Harvard Business School **N9-701-056**
January 24, 2001

Rambus, Inc.: Commercializing the Billion Dollar Idea

In January, 2001, after seven years of development and three years of production, Rambus's DRAM memory technology, RDRAM, could be found in more than 150 products. Yet Rambus had faced several unexpected obstacles in its attempts to make RDRAM the dominant memory architecture in personal computers, the single largest market for memory. The last six months had witnessed typically ambiguous events. On one hand, Rambus chalked up a significant increase in revenues during 2000, and RDRAM shipments were at an all-time high. On the other hand, DRAM manufacturers – many of them licensees of Rambus technology – continued to pursue alternate approaches to speeding up DRAMs.

As Rambus faced another challenging year, industry observers wondered how the firm could achieve its goal of becoming the dominant DRAM architecture in PCs – and whether that was still the appropriate goal.

The Semiconductor Industry

Semiconductors were first commercialized on a large scale in the 1950s, and by 2000 were credited with much of the productivity growth in Western economies. Scientists at AT&T's Bell Laboratories developed the first transistor in 1948 to replace vacuum tubes with a more reliable way of regulating electronic currents. Transistors used semiconductor material to amplify an electronic signal. Made of silicon or germanium crystal, and often measuring no larger than a human finger, transistors could control the flow of a very large electric current by means of a very small electric current at another point on their surface, serving as switches.[1]

Unlike vacuum tubes, transistors were virtually unbreakable. They marked an improvement over conventional vacuum tubes in at least three ways: they generated less heat, consumed less power, and were much smaller.[2] In the 1950's and 1960's, transistors were used in a variety of electronic products, including hearing aids, television cameras, portable and car radios, and television receivers. The devices also proved useful to the mainframe computer industry: The transistor's smaller size, cooler temperature, and lower power consumption relative to vacuum tubes helped produce a much smaller and less cumbersome machine.[3]

Research Associate Briana Huntsberger and Professor Brian Silverman prepared this case from published sources as the basis for class discussion rather than to illustrate either effective or ineffective handling of an administrative situation.

The first integrated circuit, an electronic device using semiconductors, was invented in 1958. The integrated circuit combined transistors with resistors and capacitors, other electronic devices, to make a device that was much smaller and less fragile than the transistor. Because integrated circuits could be miniaturized, they made the development of the microcomputer in the 1960's possible. [4]

In 1971, Intel created the first microprocessor, a type of semiconductor that made it possible to coordinate several integrated circuits in the same machine. The development of the microprocessor led to the appearance of the first personal computers in the mid-1970s. Over the next 20 years, semiconductors came to be used in nearly all of the most widely-used pieces of major consumer and business equipment, including video games, televisions, personal computers, microwave ovens, and cell phones. The semiconductor industry became one of the largest in the world, with global sales of semiconductors and related products reaching $149 billion in 1999.[5] This industry was composed of memory chips, microprocessors, and other devices.

Memory Chips

The main function of memory chips was information storage. Memory chips were either read-only memory (ROM) or random-access memory (RAM). A DRAM, or "dynamic random access memory" chip, stored digital information in the form of bits, and provided high-speed data storage and retrieval. DRAMs were high-density, low cost per bit devices. Due to their low cost and high density, DRAMs were the most widely used memory chips in PCs and related products. The static random access memory (SRAM) chip operated much faster than the DRAM chip, but was more expensive than the DRAM chip and therefore used less often.

Memory chip production was notorious for its boom-and-bust cycles, and particularly for the cut-throat price competition that was sparked by overcapacity during industry downturns. The DRAM market shrank from a peak of more than $40 billion in 1995 to less than $15 billion in 1998, but was forecasted to exceed $60 billion at its next peak in 2002-2003. Throughout the 1990s, the memory chip business had become increasingly concentrated, with the top ten DRAM producers accounting for more than 90% of DRAM sales (**Exhibit 1**). Approximately half of all DRAM sales were to PC customers, with video games and PC peripherals accounting for much of the remainder (**Exhibit 2**).

Each generation of DRAM sold at a high price at introduction. As improved production yields led to lower production costs, and as more capacity and new entrants entered the market, price would typically fall to less than 20% of launch price. [6]

Microprocessors

Microprocessors had a more complicated set of functions than memory chips. Used primarily in computers, where they were sometimes referred to as central processing units, microprocessors controlled and coordinated data processing from all parts of the computer. Each microprocessor contained thousands or even millions of transistors packed onto a single chip, making the microprocessor the source of a computer's power. The best known microprocessors were Intel's Pentium series.

In the course of performing operations, a microprocessor frequently needed to store and retrieve information. Each data storage or retrieval action required the microprocessor to "talk" to the DRAMs connected to it. For example, when performing an arithmetic operation such as adding 2 and 3, the microprocessor would store the numbers 2 and 3 in a DRAM, then retrieve both numbers, add them, and store the result in the DRAM as well.

Microprocessor production had been dominated by Intel Corporation since the mid-1980s. In the late 1990s, Intel microprocessors held more than 90% market share in the PC market. However,

rivals such as AMD and Cyrix had closed the product innovation gap with Intel, introducing microprocessors of comparable power concurrently with (and sometimes even ahead of) Intel.

