THE HUBER MILLS digital power report
Powering the 21st Century
April 2002/Vol. 3 Issue 4
Published by Gilder Publishing, LLC

Packing Power

Amkor leads in the technology challenge of packing logic and power into chip-scale dimensions

"It's a great company," one analyst told us. "Great people. Strong revenues. Nice cash position. A gold-plated customer list." The company's only problem, he volunteered, is that it's "not a technology play." Packaging semiconductors "isn't rocket science," so investors don't get excited.

It is, and they should. If Amkor Technology's (AMKR) business is "just packaging," then so is Dell's, so is Nokia's, so is Intel's, for that matter. Intel "packages" transistors, resistors, and capacitors in layers of silicon, metal, and plastic. Dell packages circuit boards and disk drives. Nokia packages RF chipsets. The others design and optimize digital logic, for the most part. Amkor designs and optimizes the thermal, electrical, mechanical, and chemical end of the logic-packing enterprise.

And that's where a lot of the action lies these days. The package has to get power in, to drive the logic chip, and radio waves out, if the chip is to communicate without wires, and heat out as well, to save the whole package from meltdown. Call it "packaging" if you like, but we aren't talking cardboard boxes here.

You can't easily ship naked slivers of silicon out of a fab, so all chip manufacturers do a lot of packaging in-house. Then, most typically, they hand off well-sealed chips to the next-tier packagers, the assemblers of printed circuit boards, computers, and cell phones—companies that take the chips and put them into functioning devices. The package is the insect-scale interface between electrical power on one side, logic in the middle, and communications and heat on the other side. Building such interfaces well is a very high-tech undertaking indeed. And one that grows more and more important as chips get bigger, and as complete, functioning digital systems collapse into sealed, chip-sized packages.

Amkor is the leading independent player in the market for packaging integrated circuits and complete single-package systems. The company was incorporated (under a different name) in 1970 by James Kim, eldest son of H. S. Kim. The younger Kim founded Amkor to design semiconductor packages and provide semiconductor packaging services through a supply relationship with the elder Kim's chip manufacturing company, the Anam Industrial Company (AICL). (The Kim family still owns just under half of Amkor's common stock.) Based in Arizona, Amkor currently operates thirteen factories around the world, with state-of-the-art high-density 3-D packaging capabilities in South Korea, Japan, the Philippines, China, and Taiwan. Amkor acquired a solid foothold in Japan last year, when it acquired Citizen Watch's high-density packaging unit, following on a joint venture established in 2000 with Toshiba's assembly operations; Amkor will acquire those operations entirely as well within two years.

Amkor currently provides more than 30 percent of the worldwide, outsourced market for semiconductor assembly and testing, and 40 percent of the worldwide market for outsourced laminate-based ball-grid-array packages. It offers more than 10,000 variations in packaging—the semiconductor industry's broadest line—and assembles more than 4 billion semiconductor packages a year. Intel accounts for about 10 percent of Amkor's $1.5 billion in revenues. Among the other 200 or so companies that also outsource to Amkor: International Rectifier, Infineon, IBM, Toshiba, Ericsson, Vishay, NEC, Analog Devices, Sony, Bosch, Conexant, AMD, Agere Systems, Altera, LSI Logic, Motorola, STMicroelectronics, Texas Instruments, and Philips.

Packing Power and Logic

The dense packing of logic begins with the integrated circuit—smaller gates, packed more densely, onto larger chips. And we know what this kind of packing delivers: Faster logic. The smaller the gate,

the more gates you can put on the chip, and the faster you can run them. Just push that strategy along far enough, as Intel and others have done over the last three decades, and you transform an abacus into a Pentium.

It takes power to move the bits, however. Electric power has to move in, to charge up the capacitors that store the bits, and to flip the gates that permit one bit to interact logically with another. Power has to move out as heat, too, at the tail end of the storing and processing. Thus, the technologies of logic invariably come enveloped in the technologies of power.

The power must—ultimately—be as fast as the logic elements, which means about ten times faster than the chip's clock speed; that's about how fast a gate has to flip so as to appear always "on" or "off," and never in between. Moving currents at that speed—ramping them up and down as the loads require—is monstrously difficult. You can't begin to run serious current through most ordinary copper wires at those frequencies, not if you hope to keep control of it. At those speeds, every connection, every bit of wire, every piece of solder, becomes a tiny antenna. A large share of the power radiates out as radio waves. A jumble of electric harmonics form, and cascade back and forth through the wiring.

