January 2005—Volume 4, Issue 1

Denali Memory Report

Market Analysis and Trends in the Semiconductor Memory Industry

The Denali Memory Report (DMR) is produced and distributed by Denali Software, Inc. The DMR is provided free of charge to all registrants on Denali's web site, **www.eMemory.com**. This web site contains all back issues of the DMR, and provides a host of various memory related resources including proceedings from Denali MemCon events and the worlds largest online database of semiconductor memory components.

In the DMR, readers will find memory market news, discussions of market trends, products and product strategies of memory vendors, alliances and industry technical consortia, and corporate financial performance. In each issue, we have included an interview with an important industry spokesperson, commenting on his company's status or corporate strategy, the outlook for the industry, or technical trends within his company's markets.

Denali Software Inc. is the world's leading provider of EDA tools and Intellectual Property (IP) solutions for chip interface design, integration and verification. Denali's Databahn™ Design IP products offer fully configurable design cores for complex interfaces such as SATA II and DDR-based memory systems. Denali's PureSpec™ Verification IP product supports all complex interfaces, including PCI Express, Advanced Switching Interconnect (ASI), USB, and SATA II. More than 400 companies worldwide use Denali's tools, technology, and services to design and verify complex chip interfaces for communication, consumer, and computer products. For more information, please visit Denali at **www.denali.com** or contact Denali directly at: (650) 461-7200, or email: **info@denali.com**.

In this Issue:

Register now for Denali's MemCon Taiwan 2005

www.MemCon.com

March 15, 2005—Ambassador Hotel, Hsin Chu, Taiwan

Denali MEMCON TAIWAN MARCH 15

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

MEMORY INDUSTRY UPDATE

Semiconductor Market Outlook

We won't dwell long on discussing the high degree of uncertainty and variations in market forecasts for 2005. Suffice it to say that the outlook is "concentrated" between "up 5 percent" and "down 5 percent", for the total chip industry in 2005 vs. 2004's ~$213B revenue. Such a forecast is surely "by gosh and by golly", as no one can really put their arms around the driving forces of demand, which is complicated by a weakening US dollar (the metric in which the chip industry measures itself), and the dark forbidding specter of price competition, especially in memories, which can conceivably undo any revenue growth in all other markets combined. Cell phone growth is said to be strong or weak; PCs are said to be just the opposite—weak or strong. Inventories, as they usually do, have hung around for longer than thought.

This is a problem common to all downturns and softening markets: no one can really count inventories, and count them accurately, at every stage of production and consumption, nor can they agree on what the proper amount is, and, understandably, since the flywheel of chip production has significant inertia, manufacturer's have a strong reluctance to slow down production and (thereby) increase unit cost. Few chip-makers cut production unless there are clear signs (1) that there is excess in the line and in the channel, and (2) that it will be there indefinitely. Such is the situation today—inventories did not come into balance nearly as fast as some early forecasts indicated they would, e.g Winbond's just-released January results noted a sequential decline in sales of about 10 percent, attributed it in part to inventory hangover from 2004.

Within the past two weeks, Toshiba, NEC and Hitachi all reduced their profit outlook for their fiscal year ending 31 March, for their entire corporation, noting the weakness in their chip sector's business.

Profit margins suffered in 4Q04 (see memory company financials in the article below), leading to adjustments in labor force, business withdrawals and a significant sobering up after a 6-8 quarters of good growth out of the trough of 2001-02. Of all companies reporting this past two or three weeks, only ARM Holdings was optimistic about its outlook. Second place in the "positive outlook race" was "flat in 1Q, hopeful for 2Q", but not really believing in the strengthening in 2Q05, since these companies could provide no reason for optimism except guessing that the inventories would be worked off by then—despite a weak order book for 1Q05 today. Foundries mostly reported acceptable and profitable results for 4Q04, but with declines sequentially and outlooks (in wafer starts, pricing and profit margins), which were decidedly lackluster out into 2005.

Remember that a steady and flat revenue run rate from 4Q04 compared to 3Q04, because of the strong ramp during 2004, means that 2005 would be up about 3.5 percent with stable sales, so any forecast for 2005 that is less than 3.5 percent implies an inflection point downward at some point in the year—if it has not already happened, since for the who industry 4Q04 came in at $55.1B, already down from 3Q04's $55.6B. Whether this turn down in 4Q04 is a brief respite or the onset of another downturn is anyone's guess. Pundits have already trimmed their 1Q05 estimates, setting expectations for another sequential quarterly decline. But for the whole of 2005, smart forecasters line both sides of the road, and the landscape changes daily.

Though "official" 2005 total semiconductor forecasts are blocked between about +/- 5 percent, compared to 2004, Denali believes the real spread is much wider, both on the upside and downside, perhaps as large as +/- 20 percent—it all depends on price stability and demand strength. Though we enter 2005 with the vectors pointed the wrong way, we've been surprised before, and there are several cards that remain to be played that have important consequences: reversal of the dollar weakness, stronger economies in Europe, more difficulty ramping 0.11µ line geometries, leaving supply growth constrained, further restraint in vendor Capex for 2005, etc.

Table 1. WSTS Final 2004 Data ($B)

	1Q04	2Q04	3Q04	4Q04	Yr 04	Yr 03	Growth 04/03
DRAM	5.5	7.0	7.1	7.3	26.9	16.7	60.9%
SRAM	0.77	0.80	0.82	0.64	3.03	2.60	16.5%
Flash	3.9	3.9	3.9	3.9	15.6	11.7	33.0%
Other	0.37	0.44	0.44	0.40	1.65	1.48	11.5%
Total	10.5	12.2	12.3	12.2	47.1	32.5	44.9%
Total IC	40.9	44.9	46.5	46.5	178.8	140.0	27.7%
Total Semiconductor	48.9	53.5	55.6	55.1	213.0	166.4	28.0%

Source: WSTS

Memory Market Outlook

On the memory side, bit demand continues to be pretty good, but there's still downward price inertia from the DRAM price declines over the past 3-4 months, and, with major ASP declines in 2H04 now the new "pricing baseline", both NAND and NOR flash will have to ship more than 70 percent more bits in 2005 to match 2004's revenue numbers—even if prices stay put, which is also unlikely.

