

Early Analysis of Cost/Performance Trade-Offs in MCM Systems

Vivek Garg, *Student Member, IEEE*, Darrell J. Stogner, Craig Ulmer, David E. Schimmel, *Member, IEEE*, Chryssa Dislis, *Member, IEEE*, Sudhakar Yalamanchili, *Senior Member, IEEE*, and D. Scott Wills, *Member, IEEE*

Abstract—This paper explores early analysis of the complex relationships between system architectures and the active and packaging materials from which they are implemented. The goals of this analysis are to enable the designer to specify cost effective technologies for a particular system and to uncover resources which may be exploited to increase performance of such a system, early in the design process. We describe a prototype tool called IMPACT, which will predict cost, performance, power, and reliability, and present several case studies demonstrating its use.

Index Terms—Advanced packaging, early analysis, MCM, system trade-offs.

I. INTRODUCTION

THE RESOURCE requirements of current generation integrated circuit technology are close to exceeding capabilities of traditional packaging techniques. Critical packaging resources include input/output (I/O) bandwidth, off-chip signal transmission time, system footprint, and mass. Multichip modules (MCM's) utilize chip scale packaging (CSP) techniques to eliminate the intermediate level package, enabling direct placement of dice on a substrate, which contains the interconnections to realize circuit connectivity.

This paper explores the interaction between the resources provided by the packaging technology and the system architecture. It has been reported that decisions made very early in the design cycle have a significant impact on implementation and expenses incurred in the development of a product [23]. We address the concept of *early analysis* which allows a designer to assess the effects of technology-based decisions on the cost, performance, reliability, and power metrics of a system. Early analysis is essential to provide designers with the ability to very rapidly evaluate architectural alternatives without actually implementing the system. The architecture may be specified structurally at a very abstract level, and coupled with the desired technology to compute the various

metrics which evaluate the design. The goal of early analysis is to shorten the design cycle and thereby the time to market.

The foundations of our approach are found in the rich legacy of system modeling paradigms and software. We have drawn on models developed by others including [2], [6], [11], and [23]. Software which has been developed to explore some aspect of systems includes work at Stanford, Cornell, CMU, IBM, and MCC [2], [3], [19], [20], [23].

This research is most closely related to the work of Sandborn *et. al.* in developing their multichip systems design advisor (MSDA) tool. We have similar objectives in evaluating "what if" scenarios to explore architectural alternatives. However, one important difference is in the level of abstraction allowed/required by the respective approaches. Our intention is to explore the limits of high level abstraction in modeling and interpreting the interactions between architecture and packaging technology. To this end, we specifically do not support or require lower level design information such as detailed netlists. Instead, we have striven to define the greatest degree of abstraction from which we can obtain meaningful results. An additional distinguishing feature of our approach is the use of model hierarchies. Depending on the completeness of the specification provided by the system architect, our software tool will choose a particular approach to modeling, say, substrate area, or test generation costs. If more refined or complete information is later provided, it is incorporated into the estimates. In this process the user is guided to required and optional parameters that may be specified.

The remainder of this paper is organized as follows. In Section II, we consider the effects of advanced packaging technologies on systems in general. In Section III, we detail the organization of the IMPACT tool, and in Section IV we give an overview of the models implemented in IMPACT. Several architectural studies are presented in Section V, followed by conclusions in Section VI.

II. IMPACT OF MULTICHIP MODULES ON SYSTEM DESIGN

The design of computing systems is fundamentally a process of optimizing system parameters subject to a given set of physical constraints. The advent of MCM technology is effecting a tremendous change in these physical constraints with consequent redefinition of the appropriate package and die boundaries. Our goal is the early analysis of the impact of these constraints so that the system may be designed to make the most effective use of MCM technology. The main effects of

Manuscript received February 18, 1997; revised May 13, 1997. This paper was presented in part at the IEEE International Conference on Innovative Systems in Silicon, Austin, TX, October 9–11, 1996. This research was supported in part by the National Science Foundation under Grant EEC-9402723.

V. Garg, C. Ulmer, D. E. Schimmel, S. Yalamanchili, and D. S. Wills are with the Packaging Research Center, School of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta, GA 30332-0250 USA.

D. J. Stogner is with the Mobile Communications Division, Motorola, Plantation, FL 33322 USA.

C. Dislis is with the Department of Cybernetics, University of Reading, Reading RG6 2AY, U.K.

Publisher Item Identifier S 1070-9894(97)05868-4.

using MCM technology in the design of system architectures are the following.

A. Performance

The elimination of intermediate level packaging implies that the chip is directly attached to the MCM substrate using either wirebond or controlled collapsible chip connection (C4). Modern architectures are limited by off-chip delays. The parasitics associated with an off-chip interconnect in a traditional package typically include the resistance, capacitance, and inductance of the wirebond used to attach the die to the package, and the resistance, capacitance, and inductance of the brazed pins of the through hole or surface mount package. These parasitics are several times larger than those encountered in direct chip attach technologies. The C4 parasitics are an order of magnitude less than those associated with through hole and surface mount packages [27]. This reduction in parasitics encountered by the signals going off-chip improves the signal transmission speeds. Furthermore, the conductors on the substrate behave like transmission lines resulting in faster signal transmission times than on-chip interconnects, further contributing to the performance of MCM based systems.

Incorporating the components of a printed circuit board onto an MCM substrate typically results in a reduction in footprint. This helps reduce the interconnect lengths between components, reducing the transmission time between them. Furthermore, for some MCM technologies the dielectric constant of the substrate is much smaller than that of FR4 or other materials used in the printed circuit board technology, resulting in faster signal propagation. As monolithic designs become large, the on-chip aluminum interconnect delays become larger than off-chip interconnect delays on the substrate, providing further impetus to move to MCM implementations.

B. Cost

MCM packaging has been predominantly limited to high performance applications, where cost is not the primary issue. When MCM's are used in such "niche" applications, they are not manufactured in high volume, and as a result are not economically competitive with printed circuit board (PCB) technology. The use of MCM's in automotive applications has proven that when manufactured in high volume, they can indeed be cost-effective. Other factors adding to the cost of MCM's are testing of unpackaged (bare) dice to ensure correct functionality, also known as the known good die (KGD) issue, and signal redistribution on existing dice designed for peripheral I/O to enable area array bonding. Most of these costs are related to the fact that MCM's are not widely used in applications, and as a result the process and equipment involved in its manufacturing is not cost competitive with other more mature technologies. Active research involving both industry and academia is focusing on the goal of making MCM technologies price competitive with printed circuit board technology. Several low-cost processes have been introduced and others are being proposed which may lead to a reduction of substrate costs by a factor of five and ten, respectively [12], [24], [26], [28].

Lower off-chip delays incurred in MCM technology encourages the partitioning of large monolithic dice into smaller ones. Smaller die sizes result in higher yields, and as a result decrease the cost of dice. This decrease in cost can be used to offset some of the costs incurred in the use of MCM packaging. Ultimately we expect the MCM process to mature to a point where it is price-competitive with the PCB technology. At that time, partitioning of the dice will reduce the cost of the system, possibly without appreciably affecting the performance.

C. System Yield and Reliability

System yield is a function of the yield of the individual components. System yield for MCM's is the product of the yield of all the dice, substrate, and the bonding process. Thus, system yield can be uneconomically low for complex MCM's, unless particular attention is paid to test coverage and delivered die yield for bare dice, mainly through KGD methods. Characterizing and testing bare dice is a more expensive undertaking than providing the equivalent quality levels for packaged dice. However, if the problems of low system yield can be overcome, the use of an MCM implementation may increase the reliability of the system. It has been reported that the reliability of C4 die attach mechanism is 0.5 ppm which is six times more reliable than wirebond (3 ppm) [10]. Consequently, use of C4 type die attach will increase the reliability of a given system. Furthermore, elimination of the intermediate packaging removes reliability concerns related to these components. However, this elimination leads to other problems involving thermal [coefficient of thermal expansion (CTE)] mismatches between the silicon dice and the MCM substrate. Epoxy encapsulants are often used along with C4 connections to minimize thermal mismatch problems.

