCM technology, and which can impact the computing system architecture. An example of the application of this research to the design of a high performance computing node is presented in Section 3. The models used for the analysis are discussed in Section 4. Section 5 presents the results of some example trade-off studies for modern microprocessors and specific results for the Alpha 21164. The paper concludes with a status report and the direction of ongoing and future efforts.

2 MCM Packaging Technologies

Multi-chip Modules (MCM) introduce an additional level in the packaging hierarchy wherein multiple dice can be directly placed on a module substrate that is typically a laminated interconnect (MCM-L), ceramic interconnect (MCM-C), or deposited interconnect (MCM-D). The ten year goals of the newly established Packaging Research Center at the Georgia Institute of Technology include an order of magnitude reduction in MCM cost, with an attendant factor of 5-10 improvement in performance, size, reliability and chip I/Os. The major issue for system

[†]. This research was supported by the National Science Foundation under grant EEC-9402723.

designers is to understand how computing system architectures can make the best use of these improvements in MCM technology. Characteristics of MCM packaging that impact architectural trade-offs include the following [15,5,3].

Off-chip Delay

Global interconnects on MCM substrates are lower loss than intra-chip interconnects. As a result, better performance can be obtained through the use of MCM technology. A reduction in inter-chip latency is possible due to lower parasitics. This allows the repartitioning of a monolithic design across several smaller dice.

I/O Bandwidth

MCM technology supports area I/O as opposed to perimeter I/O available using single chip packaging. Lower parasitics also support higher I/O bandwidth.

Process Constraints

If monolithic designs can be cost-effectively implemented on several smaller dice, architectural components such as cache memories can be implemented in optimized IC processes.

We are interested in how these and other MCM packaging related factors affect system design. The goal of this research is to enable the exploration of the interactions between packaging technologies and system architecture.

3 Case Study: Processor-Memory Hierarchy

This type of early analysis can be illustrated by an example. The memory hierarchy is a critical component of modern high performance RISC microprocessors. While processor clock speed has continued to increase dramatically, memory speeds have grown much more slowly. This has resulted in the need for multiple levels of cache memories to enable the processor to continue to function at the maximum speed. The large difference between intra-chip and inter-chip delays, and the limited number of I/Os available in modern packaging technology has promoted larger dice and migration of the cache hierarchy onto the die. For example, the 300 MHz 21164 Alpha processor[4] has 8 KB level 1 (L1) data and instruction caches *and* a 96 KB level 2 (L2) cache on chip. The resulting die is 18 mm square and is manufactured in 0.5 CMOS technology. As die sizes increase, yields drop, costs rise and the high resistivity of the aluminium interconnect causes intra-chip delays to become significant.

Several recent studies have begun to examine the impact of the MCM technology on the memory hierarchy [5,3,13]. Consider the options that would become available with a large number of I/Os and dramatically reduced cost of MCM manufacturing. With off-chip delays no longer dominant, chip boundaries may be re-drawn to provide better trade-offs in cost and performance. Specifically, consider moving the L2 cache off-chip in the above example of the DEC Alpha processor, resulting in the following trade-offs.

- This partitioning will result in smaller die for the processor (logic) which leads to higher yields and hence lower cost.
- An SRAM process may be used for the L2 cache rather than a logic process, leading to a denser, faster design.
- The reduced processor die cost may enable a larger L2 cache, which improves performance via a higher cache hit rate. This improvement may compensate for any nominal increase in L2 access times due to off-chip delays.
- Several smaller die vs. one large die produces a dilation in the distribution of the thermal energy generated by the devices.
- The size of the MCM substrate increases as a function of the die footprint, increasing the substrate cost.
- The increased number of I/Os due to partitioning at the L1/L2 interface might add to the die testing costs.
- The increased number of I/Os may also increase the MCM substrate testing costs.