There is something new afoot in 2005, as well. Several DRAM makers can swap DRAM capacity for NAND flash capacity (Samsung, foremost, but also Infineon, Micron, and Hynix), which can take pressure off DRAM supply growth and prices and move it to NAND flash. There is a lot of common tooling between the two products, and the larger constraint appears to be established market position (how to get rid of the NAND you make). Though Samsung has by far the most developed NAND flash business among DRAM makers (and Elpida has none at all today), clearly trading off increasingly unprofitable DRAM capacity for still-profitable NAND capacity is on every DRAM maker's mind. Micron had some interesting analysis of the NAND-DRAM capacity tradeoff in its 28 January Analyst Meeting, which offered their estimates of the relative 2Gb NAND and 256M/512M DRAM price tradeoffs that would dictate maximum profitability.

Fortunately for NAND, the market is expanding very rapidly with new, even lower price points and exploding applications. NOR may suffer a different fate, with its future tied largely to cell phones, in which the flash portion will get increasing competition from NAND. Though a large and robust market, NOR is not growing nearly as fast as the sum of all of the NAND "bulk bits for audio (MP3) and video" applications.

For sure, NOR and NAND flash have priced many companies out of the profit zone for most companies (NOR more so than NAND, where Toshiba, Renesas and Samsung still make money) over the past six months; DRAMs are teetering on the edge of more red ink for most vendors, even as they add capacity to make still more bits. Woe to all those NAND flash makers who joined the hunt in 2004—Infineon, ST Micro, Hynix, Micron—and ran right into price per MB declines on the order of 50-60 percent just since last spring.

News and New Products

Big Blue, Big Expansion

Big Blue is again the focal point of a major investment consortium in their home base in East Fishkill, New York. While they have been the "team leader" (for a price) of many partnerships, technology exchanges and joint development exercises for more

than a decade, this latest work up appears to be the broadest and most well-funded, and is helped by some funds from the State of New York, as a part of its Advanced Technology Investment Program. All of IBM's earlier partners are on board for a $2.7B multifaceted chip making and equipment technology development program in New York's Hudson River Valley: Sony, Toshiba, Infineon; AMD (IBM's entre in thwarting Intel Hegemony in microprocessors), plus Singapore's Chartered Semiconductor and Samsung. New additions from the equipment side include Applied Materials, and maybe Tokyo Electron. In addition, ASML will locate a $400M lithography development operation at University of Albany, less than 100 miles up the river, past SingSing Prison, and past the grand estates of Franklin Roosevelt and Andrew Carnegie.

It has always been hard-to-impossible to judge the outlook for the chip business, based on IBM Micro's fab investment activities. Some would dare to say they take the long view, and invest though any downturn, are not rattled by short-term perturbences in the market, keep their gaze on the fixed stars. Others feel that the decision-making process is so long and wrought with reviews, redirections and often fuzzy strategic objectives (To serve internal or external demand? To chase ASIC or foundry business? DRAM or not?), that they are usually precisely one-half cycle out of synch with the rest of the industry. So it is not really surprising to see Big Blue wade in with a major new investment announcement at just the time that others are putting all Capex through a serious review, and often putting on the breaks.

The State of New York, with some lobbying help from IBM, several years ago gave IBM generous tax holidays and grants to help IBM establish and expand it's 300mm East Fishkill fab, first using a vacant site left from its glory days in the 1980s and 1990s.

Big, Bigger, Biggest: Samsung To Invest $25B Over Six Years, But Where Will They Drive Their "Memory" Business?

Having achieved and cemented its position as the world's Number Two chipmaker in 2004, Samsung has embarked on an audacious move to grow still bigger and faster than the industry, with an "Intelesque" capital budget for the remainder of the decade of more than $25B through 2010. A look at their investment history over the past 2-3 years shows that this is not to be sniffed at; a look at their profitability and operating margins shows that they have the cash generating capabilities to make it happen. Being more than 92 percent commodity memory and its close variants (server DIMMs, GDRAMs, RDRAMs, UHS SRAMs, etc.), one would think, though, that they have to be near to topping out in the memory market—more market share brings unacceptable price risk, as there are not enough memory niches today to protect their mainstream revenue exposure. So they can build more fabs with better processes and higher yields and make more memory, or they can branch out into other silicon directions.

But, significant uncertainties surround their product (and corporate) roadmap, and strategy, for the next decade. It is not for no good reason that new Asian chipmakers—Japan, Korea and Taiwan (and China's foundry, SMIC)—entered the memory market first: One large fab making only one product, provides lots of learning but requires virtually no marketing, no product definition, none of the custom logic complexities in verification and validation, and little selling expense. Samsung has done it better than anyone before, and profitably weathered what has been the toughest five-year period in DRAM and memories ever. From humble beginnings in DRAMs two decades ago, they now own about 30-35 percent of each of SRAM and DRAM, 50 percent of the graphics DRAM business, and 60 percent of NAND flash (and 23 percent of all flash), so now what is next?

One can see how begrudgingly non-memory markets give up a $1B in revenue stream, with PLD makers Altera and Xilinx taking more than 15 years to grow to $1B/year; analog makers Linear and Maxim not there yet, though they were launched in about the same 1983 timeframe. LSI Logic, Cypress and IDT in and out of the $1B club. Will we see a TI like focus on DSPs (or some other product) to lock into a fast growth, high volume potential market (and which one)? Will they buy someone or "grow their own"? Or will they buy and dismember someone, like Oracle and PeopleSoft, merely to replace their acquisition partner's customer base and product with their own?