D. Power

The lower parasitics of the C4 or wirebond connections, as compared to brazed pins, result in smaller signal drivers for the same level of performance, leading to lower power consumption. Concentration of components closer together on the MCM substrate may give rise to more challenging thermal management problems. However, MCM-C substrates are typically better conductors of heat than printed circuit boards, and may compensate for the increased heat flow.

E. Ergonomics

After performance, the most significant drivers for MCM's are low volume and mass. With the recent explosion in portable consumer electronics such as mobile telephones, personal digital assistants, laptop computers, etc., the industry is constantly striving for smaller and lighter products. An MCM package involves bare dice, discrete components, and a substrate which houses the wiring required to connect all the circuits to be placed on the package. As a result, the MCM package can incorporate the functionality of a PCB by replacing the board with the substrate, and by attaching bare dice and discrete components directly on the substrate, reducing the overall size and weight of the system. Due to the elimination of the intermediate level packages, the effective

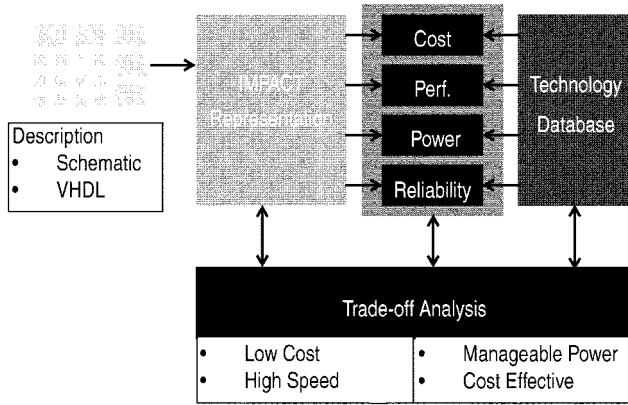


Fig. 1. Organization of IMPACT tools.

usage of the MCM substrate increases, since a die typically comprises of only 20% of the package area. This implies that larger printed circuit boards can be reduced to small MCM substrates, resulting in significant reduction in system footprints.

III. IMPACT METHODOLOGY

The Packaging Research Center (PRC) is an NSF sponsored Engineering Research Center established at the Georgia Institute of Technology. The center is multidisciplinary in nature with the common goal of developing economically viable MCM technology for consumer applications. As part of this effort, we are developing the IMPACT modeling tools to help designers perform early analysis of the impact of MCM packaging on system architectures. These tools are based on a set of core models that capture technology parameters for substrate, die, packaging, and assembly, and system specifications and enable the assessment of cost and performance related metrics such as MCM footprints, die yields, test costs, etc. A user specified architecture can be evaluated based on cost, performance, reliability, and power metrics computed using IMPACT. At present the models that are being used are from the published literature. In the future, we will incorporate models from the Packaging Research Center's manufacturing process as they become available. The models are fully interchangeable, so that designers may use technologies and processes of choice to evaluate their designs. The core set of models include cost models for die, substrate, assembly, and test. Models for on-chip and off-chip interconnects provide signal transmission delay estimates. Reliability and power dissipation models are under development.

Fig. 1 shows the organization of the IMPACT tools. The design may be entered as list of dice, a schematic, or as a description in the VHSIC Hardware Description Language (VHDL). In either case a structural description of the design is extracted to perform the analysis. It should be noted that the design is specified at a very abstract level in terms of computational units, memory units, and information channels. Each of these components have attributes associated with them to describe their physical features and functional specifications. Currently the designer is expected to provide the total number of functional units in the system, identify them as memory or

logic die, and assign them to partitions, where each partition represents a package entity, such as an MCM. The connectivity between these functional blocks is represented via information channels, which simply indicate the rate of information flow between the corresponding blocks. Information channels place constraints on the number and performance of I/Os within a die.

Once design partitions have been established, the cost and performance metrics for the system can be calculated. Other configurations of the system may be evaluated to achieve the desired specifications. An iterative process may be used to identify the appropriate technology—MCM-L, C, or D, C4 or wirebond, stacked die, etc.—to satisfy design specifications, and optimize the design space. The trade-offs to be considered include cost-effectiveness, performance, price/performance, reliability, and thermal management.

IV. IMPACT MODELS

The IMPACT tools consist of a hierarchy of models for cost, size, and various other metrics related to the dice and MCM substrates required for a user specified system. A global view of the models implemented in IMPACT is shown in Fig. 2. These models range from architectural level to process level, and they are applied hierarchically so that parameters may be specified by the user at any level. For example, consider an application-specific integrated circuit (ASIC). If the cost of the die is known, these costs may be readily used. However, if the cost is not known, lower level models of the die yield and fabrication cost are invoked to compute an approximate value of the die cost. This methodology allows for the maximum amount of flexibility for the designers using IMPACT. In this section, we present an overview of some of these models, and how they are used in IMPACT.

A. Die Models

The cost of a die is dependent on its size, the process used to fabricate the die, the size of the wafer, and probe and parametric tests conducted to validate it.

The size of a die can be determined by Donath's model which is based on Rent's rule. Donath's model relates the average wire length in units of gate pitch, R_m , to the number of gates, N_g , on the die and the parallelism factor, p , corresponding to the architecture on the die [2]

$$R_m = \frac{2}{9} \left(\frac{7N_g^{p-0.5} - 1}{4^{(p-0.5)} - 1} - \frac{1 - N_g^{(p-1.5)}}{1 - 4^{(p-1.5)}} \right) \times \frac{1 - 4^{(p-1)}}{1 - N_g^{(p-1)}}, \quad 0 < p < 1. \quad (1)$$

The gate dimension and chip size can then be calculated as

$$d_g = \frac{f_g R_m p_w}{e_w n_w} \quad \text{and} \quad x = \sqrt{N_g} d_g, \quad \text{respectively} \quad (2)$$

where f_g is the average fanout of a gate, p_w is the wiring pitch, e_w is the wiring efficiency, and n_w is the number of wiring layers.

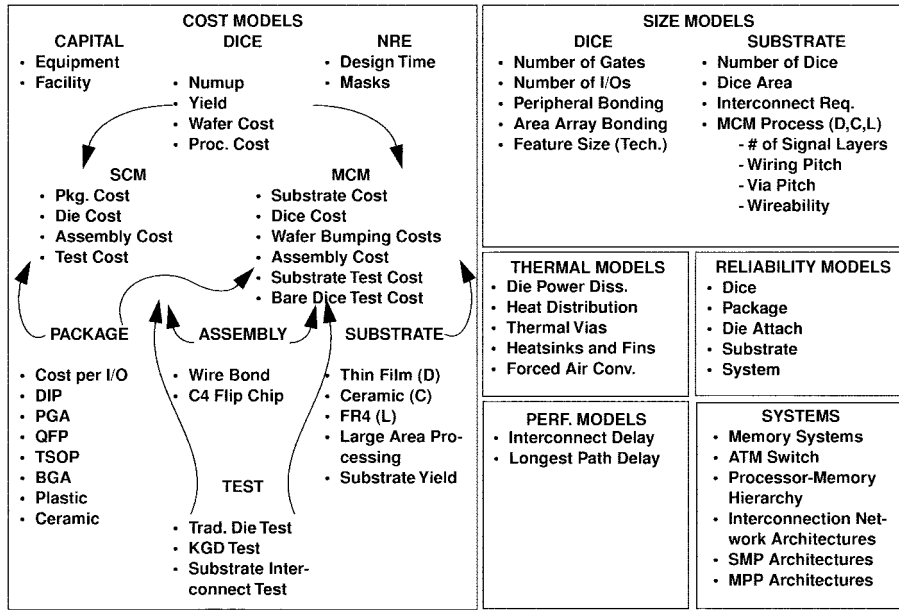


Fig. 2. Model flow in IMPACT tools.

