Analyzing Packaging Trade-Offs During System Design

PETER A. SANDBORN
University of Maryland
MIKE VERTAL
Nu Thena Systems

Integrating packaging trade-off analysis with functional verification and architectural design results in a complete virtual prototyping solution for optimizing complex electronic systems. The authors discuss the role of packaging costs in system design and present examples highlighting packaging design trade-offs.

THE PRIMARY DRIVER of today's commercial electronics market is time. Portable computers, cellular telephones, and a host of other complex, high-density systems often have design cycles of less than a year and even shorter market windows. The very existence of these products depends on finding quick design solutions to meet increasingly challenging performance and cost requirements. Central to the success of these products are highly developed design methodologies and tools that facilitate first-pass success.

Although product developers apply a great deal of effort to selecting and designing chips with manufacturing costs in mind, they often sacrifice system packaging cost savings to meet time-to-market requirements. It's not uncommon for developers to "leave money on the table" at the end of a product development cycle. That is, they spend more than necessary because there isn't time to iterate a packaging-associated design change.

Another characteristic of short development cycles is slow adoption of technological advances. Developers don't readily adopt new technologies and materials because they cannot risk lengthening their product design cycles. Even if adequate infrastructure and suppliers exist, the lack of design experience and design support tools make new technologies risky. We lose many opportunities for gaining competitive advantage through the use of new technologies because we lack adequate methodologies and tools for making design and technology trade-offs.

A key to minimizing system cost without adversely affecting time to market is combining manufacturing information with application-specific design information as early in the design process as possible. By including an analysis of manufacturing costs in the system design methodology, we can avoid leaving money on the table at the end of development without increasing design time.

The common wisdom is that we commit 80% of a product's cost, size, and performance (electrical timing, power dissipation, and reliability) in the first 20% of the design cycle[1] (Figure 1). Yet this is the portion of the cycle for which we have the least mature methodologies and tools. This portion includes requirements capture, specification, trade-off, and partitioning. The later 80%, which includes part of the detailed simulation and all the physical design, refines and implements the design.

Clearly, the earliest stages of the design process provide the best opportunity to significantly impact a system's characteristics. This opportunity has led to the concept of virtual prototyping, which allows designers to define and test system attributes prior to large design or fabrication investments. By integrating packaging trade-off analysis with functional verification and architectural design, we can create a complete virtual prototyping solution for optimizing complex electronic systems.

Role of packaging trade-off analysis

Packaging trade-off analysis extends traditional virtual prototyping to include the determination of technology implementation details (Figure 2). Exactly how far the virtual

0740-7475/98/$10.00 © 1998 IEEE

prototype extends into describing system details depends on your design horizon. Most of the electronic systems design community (driven by IC design) defines virtual prototyping as functional verification and architectural design. Functional verification determines that the system's functionality satisfies customer requirements. Architectural design develops an integrated collection of hardware, software, and interface components that implement the system's functionality.[2]

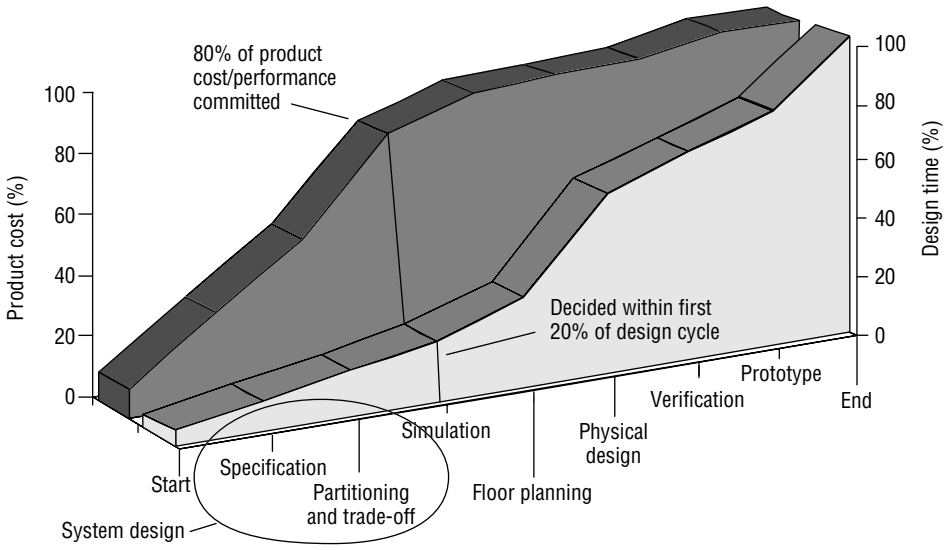

Figure 1. *A significant portion of an electronic system's cost, size, and performance is committed long before physical design (layout and routing) begins.*