The above are examples of the types of architectural trade-offs that become important with the new MCM packaging technology. Our goal is to be able to perform such trade-offs during conceptual design. This paper specifically focuses on the trade-offs between on-chip versus off-chip L1 and/or L2 caches. The cache performance numbers used in the following analysis are published figures [8].

4 Packaging Models

The following describe the models used to study the trade-offs reported here.

4.1 Cost Model

The cost model is similar to those developed elsewhere, modified to reflect the technology goals of our center [5,14]. Equation (1) estimates the cost of a single chip module (SCM) based on the wafer cost, logic process yield, memory process yield, packaging cost per I/O, and the testing cost.

$$C_{SCM} = \frac{C_{wafer} + C_{test}}{N_{yielded}} + C_{package}, \quad (1)$$

where C_{wafer} is the base cost of the wafer including the processing costs, C_{test} is the cost of testing all the yielded SCMs, $C_{package}$ is the cost of the package used for the SCM, and $N_{yielded} = numup \times Y_{logic} \times Y_{mem}$. The relevant parameters [5] used for the model are shown below.

- Cost of processed wafer = \$5000
- Cost of Packaging = \$0.10 per I/O
- Cost of testing SCM = \$0.10 per I/O
- Cost of testing bare die = \$0.15 per I/O
- Cost of substrate = \$50 per in²
- Cost of test and assembly = \$15 per module

The total number of dice obtained from a wafer, known as *numup*, is determined by (2).

$$\frac{\pi D^2}{4A} - \frac{\pi D}{\sqrt{2}A} - 4 \quad (2)$$

The logic die yield is given as

$$Y_{logic} = e^{(-A_{logic}\delta)}, \quad (3)$$

while the memory die yield is given as

$$Y_{mem} = e^{(-A_{mem}\delta)} + A_{mem}\delta e^{(-A_{mem}\delta)} + (A_{mem}\delta)^2 e^{\left(\frac{-A_{mem}\delta}{2}\right)}. \quad (4)$$

In the above expressions, D represents the diameter of the wafer, A represents the area of the die, and δ represents the defect density. Note that memory die yields are superior to logic die yields. This is due to the use of well known techniques for using redundant columns of cells to improve the yields in the presence of faults.

The MCM cost is determined based on the cost of the bare die, testing cost for each die, and the cost of manufac-

turing, testing, and assembling the MCM substrate. The expression for the MCM cost is given by (5).

$$C_{MCM} = \sum_{N_{mcm}} C_{die} + C_{substrate} + C_{test} \quad (5)$$

4.2 Interconnect Delay Model

Global interconnects within a modern electronic system exist at two levels: within a single chip (*intra-chip* interconnects) and within the packaging medium connecting multiple chips (*inter-chip* interconnects). Our analysis of pulse propagation in both types of interconnect follows that in [1]. Inter-chip interconnects on a typical MCM substrate are characterized by low-loss dielectrics and by conductors with low resistivity (e.g., copper) and large cross-section, making losses due to line resistance and shunt conductance negligible in the delay model. This allows inter-chip interconnects to be modeled as lossless, ideal transmission lines. For global interconnects within a chip, the line resistance cannot be ignored when it is comparable to or larger than the resistance of the device driving the line. The resistance of global on-chip lines becomes significant as feature size is scaled down and die size is scaled up. Because the resistance of an on-chip interconnect usually dominates its inductance, it can be modeled as a distributed RC line. The time required for the output of the line to attain 50% of the input voltage step is given by $0.4rc_l^2$, where r_{int} and c_{int} are the resistance and capacitance per unit length and l is the total interconnect length.

To compare the costs of inter-chip and intra-chip communication, the delay models cited above are shown in Figure 1 in the context of practical driver-receiver circuits.