Technological Advance and "Moore's Law"

It was Intel's co-founder, Gordon Moore, who predicted in 1965 that the number of transistors per given area of silicon would double every 18 months. Through the late 1990s, "Moore's Law" had proved a reliable predictor of the pace of technological advancement in the semiconductor industry. The density of memory chips increased from 1 kbits per chip to 256 mbits per chip, a factor of 256,000. Microprocessor performance also increased in line with Moore's predictions (**Exhibit 3**).[7]

As semiconductor designers shrank the linewidths of transistors, enabling the exponential increase in transistors per chip, the equipment required to manufacture semiconductors became increasingly complex and expensive. The cost of a semiconductor fabrication facility, or "fab," rose from less than $50 million in 1970 to more than $2.5 billion in 2000 (see **Exhibit 4**).

Beginning in the early 1990s, the Semiconductor Industry Association, a trade association for producers of semiconductors and ancillary equipment, initiated a series of meetings to identify challenges to continued innovation in semiconductor technology. These meetings led to the development of a "National Technology Roadmap for Semiconductors" in 1992, with updated versions in 1994 and 1997. In 1999, the SIA released an "International Technology Roadmap" that for the first time explicitly included input from non-U.S. firms.

The Roadmap exercise began with the question, "what do we have to do to continue improving semiconductors according to Moore's Law?"[8] Semiconductor design and production was decomposed into several topic areas such as lithography, devices, and packaging, and technology working groups composed of engineers from multiple firms worked to identify those technological obstacles that could impede continued improvement. Although the Roadmap often indicated promising lines of attack, it rarely promoted specific fixes for these obstacles, leaving such work to individual firms.

In the late 1990s, the push for faster semiconductors approached a new bottleneck. Although microprocessors continued to increase in potential speed, they could not run at top speed because memory chips were not able to supply them with data quickly enough.[9] One observer compared this problem to a house that grew ever larger and more populated, but continued to have only one door:

> A memory chip can be likened to a house with its rooms arranged in rows and columns. Each room contains a single bit of data, either a zero or a one. A megabit chip…has 1,024 rows and 1,024 columns of rooms.
>
> The problem is that the houses have quadrupled in size every few years…[yet] the hallways have not been widened commensurately nor has the number of exits increased. Many DRAMs are still like houses with only one door that is wide enough for only one inhabitant to pass through…. A microprocessor these days might go knocking on the DRAM's door for a piece of data every 20-billionths of a second. But it takes 50-billionths of a second for the requested data to emerge from the house. Like a boy arriving to pick up his date, the microprocessor is often made to cool its heels.[10]

DRAM producers developed a number of innovations to speed up the flow of information, including additional pathways or wider pathways between DRAM and microprocessor. One method involved using wider "buses," or pathways along which data could travel, to increase data flow. Another method involved adding extra memory chips that would sit idle for much of the time, but that could allow a microprocessor to draw from multiple chips simultaneously at times of peak

bandwidth demand. A third method entailed linking SRAM caches and arrays of DRAMs in a complicated hierarchy to manipulate data more rapidly. A fourth method was to use higher speed memory devices.

Yet each of these innovations ran into new constraints, limiting their use beyond the late 1990s: wider buses required more pins, which would eventually raise packaging costs to unacceptable levels; adding generally-idle chips was also costly; and there were physical limits to linking SRAMs and DRAMs and to increasing the speed of conventional DRAMs. Many industry observers agreed that, at some point before 2000, a fundamentally re-designed DRAM would be needed to achieve sufficient increases in speed.

Rambus Inc.

Founded in 1990, Rambus was born during a dinner-table conversation between Mark Horowitz, an electrical engineering professor at Stanford University, and Michael Farmwald, an entrepreneur who had already founded two small chip companies, FTL (which was subsequently sold to MIPS) and Chromatic Research.

Farmwald had noted that although chip speeds were increasing exponentially, the links between chips were not keeping pace. In his view, the primary bottleneck stemmed from the fact that interactions between a logic chip and memory chip were sent via three separate messages – the address where the impulse is to go, the data bits themselves, and the instruction that says what should be done with that data (e.g., store it, replace it, etc.). The three messages traditionally ran on different paths at different speeds, and were recombined on the memory chip. The process paced the transmission based on the slowest route among the messages.

Farmwald and Horowitz came up with a different method: package the three messages for transmission down a single electronic line. This shift in methodology would require significant design changes in both memory chips and logic chips. The resulting memory chips would be slightly larger than current DRAMs, and their production would require the development and purchase of new manufacturing and testing equipment. But if such design changes could be pulled off, it would speed up computers by at least an order of magnitude.

Over the next six months the pair raised some $2 million, in exchange for 50% of the company, from venture capitalists. They also signed up Geoff Tate, a Harvard M.B.A. graduate who was a Senior Vice President at AMD, to be president of the venture. Horowitz took a leave from Stanford, and hired away from Xerox Palo Alto Research Center a former graduate student and two of his Xerox PARC colleagues. By the end of 1990, the team had designed the requisite changes in chip layouts, and by 1992 had applied for at least seven patents (see **Appendix 1** for a technical description of Rambus technology).