What Samsung does next is one of the most interesting questions in the industry today. Can they go the next level—be a $20B company in the next 2-3 years (at $16.3B in 2004, they're about half Intel's size in terms of total semiconductor revenues). Though Samsung also seems to be managed quite profit-consciously, while growing about as fast as anyone in any product line, if sheer revenue mass is the target, then Samsung has come a long way. However, at some point, the risk of being too large a play in memories, with their volatility, and inherent ability to be destabilized in price by a small player intent on foul play. Samsung, to continue growing as they have in the past, will have to broaden out. Certainly, they have a lot of core competencies that can be applied in other product areas, or in the foundry arena. But once leaving memories, the high-volume market "low hanging" fruit is about gone—unless they want to buy AMD and tackle Intel head-on in MPUs. Nothing else can fill a fab like commodity NAND flash or DRAMs, which often fill the same fab with little change in tooling.

Errata on JLC Associates Web Link

Although it is pretty late in the day to make an amendment, it should be noted that we listed the web site of the moderator (Jim Cantore) of the Network Systems Design Conference (NSDC), Memory Subsystems Design incorrectly. It should have been **JLCAssoc@comcast.net**, not JLCAssoc@comcast.com, as we indicated in our earlier DMR. We apologize for the error, which probably led several DMR readers down an internet dead end with no place to go.

Elpida's Interesting LP DRAM Technology: Super Self Refresh (SSR) DRAMs

Everyone is driving to get power reduction technology in their chips. Down one level, right along "consumer" as the "chip business market-driver" for the next decade, is a much-elevated focus on low power systems—emphasis on the word "systems"—to cope with more feature demand, longer battery life and more mobility. Battery technology has moved fast, for sure; system power reduction and power management has moved just as fast, or faster.

Although memories are rarely the power hog in most systems when they are operating, when they're in the standby mode all that changes. For flash, that is not a problem—it consumes no power unless reading and writing. For DRAM, though, to retain its data, it has to refresh the cells, and this consumes power. Data is lost because each DRAM cell bleeds charge at a slow rate, and over small amounts of time, it may lose it data. So, in Dynamic RAMs, periodically the cells are read, and the charge is restored into the cells. "Refreshing cells" takes energy, and the less frequently one can get by without refreshing cells, the better, and lower power the DRAM is. That is why the JEDEC conventions for LP DRAMs have two intelligent and reduced refresh frequency specifications: TCSR and PASR. Temperature Compensated Self Refresh (TCSR) takes advantage of the fact that if the temperature is lower, leakages from the DRAM cell are less, and it needs to be refreshed less often. Partial Array SR (PASR) applies the rule that only active parts of the array need to be refreshed and restored during active mode, so refreshing the remainder of the chip is neglected, reducing both active and passive power in the DRAM.

In large DRAM arrays, there are wide differences across the array in the frequency of refresh needed for each cell, since each cell bleeds charge at a different rate. These differences are due to manufacturing variations in the capacitance of each cell, which may be under 0.2µ on a side, and right at the hairy edge of today's manufacturing ability. When devising a refresh regimen, the manufacturer has to pay attention to the weakest cell in the array, and refresh all cells based on the weakest cell (though some marginal cells have already been deselected through use of redundancy).

Elpida's new Super Self Refresh technology takes advantage of this wide variation in "need to refresh cells", and throttles back the refresh rate by a large amount, about 20x less frequently. But by potentially letting some cells fail by their excessive leakage, and drop their bit, one can make the refresh regimen far less frequent—and then, as Elpida has done, correct the data in the cell later with ECC s circuitry on the chip. Clever, yes?

The ECC circuitry takes up some die area, but, along with PASR and TCSR, enables the Elpida LP DRAM to achieve a 95 percent power reduction without loss of data integrity. (Though the ECC methodology does limit the application of SSR DRAMs.)

From an historical perspective, this is not the first DRAM to put ECC on board the DRAM, though it probably is the first to do so for purposes of reducing power. In the mid-1980s, Micron Technology first offered an ECC 256K DRAM, mostly to improve yield and get into volume production more quickly. These parts were rejected by the market, especially when everyone else soon came to market with "traditional" 256M DRAMs. Micron went back and redid its part, and came out with a competitive 256M DRAM like everyone else. No one likes a supplier who wanders very far from the mainstream in commodity products, and such adventures are quickly suppressed.

A better example was IBM's 16M "Luna Classic" introduced in about 1988 at the ISSCC, which had a bank of ECC circuitry on either half of the die's array, as a means both of getting early yield up and ensuring high reliability in IBM's mainframes. There was no second source, and, since it was all IBM, it was designed into IBM's computers by fiat, but not without good reason.

This is the advantage of "total system design", which only IBM has any semblance of today, but many consumer electronics supplies find themselves in need of (e.g., cell phones). IBM did not have to rely on (sometimes too-timid) industry standards, they could draw on their vast technical resources and skill set, especially in packaging, to design and build computers which optimized system price performance. They owned every important technical link in the chain. Despite IBM's retreat from "commodity" DRAMs many years ago, there remain residuals of this capability in play in its PowerPC and PowerArchitecture effort, which are driving IBM's significant market share gains in the Unix Server world today.

Along the way, Luna Classic made it to Mars on the Pathfinder Lander, as on-board ECC made this chip exceptionally rad hard. IBM sold Luna C wafers to Loral, which built them into the lander and shot them off into space, landing on Mars a few months later.

Denali Corner

Denali Press Releases

12/14/04 Ubicom Selects Denali's Verification IP Products for Chip Development Efforts

Denali announced that Ubicom™ Inc., a leading provider of communications processor and software platforms, has selected Denali's PureSpec™ verification intellectual property (IP) products for its chip design and verification efforts. Ubicom engineers now use Denali's PureSpec verification software to model and simulate interactions between its chips and other devices in the target system. "Our customers demand high-quality, ultra

high-reliability system-on-chips that utilize leading-edge interface standards," said Jon Gibbons, Ubicom's director of engineering. "Functional verification of these systems is critical, especially for the chip to chip protocols, and that's why we chose Denali. The quality and completeness of its verification IP enables us to produce a high-quality product on time."