Rent's rule relates the number of I/Os on a die to the number of gates [2]

$$N_{I/O} = \beta N_g^\chi, \quad (3)$$

It is an empirical rule parameterized for four different types of systems, each corresponding to different values for β and χ . They are (1.9, 0.5), (3.2, 0.434), (0.82, 0.45), and (7.0, 0.21) for complementary metal-oxide-semiconductor (CMOS) gate arrays, multiple integrated circuit (IC) designs, microprocessors, and functionally complete chips, respectively.

Once the size of the chip, x , is determined the number of dice fabricated from a wafer of diameter, D , is given by

$$N_c = \frac{\pi D^2}{4x^2} - \frac{\pi D}{\sqrt{2}x^2} - 4. \quad (4)$$

The chip yield, which is dependent on the type of circuitry on the chip is computed as the product of the logic and memory die yields which in turn are computed as

$$\begin{aligned} Y_{\text{logic}} &= e^{-A_{\text{logic}}\delta}, \\ Y_{\text{mem}} &= e^{-A_{\text{mem}}\delta} + A_{\text{mem}}\delta e^{-A_{\text{mem}}\delta} + A_{\text{mem}}^2\delta^2 e^{-(A_{\text{mem}}\delta/2)} \end{aligned} \quad (5)$$

where δ is the defect density of the wafer, and A_{logic} and A_{mem} are die area dedicated to logic and memory circuitry respectively. The number of good dice yielded from a wafer is given by

$$N_y = N_c Y_{\text{logic}} Y_{\text{mem}}. \quad (6)$$

The number of fabricated dice and yield models can be found in [11].

The die cost can then be computed by simply computing the ratio of the wafer processing cost for the die process and the number of yielded die

$$\text{Pre-testing Die Cost} = \frac{C_{\text{wafer}} + C_{\text{proc.}}}{N_y}. \quad (7)$$

There are several types of testing which may be conducted to verify the functionality and reliability of the dice. These are probe tests, parametric tests, and bare die test. Each test phase results in removal of dice from the yielded set, consequently adding to the cost of a die. The testing models are addressed in Section IV-C.

B. Substrate Models

In this section, we present the models related to the size and cost of MCM substrates [23]. The size of an MCM substrate depends on several parameters ranging from the type of technology being used (L,C,D) to the number of dice on the MCM, and even the thermal conductivity of the substrate. There are several wiring limited models available to determine the size of the MCM substrate for a specific design. However, the size of the substrate may not be necessarily be constrained by the wireability. Other factors constraining the size of the MCM substrate are the number of I/Os on the MCM, number of vias in the substrate, number of dice on the MCM, and thermal dissipation of the substrate. All of these constraints present a complex interrelationship which must be resolved by simply evaluating each of the models and then determining the constraining factor.

The interconnect capacity, I_c , is given as

$$I_c = \frac{n_w}{p_v} \quad \text{and} \quad I_c = \frac{1 + T_c}{p_v} n_w, \quad (8)$$

where T_c is the tracks per channel and p_v is the via pitch, for designs without and with vias respectively. The interconnect capacity is essentially a measure of the available resources for a specific size substrate.

There are three wireability based models used for determining the size of the MCM substrate. These are Seraphim's model, Bakoglu's model, which is an extension of Donath's model, and finally Hannemann's model which has been shown to be closely correlated with Bakoglu's model.

Seraphim's model is an extremely simple model which assumes that the chips are placed on the MCM substrate with a chip pitch, F_p , average connection length of $1.5F_p$, and an average fanout of 1.5. The substrate area is given by

$$A_{\text{sub}} = \frac{2.25N_{\text{chip}}F_pN_{\text{I/O}}}{2c_wI_c}. \quad (9)$$

While Seraphim's model is simple, the assumptions limit its application.

Bakoglu's model enables the computation of the average length of an interconnect on the MCM substrate in units of chip pitch

$$R_m = \frac{2}{9} \left(\frac{7N_{\text{chip}}^{(p-0.5)} - 1}{4(p-0.5) - 1} - \frac{1 - N_{\text{chip}}^{(p-1.5)}}{1 - 4(p-1.5)} \right) \times \frac{1 - 4(p-1)}{1 - N_{\text{chip}}^{(p-1)}}, \quad 0 < p < 1. \quad (10)$$

The MCM substrate area is determined by the product of the number of dice and the square of the chip footprint, F_p

$$F_p = \frac{f_c}{(f_c + 1)} \frac{N_{\text{mcm}}R_m}{N_{\text{chip}}c_wI_c}, \quad A_{\text{sub}} = N_{\text{chip}}F_p^2 \quad (11)$$

where f_c is the average fanout of the chip, and N_{mem} is the total number of chip I/Os and I/Os to/from the MCM.

Hannemann's model of the substrate is given by

$$A_{\text{sub}} = \left(c \frac{bN_{\text{mcm}}}{n_w\sqrt{N_{\text{chip}}}} \right)^2, \quad b = \frac{n_w}{I_c}, \quad c = \frac{f_c}{(f_c + 1)} \frac{R_m}{c_w} \quad (12)$$

where b is the feature size parameter, and c is the correlation between the Bakoglu and Hannemann models. For a correlation factor of 3.9, the Hannemann model provides a good approximation to Bakoglu's model for modules with 10–30 dice and an average net fanout of 1.5–2.

The MCM substrate may be limited by the number of I/Os going off the module. The constraints for peripheral and area I/O are

$$A_{\text{sub}} = \left(\left(\frac{N_{\text{I/O}}}{4} + 2 \right) p_{pb} \right)^2 \quad \text{and} \quad A_{\text{sub}} = N_{\text{I/O}} p_{ab}^2 \quad (13)$$

respectively, where p_{pb} is the peripheral pad pitch and p_{ab} is the area array pad pitch.

Via limited footprint is given by

$$A_{\text{sub}} = \frac{N_{\text{via}}}{c_v\rho_v} + A_{\text{unusable}} \quad (14)$$

where N_{via} is the number of vias, c_v is the via efficiency, ρ_v is the via density, and A_{unusable} is the area that may not be used for vias.

The MCM substrate may often be limited by the area of the dice on the MCM. The substrate area required for N_{chip} dice is given by

$$A_{\text{sub}} = \sum_i^{N_{\text{chip}}} (L_i + 2L_{\text{bond}} + s)(W_i + 2L_{\text{bond}} + s) \quad (15)$$

for dice attached using wire bonding. For flip-chip bonding, the length of the bond, L_{bond} , is 0. L and W represent the

length and width of the individual dice, and represents the minimum spacing required between chip placement sites.

Finally, the thermal properties of the substrate also place a constraint on the substrate size. A substrate must be capable of dissipating the heat generated by the dice assembled on it. The substrate area as constrained by thermal properties can be determined using

$$A_{\text{sub}} = \frac{P_{\text{chip}}}{\rho_p} \quad \text{and} \quad P_{\text{chip}} = \sum_i^{N_{\text{chip}}} P_i \quad (16)$$

where P_{chip} is the sum of the power of all dice on the MCM, and ρ_p is the power density supported by the substrate.

As mentioned previously, these models place conflicting constraints on the substrate size. IMPACT evaluates all the constraints due to these models, and the most limiting constraint is selected.