Rather than build the chips itself, either through construction of its own fab or through contracting with a foundry, Rambus chose to license its intellectual property to chipmakers. The company would then collect a royalty on each chip sold that included Rambus technology. Licensees paid an up-front contract fee ranging from a few hundred thousand dollars for a narrow license up to several million dollars for a broad license. Licensees also paid a per-chip royalty. The precise royalty rate was negotiated individually with each licensee, but was capped at 2.5% of the chip's selling price for DRAM and 5% of the chip's selling price for logic chips. In addition, Rambus required that licensees grant to Rambus a royalty-free cross-license on any patented improvements they developed that were related to the Rambus technology. In turn, Rambus would freely sublicense these improvements to all other licensees.

In March, 1992, Rambus announced that it had signed licensing agreements with three of the world's five largest memory chip producers: Toshiba, NEC, and Fujitsu. Each of the three licensees planned to develop, produce, and sell 4Mb and 16Mb DRAMs based on Rambus technology. Over the next four years, Rambus signed up more than a dozen more licensees, including nine memory chipmakers that accounted for more than 90% of all memory chips, as well as several logic chip manufacturers (**Exhibit 5**).

At the same time, Rambus focused its attention on courting "systems firms," those companies whose products required the use of memory and logic chips. In July, 1994, Nintendo announced that its Nintendo 64 video game system would rely on Rambus-based chips built by NEC and Toshiba.

Memory-Bus Battle: SYNCLINK and Intel Deal

By September, 1996, the Rambus chip was used in four products: Nintendo's Ultra64 game console; SGI's Indeo graphics workstation; Cirrus Logic's Laguna graphics board; and Creative Labs' GraphicsBlaster graphics accelerator board. These products generated roughly $100 million in Rambus chip sales in 1996 for DRAM vendors, the bulk of which were attributable to Nintendo shipments.

At the same time, an alternate high-speed memory technology known as SyncLink, or SLDRAM, began to generate interest among DRAM producers. In early 1996 a number of DRAM producers – including Fujitsu, Hyundai, Micron, NEC, and Samsung – formed the SyncLink Consortium to develop a synchronous DRAM approach to the memory bandwidth problem. Industry analysts speculated that the consortium was formed in part as a backlash against the royalties required by Rambus for its RDRAM technology.[11] In July, the consortium submitted its proposed DRAM architecture as an open standard to the IEEE, a governing association of electrical engineers (see **Appendix 2** for a brief description of synchronous DRAMs and SyncLink technology).

Rambus CEO Geoffrey Tate dismissed SyncLink as an uneconomical reinvention of the wheel:

> "The SyncLink group has discussed some interesting ideas, but they have not worked out the details or implemented them. To develop a high-bandwidth memory-interface technology takes years, tens of millions of dollars and very good engineering on the chip and the systems side. What is SyncLink's development schedule? Where is its engineering staff?" [12]

In response, SyncLink's proponents argued that Rambus's approach threatened the very health of the semiconductor industry:

> "If Rambus succeeds, it would give a very small group of engineers a technology hegemony…. It would be like creating another Microsoft. Rather than concentrate power in the hands of a few engineers, it is much better to keep it diversified among all of the DRAM engineers working around the world."[13]

As this battle heated up in late 1996, Intel and Rambus announced the signing of a development and license agreement under which both parties would cooperate to develop a specification for an extension of the RDRAM optimized for PC main memory applications. The contract also called for Intel to use "reasonable best efforts" to develop a PC main memory controller designed for use with such RDRAMs. The partnership's goal for the next-generation RDRAM, to be called "Direct RDRAM," was to achieve 1.6Gbps performance by early 1999.

Following this announcement, Intel met with major DRAM vendors to encourage them to use Rambus technology in future DRAM generations, and to determine precisely what needed to be done

to the RDRAM design to bring it in line with Intel's needs. Although Intel released little information in its public statements, managers privately noted that Intel decided that it had to take a leadership role by the end of 1996 to assure that new chips would be ready by 1998-1999:

> "Intel knows that the SyncLink DRAM will never be ready in time; it is a design-by-committee thing that got started too late to meet Intel's needs. But to take the Rambus architecture to the speeds that Intel needs, a lot of work still needs to be done. At those speeds, nothing is trivial. Intel will set the spec, and Rambus will implement it."[14]

Buoyed by the Intel deal, Rambus went public in May 1997. Its prospectus provided additional information regarding its arrangement with Intel. Specifically, Intel received a warrant to buy up to 1 million shares of Rambus common stock for $10 each. But Intel could not exercise the warrant until Rambus' technology comprised at least 20% of the chip sets that Intel sold for two consecutive quarters. The warrant would expire on December 31, 2000, if Intel did not reach the quota by then. Otherwise, it would expire on December 31, 2004.[15]

Once it reached the 20% figure, Intel could also send a representative to observe Rambus' board meetings. And upon buying at least 500,000 shares of Rambus stock, Intel would gain a seat on Rambus' board.

In addition, the Intel-Rambus deal gave Intel the right of first refusal if any third party offered to acquire Rambus. Thus, in the event of a bid to acquire Rambus, Intel could acquire the Rambus by matching the offer.