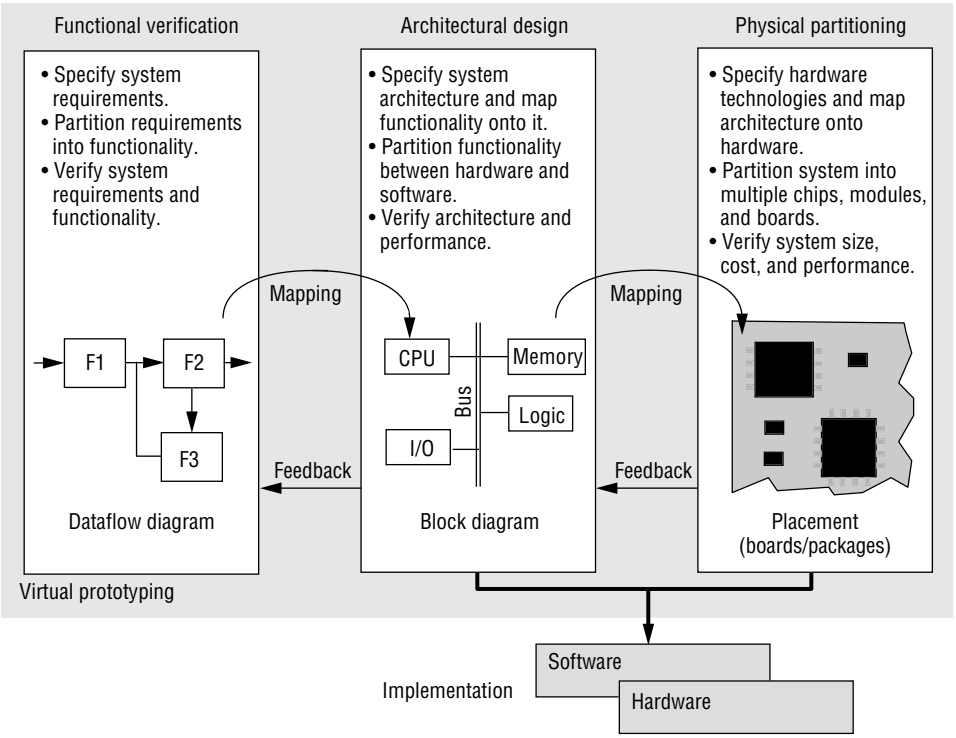

Figure 2. *Relationship of physical partitioning to functional verification and architectural design activities associated with traditional virtual prototyping. All three portions of the virtual prototyping space consist of specification, partitioning, and verification activities.*

The physical partitioning activity in Figure 2 takes inputs from architectural design in the form of a specification of the number and type of gates, bits, functional blocks, or dies, and their interconnections. Physical partitioning transforms the results of the architectural mapping into bare dies, packaged chips, multichip modules (MCMs), boards, and multiboard systems. Designers specify implementation technologies, materials, and design rules with the help of libraries and analyze the resulting physical partitioning and implementation with a suite of estimators and simulators. This analysis produces a set of performance, size, cost, and manufacturing metrics that designers can use to compare multiple design candidates.

Like functional verification and architectural design, physical partitioning has a hierarchical organization. This means that we can define physical containers (dies, single-chip packages, few-chip packages, MCMs, boards) to arbitrary depths. For example, boards can contain other boards or MCMs, boards and MCMs can contain chip packages, chip packages can contain one or more bare dies, dies can contain one or more functional blocks, and functional blocks can contain one

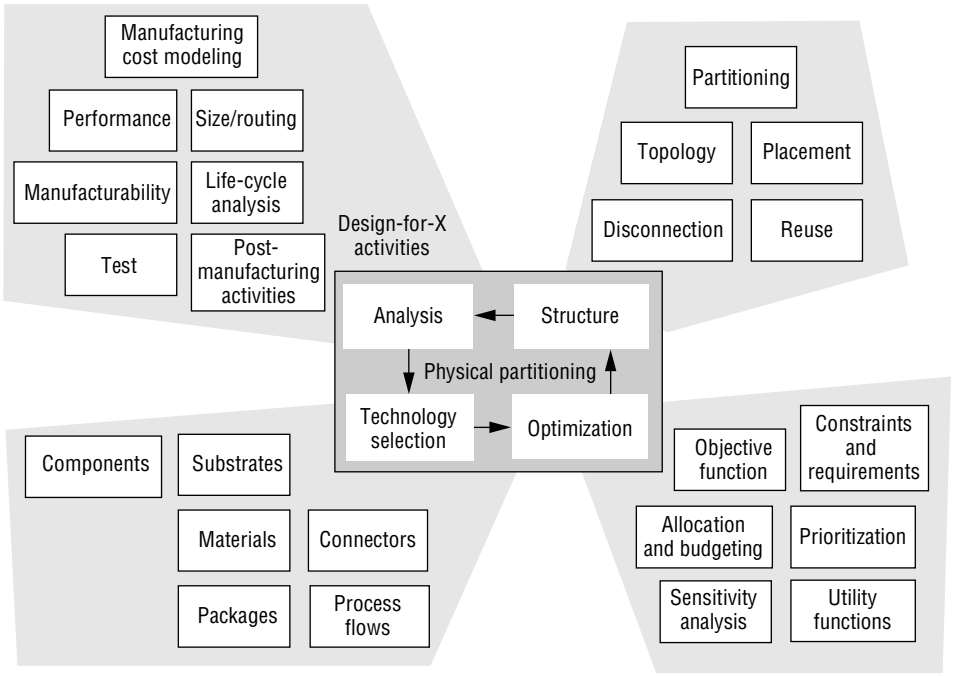

Figure 3. *Enabling activities for system-level physical trade-off analysis.*

or more gates.