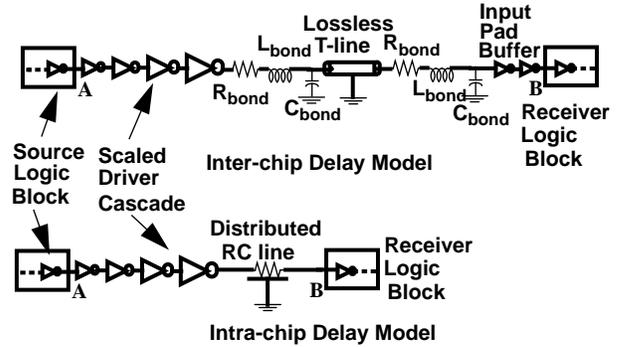


Figure 1. 50% Delay Models for Inter-chip and Intra-chip Interconnects

In each circuit, a minimum-sized CMOS inverter within a source logic block produces a signal that must be transmitted to a receiver logic block via an interconnect. The output of the source is amplified by a cascade of optimally-

sized drivers. In the inter-chip delay model, the source and receiver are on separate chips. The interconnect between the chips is modeled as a lossless transmission line with a specified time-of-flight delay and characteristic impedance. At each end of the line, lumped RLC elements are used to model the parasitics associated with connections between the chip and the next level in the packaging hierarchy. Assuming the die is attached directly to the chip carrier, the chip-to-package connection could represent either a wire-bond or a solder bump bond. The transistors driving the output pad are sized so that their driving resistance matches the characteristic impedance of the transmission line. Driving an off-chip interconnect in this way decouples rise/fall times at the driven end from the total capacitance of the line and allows signal propagation to occur at the speed of light. In the intra-chip model, the source and receiver are on the same chip and are connected by a global interconnect modeled as a distributed RC line. Although the intra-chip signal path avoids the package parasitics in the inter-chip delay model, the signaling delay is quadratic in the interconnect length implying that intra-chip delays can actually exceed inter-chip delays for long lines.

5 Results

Cost and interconnect analyses were performed for the Alpha 21164, MIPS R10K, PowerPC 604, and the PowerPC 620. In each case, the system was re-partitioned along the processor to memory hierarchy interface. For the MIPS and PowerPC implementations, this involved moving the L1 caches off-chip. For the Alpha 21164, two alternatives were examined: moving only the L2 cache off-chip, and moving both the L1 and L2 caches off-chip.

The cost analysis was based on a defect density of $0.9/\text{cm}^2$. The cost comparison takes into account the cost of testing and packaging the SCM, and the costs of the substrate, and test and assembly for the MCM. It should be noted that partitioning at the cache boundary results in an increase in the number of I/Os required in the logic and memory portion of the microprocessor. As a result the die used on the MCM are assumed to be area bonded. The comparison of costs for the microprocessors when packaged as an MCM instead of SCM is shown in Figure 2a. We note that there is a cost advantage resulting from the re-partitioning in all the processors except the PPC604. This is because most of the savings are derived from the area reduction in the logic die when the L1 cache is moved off the processor die. The area of the L1 cache in PPC604 is relatively small, hence reducing the benefits obtained.

A cost/performance analysis using an Alpha 21164 as the base case and varying L1 cache sizes was also performed. Since memory traces for these relatively recent processors were unavailable, we used results from Jouppi et al. [8] where the effect of varying cache sizes is presented in terms of the impact on the average time per instruction (TPI), the average time to execute an instruction for the SPEC benchmark traces. The cost was computed for the SCM and MCM implementations using the models in Section 4.1. The cost/performance results are shown in Figure 2b. As expected, we see advantages for moderate to large cache sizes. For smaller cache sizes the increase in the area of the processor die is not large enough to offset the costs of an increased number of I/Os, and the