The full text of this press announcement is available at:

www.denali.com/news_pr20041214.html

2/9/05 Denali Chief Verification Officer to Speak at DVCon February 16 in San Jose

Denali announced that chief verification architect Sean W. Smith, will present a paper at the Design and Verification Conference and Exposition (DVCon) on Wednesday, February 16, 2005 from 10:30 a.m.- noon PST in San Jose, California. The paper, entitled *Function Verification of Design IP—Trust or Hard Work?*, will address the functional verification challenges associated with using design IP in SoC chip design efforts. Denali is also an exhibitor at the conference, and will demonstrate its entire suite of chip interface design and verification products during DVCon in Booth Number 301.

The full text of this press announcement is available at:

www.denali.com/news_pr20050209.html

Denali CEO Invited to EDA Consortium's Annual CEO Forecast and Industry Vision Panel

On February 21, Denali will announce that its President and CEO, Sanjay K. Srivastava, is an invited panel speaker for EDAC's annual CEO Forecast and Industry Vision Panel moderated by Jay Vleeschhouwer of Merrill Lynch. Topics will include trends affecting industry growth as well as forecasts for the upcoming year. Other CEO panel members include: Aart de Geus of Synopsys, Mike Fister of Cadence, and Wally Rhines of Mentor Graphics. The event will be held on Thursday, February 24, 2005 at Hewlett-Packard in Palo Alto.

For event details and registration, please visit:

www.edac.org/invites/ceo_forecast.html

Denali Webcasts

Coming Soon! "MultiMediaCard Association—Enabling the Mobile Revolution"

On February 24, Denali will host a webcast with a guest speaker from the MultiMediaCard Association. The webcast includes and introduction and market overview from Memory Market Analyst Lane Mason, and is followed by an interactive presentation by Andy Yang, MultiMediaCard Marketing Committee member and Strategic Marketing Manager at Samsung Semiconductor. The webcast begins with an overview of the MMC Association and MMC technology, then provides technical details and enablement efforts for the latest MMCA technology, including MMC*plus*™ and MMC*mobile*™.

This webcast event will be broadcasted live on February 24, 11:00 AM Pacific time. DMR readers can preview this event at:

www.denali.com/tv/SamsungMMC20050224

Available on Demand

On December 15 and 16, Denali Chief Verification Architect, Sean W. Smith, hosted a two-part webcast series on functional verification for SoC chip development. Part 1, *Verification Infrastructures for IP Based Methodologies* addressed topics such as configurable verification IP requirements, coverage techniques, test generation, compliance, and strategies and tactics for implementing a verification reuse methodology. Part 2, *Verification Reuse and Beyond: System Level Verification*, continued the theme of Part 1 to explore where leading edge verification is going, and what is possible in terms of attacking certain aspects of design quality and productivity, and deliver the next big step in verification productivity.

These archived webcast events can be viewed by DMR readers anytime at:

www.denali.com/verification_series_200412.html

On December 7, Denali Chief Verification Architect, Sean W. Smith, presented a webcast seminar addressing advanced techniques for modeling and simulating memory devices for functional verification. The webcast, *Denali MMAV—Advanced Applications and Techniques for Today's Testbenches* is targeted at design and verification engineers, and provides detailed techniques and methods for using MMAV™ to increase functional verification productivity for chip designs that interface to external memory devices. Topics include: Model quality and completeness, integration with commercial Verilog/VHDL/C-based simulators, leveraging verification/testbench languages, memory transaction analysis, and debugging. Examples are given, and apply to of all memory types, such as DDR-DRAM, SRAM, Flash, Flash-Cards, and more.

This archived webcast event can be viewed by DMR readers anytime at:

www.denali.com/tv/Verification20041207

Denali MemCon 2005

Denali MEMCON

TAIWAN	MARCH 15, 2005
AUSTIN	APRIL 28, 2005
TOKYO	SEPTEMBER 15-16, 2005
SAN JOSE	OCTOBER 12-13, 2005

Focused on semiconductor memory and other critical chip interfaces such as PCI Express, Serial ATA, and USB, Denali MemCon provides a unique forum for exchanging information on product roadmaps, technology for power and performance, market trends, design tools, IP, and advanced system design issues. MemCon is held annually in several cities including Austin, San Jose, Tokyo, and Taiwan. We encourage you to attend an event near you and learn more about the very latest developments in memory, storage, and chip-to-chip communications.

For speaking opportunities and sponsorship information, please visit: **www.denali.com/memcon/sponsor2005.html**

For general information, including event registration, presentation materials from previous events, or to view upcoming conference agendas, please visit: **www.MemCon.com**

MemCon Sponsors

Samsung's Jim Elliott (left) and Micron's Mike Black speak out on their experiences at Denali MemCon San Jose 2004

View Video

Denali Memory Report

Denali MemCon Taiwan 2005

Seating is limited, register today at www.MemCon.com!

MemCon Taiwan 2005 will be held on March 15 in Hsinchu, and is expected to draw well over 600 attendees to participate in this exclusive event featuring: keynote presentation from Pat Hays, Vice President of Engineering for MIPS Technologies, Inc., a panel discussion on *Consumer Electronics: The New Era for Memory and Storage*, and over 11 presentations from industry experts on topics such as:

- PSRAM, CellularRAM™ and LP DRAM for mobile applications
- Controller solutions for Flash, DDR, and hard drive interfaces
- New memory and storage options for mobile consumer products
- Memory interface considerations for DDR and XDR systems
- Using Verification IP to accelerate functional verification
- Configurable SOC interconnect IP for wireless handsets
- System-level design and integration for PCI Express systems

Breakfast and lunch will be served, and Denali will host an evening cocktail reception immediately following the last presentation. Registration is free for industry professionals.