Power must—ultimately—be as fast as the logic

One part of the solution is to array very small, very fast power capacitors and other power circuitry very close to the fast-changing payload. A Pentium on a circuit board, for example, is surrounded by "decoupling" capacitors, mounted all around it like a school of pilot fish. These capacitors get some help—ironically—from the inherent electrical capacitance of the on-board memory, the DRAM—the relatively slow logic buffers the power transients just enough to make the chip as a whole stay properly powered. Within a few years, however—certainly by the time clock speeds exceed 10 GHz—decoupling capacitors will have to be built right into the chip at multiple points across its surface. External power circuitry alone will not be able to deliver power down into the deepest recesses of the chip fast enough to keep pace with a clock running at such speeds. The 72-watt 600 MHz Alpha chip already has a minuscule 250 nanoFarad decoupling capacitor on the chip itself.

Whether built into the chip or arrayed immediately outside it, decoupling capacitors store only a few cycles worth of charge, they must be constantly replenished by power circuitry further up the line. A cell phone's battery has to be recharged every day or two; the decoupling capacitors have to be recharged at about the clock speed of the chip they serve. The cell phone's charger can trickle a fairly steady charge into the battery over a couple of hours; the high-speed chip's charger, by contrast, has to respond fast to changes in the load that occur when the logic elements are suddenly called upon to recalculate a large spreadsheet, or to create a high-resolution graphical image.

This front-end power circuitry itself has to be very compact, and packed in very close, to make it fast enough to keep pace with the digital logic. Electricity is incredibly fast, but even millimeters of wires can add electrical inertia that makes the system far too slow. So at the front end, just before the power transmutes into logic, the power circuits have to shrink to nearly the same scale as the logic circuits they serve. And they have to be designed and engineered with extraordinary care and precision—perhaps as much, ultimately, as is required in the logic circuits themselves.

What we are looking at here, in short, is one of the most fundamental interfaces in physics, information theory, and high technology—the interface between power and logic, the boundary between "bits" on the one hand, and "electrons," or "waste heat," or "radio-frequency power," on the other. All bits are packets of electrons or photons, but we don't mistake a wall socket for the electrically charged contents of a RAM chip or a Pentium. Somewhere along the line, well-ordered "power" becomes "logic," and then, somewhere a bit farther down the line, much of the "logic" becomes chaotic "power" once again, in the form of waste heat. Where exactly? The north side of the boundary is defined by the decoupling capacitor and power circuits connected directly to them, immediately outside the chip, or in the chip's package, or built into the chip itself. The south side is defined by memory capacitors and logic circuits.

Unpacking Logic and Heat

Getting power in is important; equally important, is to get it back out, after it has turned into heat. This second power-logic interface—the thermal one—begins in the individual gates and capacitors, where

Authors Peter W. Huber, Mark P. Mills
Publisher Lauryn Franzoni
Designer Julie Ward
Subscriber Services Melissa McNally
Chairman George Gilder

The Digital Power Report is published monthly by Gilder Publishing, LLC. Editorial and Business address: 291A Main Street, Great Barrington, MA 01230. Editorial inquiries can be sent to: peterhuber@gilder.com or markmills@gilder.com. Single-issue price: $50. **To subscribe call 800-223-0491**, e-mail us at hmpowerreport@gilder.com, or visit our website at www.digitalpowerreport.com

currents decay continuously into heat. The heat then flows through the semiconductor, through its packaging, and on out from there. Fail to move it out fast enough, and your data will turn into noise, and then your chip will turn into a puddle of molten sand. Heat is chaos, the antithesis of logic.

The best way to get rid of heat quickly isn't to pack things in close together, it's to spread them far apart. Give each gate more elbow room, separate the power circuits from logic elements—do the exact opposite, in other words, of what you do to optimize the speed and performance of the chip's logic. There's no clever way to dodge this fundamental dilemma. Making the individual gate smaller does improve both logical and thermal performance—the gate can run faster and dissipate less power too—but the whole point of making the gate smaller is to pack more gates more tightly, and run them faster, and in the end, smaller gates make the overall thermal problems worse, not better.

Unpacking elements is also the easiest way to solve a second closely related problem—electromagnetic interference. The unsteady flow of current in a wire generates electromagnetic emissions that can corrupt the flow of current in other wires nearby. Electrical noise tails off rapidly with distance, however, so once again, the imperative here is to pack things less tightly, not more so.

Mechanical considerations supply yet another set of antagonistic factors that the package designer must accommodate. Smaller components and tight packing generally make structures less vulnerable to ordinary vibration and mechanical shock—pound for pound, insects are a stronger than elephants. But the chip is made even more secure by packing a thick shell around it too. The thicker the package, however, the worse it sheds heat. Protect the chip too well from physical assault, and it dies of brain fever instead.

