May 2002 - Volume 1, Issue 4

Memory solutions defined.

Market Analysis and Trends in the Semiconductor Memory Industry

Denali Memory Report

The Denali Memory Report is produced and distributed by Denali Software, Inc. This monthly publication provides trends, analysis, and news in the semiconductor memory industry. The Denali Memory Report is one component of Denali's Memory Research Services (DRS), and is designed to provide practical and unbiased analysis of the memory market including:

- Industry overviews
- Vendor profiles
- Alliances and second sources
- Pricing outlooks
- Supply assurance issues
- Technology roadmaps
- Custom analysis

The Denali Research Service is intended to be synergistic with other products from Denali which include high-speed memory controller IP, and verification IP for memories and memory systems. For more information about these and other Denali products please contact info@denali.com or visit www.denali.com. For information about the Denali Memory Report or content herein, please contact us at: research@denali.com.

In This Issue:

MEMORY INDUSTRY UPDATE

FEATURE ARTICLES

INTERVIEW

Denali Software, Inc.
info@denali.com · www.denali.com · www.eMemory.com

Analyst: Lane Mason · Managing Editor: Jonah McLeod · Production Designer: Alissa Wyffels
Publisher: Kevin Silver

Denali Memory Report

MEMORY INDUSTRY UPDATE

Memory Industry Outlook:

Overall, the memory industry continues to deliver mixed signals about the nature of the recovery. DRAM price increases in the first quarter accounted for the total semiconductor industry's growth in the 1Q02 period. By contrast, the revenue run rate for SRAMs and flash continued to shrink slowly compared to 4Q01.

Without demand from still-weak end markets, there can be no end to excess capacity and no permanent price increase can be sustained. See Table 1 on page 3 for the Q102 financial results of a selected group of memory vendors.

As 1Q financial reports came in, many system vendors issued cautionary tales in the form of earning warnings or outright disappointments. Only perennial optimist, Cisco's John Chambers, saw any reason for 2Q cheer.

Growing louder is the chorus of analysts and network communications companies painting an ever-darker outlook further into the future for the networking business. Virtually no one today still expects the networking sector to contribute to any rebound in 2002.

While an emerging consensus agrees the market has turned the corner and will rise steadily from this point forward, the evidence from the demand side is hardly compelling. To support this assertion consider that strategic end markets have yet to get growth traction:

1. Networking facing tough outlook
2. Cell phones sales flat
3. PCs sales lackluster
4. Auto and consumer—better and best, respectively.
5. IT spending remains weak

Perhaps, a "new consumer" will emerge to save an industry that is offering dozens of electronics-rich products and forging new ground in a market that is hard to quantify in size and long-term outlook.

DRAMs:

DRAM prices have stopped rising. Indeed, there remains an abundance of lower-priced DIMMs and loose modules throughout the industry, nominally priced at about $3.00-$3.25 per 128M equivalent. By contrast, RDRAM and RIMMs are available on the spot market, but at a far greater premium than they had been on the contract market only recently. Earlier this year, the spot market premium had narrowed to under 10-15 percent for RDRAM over SDR and DDR. Today, that premium is closer to 100 percent again.

Furthermore, in their Dec '01-Feb '02 quarter Micron took capacity off line for upgrading and experienced a 30 percent quarter-to-quarter decline in bit production rate. This compares with an average growth of over 15 percent in bit production rate quarter to quarter for the past 12 quarters.

Despite the transition from 128MB to 256MB as a standard PC system configuration, which drove considerable volumes, no one in the mainstream DRAM user community is showing strong business.

The Micron-Hynix acquisition remains on-again, off again. Right now it appears off. However, the Korean government remains interested in seeing this deal, or a similar one, concluded. Hynix cannot survive without a major cash infusion. Hynix labor, Hynix creditors and the Hynix board apparently cannot appreciate this simple fact. Without billions of dollars in the next two years, Hynix is dead. Today's recovery of 35 cents on a dollar for stakeholders is better than tomorrow's price of 5 cents on a dollar.

Table 1. Selected Memory Makers' 1Q02 Financials

Vendor	Quarter End	Sales $M	Sequential Change	Y/Y Chg	Profits $M	Prior Qtr Profits $M	Same Qtr Prior Yr Profits $M	Business Model	Memory Products
Vanguard	31-Mar	57	42%	-47%	-12	-68	-50	IDM, foundry	DRAM
Cypress	31-Mar	193	1%	-26%	-40	-38	11	IDM	SRAM
IDT	31-Mar	87	8%	-25%	-20	-0.2	-28	IDM	SRAM
ISSI	31-Mar	16.6	10%	-68%	-10.6	-9.9	4.5	IDM	SRAM, DRAM
SST	31-Mar	75	6%	-13%	1.6	-9.1	5.4	FSM	Flash
Atmel	31-Mar	276	-3%	-48%	-44	-32	56	IDM	NVM
Etron	31-Mar	NA	NA	NA	NA	NA	NA	FSM	SRAM, DRAM
Nanya	31-Mar	225	113%	152%	34.9	-115	-50.8	IDM	DRAM
ProMOS	31-Mar	90	81%	-28%	14.6	-113	5.5	IDM, Foundry	DRAM
Mosel-Vitelic	31-Mar	93	75%	13%	49.3%	-134	-64	IDM	DRAM, SRAM
Powerchip	31-Mar	90.6	71%	-2%	1.6	-115	5.8	Foundry	DRAM
Winbond	31-Mar	194	17%	-5%	4.3	-119	-12.3	IDM, Foundry	DRAM, SRAM
Micron	28-Feb	646	52%	-21%	-30	-266	-313	IDM	DRAM, SRAM, Flash
Hynix	31-Mar	660	34%	-62%	27	-1140	-469	IDM	DRAM, SRAM
Macronix	31-Mar	82.2	-34%	-59%	-8.0	-70	-58.6	IDM	NVM
Samsung*	31-Mar	2315	47%	20%	771	-166	805	IDM	DRAM, SRAM, Flash
Rambus	31-Mar	24	-5%	-25%	6.7	6.2	7.8	IP	DRAM IP
Sandisk	31-Mar	87	4%	-2%	-17	27	-16	FSM	Flash
Mosys	31-Mar	6.4	-5%	40%	2.8	2.8	0.7	IP, FSM	SRAM, eDRAM IP