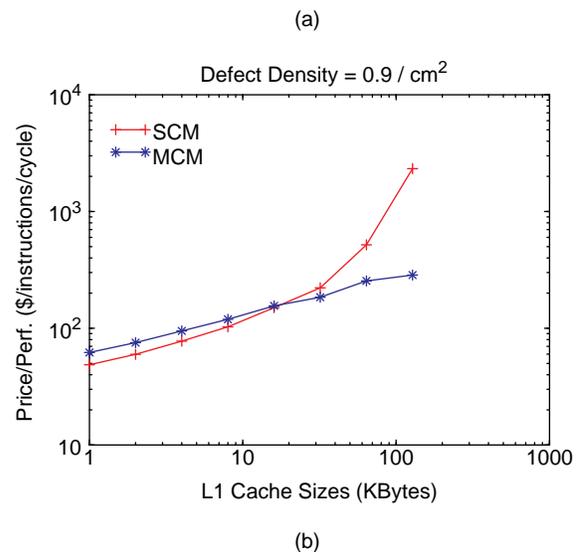
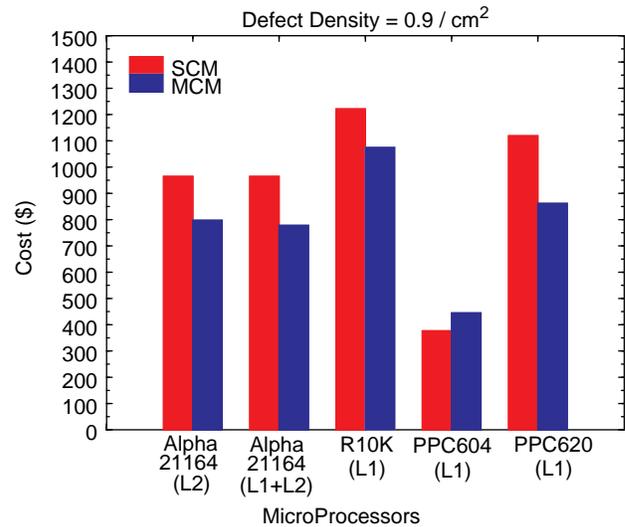


Figure 2. (a) Comparison of SCM and MCM costs for modern microprocessors; (b) Comparison of SCM and MCM Price/Performance

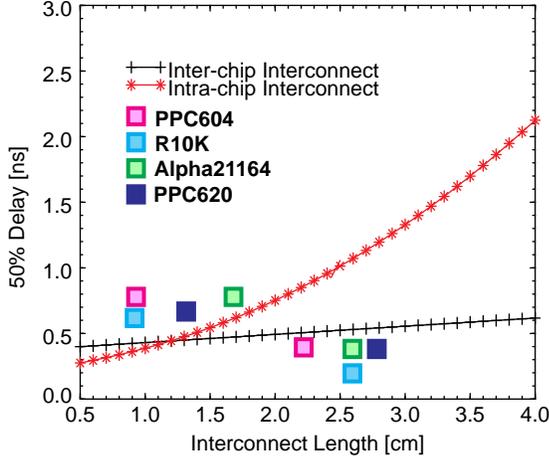


Figure 3. Comparison of intra-chip and inter-chip interconnect delays.

cost of a larger substrate. As cache sizes increase, the area increase in the microprocessor die becomes significant with significant reduction in the yields leading to increasing cost. We observe that the crossover point occurs when the cache size is approximately 20 KB. Most modern microprocessors use L1 caches of size 32 KB, which favors the MCM implementation.

Since on-chip delay increases more rapidly than off-chip delay with longer interconnects, a monolithic solution does not always represent the best cost-performance trade-off. To illustrate this point, analytical approximations of intra-chip and inter-chip 50% delays from point A to point B (Figure 1) are plotted in Figure 3 as a function of interconnect length. The delay equations are derived from expressions given in [1].

$$t_{\text{intra}}(l) = t_{\text{driver}} + 0.4r_{\text{int}}c_{\text{int}}l^2 + 0.7r_{\text{int}}C_{\text{rev}}l \quad (6)$$

$$t_{\text{inter}}(l) = t_{\text{driver}} + 1.4r_{\text{min}}c_{\text{min}} + \frac{l}{c_o}\sqrt{\epsilon_r} \quad (7)$$