Although Intel subsequently indicated a willingness to support both RDRAM and synchronous DRAM in future generations of Intel products – and provided specifications for a converter chip in early 1997 – in August 1997 the company announced that its next generation chip set, to come out in mid-1999, would be designed to work only with Rambus-based DRAMs. Shortly after Intel's announcements of support for Rambus, several DRAM producers that had been supporting SyncLink or other rival DRAM technologies – notably, Micron, Mitsubishi, and Infineon (Siemens' DRAM subsidiary) – signed licensing agreements with Rambus (see **Exhibit 6**). By early 1998, Rambus had signed up the 13 largest DRAM suppliers in the world, and was endorsed by systems manufacturers including Compaq, Dell, and Sony. NEC continued to ship RDRAMs in large quantities, with Toshiba also shipping modest quantities. Two other licensees, LG Semicon and Samsung, prepared to enter production.

Direct RDRAM, DDR, and new challenges

By the end of 1998, six DRAM producers began shipping samples of Direct RDRAM, Rambus's second generation product. Designed in conjunction with Intel, Direct RDRAM was slated to appear in PC in mid-1999 and to take the lion's share of the PC market by 2000. Direct RDRAM was estimated to cost more than conventional DRAM to produce (estimates of the extra cost ranged from 10% to 100%), and initial sample shipments indicated that the price to OEMs would be 15-20% higher than that of conventional DRAMs.[16] Industry analysts noted that such a premium might slow the adoption of Direct RDRAM for mainstream PC applications – which would not be particularly onerous, since existing Direct RDRAM fab capacity could only produce 33 million units per year, and demand was expected to exceed 170 million units within 18 months.[17]

Concurrently with the debut of Direct RDRAM, an alternate technology designed to accelerate information flow also made its debut. By the end of 1998, several companies had shipped samples of Double Data Rate synchronous DRAM (DDR SDRAM), which effectively doubled the speed of a SDRAM. Where a SDRAM was synchronized to an external clock that made sure that it transferred data in synch with a microprocessor, DDR enabled the SDRAM to transfer data twice per

clock cycle. As a result, DDR could achieve peak bandwidth speeds comparable to those of Direct RDRAM (see **Appendix 2** for more on DDR technology). DDR provided this speed without requiring a major architecture shift from SDRAMs, so that much of the same production and test equipment could be used and the production cost would be comparable to that of SDRAM.

In late 1998, more than a dozen DRAM vendors – all of which were also Direct RDRAM licensees – announced their commitment to bringing DDR to market in 1999. One participant stated:

> "This was an effort to prove to people that the DRAM industry can get together and agree on a common spec. We expect DDR to go from near-invisible to the forefront of people's thinking in the next six months or so. Intel is going to do what they're going to do. But I don't think they can ignore what the rest of the industry is doing."[18]

Intel invested $500 million in Micron in October 1998 and $100 million in Samsung in January 1999, in exchange for their commitments to produce RDRAMs. Intel also continued its series of "Rambus seminars" around the world to encourage PC manufacturers to make the switch to Direct RDRAM. At the same time, Rambus offered to all of its licensees a warrants-based incentive program similar to the one included in its agreement with Intel. Licensees would become eligible to buy a limited amount of Rambus stock at $10/share once they achieved certain product qualifications and production targets.

Intel's first chipset designed to use two 800 GHz Direct RDRAMs – and hence to achieve 1.6 GHz -- was initially scheduled for release in July 1999. However, Intel delayed its launch three times. In February, the firm postponed the launch for two months in anticipation of a shortage of RDRAM chips. In April, Intel discovered a flaw in the RDRAM interface that necessitated a three- to six-month delay to fix, indicating a further delay in launch of at least one month.

In response to these delays, several DRAM producers postponed their plans to dedicate manufacturing resources to RDRAM. Observers noted that even staunch Rambus supporters such as NEC were leery of being saddled with excessive inventory should Intel delay further. These firms proceeded cautiously, hoping to ensure that ample core-logic chipset support would exist before committing entirely to Rambus. Although Intel continued to promote a September 1999 launch date, several DRAM producers aimed for production to begin in early 2000. Industry analysts conjectured that RDRAM capacity could reach 10 million units/month by March 2000.

In October, just a few weeks before launch, Intel postponed the release of its 820 chipset until first quarter 2000 due to a newly discovered flaw associated with a memory slot on the Rambus motherboard – although not with the memory device itself.

Concurrent with these developments, a number of DRAM manufacturers introduced a new SDRAM chip that could reach speeds of 133 MHz. Rambus proponents argued that this chip, known as PC133, did not offer the performance benefits Direct RDRAM does. DRAM manufacturers countered that the transition to the PC133 from the 100MHz SDRAM standard (PC100) was "easy and cost-effective compared with Rambus," forming a natural evolution as DRAM manufacturers shrank the die size rather than requiring a new architecture. PC133 proponents further claimed that PC133 cost about half as much as Direct RDRAM. Citing concern over the high price of RDRAM – whose 128 MB module was believed to cost about $230, compared to about $125-150 for SDRAM or DDR modules[19] – Intel announced that it would support PC133 SDRAMs "as an alternative and complementary memory architecture" by the first half of 2000.

In early January, 2000, rumors spread that Intel and the five largest DRAM manufacturers were meeting to identify an architecture to serve the PC market in 2003 and beyond. The Next Generation DRAM Alliance, as this effort was called, had the broad target of creating an open

memory architecture that would conform to the standards of JEDEC (Joint Electron Device Engineering Council), perhaps the foremost standards development organization in the semiconductor industry. Rambus was not included in this effort. Intel was also rumored to be working with third-parties to support DDR SDRAMs.