Platinum Sponsor:

OCP International Partnership

Gold Sponsors:

LSI LOGIC

Micron

M-Systems Flash Disk Pioneers

Silver Sponsors:

ARM

Rambus

VIRAGE LOGIC

Rambus Design Seminar:

"Complete Vertical Solutions for PCI Express Designs"
Hosted by Cadence, Catalyst, Denali, and Rambus

Ontario Canada—Thursday, February 24, 2005 at 8:00 a.m.

Register online at: www.rambus.com

Coming Soon! MultiMediaCard Webcast

February 24, 2005 11:00AM PST

www.denali.com/tv/SamsungMMC20050224

Introduction by Denali's Memory Market Analyst Lane Mason
Interactive Presentation by Andy Yang, MultiMediaCard Marketing Committee member and Strategic Marketing Manager at Samsung Semiconductor

Denali Memory Report

MemCon Taiwan 2005 Agenda: March 15, 2005		
8:00AM	Registration	
9:00AM	Welcome, Opening Remarks	Sanjay Srivastava President, CEO Denali Software, Inc.
9:05AM	MIPS TECHNOLOGIES	Keynote: Enabling the New Revolution in Consumer Products Pat Hays Vice President of Engineering MIPS Technologies, Inc.
9:45AM	OCP International Partnership	Presentation: System Level Design Using OCP Based Transaction Level Models Chien-Chun Joe Chou & Anssi Haverinen Director of Architecture & Specialist Nokia Platforms OCP International Partnership & Nokia Technology Platforms
10:15AM	Break	
10:30AM	Micron®	Presentation: A Performance Comparison of PSRAM, CellularRAM™ and LP DRAM in Mobile Applications Peter Feeley Applications Manager Micron Technology Inc.
11:00AM	denali	Presentation: Databahn™ Controller Solutions for Tomorrow's DDR, Flash and Hard Drive Interfaces Brian Gardner Director of IP Product Marketing Denali Software, Inc.
11:30AM	M-Systems Flash Disk Pioneers	Presentation: To Disk or To Flash? What Storage Technology Will Dominate in Music-Centric Handsets Arie Tal Director of Marketing M-Systems
12:00AM	Lunch	
1:00PM	Panel Discussion	Consumer Electronics: The New Era for Memory and Storage Moderator: Lane Mason, Memory Market Analyst, Denali Software, Inc. Panelists: Ian Mackintosh, President & Chairman, OCP International Partnership Brian Gardner, Dir. of IP Product Marketing, Denali Software, Inc. Harmel S. Sangha, Director of Coreware Marketing, LSI Logic Corp. Arie Tal, Director of Marketing, M-Systems Peter Feeley, Applications Manager, Micron Technology Inc.
1:45PM	Rambus®	Presentation: DDR and XDR Memory Interface Considerations Rich Warmke Marketing Director Rambus
2:15PM	VIRAGE LOGIC	Presentation: Yield Management for Cost-Effective Manufacturing of Next Generation Consumer Application Krishna Balachandran Senior Product Marketing Manager Virage Logic Corp.
2:45PM	LSI LOGIC	Presentation: LSI Logic Eases Your Way Onto The PCI Express Bus Harmel S. Sangha Director of Coreware Marketing LSI Logic Corp.
3:15AM	Break	
3:30PM	denali	Presentation: Verification IP: Accelerating Functional Verification of Data Interfaces and System Designs Sean W. Smith Chief Verification Architect Denali Software, Inc.
4:00PM	OCP International Partnership	Presentation: Configurable SOC Interconnect IP for Tomorrow's Wireless Handsets Chien-Chun Joe Chou & Steve Hamilton Director of Architecture & Field Application Engineer OCP International Partnership & Sonics, Inc.
4:30PM	ARM®	Presentation: PCI Express PHY System Level Integration Considerations Jeff Reynolds Technical Marketing Manager ARM
5:00PM	Hosted Bar Reception	MemCon Sponsor Exhibitions

Company Financials for 4Q FY 2004

Overall, 2004 was a good year, with the chip market growing about 28 percent ($164B to $213B), and profits improved dramatically from 2003. But it ended on a sour and uncertain note, which detracted from some of the cause for celebration. Memories grew nearly 46 percent from 2003, and turned from loss to profit for most vendors. But the course of the year was frighteningly like the 1999-2000 upturn, which crashed into bath of red ink soon after the year closed; let's hope the future looks better than what followed that strong cyclical market. Both markets built out of a major downturn (1997-98 and 2001-03), both had barely recovered profit equilibrium before problems appeared (this time, inventory buildup, and price pressure), and premonitions of tragedy from mid-year—2000 and 2004. After the year closed, 2000 had the worst fears exceeded in the downturn of the following three years. The pessimists proved to be optimists. This time, 2004's rising market petered out in the fall, and opens 2005 with layoffs, Capex uncertainties, facility and manufacturing rationalization (just the beginning), and preemptive layoffs to reduce costs, and recover profitability.

The biggest difference is that 2005 promises to be a good global economy, while 2001-2 was a disaster amid the bursting bubble economy of vast but still not fully appreciated consequences. Though the chip business marches to a somewhat different drummer than the overall economy, it certainly benefits from strong demand, and decidedly from Americans who can buy foreign goods with cheap dollars and have no apparent reluctance to go into debt and overspend their incomes, individually or collectively.

