MICROMACHINING SCULPTS SENSORS IN SILICON

BUILDING SPARC'S FAST FAMILY OF WORKSTATIONS

ELECTRONIC DESIGN®

A PENTON PUBLICATION U.S. $5.00

OCTOBER 26, 1989

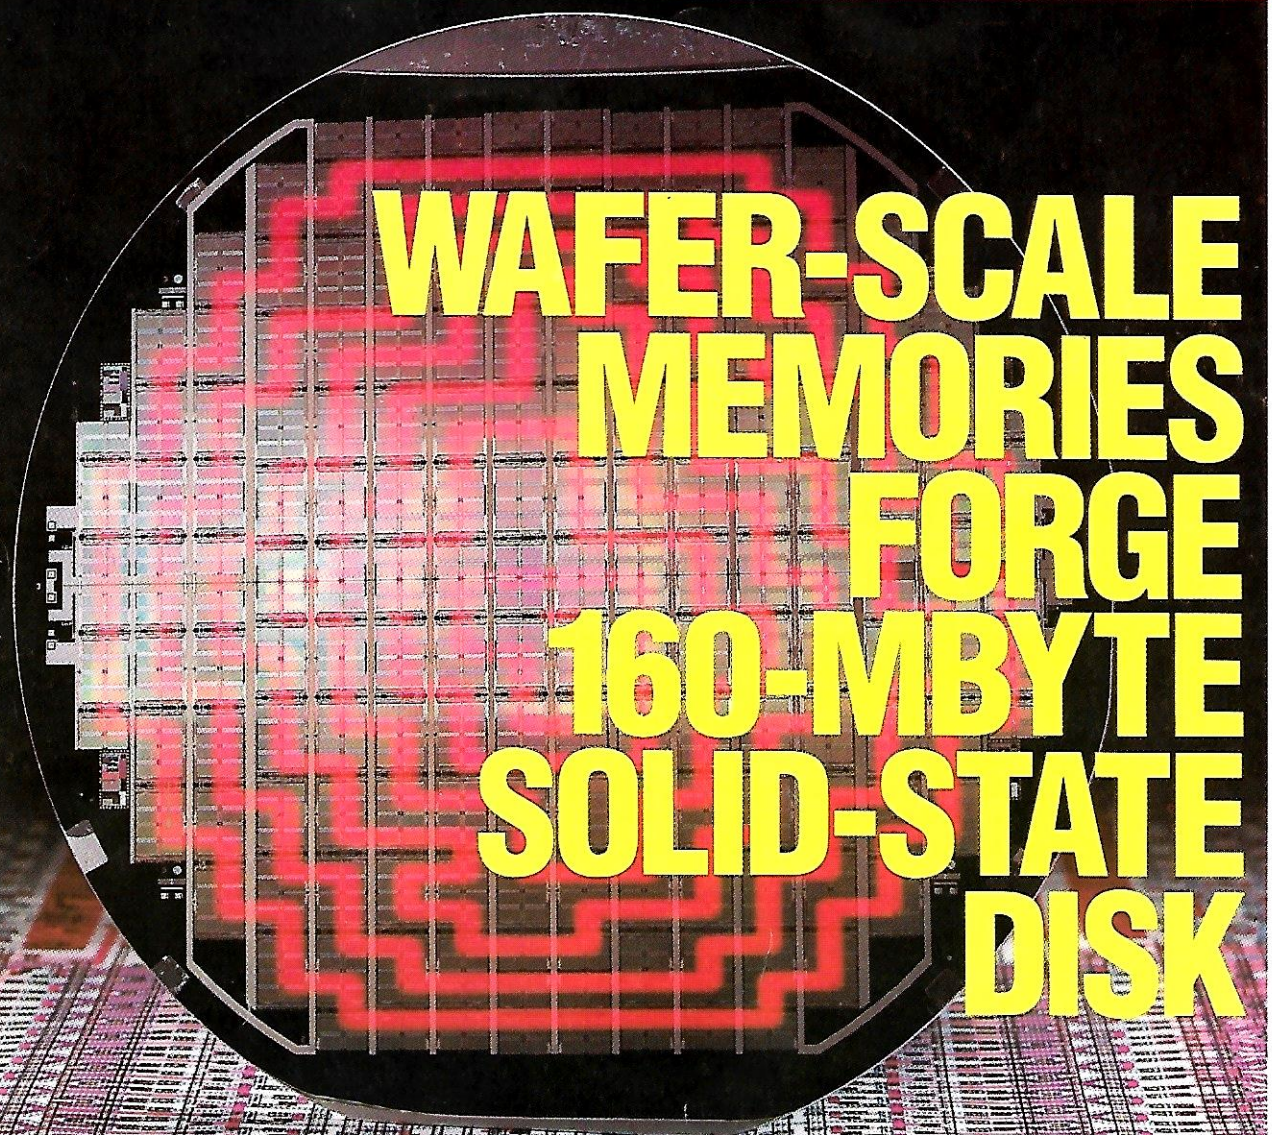

WAFER-SCALE MEMORIES FORGE 160-MBYTE SOLID-STATE DISK

SPECIAL REPORT: THE EMERGING OPTICAL-DISK DRIVE

ANAMARTIC

ELUSIVE SOLID-STATE MEMORY CONCEPT COMES TO FRUITION IN 160-MBYTE STORAGE MODULE USING 6-IN. WAFERS.

WAFER-SCALE INTEGRATION ARRIVES IN "DISK" FORM

LAWRENCE CURRAN

A dream that has tantalized system designers for more than 20 years—the use of full uncut wafers instead of diced and packaged semiconductor devices as logic or memory—has reached commercial reality. Hailing from Anamartic Ltd., a British firm, Wafer Stack is a solid-state storage subsystem that takes the place of conventional rotating disk drives in a computer system.

Wafer Stack enables computer system designers to realize extremely fast access times compared to conventional disk drives—an average of 200 μs with Anamartic's proprietary native-mode interface, or less than 1 ms using a SCSI interface. That's an improvement of about 200 times compared to rotating disk drives with their typical 20-ms SCSI access times.

Such a boost makes the Wafer Stack attractive for use in various computers, especially high-speed on-line transaction-processing (OLTP) systems used in banking, insurance, manufacturing, and communications applications. Those systems make frequent accesses to short blocks in a database. In fact, one of the companies that invested in Anamartic is Tandem Computers Inc., the $1.3-billion Cupertino, Calif. manufacturer of OLTP computers. Tandem is evaluating the Wafer Stack, with an eye toward having it in the product line within 18 months, according to Larry Laurich, Tandem's vice president for engineering. He reports that Tandem is very excited about the Wafer Stack based on early results to date. "Our team is writing to and reading from devices in the lab, and is knocking off benchmarks on schedule," Laurich says.