Note: *Samsung Revenues are Semiconductor Device Revenues only; Profits are operating profits for semiconductor devices only.

DRAM Pricing Outlook

One can see in our forecast (Table 2 on page 4) that we believe the most recent peak in DRAM pricing has passed. We expect to see a steady but real decline in average prices for the next few quarters. Vendors wanting to improve gross margins will have to reduce cost and ramp higher volumes of dimensionally smaller chips. They cannot rely on the rising prices of the past four months for any gross margin improvement.

SRAMs:

Given that about 40 percent of SRAM bits go into cell phones, and that network systems drive most of the design virtuosity in QDR, four-ports, FIFOs, SigmaRAMs, and HS DDR SRAMs, SRAM growth is far more dependent on communications than DRAMs. How much SRAM players' health depends on the well being of communications system suppliers is illustrated in the 1Q02 results of Cypress, IDT, ISSI, and other SRAM companies contained in Table 1.

Table 2. DRAM Pricing Outlook, 2001-02 Dollars per unit, large volumes

		3Q01	4Q01	1Q02	2Q02	3Q02	4Q02
128M	PC133	1.25	1.65	2.90	3.00	3.10	3.10
	DDR266	1.40	2.00	3.70	3.30	3.25	3.20
	DDR333	NA	NA	4.00	4.00	3.90	3.60
	DDR400	NA	NA	4.25	4.50	4.35	4.00
	RDRAM	5.50	5.00	4.75	4.75	4.75	4.75
256M	PC133	3.25	4.50	7.00	7.25	6.90	6.70
	DDR266	4.00	5.25	8.00	7.50	7.00	6.70
	DDR333	NA	NA	NA	10.00	9.00	7.50
	DDR400	NA	NA	NA	12.00	10.00	8.50
	RDRAM	12.00	11.50	10.50	9.50	9.00	9.00
512M	SDR	NA	NA	75.00	50.00	37.00	30.00
	DDR	NA	NA	75.00	50.00	37.00	30.00
	RDRAM	NA	NA	NA	NA	75.00	45.00

The recently issued WSTS data illustrates the SRAM market's degree of distress. WSTS reported 1Q02 SRAM revenues down 10 percent sequentially to $532M, a slowing rate of decline in revenue from the previous quarter. However, considering the 72 percent fall from its quarterly peak of $1,922M in 4Q00, any SRAM vendors with only a 25 percent revenue reduction from early last year should count themselves successful, indeed.

Flash:

In the flash world, while mid-density 16M-64M parts enjoy price stability, higher density parts entering the market are facing a declining price. With the words "flash price stability" appearing more frequently in the press, one might wonder why. While most flash vendors increased capacity considerably in 2001, a modest 17 percent uptick in bit growth left the industry with more capacity than it could use. In addition, flash is now commanding leading edge 0.18um and some 0.13um process capacity, that it never had access to before. Compounding the excess capacity problem is the industry-wide transition underway to multilevel cells—more bits in the same silicon area—and growth in NAND flash, with its inherent lower price per bit. Thus, supply in 2002 should exceed 2001 by a wide margin.

Against this backdrop is the moderating demand from the diverse flash applications of the PC and sluggish sales of cell phones, despite the increase in flash bits per phone. About the only good news is the growing digital photography and MP3 markets, both hungry flash consumers, but lacking huge volumes. Thus, are the price stability and forecasts of further tightening later in the year real or merely transient and precursors of lurking price reductions across the board later in 2002?

New Product Introduction:

Two new Network RAMs were described at Denali's MemCon in Boston, 30 April. One was from Hynix, another from Elpida. Both used the DDR2 technology roadmaps, with on-die termination, reduced voltage (down to 1.8V), and OCD (off-chip driver calibration) to support improved signal integrity as clock rates increased.

Denali Memory Report

FEATURE ARTICLE

A Break from Bashing Rambus

It has been a tough road for Rambus as they cast themselves as a technology developer and licenser. One reason is that no one liked the IP-only business model. Rambus arose in the early 1990s amid a flurry of illegitimate and scurrilous patent offensives. Jerome Lemulson used legal but devious means to collect hundreds of millions of dollars from everyone for his image recognition patents. And Gilbert Hyatt—simultaneously defying both logic and memory—claimed to have invented the microprocessor, not Federico Faggin, Ted Hoff and Stan Mazor of Intel's 4004 fame or Gary Boone of TI, or any of the recognzed contributors to its development.

Thus was the IP climate poisoned, especially among the veterans in Silicon Valley. Non-US companies, who already had to pay the DRAM royalties demanded by TI, Micron, IBM and STMicro, were more compliant. They had joined the party later and already bore tremendous IP royalty loads payable to primarily US companies.

