NEW DRAM TECHNOLOGIES

A Comprehensive Analysis of the New Architectures

SECOND EDITION



Steven A. Przybylski

MICRODESIGN RESOURCES

From the publisher of MICROPROCESSOR REPORT

Samsung Ex. 1097, p. Cover
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

New DRAM Technologies

A Comprehensive Analysis of the New Architectures

Second Edition

Steven A. Przybylski

with

Michelle M. Deatrick

MicroDesign Resources

Samsung Ex. 1097, p. i
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

Author:
Steven A. Przybylski
Consultant
3281 Lynn Oaks Dr.
San Jose, CA 95117
408.984.2719, FAX: 408.246.2412
sp@acm.org

With:
Michelle M. Deatrick

TK 7895
.M4 P7924
1996
Copy 1

Publisher:
MicroDesign Resources
874 Gravenstein Hwy. So.
Suite 14
Sebastopol, CA 95472
707.824.4004, FAX: 707.823.0504

Copyright © 1994,1996 Steven A. Przybylski

ISBN 1-885330-06-5

All Rights Reserved. No part of this report may be reproduced, stored in a retrieval system, or transmitted in any form or by any means without prior written permission.

MicroDesign Resources is a trademark of Ziff Communications Co.

This report contains and analyzes information from publicly available sources and from industry contacts. While we consider these sources reliable, we cannot guarantee the accuracy, whether past, present or future, of the information or analyses contained herein. Readers assume full responsibility for any use made of this report or its contents, in whole or in part.

Trademark names are used throughout this report. These trademarks are the property of their respective owners.

Please inform the author of any errors of fact or omission so that they may be corrected in future editions of the report.

Printed on recycled paper.

99-231315

Samsung Ex. 1097, p. ii
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

Executive Summary

The Dynamic RAM (DRAM) interface has changed very little in two decades. The growing inadequacy of this interface is driven by two long-term trends:

- **DRAM capacities have consistently quadrupled every three years.** The next-generation 64-Mbit devices, expected to be price per bit leaders by 1997 or 1998, will contain 8 Mbytes each. Consequently, by the end of the decade, most computer systems will be shipped with four or fewer DRAMs.
- **DRAM performance has not kept pace with processor performance.** DRAM access times have improved by only 7% per year while processor performance has grown at about 60% per year. While the available bandwidth per bit stored has declined continuously and dramatically, the demands placed on memory have increased.

These two trends result in a dilemma. In an era with strong motivations for both greater performance and fewer DRAMs per system, the existing DRAM interface is capable of improving either bandwidth or memory granularity (increment size), but not both simultaneously:

- **High-bandwidth memory systems cannot be made with only a few DRAMs.**
- **Increasingly, adequate performance only comes with excessive granularity.**

The need for change in the DRAM interface is pressing. The semiconductor industry is responding with a variety of new DRAM architectures that offer, in varying degrees, shorter access times, more bandwidth, and finer memory system granularity—all at a price. This report explores the need for change, contrasts the alternatives offered, and discusses the markets in which they are likely to prosper. Depending on the degree of change to the interface, alternative DRAMs are classified as conventional, evolutionary, or revolutionary:

- **Conventional choices** involve relatively minor changes to the generic interface: increased widths and/or extended data-out times.
  - **Extended Data-Out (EDO) DRAMs**, also called **Hyper-Page-Mode DRAMs**, incorporate a small change to the semantics of one control signal. This change allows the operating frequencies of memory systems to increase from approximately 25 MHz to 33–40 MHz. Even in conjunction with an increase in device width—from ×4 up to ×8 and ×16—this performance improvement is too small to solve the bandwidth problem in the long term. However, the EDO DRAMs' large number of suppliers and very small price premium will make wide EDO DRAMs a good fit for many applications during the near term.
- **Evolutionary architectures** provide substantially more bandwidth than do the generic DRAMs. Evolutionary DRAMs continue to rely on the same separate data bus and multiplexed address bus as the old generic DRAMs.
  - **Burst EDO (BEDO) DRAMs** allow pseudosynchronous bursts of four at frequencies up to 66 MHz. Burst EDO devices can be built on the same die as generic and EDO DRAMs. Consequently, Burst EDO DRAMs should command only a negligible price premium. However, since faster-than-average DRAMs are needed to run at 66 MHz, only a few DRAM manufacturers are committed to producing Burst EDO devices.
  - Enhanced Memory Systems' **Enhanced DRAMs (EDRAMs)** offer dramatically reduced access times and an improved data-write interface. The 25% cost premium is substantial. However, in small, tightly coupled systems, the cost premium can be offset by eliminating a secondary cache. Initial 4-Mbit offerings are in production though the supply will remain severely constrained until at least 1Q1996.

Continues next page

© 1996 Steven Przybylski

Samsung Ex. 1097, p. iii
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

- **Synchronous DRAMs (SDRAMs)**, developed and standardized by JEDEC, are well placed to become the next commodity part. Their high-speed synchronous interface and optional dual-bank architecture offer substantially higher bandwidths than are possible with generic DRAMs. SDRAMs are in production at the 4- and 16-Mbit levels. The initial adoption of SDRAMs was inhibited by a host of factors, including high price premiums, lack of standardized SIMMs, incompatibilities among the vendors' offerings, and confusion over the choice of a high-speed alternative to TTL levels. Though some of the incompatibilities persist, many of the initial barriers, especially high premiums, are disappearing.

- Mitsubishi's **Cache DRAMs (CDRAMs)** add a small block of SRAM and a synchronous interface to a generic DRAM. The result is a high-performance device that offers low latency and high sustained bandwidth. CDRAMs suffer from a high pin-count and a 15% cost premium. They are well suited to graphics applications and to smaller computer systems in which main memory and graphics memory are combined into a unified memory structure. As with EDRAMs, CDRAMs' combination of an effective row cache and low latency facilitates elimination of an SRAM-based secondary cache.

- **Revolutionary alternatives** offer a complete redesign of the DRAM interface. The redesign provides high bandwidth and minimum granularity. The revolutionary devices replace the separate address and data bus of the generic DRAM with a byte-wide interface that carries address, data, and signaling information.

  - Vertically or horizontally mounted **Rambus DRAMs (RDRAMs)** connect to a processor or memory controller through a physically short, 533 Mbytes/s bidirectional bus. Rambus Inc. has defined all aspects of the physical, electrical, and logical interface and has licensed a variety of DRAM and logic product vendors to produce Rambus-compatible devices. Complete specification of the Rambus Channel facilitates very high-speed systems with a minimum of design effort by system vendors. As of 1Q1996, 16- and 18-Mbit RDRAMs are available, with 8-Mbit devices sampling. RDRAM latencies for short transfers are somewhat higher than the latencies of some alternatives. However, RDRAMs facilitate small high-bandwidth memory systems; they are well suited for graphics and video controllers, unified main memories, and high-performance portable systems. Future 64-Mbit RDRAMs will offer an excellent combination of high bandwidth and moderate memory granularity to be effective main memory devices across a broad range of system types.

  - The **RamLink** interface specification, developed by the IEEE Computer Society P1594.6 Working Group, defines a protocol and ring topology for communication with DRAMs and other devices over a high-speed, byte-wide interface. A derivative called **SyncLink** migrates the RamLink protocol onto a dual bus physical implementation with 1 Gbyte/s peak bandwidth. SyncLink is moving through the standardization process with the support of a consortium of DRAM companies. Regardless, no RamLink- or SyncLink-based products are expected for at least two to three years.

Generic DRAMs are becoming increasingly inadequate. Each new alternative has specific strengths and weaknesses that make it a good fit for some DRAM applications and a poor fit for others. Toward the end of the decade, when the 64-Mbit generation becomes the cost leader, all significant market segments will have shifted away from generic fast-page-mode devices. The result will be a fragmented and unsettled marketplace with a few winners and many niche architectures.

This report uses a combination of new and old tools to examine the selection of DRAM architecture for main memory applications through the year 2000. In addition to looking at the characteristics of the various devices both individually and comparatively, this report provides accurate and in depth comparative analysis of the memory systems that these new DRAMs are capable of producing. A powerful new metric—the fill frequency—is used to identify which DRAMs are simply unacceptable in particular systems marketplaces.

Samsung Ex. 1097, p. iv
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

Acknowledgments

The authors would like to thank the many people who have directly and indirectly helped during the preparation of this report. They include:

Betty Prince of Memory Strategies International
Russ Tuck and Tom Blank of MasPar
Bob Proebsting and Bill Scharrenberg of Hyundai
Gene Cloud, Bob Fusco, and Sharron Bittick of Micron
Dave Mooring, Steve Ray, Billy Garrett, Richard Crisp, Julie Cates, and the other reviewers at Rambus
David Bondurant and Todd Oseth of Ramtron
Rajit Shah of Mosel-Vitelic
Satoru Kobayashi and Bart Ladd of NEC
Charles Hart, Sam Chen, Doug Gates, and Sherry Hill of Mitsubishi
Cormac O'Connell of MOSAID
Peter Cheng, D.G. Kim and Mark Ellsberry of Samsung
Mark Horowitz and Mehrdad Heshami of Stanford University
Hans Wiggers of Hewlett Packard
Jim Handy of Dataquest
Wolf-Dietrich Weber of HaL Computer Systems
Akhil Duggal of Alantec
Melanie Sanborn, Leslie Davidson, Ron Liskey, Bruce Koball, John Snell, Jeff Deutsch, and Linley Gwennap of MicroDesign Resources
Gail Anderson of Pine Cone Communications
Liz Longsworth of PC Computing
Ellen Clements

In addition to the above, the authors would like to thank the following people who provided information or helped review the second edition:

Craig Hampel, Julie Cates, Subodh Toprani, and Fred Ware of Rambus
Peter Cheng and Joe Kraus of Samsung
David Gustavson of SCIzzL
Joseph McIntyre and Ron Liskey of MicroDesign Resources
Jeff Mailloux, Brett Williams, and Sharron Bittick of Micron
Mark Levy of EDN
Sam Chen, Steve Forman, and Sherry Hill of Mitsubishi
Charles Furnweger and Tina Lim of NEC
David Bondurant of Enhanced Memory Systems
Michael Jahed of Toshiba
Betty Prince of Memory Strategies International
Jim Sogas of Hitachi

© 1996 Steven Przybylski

Samsung Ex. 1097, p. v
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

Samsung Ex. 1097, p. vi
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

About the Publisher

Founded in 1987, MicroDesign Resources is an analytical information services company focusing on the microprocessor and PC industries. MicroDesign Resources publishes *Microprocessor Report*, widely regarded as the leading independent newsletter in the field, and sponsors two conferences, the PC Tech Forum and the Microprocessor Forum. The PC Tech Forum will debut in May 1996 and will bring a sophisticated understanding of systems-level technologies together with a grasp of where the industry and the markets are headed. Now in its ninth year, the Microprocessor Forum is the most important annual gathering of microprocessor technologists in the industry. Michael Slater, President of MicroDesign Resources, is an international authority on microprocessor technology and trends. *Fortune* magazine has called Slater "the one expert to whom executives from Intel and the PowerPC consortium defer when asked to reconcile their wildly conflicting claims about the merits of their chips." Other analysis reports in MicroDesign Resources' Technical Libary are *Selecting a High-Performance Embedded Microprocessor* and *Buyer's Guide to DSP Processors*.