Legal Action

In mid-January, Rambus filed a lawsuit against Hitachi for patent infringement. Specifically, Rambus's complaint charged Hitachi with violating its licensing agreement by taking technology intended for the development of Direct RDRAM devices and applying it to a rival SDRAM architecture. Rambus sought punitive damages and a court injunction preventing Hitachi from making or selling most of its DRAM product line. Rambus later named Sega, a customer of Hitachi's SDRAMs, as a party to the suit.

In June, 2000, Toshiba announced that it had agreed to pay Rambus royalties on SDRAM, DDR, and controllers that interface with those types of memory – technologies that relied on "patents for fundamental high-speed memory interfaces invented by Rambus," according to Toshiba's statement. One week later, Hitachi announced that it had settled its suit with Rambus by agreeing to similar royalty terms. The royalty rates for synchronous memory were believed to be higher than those charged for Direct Rambus memory, although some observers speculated that Rambus offered Hitachi and Toshiba "sweetheart" deals with low royalty rates in order to gain the precedent of royalties for non-RDRAM chips.

Rambus immediately requested that several other DRAM vendors agree to pay royalties on SDRAMs. By October, the firm had also approached logic firms, including AMD and Transmeta, about similar arrangements. Avo Kanadjian, vice president of worldwide marketing at Rambus, explained, "We will seek smaller royalty amounts for conventional [single-data-rate] SDRAM, but greater fees for DDR."

Although some DRAM manufacturers agreed to pay royalties on SDRAMs to Rambus, several firms responded by meeting to discuss filing an antitrust complaint against Rambus. In October, Micron and Hyundai leveled antitrust charges against Rambus Inc. at the Federal Trade Commission, complaining that Rambus' synchronous DRAM patents were based on technology the design company learned of while participating in JEDEC deliberations on industry open standards between 1993 and 1995.

During the second half of 2000, Intel decided to accommodate SDRAM as well as Direct RDRAM in its Pentium 4 systems, and raised the possibility that its next-generation processor would be tied to DDR SDRAM instead of Direct RDRAM. Several observers saw this as a blow to Rambus's hopes to gain significant PC market share in 2001. Yet a few weeks later Intel began to offer OEMs a $60 rebate for every shipment of a high-end Pentium 4 desktop that used Direct RDRAM.

Into 2001 and beyond

By January, 2001, RDRAM was used in more than 150 products, from PCs to video-game consoles (including the Sony Playstation, on which Rambus was believed to earn a $2.40/console royalty) to communications equipment. Rambus had signed SDRAM licensing agreements with seven of the ten largest DRAM companies, with only Micron, Hyundai, and Infineon fighting back in court. The firm posted a dramatic increase in royalties during late 2000, and Direct RDRAM shipments had reached levels of millions (see **Exhibit 7** for financial statements). What success Direct RDRAM had achieved lay in the workstation and high-end PC market, and Rambus hoped to compete in the $1,500-$2,000 PC market in 2001. The firm expected that entry into the lower-priced,

higher-volume PC segments would generate sufficient demand to drive RDRAM production costs from their current 40% premium above that of SDRAM to less than a 20% premium.

Rambus had also recently demonstrated its next-generation technology, which could deliver speeds an order of magnitude greater than Direct RDRAM or DDR technology. Industry observers agreed that DDR and other rival technologies would be unable to achieve such speeds (**Exhibit 8**).

Yet the future course of DRAM technology remained hotly contested and difficult to discern, leading to continued battles to gain "design wins" for Rambus technology. Industry forecasts varied widely in their predictions of which technology would dominate in the market, with Rambus's predicted share of the market in 2002 varying from just over 1% to nearly 50% (see **Exhibit 9** for some forecasts).[20] Geoff Tate mused:

"[I'm fed up with][the spreading of fear, uncertainty and doubt [FUD] by certain parties who feel threatened by Rambus' technology and business model. The FUD view is that RDRAM is a failed technology with no future. The facts as we see them are that after a great deal of effort on the part of Rambus and many of our partners in the semiconductor industry, RDRAM is solidly established as a mainstream standard. In the 30-year history of the DRAM, RDRAM is just the fifth standard to reach high-volume status."[21]

Appendix 1: Rambus technology[22]

The Rambus solution is a system-level approach, rather than a chip-level approach, to solving the data flow bottleneck problem The Rambus approach encompasses three main elements: the Rambus Channel, the Rambus Interface, and the RDRAM.

The Rambus Channel is an interconnect between Rambus "masters" (e.g., microprocessor) and "slaves" (e.g., DRAM). The Rambus channel uses a small number of high speed signals to carry all addresses, data, and control information in blocks. The key feature of this channel is its ability to transfer data at rates up to 600 Mbps even though it is only 9 data bits wide. Conventional DRAM designs would require a wide, complex bus using a large number of DRAMs to match this performance level. The Rambus Channel also relies on only 32 pins to achieve its connection, far fewer than would be required with traditional DRAM technology (chip packaging costs rise steeply with the number of pins).

The Rambus Interface Cell contains the circuitry that connects custom logic on a master or a slave to the Rambus Channel. It is made to reside on an ASIC, and takes up an extremely small amount of space. The Interface uses a unique Rambus Signaling Logic technology to allow 600 MHz – 800 MHz data transfer rates with conventional system and board designs.