However, Rambus persisted and gained a key design win in Silicon Graphics' Indigo workstation. SGI's MIPS division also had a stake in the chipset for the Rambus memory, which helped. Rambus then acquired the Nintendo N64 design win with Concurrent Rambus, the predecessor to today's Direct RDRAM. Rambus was also designed into graphics boards of which Cirrus Logic shipped thousands with Concurrent Rambus RDRAMs in the mid-90s.

But the real coup was Intel's endorsement of Rambus technology in September 1996. It was Intel's remedy for bandwidth-throttled MPUs and the arterioscherotic JEDEC standards process, which yielded memory system performance improvements in inches when systems cried out for miles. Forced to buy licenses to continue to make DRAMs for PCs, memory vendors were put off. They were also forced to invest heavily in new capital equipment, especially testers to test RDRAM chips. The coup's success rested on the ability of Rambus and Intel to put in place the whole infrastructure for high-volume chip-socket-RIMM-motherboard system qualification and production by launch date, January 2000.

The failed coup is history, but Rambus is not dead. There is merit to its technology and its business model, despite being soiled by lawsuits, allegations and court rulings, and general rancorous feelings throughout the industry. The pendulum has swung too many times for Rambus to be written off prematurely. Even today, the industry's high-performance, high bandwidth problem has still not yet been demonstrably solved by any other technology.

So while it is not our place to be an apologist for any misdeed Rambus may have done, we would like to take a few lines to comment on important positive contributions Rambus made to the industry over the past decade, which the industry would do well to heed.

Rambus IP Business Model

Rambus came forth with a pure IP business model, based on technology patents two Stanford professors developed. Though unorthodox, Rambus really did not attract much attention and venom until the company was six years old.

Licensing fees and royalties for technologies and patent have been a part of the DRAM business since about 1987, when TI garnered a total of nearly $190M in fees from all DRAM makers combined. This amount rose to about $700M in the DRAM market peak of 1995, a year in which the total DRAM market was about $42B.

Denali Memory Report

In addition to TI, STMicro became a major player in DRAM IP licensing after acquiring Mostek's DRAM and SRAM patents in 1987. That year alone, STMicro earned $450M in revenues from licensing. STMicro also acquired Inmos in 1987, but ended up sharing the patent portfolio with major Inmos shareholder Thorn EMI (the record label). STMicro acquired the DRAM patents and Thorn EMI the SRAM patents of then-defunct Inmos.

Before Micron bought TI's DRAM business in 1998, each year about 10 percent of their revenues went to a variety of patent holders (including TI) for product and process rights. The sum roughly equaled Micron's R&D expenses then. The TI purchase gave Micron royalty-free access to TI patents for ten years. In the early 1990s, TI licensed, successively, Samsung, Toshiba and Hynix for ten-year contracts estimated to be far in excess of $1B in fees from each. TI continues to reap substantial revenue from its vast patent portfolio. Throughout the 1990s Korean and Taiwanese DRAM vendors paid out about 10 percent of their DRAM revenues in royalties, about twice the rate of the Japanese DRAM-making contingent.

Today, IBM gets about $1.7B in technology licensing fees and royalties, and about half of their world-leading 4000 patents granted yearly are in the microelectronics area.

Big money flows to those who have more IP, from those who have less. Later entrants to the market pay those who got in earlier.

Thus, the IP business model was established when Rambus emerged, though they were the first to adopt a pure IP model in the semiconductor space. Also, by being a pure IP house, they could not be induced to cross license for others' IP, since they had no product that needed other company's technology. (However, if Rambus' licensees had patents essential to implementing the Rambus solution, they had to cross-license the patents to Rambus so it could deliver an unconstrained, total RDRAM solution.)

Total System Performance solution

The semiconductor industry and its end markets, especially the PC market, suffer from excessive dis-integration in key areas. The PC has layer after layer of separate industries—chip, chipset, package, motherboard, and system—whose lengthy validation process necessarily delayed use of today's technology in today's products. For example, memory vendors qualified DDR SDRAMs two years before a PC with a 133MHz DDR DRAM memory chip in main memory shipped. Meanwhile, graphics cards today routinely use 250-325MHz clocks and adopted DDR in late '99.

The slowness in bringing DDR to market up through the labyrinthine PC and computer industry structure (excluding graphics cards), illustrates the bottleneck to adoption of the most advanced DRAM technologies. Stable and well-known silicon roadmaps, stretching far into the distance, preclude most technology differentiation being quickly adopted in PCs, other than faster MPUs. This institutional inertia constrains innovation spilling from DRAM R&D budgets from being proliferated widely except in isolated cases, and after long delays.

Rambus knew they would have to overhaul this technology delivery infrastructure as well as the associated technologies. They developed or contracted to develop, memory chip, RIMM module, socket, clocks and timers, SAW, reference designs manuals, and a licensing strategy to glue it all together. Without a major backer like Intel, such an overhaul would have never gotten off the ground. Only system houses with broad technology portfolios or market clout, such as IBM mainframe computers in the '60s and '70s or today's giant consumer electronics companies, have been able to construct such a bumper-to-bumper