© 1996 Steven Przybylski

Samsung Ex. 1097, p. vii
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

© 1996 Steven Przybylski

Samsung Ex. 1097, p. viii
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

Table of Contents

Executive Summary iii

Acknowledgments v

About the Publisher vii

Table of Contents ix

List of Figures xv

List of Tables xix

1. Introduction 1
- 1.1 Outline 2
  - *Introductory and Tutorial Material* 2
  - *Individual Alternatives* 3
  - *Comparative and Market Analysis* 4
  - *Supporting Materials and Resources* 4
- 1.2 Changes Since the First Edition 5
  - *New Alternatives and Significant Events* 5
  - *Market Changes* 6
  - *Other Changes since the First Edition* 7
- 1.3 Which Chapters Should You Read? 7
  - *Engineers and Engineering Management* 7
  - *Marketers and Market Analysts* 8
  - *General Interest Readers* 8
- 1.4 Conventions in This Report 8

2. Industry Trends and Their Implications 11
- 2.1 New In The Second Edition 11
- 2.2 DRAM Trends 11
  - *DRAM Speed Trends* 13
  - *Technology and Economic Trends* 15
- 2.3 Computer System Trends 19
  - *Computer CPU Trends* 19
  - *Computer Memory System Trends* 21
  - *Computer Graphics Trends* 25
- 2.4 Combining DRAM and Systems Trends 28
- 2.5 Fill Frequency: A New DRAM Metric 31

© 1996 Steven Przybylski

Samsung Ex. 1097, p. ix
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

3. DRAM Basics 41
3.1 New In The Second Edition 41
3.2 Introduction 41
3.3 What Is a DRAM? 42
The Basics 42
DRAM Internals 48
3.4 DRAM Fabrication Technology 54
3.5 Advanced DRAM Concepts 60
Soft Errors 60
Supply Voltages 60
Testing and Redundancy in DRAMs 62
Speed-Enhancing Techniques 64
3.6 Video RAMs (VRAMs) 65

4. Memory System Design 69
4.1 New In The Second Edition 69
4.2 Basic Concepts 70
Single In-Line Memory Modules (SIMMs) 70
Memory System Organization: Granularity and Expandability 76
Bandwidth and Latency: Peak, Sustained, and Sustainable 81
Performance and Bandwidth Considerations 84
Parity and Error Correction Codes (ECC) 89
Refresh 91
4.3 High-Speed Design Issues 92
High-Speed Signals and Terminated Transmission Lines 92
Clock and Data Skew Issues 96
Issues Specific to Memory Systems 98
4.4 Memory Systems for PCs 100
x86 Machines 100
The Macintosh Family 103
Portables 104
4.5 Memory Systems for Workstations 104
4.6 Memory Systems for Large Machines 107
4.7 Graphics Memory Systems 110
4.8 Unified Memory Architecture (UMA) 117

5. New DRAM Architectures 129
5.1 New In The Second Edition 129
5.2 Potential Improvements to the Existing Interface 129
The Physical Interface 130
The Electrical Interface 136
The Pin-Level Interface 147
The Temporal Interface 148
The Logical Interface 148
5.3 Generic Narrow DRAMs 150
5.4 Conventional Approaches 151
5.5 Evolutionary Approaches 151
5.6 Revolutionary Approaches 153

Samsung Ex. 1097, p. x
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

6. Conventional Alternatives 157
6.1 New In The Second Edition 157
6.2 Wide Conventional DRAMs 158
6.3 Extended Data-Out DRAMs 164

7. Evolutionary Alternatives 171
7.1 New In The Second Edition 171
7.2 Burst EDO (BEDO) DRAMs 172
Physical, Electrical, Pin-Level Interface 172
Temporal and Logical Interface 173
System Design Issues 178
Burst EDO Fill Frequencies 180
History, Current Status, and Future Plans 181
Strengths and Weaknesses 182
7.3 Enhanced DRAMs (EDRAMs) 183
Physical, Electrical, and Pin-Level Interfaces 184
Temporal Interface 186
Logical Interface 189
EDRAM Fill Frequencies 198
Miscellanea 199
Future Developments 202
Strengths and Weaknesses 203
7.4 Synchronous DRAMs (SDRAMs) 204
Physical, Pin-Level, and Temporal Interfaces 204
Logical Interface 212
Electrical Interface 229
SDRAM Fill Frequency 230
SDRAM-Lite 233
History and Barriers to Acceptance 235
Prices, Costs, and Design Wins 238
Current Offerings 241
Strengths and Weaknesses 242
7.5 Cached DRAMs (CDRAMs) 243
Physical, Pin-Level, Temporal, and Electrical Interfaces 245
Logical Interface 248
The 16-Mbit CDRAM 252
CDRAM Fill Frequencies 253
Miscellanea 255
Strengths and Weaknesses 257

8. Revolutionary Alternatives 259
8.1 New In The Second Edition 259
8.2 Rambus DRAMs (RDRAMs) 260
Pin-Level and Physical Interfaces 261
Electrical Interface 267
Temporal Interface 271
Logical Interface 274

© 1996 Steven Przybylski

Samsung Ex. 1097, p. xi
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

Miscellanea 289
First Generation 4-Mbit RDRAMs 295
The 64-Mbit RDRAMs 296
RDRAM Fill Frequencies 299
Current Status and Business Issues 300
Strengths and Weaknesses 305
8.3 RamLink DRAMs 306
Pin-Level Interface 307
Physical, Electrical, and Temporal Interfaces 309
Logical Interface 310
Current Status and Future Developments 317
Strengths and Weaknesses 319
8.4 SyncLink DRAMs (SLDRAMs) 320
Bus Structure and Pin-Level Interface 322
Physical, Electrical, and Temporal Interfaces 324
Logical Interface 325
Changes from RamLink 332
SLDRAM Fill Frequency 333
Current Status and Business Issues 334
Strengths and Weaknesses 335

9. Chip-Level Comparison 337
9.1 New In The Second Edition 337
9.2 Pin-Level and Physical Interfaces 338
9.3 Logical Interface 339
9.4 Electrical and Temporal Interfaces 345
9.5 Peak Fill Frequencies 348
9.6 Power Consumption 349
9.7 Price and Supply 350

10. System-Level Comparison 353
10.1 New In The Second Edition 354
10.2 High-Performance Memory System 355
Memory Size and Granularity 366
Row-Cache Organizations 367
Access Latencies 372
Sustainable Bandwidths 377
Fill Frequencies 381
Miscellanea 383
10.3 PC Memory System 386
Memory Size and Granularity 389
Row-Cache Organizations 396
Access Latencies 397
Sustainable Bandwidths 405
Fill Frequencies 408
Unified Memory Architecture PCs 410
Miscellanea 411

Samsung Ex. 1097, p. xii
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

10.4 Memory Systems for Portable Systems 413
Memory Size and Board Area 415
Row-Cache Organization, Bandwidth, and Access Time 416
Power Considerations 419
10.5 Embedded Applications 425
Video Games and Other High-Performance Embedded Systems 425
Other Consumer Electronics and HDTV 427
10.6 Other DRAM Applications 428
Frame Buffer Memories 428
Mainframes, Supercomputers, and Massively Parallel Machines 430

11. Market Forces and Opportunities 433
11.1 New In The Second Edition 433
11.2 Market Segments and the DRAM Alternatives 434
DRAM and VRAM Market Segments 434
Strengths and Weaknesses of the DRAM Alternatives 438
Matching Market Segment Needs To DRAM Strengths 440
Changes at the 64-Mbit Density 450
Comparison of Peak Fill Frequencies 451
11.3 DRAM Costs and Pricing 454
Cost of Goods 454
DRAM Pricing 457
Expected Pricing for New Alternatives 463
11.4 Changing DRAM and VRAM Market Dynamics 468
11.5 The Coming Fragmentation of the DRAM Market 473
11.6 Barriers to Market Acceptance 475

A. List of Acronyms 479

B. Glossary 485

C. List of Vendors 507

D. References 511

E. Reprinted Data Sheets 517
E.1 A Generic 4M×4 Fast-Page-Mode DRAM Data Sheet 517
E.2 A 1M×16 Fast-Page-Mode DRAM Data Sheet 537
E.3 A 4M×4 Extended Data-Out DRAM Data Sheet 577
E.4 A 1M×16 Burst EDO DRAM Data Sheet 593
E.5 A 1M×4 Enhanced DRAM Data Sheet 613
E.6 A 4M×4, 2M×8, and 1M×16 Synchronous DRAM Data Sheet 637
E.7 A 1M×16 Cached DRAM Data Sheet 741
E.8 A 16-Mbit and 18-Mbit Rambus DRAM Data Sheet 809

Index 841

© 1996 Steven Przybylski

Samsung Ex. 1097, p. xiii
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

Samsung Ex. 1097, p. xiv
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

List of Figures

Figure 2-1 DRAM density versus time. 12
Figure 2-2 Row and column access times versus time. 14
Figure 2-3 Minimum feature size versus time. 16
Figure 2-4 Fab line capital cost versus time 17
Figure 2-5 Memory cost versus time. 17
Figure 2-6 CPU performance versus time 19
Figure 2-7 Cache miss ratio versus cache size. 21
Figure 2-8 Interaction effects that determine system memory size 22
Figure 2-9 Memory size and number of DRAMs per system versus time. 29
Figure 2-10 Memory size and number of DRAMs per system versus time to 1999. 30
Figure 2-11 Traditional and modified view of the DRAM performance gap 32
Figure 2-12 Long-term trends in DRAM and system fill frequencies 37

Figure 3-1 DRAM storage cell 43
Figure 3-2 Subdivision of the DRAM address 44
Figure 3-3 Basic DRAM timing diagram 46
Figure 3-4 DRAM internals 49
Figure 3-5 Internal voltage levels during a data read and refresh 51
Figure 3-6 Typical sense amplifier 52
Figure 3-7 Physical organization of a 4M×1 DRAM 55
Figure 3-8 Physical organization of a 512K×8 DRAM 56
Figure 3-9 A planar DRAM cell cross-section 58
Figure 3-10 A trench DRAM cell cross-section 58
Figure 3-11 A stacked DRAM cell cross-section 59
Figure 3-12 Soft error rate versus time. 61
Figure 3-13 Number of soft errors over 10 years of use versus time 61
Figure 3-14 Typical graphics controller block diagram 65
Figure 3-15 VRAM block diagram 66

Figure 4-1 Standard 30-pin SIMM physical diagram 71
Figure 4-2 Four-Mbyte byte-wide SIMM block diagram 72
Figure 4-3 2M×36 SIMM block diagram 73
Figure 4-4 Buffered SIMM block diagram 74
Figure 4-5 Typical low-end memory system block diagram 77
Figure 4-6 Ranks, banks, and sub-banks 79
Figure 4-7 Burst transfer timing diagram 81
Figure 4-8 Interleaved burst transfer timing diagram 84
Figure 4-9 Interleaved memory system block diagram 85