In (6) and (7), r_{int} and c_{int} are the resistance and capacitance of the distributed RC line per unit length, C_{rev} is the input capacitance of the receiver circuit, r_{min} and c_{min} are the driver resistance and gate capacitance of a minimum-size inverter, t_{driver} is the delay through the driver cascade (approximately 0.3 ns for both model), and l is the interconnect length. The curves in Figure 3 were generated using device parameters from a 0.5 micron 3.3V process. On-chip interconnects have a height of 1 micron and a width of 2 microns, yielding a r_{int} value of $140\Omega/\text{cm}$. Given the effects of fringing fields, a limiting value for c_{int} of 2 pf/cm is used [1].

As Figure 3 shows, the break-even interconnect length is approximately 1.25 cm, i.e. signal paths longer than 1.25 cm should be routed via the MCM substrate. However, on-chip interconnects in a monolithic system will typically be shorter than off-chip interconnects in a partitioned, MCM-based system with identical functionality. In Figure 3, the cluster on the left indicates the signal path lengths for the monolithic implementation of 4 commercial microprocessors. The cluster on the right indicates the corresponding inter-chip lengths when the caches are moved off-chip in the MCM solution. For the Alpha21164, the interconnect between the L1 and L2 caches is the worst-case length. For the other systems, the worst-case interconnect length is between the fetch unit and the L1 cache. Figure 3 shows that the worst-case delays are comparable for the PPC604, R10K, and PPC620 systems. The delay for the partitioned Alpha system is significantly lower than the delay for the monolithic Alpha implementation. As future processors become increasingly complex and larger, and MCMs become less expensive, the MCM solution should become increasingly effective.

To obtain more accurate results for the Alpha, HSPICE models using the process parameters given above were constructed for the circuits in Figure 1. The package bond parasitics used in the simulations were typical of flip-chip C4 bonding process. These simulations show a significant performance benefit using separate die for the processor and the second level cache. The 50% propagation delays from point A to point B in Figure 1 are $t_{\text{PH}} = 0.54$ ns and $t_{\text{PL}} = 0.60$ ns for the partitioned system compared to $t_{\text{PH}} = 0.68$ ns and $t_{\text{PL}} = 0.75$ ns for the monolithic system. The rise and fall times at the receiver input (point B) for the partitioned system ($t_{\text{RISE}} = 0.04$ ns and $t_{\text{FALL}} = 0.06$ ns) are also much smaller than those in the monolithic system ($t_{\text{RISE}} = 0.98$ ns and $t_{\text{FALL}} = 1.06$ ns). Further simulation details can be found in [6].

Figure 4 shows that the MCM cost increases linearly with the die test costs and quadratically with the defect density. Hence, there is a need for extensive, detailed models of testing costs (both MCM substrate and die) to accurately predict MCM costs.

6 Conclusion

The goal of this work is an understanding of the impact of MCM packaging technology on system design. Preliminary results suggest that MCM technology can be exploited to realize a new class of cost effective system designs. The partitioning of the memory hierarchy is sensitive to packaging technology. MCM packaging can

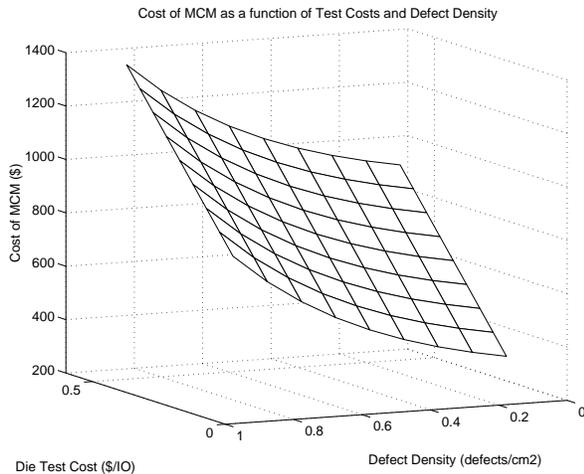


Figure 4. Sensitivity of MCM costs to die test costs and defect density.