Rambus Financials **FYE 9/30**

	Mo.	Contract Revenues $M	Royalties $M	Tot Rev $M	R&D $M	MSG&A $M	Pretax $M	PAT $M	Shares M
FY 1997		**20.2**	**5.8**	**26.0**				**2**	**86.0**
1Q	12								
2Q	3								
3Q	6								
4Q	9								
FY 1998		**28.7**	**9.1**	**37.8**				**6.8**	**97.5**
1Q	12								
2Q	3								
3Q	6								
4Q	9								
FY 1999		**35.4**	**8.0**	**43.4**				**8.7**	**100.2**
1Q	12								
2Q	3								
3Q	6								
4Q	9	10.6	1.7	12.3	1.4	3.7	4.2	2.65	
FY 2000		**39.7**	**32.6**	**72.3**				**21.55**	**107.8**
1Q	12	9.3	2.6	11.9	2.1	3.4	3.9	2.60	
2Q	3	12.1	3.5	15.7	2.7	4.4	6.1	3.97	
3Q	6	11.1	6.6	17.8	2.6	6.2	7.3	4.76	
4Q	9	7.0	19.9	26.9	3.3	6.5	15.7	10.22	
FY 2001		**21.8**	**95.4**	**117.2**				**31.84**	**106.0**
1Q	12	7.9	26.8	34.7	3.3	5.0	22.0	13.17	
2Q	3	7.6	23.7	31.3	4.5	5.3	13.6	8.18	
3Q	6	3.5	19.8	23.3	4.6	5.1	4.4	3.86	
4Q	9	2.8	25.1	27.9	4.8	5.9	10.5	6.63	
FY2002									**104.0**
1Q	12	3.1	21.8	24.9	5.1	5.5	9.5	6.2	
2Q	3	1.7	21.8	23.5	5.2	6.1	10.4	6.7	

solution. Rambus with Intel's considerable assistance almost succeeded with the RDRAM rollout.

The industry would do well to realize that lacking an optimized system solution has hamstrung the PC's evolutionary development curve. Unfortunately, the networking world appears to be traveling the same path as the PC, divorcing elemental: (1) Silicon technologies, (2) Leading edge packaging, (3) System power management, and (4) Optimized system performance methodologies, from total system requirements. (There will be more on this in subsequent *DMR*s.)

Incrementalist Institutions not up to the task

Manufacturers in the semiconductor space are not satisfied with the pace and direction of the JEDEC standards committee process. Some see its open IP model conflicting with others' IP-based business model. Others refuse to disclose their own advances to gain JEDEC's stamp of approval. QDR, SigmaRAM, Rambus, the Intel-led ADT, Infineon's RLDRAM and Fujitsu's FCRAM all occurred outside JEDEC.

In the early days Intel, TI, Mostek, and many others who pioneered the memory marketplace in ROMs, EPROMs, SRAMs, and DRAMs, were willing to let the market choose the winner. After all, users want supply assurance and competitive prices—both assumed to follow from multiple sourcing, which non-JEDEC consortia fulfilled—more than they want the JEDEC stamp of approval. Rambus' unique signaling technology development took place outside of JEDEC for just these reasons.

Rambus Technology Development Expenses

From early 1985 through 1987, the DRAM market was ravaged by its periodic pricing suicide, dumping complaints, fab write-downs, and shake-ups among major

DRAM players. During this time Intel and Mostek completely left the market, along with minor DRAM players like National Semiconductor, Fairchild, AMD, Inmos, and Standard Telecom and Cable.

TI launched its IP offensive in 1987 after a change of heart on licensing and patent royalties. TI was not receiving sufficient return from DRAM component sales to justify its large development expenses sustained over years, said TI VP, Kevin McGarrity. "We either have to make money in DRAMs or, alternatively, in licensing the DRAM technology itself, in order to get a fair return for our shareholders," he said.

What contribution has Rambus made to the state-of-the art in DRAM technology and are they entitled to a fair return? Here we enter the sticky area currently under review by several courts of what technology is owned by Rambus. To date, Rambus has spent over $60M on R&D, and undertaken contract development programs costing over $125M. Both tracks pursued a narrow range of high-performance signaling technologies and supporting enablement capabilities. Rambus licensees and importantly, Intel themselves, have spent hundreds of millions of dollars more. Whatever the courts finally rule, Rambus' patent portfolio will be seen to contain major, valuable elements not found anywhere else.

For the record, Rambus' RDRAMs garnered about 11 percent of the DRAM market in 2001, largely in the Pentium 4 and Playstation 2 systems. They announced that more than 300 products contain RDRAMs. Sales are expected to drop off markedly with Intel's chipset support transitioning to DDR PC systems. Even the Intel workstation 850 RDRAM chipset, which met with considerable success, will be replaced with a DDR version on the next spin. Finally, Rambus' legal woes make it difficult for potential users to see through the poisoned environment surrounding the company to the value in the Rambus technology, which remains compelling (see www.rdram.com). Whether Rambus can improve its image to achieve a competitive comeback into the industry remains to be seen. It will not be easy.

Fortunately for Rambus, its technical critics have nothing to offer that matches Rambus' performance in many ways. The DDR266 2.4GB/s bandwidth is just now supported in the PC. Meanwhile Rambus has been able to turn their core clocks to 600MHz and add two more channels to put a 4.8GB/s data point on its roadmap today.

While the final word is yet to be written, Rambus' IP model is viable, its contribution to the technical state of the art is noteworthy, and its legal entitlement to some reward for its efforts and expenses is reasonable. Its "total system mentality," if adopted more widely, would greatly benefit the industry.

Outlook Bright for Embedded DRAM! Honest!

Rarely has so much money been spent on a technical marvel, so enticing at first glance, but so unprofitable in the long haul. Embedded DRAMs first arose to consolidate several chips into one. Applications requiring only a few megabits of DRAM and system companies afraid trailing DRAM technology would vanish or low-density DRAM chips would cost too much, also drove eDRAM development. "System-on-a-chip" (SOC) had yet to be invented as a concept or a design goal, though higher transistor budgets were demanding more and more memory on a logic chip. Examples include MPUs with L2 caches, and complex ASICs containing hundreds of small memory arrays.