© 1996 Steven Przybylski

Samsung Ex. 1097, p. xv
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

Figure 4-10 Physical address fields 86
Figure 4-11 Block diagram of a typical Pentium® PC 101
Figure 4-12 Block diagram of a high-end workstation 106
Figure 4-13 Low-end graphics controller block diagram 112
Figure 4-14 Mid-range graphics controller block diagram 113
Figure 4-15 Traditional untiled frame buffer 114
Figure 4-16 A tiled frame buffer makes better use of DRAM row accesses 115
Figure 4-17 Traditional PC graphics architecture 117
Figure 4-18 Dual-supplier unified memory architecture (UMA) PC 120
Figure 4-19 Single-supplier unified memory architecture (UMA) PC 121

Figure 5-1 DRAM package types 132
Figure 5-2 Timing diagram showing excessive ringing 138
Figure 5-3 Gunning Transceiver Logic (GTL) driver circuit diagram 140
Figure 5-4 HSTL defines a variety of electrical environments 142
Figure 5-5 Stub-series termination mitigates the impact of stub reflections 144
Figure 5-6 Rambus electrical interface circuit diagram 146

Figure 6-1 Block diagram of a ×16 DRAM 159
Figure 6-2 Package types for ×16 DRAMs 160
Figure 6-3 Horizontally versus vertically organized memory systems 161
Figure 6-4 Wide DRAM fill frequencies 164
Figure 6-5 Extended Data-Out (EDO) DRAM timing diagram 165
Figure 6-6 EDO DRAM fill frequencies 168

Figure 7-1 Burst EDO read timing diagram 174
Figure 7-2 Burst EDO write timing diagram 174
Figure 7-3 Burst EDO fill frequencies 180
Figure 7-4 EDRAM packages 185
Figure 7-5 EDRAM timing diagram 188
Figure 7-6 EDRAM block diagram 191
Figure 7-7 EDRAM timing diagram with overlapped row-cache hit and precharge 193
Figure 7-8 EDRAM ×8 block diagram 196
Figure 7-9 EDRAM burst-mode timing diagram 197
Figure 7-10 EDRAM fill frequencies 198
Figure 7-11 SDRAM packages 205
Figure 7-12 SDRAM basic timing diagram 207
Figure 7-13 SDRAM timing makes at-speed operation difficult 209
Figure 7-14 SDRAM block diagram 213
Figure 7-15 SDRAM addressing across families 215
Figure 7-16 SDRAM command register format 216
Figure 7-17 SDRAM timing diagram 219
Figure 7-18 SDRAM timing diagram showing accesses to both banks 220

Samsung Ex. 1097, p. xvi
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

Figure 7-19 SDRAM timing diagram showing long alternating bursts 221
Figure 7-20 SDRAM state diagram 223
Figure 7-21 SDRAM fill frequencies 231
Figure 7-22 CDRAM block diagram 244
Figure 7-23 CDRAM package 245
Figure 7-24 CDRAM output modes 247
Figure 7-25 CDRAM timing diagram 250
Figure 7-26 CDRAM fill frequencies 254
Figure 7-27 CDRAM unified memory architecture block diagram 256

Figure 8-1 RDRAM package and pin descriptions 263
Figure 8-2 Board layout of a segment of the Rambus Channel 262
Figure 8-3 Populated Rambus Channel 266
Figure 8-4 Rambus Channel high-level block diagram 267
Figure 8-5 Rambus Channel circuit diagram 268
Figure 8-6 Rambus Channel detailed block diagram 271
Figure 8-7 RDRAM block diagram 279
Figure 8-8 Rambus Channel Request Packet format 277
Figure 8-9 Rambus Channel timing showing a row-cache hit 279
Figure 8-10 Rambus Channel timing showing a row-cache miss 280
Figure 8-11 Serial address packets facilitate random octet bursts 282
Figure 8-12 RDRAM address mapping 286
Figure 8-13 RAC block diagram 292
Figure 8-14 RDRAM fill frequencies 300
Figure 8-15 RamLink system block diagram 307
Figure 8-16 RamLink slave block diagram 308
Figure 8-17 SyncLink system block diagram 310
Figure 8-18 RamLink packet format 312
Figure 8-19 New SyncLink system block diagram 323
Figure 8-20 SyncLink pinout 324
Figure 8-21 SyncLink DRAM block diagram 326
Figure 8-22 SyncLink packet format 327

Figure 10-1 High-performance memory system block diagram 356
Figure 10-2 High-performance memory system latencies 373
Figure 10-3 High-performance memory system initial latencies 378
Figure 10-4 High-performance memory system sustainable bandwidths 379
Figure 10-5 High-performance memory system sustainable fill frequencies 382
Figure 10-6 PC memory system block diagram 387
Figure 10-7 PC memory system latencies 398
Figure 10-8 PC memory system initial latencies 402
Figure 10-9 PC memory system sustainable bandwidths 407
Figure 10-10 PC memory system sustainable fill frequencies 409

© 1996 Steven Przybylski

Samsung Ex. 1097, p. xvii
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

Figure 10-11 The effect of memory technology on PC performance 412

Figure 11-1 Market affinities 442
Figure 11-2 Comparison of DRAM capabilities and market requirements 469
Figure 11-3 DRAM price history 458
Figure 11-4 Historical generational crossover 469
Figure 11-5 Contemporary generational crossover 470
Figure 11-6 Future generational crossover 472
Figure 11-7 DRAM market fragmentation 475

Samsung Ex. 1097, p. xviii
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

List of Tables

Table 2-1	Graphics standard's display bandwidth requirements	25
Table 2-2	System bandwidth and granularity requirements	34
Table 2-3	DRAM bandwidth and size capabilities	35
Table 3-1	Primary DRAM timing parameters	48
Table 4-1	Linear and sub-block ordered burst transfers	88
Table 4-2	Typical frame buffer sizes and display bandwidths	111
Table 4-3	Characteristics of different UMA architectures	123
Table 5-1	DRAM package types	135
Table 5-2	Electrical signaling standards	147
Table 5-3	Interface changes in the new DRAM alternatives	154
Table 6-1	Comparison of EDO and fast-page-mode cycle times	166
Table 7-1	Comparison of Burst EDO, EDO, and fast-page-mode timing parameters	177
Table 7-2	Comparison of Burst EDO, EDO, and fast-page-mode in PCs	179
Table 7-3	EDRAM timing parameters	187
Table 7-4	EDRAM PC system response times	190
Table 7-5	SDRAM timing parameters	207
Table 7-6	Typical 66-MHz system timing parameters	210
Table 7-7	SDRAM access times	217
Table 7-8	SDRAM command encodings	222
Table 7-9	SDRAM-Lite feature set	234
Table 7-10	CDRAM timing parameters	246
Table 7-11	CDRAM operations	249
Table 8-1	RDRAM configuration registers	276
Table 8-2	RDRAMs' power consumption	287
Table 8-3	RamLink signal descriptions	308
Table 8-4	RamLink command field encoding	314
Table 8-5	SyncLink command field encoding	327
Table 8-6	SyncLink byte ordering	330
Table 9-1	Chip-level comparison	340
Table 10-1	High-performance memory system comparison	360

© 1996 Steven Przybylski

Samsung Ex. 1097, p. xix
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

Table 10-2 PC memory system comparison 360
Table 10-3 PC memory system access time comparison 360
Table 10-4 Portable system comparison..... 360

Table 11-1 DRAM market requirements. 435
Table 11-2 DRAM architecture's strengths and weaknesses..... 439

Samsung Ex. 1097, p. xx
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

1 Introduction

For two decades, one of the few constants in the semiconductor and computer industries has been the standard dynamic RAM (DRAM) interface. This interface—the physical, electrical, and logical definition of the generic commodity DRAM—has played an important and lasting role in shaping and invigorating the semiconductor and electronic industries.

After 20 years, the interface that served eight successive generations of DRAMs—from 4 Kbits to 16 Mbits—is increasingly inadequate. In another generation, the old standby asynchronous interface in the ×1 and ×4 organizations will be essentially unused. These organizations will be unavailable at the 256-Mbit level.

In response to the need for change, the semiconductor industry is offering a variety of new interfaces. Some incorporate only limited changes from the existing generic interface; others offer a more radical redesign. This report is about the new interfaces: their strengths and weaknesses; their proponents and availability; their market affinities and longer term prospects.

Intended audiences

Three groups of people need to know about the coming changes in the DRAM industry:

- Engineers and their managers in system companies who need to choose a DRAM architecture for the short and medium term.
- Marketing personnel in system and semiconductor companies who need a broad-brush understanding of the

Samsung Ex. 1097, p. 1
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

significant changes taking place in this $32 billion industry sector.

- Other technical and nontechnical individuals interested in learning more about DRAMs, memory systems, or the forces at work in this market.

This report is designed to meet the needs of all three groups. Whatever your needs, there is a lot here for you.

1.1 Outline

The 11 chapters of this report fall into three groups: introductory and tutorial material; discussion of the individual DRAM alternatives; comparative technical and market analysis.

1.1.1 Introductory and Tutorial Material

The first four chapters explain why changes in the DRAM market are happening now and provide background material on DRAM internals and the use of DRAMs in memory systems.

Chapter 1. Introduction
Lays out the report's structure and directs readers to the sections relevant to their needs.

Chapter 2. Industry Trends and Their Implications
Brings together 20 years of DRAM and computer system development into a set of easily understood trends. In particular, a discussion of the differences between system requirements and DRAM capabilities clearly and dramatically explains the growing inadequacy of the generic narrow DRAM interface and many of the proposed alternatives.

Chapter 3. DRAM Basics
Describes how DRAMs work, how they are built, and how the interface gets information in and out. This discussion progresses from a simple overview through the goriest of technical details.

Chapter 4. Memory System Design
Focuses on how DRAMs are put together to form memory systems. Explains basic and advanced concepts, then illustrates the concepts with examples of memory systems for personal computers, workstations, and

Samsung Ex. 1097, p. 2
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

mainframes. Section 4.8 is devoted to the rationale, strengths, and weaknesses of Unified Memory Architectures (UMAs).

1.1.2 Individual Alternatives

The middle four chapters describe all of the alternative DRAM architectures in detail.

Chapter 5. New DRAM Architectures
Sets the framework for a detailed discussion of DRAMs by breaking down the interface into its physical, electrical, pin-level, temporal, and logical components. Describes the most common options for each component. Introduces the alternative architectures and—depending on the degree of change from the traditional generic interface—classifies the architectures as conventional, evolutionary, or revolutionary.

Chapter 6. Conventional Alternatives
Describes the conventional alternatives. These DRAMs incorporate only relatively minor changes to one or two components of the entire standard interface. Conventional options that are discussed in detail are **wide DRAMs** and **Extended Data-Out (EDO) DRAMs**.

Chapter 7. Evolutionary Alternatives
Discusses the evolutionary alternatives. Evolutionary DRAMs embody changes to the DRAM interface that are more substantive than those of the conventional alternatives. However, the DRAMs still preserve the separate data and multiplexed address buses. The evolutionary alternatives are **Burst EDO (BEDO) DRAMs, Enhanced DRAMs (EDRAMs)**, **Synchronous DRAMs (SDRAMs)**, and **Cached DRAMs (CDRAMs)**.