For memory makers, like the rest of the industry (though more consistent in their performance), 4Q04 was the time to say goodbye to good times. In the largely communications segment, Cypress lost money and IDT also lost money. They both announced layoffs. In NOR flash, Intel's aggressive flash pricing gained it market share but pulled AMD/Spansion into the red in 4Q04. Intel has been red in flash for four years, and stayed there, though it has improved the profitability of the recently combined Communications Division, in which flash is organized. SST

Figure 1. Samsung 4Q04 Results Summary

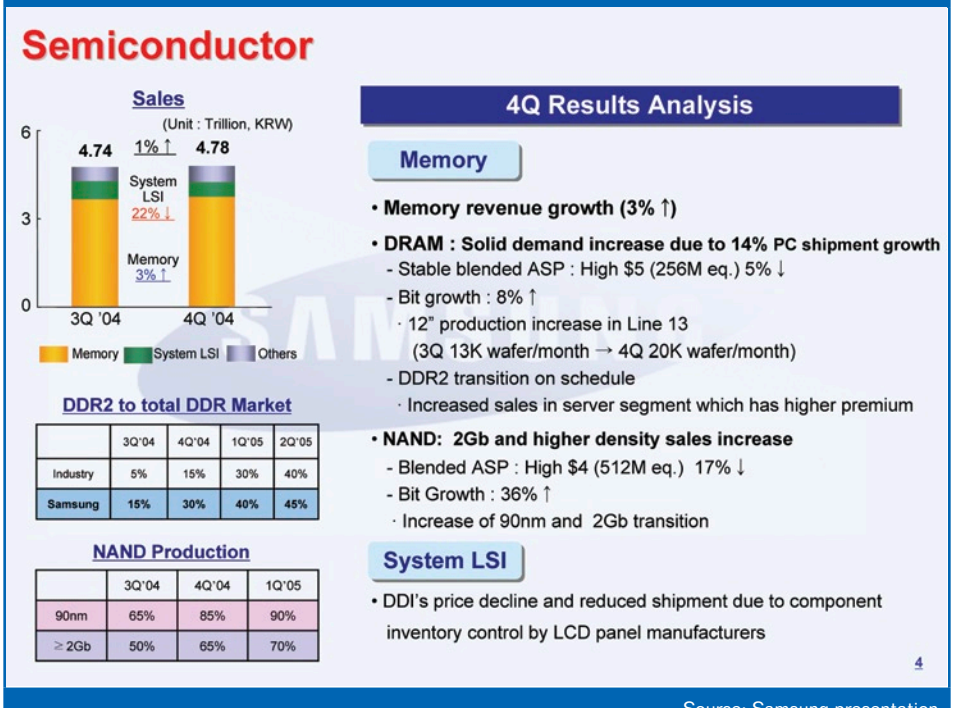

Source: Samsung presentation

lost money in NOR, Atmel substantially reduced its overall loss from 3Q04, but we don't know how well they did in their flash business unit. ST Micro, always tightly managed, made money in memories overall, which are largely NOR flash.

Samsung was, as usual, unique, showing good profits in its memory division

One analyst thinks Toshiba is more profitable in NAND flash than Samsung; as noted above, it is unlikely that any of the NAND flash newcomers have been able to cost reduce as fast as prices have dropped. Sandisk noted a 30 percent price drop in card prices in 4Q04, after a similarly large 22 percent in the summer quarter. Who can bring costs down that fast?

Taiwan DRAM makers, excluding ProMOS, which had not reported at press time, reported good profit results, while Micron continued to struggle through its newly adopted "avoid commodity DRAMs" strategy, instead focusing on Cellular RAM, image sensors (!), and legacy SDR DRAMs, and trying to get into the fast-growing NAND flash market. They had reported their 1Q05 about a month earlier, as their fiscal year ends at the ends of August.

Licensing and Royalties

On another note, Denali has not always been consistent in its treatment of companies' licensing and royalty income, though

Table 2. Memory Company Financials, 4Q04 and FY2004

	Sales					Profits			
	4Q04	3Q04	Yr 04	Yr 03	Growth	4Q04	3Q04	Yr 04	Yr 03
AMD Flash	504	538	2343	1419	65.1%	-39	15	35	-3
Atmel	408	413	1652	1329	24.3%	-7	-28	-45	-118
Cypress	210	220	948	837	13.3%	-18	4	35	-6
Hynix	NA	1350	NA	3080	74.4%	NA	466	NA	-2099
Intel/Comm	1364	1327	5027	3928	28.0%	-196	-251	-792	-824
IDT	96	97	388	332	16.9%	6.3	9.5	19	-243
Infineon	996	1020	3794	3026	25.4%	255	188	570	31
ISSI	31	31	171	117	46.9%	-13	-16	-11	-15
Macronix	190	192	728	514	41.6%	15	15	46	-238
Micron	1260	1189	4557	3514	29.7%	155	94	312	-956
MoSys	1.5	1.7	11	19	-42.5%	-4	-5	-12	3
Nanya	346	311	1229	825	49.0%	37	64	220	-42
Powerchip	533	476	1740	665	161.6%	187	207	647	5
ProMOS	NA	324	NA	729	81.8%	NA	95	NA	-43
Rambus	39	39	145	118	22.5%	6.5	10	34	23
Samsung	3603	3190	12579	7922	58.8%	1326	1401	4981	1653
Sandisk	549	408	1777	1079	64.7%	78	54	267	205
ST Micro	508	505	1974	1358	45.4%	16	27	87	-46
SST	104	112	449	295	52.3%	-27	14	24	-65
Vanguard	125	129.1	492	317	55.2%	30	38	131	11
Winbond	208	224	970	859	12.9%	3	26	124	-32
Sum	12874	12097	47669	30211	57.8%	2230	2428	8666	-2412
All others	2225	2025	7655	3956	93.5%	325	350	1125	121
Total	15099	14122	55324	34167	61.9%	2555	2778	9791	-2291

Source: Company reports and Denali

we've footnoted it from time to time, and remarked on its relative stability compared to the volatility of sales and profits. This is due in part to the fact that companies rarely (or often) do not break out their royalty payments (and indeed, may not even include them in their COGS or expenses in their reports). This quarter, Infineon takes yet another pretax gain of 118M Euros (about $156M), as a technology-licensing payment from ProMOS, which payment dwarfs its operating income from DRAM sales. This time we've included it as a genuine (though non-recurring) operating profit. ProMOS, which has not reported yet, in the past has not called out licensing payments as a standalone factor in their earnings release (though perhaps they do filings with the Taiwan Stock exchange).