The Rambus DRAM consists of a standard DRAM core and a Rambus Interface Cell. The Rambus DRAM can deliver bandwidth ten times that of a conventional DRAM interface, thanks to its ability to send information in blocks and to its use of a synchronous transfer protocol in which the DRAM is controlled by the same clock that controls the microprocessor, thus reducing waiting times.

Appendix 2: Synchronous DRAM, SyncLink (SLDRAM), and DDR technology[23]

Synchronous DRAM (SDRAM) is synchronized to the system clock that controls the microprocessor. This synchronization enables the memory controller to know on which clock cycle data requests will be available. Data is then input to the "rising edge" of the clock (once per clock cycle) instead of with every two or three clock cycles, which is the norm for conventional types of DRAM. SDRAM also relies on up to four memory banks that operate simultaneously, in addition to a "burst mode" feature that addresses an entire block rather than just one piece of data. These features speed the DRAM further.

SyncLink, also called SLDRAM, is an extension of SDRAM architecture that extends the SDRAM's four-bank design to as many as 16 banks. In addition, SLDRAM adds a 400MHz packet-orietned bus and a new interface that reduces pin count to 50 to 60. With a 16-bit data path, a SyncLink interface could burst data at 1.6 Gbps.

Double Data Rate SDRAM (DDR) is an extension of conventional SDRAM. DDR adheres to the same design core as SDRAM, but includes two basic differences that enhance its speed. First, DDR uses more advanced synchronization circuitry than that used in conventional SDRAMs. Second, DDR uses more advanced signal technology to better handle incoming data from different memory modules. As a result, DDR essentially doubles the memory speed as compared to conventional SDRAMs. Consequently, DDR allows data to be read twice per clock cycle (on both the "rising" and the "falling edge" of the clock), thus providing twice the bandwidth of standard SDRAMs. DDR offers comparable performance to RDRAM in achieving speeds necessary for the current and near-term generations of computing technology. However, industry observers generally agree that DDR will reach its technological limit before RDRAM does.

Exhibit 1 Sales and market shares of top DRAM producers, 1995-1999

	1999	1998	1997	1996	1995
Memory					
Samsung	$4.75 (22.9%)	3.35 (24.1%)	3.70 (18.1%)	$4.12 (16.4%)	$6.46 (25.7%)
Hyundai	4.21 (20.3%)	1.48 (10.7%)	1.93 (9.9%)	2.30 (9.2%)	3.50 (13.9%)
Micron	3.32 (16.0%)	1.80 (13.0%)	1.65 (8.5%)	1.57 (6.3%)	2.48 (9.9%)
NEC [a]	1.72 (8.3%)	1.16 (8.3%)	2.30 (11.8%)	3.17 (12.6%)	4.74 (18.9%)
Infineon (Siemens)	1.67 (8.0%)	1.01 (7.3%)	0.99 (5.1%)	[b]	[b]
Hitachi [a]	1.08 (5.2%)	0.92 (6.6%)	1.90 (9.7%)	2.80 (11.2%)	4.44 (17.7%)
Toshiba	0.93 (4.5%)	0.73 (5.2%)	1.16 (5.9%)	2.23 (8.9%)	3.72 (14.8%)
Mitsubishi	0.68 (3.3%)	1.00 (7.2%)	1.22 (6.3%)	1.40 (5.6%)	2.21 (8.8%)
Fujitsu	0.55 (2.7%)	0.86 (6.2%)	1.15 (5.8%)	1.35 (5.4%)	2.06 (8.2%)
LG Semicon	[c]	0.89 (6.4%)	1.16 (5.9%)	2.00 (8.0%)	3.00 (11.9%)
TI	[d]	[d]	1.28 (6.6%)	1.60 (6.4%)	3.20 (12.7%)
Total top 11	$19.40 (93.6%)	$13.20 (94.9%)	$18.44 (94.1%)	$22.57 (87.8%)	$35.83 (89.8%)

[a] In 2000, NEC and Hitachi merged their DRAM operations into a new venture called Elpida.
[b] Siemens did not report separately its DRAM business until 1997.
[c] LG Semicon merged its DRAM business into Hyundai in 1999.
[d] TI's DRAM business was acquired by Micron in 1998.

Source: Semico, Inc.; IC Insights Inc.; Dataquest; casewriter estimates

Exhibit 2 Sales of selected products that incorporate DRAMs (million units)

	PCs	Video game consoles	3D Graphics controllers
1995	58.9	13.6	70.7
1998	89.8	16.7	79.0
2001 (est.)	136.5	24.3	107.5

Source: Dataquest, Inc.