Physical partitioning must consider the interrelated components shown in Figure 3. The following four categories of activities must seamlessly interact to provide effective physical partitioning:

- *Structure.* Structure represents the architecture of the system's physical implementation. The following activities are associated with structure: placement (orientation of adjacent components to each other), topology (orientation of adjacent boards to each other), disconnection, (designing disconnectable subsystems), and reuse (dividing systems into parts that can be reused in other systems). During physical partitioning, the designer can evaluate different physical implementations by moving architectural elements among physical partitions and changing component and/or subassembly placement within partitions.
- *Technology selection.* Technology characterization is at the core of physical partitioning. Physical partitioning is highly dependent on technology decisions; the other portions of the virtual prototyping solution are less sensitive to technology choices. Technologies that support physical partitioning include components, substrates, materials, connectors, packages (for bare dies), and process flows for fabricating, assembling, and testing systems. These technologies should be accessible in libraries. In addition to the data itself, technology selection requires data management methodologies that allow physical partitioning to access the necessary technology data quickly and accurately.
- *Analysis.* Physical partitioning must include analysis of all the "design-for-X" activities: design for manufacturability, environment, testability, cost, and so forth. It also includes all performance estimation activities: electrical, thermal, and reliability. This portion of the virtual prototyping system can use point design tools (simulators) and/or estimation-level advisors. Each analysis should lead to information about economic impact and must be tightly coupled to manufacturability. This category should also include analysis of design for postmanufacturing activities such as recycling, disassembly, servicing, maintenance, and upgrading.
- *Optimization.* Optimization is a management framework within which all the partitioning and design analysis activities are performed. Optimization does not necessarily mean that the system must automatically choose the optimum design specification without user involvement. Rather, it represents tools that collaborate with the user to optimize the design's physical implementation. Optimization includes objective function formulation, allocating and budgeting, sensitivity analysis, prioritization, and constraint and requirement management.

Many companies see formalizing and automating the trade-off analysis process as a key to continuous improvement strategies. It is one more piece in the concurrent-engineering puzzle and a proactive approach to product optimization. For trade-off analysis and its associated design components to be fully effective, we must adopt metrics that unify diverse design concerns. Most design characteristics, such as testability, manufacturability, and reliability, have clear-cut metrics, usually tied to customer expectations and demands. These metrics are typically parametric (that is, they have independent and dependent variables), and their impact on a product's manufacturing life-cycle cost is quantifiable. More important, most compa-

nies understand and appreciate the critical nature of these quality indicators at the product design level.

Several notable efforts have made progress in integrating estimation and design analysis activities to create packaging trade-off analysis methodologies and software tools.[3-9]

Estimating system economics

Original equipment manufacturers use a hierarchy of analysis methods to estimate costs. Current capabilities available for packaging trade-off analysis are limited to the manufacture of a single product. Nevertheless, engineers must remember that the product's economic viability depends on their understanding its entire life cycle—and potentially the life cycles of all the other products the company produces. Table 1 shows the hierarchy of cost estimation methods companies use today.

Table 1. *Cost estimation hierarchy.*

Product development level	Scope	Key attribute	Popular methods
Corporate	All company activities (multiple products)	Accurate allocation of overhead to specific products	Activity-based cost analysis Traditional cost accounting
Program	Single product	Entire product life cycle including hardware and software	Parametric Traditional cost accounting
Manufacturing	All or part of single product	Hardware only	Cost of ownership Process flow modeling Factory simulation Parametric
Process	Single processing activity	Hardware only	Analytical models

Companies perform most trade-off analysis at either the program or manufacturing levels, using parametric, cost-of-ownership, or process flow modeling. The methods shown in Table 1 are applicable to five different system components: IC fabrication, IC assembly preparation (bumping, test and burn-in, single-chip package fabrication), substrate fabrication, assembly (including test and rework), and software development.

A key to correctly modeling system economics is the integration of manufacturing, program, and life-cycle analysis. *Life-cycle analysis* refers to a group of methods for assessing materials, services, products, processes, and technologies over their entire life. It includes inventory analysis, which creates an inventory of energy and material used and wasted in the creation of a component. It also includes impact analysis, which determines the inventory's impact on system metrics. Life-cycle analysis is a necessary part of IC fabrication because environmental health and safety (EHS) and waste disposition costs constitute about 10% of the final product cost. Life-cycle analysis is necessary in substrate fabrication because costs of materials and waste disposition dominate board costs. Life-cycle costs included at the program level include design, inventory, learning curve, transportation, liability, and support.

There are several manufacturing cost analysis methods. The development of cost-of-ownership models for IC fabrication[10] and electronic system assembly has led to an understanding and appreciation of activity-based cost estimation methods directly tied to specific process steps. The printed wiring board (PWB) fabrication process has similarities to IC fabrication and electronic system assembly. All three manufacturing activities are process flow and activity oriented; they require labor, material, tooling, equipment, and facilities; and they are applied to single parts or batch formats.

However, the significant cost drivers of PWB fabrication differ from those of IC fabrication and electronic system assembly. Cost-of-ownership approaches for IC fabrication focus on computing the lifetime cost of owning and operating specific equipment and the equipment's impact on the fabrication process. These are the cost drivers in the IC industry. Although IC fabrication models include the cost of materials, it is not the focus of the analysis, which typically does not emphasize the ability to perform detailed material manipulation. In contrast, materials are the main component of PWB fabrication cost—in some cases, more than 50%. Labor is the second-largest cost driver, and equipment and facilities are a distant third. Even significant changes in equipment and facilities costs (such as maintenance and downtime) typically have little impact on final PWB cost. Therefore, for PWBs, we have developed an alternative cost-of-ownership modeling approach based on material use and waste. This approach, which we call the material-centric ap-

Material-centric PWB fabrication model

A material-centric PWB fabrication model[1] defines each activity or process step in terms of what it does to the materials associated with the substrate being fabricated. Fundamentally, the model defines five such activities:

- material addition to the product: plating, coating, lamination, filling
- material subtraction from the product: etching, stripping, drilling, trimming
- waste disposition
- scrapping defective parts
- no material manipulation

The first four activities may have associated consumables—materials attached to the process as opposed to the product. Consumables are used (and wasted) by the activities, but at no point in the process do consumables reside in the product. Examples of consumables are water, artwork, and drill bits.