Every player in the memory business gets or pays royalties, and if they do not, they do wield their IP to gain equity in manufacturing ventures (like Nanya's and Infineon's DRAM manufacturing JV, Inotera, or ProMOS recent DRAM technology deal with Hynix). We have often commented on the flash technology royalty income streams on SanDisk and SST, which are called out in their financial statements. Rambus is almost pure IP payments, with a small fraction of income being engineering and design services. MoSys, though their star has now passed into retreat, also derived most of its income from IP. AMD and Infineon pay Saifun for "mirror bit" flash technology. Intel gains an uncertain amount of income from its memory patents; Micron, before the acquisition of TI's memory business in 1998, used to call out its royalty payments (which ran about 10 percent of revenues) each quarter, but, being granted a 10 year royalty "holiday" as a part of that TI Business unit acquisition, soon stopped such reporting.

When I was an employee of Vanguard, we had bought the Mitsubishi 0.18µ DRAM technology license, and 64M and 128M designs, for amounts comparable to those recently paid by ProMOS to Infineon. While at IBM, we had licensed to Nanya our DRAM technology (and, as a matter of fact, IBM's trench DRAM technology, later propagated to Toshiba, Infineon, Nanya, Winbond, and now used in SMIC's manufacturing process, was all originated in IBM's DRAM business in the late 1980s, but now only remains in IBM's eDRAM technology, in a performance optimized-mode).

For the IP heavies, 2004 was a good year; Sandisk brought in $174M in licensing and royalty income; SST brought in $38.5M—both all in flash technology. Rambus brought in $120M for DRAM technology licenses (SDR, DDR, RDRAM, XDR) plus some incidentals.

INTERVIEW

DMR Interviews Bill Gervasi, Netlist, "The Industry's Move from Parallel to Serial Interfaces"

Over the past several years, many of the widely-used industry standard interfaces have adopted a new tack, to accommodate higher data rates amid the constraints of chip and board cost, legacy interfaces and long-evolved system configurations.

Just within the past 18 months, we've seen a kick-start and rapid acceptance of PCI Express (PCIe), to replace PCIx and PCI 2.0 in many parts of the standard PC. Intel, of course, was instrumental both in developing and promoting this spec. We've also seen a close cousin moved forward, ASI, or Advanced Switching Interface, into the server and communications domain.

For hard drives, Serial ATA is set to replace ATA over the coming year or so; the now ubiquitous USB port, first launched in about 2001, provides an easy to use port for printers and, with low cost flash drives, has virtually replaced floppy disks as the most commonly used method of transferring files from one computer to another.

Denali Memory Report

Serial RapidIO is replacing RIO interfaces; parts of the powerful PCIe spec are being adopted in the emerging Fully Buffered DIMM (FB DIMM) spec and as a base of ASI (Advanced Switching Interface).

Everything is changing from one relatively slow clock rate/low bandwidth, wide parallel interface to a much higher clock rate narrow interface.

In this issue of *DMR*, we have with us Bill Gervasi, Senior Technologist of Netlist, Inc., to talk some about this interface trend. Bill is a leading participant within the JEDEC standards committee, which is responsible for the development of many of these new interface protocols. He'll be explaining some of the events, which are driving these changes, and elaborate on where it all might end.

Bill was a long-time Member of Intel's Senior technical staff, before he joined graphics chipmaker, S3 in 1997. From there he went to Transmeta, which had the audacity to tackle Intel head-on in the laptop MPU space with their portfolio of low-power technologies. He joined Netlist in 2004, where he is now Senior Technologist of Netlist, an Orange County based memory modules supplier.

Here's what Bill had to say, to help understand this major industry transition.

Denali Memory Report (DMR): What is the Number One reason that systems have, over the past 2-3 years migrated to serial interfaces?

Bill Gervasi (BG): First of all, the interfaces are not "serial" so much as "narrow". The reason for narrowing the buses is to be able to transition to uni-directional differential interfaces that can be boosted to very high signaling rates. Differential signaling gives a nominal 2x improvement in switching edge rates; unidirectional signaling separates the transmitting operation and silicon from the receiving end, also making for improved signal integrity (although in this case, with the addition of an I/O).

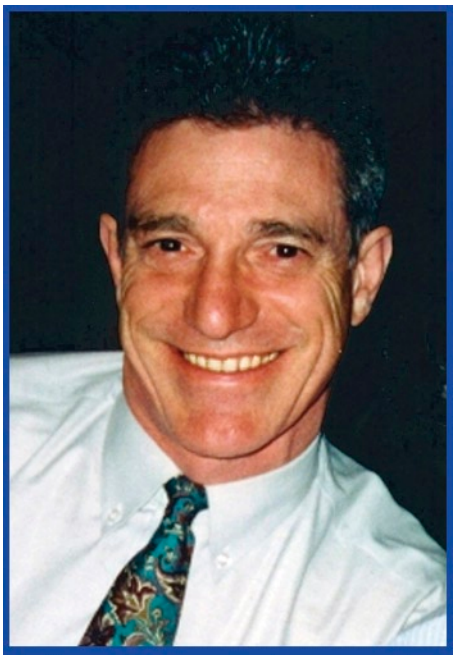

Bill Gervasi
Senior Technologist
Netlist, Inc.

DDR2 DRAM strobes are differential and remain bi-directional. DDR3 looks like its headed the same way, so at least DRAMs have not yet moved to unidirectional interfaces.

DMR: In some applications that have been brought to Denali's attention by customers, low pin-count and low-density serial flash are attractive because the actual package was physically smaller—fewer pins due to serialization of interface. It was not just a high pin count issue.

BG: Once the transition to high-speed buses is made, the pin count can decrease, though not always. The tradeoff is that power consumption can increase due to the speed and often, the need for a PLL or similar locking mechanism to track the bit rate.

DMR: Is the apparent move to FB DIMMs truly a part of this industry-wide trend or is it to more to solve the limited ability of DRAM controllers to address large arrays in servers using traditional Registered DIMMs?

BG: The FB-DIMM is needed when very large capacity memory subsystems are needed. Direct point-to-point buses will dominate commodity systems for the foreseeable future, i.e., DDR3.