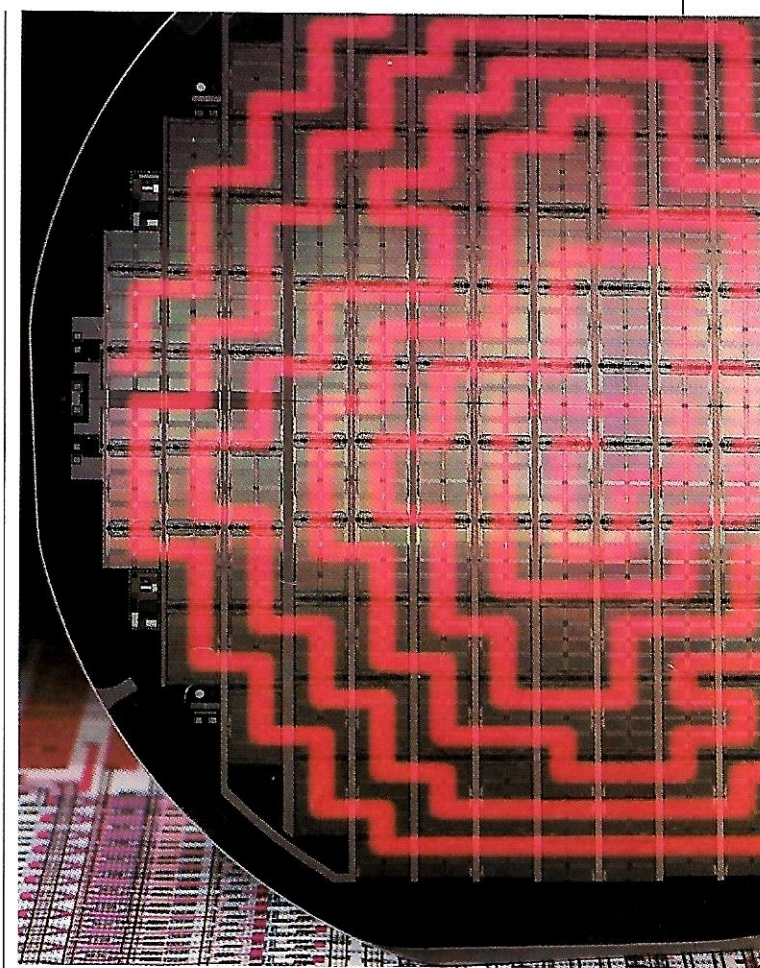

The Wafer Stack fills the gap between ever-faster semiconductor main memory and slower disk drives, a bottleneck that has existed for the last two

COVER: SOLID-STATE DISK

decades. It cracks that bottleneck because of its inherent zero-latency performance, resulting in extremely fast access times and sustained transfer rates ranging from 1.6 Mbytes/s for a native-mode interface to 3 Mbytes/s for a SCSI interface. Conventional disk drives offer similar sustained transfer rates, but it's in the access times that the Wafer Stack shines.

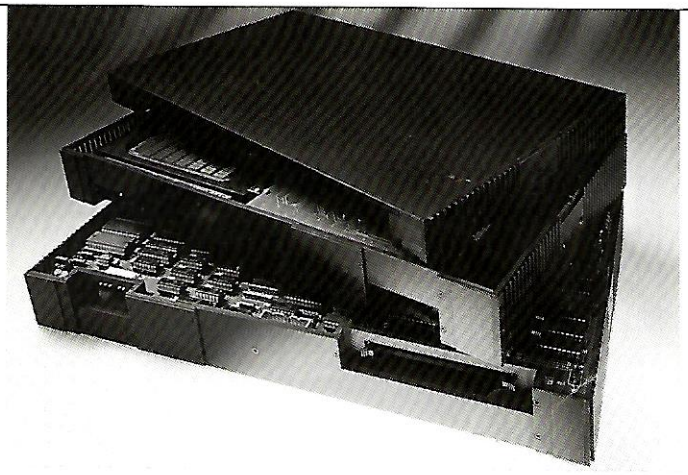

1. TWO WAFERS FACE EACH OTHER in a clamshell-like arrangement to form a 40-Mbyte storage module in Anamartic's Wafer Stack.

The Wafer Stack uses two 6-in.-diameter wafers to form one module containing 40 Mbytes of storage in 1-Mbit CMOS dynamic RAMs (DRAMs) *(Fig. 1)*. As many as four such modules can be stacked atop one another to supply a maximum of 160 Mbytes of storage. A wafer controller board is also required to manage data storage and transfer, error correction, and "scrubbing" on the wafers.

John Scandalios, vice president for marketing at Anamartic, explains that scrubbing assures that any memory cells developing random errors are purged and replaced, even if they occur long after the system was working at a customer site. One controller can handle from one to four modules, with the entire stack based on the form factor of an 8-in. disk drive—8.5-in. wide by 5-in. high by 15-in. deep.

The standard Anamartic native-mode interface can be augmented by a SCSI or another interface added by a buyer—an OEM or end user. When connected to a host computer through a SCSI interface, the Wafer Stack emulates a disk drive.

Although the Wafer Stack fits into the same footprint as a conventional 8-in. Winchester disk drive, the similarity ends there. Anamartic's considerable contribution comes inside the box housing the Wafer Stack. Where other solid-state disks use individual memory chips wired to a pc board, Anamartic's wafers are mounted to a carrier intact. Using the full wafer eliminates several processing steps, including as much as 90% of the costly wiring and soldering, which can reduce wafer yield.

There are 202 dice per wafer, which is more than necessary to obtain 20 Mbytes of data per wafer. The additional devices are included to allow for the loss of device portions that test out as flawed, and to supply spares to replace those failed elements. In addition to containing a 1-Mbit DRAM, each die also carries additional programmable configuration logic, which Anamartic calls Conlog, to connect each Conlog element to its four neighboring dice to the north, east, south, and west over signal lines that form logic networks on the wafer *(Fig. 2)*.