The limiting factors to eDRAM's widespread adoption have evolved slowly since its inception, but remain dominated by the following issues. First, the volatile and unpredictable nature of the competing standalone DRAM market can threaten eDRAM solutions with a cheaper, more flexible two-chip solution. Second, to incorporate eDRAM required extra processing and mask steps not found in the standard logic flow. Finally, though far denser, DRAM is more difficult than SRAM to successfully incorporate into a mostly-logic chip.

Early eDRAM Applications

The integrated frame buffers (IFB) for laptops was the first application that found an eDRAM solution attractive. NeoMagic built a profitable, though short-lived business, embedding a few MB of DRAM with its Graphics Processor (GPU) logic. Eventually, the eDRAM solution was out-gunned by several developments.

1. It was expensive to upgrade, and evolve the GPU, without tearing up the whole chip, when frame buffers demanded more memory and higher performance.
2. Standalone DRAM prices plummeted and competitors began building 4MB frame buffers, with higher performance and lower cost than the NeoMagic IFB with only 1.2MB of eDRAM.
3. To add impetus to the demise of eDRAM in IFBs, Multichip Packages fell in cost further, making the separate-chip solutions more flexible and more cost competitive. Also contributing was a growing known-good die (KGD) business, aided by changes in DRAM industry manufacturing logistics to facilitate a KGD market.

As a result, the ranks of committed eDRAM vendors thinned when development budgets were pared back. The result was that eDRAM technical progress was gradual and limited. Toshiba, Infineon and Mitsubishi all had strong eDRAM activities and had done a good deal of business in the past. All three have since scaled backed its eDRAM programs. NEC was an early eDRAM player, with design wins in printers, monitors, hard disk drive controllers, but has now pulled back as well. LSI Logic's joint technology development venture with Micron a few years ago yielded nothing to the market. Samsung's eDRAM efforts are now defunct, too. Finally, eDRAM appears wholly outside the interest of Taiwan's DRAM makers (though not foundry TSMC), perhaps because most of its technology roadmaps are licensed from outside.

Renaissance for eDRAM

The eDRAM market outlook dimmed steadily from about 1995 until just recently. However, several developments within the past few months offer promising signs that the technology may be experiencing a resurgence. One sign appeared at the TSMC Technology Symposium in April, where TSMC rolled out the MoSys 1T SRAM cell (a DRAM solution) and the Atmos planar DRAM embedded cell. Another sign is the growing success of the MoSys 1T SRAM. Finally, at the Denali MemCon event in Boston April 30, 2002, IBM presented the status of its continuing eDRAM development along with the roadmap for its eDRAM technology.

To implement an eDRAM within an industry standard logic process, MoSys relied on creativity in its cell circuit with hidden refresh. With a 24MB DRAM array in the Nintendo GameCube, MoSys has set a record for both a high volume design, and an extremely large array. Once a Multibank DRAM innovator, then an SRAM supplier, MoSys has now become a memory IP provider. In the most recent quarter, they were generating more than $2.8M per quarter in net profit, and recently announced five design wins. (See interview with MoSys VP Mark-Eric Jones p. 12.)

For the 0.13um generation, the MoSys cell size is 1.1 um^2, one-sixth the size of a 6T SRAM cell. The company expects to ship parts with clock speeds up to 350MHz this year. An extended version of the MoSys high-density memory array will add 1-2 extra masks. (In constrast to the MoSys, IBM's eDRAM cell size is 0.31 um^2 in 0.13um design rules.) The basic Atmos technology will add no additional masks, though products incorporating the technology have yet to ship.

For all designs fabricated on the TSMC process to date, 54 percent are from networking applications, 27 percent consumer, and 19 percent all others. IBM has a similar profile of designs going through its fab. Like TSMC, IBM has developed separate processing flows for these major application areas, providing capability for low power, high density and high performance.

IBM's eDRAM activity

Art Kilmer gave a solid overview of IBM's eDRAM program at the recent Denali MemCon event in Boston, and offered several salient points about the eDRAM market.

While there is no 'killer application', there are many reasons that individual design problems prefer eDRAM to eSRAM or standalone DRAM. Performance, power reduction, array size, and cost are all common issues and IBM has tailored technical solutions to address each of these issues. In addition, eDRAM technology progress has been incremental and steady including key developments in the use of trench cell (discussed more below) and adaptation of eDRAM to IBM's high-performance logic process. Today, IBM offers eDRAM both as a foundry option and as part of its ASIC cell library. IBM foundry is in production with eDRAM using a 0.18um process and has begun accepting designs for the 0.13um process. The foundry is also developing a eDRAM spec for the 0.10um process. In addition, IBM's ASIC cell library business has more than a dozen eDRAM designs in production.

IBM's logic-optimized eDRAM cell is far higher performance than its standalone DRAM cell, but about 40 percent larger than a commodity DRAM cell using the same design rules. The eDRAM macro has a standard 256-bit bus and read-write performance that matches all but the fastest SRAMs: a tRC of 12ns for the 0.13um process and targeting a tRC of 5ns for the 0.10um process. This is truly SRAM speed. For the 0.10um process, peak bandwidth is 16GB/s, about twice that of the 0.13um process.