Chapter 8. Revolutionary Alternatives
Describes the revolutionary alternatives, which offer a complete redesign of the DRAM interface. Both **Rambus DRAMs (RDRAMs)** and **RamLink/SyncLink DRAMs** replace the data and multiplexed address buses with one or two high-speed byte-wide interfaces that carry all address, data, and signaling information.

© 1996 Steven Przybylski

1.1.3 Comparative and Market Analysis

The final three chapters compare the various alternatives and match them up with market segments. The comparison is made at two levels: 1) DRAM-to-DRAM and 2) memory system-to-memory system.

Chapter 9. Chip-Level Comparison
Systematically reviews each alternative from the perspectives of the five components of the entire interface: physical, electrical, pin-level, temporal, and logical.

Chapter 10. System-Level Comparison
Analyzes the memory systems that can be built with each DRAM architecture. Because the memory systems of various computer classes differ so greatly, the analysis is repeated four times: for workstations, personal computers, portable systems, and high-performance embedded systems.

Chapter 11. Market Forces and Opportunities
Moves beyond the report's primary focus on the new DRAMs' technical attributes. Deals with market forces and discusses which alternatives will likely succeed in each of ten market segments. Highlights the major interactions that drive DRAM pricing. Forecasts the coming fragmentation of the DRAM marketplace.

1.1.4 Supporting Materials and Resources

Appendix A. List of Acronyms
Lists the 106 acronyms used in the report.

Appendix B. Glossary
Defines 125 commonly used terms.

Appendix C. List of DRAM Vendors
Lists the names, addresses, and phone numbers of 18 DRAM manufacturers.

Appendix D. References
Comments on books, articles, conferences, and patents relating to DRAMs, the new DRAM architectures, and computer and memory systems.

Samsung Ex. 1097, p. 4
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

Appendix E. Reprinted Data Sheets
Reproduces eight data sheets covering the spectrum from generic through revolutionary DRAMs.

Index

1.2 Changes Since the First Edition

A great deal happened between May 1994, when the first edition of this volume went to press, and January 1996, when the second edition was finalized. In addition to reporting on those events and changes, this second edition adds a number of important analytical tools that clarify the differences among the alternative DRAMs. Major improvements in this edition include updates relating to changes to individual alternatives and to marketplace and systems requirements. The second edition also implements important improvements in the analysis and presentation of data for all alternatives. The following subsections describe the changes in more detail.

1.2.1 New Alternatives and Significant Events

Burst EDO DRAMs:
Burst EDO DRAMs, first publicly announced in January 1995, are covered in depth in Section 7.2 and contrasted to the other alternatives in Chapters 9 through 11. A Burst EDO data sheet is provided in Appendix E.4.

SDRAMs:
Changes to the SDRAM world since May 1994 include:

- Standardization of both 200- and 168-pin DIMMs
- Progress towards defining a workable high-speed electrical interface
- Reduction but not elimination of the incompatibilities across vendors
- The proposal of the reduced-functionality SDRAM-Lite
- Unveiling of 64-Mbit SDRAM architectures by several vendors

© 1996 Steven Przybylski

- The success of 4-Mbit SDRAMs in the video game and GUI accelerator markets.

Rambus:
In September 1995, Rambus, Inc. began to disclose its 64-Mbit generation. Significant changes to both the protocol and DRAM architecture substantially address most of the criticisms leveled at previous generations. Also, with the first shipment of Silicon Graphics Inc.'s Indigo2 IMPAC workstation, RDRAMs are now in end-user products. Rambus is covered in depth in Section 8.2.

SyncLink:
In late 1994, portions of the DRAM vendor community noticed the existence of RamLink and took an interest in it. Since then, a derivative of RamLink, called SyncLink, has been partially defined. An industry consortium has formed to commercialize and standardize a specification in advance of SyncLink DRAMs (SLDRAMs) devices coming to market. SyncLink is described in Section 8.4.

64-Mbit devices:
Though cost-effective 64-Mbit devices are still at least 18 months off, samples of this next generation are available in several architectures. The 64-Mbit devices that have been defined are documented in their respective sections.

1.2.2 Market Changes

Unified Memory Architectures:
The analysis of unified memory architectures—those in which the frame buffer and main memory coexist in the same memory array—awkwardly straddles the traditional division between main memory and graphics applications of DRAMs. UMA systems are covered in both chapters 4 and 10.

High-Performance Embedded Systems:
Chapter 10 now contains significant analysis focussed on memory system design for high-performance embedded systems, including set-top boxes, video games, and DSP systems. These systems are all characterized by small memory sizes as well as high peak and sustained bandwidth requirements.

Samsung Ex. 1097, p. 6
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

1.2.3 Other Changes since the First Edition

Fill Frequency—A Powerful New Metric:
The fill frequency, defined and explored in Chapter 2, is now used throughout the volume to clarify the important performance differences between various alternatives and to highlight instantly which devices are appropriate candidates for particular applications.

General Update:
All information regarding the status of both individual devices and general and specific industry trends has been updated to take into account the most recent announcements.

1.3 Which Chapters Should You Read?

1.3.1 Engineers and Engineering Management

The most advanced sections of the tutorial material in the first several chapters will benefit system engineers and their managers. However, the primary benefit to this group comes from the later chapters.

Chapters 5 through 8 provide the detailed information needed to narrow the search for the most appropriate DRAM architecture. Chapter 5, New DRAM Architectures, introduces much of the terminology and structure for these and the final comparative chapters.

Chapter 9, Chip-Level Comparison, provides a quick overview of the alternatives. Chapter 10, System-Level Comparison, is most valuable for comparing the alternatives in the context of specific systems. Chapter 11, Market Forces and Opportunities, discusses which DRAMs are likely to gain market acceptance and achieve the high volumes and multiple sources of supply needed for aggressive price competition. Those pursuing a detailed study of the subject are likely to find useful information in both the annotated list of references in Appendix D and the data sheets reproduced in Appendix E.

© 1996 Steven Przybylski

1.3.2 Marketers and Market Analysts

The first four chapters present important technical background with an emphasis on accessibility. The discussion of market and technology trends in Chapter 2 offers valuable insight into why the DRAM marketplace is changing now and the degree to which it will change over the coming decade. The tutorial chapters on DRAMs and memory systems (Chapters 3 and 4) will satisfy professional and personal curiosity regarding the manufacture and use of DRAMs. The latter half of Chapter 5, New DRAM Architectures, and all of Chapter 9, Chip-Level Comparison, discuss and contrast the old and new DRAM interfaces. Finally, Chapter 11, Market Forces and Opportunities, explains the dynamics of the DRAM marketplace and provides insight into which alternatives are likely to succeed in which market segments. Marketers and market analysts will find the list of acronyms (Appendix A), the glossary (Appendix B), and the annotated references (Appendix D) especially useful.

1.3.3 General Interest Readers

The sections of interest to the general reader vary significantly depending on the reader's technical expertise and degree of interest. The first few sections of both Chapters 3 and 4 introduce important terms and concepts. Chapter 2 discusses the trends that will shape the future DRAM marketplace. Short overviews of current and future DRAM organizations are given in the latter half of Chapter 5. Chapter 11 summarizes the strengths and weaknesses of the DRAM types and discusses which market segments each DRAM is best suited to. For those readers unfamiliar with industry jargon, the list of acronyms in Appendix A and the glossary in Appendix B will prove handy.

1.4 Conventions in This Report

K, M, and G

In standard scientific nomenclature, the metric modifiers k, M, and G to refer to factors of 1,000, 1,000,000, and 1,000,000,000, respectively. As a matter of convenience, computer engineers have adopted K as the symbol for a factor of 1,024 (2^{10}). Unfortunately, M and G are also used for factors of 1,048,576 (2^{20}) and 1,073,741,824 (2^{30}), respectively. This can lead to confusion in circumstances where it is unclear whether an M refers to a factor of 1,000,000 or 1,048,576. In this report, M and G refer to

Samsung Ex. 1097, p. 8
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

the binary multiplier—that is, a multiplicative factor of an integral power of 2—when dealing with DRAMs or units of information (bits and bytes). However, their metric incarnations are used in the rare cases when they are combined with units of length, area, other physical units, or cost.

Glossary terms in bold

Where glossary terms are used for the first time, or where the reader might first encounter them, they are highlighted in bold.

Signal and register names in helvetica

In the text, tables, and figures, references to specific signals, buses, and registers are set in helvetica. An overbar is used to indicate a signal or bus that is active low. Where appropriate for clarity, bus widths are indicated as a suffix: Data0..31 is a 32-bit bus. Rarely, a following * is used in a figure to indicate an active low signal or bus.

Buses and signals in figures

In block diagrams, a thin line is used to indicate a single signal; a thick line is used for buses. Bus widths are either explicitly indicated, inferred from a signal name, or indeterminate. A solid line of medium width indicates a chip or module boundary; the interior of the chip or module is shaded.

Die and dice

Contrary to colloquial speech, the dictionary plural of *die* is *dice,* not *dies* or *die*. All hail Webster.

Dimensions

In physical diagrams of DRAM packages and SIMMs, dimensions are given in inches; dimensions in millimeters are given in parentheses.

Dates

Nonspecific dates are often abbreviated: 1Q1992 refers to the first quarter of the 1992 calendar year; 3Q1996 refers to the third quarter of 1996; and 1H1998 refers to the first half of 1998.

Bits and bytes

Bits and *bytes* are units of information, just as *seconds* are units of time. Normally, bits and bytes are abbreviated *b* and *B,* respectively. However, because both units are used heavily in this report, for clarity's sake the abbreviations are not used. The usage of bits and bytes differs from the usage of meters and seconds in that engineers often refer to specific bits and bytes. In these cases, the terms are typically the subject or object of a sentence and are not prefaced by a number.

Capitalization and hyphenation

Among those who care, the conventions for capitalization and hyphenation are hotly debated. In regard to capitalization, this report adopts a relatively aggressive style. To minimize

© 1996 Steven Przybylski

ambiguity, technical phrases are in title (headline) style, especially when they refer to specific instances or entities. For example, many of the new DRAMs are synchronous, but "Synchronous DRAM" refers to a particular DRAM architecture. Similarly, "Enhanced DRAM" refers to a specific DRAM architecture, rather than to all of the alternatives that are enhanced.

Writers and editors alike tend to aggressively hyphenate the long compound phrases that are common in technical writing. In fields such as this, both technology and language are changing so rapidly that even the most recent dictionaries provide little guidance. As a result, authors necessarily, though often unconsciously and inconsistently, make their own decisions with regard to hyphenation. In this report, the guiding principles for hyphenation are clarity, consistency, and adherence to the formal rules of hyphenation. Where there is latitude, the tendency is to not hyphenate.

Samsung Ex. 1097, p. 10
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

2 Industry Trends and Their Implications

2.1 New In The Second Edition

Changes to this chapter from the first edition include updates to all of the trend charts and related commentary to reflect the events of 1994 and 1995.

Section 2.5 introduces a new metric, the fill frequency. The fill frequency measure is used throughout the second edition to identify the application domains for which each new alternative is appropriate. In addition, long-term trends in the fill frequency, illustrated in Figure 2-12, graphically and dramatically illustrate the powerful forces driving the proliferation of new DRAM architectures.

