

# The Lowdown on High-Rise Chips

George Lawton

**A**s today's silicon technology reaches its technical limits in many areas, developers are looking for new ways to design better chips. Typically, they look for ways to make chips faster, more energy efficient, and cost effective, while still using current fabrication techniques and materials.

Traditionally, manufacturers have added functionality or increased performance by making chips bigger, shrinking transistor sizes, adding dopants, or using exotic materials instead of silicon. But each of these approaches has limits: Larger chips have lower per-wafer yields; reducing transistor sizes faces engineering obstacles, particularly as gates approach the size of the electrons that must pass through them; and dopants and exotic materials can be expensive and create manufacturing challenges.

In light of this, companies are working on high-rise, multilayer chips. Instead of laying out the vast majority of a chip's circuitry on a single layer of silicon, they are making chips with circuitry arranged vertically as well as horizontally.

One key to high-rise chips' improved performance is that all of the wiring is closer together than on typical flat processors, where circuitry can be separated by the entire length of a large substrate. Closer wiring permits the use of short interconnects running throughout the chips' many layers. The ability to use short interconnects enables



higher clock rates, lowers energy usage, and reduces overall chip sizes.

John Trezza, chief technology officer of chip designer Xanoptix, said high-rise technology has made his company's optical transceiver faster than modules based on flat chips.

While researchers have experimented with high-rise chips in the past, companies are now beginning to sell them, initially for networking and memory applications. Proponents say high-rise memory chips could make digital recording media inexpensive and convenient enough to replace photographic film and audiotape.

Irvine Sensors is producing stacked-chip assemblies for miniaturized sensors, cameras, and image-processing applications. Matrix Semiconductor is now selling high-rise memory/data-storage chips, as shown in Figure 1. Xanoptix is selling a vertical chipset for optical networking. And Ziptronix is producing a high-rise networking chip.

Eventually researchers hope to use the technology to support other types

of chips, including field-programmable gate arrays.

But the technology is still in early phases, and manufacturers must overcome several hurdles including design issues, thermal constraints, and competition from alternative approaches.

## INSIDE HIGH-RISE CHIPS

Associate Research Professor James Jian-Qiang Lu, who leads wafer-level high-rise-chip research programs at Rensselaer Polytechnic Institute (RPI), said the industry has pursued vertical-chip approaches since the 1980s.

Manufacturers have already been making forays into multilevel chip processes. For example, Intel's Pentium 4 uses several layers of wiring, which places the transistors used for logic and memory closer together. Local wires at lower chip levels, next to the transistors in the substrate, carry signals from nearby transistors within a circuit or functional block, such as a video decoder. Long-distance interconnects in the upper chip layers carry signals from transistors spaced far apart and thus sometimes experience latency.

This isn't a true high-rise chip because the transistors and diodes are only on the silicon substrate.

Some manufacturers, such as Irvine Sensors, have made *3D packages*, in which different functions, such as memory and logic, are put on different chips, which are wired together only at the edges.

In high-rise chips, on the other hand, a larger number of interconnects pass through *vias*, tiny holes in the circuitry layers, enabling high performance levels, said RPI Professor Ronald Gutmann.

Lu said the use of thin interconnects could enable up to half a million wires between layers while the vias would use only 3 percent of the silicon area, leaving room for transistors, diodes, wiring, and other elements.

These wires connect the various elements that would normally be spread out across a flat chip but that in high-rise architectures are much closer

together. And because the elements are closer together, the interconnects between them can be shorter than on traditional chips.

This increases clock rates because signals don't have to travel so far. Lu said studies have shown that with current flat-chip trends, the signal delay across the length of a chip will be more than two clock cycles by 2008 and greater than four clock cycles by 2014. When the delay is more than one clock cycle, chips need additional components such as repeaters to manage signals so that they arrive in synch with clock cycles. High-rise chips can eliminate or minimize such problems.

High-rise chips' shorter interconnects can also reduce power usage, again because signals don't have to travel far. Because of this, they can use smaller, less powerful, less expensive transistors to send signals.


In addition, because many of their elements rise above the silicon substrate, the chips are smaller horizontally. They thus use smaller substrates, which enables larger yields from a single wafer. And largely because they use smaller substrates, high-rise chips' overall package size is smaller than comparable flat chips.

This size differential lets high-rise chips fit more easily in smaller devices than flat chips of comparable power. This also lets manufacturers create high-rise chips that are more powerful than comparably sized flat chips.