A chip fabricated in the 0.10 um process, with half the die area devoted to memory, can contain more than 500 Mb of eDRAM. Importantly, SRAM soft-error rates (SERs) are increasing with the transition to finer design rules. With its 100-fold to 1000-fold better SER immunity than eSRAM, eDRAM is an even more preferred embedded memory alternative.

Since IBM is a trench-cell devotee, they still require an additional four masks in addition to the logic process. The added masks require about a 20 percent cost premium. It's large, but far less than a few years ago, and now acceptable for a broad range of performance-sensitive applications that could benefit from having a large amount of RAM on the ASIC/SOC.

Despite IBM's withdrawal from the Standalone DRAM business, DRAMs based on the IBM trench cell construction are still being produced at Nanya, ProMOS, Mosel-Vitelic, Winbond, and Infineon. Trench-cell DRAMs had been produced on certain Toshiba fabs for a while as well. The trench construction has advantages for

Denali Memory Report

eDRAM. The smooth topographic of the cell puts the capacitor into the substrate along with the control transistor, which is sideways in the same trench. These subterranean features ease metal wiring over the top. The trench cell also has lower leakage currents, thus leading to an improved refresh regimen.

Future Bright for eDRAM

It appears eDRAM is finally poised for a much larger role in future SOC and ASIC implementations for a number of reasons;

1. Transistor budgets are becoming larger with each generation of technology.
2. Logic and system performance is throttled by inadequate amounts of memory embedded in MPUs/SOCs.
3. The International Technology Roadmap for Semiconductors shows SOCs dominated by memory by 2006 and accounting for 90 percent of an SOC by 2015.
3. Designers are finding scaled SRAM solutions woefully inadequate, both technically and economically.
4. Concurrently, embedded DRAMs are steadily making progress in faster read and write cycle times, reduced process complexity, reduced power consumption, and lower soft-error rates than scaled SRAMs.

However, compelling technology always comes with caveats, so too for eDRAM. First, different eDRAM process flows are not alike, though there are similarities. Second, multichip package solutions are increasingly cost competitive against eDRAM. Finally, most designs are hardly static enough to ensure long running eDRAM solutions. Nevertheless, some combination of design problem and technical solution that eDRAM can serve seems to have achieved critical mass. We expect to see more eDRAM applications and profitability for the memory class, in the next few years.

Denali MemCon Boston Unqualified Success

Denali hosted a MemCon Boston ‘02, focused totally on semiconductor memories, in Boston on April 30th. Ten memory companies presented their new technologies, including Denali plus a presentation from JEDEC Chairperson Desi Rhoden. The conference was a great success, more than 180 ASIC designers and others interested in the outlook for memory technology attended. The first-of-its-kind show was a great opportunity for designers to learn of new technologies from vendors and for vendors to describe their technical and market requirements to designers.

Denali plans to host Denali MemCon Japan in Tokyo and Denali MemCon Silicon Valley in October. The Tokyo conference will focus on PSRAMs, Flash, packaging, and consumer applications of memories. The Silicon Valley Conference will feature discussions of Network DRAMs, QDR and Sigma RAMs, and Flash for communications.

Visit www.eMemory.com for more information about upcoming events.

Errata:

In the last table of the April issue of Denali Memory Report (Vol 1, Issue 3), we understated Cypress presence in both the SRAM market and the Other Memory market. Cypress revenues for those product classes for 2001 should have been listed as $356M for SRAMs, and $103M for the Other Memory market. We apologize for the error.

Denali Memory Report

INTERVIEW

Manufacturability and High-Yield Key to Embedded 1T-SRAM® Success

Mark-Eric Jones describes how the MoSys embedded 1T-SRAM succeeds where other embedded memory technology has failed.

As Vice President and General Manager - Intellectual Property, Mark-Eric Jones is an unapologetic cheerleader for the MoSys embedded 1T-SRAM technology. He believes MoSys will be to embedded memory what ARM Ltd. of Cambridge, UK is to embedded processors in system-on-chip designs. The company's financial performance—highly profitable and growing—and continuing customer design wins suggest that the MoSys technology is living up to its marketing. The ***Denali Memory Report*** caught up with Mark-Eric right after the company's last quarterly conference calls at the end of April. The company had just completed launching five new design wins and strategic alliances. Here is what he had to say.

***Denali Memory Report*: Can you provide our readers some background on MoSys?**

Mark-Eric Jones: Our CEO Dr. Fu-Chieh Hsu and CTO Dr. Wing-Yu Leung, who had both been involved in advanced memory design for some time, founded Monolithic Systems Technology Inc., MoSys, 10 years ago. They believed that for very deep submicron designs, a new memory architecture was required to achieve the best results for large memory arrays.

The first product based around that architecture, was called MultiBank® DRAM (MDRAM). It hit the market at the end of 1995 in a fast DRAM, which even by today's standards is still very high performance. The unique MoSys multi-banking architecture divides the memory into a large number of very small banks without a significant silicon area penalty. The single-bit-cell architecture achieved an order of magnitude higher performance than a conventional DRAM configured as a pseudo static RAM.

DMR: How did the MDRAM become the 1T-SRAM?

MEJ: To make the access and the interface of the memory look like an SRAM, we developed the MCache, which didn't make the device completely SRAM-compatible. It required the designer using the device to make a slight change on the interface. It's a

bit like some of the pseudo SRAMs that various companies have announced recently.

However, the company realized that forcing designers to adopt a unique design was not an effective strategy. We had to make the interface and the characteristics absolutely industry standard. That's when we came out with the first true SRAM-compatible devices around the 1997-1998 time frame.