C. Test Models

Testing is an important factor at all levels of design hierarchy ranging from the die to the MCM. Die testing can be performed at several different stages in the design cycle. Parametric and probe testing are generally performed on the dice by the foundry before they are diced and packaged. As a result the dice failing the tests add to the cost of the fabricated dice yielded from the wafer. In addition, a burn-in cycle can be used to force early life failures to become apparent. These failures are then identified by a subsequent test, thus increasing the die quality but also increasing the cost per die.

In the cases of bare die, either the manufacturer or the designer must ensure that the bare die are known good. This process is well known in the industry as the KGD problem. To perform testing of bare dice, there are several strategies in use. Reusable carrier-based techniques require placement of the bare die into a carrier which uses a small amount of force to create a temporary contact with the bond pads on the chip. The test vectors may then be applied to perform functional testing. An alternate method for KGD testing involves making wire-bond or ball-bond contacts from the chip to a test carrier. Upon completion of the testing the wire-bonds can be shaved off with a laser, or alternatively the ball-bonds may be removed by reflowing the solder balls. While the temporary contact approach may require the dice to be repositioned to ensure good contact, it is gentler on the dice under test. Also the carrier lifetimes for the soft contact methods may be longer than those requiring hard contact. Reusable temporary carriers can also be used for the burn-in stage, although the maximum number of uses is negatively affected by elevated temperatures and thermal cycling.

MCM testing involves testing of the interconnects on the substrate, and the functional testing of the assembled module. Functional testing of the MCM is equivalent to the testing of the individual IC's on the MCM, and the testing of the overall functionality of the module. As far as testing is concerned, testing the module is approximately equivalent to testing a complex chip, with the associated problems of ensuring good observability and controllability. It is essential therefore to use design for testability methods to ensure the module is testable.

This is accomplished typically via boundary scan or BIST methods to ensure good observability and controllability. The importance of these methods cannot be underestimated, given the relatively low system yield expected for complex modules unless KGD methods can guarantee a very high fault cover (99% or more).

Dislis *et al.* [8] have created a complex cost and quality model to evaluate the cost effectiveness of MCM test strategies, incorporating KGD methods, burn in, sample testing at the die level, module test with and without the use of boundary scan, and varying rework scenarios. These models are very detailed as they track both cost and quality throughout the process and incorporate a large number of primary parameters related to test economics, and several analytical expressions (secondary parameters) to assess the effect of test methods on the cost of die population yielded from a wafer. A subset of these models has been used in IMPACT, and the test methodology invoked is presented in Fig. 3. The methodology incorporates a probe test stage, a KGD stage (using reusable temporary carriers) which can incorporate an optional burn-in stage, and a module test stage (assuming boundary scan), again with an optional burn-in test stage. A burn-in stage by itself is not meaningful unless followed by a test stage. The costs accounted for at each stage include non-recurring engineering (NRE) costs, test generation and application costs, and equipment costs. Each test stage takes in as its input the current defect level, a weak die population set (which will be depleted by burn-in if applied) and the cumulative cost (from previous stages). The output of the test stage is a newly classified population with a different (better) defect level, and higher cumulative cost. An iterative process can be utilized to continually burn-in and test the dice until an acceptable defect level has been reached.

In order to illustrate the test modeling philosophy, generic test generation, and test application models will be described, which will then be related to the test stages. The generic models are modified to take account of varying limitations inherent in each test stage.

Test pattern generation (TPG) modeling: The TPG model used here is based on a set of TPG cost models developed for ASIC's [7], and is based on the observation that the effort required to reach a fault cover of between 80% to 90% is relatively limited and can be modeled by an exponential curve; after that, harder faults remain, making test generation expensive and slow, modeled by a linear region. In practical terms, the exponential (cheap) region relates to automatic test pattern generation (ATPG), while the linear (expensive) region can be seen to relate to manual test pattern generation (MTPG). The breakpoint where MTPG has to be used relates to the complexity of the circuit, in terms of gates and sequentiality. The evaluation of automatic TPG data measured for scan based circuits showed that there is no significant correlation between the CPU time and circuit characteristics other than the gate count. The ATPG cost itself is a power of the gate count and the sequential depth of the circuit

$$atg\ cost = ktgs[gates \times (av-s + 1)]^{\exp tgs} \quad (17)$$

where $ktgs$ is a linear normalization factor, $av-s$ is the average

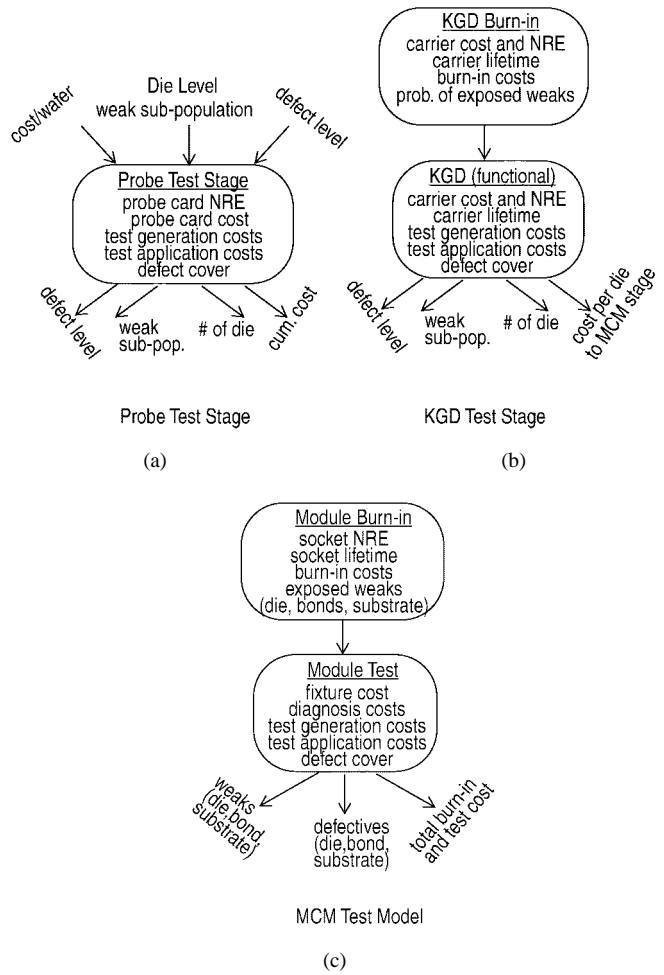


Fig. 3. Test methodology for Probe, KGD, and MCM testing.

sequential depth of the TU, and $\exp tgs$ is an exponential factor linking the automatic TPG time to the gate count, gates. The automatic test pattern generator is only run until it becomes impractical to continue, and vectors for the remaining faults, if any, are generated manually.

The fault cover achieved by the automatic TPG is calculated by

$$f_{cach} = [1 - 2.72^{(-4.6 \times sdim/av-s)}] \times [1 - 2.72^{(-4.6 \times glim/gates)}] \quad (18)$$

where $sdim$ and $glim$ are the typical sequential depth and gate count, respectively, for which the fault cover of the automatically generated test patterns remains under 99%.

The manual test pattern generation cost is related to the number of remaining faults. The MTPG stage is only invoked if the achievable fault cover is below the required target. The number of remaining faults is simply a function of the estimated total number of faults and the fault cover already achieved. The time taken for MTPG is proportional to the remaining number of faults (with an empirical proportionality factor of the average time to generate a vector per fault), and the cost is the time taken multiplied by the appropriate engineering cost rate. Algorithmic TPG, used for memory, incurs negligible cost. In this model, memory, even if part

of a particular die or testable unit, is treated as a testable unit in its own right, as it does not add to TPG cost.