2.2 DRAM Trends

For the past two decades, DRAMs have consistently quadrupled in size every three years (see Figure 2-1). The current 16-Mbit devices store 4,096 times more bits than the 4-Kbit parts introduced in 1974.[1] The cost of one megabyte of memory has declined from $8,000 in 1975 to as little as $23 in early 1993.[2] This phenomenal growth rate of 60% per year in memory

1. Although 1-Kbit parts were introduced in the early 1970s, the one-transistor storage cell became the industry standard for dynamic RAMs with the 4-Kbit generation.

2. B. Prince, *Semiconductor Memories: A Handbook of Design, Manufacture and Application*, John Wiley, 1991; and *Microsystems Insider*, Volume I, Number 1, October 1993.

Samsung Ex. 1097, p. 11
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

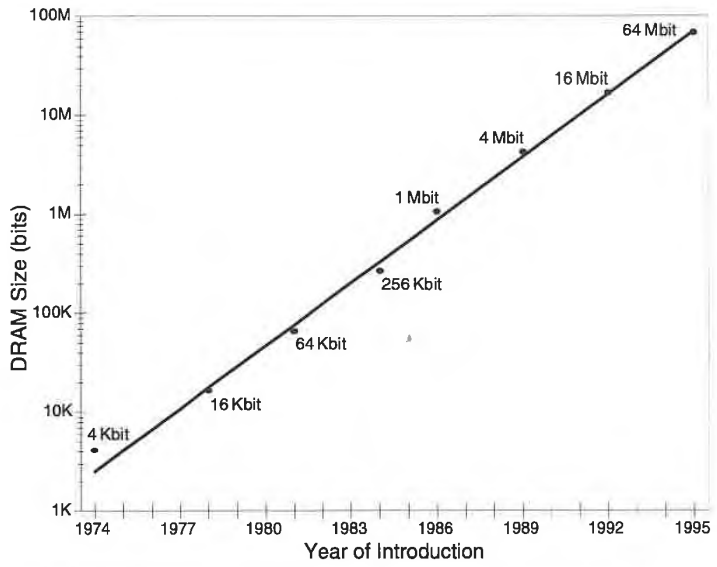


Figure 2-1 For 20 years, a new DRAM generation has become available every 3 years.

capacity per chip, coupled with a 30% yearly price reduction over a 20-year period, is unmatched in the history of commerce. The tremendous improvements in integrated circuit (IC) fabrication technology that made memory so cheap have facilitated the propagation of electronics into all aspects of our lives.

DRAM densities will continue increasing

DRAM densities are likely to continue increasing until at least the end of this decade. The next two generations—the 64- and 256-Mbit devices—are progressing well and on schedule. The first 64-Mbit devices began sampling in early 1995 and should be offered at competitive per-bit price levels by 1998. The 256-Mbit devices, which will contain 32 Mbytes per DRAM, should be widely sampled by 1997.

Despite the incredible growth in density, the basic **interface** by which the DRAMs communicate with the rest of the system has changed very little over the decades. Not surprisingly, the old interface has become inadequate for a significant and growing group of applications. In the coming two generations, this group will mushroom to include some of the largest DRAM market segments. Consequently, the DRAM industry has proposed a variety of new interfaces that address the growing demands on the DRAM interface. This chapter discusses the reasons for the

Samsung Ex. 1097, p. 12
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

standard DRAM interface's inadequacy. This discussion provides important background for examining the interface's ability to meet the coming challenges.

This section (Section 2.2) explores how DRAMs have improved over the past 20 years in their ability to store and supply information. Section 2.3 comments on changes in the demands placed on DRAMs by processors. Section 2.4 pulls the supply and demand aspects together. And Section 2.5 introduces a metric for the comparison of DRAM capabilities and memory system requirements. The fill frequency—the ratio of a DRAM's bandwidth to its size—is used throughout this report to identify the applications for which individual DRAM architectures are appropriate.

2.2.1 DRAM Speed Trends

The data stored in DRAMs is organized in a two-dimensional array of rows and columns. To access a specific bit of information, the memory controller first performs a **row access** to make the entire **row** containing the desired data available. Once the row access is complete, any of the data in the accessed row can be retrieved via a **column access**. This structure creates a natural internal memory hierarchy: data in the currently accessible (i.e., most recently accessed) row is more quickly readable than data in the other rows. When both a row access and a column access are needed to fetch a piece of data, the elapsed time is called the **row access time**. If only a column access is needed, the elapsed time is called the **column access time**.

Figure 2-2 shows steady decreases over the past two decades in both the row and column access times. However, the declines have not matched the rate of increase in DRAM capacity. In fact, over that period the row access time has declined at a relatively modest 7% per year, while the column access time has declined roughly 25% per year. In recent years, column access times and row access times have declined at roughly the same rate.

The maximum **bandwidth** a DRAM can provide is related to its **column cycle time**: the time from the start of one column access to the start of the next. As column access times have declined, so have column cycle times. The total bandwidth per DRAM data pin has consequently increased at about 25% per year. However, since the number of bits inside each DRAM has

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 13
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

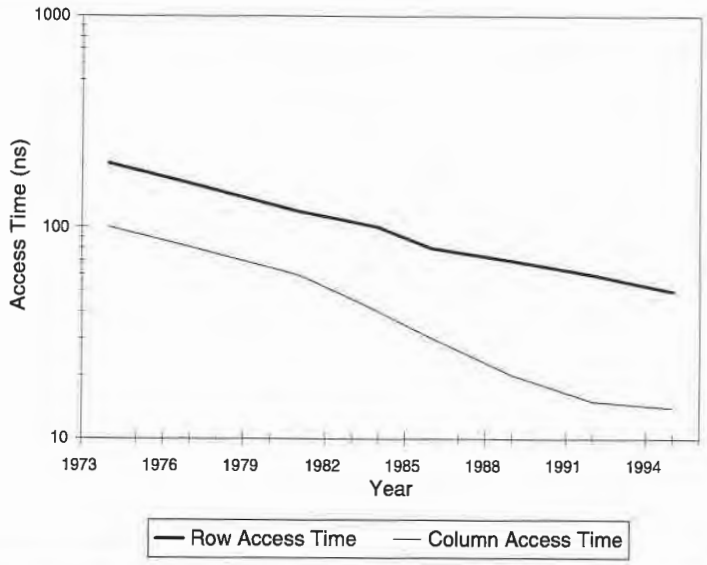


Figure 2-2 Since 1973, row access times have declined at roughly 7% per year, while column access times have declined at close to 25% per year.

been increasing continuously, and at a higher rate than the peak bandwidth, the available bandwidth per bit of memory has declined consistently, also at about 25% per year.[3]

Row and column access times attainable in real memory systems are often substantially lower than the chip-level access times found in DRAM data sheets. The difference between system-level and chip-level access times constitutes the temporal overhead in the system. This overhead stems from many causes, including:

- The consequences of operating many DRAMs in parallel
- The contribution of nonmemory components to system cycle time
- The physical size of the memory system
- Timing margin inserted to ensure reliable operation across a variety of operating conditions

3. Based on a constant DRAM width, the available bandwidth per bit of fast-page-mode memory has been declining at 25% per year. Assuming that the width of the most popular DRAM doubles every two to three generations (six to nine years), the effective rate of decline in the available bandwidth per bit has been between 15 and 20% per year.

Samsung Ex. 1097, p. 14
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

The large gap between minimum DRAM access time and actual cycle time in the system has meant that over the last 20 years system-level access times have declined even more slowly than chip-level times.

Further decreases in system-level column access times with the current fast-page-mode interface will be limited by the external board-level signal paths rather than by DRAM capabilities. The **TTL** voltage levels uniformly used for communication between ICs were standardized over 30 years ago. Though TTL levels are more than adequate at 1–10 MHz, the standard is increasingly difficult to work with above 50 MHz; it is essentially untenable above 100 MHz. At high frequencies, the inevitable temporal overhead makes it impossible to take full advantage of devices with short access times. New signaling standards must be used to benefit fully from ongoing improvements in semiconductor technology.

2.2.2 Technology and Economic Trends

Unparalleled improvements in basic fabrication technology have driven the tremendous sustained increase in DRAM capacity and speed. Figure 2-3 shows the **minimum feature size** for each DRAM generation at the time of its introduction. The figure shows an annual rate of decline in feature size in the 11–18% range. This rate of decline has diminished as process engineers have pushed feature sizes well below 1 μm; the rate of decline is likely to slow even further as the limits of photolithography are tested.

Because the number of transistors or storage cells per square millimeter is proportional to the square of the minimum feature sizes, the number of devices in a constant-sized integrated circuit has grown at 28–45% per annum for the last two decades. The difference between this growth rate and the 60-odd percent growth rate in DRAM density is accounted for by steady growth in the size of an economically manufacturable die. The 4-Kbit DRAMs were on the order of 10 to 15 mm^2, while 16-Mbit devices are between 115 and 140 mm^2.

Physical plant investments skyrocket

Unfortunately, the improvements in fabrication technology have not been without consequences. The cost of a DRAM fabrication facility has grown dramatically as well. Figure 2-4 shows that the cost of a fab line grew at an average rate of 32% per year over the past 20 years. Currently, building and outfitting a new

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 15
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

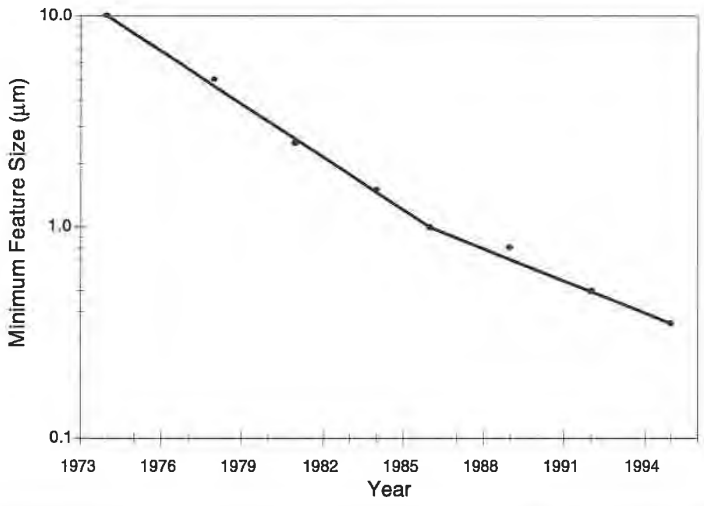


Figure 2-3 The driving force behind the growth in DRAM sizes and the improvements in speed has been the constant improvement in minimum feature size over the years.

manufacturing facility costs more than $1 billion.[4] In the past, growing demand for DRAMs has justified and offset this growing cost. Despite the continued decline in the cost of memory shown in Figure 2-5,[5] revenues from each DRAM generation have grown at roughly the same rate as capital expenditures. These roughly comparable growth rates result both from increases in unit volume and, more particularly, from the quadrupling in the number of bits per device every three years. The net result is that even companies managing to maintain market share have ratios of revenue to capital investment that are either flat or slowly declining.