In 1998, the company also realized that the standalone SRAM market was not particularly attractive since it was hardly growing and at one point even declining.

Denali Memory Report

By contrast, two forces were driving embedded SRAM expansion. One, the proportion of a SOC occupied by memory was increasing, and, secondly, the SOC market itself was escalating rapidly. In 1998-1999 the company began licensing its technology for use in SOCs to play in this much larger market.

DMR: How does the standalone SRAM business fit into the company's overall strategy?

MEJ: We believe that there is more to memory than raw performance. The trick is to make it manufacturable while keeping the yield high—what I would call the quality aspects. Embedded memory today tends to have far more performance than the accompanying logic really requires. Certainly, that's the case with 1T-SRAM.

Our standalone business is very important to ensuring the quality aspects of our embedded memory business. Because we have shipped 40 million standalone chips, we understand the manufacturing issues, the test issues, the yield issues, not just from a theoretical point of view, but also from a real practical point of view.

And we will continue to stay in the standalone business as a development platform for new memory technologies. For example we recently announced a new 1T-SRAM R™ (reliability) technology.

DMR: What applications use the MoSys standalone SRAMs?

MEJ: Where these high-density SRAM-compatible parts end up are in designs at high-end communications companies. For example if you look at our published quarterly earnings results, you'll see that some of the top customers of our standalone business are companies like Cisco.

Our standalone SRAMs serve at the high end of the memory market, not at the commodity end. For example, last year, we announced what was, at the time, the largest absolutely SRAM-compatible 36Mb device. And we were shipping last year as well. We announced our product before the 72Mb pseudo static announcements from other companies, which were not fully compatible SRAM devices. This year, we will be coming out with an absolutely SRAM-compatible 72Mb device.

DMR: Does MoSys have plans to license its 1T-SRAM technology to a third party to produce standalone SRAMs?

MEJ: We've chosen not to do that at this time. In fact, our normal licenses specifically preclude pure memory devices. Who knows what will happen in the future? For example, we have licensed NEC to make a custom 1T-SRAM memory device used in the Nintendo GameCube. In addition, we have licensed NEC to embed the 1T-SRAM in SOC ASIC designs.

DMR: How long after you ship a next generation cutting-edge, high-performance SRAM before it is offered in embedded format?

MEJ: That depends on specific customer requirements. In other words, there isn't a fixed time period. Some customers will want to be at the bleeding edge themselves. Others will want a significant amount of production history before they're ready to license any technology.

DMR: Let's talk about your business model. How would you characterize MoSys as a company?

MEJ: The MoSys business model is closer to a silicon model in the sense that we're developing unique memory technology. We've been granted 40 U.S. patents on this technology and there are many more pending. Other embedded memory vendors provide good tools for efficiently configuring, tailoring, and delivering commodity memory designs to customers. Thus, if you think of IP memory companies

Denali Memory Report

in general as lying closer to the EDA side of the EDA-semiconductor spectrum, I would characterize MoSys as being closer to the semiconductor end.

DMR: Can you contrast your embedded memory IP business model to that of other IP companies?

MEJ: Companies such as ARM, MIPS, and Rambus have taken up to 10 years to get its IP revenue to a decent number. I haven't plotted our IP revenue against that of other public IP companies, but I think you'll find ours has grown much faster. Recently, Dataquest confirmed that we had the fastest growing IP revenue of the top 20 IP companies worldwide.

We only started our IP business in 1999 and already licensing and royalties represent the majority of our revenue. Last year our IP licensing business exceeded both the revenue and profits of our standalone products.

DMR: You claim that you are to the embedded memory market, what ARM is to the embedded processor market. Can you charge as much for your designs and your royalty?

MEJ: Actually, I would argue that memory can charge more for both. The International Technology Roadmap for Semiconductors (ITRS) shows memory already accounting for over 50 percent of a typical IC SOC, growing to 94 percent by the year 2014.

Furthermore, which is higher value to a customer: a processor that is 10 percent better than some other processor, or memory that's 10 percent better than another memory? Which will command the highest royalty-memory occupying 94 percent of the die area or a CPU occupying just 1 percent?

DMR: What is the target application for your embedded memory?

MEJ: We look at applications where memory represents a large percentage of the die.

Our initial focus is on consumer and communications. In the consumer market the volume is typically very high, therefore we can save money for our customers. In several cases, we have saved them up to 30 to 40 percent of their chip cost, compared to a 6T SRAM.

In advanced communications chips, having an optimal memory technology that provides SRAM performance, but at much higher density, much lower cost and also much better yield than is possible with 6T SRAM, is a big advantage.

DMR: In what high volume applications have you achieved design wins?

MEJ: We have multiple design wins in cell phones, digital cameras and camcorders, and DVD players. And we have a significant design win in the Nintendo GameCube. I think in the end, what's important to our business model is royalties, which already account for over 40 percent of our total revenue. And this demonstrates that our products are being used in high-volume applications.

DMR: Of the design wins in the consumer and communications space you mentioned, how much MoSys embedded memory is on a typical die??

MEJ: If you look at these projects, the amount of memory varies across these projects from 1 megabit to well over 20 megabits.

DMR: When do you expect to see the products that contain your licensed technology showing up in electronics retailers?

MEJ: I look at a slightly different measure, which is "when do we expect to see the royalties?"

Denali Memory Report

Our royalties are reported a quarter after they are paid. I expect that from signing an agreement to seeing appreciable royalties—I don't mean a first bit of pre-production, but volume product—is typically 18 months to two years.