result in a large monolithic die being replaced by several smaller dice. This reduction in individual die size leads to improved yield. The separation of memory and digital logic allows the use of optimized IC processes. With the L2 cache on a separate die, it can be enlarged to further improve system performance. Partitioning may also help in addressing thermal management issues. These advantages are achieved at the expense of the added cost of a larger substrate, increased test and assembly of the MCM and the wafer bumping. The right balance is established by system performance goals. We are currently incorporating these models into IMPACT: a set of modeling and analysis tools for assessing the impact of MCM technology on the design of the next generation of computing systems.

As MCM technology advances will become an increasingly favorable option for a wide range of applications. We hope to facilitate this process with early analysis tools that can reliably predict the impact of packaging options on system level metrics.

7 Acknowledgments

The authors gratefully acknowledge helpful discussions with Dr. Peyman Dehkordi and Dr. Vivek De, and their many useful suggestions, particularly with respect to modeling and analysis techniques.

8 References

[1] Bakoglu, H.B., *Circuits, Interconnections, and Packaging for VLSI*, Addison-Wesley Publishing Company, Inc., 1990.

[2] Bowhill, W.J. et al., "Circuit Implementation of a 300-MHz 64-bit Second-generation CMOS Alpha CPU," *Digital Technical Journal*, June 1995.

[3] Dehkordi, P., Ramamurthi, K., and D. Bouldin, D., "Early Cost/Performance Cache Analysis of a Split MCM Based MicroSparc CPU," *Proceedings of the MCM Conference*, February 1996.

[4] Edmondson, J.H. et al., "Internal Organization of the Alpha 21164, a 300-MHz 64-bit Quad-issue CMOS RISC Microprocessor," *Digital Technical Journal*, June 1995.

[5] Franzon, P.D., Stanaski, A., Tekmen, Y., and Banerjia, S., "System Design Optimization for MCM," *Technical Report NCSU-ERL-94-16*, North Carolina State University, 1994.

[6] Garg, V. et.al., "Impact of Packaging Constraints on System Design: A Case Study of the Memory Hierarchy", Technical Report CSRL-95/08, School of Electrical and Computer Engineering, Georgia Institute of Technology, September 1995.

[7] IBM, *Packaging Technology: Product Description Literature*.

[8] Jouppi, N.P. and Wilton, S.J.E., "Tradeoffs in Two-Level On-Chip Caching," *Proceedings of International Symposium on Computer Architecture*, Chicago, Illinois, April 18-21, 1994.

[9] MicroModule Systems (MMS), MCM-D Process Parameters.

[10] Mips, *R10000 Microprocessor Product Overview*, October 1994.

[11] Motorola, *Technical Brief on PowerPC 604 and 620*.

[12] Mulder, J.M., Quach, N.T., and Flynn, M.J., "An Area Model for On-Chip Memories and its Application," *IEEE Journal of Solid State Circuits*, 26(2):98-106, Feb. 1991.

[13] Roberts, J.D. and Dai, W.W.-M., "Early System Analysis of Cache Performance for RISC Systems: MCM Design Trade-Offs," *Technical Report UCSC-CRL-92-02*, University of California, Santa Cruz, March, 1992.

[14] Sandborn, P.A., Abadir, M.S., and Murphy, C.F., "The Trade-off Between Peripheral and Area Array Bonding of Components in Multichip Modules," *IEEE Transactions on Component, Packaging, and Manufacturing Technology*, 17(2):249-256, June, 1994.

[15] Sandborn, P.A. and Moreno, H., *Conceptual Design of Multichip Modules and Systems*, Kluwer Academic Publishers, 1994.

[16] Tummala, R.R. and Rymaszewski, E.J., Eds., *Microelectronics Packaging Handbook*, Van Nostrand Reinhold, New York, 1989.