

The Road Ahead

Andrew B. Kahng

University of California at San Diego

The significance of packaging

■ **PACKAGING** has become a key driver for design and test technologies, as they advance to satisfy the requirements of the *International Technology Roadmap for Semiconductors* (<http://public.itrs.net>). The package is silicon's interface to the world and must satisfy many demands. For example, it must

- fit onto line cards and satisfy other form factor constraints;
- maintain conditions with respect to heat, moisture, electrical isolation, and ionizing radiation;
- provide critical design flexibility for silicon technology and system architecture optimizations; and
- hold the line on cost as a dominant component of system cost.

Die package codesign is a familiar challenge to longtime readers of the *ITRS* chapters on Design and Assembly, and Packaging. But despite the *ITRS*, both academia and industry have devoted insufficient resources to this challenge for too long.

Cost

The cost of silicon in a 200-mm wafer for a generic CMOS logic process averages about 10¢/mm². This cost grows superlinearly with die area and also depends on process complexity and yield ramping. Thus, a large die might cost up to several tens of dollars. With the move to 300-mm wafers, estimates of die cost per square millimeter drop to as low as 5 cents at the 90-nm technology node. On the other hand, pack-

age costs range from 1 to 7 cents per pin, generally increasing with area array I/O (flip-chip) technology or with the high pin counts found in networking and high-performance computing products. For a high-end part, package cost today can easily exceed \$100.

Critical challenges include controlling thermal effects and assembly costs. Package materials (for example, ceramic versus plastic) become more expensive with increasing power budgets. Estimates of this dependence are on the order of \$1/watt. That means just 1 watt or 20 extra package pins will add as much cost to the system as 20 mm² of die area. Two limiting cases are small die with very high thermal flux (as in silicon-germanium for handheld applications) and large die with very large thermal mass (as in high-end server chips). In each case, the problem goes beyond energy storage or local heat removal: Novel technologies such as fuel cells or thermoelectric cooling cannot get around practical system limits on total cost of goods and total power budget—for example, the number of watts per square foot that a data center can handle. With the long-foreseen transition to flip-chip from wire bond technology, assembly cost has emerged as another critical challenge, because the flip-chip joint is inherently more expensive than its wire bond equivalent.

Bandwidth

In recent years, networking system bandwidth requirements have grown at multiples of 10× or 4× between generations. Contrast this increase with the far smaller rates of advance in

- signaling standards,
- number of available die bumps and package pins, and
- materials and manufacturing technologies for connectors and backplanes.

Furthermore, one high-speed port often implies eight or more lines when you take into account bidirectionality, differential signaling, and shielding. Today, package technologists continually juggle the roadmaps for bump and pin pitches to meet on-silicon interconnect pitches, current delivery requirements, and other constraints. Tighter pitches increase the number of available bumps and pins and help alleviate *IR* drop or current delivery concerns arising from increased power and frequency with lower supply voltages. However, the resulting cost increases, and even the eventual impact of trace resistance, countermand this increase in the number of bumps and pins. Thus, a continued bandwidth crunch into and out of the chip seems unavoidable without fundamentally new paradigms for packaging and on-chip architectures. Until such paradigms emerge, the industry will continue to see anomalous, challenging chips that have tiny dies in large packages with many high-speed I/Os.

SoC versus SiP

System-on-a-chip (SoC) integration has garnered much attention throughout the electronics industry. But considering the integration and cost drivers for SoCs, *system in a package* (SiP) is a more relevant term. Multitechnology process integrations, like flash or RF with logic, are occurring more slowly than originally anticipated, because they force tremendous technology compromises. Thus, high-performance memory coexisting with high-performance logic in the same carrier, for example, is often more cost-effective and yields better performance than an embedded-memory solution. Amortization of nonrecurring mask and design costs also suggests the reuse of a given silicon die within multiple (single-package) contexts. Despite technical hurdles such as the known-

good-die problem, multichip integrations appear to have a solid foothold. Stacking (for example, flipping memory atop logic) greatly expands the architectural solution space by adding a third dimension, and it also appears to be well established. Of particular interest in the multichip context is so-called *bumpless packaging*, which relegates package-level interconnects to wafer fab processes through coarse top-level traces. (For bumpless packaging, some issues still need to be clarified, notably whether current delivery is feasible with a lithography-based approach.)

THERE ARE SEVERAL other issues related to packaging that the design and test community should address. These include

- die area and other costs of decoupling capacitance,
- dependence of test cost on pin count and signaling speed,
- impact of packaging technology on IC reliability requirements, and
- integration of optoelectronic and RF elements.

Due to space constraints, I haven't addressed these issues here. However, it should be clear that packaging and assembly challenges are increasing and are urgently entwined with design and test challenges. Die package co-design requires solid new thinking from all parts of the design and test community. ■

■ Direct questions and comments about this department to Andrew B. Kahng, Computer Science and Engineering Dept., University of California at San Diego, 9500 Gilman Dr., MC 0114, La Jolla, CA 92093-0114; abk@ucsd.edu.

For further information on this or any other computing topic, visit our Digital Library at <http://computer.org/publications/dlib>.