Process steps that model these activities contain information about the equipment and facilities required, but the equipment and facilities do not define the process step. Defining process steps in terms of their material treatment enables straightforward modeling of material and waste costs.

We implemented the material-centric cost model in SavanSys, an existing multidisciplinary tool for multichip packaging trade-off analysis. Figure A summarizes the tool's basic process flow model.

During process execution, the tool creates and manipulates inventories of material in the product, material in the waste stream, and energy consumed. Each material inventory catalogs material volume at standard temperature and pressure. As each process step executes, the tool computes its material and energy requirements and adds them to or subtracts them from the inventories. Some activities transfer materials between inventories. For example, if a step produces waste materials by removing material from the product, the waste quantity is subtracted from the material-used inventory and added to the material-wasted inventory. All inventories are normalized to one instance of the part being processed. That is, the inventories record the used and wasted materials for a single panel or board. To determine the total waste, the user must multiply the waste inventory contents by the number of panels or boards processed.

Reference

1. P.A. Sandborn and C.F. Murphy, "Material-Centric Modeling of PWB Fabrication—An Economic and Environmental Comparison of Conventional and Photovia Board Fabrication Processes," *IEEE Trans. Components, Packaging, and Manufacturing Technology, Part C*, Vol. 21, No. 2, Apr. 1998, pp. 97-110.

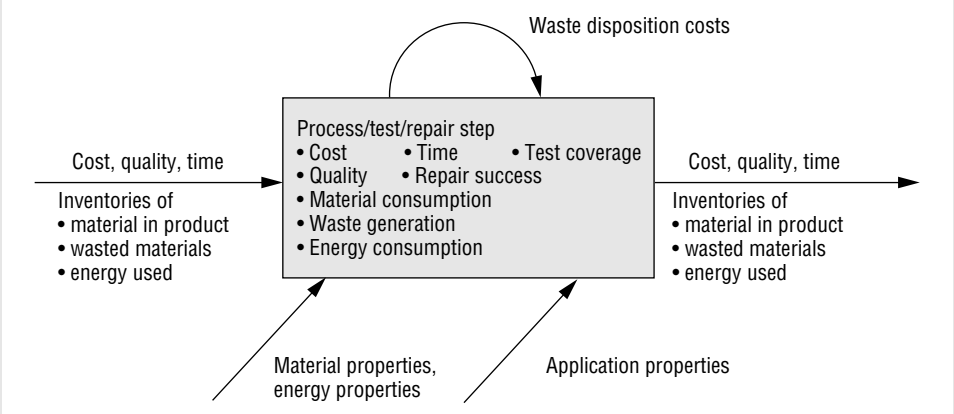

Figure A. *SavanSys trade-off analysis tool's process step model used in material-centric PWB cost modeling.*

proach (see box), integrates life-cycle analysis and manufacturing cost analysis.

Another manufacturing cost analysis method is factory simulation, prevalent in the IC and board fabrication communities. Factory simulators aim at optimizing factory floor operations—a valuable activity, but not generally applicable to application-specific system design focused on trade-offs.

Cost modeling alone is insufficient for making accurate economic trade-offs; we must also consider system quality (manufacturing yield and lifetime reliability). System quality is a function of fault probability and test coverage. Fault probability measures the likelihood of at least one fault per component. We determine fault probability from the distribution of defects introduced by various processing activities. Test coverage is the probability that a particular test activity will detect a defective component. Models of the cost and effectiveness

of testing activities applied to electronic systems are available.[11]

Application examples

High-density systems are characterized by strong interdependencies among size, cost, manufacturing, and performance (Figure 4). In the days when systems used only through-hole components in large I/O pitch packages, divide-and-conquer and correct-by-verification design approaches were appropriate. Today's high-density PCMCIA (Personal Computer Manufacturers Communications Interface Adapter) cards, chip-scale packages, MCMs, and microvia boards require correct-by-design approaches that concurrently address strongly coupled size, performance, manufacturing, and cost issues.

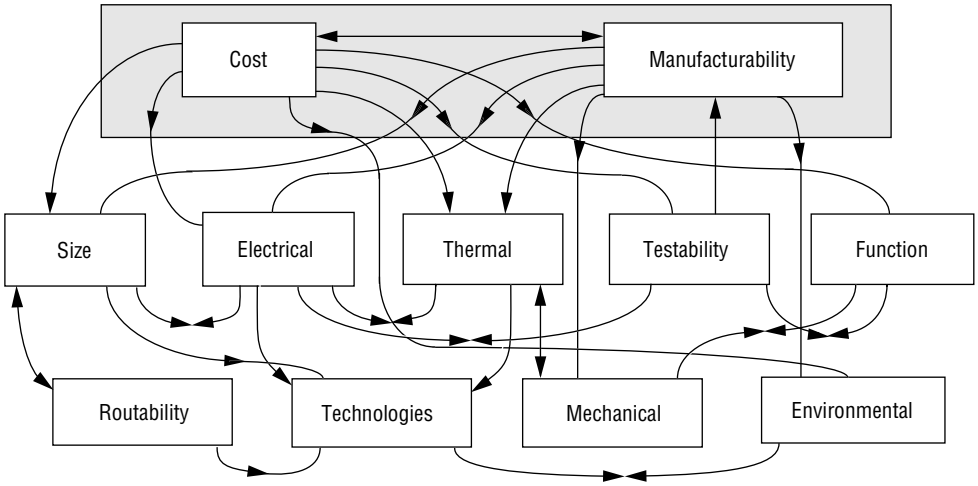

Figure 4. *The interdisciplinary nature of high-density system design.*

The examples of packaging design trade-offs presented in this section illustrate the following points:

- The least expensive set of components does not always lead to the least expensive system.
- Lack of communication between design and manufacturing makes finding optimum economic solutions impossible.
- To find the optimum technology implementation solution, the trade-off analysis may need to consider post-manufacturing life-cycle costs.