Using proprietary software, an external controller tests each die, then programs the Conlog elements to interconnect the good dice into one continuous bidirectional data path. The path takes the form of a spiral, running counterclockwise from the wafer's edge to the center *(Fig. 3)*. Portions of bad dice, but not the entire DRAM, are routed out of the spiral by the software. Anamartic partitions each DRAM into 32 tiles of 32,000 bits each, so that only a failed tile is discarded. In conventional practice, if just one bit of a 1-Mbit DRAM is flawed, the entire device has to be scrapped.

PROM Holds Map

Each wafer-carrier board holds a CMOS flash PROM. The PROM holds the map of the Conlog spiral, including the locations of failed tiles and spare memory cells. It's called into play by the controller to access only good cells after the Conlog is established. It also activates spare cells if random errors are identified during scrubbing, long after the Wafer Stack has been in system use.

David Hall, chief executive officer of Anamartic in England, says that because of this ability to retain most of the memory tiles in a DRAM with some bad bits, "We can use about 80% of the dice on a wafer. This means that the Wafer Stack can be priced in volume at 60 to 75% of the price of equal-capacity solid-state disk drives based on traditional individual chips." He adds that conventional semiconductor manufacturing techniq... may result in as many as 75... the chips being rejected.

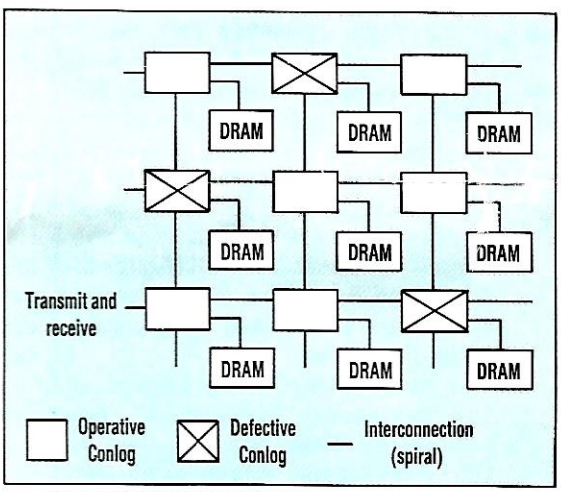

2. EACH DRAM DIE ...lso contains associated configuration logic, w... ...rammed to disable defective memory eleme...

The 1-Mbit DRAMs are produced with a 1.3-μm n-well CMOS process. The chips have dimensions of 13.65 × 4.4 mm, which is about 20% larger than a standard device, in order to accommodate the spare cells. The DRAMs are organized as 256-kword by 4-bit fast-page units, requiring a 52-μs refresh time.

The Wafer Stack is the result of a team effort. The proprietary interconnect software and associated Conlog logic were designed in Cambridge, England, at Anamartic's research and development labs. Wafer development is a joint effort involving Fujitsu, which is producing the wafers in Japan and distributing the product there. Besides U.S. and European sales and marketing, Anamartic's U.S. operation in San Jose, Calif. is also responsible for the engineering of customer interfaces to the Wafer Stack.

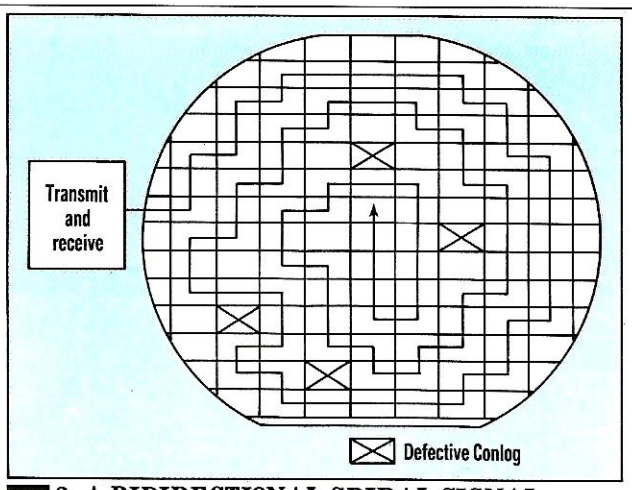

3. A BIDIRECTIONAL SPIRAL SIGNAL network runs from the wafer periphery to the center.

Like A Circuit Board

Anamartic's Scandalios says that the wafer carrier is similar to a conventional circuit board "that's a bit larger than the 6-in. wafer." A silicon nitride layer protects the top layer of the wafers, over which a silicon gel is applied, to protect the wire bonds from moisture. Two wafers are placed face-to-face in a hermetically sealed clamshell-like arrangement to form a 40-Mbyte module.