Exhibit 3 Moore's Law

Chip	Year of Introduction	Transistors
4004	1971	2,250
8008	1972	2,500
8080	1974	5,000
8086	1978	29,000
286	1982	120,000
386	1985	275,000
486	1989	1,180,000
Pentium	1993	3,100,000
Pentium II	1997	7,500,000
Pentium III	1999	24,000,000
Pentium IV	2000	42,000,000

Source: Intel Corporation, http://www.intel.com/research/silicon/mooreslaw.htm

Exhibit 4 The cost of a semiconductor fabrication facility, 1970-2000

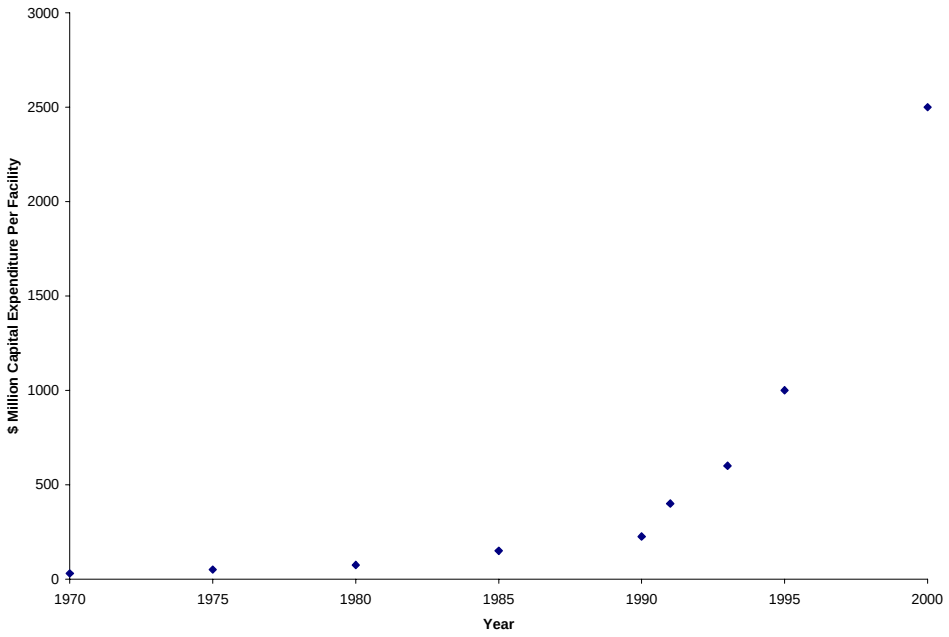

Source: Jonathan West, "Institutional diversity and modes of organization for advanced technology development: Evidence from the semiconductor industry," DBA thesis, Harvard Business School, 1996; casewriter estimates.

Exhibit 5 Rambus licensees

Licensee	Date of 1st License[a]
MEMORY	
Toshiba	1990
NEC	1991
Fujitsu	1992
Oki Electric	1993
LG Semicon	1994
Hitachi	1994
Samsung	1994
Hyundai	1995
Monolithic System	1996
IBM	1997
Micron Technology	1997
Mitsubishi	1997
Infineon (Siemens)	1997
Matsushita	1998
Vanguard	1998
Winbond	1999
LOGIC	
Toshiba	1990
NEC	1991
Cirrus Logic	1993
Chromatic Research	1994
LG Semicon	1994
LSI Logic	1994
IBM	1995
Intel	1996
SGS-Thomson	1996
Fujitsu	1997
TI	1997
AMD	1998
Compaq	1998
S3	1998
Matsushita	1999
H-P	1999
Agilent	2000
PMC Sierra	2000

[a] Some licensees did not renew when their initial licenses lapsed.

Source: Rambus 10-K reports; Rambus S-1 filing.

Exhibit 6 Rambus Licensees, May 1997

Licensee	Rambus	SyncLink	DDR
Fujitsu [a]		X	X
Hitachi	X	X	X
Hyundai	X	X	
IBM	X	X	
Intel	X		
LG Semicon	X	X	
LSI Logic	X		
Micron	X	X	
Mitsubishi	X	X	
NEC	X	X	
Oki	X	X	
Samsung	X	X	X
Infineon (Siemens) [a]		X	
TI [a]		X	X
Toshiba	X	X	
TSMC	X		

[a] Fujitsu, Siemens, and TI subsequently signed licensing agreements with Rambus by May 1998.

Source: "Rambus, Inc." Hambrecht & Quist research report, May 1997.

Exhibit 7 Rambus, Inc. Financial Statements ($000)

	2000	1999	1998	1997	1996	1995	1994	1993	1992
Contract revenues	$39,683	$35,353	$28,727	$20,186	$11,205	$7,364	5,000	3,371	1,916
Royalties	32,628	8,017	9,137	5,829	65	0	0	0	0
Total sales	72,311	43,370	37,864	26,015	11,270	7,364	5,000	3,371	1,916
Cost of contract revenues	12,093	12,232	8,988	5,491	4,821	5,236	3,844	1,950	1,053
R&D	11,501	8,123	9,649	9,815	5,218	3,117	3,067	4,291	3,546
SG&A	21,140	13,516	11,260	8,755	5,799	5,064	4,286	3,092	3,064
Stock-based comp. Expense [a]	171,085	0	0	0	0	0	0	0	0
Total expenses	215,819	33,871	29,897	24,061	15,838	13,417	11,197	9,333	7,663
Operating income (loss)	(143,508)	9,499	7,967	1,954	(4,568)	(6,053)	(6,197)	(5,962)	(5,747)
Net interest (interest expense)	4,714	4,339	3,361	1,342	439	322	(81)	(123)	115
Pretax income (loss)	(138,794)	13,838	11,328	3,296	(4,129)	(5,731)	(6,278)	(6,085)	(5,632)
Income tax	(32,667)	5,120	4,540	1,315	286	1,289	351	251	962
Net income (loss)	($106,127)	$8,718	$6,788	$1,981	($4,415)	($7,020)	($6,629)	($6,336)	($6,594)
Current assets	$142,937	$98,478	$91,793	$80,757	$10,145	$16,196	NA	NA	NA
PPE, net	6,724	4,232	3,989	4,338	2,340	1,598	NA	NA	NA
Deferred taxes	55,404	4,123	4,720	1,001	0	0	NA	NA	NA
Other assets	14,566	8,940	10,485	1,782	383	513	NA	NA	NA
Total assets	219,631	115,773	110,987	87,878	12,868	18,307	8,395	7,807	5,300
Current liabilities	33,187	36,704	32,175	31,021	15,072	11,763	NA	NA	NA
Deferred revenue	24,122	17,505	37,020	30,196	9,940	14,480	NA	NA	NA
Total liabilities	57,309	54,209	69,195	61,217	25,012	26,243	18,401	16,158	12,361
Common stock; paid-in capital	285,982	78,598	67,640	59,887	22,347	22,112	NA	NA	NA
Deferred stock-based comp.	(571)	0	0	0	0	0	NA	NA	NA
Accumulated deficit	(123,132)	(17,005)	(25,723)	(32,511)	(34,492)	(30,077)	NA	NA	NA
Accumulated other gain (loss)	43	(29)	(125)	(715)	1	29	NA	NA	NA
Total equity	$162,322	$61,564	$41,792	$26,661	($12,144)	($7,936)	($10,006)	($8,351)	($7,061)