System on a Chip

Up to a point, the best way to reconcile all these conflicting packing imperatives is to collapse things onto a single chip. By and large, single-chip solutions are faster, cheaper, stronger—better all around.

This is why the number of chips on a PC motherboard, or inside a cell phone, keeps dropping year by year, as formerly discrete components get folded into each other. Individual chips keep getting physically bigger, and incorporating more functionality. The substantial arrays of chips once required to provide multimedia functionality, or high-bandwidth communications channels, or memory and processor combinations, or optical-to-electron conversion, or air-bag actuators, are now either routinely implemented, or emerging as, single-chip solutions. There is a strong trend toward integrating both analog and digital components on a single chip—the analog circuitry generally providing certain timing and power and control functions, alongside the digital circuits that perform logic operations on the same substrate. RF circuitry is now being built directly onto logic chips too, for use in cell phones, high-bandwidth wireline systems (like phone-line DSL), and so forth. Performance rises, costs drop, and the power of the expensive PC of 1990 ends up in the disposable toy of 2000.

Power and thermal considerations aside, there are other limits to how far the SoC can be pushed

As gates get smaller and chips get bigger, more functionality gets packed into the single chip—the point where a chip ceases to be a "component" and becomes a "device." Chip-level integration has indeed been pushed to the point where it has earned its own acronym: System on Chip (SoC). SoC shipments have vaulted from a few hundred million units a few years ago, to over a billion today, and are rising rapidly. And SoCs have spawned an entirely new sub-industry of design houses and fabless innovators.

But the acronym conceals as much as it reveals. What does it take to qualify as a "system"? Not as much as one might suppose. Typically, it doesn't take anything in the way of power circuitry. Just assume the power; and define the "system" from there. That's assuming a lot. As more elements get packed on to a chip, and the chips themselves get bigger, it gets progressively harder to get the electric power in, and the heat out.

Several of the most important chip-packaging innovations of recent years have centered on those two challenges. Laminated organic packages have displaced ceramics (in the packaging of all of Intel's chips, for example) because the organics' superior dielectric properties substantially improve chip-level power distribution and signal transmission. "Flip-chip" designs mount the junction side (the active region) of the semiconductor face down on the substrate, and shorten the wiring that feeds power into the chip—this improves power delivery by lowering electrical resistance, and improves thermal conductivity even more, because the hot surface of the chip is now close to the heat-dissipating substrate on which the chip is mounted. A ball grid array (BGA) package architecture shortens the electron-carrying wires and improves heat transfer still more. A conventional lead-frame chip looks like a bug splayed flat, with long wire leads emerging like legs from all around the edges of the package. The BGA uses instead hundreds of tiny bumps of metal arrayed in a grid across the base of the

package; these "balls" bond directly to a mirror-image array of flat contacts on the printed circuit board.

Most of Amkor's business and high-tech expertise remains focused on the most basic packaging challenge—packaging individual chips, especially high-performance chips. But "there's really no such thing as a standard package anymore," Richard Groover, Amkor's VP of Emerging Products, told us. The package is the power/logic interface; both power and logic must be designed and optimized for specific functions; so the interface—the package—must be, too.

Amkor's cross-cutting expertise and enormous library of thousands of package designs provide opportunities for rapid customization and optimization not easily matched within specialized chip houses. Amkor's basic library includes designs that have been optimized to accommodate different power levels, semiconductor types (RF, ASIC, CPU), and operating environments, from children's toys to military tanks. The company designs standard surface mounts, thinner packages to fit in tiny handhelds, and high-performance thermal and electrical packages for hot ASICs and CPUs. Optimized for pagers, disk drives, and wireless devices, Amkor's Super BGA package has a copper layer to improve heat transfer, and an exposed pad that directly mates the packaged chip's substrate to the heat sink. Amkor builds packages with organic or ceramic substrates; with standard low-cost wire bonds; and ultra-high-density wire bonds. The company also builds packages that incorporate the technically demanding bump bonding critical for the 1,000 to 2,000 input/output contacts needed for high performance ASICs, DSPs, and graphic chips. The BGA bonds the package to the board; bump bonding is architecturally similar, but places the solder bumps on the silicon itself, and attaches the silicon die directly to a precise conductor array on the substrate.