Finally, each layer of a high-rise chip's circuitry can be optimized for a different technology, meaning a single chip could offer high levels of various types of functionality. As Figure 2 shows, this represents an alternative to system-on-chip approaches.

### HIGH-RISE-CHIP INTEGRATION

According to RPI's Lu, there are three main high-rise-chip integration techniques: multichip stacks, recrystallized silicon, and monolithic wafer-level 3D integration. Some approaches use variations or combinations of these.



**Figure 1.** This photograph clearly shows the circuitry layers in Matrix Semiconductor's four-level memory chip.

### Multichip stacks

A simple technique is a face-to-face interconnection of two chips. Using the *flip-chip-on-chip* approach, bond pads on each chip are lined up and soldered together to provide interconnections.


Multiple chips can also be stacked and formed into a unit by wire bonding the edges together, a technique that permits a few thousand interconnects.

Irvine Sensors is using a technique that entails punching several rows of vias near the chip's edges, through which up to 10,000 interconnects can pass.

The company is also looking at creating vias within the interior of the chip, noted Volkan Ozguz, Irvine's chief technology officer. Vias in more locations would allow shorter connections between functional elements on different layers.

### Recrystallized silicon

In the *recrystallized silicon* approach, chip makers grow the silicon on which the various layers of elements sit. On top of the substrate, they lay down polycrystalline silicon and use lasers or other techniques to recrystallize the



**Figure 2.** Irvine Sensors' Neo-Stacked computer system on a high-rise chip includes a CPU, circuitry for RAM and flash memory, and interfaces for an LCD display, mouse, keyboard, USB host controller, and other types of external input and output.

material and thereby create new silicon layers on which they can add transistors and other elements. Researchers originally developed this technology to grow the display elements on top of the controlling circuitry in flat-panel LCD screens.

Stanford University Associate Professor Thomas Lee, director of advanced development at Matrix, said the company has started producing 3D memory chips in quantity using this technique.

He said this approach is less expensive than other high-rise chip techniques because growing new silicon layers is less costly than creating wafers or dies in fabrication plants and bonding them together. Also, the layers of grown silicon are thinner than the substrates of dies typically bonded together in high-rise chips. They thus yield smaller high-rise chips with more densely packed circuitry.

However, the technology is limited because electrons don't flow quite as easily in the polysilicon, due to its molecular structure, as they do in single-crystal silicon. This slows performance. Manufacturers can't solve this problem by growing single-crystal silicon because it requires high temperatures that would damage circuits on the chip's bottom layer.

### Monolithic wafer-level integration

Monolithic wafer-level integration delivers high performance for logic and memory applications while still using existing fabricating technologies, said Kathryn Guarini, a researcher at IBM's T.J. Watson Research Center.

With this technique, manufacturers glue entire wafers, rather than individual dies, together and run wiring through vias between the layers. The wafers are then cut into multiple high-rise chips.

Bob Markunas, Ziptronix's vice president of market development, said wafer-scale processing offers greater economies of scale than chip-level integration because an entire wafer's worth of chips can be handled in one process.

This will yield even more benefits as wafers get bigger. For example, Markunas said, as wafers grow from 8 to 12 inches in diameter, the number of chips produced will grow 225 percent, while the processing cost will increase only 40 percent.

However, bonding entire wafers can cause problems. Typically, numerous dies on any given wafer are bad. When a bad die is connected to a good one, the entire high-rise chip is bad. This

becomes a bigger problem as the number of layers on a chip grows.

IBM, the Massachusetts Institute of Technology, RPI, and other institutions are pioneering a promising approach in which they remove excess silicon from the bottom of a wafer or die via chemical-mechanical polishing (CMP) so that elements can be stacked closely together with less intervening silicon. Reducing the amount of silicon used permits smaller chips. Also, less silicon reduces thermal resistance and enables chips to more easily dissipate heat.

**High-rise chips promise to be faster, smaller, and cheaper.**

In CMP, a wafer or die sits on a pad. Spindles rotate the pad and wafer in opposite directions while abrasives and reactive alkaline chemicals pass between them, removing excess silicon and flattening the surface.

Monolithic wafer-level integration and recrystallized silicon promise the highest performance gains because they use less silicon above the substrate and thus let manufacturers pack more circuitry into smaller areas.