Test application modeling: The test application costs for different test stages use broadly the same modeling approach. Test costs are related to the number of test vectors, which are calculated from the TPG model, on the following basis: each pattern generated by automatic TPG will detect at least one, but typically more than one, fault (the number of test vectors per fault is a primary parameter). Each of the faults not detected by automatic TPG will require one test vector to detect. (note that the number of vector per faults for manual TPG can easily be altered to reflect different circuit complexities). The number of test vectors required to test any memory on the chip is added to give the final test vector number. Test application costs relate to the use of the ATE. One element is the test time per TU (also including the setup time, t_{setup} , and related to $tpats$, the number of vectors), but the other is related to the ATE memory and the number of times it has to be reloaded. ATE charges often incorporate an extra cost for large test vector sets, over and above the actual test time. The test time per TU is given by

$$test.t.TU = \frac{tpats}{ATEspeed \times 10^6} + t_{setup}. \quad (19)$$

This test time has to be multiplied by the number of TU's under consideration to give the total test time, $test.t$. The test time in hours multiplied by the ATE cost rate is the test time related ATE cost. The ATE memory related cost per TU is given by

$$ROUNDDOWN\left(\frac{tpats}{1024 \times mem.lim} \times vec.cost\right) \quad (20)$$

where $mem.lim$ is the ATE pin memory (in kB) and $vec.cost$ is the recurring test vector cost per $mem.lim$ patterns.

Parametric and probe test: The cost of probe and parametric testing, C_{pp} , is a function of the probe test vector setup cost, $p_{vect.NRE}$, fixture costs, $fixt.c$, pattern application cost, $p_{app.cost}$, parametric test cost, $param.c$, and the manufacturing volume. The test pattern generation model is invoked, but it is assumed that test pattern generation effort is reusable. The equipment cost relates to the cost of the probe cards, and the test application process is in practice time limited, as this is a slow test application process. As a result, a high fault cover may not be achievable.

Burn-in: Burn-in was modeled at both the die stage, as part of KGD, and at the MCM stage, as part of module test. The modeling is similar in both cases. The cost of burn-in arises from handling costs, the cost of carriers (for KGD) or sockets (for MCM's), number of burn-in boards, and the equipment running costs. The handling costs are related to the time to load the burn-in boards. For die, it includes the insertion time of die into temporary carriers, and an iteration factor is included as several insertions may be required. The number of carriers required is throughput related, and both recurring and nonrecurring costs are taken into account. The number of burn-in boards is also throughput related, as enough boards are required to achieve the required throughput. Throughput was included in the burn-in modeling due to the length of

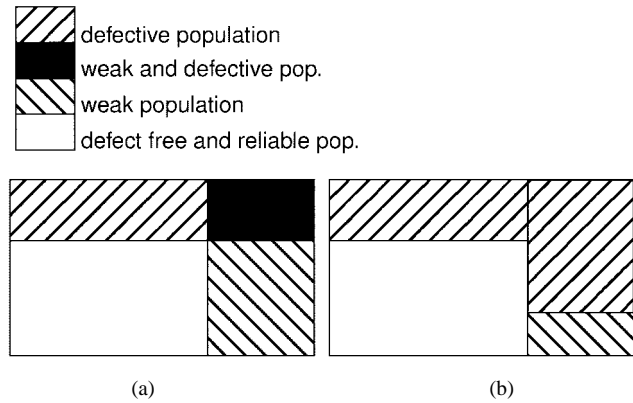


Fig. 4. Modeling weak and defective populations.

time involved, which is an obvious bottleneck in the test flow. The burn-in process forces weak units to become defective and therefore detectable by functional test. The modeling is summarized in Fig. 4. Note that the defective and weak populations are not to scale. There are initially four groups of units: fault free and reliable, defective (but not weak), weak (but not defective and therefore not detectable by functional test), and defective and weak at the same time. After burn-in, most of the weak units will have joined the defective population, which then forms the population which will be depleted by the test stage. The model governing the expected number of weak units which will be forced to fail, and which is related to temperature and burn-in time is taken from [1] and [17], but alternate models can be used, more closely related to experimental data and varying burn-in methods.

KGD test: KGD test involves test generation to a high fault cover (using the generic model outlined above), as well as test application of the test vectors. Added to these generic costs is the cost of the reusable carriers, related to the purchase cost, the number of uses possible, and the required throughput. The number of uses possible will be lower if the dice are subjected to burn-in. The extra test related costs increase the cost per die, which is further increased by the fact that the defective die population is depleted.

MCM test models: The MCM test model involves the test of the dice, the substrate and the interconnect. If boundary scan (IEEE 1149.1) is used, the modeling is straightforward: test vectors can be reused from the KGD process, so only one set of TPG costs needs to be invoked. The total number of vectors can be approximated from the total vectors for all die and the length of the scan chain, and the generic test application model invoked, with an added cost of the appropriate fixture. Furthermore, diagnosis costs can be assumed to be part of the test cost, and an interconnect test is also a part of the boundary scan test (if there is software support, there is little additional test generation effort involved). Modeling of burn-in is the same as for die.

D. Interconnect Models

Global interconnects within a modern electronic system exist at two levels: within a single chip (*intrachip* interconnects) and within the packaging medium connecting multiple chips

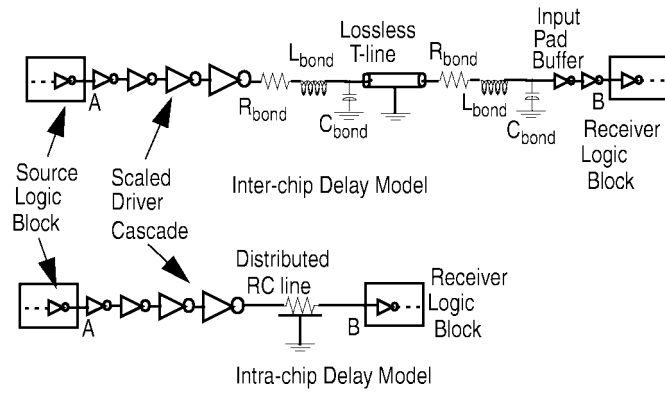


Fig. 5. 50% Delay models for interchip and intrachip interconnects.

(*interchip* interconnects). Our analysis of pulse propagation in both types of interconnect follows that in [2]. Interchip 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 interchip 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 resistance capacitance (RC) line. The time required for the output of the line to attain 50% of the input voltage step is given by $0.4r_{\text{int}}c_{\text{int}}l^2$, 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 interchip and intrachip communication, we utilize the delay models cited above and shown in Fig. 5 in the context of practical driver-receiver circuits. 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 interchip 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 resistance, inductance, capacitance (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 intrachip 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 intrachip signal path avoids the package parasitics in the interchip delay model, the signaling delay is quadratic in the interconnect length implying that intrachip delays can actually exceed interchip delays for long lines.

V. CASE STUDIES

In this section, we illustrate the application of IMPACT to three different systems, each representing a distinct analysis. The first study evaluates the memory hierarchy in modern reduced instruction set computers (RISC) microprocessors in light of the newly available packaging options. The second experiment considers the problem of partitioning the number of processors in a parallel computer system into several die configurations. Finally, the last case study applies the IMPACT tools to an asynchronous transfer mode (ATM) switch to perform a feasibility study of several architectural alternatives.

A. Case Study I: Processor-Memory Hierarchy

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 at a much slower pace. This has resulted in an imbalance in the memory and processor speeds, which requires multiple levels of cache memories to enable the processor to continue to function at the maximum speed. The large difference between intrachip and interchip 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 21 164 Alpha processor [9] 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 mm² and is manufactured in 0.5 CMOS technology. As die sizes increase, yields drop, costs rise and the high resistivity of the aluminum interconnect causes intrachip delays to become significant.

Several recent studies have begun to examine the impact of the MCM technology on the memory hierarchy [6], [11], [22]. 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, we consider moving the L2 cache off-chip in the above example of the DEC Alpha processor, which results in the following trade-offs.