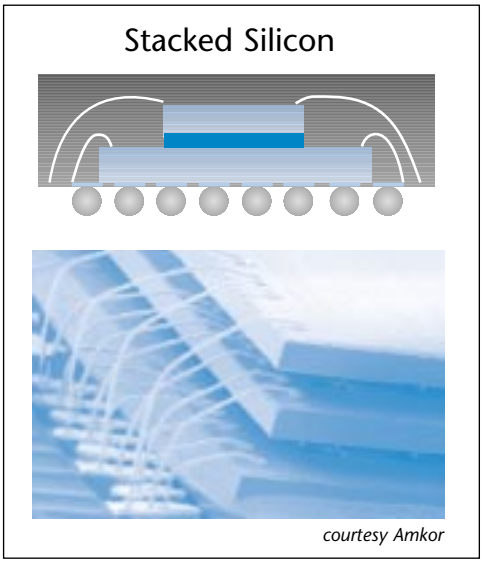

courtesy Amkor

Stacking and Packing Chips and Peripherals

Power and thermal considerations aside, there are other limits to how far the SoC can be pushed. Gearing up a chip fab to manufacture single-chip solutions often takes more time than fast-changing markets allow. For relatively low-volume, but lucrative custom markets, the high cost of mask design and layout can make a SoC prohibitively expensive. SoC designs raise other practical problems too; the bigger the chip, for example, the worse the yields at the chip fab, because any one flaw can kill the whole device. Some areas of silicon real estate are more expensive than others; cache memory built into the Pentium itself is a lot more expensive than the RAM chips that can provide gigabytes of memory alongside. And the more you pack into a single chip, the less flexibility you retain to vary the mix of processor, memory, graphics capabilities, and so forth. A graphics-intensive workstation requires a different mix of CPU, memory, and I/O chips than a server optimized to retrieve files or search databases.

A second set of SoC limits centers on the inescapable tensions between power and logic. In the delivery of more power, higher voltage and fatter wires are generally better than the alternative; the opposite is true in the pursuit of more digital logic. Smaller logic gates demand lower voltages, otherwise the current will punch through the thinner walls of the smaller element. But when more gates get packed closer together, they collectively require more power, and require it faster—and the only way to pump more power into less space, faster, is to raise the voltage just up the line. Much of the time, voltage considerations alone thus force the physical separation of power and logic devices, at least at any level above decoupling capacitors built into the logic chip itself.

Power/logic divisions spawn many serious chemical and physical antagonisms, too. Silicon carbide is an especially attractive material for power applications (*Quantum Power, Digital Power Report, May 2001*)—but it is much more difficult to work with, and therefore far from being a cost-effective option for building a logic device. Gallium Arsenide is the dopant of choice for building the key power components—the RF amplifiers—in cell phones, but gallium's chemistry and physical lattice are utterly incompatible with silicon. Indium Phosphide has many advantages in manufacturing devices that can project high-power millimeter waves (*The Power of Millimeter Waves, Digital Power Report, November 2001*)—but InP wafers are far more fragile than their gallium and silicon cousins, requiring different handling. Some semiconductors expand more than others when they get hot, which limits, once again, how close together they can be packed.

Electron/photon boundaries give rise to yet another set of challenges. An eyeball-on-a-chip needs a lens above

the semiconductor that transforms inbound photons into electric currents. An RF transmitter needs an antenna to transform currents into photons. But lenses and antennas can only be shrunk so far, and for the most part they can't be built right on to the logic or sensor chip.

So, when you reach the limit to packing more on to a single chip, you begin packing together separate chips, along with "passives" (capacitors, inductors, resistors), power bricks (*Cisco of the Powercosm?, Digital Power Report, May 2000*), power supplies, and other components. The printed circuit board (PCB) is the most familiar packaging scaffold used at this level of things. But the PCB's wires are tremendously long, and thus horribly slow, compared with those inside a chip. The sluggish speed of PCB wiring limits how fast you can move power into the logic elements, and how fast the chips can exchange bits.