We performed all the trade-off analyses described here long before physical design (layout and routing) activities or the existence of a netlist, using the SavanSys tool from Nu Thena Systems (see box, next page).

Which bare die should I use? A major challenge associated with building systems that include bare dies is die test and burn-in. We refer to this as the known-good-die problem. Although single-chip packages require larger board areas and add unwanted electrical parasitics, they have the advantage of allowing inexpensive die test and burn-in prior to assembly. This opportunity for test and burn-in is not necessarily available to dies that are never packaged.

In our first example, we must construct a data accumulator/storage module using a combination of wire-bonded bare dies, packaged chips, and discrete components mounted on both sides of a PWB. Eight of the bare dies are static RAMs, available at various prices depending on quality level. The challenge is to determine the best cost-and-quality combination for this application, considering module assembly, test, and rework costs. We can purchase the SRAMs for $100 each, untested, with a yield of 80%. If we purchase the untested SRAMs, the board must be assembled, tested, and diagnosed, and defective SRAMs replaced as needed. A second option is to buy 100%-yield bare-die SRAMs for $500 each from a third party that has performed the test and burn-in. (A third option, which we do not consider in this case, also exists: System OEMs can invest in the capability to test and burn in bare dies themselves.) Which of these two options results in the least expensive system?

To determine the best combination for this application, we vary the SRAM cost and yield (at assembly) in accordance with the available dies. In each case, we have the resulting module tested and repaired (if necessary) after assembly to bring the module yield to approximately 100%. Figure 5 (next page) shows the results of this analysis. The diagonal line shows module cost as a function of the price paid for a known-good SRAM (100% yield at assembly). The horizontal line indicates the cost of a module assembled with 80%-yield dies bought for $100 each (the module cost includes diagnosis and rework to replace defective SRAMS after assembly). The intersection of the two lines indicates that we should use known-good SRAMs if we can purchase them for less than $600 each. If tested and burned-in SRAMs cost more than $600 each, we can manufacture the system more economically by assembling the module using lower-yield, untested dies and repairing the system later.

Finding technology windows of opportunity. A critical element in estimating the cost of electronic packaging sys-

SavanSys trade-off analysis tool

The trade-off platform we used to implement the analysis in this article is SavanSys from Nu Thena Systems. SavanSys is a software tool for enhancing the manufacturability and decreasing the design risk of IC packaging technologies. The tool analyzes system-packaging trade-offs by concurrently computing physical, electrical, thermal, reliability, and cost/yield metrics for multichip systems.

MCMs and traditional packaging (through-hole and surface-mount) technologies handled by SavanSys include traditional and microvia PWBs, low-temperature cofired ceramic, and thin film (chip-first and chip-last). Component assembly approaches include wire bonding, tape-automated bonding (TAB), flip-chip, and single-chip packages. SavanSys includes a materials database for bare-die attachment, encapsulation, extrusion attachment, and bonding and substrate technology definition.

SavanSys enables the user to compute the following costs of assembled electronic systems:

- components
- component preparation (wafer and die-level burn-in, bumping)
- single-chip packages
- surface-mount and through-hole assembly
- bare-die attachment (TAB, wire-bond, flip-chip)
- tooling
- substrate
- repair and rework
- testing

In addition, users can define optional forecasting functions and learning curves for any or all process steps, and handling costs for all steps that insert components into the process flow.

The SavanSys trade-off analysis tool evaluates the impact of technology, material, and design rule variations on the cost, size, manufacturability, and performance of a board or system of boards. It enables designers to choose optimum physical partitionings early in the design process to facilitate successful implementations. SavanSys is integrated into the Mentor Graphics and Cadence physical design frameworks and is compatible with the Aspect and DIE (Die Information Exchange) format databases.

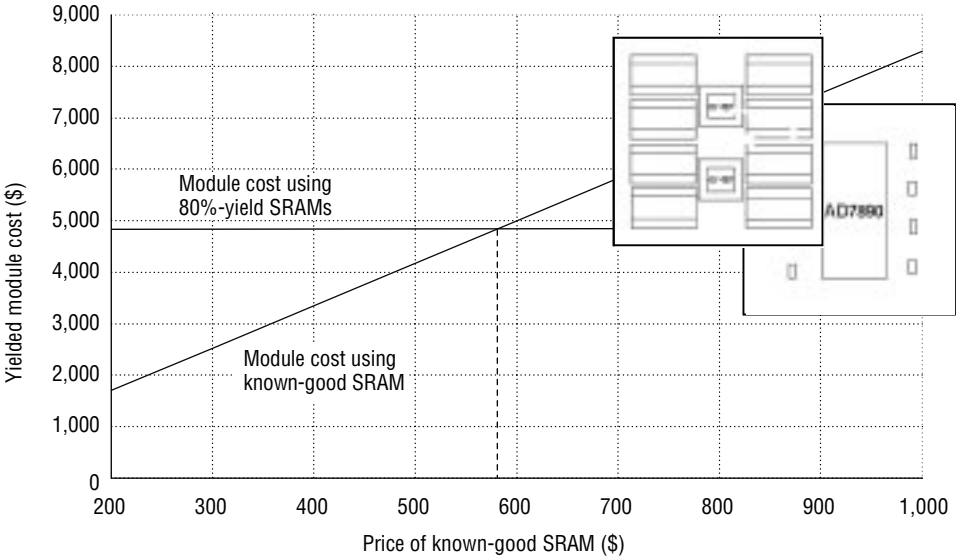

Figure 5. *Yielded module cost versus the price of tested and burned-in (100%-yield) SRAMs. Dashed line indicates price at which known-good SRAM becomes uneconomical. The inset shows the front and back layouts of the data accumulator/storage module.*