Many factors influence the pricing of DRAMs over the short, medium, and long term. Over the short and medium term, the price of memory is very volatile. This is because 1) DRAMs from various manufacturers are interchangeable, 2) no substitute exists for DRAMs, and 3) DRAM vendors require a long lead

4. B. Prince, *Semiconductor Memories: A Handbook of Design, Manufacture and Application*, John Wiley, 1991; and J. Hutcheson (VLSI Research), "A New Way to Gauge the Real Cost of Fab," *EE Times*, December 6, 1993, p. 24.

5. Figure 2-5 shows the long-term trend evident in several public domain sources of DRAM prices. Chapter 11 discusses DRAM pricing in greater detail. In *Semiconductor Memories* (John Wiley, 1991), Betty Prince offers an excellent discussion of the factors influencing the cost of memory. More comprehensive and accurate price histories and projections are available for a fee from market research firms such as Dataquest, ICE, and In-Stat.

Samsung Ex. 1097, p. 16
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

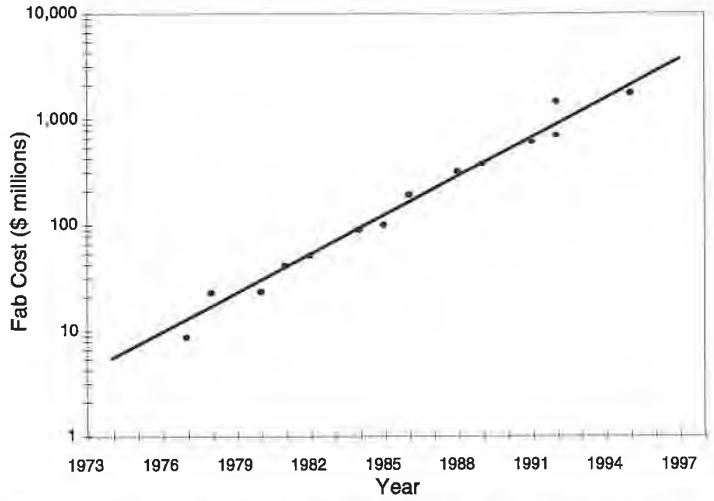


Figure 2-4 The cost of building integrated circuit fab lines has grown approximately 30% per year for almost 20 years. (Sources: B. Prince, *Semiconductor Memories*, John Wiley, 1991; H. Komiya, *Future Economic and Technologic Prospects for VLSI*, 1993 ISSCC; and J. Hutcheson (VLSI Research), "A New Way to Gauge the Real Cost of a Fab," *EE Times*, December 6, 1993, p. 24).

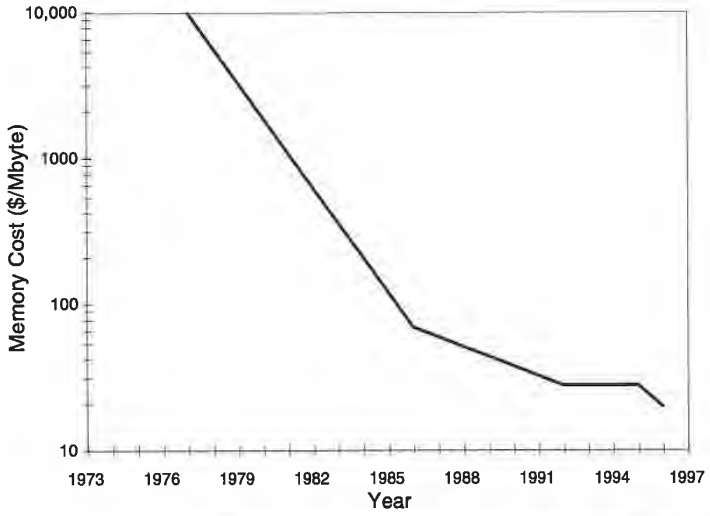


Figure 2-5 Until the mid 1980s, the price of memory declined about 40% per year. Between 1992 and 1995, DRAM prices did not decline at all. Due to an extended shortage, over much of that period, the price of memory was substantially above its 1993 low.

© 1996 Steven Przybylski

time and a great deal of capital to increase production significantly. Over the long term, the persistent growth in demand for memory (largely facilitated by ongoing declines in memory prices) and the steady rate of improvement in manufacturing technology have produced the dramatic long-term decline in prices shown in Figure 2-5. The changed rate of decline in the late 1980s coincides with the introduction of international trade politics as a new factor influencing pricing rather than with a major shift in the economics of supply or demand.

Prolonged shortage in early 1990s

Since the early 1990s, the PC industry has surged dramatically as PCs became ubiquitous in both business and home. This pervasiveness, combined with a steady growth in the average main memory size (see Section 2.3.2), has meant that the demand for DRAMs has exceeded the supply for several years. This persistent shortage has contributed to a reduction and even a temporary reversal of the long-term decline in the price of memory.

However, heavy new investment in DRAM manufacturing capacity worldwide resulted in substantial amounts of new memory capacity coming on line in 1995 and 1996. The increased supply combined with a slackening in the rate of growth in the demand for memory to produce a significant reduction in memory prices during 4Q1995 and 1Q1996. Over this period, the spot price for 4-Mbit DRAMs dropped by close to 50%—the first substantial reduction in the cost of memory for several years. Whether this dramatic decline in pricing is a temporary blip or a return to a traditional rate of decline in the cost of memory will depend primarily on the PC marketplace's growth rate relative to the planned growth rate in the supply of DRAMs.

In the longer term, as the industry looks forward to the 256-Mbit and 1-Gbit generation, concern is growing that economic factors—particularly exponential growth in the cost of IC factories and slackening demand for ever larger DRAMs—will dampen growth in DRAM densities before any insurmountable physical limits have been reached. A significant part of the overall challenge facing semiconductor vendors is to find the DRAM architecture and interface that will make single-chip 32- and 128-Mbyte DRAMs into salable products. This report explores that challenge as well as the current confusion at the 16-Mbit level.

A more detailed discussion of DRAM pricing, product life cycles, and overall industry changes is presented in Chapter 11.

Samsung Ex. 1097, p. 18
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

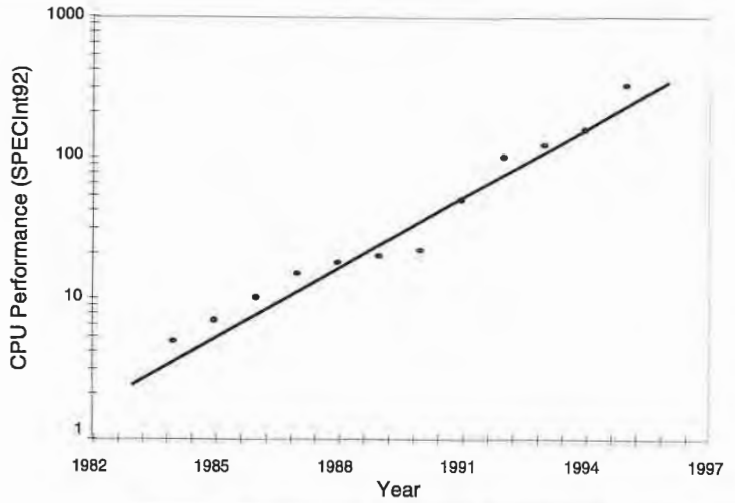


Figure 2-6 CPU performance has climbed at a rate of approximately 50% per year for over the past decade. (Source: D. Patterson and J. Hennessy, *Computer Architecture: A Quantitative Approach*, Second Edition, Morgan Kaufman, 1995).

2.3 Computer System Trends

2.3.1 Computer CPU Trends

The long-term improvements in semiconductor fabrication technology that have caused growth in DRAM sizes have similarly impacted processors. In the early 1970s, the first single-chip microprocessors contained around 3,000 transistors and were able to execute fewer than 500,000 instructions per second. In 1993, the largest microprocessors have on the order of 3 million transistors and are capable of peak execution rates up to 400 million instructions per second. Since the mid 1980s and the advent of RISC processor design techniques, the processing power of microprocessors has grown at up to 50% per year, thus doubling every 18 to 24 months. This trend is shown in Figure 2-6, which shows the reported performance of various workstations, as measured by the now-defunct SPECInt92 benchmark suite.

Caches increasingly important

Increasing demands on the computer system's main memory have accompanied the increase in processor execution speed. To a large extent, the increased use of caches has caused the

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 19
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

growth rate in the number of **references**—that is, requests for data—that reach main memory to be lower than the growth rate in the number of references generated by the processors. **Caches** are small, fast memories that are physically and logically close to the processor. The ability of caches to satisfy a large proportion of the processor's requests for information decreases the frequency of main memory accesses.

In the early 1980s, only the larger minicomputers and mainframes contained caches. With the advent of the Motorola M68020 and the Intel i486 with their on-chip caches, caches became common in personal computers. Today, most desktop PCs and workstations have external SRAM-based secondary caches in addition to the primary caches built onto the microprocessor ICs. These internal primary caches are typically between 8 and 32 Kbytes. In PCs, external caches are often as large as 512 Kbytes; in workstations, external caches are frequently as large as several megabytes.

Figure 2-7 illustrates the effectiveness of caches at reducing the number of references reaching main memory. The **cache miss ratio** is the proportion of all CPU memory references not satisfied in the cache that must be referred to main memory. On average, each doubling of the cache size reduces the number of main memory accesses by about 30% when executing general purpose, integer-intensive code.

Cache sizes continue to increase

Since the introduction of caches into PCs and workstations about a decade ago, average cache sizes have increased at about 50% per year. This increase has been driven by the performance increases facilitated by larger caches and by a cost per bit reduction in SRAMs similar to that for DRAMs. The net effect of this trend and the doubling in processor performance every 18 months is that the number of processor-generated main memory references has expanded at roughly 30% per year. I/O references have increased dramatically with larger program and problem sizes. In particular, the use of virtual memory and networks has greatly expanded the amount of I/O-to-memory traffic in PCs. However, the growth in I/O references to memory cannot be moderated by increased cache sizes.

As caches become larger, the number of misses declines until it hits a minimum value that varies from application to application. Well before the minimum is reached, though, the law of diminishing returns kicks in. Once the cache size has reached an intermediate value, the performance lost to cache misses is small, so further reductions in the number of misses yield only

Samsung Ex. 1097, p. 20
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

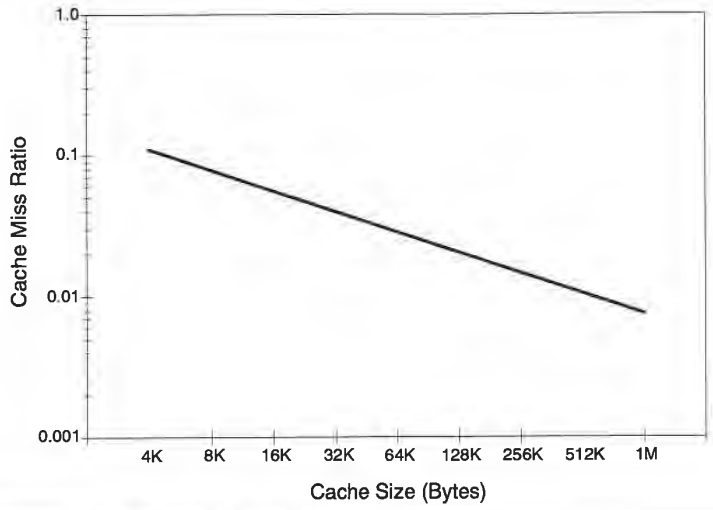


Figure 2-7 Averaged over many applications, each doubling of cache size reduces the number of memory references reaching main memory by approximately 30%.

small improvements in performance. In fact, with very large caches, system-level performance can decline as cache sizes increase if the larger cache has a slower response time to cache hits. But from a marketing perspective, bigger is almost always better. From a technical or cost-performance perspective, a maximum desirable cache size definitely exists.[6]

2.3.2 Computer Memory System Trends

Another important consequence of falling memory prices has been the increasing average memory size of PCs and workstations. As memory became cheaper, users could afford more memory. The larger and more sophisticated PC applications and operating environments became salable. Soon the needs of the new applications were driving end users to purchase more memory for their old and new machines. Microsoft Windows really took off only when 4-Mbyte main memories became affordable for a significant proportion of PC buyers. The explosive adoption of Windows fueled a surge of demand for DRAMs

6. The January 26, 1993 issue of *PC Magazine* contains a review of the performance of 74 PCs based on the Intel 486DX2-66 processor; 71 of the PCs had secondary caches. The caches ranged in size from 64 to 256 Kbytes. Simple correlation analysis shows that in this class of machine, cache size is a very poor predictor of any reported measure of system-level performance. A similar result follows from an analysis of 49 Pentium systems reviewed in the December 6, 1994 issue of PC Magazine.