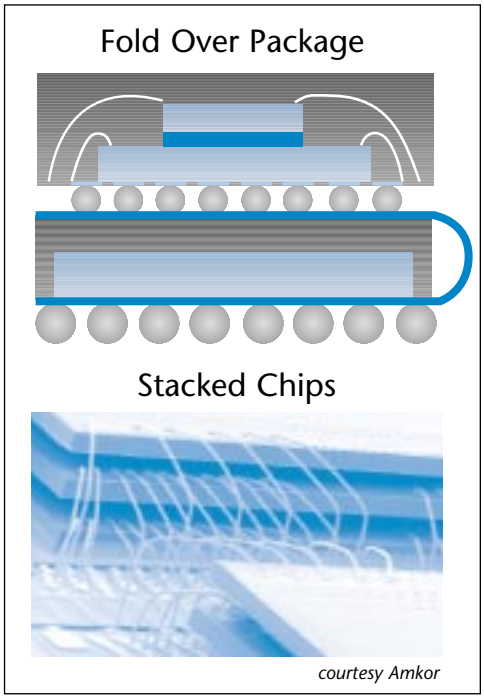

courtesy Amkor

Introduced in 1998 (by Sharp, an Amkor JV partner), the stacked chip-scale package (S-CSP) is a first response to these problems. Bare die from multiple chips—ASICs and memory chips, for example—are stacked up vertically and interconnected within a single very compact package; the PCB's suburban sprawl gives way to the high rise. Numerous companies now manufacture S-CSPs of various kinds for portable communication devices—among them Fujitsu, Hitachi, Mitsubishi, NEC, ASE, Toshiba, Dense-Pac Microsystems, and Amkor. Amkor's 3-D stack offerings incorporate many different combinations of wire bond, ball grid array, and flip-chip configurations. The company is both developing and exploring licenses from others for various folded-stack technologies that array chips vertically, in accordion-like configurations. One such approach puts chips on a flexible-PCB-like tape, which can then be rolled up into a "silicon burrito." These designs offer extremely compact packaging; they also present exceptionally difficult thermal challenges.

System in Package

Pack additional passives—resistors, capacitors, and inductors—into the same package and you now you have a System in Package (SiP). The passives are needed mainly to ensure supplies of sufficiently fast, clean power. A SiP may incorporate elements that run at up to five different voltages (3.3V, 2.8V, 2.4V, 1.8V, and 1.4V); passive electronic components within the package distribute the power accordingly. A tremendous amount of expertise must go into the architecture and electrical design of the package as a whole, to prevent current flows into one device from degrading the power that reaches its neighbor. Even larger components like micro lenses or antennas, radio-frequency shields, and batteries can migrate into the same casing as well.

S-CSP and SiP architectures shorten wires, save real PCB estate, boost speed, and improve the overall performance of the logic, for much the same reason that smaller, more tightly packed gates improve the performance of an integrated circuit. But the power-in and heat-out problems get worse. The 3-D structure of an S-CSP is inherently more difficult to power and cool than the 2-D alternative, precisely because 3-D structures are inherently denser. And by packing power circuitry in very close to the logic, a SiP moves more electromagnetic noise and waste heat right in close to where it can cause the most trouble, in the logic elements alongside.

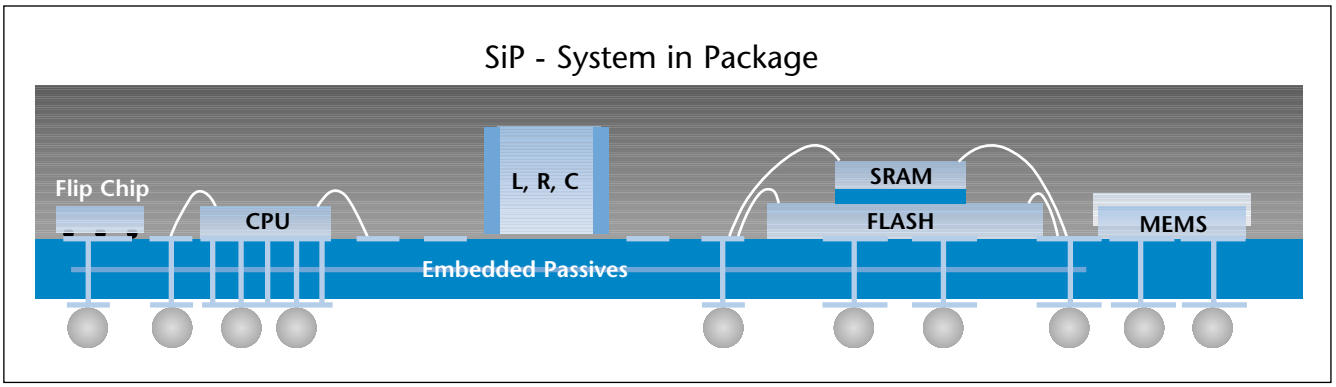

The more power you pack into individual chips, the more attractive the multi-chip SiP package becomes. If you're working with hundreds of chips and passives, you have to use a printed circuit board; but as the number of components falls, the single-package solutions grow increasingly feasible and attractive. SiP sales are starting to take off, blowing past and, of course, incorporating SoCs; over 3 billion units shipped last year, projected to reach 10 billion in a few years, with wireless telecom applications accounting for the biggest share.