tems is understanding the manufacturing process to be used and applying that knowledge during the design process. A simple manufacturing reality often overlooked by system designers is panelization. Most substrate fabrication approaches can fabricate multiple boards at the same time on standard-size panels. The board's size and shape, the panel's size and shape, and the substrate technology determine how many boards can be fabricated on a single panel. The number of boards fabricated on a panel, which we call the "number-up," is a major cost driver we can use to determine the cost of assembling a system.

Consider an application whose wiring requirements dictate the need for an eight-layer board. At present, for laminate substrates with eight layers or fewer, the least expensive board fabrication approaches appear to be those using mechanically drilled

vias (Figure 6). Above eight layers, microvia technologies, which use photolithography, plasma etching, or laser drilling to create vias, appear to provide a less expensive approach.

PWB fabrication cost is the cost of fabricating a panel divided by the number-up. Unlike IC fabrication, which depends on step-and-repeat photolithography that requires all parts to be oriented identically, board fabrication can support 90-degree relative rotations of boards on the same panel (nonhomogeneous panelization). The ability of a substrate technology (and/or fabrication facility) to support nonhomogeneous panelization depends on two things. First, materials must be dimensionally homogeneous; that is, during lamination they must change shape the same way in the x direction as in the y direction. This characteristic allows boards to be arranged on the panel with varied rotations without artwork compensation problems. The second requirement is fabrication equipment that can be programmed for multiple board orientations.

Figure 7a shows the number-up as a function of the edge scrap allowance, the amount of space needed around the edge of the board. Obviously, for some board sizes and edge scrap allowances, no additional boards can be fabricated on the panel through nonhomogeneous panelization. In other cases, the number-up increases. For this example, we assumed that the less expensive, conventional board technology we considered allows only homogeneous panelization, and that the more expensive, microvia technologies allow nonhomogeneous panelization. Figure 7b shows the cost per board for the panelizations shown in Figure 7a, using the cost of a 140-inch-per-square-inch interconnect capacity from Figure 6. The analysis shows that in at least one case (the 10 × 5-inch board), the more expensive, microvia approach results in a less costly board because materials and manufacturability allow nonhomogeneous panelization during manufacturing. For simplicity, this analysis ignores potential system density improvements (layer reduction) that may also be possible with microvia technology.

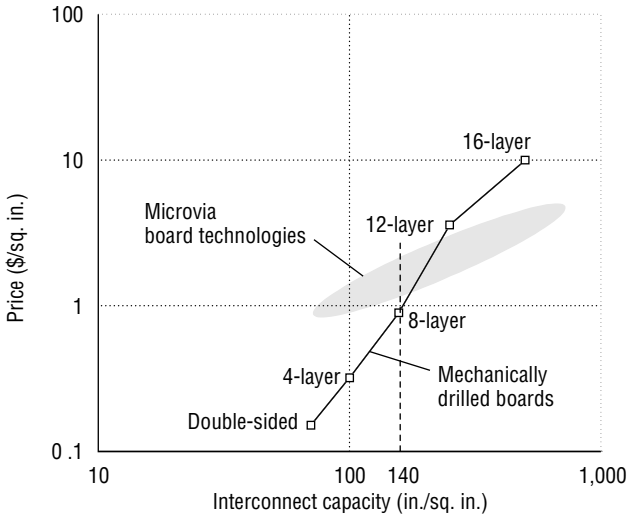

Figure 6. *Relative costs of conventional, mechanically drilled laminate substrates and microvia substrates (source: Fisher[12]). The gray "cloud" represents the prices of microvia board technologies as a function of their interconnect capacities. Prices may differ for the same capacity due to variations in the technology used to fabricate the boards.*