Tandem's Laurich characterizes his company's work with its early Wafer Stacks as prototype and alpha testing. "We don't have a full configuration yet, but we expect deliveries of those in the next 30 days or so." The early experience prompts him to say that he's very excited about the technology and happy to work with Anamartic. "It's been a real team effort. We started with an essentially undefined project, but we've been knocking off the benchmarks quite well since getting it defined," Laurich says. "If things continue to go the way they have, 12 to 18 months seems to be a logical timeframe to have the Wafer Stack in a Tandem system."

Laurich isn't ready to elaborate on the benchmarks, but points out that Tandem plans for the Wafer Stack to eventually replace the company's V80 and XL80 Storage Facilities, which serve Tandem's own OLTP systems and are sold for use with other computers. The V80 holds 265 Mbytes of formatted data for each 8-in. disk drive. Eight of these disk drives can be put into a system to supply 2.1 Gbytes of formatted data. The XL80 uses 9-in. disk drives to supply a maximum of 7.16 Gbytes of formatted data. This large amount of storage is obtained by ganging eight disk drives.

A major attraction for Laurich is the "no-latency/fast-access" feature of the Wafer Stack applied to OLTP systems, which typically make frequent accesses to short blocks of data in a database. The database can be partitioned to take advantage of those frequent accesses to greatly boost system throughput. Finally, Laurich is convinced that although the Wafer Stack storage subsystem carries a high initial cost compared to magnetic mass storage, it still holds the promise of being at least 50% cheaper than it would be to use standard 1-Mbit DRAMs in a conventional solid-state disk.

So-called "solid-state disks" using dice that were separated from the wafer and mounted to substrates or circuit boards are currently being used by a few computer companies. But 20 years of effort to come up with a viable commercial approach to true wafer-scale integration hasn't succeeded until now.

A 20-Year Effort

The most notable recent attempt, a project at Trilogy Ltd. that also involved development partners, failed to yield a viable product despite an investment estimated at about $100 million. Research and development into the concept goes back more than 20 years.

Texas Instruments Inc. had a project in the late sixties (funded by the then Avionics Laboratory at Wright-Patterson Air Force Base) that applied "discretionary wiring" in an airborne computer. The idea was to leave the top metallization layer of a wafer uncommitted until the wafer probing stage at which good devices are identified. Then a mask was to be created that interconnected only those good dice with final metal. However, the effort didn't evolve into commercial production. That dream was left for Anamartic to realize with the Wafer Stack storage subsystem.□

Price And Availability

A 40-Mbyte Anamartic Wafer Stack module sells for $11,680. Four modules can be stacked to deliver a 160-Mbyte configuration priced at $28,760. Prices include a wafer controller and SCSI interface. Limited quantities are available now.

Anamartic Inc., 1054 Saratoga/Sunnyvale Rd., San Jose, CA 95129; (408) 973-8088.

ANAMARTIC

Product of the Year Awards

boundary chip). Under external command, it forms a transmit link to one good neighbor and a receive link from another. If its DRAM is not usable, it shuts off the local V_{CC} and makes no links.

The links start at an edge cell and follow the edge of the area of good cells, making a rough spiral to the center (see *inset in photo*). Because the links only go to neighbors with a common boundary, but not diagonally, some good cells become blocked by bad ones.

Cells that have a few bad addresses, but work otherwise, are usable. A small EPROM stores bad-spot information, just as the boot record on a hard disk stores bad sectors. Data transfer is bit-serial through the spiral, but transfers in the local DRAM chips are nibble wide, converted by multiplexers in the control logic.

Other full-wafer schemes may eventually do fancier things, such as whole systems on a wafer, but this one is here and now. (Evaluation units of the Wafer Stack, including SCSI interface, $11,680 for 40 Mbytes, $28,760 for 160 Mbytes.)

Anamartic Inc., San Jose, CA
John Scandalios 408-973-8008

EEM FILE 3130 ▶ **CIRCLE 307**