DMR: Can you describe where you are on the TSMC embedded road map revealed the latter part of April at the 2002 TSMC Technology Symposium in San Jose, Calif.?

MEJ: We have shipped and are already receiving royalties on 1T-SRAM-M. And we have already taped out test chips for our 1T-SRAM on the TSMC 90-nanometer (nm) or 0.10-micron process technology. Furthermore, TSMC has committed that our 1T-SRAM will be available the same time that 6T SRAMs are available from the TSMC 90 nm process. Thus, there will be no lag between the availability of standalone 6T SRAMs and our embedded 1T-SRAM.

DMR: You mentioned that April was also a busy month for new announcements. Can you elaborate?

MEJ: We've had five new licensing announcements in the last full week of April: Matsushita, LG Electronics, STEPMIND, Pixim, and Philips Semiconductor. Matsushita licensed our technology for ultra-dense high performance memory for future communications IC products. The LG license is for embedded memory in its DVD IC product lines. STEPMIND plans to incorporate the 1T-SRAM into future wireless networking ICs. Pixim licensed the technology for its imaging ICs. Finally, Philips plans to use the memory in ASICs and SOCs for cellular infrastructure.

DMR: In these five design wins, what were the critical factors in terms of power, bandwidth, cost, and performance?

MEJ: The primary reason for most of these wins was cost. In some cases, power was significant, but I think you'll find that cost, manufacturability, and yield were the foremost reasons.

Data TSMC published shows the yield for 1T-SRAM is far greater than for 6T SRAM, even per square inch. And of course the yield is even higher when measured per megabit due to the much smaller area used by 1T-SRAM memory. Some of these announcements will incorporate our new Transparent Error Correction™ technology. Also, all five will be implemented on either 0.18- or 0.13-micron processes.

DMR: How well does the MoSys embedded 1T-SRAM technology scale?

MEJ: We've demonstrated the scaling on six standard logic process generations: 0.25-micron to 0.22, 0.18, 0.15, and now 0.13 micron. In the latter part of April, we extended our collaboration agreement with TSMC to include its new 90-nm process technology.

We've proven it's scalable. In fact the 1T-SRAMs advantages increase with scaling when compared with a 6T SRAM. The 6T SRAM requires a large amount of interconnect. Making smaller transistors is relatively easy, but scaling complex interconnect is very difficult. The 1T bit cell has virtually no interconnect and therefore we see the advantages of 1T over 6T growing, as we migrate to even deeper submicron processes.

Another very important scaling factor is quality. Our emphasis is more on memory quality than absolute performance since 1T-SRAM already offers higher performance than is needed by most applications. By quality we mean yield, reliability and soft error rate, the three most important measures of memory technology quality.

DMR: What about the soft error rates? Being a single transistor, the soft error

Denali Memory Report

rates should be higher, as a result of just having one element?

MEJ: Actually, that would intuitively seem right, but the data shows exactly the opposite. We have measured and published soft error rates for 1T-SRAM for multiple process geometries including 0.13-micron. The measured soft error rate is 1000 FITs (failures in time) per megabit. With 6T SRAM we've seen figures of up to a 100,000 FITs per megabit in 0.13-micron.

In early Q2, 2002 we announced our new 1T-SRAM-R technology, which has set some new quality standards for embedded memory. It has a soft error rate of just 10 FITs per megabit, four orders of magnitude better than 6T SRAM in the 0.13-micron process.

DMR: Did this improvement result from a design improvement, a process improvement, or onboard error correction?

MEJ: It's a combination of design and what we call onboard Transparent Error Correction. It involves no change to the process. Our philosophy is to make the memory fit the logic process not force a different non-optimal process onto the designer.

Unlike ECC, which typically adds 20 to 30 percent more area for extra memory bits, the 1T-SRAM-R Transparent Error Correction doesn't increase the memory area at all. It does actually add 20 percent more bit cells, but the bit cells are 20 percent smaller.

You can't do that with a 6T SRAM because the bit cell is already the smallest possible in that process geometry. You can't do it with embedded DRAM because, again, the process sets the size of the bit cell.

With 1T-SRAM, we were considerably above minimum design rules for the process. Thus, by running test chips with 30 percent smaller bit cells to check yield, we found that we could safely implement the reduced bit cell size for 1T-SRAM-R.

DMR: But, doesn't the smaller cell size itself increase soft error rate?

MEJ: Yes, reducing the bit cell 20 percent increases soft error rate slightly, by about a factor of 3, which is the reason, we do not implement it on our regular 1T-SRAM. On the other hand, because we're adding six error correction bits for every 32-bit word, the improvement in soft error rate is increased by orders of magnitude. Taking a small step back achieves a giant leap forward.

And although today we offer the 1T-SRAM-R as an option, I'm sure that within a couple of years it will be our standard offering. The customer gets quality improvements without any extra cost—the silicon is no larger and there is no additional royalty. In addition, we've eliminated the need for laser repair totally, which saves manufacturing costs.

DMR: You don't have redundancy in your embedded 1T-SRAM R arrays?

MEJ: That's correct. If you work it out, redundancy can normally only correct 2 bits per megabit of error. The correction technique in the 1T-SRAM R can theoretically correct up to 32,000 bits per megabit, if the errors are evenly spaced, one per word. Notwithstanding that in practice errors are random, transparent error correction easily corrects tens or hundreds of bits of error, making it much more powerful than redundancy.

It's a good technology to use when manufacturing on a leading edge process, which may be immature with a very high defect density.

DMR: We thank Mr. Jones for this insight into the MoSys embedded memory technology and for providing insight on future developments.