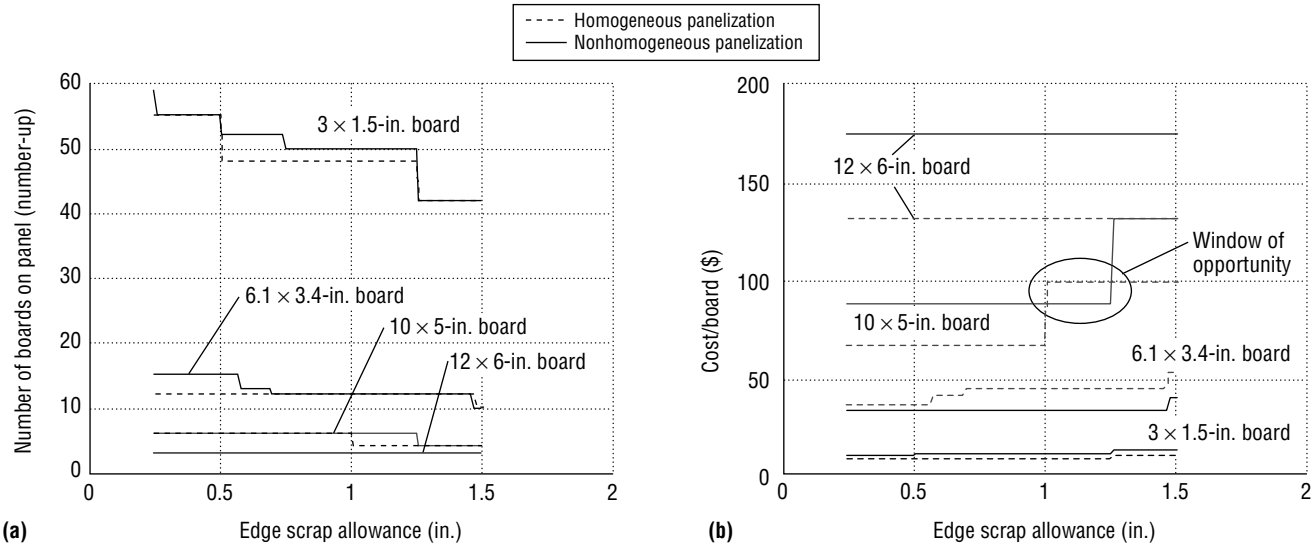

Figure 7. *Number of boards per 18 × 24-in. panel (number-up) as a function of edge scrap allowance* ***(a)****; cost per board as a function of edge scrap allowance* ***(b)****. "Window of opportunity" indicates conditions under which a more expensive technology yields a less expensive system.*

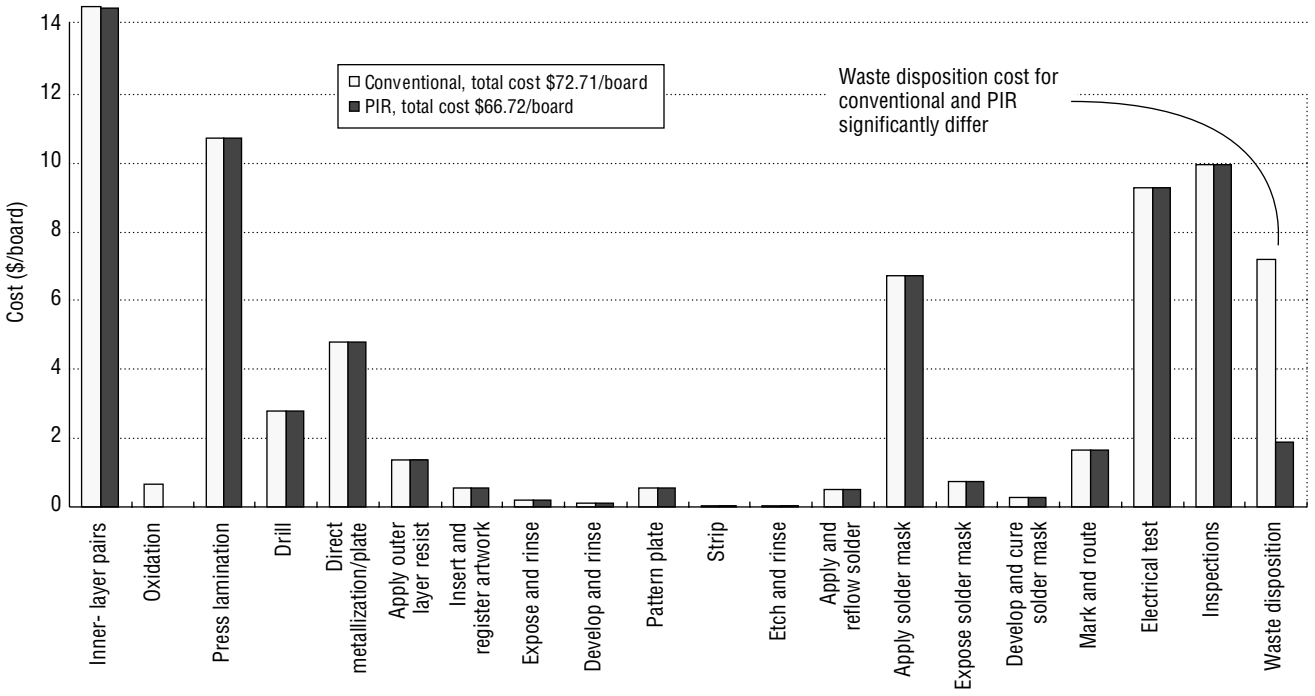

Figure 8. *Costs of conventional and PIR full multilayer fabrication including waste disposition.*

Expanding trade-off analysis into product life cycle. Design rules, board shape and area, number of layers, and materials drive PWB fabrication cost. Many different processes are in use today. Consider the trade-off between two alternative resists in PWB manufacture. Resists are photosensitive materials whose solubility in a developing solution changes on exposure to light, allowing the metal on a panel's top layer to be etched in a desired pattern. For a particular application, should you use a PWB fabricated with a conventional resist such as Riston or a newer, experimental resist called Permanent Innerlayer Resist (PIR)? When using PIR, the fabricator leaves the photoresist used in patterning the internal layers on top of the copper circuits after development and etching of the layer pairs. This PIR film acts as an adhesion promoter, a function typically achieved by oxidizing the copper traces. With PIR, a board shop can eliminate the stripping and oxide treatment processes.

Figure 8 compares the costs of PIR and the conventional approach for a 6×9-inch, eight-layer board on an 18×24-inch panel. For simplicity, the figure groups approximately 230 process steps into 20 major activities. The total cost for the PIR board is virtually identical to that of the conventional board. The main cost difference (excluding waste disposition for the moment) is the oxidation process at $0.70 per board. With a known uncertainty of ±$0.80 per board, Figure 8 indicates that the PIR approach may be less expensive than the conventional approach—but by a margin smaller than the uncertainty. So far, therefore, this trade-off analysis yields questionable answers at best.