[a] employee stock-based compensation expense results from paying stock options.

Source: Rambus 10-K reports; Rambus S-1 filing.

Exhibit 8 Expected life span of existing high-speed memory technologies

	SDRAM	DDR SDRAM	SyncLink SDRAM	Rambus DRAM
Current peak bandwidth	133MHz	200MHz	400MHz	1.6GHz
Current MHz	133MHz	200MHz	400MHz	800MHz
Year of obsolescence	2001	2003	2003	> 2006

Source: Toshiba, Rambus, casewriter estimates

Exhibit 9 Divergent forecasts, market shares of competing DRAM technologies (%)

IC Insights, Inc.	PC100 SDRAM	PC133 SDRAM	DDR	RDRAM	Other
2000	40	37	6	8	9
2001	15	55	13	10	7
2002	8	44	23	20	5
Seimco Research Corp.		SDRAM	DDR	RDRAM	Other
2000		84	3	2	11
2001		70	22	3	5
2002		61	35	1	3
Cahners In-Stat Group		SDRAM	DDR	RDRAM	Other
2000		75	3	6	12
2001		56	15	21	6
2002		40	19	36	3

Source: Day, J., "Navigating the DRAM market, technology maze," *Electronic Buyers News*, August 21, 2000, p. 57.

[1] Encyclopedia of Emerging Industries: Semiconductors (One Source).

Encyclopedia of Global Industries, #3674-One Source

[3] Ernest Braun and Stuart MacDonald, Revolution In Miniature (Cambridge: Cambridge University Press, 1978) pg. 78

[4] Encyclopedia of Emerging Industries: Semiconductors (One Source).

[5] Standard & Poors Industry Surveys: Semiconductors. Pg. 1

[6] Collis, D.J., G.P. Pisano, and P. Botticelli, "Intel Corp.--1968-97," Harvard Business School Case # 797137.

[7] "Rambus: Company Report" (Hambrecht & Quist, 6/9/1997)

[8] Semiconductor Industry Association, *International Technology Roadmap for Semiconductors*, 1999.

[9] Pollack, A. "Gridlock at the chip-to-microprocessor intersection," *The New York Times*, March 22, 1992, p. 11.

[10] Ibid.

[11] Lammers, D., "The Synclink DRAM finds new support," *Electronic Engineering Times*, July 1, 1996.

[12] Robertson, J. "Superfast DRAM Spec OK," *Electronic Buyers' News*, July 1, 1996.

[13]Ibid.

[14] Lammers, D., "Predicts It'll be PC standard by '99 -- Intel champions Rambus DRAM Electronic Engineering Times, " November 11, 1996.

[15] Sheerin, M., "Intel-Rambus relationship complex," *Electronic Buyers' News*, September 1, 1997.

[16] MacLellan, A., "Direct RDRAM has premium," *Computer Reseller News*, October 19, 1998, p. 175.

[17] Howle, A., "Direct Rambus to hit in mid-1999," *Computer Reseller News*, November 23, 1998, p. 36

[18] MacLellan, A., "DDR tugs tiger's tail," *Electronic Buyers' News*; Manhasset; Dec 7, 1998.

[19] Harbert, T., "Memory the Rambus way," *Electronic Business*, September, 2000, p. 87+.

[20] 1% figure from Semico, Inc. 50% figure from Dataquest, Inc.

[21] "Rambus Inc. fights 'FUD,'" *ENEWS*, December 4, 2000, p. 4.

[22] This appendix draws heavily on Rambus's S-1 filing, May, 1997; Chaplinsky, R.C., "Rambus – Company report," Hambrecht & Quist Inc., June 9, 1997; and Atkins, S., "SDRAM memory: DRAM and beyond," *Computer Technology Review*, Second Quarter, 1999.

[23] This appendix draws heavily on Atkins, S., "SDRAM memory: DRAM and beyond," *Computer Technology Review*, Second Quarter, 1999.