Advantages

- 1) This partitioning will result in smaller die for the processor (logic) which leads to higher yields and hence lower cost.
- 2) An SRAM process may be used for the L2 cache rather than a logic process, leading to a denser, faster design.
- 3) 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.

- 4) Several smaller die versus one large die produces a dilation in the distribution of the thermal energy generated by the devices.

Disadvantages

- 1) The size of the MCM substrate increases as a function of the die footprint, increasing the substrate cost.
- 2) The increased number of I/Os due to partitioning at the L1/L2 interface will add to the die testing costs.
- 3) The increased number of I/Os may also increase the MCM substrate testing costs.

These are examples of the types of architectural trade-offs that can be explored with modern MCM packaging technology. Our goal is to be able to perform such trade-offs during conceptual design. This study 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 [18].

Cost and interconnect analyzes 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 single chip module (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 Fig. 6(a). 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. [18] 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 for the SCM and MCM implementations were determined using the models in Section IV. The cost/performance results are shown in Fig. 6(b). As expected, we see that for moderate to large cache sizes it is advantageous to use MCM's. 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 cost of a larger substrate. As cache sizes increases, the area increase in the microprocessor die becomes

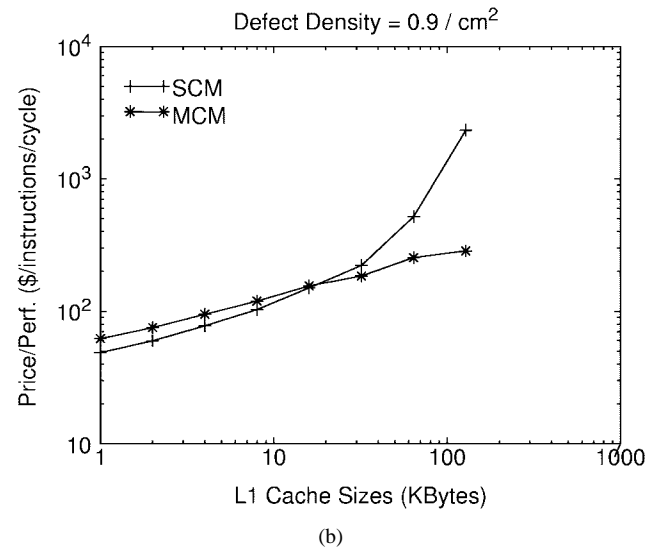
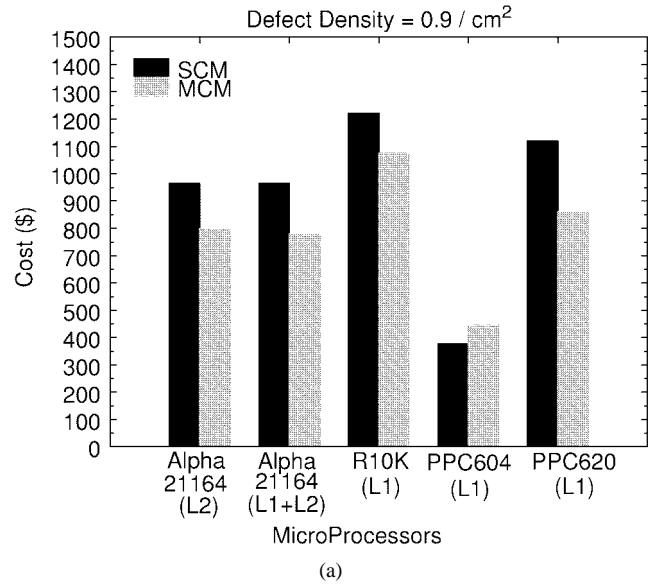


Fig. 6. (a) Comparison of SCM and MCM costs for modern microprocessors and (b) comparison of SCM and MCM price/performance.

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 intrachip and interchip 50% delays from point A to point B (Fig. 5) are plotted in Fig. 7 as a function of interconnect length. The delay equations are derived from expressions given in [2]

$$t_{\text{intra}}(l) = t_{\text{driver}} + 0.4r_{\text{int}}c_{\text{int}}l^2 + 0.7r_{\text{int}}C_{\text{rev}}l \quad (21)$$

$$t_{\text{inter}}(l) = t_{\text{driver}} + 1.4r_{\text{min}}c_{\text{min}} + \frac{l}{c_o}\sqrt{\epsilon_r}. \quad (22)$$

In (21) and (22), r_{int} and c_{int} are the resistance and capacitance of the distributed RC line per unit length, C_{rev}

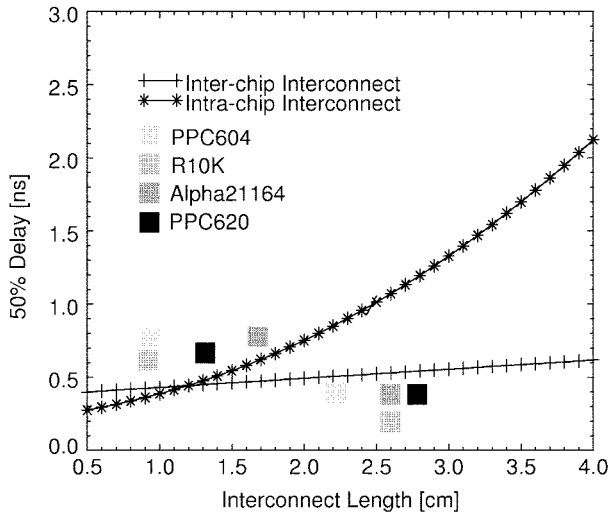


Fig. 7. Comparison of intrachip and interchip interconnect delays.

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 Fig. 7 were generated using device parameters from a 0.5 μm 3.3 V process. On-chip interconnects have a height of 1 μm and a width of 2 μm , yielding a r_{int} value of 140 Ω/cm . Given the effects of fringing fields, a limiting value for c_{int} of 2 pf/cm is used [2].

As Fig. 7 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 Fig. 7, the cluster on the left indicates the signal path lengths for the monolithic implementation of four commercial microprocessors. The cluster on the right indicates the corresponding interchip 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. Fig. 7 shows that the worst-case delays are comparable for the PPC604, R10K, and PPC620 systems. The delay for the partitioned Alpha system is significant lower than the delay for the monolithic Alpha implementation. As future processors become increasingly complex and larger, and MCM's become less expensive, the MCM solution should become increasingly effective.

B. Case Study II: SIMPIL—A SIMD Pixel Array Processor

SIMPIL is a single instruction stream, multiple data stream (SIMD) array processor designed to be used for embedded image processing and computer vision applications [4], [5]. The architecture is scalable to several thousand processing elements (PE's) interconnected in a two-dimensional (2-D) array topology. This analysis answers the question, where should the die boundaries be placed for a single MCM design. This is a trade-off between the number of I/Os and the chip size, and the MCM cost is used for evaluation. We also

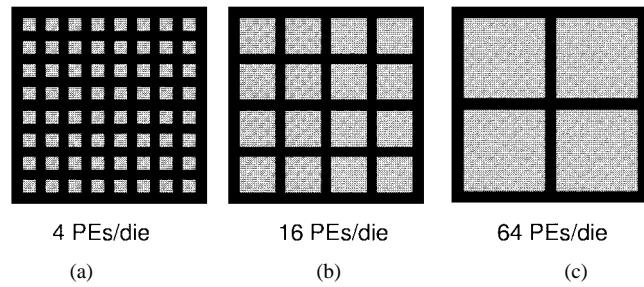


Fig. 8. Various die partitionings for a single MCM 256-node system.