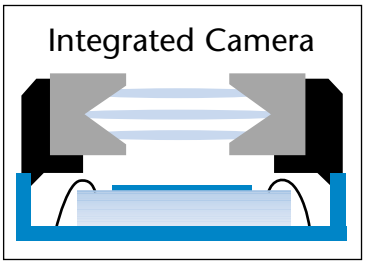

Amkor moved the SiP group out of its Emerging Products division in early 2000, to establish those operations in their own division. The target applications: Power conversion, high-speed mixed-signal chip sets, imaging/sensing, and automotive. The company has ramped up its SiP production capability at operations in the Philippines (memory cards, RF), Korea (networking, computing, and opto-electronics), Taiwan (sensors, cameras), and their China factory (cell-phone SiPs) later this year.

The more power you pack into individual chips, the more attractive the multi-chip SiP package becomes

A decade ago, the brick-sized cell phone was packed with a jumble of discrete components; today's cell phone is built around a much smaller handful of SoCs and passive and analog components; these are destined, in turn, to collapse into a few SiPs. Bluetooth is already a single SiP for short-range wireless communication networks. So too are GPS navigation systems. And Amkor is an established leader in manufacturing SiPs for the fast growing RF and wireless applications. These SiPs pack into a single device all the components required for a truly functional wireless telecom system: not just logic but RF filter capacitors and resistors, the transmit/receive switch, power amplifier, balancing circuits (transformers for the tricky RF noise reduction), flash memory, and a system clock (which is often one of the biggest off-the-die components).

Successful packaging at this level of density and complexity, especially for RF applications, depends on very close collaboration between Amkor and its customers. Amkor and Philips, for example, recently collaborated in the development of a high-power voltage regulation SiP for cell phones, that achieves higher performance in 60 percent less space on the shrinking cell-phone circuit board.

In networking and computing applications, Amkor has developed packages in which a flip-chip ASIC is mounted on the same substrate as surface-mounted memory devices, along with critical capacitors and other passives. All routing between the ASIC and memory is done on the first or second layer of the ASIC package substrate rather than on the motherboard. This level of ASIC/memory integration in a SiP significantly reduces both size and cost, and improves yield. The same architecture shows up in Amkor packages for ADSL, opto-electronics, Internet routers, and low power DC/DC converters. An xDSL chip set, for example, normally contains 4 chips and about 10 passive components; all collapse into a single 27 mm square SiP.

Sensors present a particularly important area of growth for SiP manufacturers, and Amkor is seizing the opportunity. High-resolution low-cost mega-pixel CMOS image sensors are now poised to move beyond digital cameras into PDAs, toys, and cell phones. Some 2 million camera-equipped phones were sold in Japan last year; Nokia is introducing one soon, as is Motorola. The market for CMOS SiP cameras in cell phones is forecast to break 50 million unit sales in a few years. It is the SiP that makes possible the solid-state eyeball, makes it cheap, robust, and easy to deploy ubiquitously—the total SiP camera market is expected to grow ten-fold in five years. Amkor's imaging SiP incorporates all the key power electronics, the IC, CMOS engine, and lens. Other SiP-based sensors include MEMs devices such as the accelerometers that are used to deploy air bags. Fingerprint recognition presents another huge opportunity. Until very recently, such systems remained prohibitively expensive. Working in collaboration with STMicroelectronics, Amkor manufactures a SiP that collapses all the key elements into a single package, the TouchChip

Raw storage capacity is collapsing into SiPs as well, and Amkor is again a player. Poker-chip-sized

System in Package

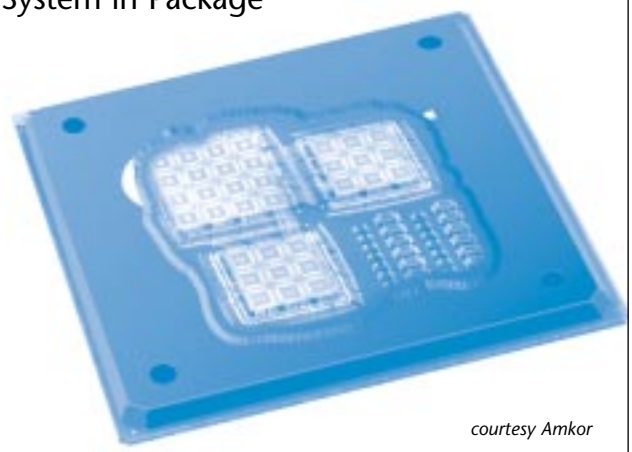

courtesy Amkor

memory cards now facilitate the sharing and transfer of data among digital cameras, audio players, computers, and other platforms. Amkor's MultiMedia Card (MMC) collapses the memory, critical ASIC and peripherals into a single, ultra-thin, highly reliable, and durable consumer-friendly package. Ultra-thin digital ID and credit cards are also in Amkor's pipeline; they will incorporate integrated circuits and support modest levels of programming. And they will be produced on a fully automated production line that will bring to high density SiPs the same economies of scale that chip fabs brought to CPUs.