The trade-off between these two resist technologies becomes clear only when we consider additional life-cycle costs. Specifically, a large fraction of the cost of conventional PWB fabrication is waste disposition cost. It turns out that oxide treatment for improved copper-to-prepreg adhesion is one of the "dirtier" PWB fabrication processes. The strong caustic at high temperature required for the oxide process is costly to use and dispose of due to its high pH, its oxidizer content, and its build-up of dissolved metals. Eliminating the oxide process reduces water use. In addition, using PIR eliminates stripping and its associated nonhazardous waste and water use. The difference between waste treatment costs for the two approaches is striking. At over $7 per board, the waste disposition cost of the conventional process is more than three times that of the PIR approach at just under $2 per board. The primary driver is the amount of water required by the conventional process, which increases sewage and sludge disposition costs.

After considering waste disposition, we find that the PIR approach offers potentially almost 8% in cost savings for a multilayer board, the primary benefit coming from reduced water use. Total waste decreased from 427 liters per conventional board to 41 liters per PIR board; again, most of the reduction was water. Waste excluding water and gas decreased from 5.3 to 3.8 liters per board (28%). A manufac-

turing cost model alone would not have identified the major cost differentiation for this case. Understanding this trade-off required both manufacturing cost modeling and detailed waste material inventory and waste disposition modeling—in other words, the integration of life-cycle and manufacturing cost analyses.

THE DEEPER the product life-cycle trade-off analysis goes, the more likely it will find the best economic solution. At a corporate level, no single product can be optimized for cost independently of the rest of its product family. Often, we must sacrifice one product's optimization to optimize the family as a whole. In other words, we may use a more expensive technology than necessary in the first product of a family simply to learn how to design and manufacture in that technology. This paves the way for later products in the family for which the new technology will be absolutely necessary. Product life-cycle costs also include design and software development costs. Therefore, future trade-off analysis solutions must deal with hardware and software concurrently. D&T

References

1. L.K. Keys, "System Life Cycle Engineering and DF'X'," *IEEE Trans. Components, Hybrids, and Manufacturing Technology*, Vol. 13, No. 1, Mar. 1990, pp. 83-93.
2. M. Vertal and J. Crowley, "An Approach to Hardware/Software Co-Modeling for Rapid Design Exploration," *Proc. High-Level Electronic System Design Conf.*, 1997, pp. 140-147.
3. W.E. Pence, *Electrical, Thermal, and Architectural Aspects of VLSI Packaging and Interconnects for High-Speed Digital Computers*, PhD dissertation, Electrical Eng. Dept., Cornell Univ., Ithaca, N.Y., 1989.
4. H.B. Bakoglu, *Circuits, Interconnections, and Packaging for VLSI*, Addison Wesley, Reading, Mass., 1990.
5. L.L. Moresco, "Electronic System Packaging: The Search for Manufacturing the Optimum in a Sea of Constraints," *IEEE Trans. Components, Hybrids, and Manufacturing Technology*, Vol. 13, No. 3, Sept. 1990, pp. 494-508.
6. D.P. LaPotin, T.R. Mazzawy, and M.L. White, "Early Package Analysis: Considerations and Case Study," *Computer*, Vol. 26, No. 4, Apr. 1993, pp. 30-39.
7. P.A. Sandborn and H. Moreno, *Conceptual Design of Multichip Modules and Systems*, Kluwer Academic, Norwell, Mass., 1994.
8. *EDA Navigator*, Xynetix Design Systems, Inc., Fishers, N.Y.; www.xynetix.com.
9. V. Garg et al., "Early Analysis of Cost/Performance Trade-Offs in MCM Systems," *IEEE Trans. Components, Packaging, and Manufacturing Technology, Part B*, Vol. 20, No. 3, Aug. 1997, pp. 308-319.
10. D.L. Dance, T. DiFloria, and D.W. Jimenez, "Modeling the Cost of Ownership of Assembly and Inspection," *IEEE Trans. Components, Packaging, and Manufacturing Technology, Part C*, Vol. 19, No. 1, Jan. 1996, pp. 57-60.
11. C. Dislis et al., *Economics of Design and Test of Electronic Circuits and Systems*, Ellis Horwood, New York, 1995.
12. J. Fisher, "Roadmapping the High Density PWBs," *Proc. IPC National Conf.: Solutions for Ultra High Density PWBs*, Institute for Interconnecting and Packaging Electronic Circuits, Northbrook, Il., 1996, pp. 1-20.

Peter A. Sandborn is an associate professor in the Department of Mechanical Engineering at the University of Maryland. Previously, he was a senior member of the technical staff at Microelectronics and Computer Technology Corp., a founder and the chief technical officer of Savantage, Inc., and a technical contributor at Nu Thena Systems, Inc. His research experience and interests include developing methodologies and tools for technology trade-off analysis of electronic systems with a focus on design to cost, design for environment, and design for manufacturability. Sandborn received the BS degree in engineering physics from the University of Colorado and the MS and PhD degrees in electrical engineering from the University of Michigan. He is a member of the IEEE.

Mike Vertal, vice president of product development at Nu Thena Systems, has been active in the system-level design of complex electronic systems for eight years. He led the development of the Foresight design tool and applied it to a wide variety of hardware/software designs in telecommunications, high-speed data communications, and aerospace. He has published numerous papers on embedded control systems, system modeling and simulation, and hardware-software codesign. Vertal received his bachelor's and master's degrees in electrical engineering from Case Western Reserve University. He is a member of the IEEE.

Send questions and comments about this article to Peter A. Sandborn, University of Maryland, Dept. of Mechanical Engineering, College Park, MD 20742-3035; sandborn@calce.umd.edu.