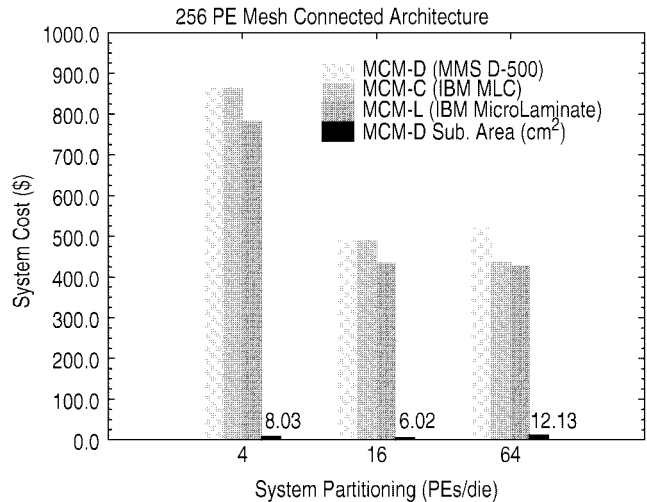


Fig. 9. MCM-D, C, L costs for various system partitionings.

explore the cost impact of using various MCM fabrication technologies, and various semiconductor technologies. Fig. 8 illustrates the three configurations which are considered. At one extreme we have 4 PEs/die which provides high yield due to smaller die area, but also restricts the number of I/Os, and since a large number of dice need to be assembled onto the MCM, the overall system reliability may be adversely affected. At the other extreme, fabricating 64 PE's on a die increases the die size rather drastically which in turn increases the MCM substrate area as well. As a result the system cost for either of these cases is quite large. An economically superior strategy uses 16 PEs/die which yields better system cost than the other two partitionings. These results are graphically illustrated in Fig. 9.

Fig. 10 shows the cost of the same system taking into account projection of semiconductor evolution. It indicates that in the years 1995–1998, when the integration levels are relatively moderate, the strategy using 16 PEs/die is economically better than the other two. But as the levels of integration improve over the years, i.e., as the defect density decreases, the effect of die area on the yield and hence cost of the die is minimized, and as a result systems with a higher degree of integration become economical once again.

This cost analysis we have performed does not take into account the NRE costs associated with the design and fabrication process. The system cost is the sum of the cost of bare dice, calculated based on die yield models and number of dice fabricated on a wafer; cost of the substrate, cost of C4 die

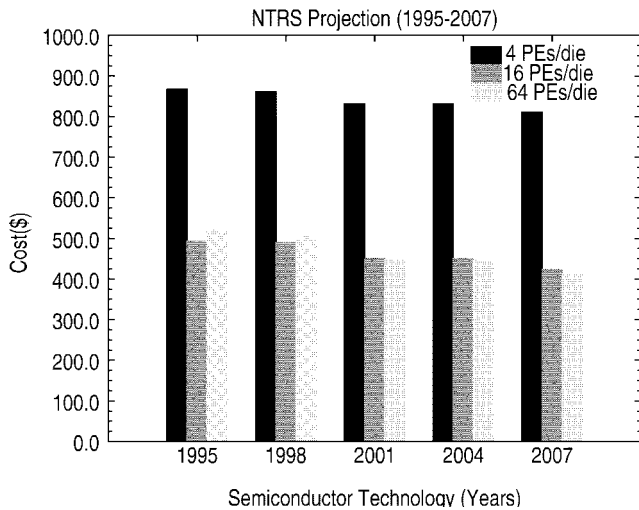


Fig. 10. MCM-D cost for various system partitionings, assuming fixed size die.

attach, and cost of testing dice and substrate. The substrate area calculations were obtained from the maximum of the surface area required to accommodate the dice, and wireability analysis using Bakoglu's extension of the Donath's model [2]. All die process related parameters were obtained from [25], and the MCM process related parameters were obtained from the available commercial process specifications. The models and parameters used to analyze the design can be found in Section IV and in [15].

VI. CONCLUSION

The goal of this work is an understanding of the impact of MCM packaging technology on system design. Our case studies suggest that MCM technology can be exploited to realize a new class of cost effective system designs. We have developed a suite of tools called IMPACT to help designers assess the effects of packaging on system architecture and design. The goal of IMPACT is to provide decision support for designers very early in the design cycle. Having the ability to predict the effects of packaging on system design early in the cycle can help shorten design cycles leading to higher profitability. Early decision support mechanisms also promote cost-effective use of packaging technologies, and provides the designers with a venue to evaluate alternate architectures. As packaging technologies advance, the traditional limitations such as limited I/O pads and slower off-chip bandwidth, etc. are no longer applicable. As a result, traditional architectural styles may be altered to realize more cost-effective designs which provide better or comparable performance.

As MCM technologies advance and mature, they will become an increasingly viable option for a wide range of applications. Our objective is to facilitate this process with early analysis tools that can reliably predict the impact of packaging options on system level metrics.

REFERENCES

- [1] A. F. Alani, C. Dislis, and I. P. Jalowiecki, "Burn-in economics model for multi-chip modules," *Electron. Lett.*, vol. 32, no. 5, pp. 2349–2351, Dec. 1996.
- [2] H. B. Bakoglu, *Circuits, Interconnections, and Packaging for VLSI*. Reading, MA: Addison-Wesley, 1990.
- [3] W. P. Birmingham, A. P. Gupta, and D. P. Siewiorek, "The MICON system for computer design," in *Proc. 26th Design Automat. Conf.*, 1989, pp. 135–140.
- [4] H. H. Cat, J. C. Eble, D. S. Wills, V. K. De, M. Brooke, and N. M. Jokerst, "Low power opportunities for a SIMD VLSI architecture incorporating integrated optoelectronic devices," in *Proc. GOMAC'96 Dig. Papers*, Orlando, FL, Mar. 1996, pp. 59–62.
- [5] H. H. Cat, M. Lee, B. Buchanan, D. S. Wills, M. A. Brooke, and N. M. Jokerst, "Silicon VLSI processing architectures incorporating integrated optoelectronic devices," in *Proc. 16th Conf. Advanced Research in VLSI*, Chapel Hill, NC, Mar. 1995, pp. 17–27.
- [6] P. Dehkordi, K. Ramamurthi, and D. Bouldin, "Early cost/performance cache analysis of a split MCM based MicroSparc CPU," in *Proc. MCM Conf.*, Feb. 1996.
- [7] C. Dislis, J. H. Dick, I. D. Dear, and A. P. Ambler, *Economics of Design and Test of Electronic Circuits and Systems*. New York: Ellis Horwood, 1995.
- [8] C. Dislis, A. F. Al-Ani, and I. P. Jalowiecki, "MCM quality and cost analysis using economic models," in *Proc. IEEE Int. Test Conf.*, 1995.
- [9] J. H. Edmondson *et al.*, "Internal organization of the Alpha 21164, a 300-MHz 64-bit Quad-issue CMOS RISC microprocessor," *Digital Tech. J.*, June 1995.
- [10] P. Elenius, "Why flip chips," presented at *Kulicke and Soffa Industries Seminar*, Jan. 31, 1996.
- [11] P. D. Franzon, A. Stanaski, Y. Tekmen, and S. Banerjia, "System design optimization for MCM," Tech. Rep. NCSU-ERL-94-16, North Carolina State Univ., Raleigh, 1994.
- [12] D. C. Frye, M. P. Skinner, and R. H. Heistand II, "Cost implications of large area MCM processing," in *Proc. 1994 Int. Conf. Multichip Modules*, Denver, CO, Apr. 13–15, 1994, pp. 69–80.
- [13] V. Garg, D. E. S. Lacy, D. E. Schimmel, D. Stogner, C. Ulmer, D. S. Wills, and S. Yalamanchili, "Impact of packaging constraints on system design: A case study of the memory hierarchy," Tech. Rep. CSRL-95/08, School Elec. Comput. Eng., Georgia Inst. Technol., Atlanta, Sept. 1995.
- [14] ———, "Incorporating packaging constraints into system design," in *Proc. Euro. Design Test Conf.*, Paris, France, Mar. 11–14, 1996, pp. 508–513.
- [15] V. Garg, D. E. Schimmel, D. Stogner, C. Ulmer, D. S. Wills, and S. Yalamanchili, "IMPACT models and parameters," Tech. Rep., School Elect. Comput. Eng., Georgia Inst. Technol., Atlanta, 1996.
- [16] G. L. Ginsberg and D. P. Schnorr, *Multichip Modules and Related Technologies: MCM, TAB, and COB Design*. New York: McGraw-Hill, 1994.
- [17] F. Jennesen and N. E. Petersen, *Burn-in: an engineering approach to the Design and Analysis of Burn-in Procedures*. New York: Wiley, 1982.
- [18] N. P. Jouppi and S. J. E. Wilton, "Tradeoffs in two-level on-chip caching," in *Proc. Int. Symp. Computer Architecture*, Chicago, IL, Apr. 18–21, 1994.
- [19] P. J. Krusius, "System interconnection of high density multi-chip modules," in *Proc. SPIE: Intl. Conf. Advances Interconnects Packag.*, 1990, vol. 1390, pp. 261–270.
- [20] D. P. LaPotin, "Early analysis of multichip packages," in *Proc. 10th Int. Electron. Packag. Conf.*, 1990, pp. 557–563.
- [21] J. H. Lau, Ed., *Chip On Board: Technologies for Multichip Modules*. New York: Van Nostrand Reinhold, 1994.
- [22] J. D. Roberts and W. W.-M. Dai, "Early system analysis of cache performance for RISC systems: MCM design trade-offs," Tech. Rep. UCSC-CRL-92-02, Univ. California, Santa Cruz, Mar., 1992.
- [23] P. A. Sandborn and H. Moreno, *Conceptual Design of Multichip Modules and Systems*. Norwell, MA: Kluwer, 1994.
- [24] C. M. Schreiber, "MCM substrates fabricated using flexible polyimide film," in *Proc. 1995 Int. Conf. Multichip Modules*, Denver, CO, Apr. 19–21, 1995, pp. 567–572.
- [25] Semiconductor Industry Association, *The National Technology Roadmap for Semiconductors*, 1994.
- [26] P. A. Trask and V. A. Pillai, "Large format MCM-D processing," in *Proc. 1994 Int. Conf. Multichip Modules*, Denver, CO, Apr. 13–15, 1994, pp. 94–99.
- [27] R. R. Tummala and E. J. Rymaszewski, Eds., *Microelectronics Packaging Handbook*. New York: Van Nostrand Reinhold, 1989.
- [28] G. E. White, E. Perfecto, T. DeMercurio, D. McHerron, T. Redmond, and M. Norcott, "Large format fabrication—A practical approach to low cost MCM-D," in *Proc. 1994 Int. Conf. Multichip Modules*, Denver, Colorado, Apr. 13–15, 1994, pp. 86–93.