### OBSTACLES TO RISE ABOVE

High-rise-chip technology faces a number of obstacles to success. For example, adoption will require the development of new electronic design-automation tools optimized for high-rise chips. The tools will use many of the same principles found in tools for single-level chips but will still need some adaptations, explained Rafael Reif, head of MIT's Department of Electrical Engineering and Computer Science.

### Construction complexity

The integration and fabrication of high-rise chips' many elements are two key areas of complexity for their large-scale production, noted Irvine's Ozguz.

Another issue is that with high-rise-

chip integration, interconnects get shorter. However, Gutmann noted, shrinking the wires and making them thinner threatens signal integrity because at high clock frequencies, signals can leak from them. Also, lithography becomes more complex with shorter, thinner interconnects.

And chip makers are used to working with I/O drivers designed to drive interconnects on traditional chips that are much longer than those on high-rise chips.

In addition, Lu said, for wafer-scale integration, fabricators will have to introduce four new steps at once: wafer alignment, bonding, thinning, and interconnect formation. Doing this successfully could be complex and expensive.

### Heat and interconnect issues

A key challenge for high-rise chips is dissipating the heat from the closely stacked circuitry. In traditional chips, heat is carried up through the surfaces of the relatively conductive elements. In high-rise chips, though, the silicon between layers of elements creates substantial thermal resistance and thereby reduces heat dissipation.

Reif said heat buildup could make high-rise approaches impractical for chips—such as CPUs and graphics and networking processors—that deal with particularly demanding applications.

One strategy for coping with heat buildup is using some of the copper wires running through the vias to conduct much of the heat away from the chip.

Meanwhile, the more layers a chip has, the more silicon is lost to vias for interconnects, and this reduces the area available for transistors, diodes, and circuitry, Reif said. And the more layers a chip has, the more difficult and expensive it becomes to deliver power effectively to all elements.

These factors limit high-rise chips' height.

Thus, manufacturers are considering chips with up to four layers for most applications, although, Gutmann noted,

some researchers have built proof-of-concept chips with dozens of layers.

### Competition from alternative approaches

Users may not quickly switch to high-rise chips, which would require new manufacturing approaches, if researchers improve traditional flat-chip technology. Two industry projects are designed to accomplish this.

First, the X Initiative's ([www.xinitiative.org](http://www.xinitiative.org)) X Architecture would let designers connect transistors at 45-degree, in addition to the customary 90-degree, angles. This would shorten wire lengths.

Second, Sonics Inc.'s network-on-a-chip approach would put high-bandwidth, low-latency signals on short interconnects, thereby reducing delays. NoCs would use the longer wires only for low-bandwidth signals that better tolerate latency.

**A**lthough high-rise-chip adoption is just beginning, manufacturers are already looking at new ways to work with the technology. For example, in cases where various components—such as logic and memory—are different sizes and thus won't easily bond, Ziptronix is looking at ways to attach various-sized dies to a larger base layer.

According to Irvine's Ozguz, high-rise-chip technologies are still in the early stages of development and probably won't be popular for many years. Based on current research initiatives, he said, the chips probably won't be used widely for 10 to 20 years, although some aspects of the technology will be adopted sooner.

However, predicted Reif, serious use of high-rise chips will begin in perhaps five years, once it is less expensive to bond multiple dies than to simply package them together.

Referring to high-rise chips, Sonic CEO Grant Pierce noted, "That is pretty exotic stuff and takes a long time for acceptance."

Nonetheless, Lu said, the chip industry is increasingly interested in high-rise chips. He concluded that manufacturers are transitioning from asking why they should make high-rise chips to asking what types of high-rise chips they should build. ■

*George Lawton is a freelance technology writer based in Brisbane, California. Contact him at [glawton@glawton.com](mailto:glawton@glawton.com).*

Editor: Lee Garber, *Computer*,  
[l.garber@computer.org](mailto:l.garber@computer.org)

To receive regular updates, email

[dsonline@computer.org](mailto:dsonline@computer.org)

**VISIT IEEE'S  
FIRST  
ONLINE-ONLY  
DIGITAL  
PUBLICATION**

IEEE

**distributed systems**

ONLINE

Expert-authored articles and resources

*IEEE Distributed Systems Online* brings you peer-reviewed features, tutorials, and expert-moderated pages covering a growing spectrum of important topics:

**Security**

**Grid Computing**

**Mobile and Wireless**

**Middleware**

**Distributed Agents**

[dsonline.computer.org](http://dsonline.computer.org)