There may be an alternative to the FB-DIMM based on the Registered DIMM concept that will challenge the middle ground as well.

DMR: On the chip level, what does it entail to change the interface from parallel to serial, so that the output is serialized already?

BG: Nothing special—standard SERDES logic, although SERDES design skills are not found of use any where.

DMR: What can you say about incremental die area to add serdes on the "DRAM"?

BG: The die area needed for a serdes is not significant; however PLLs for high speed operation can be a significant die area increment, especially given the temperature range over which it needs to operate

DMR: Is there a chance (or, What are the pros and cons?) of serializing the data ON the chip, like RDRAM, vs. having parallel outputs and serializing the data from several chips in an "AMB-like" separate chip, such as is being developed for Fully buffered DIMMs? .

BG: Why pay for it in every DRAM? There are some very well-established economics in they way things are done today, based on the wide variety of system architectures, the realities of the industry's diverse technical skills and system manufacturing costs, and the economics of (low-cost) system design. These are highly evolved over decades and will not change overnight. To serve the broad market, a legacy of yesterday's broad market, standards must recognize "given cost effective constraints"—in this case, DRAMs are as low cost as they are because the industry has resisted every attempt to put "optional" and "specialized functionality" on each chip for thirty years. For the industry as a whole, with many diverse applications, it is better to keep the DRAM on the steep cost reduction curve, and put the special functionality in a separate chip that can serve many DRAMs. VRAMs, SGRAMs, even RL DRAM and FC RAM have limited cost-reduction potentiality because they add special functionality on a standard DRAM, which limits their market and cost reduction potential.

DMR: Except for RDRAM, are there any RAMs (SRAM or DRAM) that appear headed to serial interfaces: DDR2 is parallel, and DDR3.targeted for the market in 2007-2010, also will be parallel. After that?…DDR4…SRAM and flash?

BG: We're not likely to see serial interfaces in DRAMs. As long as the ratio of DRAMs to controllers remains high, it is better to centralize the serialization function. In flash, there is a nascent move to serial interfaces, but this seems driven as much by getting the chips into smaller package, and packages with fewer pins in low flash density and cost and space sensitive applications. Most serial flash applications require only a few Mb, probably 8-16Mb, which is a small die that is often pad limited.

DMR: Do you have insight into whether the proprietary interface standards, like Sun HP Cisco, and IBM may have in their systems, are moving to serial as fast as the open standards, like (S)ATA, PCIe, (S)RIO, USB….?

BG: For very high capacity systems, this seems to be a likely future. For mainstream systems, like we see in most of the consumer space and in PCs, alternatives are likely to dominate in the memory interface. The other functions coming off the South Bridge are running to their own roadmap, which is moving in the narrow-channel, high clock rate direction, also.

In the hard drive arena, the move from ATA to SATA seems largely to have been driven by the limitations of the wide HDD ribbon cable. SATA will make easier management of the mechanical aspect of hooking up HDD

and enabling it to function at the higher datarates with today's drives.

DMR: With the ability to do wide busses on-chip (like in eDRAM), is there any reason to consider serial links, and serial data transfers in embedded applications?

BG: The answer is yes and no. Embedded memories can just as easily, or more easily use parallel interfaces to achieve the same bandwidth requirements while keeping power under control. The motivation to interface serialization in embedded memories is not at all the same as in systems containing a multiplicity of chips.

If we look at the major DRAM applications, we see that the ratio of 18 DRAMs (or fewer, on "value systems") to each CPU/controller seems to be pretty constant. Packing this kind of memory density onto the main chip seems unlikely.

DMR: What signaling changes will be optional and which will be mandatory as serialization takes place, and signaling clock rates go up by multiples of 6-8-10 into the GHz zone, or more? Clock sent with signal, differential signaling?

BG: I suspect this will be driven by PCB technology. If it is too expensive to mass produce a PCB that can carry 10GHz, then the bus will simply need to go wider to accommodate the limitations of the carrier. There remains always the balance of clock rate and bus width—sometimes wider is cheaper, sometimes faster clocks is cheaper. Cost matters, immensely, and there are always costs to forcing change on the industry from the established way of doing things (though rewards might be down the road).

Systems designed and made to custom specs, in high-end system houses, may deviate significantly from the "down and dirty" lowest possible cost commodity system marketplace. But their system economics are different, and their market is more specialized. You saw what happened with Rambus' RDRAMs—tough board design issues limited it take up in the marketplace.

DMR: What do you see as the major technical challenges in this transition…any? Many?

BG: I don't anticipate changes that will require a revolution. The controller complexity will likely continue to increase with things like per-bit deskewing, more pre-emphasis or de-emphasis, edge rate control, etc. Smarter packaging at the device level will help delay the need for complex packaging at the system level.

DMR: How about a capsule summary?

BG: The move to narrow channel, higher speed interfaces is driven in different segments by differing needs and constraints. Maintaining low system cost and facing up to the higher bandwidth requirements are the most important driving factors, but how these forces drive and constrain interface evolution depends on the specifics of the system and interface—and on the larger industry in which it is embedded.

For the time being, narrow channels and higher speed interfaces have proven their value. They are technically doable, and promise rewards of high bandwidth and simpler packaging and board design and layout. But their adoption is by no means universal, or on the same track across all segments, or should it be. In addition, the potential, in some applications of increased power due to those same higher clock rates, and the potential need to add PLLs on the interface, pose some potentially costly obstacles that are to be avoided if at all possible.

DMR: We'd like to thank Mr. Gervasi for sharing his insight with us.

DATABAHN

THE GOLD STANDARD FOR DDR INTERFACES

Over 120 design-wins, 50 chips in production

100% first-pass silicon success

Supports all SDR/DDR1/DDR2 & Mobile SDR/DDR devices

Configurable for Power, Performance, Multi-Port Access

Includes Integrated PHY

denali

www.denali.com/databahn