Because "packaging" is such a loosely defined term, Amkor can be said to face off against countless competitors around the globe too, including divisions within major chip companies. At the leading edge however—in the realm of the SiP—the competition thins out quite a bit. Only a handful of companies, including Pacific Rim player MicroRoutes, Alpine MicroSystems, and Switzerland's Valtronic, rank alongside Amkor as leading SiP manufacturers in the global arena. In the United States there are also the small but impressive DPAC Technologies and Irvine Sensors (the latter with military-grade ultra-high-density system-in-a-cube), ChipPAC, and Tessera Technologies, a pioneer of micro-BGA and folded packaging.

Packing Power

The flash of genius that led to the integrated circuit was the thought that other components—capacitors and resistors in particular—could be etched on to the same sliver of semiconductor as the transistor gates. But there are definite limits to how far the single-chip gate-packing can be pushed. Thermal, mechanical, and chemical limits; engineering limits; and economic limits, centered on the simple fact that with complete systems, one size doesn't fit all. The packaging company takes over wherever the single-chip solutions reach these inevitable stopping points.

If this is more apparent today than it used to be, it is only because digital logic has grown so cheap, and we are buying so much more of it, and packing it into so many more places where chips never used to land before. There was no demand for wireless SiPs when wireless markets were small, wireless service was expensive, and wireless phones required hundreds of discrete components scattered across multiple circuit boards. Now that the phone has been reduced to a much smaller set of much larger chips, the inevitable next step is to dispense with the circuit board entirely, and pack them into an even smaller, denser SiP. The old, smaller, simpler chips spawned the printed circuit board, which was—and remains—the right packing structure for systems that require hundreds of discrete devices and components. The new larger, more complex chips spawn the SiP, which is fast emerging as the optimum packing structure for systems that can be built out of smaller sets of larger chips.

The packaging company takes over wherever the single-chip solutions reach these inevitable stopping points

"The technology boundaries between semiconductor technology, packaging technology, and systems technologies are blurring," observes the Semiconductor Industry Association in the 2001 edition of its *Technology Roadmap for Semiconductors*. Just a few years ago, packaging was little more than an aside in the SIA's annual reports. It isn't now. "Packaging technology is now a critical competitive factor, as it affects operating frequency, power, complexity, reliability and cost," the SIA concludes. Cost and manufacturing considerations will always be very important, the SIA notes, but there are fundamental technological advances to be pursued as well. Among the key goals: improve organic substrates to support higher frequencies; improve and augment the use of embedded passives; address the challenges associated with very high frequency (above 10 GHz) performance. "Dramatic improvements in materials properties—to address high frequency, higher power density, and increased mechanical stress—will be required." Thermal solutions must continue to focus "on complete integration from the chip through the system." SiP technology, which only began to be commercialized two years ago, is now the fastest growing segment of packaging.

In the end, the packing of power and logic is a game of Russian dolls, each layer enveloped by the next. In every layer, smaller, faster, denser is better than the alternative. And in every layer, smaller, faster, denser requires more power pumped in faster, upstream, and more heat, sucked out faster, downstream. With power and logic alike, the race goes to the swift, and the swiftest are the smallest.The smallest are the ones that are most densely packed.

Yes, Amkor is a "packaging" company. But the packaging of semiconductors has nothing in common with the packaging of strawberries. The chip-level packagers are the front-end managers of digital power and its byproducts. The in-house packaging operations of major chip companies will continue to play major roles in this market. But on the strength of its very sophisticated technology, Amkor is well positioned to emerge as the dominant independent provider of power-handling solutions on the doorstep of the microcosm.