© 1996 Steven Przybylski

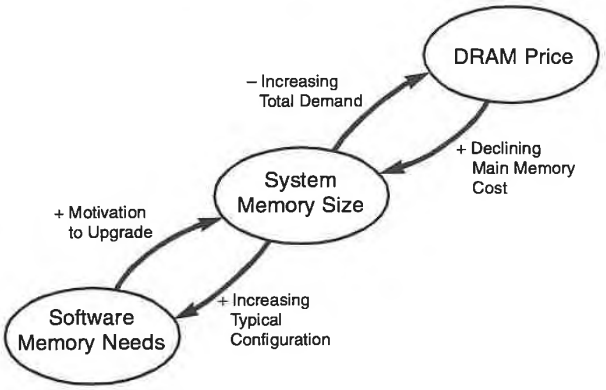


Figure 2-8 Software needs and DRAM prices affect system memory sizes—and vice versa.

in 1992 and 1993 that in turn led to a mild DRAM shortage and a firming of prices.

Main memory sizes are increasing

The factors that influence system main memory size are many and closely interdependent. Most significantly, there is the three-way tug-of-war between software needs, price elasticity of the system markets, and DRAM price. These are shown schematically in Figure 2-8.

Software demands push memory sizes up

On the software side, software vendors grow by selling new copies of their applications to new and existing customers. Staying ahead of the competition requires an ever larger feature list. Additional functionality translates into additional code space and larger file formats. The mantra of the PC software industry appears to be that new software need only hobble along on a minimum configuration system: it does not need to run well. The market only demands reasonable response rates on "typical" configurations. Consequently, users with minimal machines are strongly motivated to upgrade both processor speed and memory size in order to obtain reasonable performance on new software. This motivating factor stems from all types of software: games, edutainment, word processing and other productivity packages, as well as vertical applications. As soon as the majority of people with a minimum configuration have upgraded to a more capable, typical, configuration, that typical

Samsung Ex. 1097, p. 22
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

configuration becomes the new minimum configuration—and the cycle begins all over again. If it were not for the moderating influence of the price of DRAMs, memory sizes would grow rapidly.

Supply and demand limit system memory sizes

There is a subtle relationship between the price of memory and main memory size. As main memory sizes increase, the total demand for DRAMs increases. The market for commodity DRAMs is sufficiently elastic that the total demand comparable to the total supply. Software demands can increase the memory system size only within the constraints of system-level affordability.

Between 1980 and 1991, the total memory cost in an entry-level PC ranged between $50 and $100. As the cost of memory decreased, the amount of memory per system increased, so that the total memory cost per system remained roughly constant. The growth rate in system memory sizes is directly coupled to the price of DRAMs and the many factors that determine the supply and demand of memory.

Since 1991, the cost of memory in an entry-level PC has increased from typically less than $100 to a high of over $200. This increase is because software demands have increased the minimum memory size that yields acceptable performance on low-end processors to 8 Mbytes. The increase in system cost for memory has been possible because system-level prices have remained relatively constant and the cost of a number of other significant system-level components—for example, disks, graphics controllers, and monitors—have come down.

Given that PC consumers have come to expect more capability for less from year to year, it is unlikely that system-level memory costs can rise substantially above $200 for an entry-level system. This will be especially true as 3-D, multimedia, and networking functions increasingly compete with memory, processor, and other existing units for a limited cost budget. For the next several years, the industry will return to a mode in which increases in memory size derive primarily from the ongoing decline in the price of memory bits.

Memory sizes increase in all system classes

As the price of memory has declined, total memory sizes have increased in all computer classes. For most classes of computers, the size of main memory has increased at a rate that is slightly less than the rate at which the price of a bit of memory has come down. Main memory sizes have traditionally doubled approximately every three years.

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 23
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

Even though the total memory size has been increasing, the number of DRAMs in memory systems has steadily declined because the capacity of DRAMs doubles every 18 months. For example, in 1980 a 32-Kbyte PC memory could be built with sixteen 16-Kbit DRAMs.[7] In 1993, a 4-Mbyte PC memory could be constructed with just two 16-Mbit DRAMs. By 1997, a single 64-Mbit DRAM will form an 8-Mbyte PC memory.

Maximizing use of total available bandwidth

Even as the number of DRAMs in a memory system has been declining, the total amount of bandwidth available within the memory system has remained relatively unchanged because the bandwidth per DRAM pin has been increasing. Until recently, designers met growing demands on the memory subsystem by using a higher proportion of the total available bandwidth. They accomplished this either by increasing memory system width or by **interleaving** the banks. Both techniques have the side effect of expanding the minimum memory size and the **granularity** of the memory system.[8]

Memory granularity problems

Memory granularity is becoming a problem, though. As each new generation of DRAMs quadruples the depth of each memory bank, the minimum memory size and minimum memory increment also quadruple. Though the largest acceptable granularity in each market segment grows with the average memory size, the granularities that result from using the cheapest narrow DRAMs are reaching the limits the markets will bear. For example, with 16-Mbit ×4 DRAMs, the granularity of a 32-bit memory system is 16 Mbytes—more than is currently acceptable in the extremely cost-sensitive low-end PC marketplace.

In addition, a number of new application areas require substantial amounts of bandwidth from relatively small memories. Printers, set-top digital cable converters, high-performance personal digital assistants (PDA), and several types of computer graphics controllers need memories of only one to four megabytes to provide tens to hundreds of megabytes per second of bandwidth.

Memory prices keep dropping, but erratically

Despite the growth in main memory sizes, with the per-megabyte cost of memory falling at its traditional rate of 40% per year, the total cost of main memory fell dramatically through the 1970s and 1980s in all classes of computers. For PCs in particular, the drop in total memory costs factored heavily in the decline of end-user prices. In the last few years, memory prices

7. Ignoring the DRAMs used for parity protection.
8. Assuming a constant DRAM width.

Samsung Ex. 1097, p. 24
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

Name	Resolution	Bits per Pixel	Colors	Frame Buffer Size	Refresh Rate	Display Bandwidth
VGA	640×480	8	256	300 Kbytes	60 Hz	18 Mbytes/s
	640×480	16	64K	600 Kbytes	72 Hz	44 Mbytes/s
SVGA	800×600	8	256	480 Kbytes	60 Hz	29 Mbytes/s
	800×600	16	64K	960 Kbytes	70 Hz	67 Mbytes/s
	1,024×768	4	16	393 Kbytes	70 Hz	28 Mbytes/s
	1,024×768	8	256	786 Kbytes	70 Hz	55 Mbytes/s
	1,024×768	16	64K	1.6 Mbytes	70 Hz	110 Mbytes/s
	1,024×768	24	16M	2.4 Mbytes	72 Hz	170 Mbytes/s
	1,280×1,024	8	256	1.3 Mbytes	72 Hz	94 Mbytes/s
	1,280×1,024	16	64K	2.6 Mbytes	72 Hz	189 Mbytes/s
	1,280×1,024	24	16M	3.9 Mbytes	72 Hz	283 Mbytes/s
	1,600×1,200	24	16M	5.8 Mbytes	72 Hz	414 Mbytes/s

Table 2-1 The display bandwidth grows with the screen resolution, color depth, and refresh rate.

have fallen more slowly, causing total memory costs to increase as average memory sizes increase. As dramatic price jumps in mid 1993 demonstrated, DRAM prices have always been erratic over the short term. The dramatic price declines of late 1995 and early 1996 have underscored that lesson.

2.3.3 Computer Graphics Trends

Over the past few years, personal computer users have demanded much better graphics performance. The resolution, size of the color palette, and display responsiveness for complicated images have become significant decision factors for many buyers. As the industry transitioned from VGA and its lower resolution predecessors (MDA, CGA, and EGA) to SVGA and higher resolution modes, the amount of memory needed to store the graphics image has increased significantly. Table 2-1 shows the frame buffer sizes needed at a variety of resolutions. During the early 1990s, the most popular resolutions were up to 1,024×768, with 16 or 256 colors—resolutions that allow the frame buffer to fit into 512 Kbytes of **graphics memory**.

As more end users expect the ability to manipulate true color images, and as affordable processors become able to manipulate those images responsively, graphics controllers with 16 and 24

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 25
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

bits per **pixel** will become increasing popular. Images and systems with 24 bits per pixel are called true-color images and systems. However, images displayed with 16 bits per pixel are almost indistinguishable from their true-color counterparts, especially at higher resolutions. With more bits per pixel, the amount of graphics memory grows to 2–4 Mbytes. More significant is the increased amount of bandwidth needed just to continually dump the contents of the **frame buffer** to the screen. The display bandwidths shown in Table 2-1 are average bandwidth requirements. Peak bandwidth demands made on a frame buffer are typically about 50% higher. In addition to increasing processor demands on the frame buffer, ever larger display bandwidth requirements will soon make bandwidth the main concern of designers of PC video controllers.

Multimedia requires more bandwidth

Users increasingly expect multimedia PCs to display video images in a window. This means that additional video memory bandwidth is needed to write the video image continuously into the frame buffer. The bandwidth content of an NTSC video signal is approximately 25 Mbytes/s at 24 bits per pixel and 30 frames per second.

Another important traditional difference between the video and computer graphics worlds is that of color space. Computer graphics are generally defined in the RGB color space, in which the intensity of the red, green, and blue color components are individually specified. In the video world, each pixel is specified in the YUV color space or one of its derivatives. In this equivalent but different color space, Y represents the luminance or black-and-white intensity while U and V encode the color components. When displaying a video signal in a computer graphics system, the video signal must either be converted to the RGB color space before the video information is written into the frame buffer, or the pixels in the frame buffer must be marked as YUV pixels and displayed as such by the RAMDAC.