Vivek Garg (S'97) received the B.S.E.E. degree from the University of Delaware, Newark, in 1990, the M.S.E.E. degree from the Georgia Institute of Technology, Atlanta in 1992, and is currently pursuing the Ph.D. degree in the School of Electrical and Computer Engineering, Georgia Institute of Technology.


His research is focused on architectural innovation for high performance SIMD machines. Other areas of research interest include parallel algorithms and architectures, interconnection networks, electronic packaging, and VLSI design. He has been a member of the System Integration, Design, and Test group of the Packaging Research Center since 1995. He held a DuPont Graduate Fellowship and a Presidential Fellowship at the Georgia Institute of Technology from 1990 to 1994.

Mr. Garg is a member of IMAPS, Eta Kappa Nu, and Tau Beta Pi.



Darrell J. Stogner received the B.S. and M.S.E.E. degrees from the Georgia Institute of Technology, Atlanta, in 1995 and 1996, respectively.

From 1995 to 1996, he was a member of the Systems Integration group, Packaging Research Center, where he worked on tradeoff analysis for MCM based designs. Since 1997, he has been a Software Engineer with the DSP group, Motorola's Land Mobile Products Sector, Plantation, FL.



Craig Ulmer received the B.S. degree from the Georgia Institute of Technology, Atlanta, in 1995, and is currently pursuing the M.S.E.E. degree.

He has worked with the Packaging Research Center at Georgia Tech since 1995, focusing on rapid predictive modeling of MCM's. His other research areas include real-time communication networks as well as computer based education in DSP.



David E. Schimmel (S'82-M'90) received the B.S.E.E. (with distinction) and Ph.D. degrees from Cornell University, Ithaca, NY, in 1984 and 1991, respectively.

He is an Associate Professor in the School of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta. During the Spring 1991 term he was a Visiting Researcher at the University of Linköping, Sweden. He was a Summer Faculty Fellow at NASA's Jet Propulsion Laboratory, Pasadena, CA, in 1995 and 1996.

He has also been a Visiting Engineer and a Consultant to IBM Almaden Research Center. His research interests include parallel computer architecture, algorithms and interconnection networks, asynchronous systems, VLSI design, and the impact of packaging technology on systems. He is author or coauthor of more than thirty refereed technical publications.

Dr. Schimmel was the 1993 Chair of the Atlanta Chapter of the IEEE Computer Society and is a member of ACM, Tau Beta Pi, and Eta Kappa Nu.



Chryssa Dislis (S'89-M'91) received the B.Sc. degree in electronic engineering from the University of Sussex, U.K. and the Ph.D. degree in digital systems from Brunel University, U.K.


She is a Lecturer in the Department of Cybernetics, University of Reading, U.K., and has been working in the area of Design for Test and Test Economics for the past eight years. Her work encompasses test strategy evaluation and optimization for ASIC's, VLSI boards, and Multichip Modules. Her other research interests include test pattern generation, mixed signal test, and autonomous agents. She is the author of a book on the Economics of Test and more than 30 papers and book chapters, and co-presents the IEEE International Test Conference Economics of Test tutorial.



Sudhakar Yalamanchili (S'79-M'79-SM'91) received the B.E. degree in electronics from Bangalore University, India in 1978, and the M.S. and Ph.D. degrees in electrical and computer engineering from the University of Texas at Austin in 1980 and 1984, respectively.

He is currently an Associate Professor in the School of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta. Prior to joining Georgia Tech in 1989, he was a Principal Research Scientist at the Honeywell Systems and Research Center, Minneapolis, MN, from 1984 to 1989. At Honeywell he was the Principal Investigator for projects in the design and analysis of multiprocessor architectures for embedded applications. During that time he was an Adjunct Faculty and taught in the Department of Electrical Engineering, University of Minnesota. He is a co-author of the forthcoming text *Interconnection Networks: An Engineering Approach* (New York: IEEE Press) and the author of the upcoming text *VHDL Starter's Guide* (Englewood Cliffs, NJ: Prentice-Hall). His current research interests are in the design and analysis of reliable multiprocessor interconnection networks, algorithms for dynamic resource management in multiprocessors, and the next generation of high density packaging technologies for high performance computer architectures.

Dr. Yalamanchili is a member of the ACM and SCS.



D. Scott Wills (S'80-M'90) received the B.S. degree in physics from the Georgia Institute of Technology, Atlanta, in 1983, and the S.M., E.E., and Sc.D. degrees in electrical engineering and computer science from the Massachusetts Institute of Technology, Cambridge in 1985, 1987, and 1990, respectively.

He is an Associate Professor of Electrical and Computer Engineering at the Georgia Institute of Technology. His research interests include VLSI architectures, high throughput portable processing systems, optoelectronics-enabled systems, and multicomputer interconnection networks.