Peter Huber and Mark Mills
April 2, 2002

The Power Panel

Ascendant Technology	Company (Symbol)	Reference Date	Reference Price	4/02/02 Price	52wk Range	Market Cap
System Integrators**	Amkor Technology (AMKR)	4/2/02	21.85	21.85	9.00-26.24	3.5b
	Emerson (EMR)	5/31/00	59.00	56.88	44.04 - 72.09	23.9b
	Power-One (PWER)	4/28/00	22.75	7.99	5.32 - 27.35	630.0m
Electron Storage & Ride-Through	Wilson Greatbatch Technologies (GB)	3/04/02	25.36	26.32	17.26 - 39.00	546.9m
	C&D Technologies (CHP)	6/29/01	31.00	21.00	16.35 - 38.60	548.0m
	Maxwell Technologies (MXWL)	2/23/01	16.72	9.25	5.81 - 22.50	94.0m
	Beacon Power (BCON)	11/16/00	6.00*	0.56	0.50 - 8.20	23.9m
	Proton Energy Systems (PRTN)	9/29/00	17.00*	6.30	4.00 - 15.12	209.3m
	Active Power (ACPW)	8/8/00	17.00*	5.04	3.13 - 30.20	206.0m
Project, Sense, and Control	Danaher Corp. (DHR)	1/29/02	61.56	70.47	43.90 - 74.25	10.5b
	FLIR Systems (FLIR)	1/9/02	41.64	48.54	7.75 - 59.50	807.2m
	Analogic (ALOG)	11/30/01	36.88	42.76	33.40 - 50.00	564.4m
	TRW Inc. (TRW)	10/24/01	33.21	51.50	27.43 - 53.00	6.5b
	Raytheon Co. (RTN)	9/16/01***	24.85	41.75	23.95 - 42.30	16.6b
	Rockwell Automation (ROK)	8/29/01	16.22	19.89	11.78 - 47.20	3.7b
	Analog Devices (ADI)	7/27/01	47.00	43.52	29.00 - 53.30	15.9b
	Coherent (COHR)	5/31/01	35.50	34.56	25.05 - 45.55	991.1m
Powerchips	Cree Inc. (CREE)	4/30/01	21.53	13.90	12.21 - 36.65	1.0b
	Microsemi (MSCC)	3/30/01	14.00	16.11	10.90 - 40.10	459.2m
	Fairchild Semiconductor (FCS)	1/22/01	17.69	28.24	11.86 - 30.50	2.8b
	Infineon (IFX)	11/27/00	43.75	21.75	10.71 - 44.40	15.1b
	Advanced Power (APTI)	8/7/00	15.00	12.45	6.50 - 18.00	108.6m
	IXYS (SYXI)	3/31/00	6.78	11.73	4.27 - 19.45	314.7m
	International Rectifier (IRF)	3/31/00	38.13	44.70	24.05 - 69.50	2.8b
Network Transmission	ABB (ABB)	9/29/00	24.24**	7.95	6.10 - 18.95	9.4b
	American Superconductor (AMSC)	9/30/99	15.38	7.84	6.50 - 27.90	160.4m
Distributed Power	General Electric (GE)	9/29/00	57.81	37.10	28.50 - 53.55	368.6b
	Catalytica Energy Systems (CESI)	9/29/00	12.38	3.52	3.25 - 24.00	61.1m
	FuelCell Energy (FCEL)	8/25/00	24.94	16.91	10.48 - 46.72	661.4m
	Capstone Turbine Corp. (CPST)	6/29/00	16.00*	3.82	2.75 - 38.25	294.9m

Note: This table lists technologies in the Powercosm Paradigm, and representative companies that possess the ascendant technologies. But by no means are the technologies exclusive to these companies. In keeping with our objective of providing a technology strategy report, companies appear on this list only for the core competencies, without any judgment of market price or timing. Reference Price is a company's closing stock price on the Reference Date, the date on which the Power Panel was generated for the Digital Power Report in which the company was added to the Table. All "current" stock prices and new Reference Prices/Dates are based on the closing price for the last trading day prior to publication. IPO reference dates, however, are the day of the IPO. Though the Reference Price/Date is of necessity prior to final editorial, printing and distribution of the Digital Power Report, no notice of company changes is given prior to publication. Huber and Mills may hold positions in companies discussed in this newsletter or listed on the panel, and may provide technology assessment services for firms that have interests in the companies.

** Offering price at the time of IPO.*

*** Effective April 6, 2001, ABB was listed on the NYSE. The reference price has been adjusted to reflect this change. The 52-week range covers the period from April 6, 2001 only.*

**** The October 2001 issue closed on September 16, 2001 and was posted at 8 a.m. on September 17, 2001. Due to the markets' close in the week after September 11, our reference price reflects Raytheon's closing price on September 10, 2001.*

***** The former categories of "Silicon Power Plants" and "Motherboard Power" have been rolled into this new broader class of technologies.*

More information about the Powercosm and its technologies is available on www.digitalpowerreport.com