3-D Graphics increases bandwidth and frame buffer requirements

3-D graphics is substantially more complex than 2-D graphics. One of the fundamental complexities is deciding whether one object is in front of—and therefore occludes—another. This problem is typically solved by rendering (i.e., drawing into the frame buffer) both objects and deciding on a pixel by pixel basis what is closest to the observer and should be displayed. To accomplish this, in addition to a color value, a depth or z-value is stored for each pixel. Depth values typically range between 16 and 32 bits, and are stored in either floating point or fixed point formats. To draw a pixel of an object, the controller first checks the depth of the current pixel, and overwrites it with the new

Samsung Ex. 1097, p. 26
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

color and depth only if the new depth is less than the previous depth. Naturally this takes substantially more bandwidth into the frame-buffer/z-buffer combination than is typical for 2-D systems.

3-D graphics systems are generally double-buffered. This means that there are two complete copies of the frame buffer. While one is used for rendering the scene, the other is used for displaying the previous rendition. Once the scene is completely rendered, the roles played by the two frame buffers are switched, and rendition begins again. Double-buffering facilitates smooth graphics even when the rendition rate is substantially slower than the screen refresh rate. Specifically, without double-buffering, an occluded object may be occasionally and accidentally displayed when a portion of a frame buffer is displayed to the screen between the rendition of a far object and that of a near occluding object.

High-performance graphics

In the workstation world, display resolutions of $1,024 \times 768$ and $1,280 \times 1,024$ have been common for some time. One important application has been the real-time or near-real-time rendering of complex two- and three-dimensional objects and scenes. In systems targeted at these applications, a new image is drawn into the frame buffer on every frame: the bandwidth needed to write the image into the frame buffer must at least equal the bandwidth needed to dump it to the display. High-performance graphics architecture's principal focuses are on 1) providing the computational elements required to render complex images, and 2) mating the processors with the video memory so as to provide adequate bandwidth and computation across the image.

Overall, the appropriate resolution for a computer graphics system is closely related to the display technology. The cost of computer monitors are more determined by the cost of glass and metal than by the cost of electronics. Consequently, price reductions for color monitors have been more modest than for other components of the computer system. As a result, the industry is transitioning slowly to 17" monitors and the 1,280 x 1,024 resolutions that they facilitate. Although there is a quadratic relationship between the linear pixel density and the frame buffer size and, consequently, the screen refresh bandwidth, all three quantities are expected to grow fairly slowly. The expectation of a slow rate of growth in the screen refresh bandwidth is further justified by the recognition that pixel depths beyond 24 bits add nothing to the visual perception of the rendered image.

© 1996 Steven Przybylski

2.4 Combining DRAM and Systems Trends

So far, we have noted with respect to DRAMs:

- DRAM sizes are growing at 60% per year.
- Row access times have been declining slowly (7% per year).
- Column access times have been declining at close to 25% per year.
  - Bandwidth per *pin* has been increasing.
  - Bandwidth per *bit* has been decreasing.
- TTL signaling levels are becoming inadequate as board-level operating frequencies approach and surpass 100 MHz.
- The cost of DRAM fabrication facilities has been growing, on average, at 30% per year.
- Over the long term, the cost of memory continues to decline.

With respect to systems:

- Memory sizes have been rising at roughly 25% per year.
- Total main memory costs have been falling.
- The number of DRAMs per system has declined continuously.
- The total amount of bandwidth available to designers within the memory system has been roughly constant.
- Processor and I/O demands on main memory have been growing despite heavy use of CPU caches.
- The granularity of a memory system quadruples with each new DRAM generation unless the DRAMs get wider or the memory system gets narrower.
- Graphics memory sizes and bandwidths have grown significantly in recent years.

Figure 2-9 illustrates three of the most important trends. For several classes of machines, the number of DRAMs needed as a function of time is shown. As each new generation of DRAMs comes into use, the number of DRAMs declines, even as memory size increases. With the transition to 16-Mbit devices, the

Samsung Ex. 1097, p. 28
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

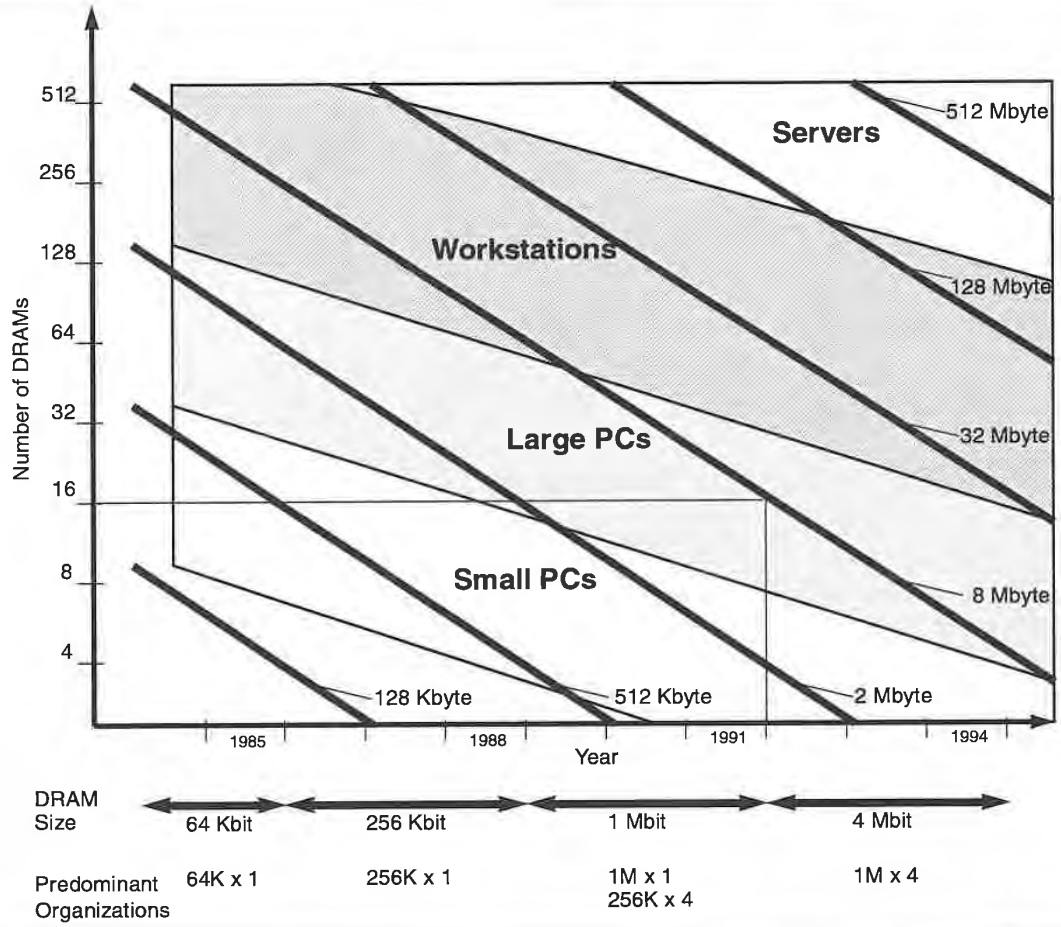


Figure 2-9 Because memory sizes are growing more slowly than DRAM sizes, the number of memory chips in a system decreases over time.

memory systems of many smaller machines and embedded systems can be built with a single DRAM.

Figure 2-10 extends this graph to the end of the decade. By the time 256-Mbit DRAMs become commonplace, most machines will contain four or fewer DRAMs. By the year 2000, only the largest machines will contain more than one hundred DRAMs.

Memory system bandwidth lagging demand

One aspect of this projection is troubling: total available bandwidth within the memory system is lagging behind the increasing demands placed on memory. Until recently, generic DRAMs have met these demands at the expense of increasing memory system granularity. In the future, DRAMs will have to provide

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 29
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

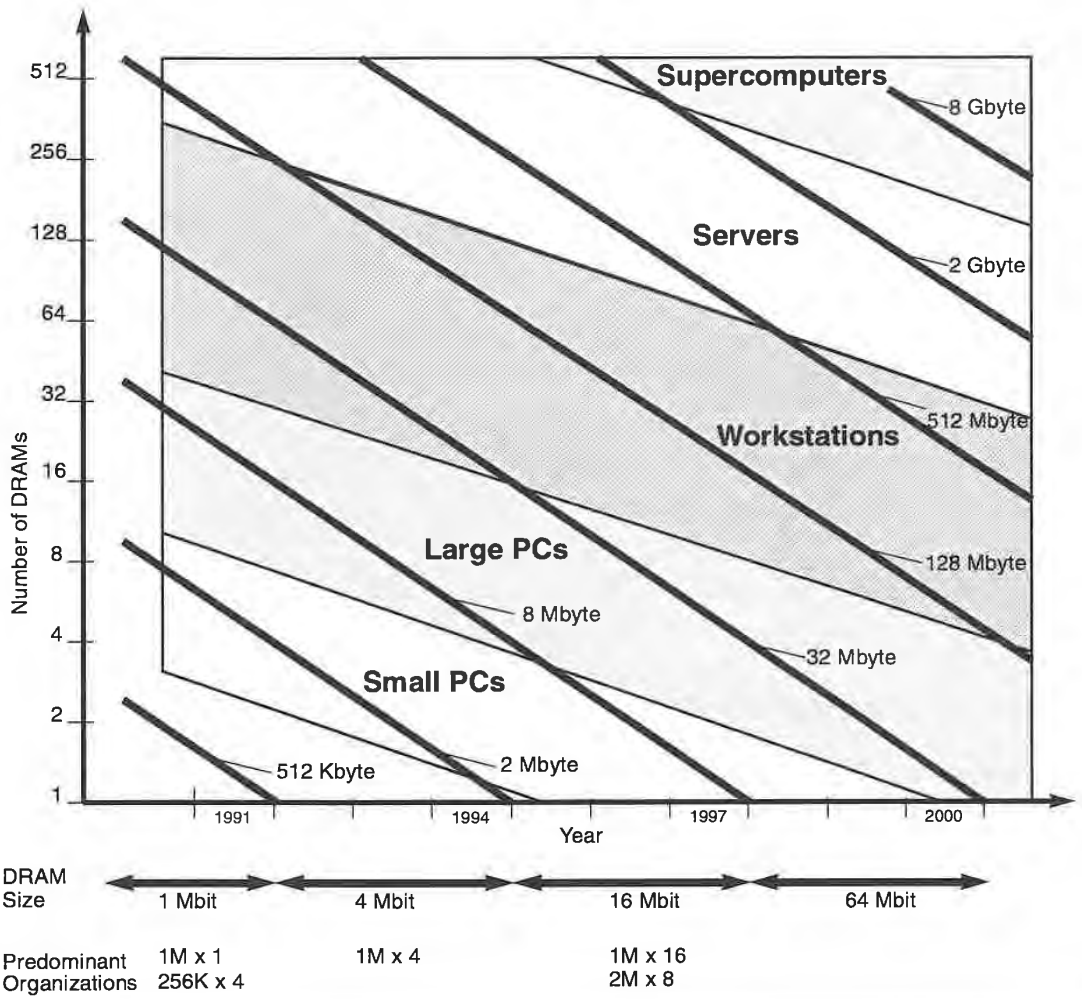


Figure 2-10 By the end of the decade, most machines will be sold with only a few DRAMs. Only the largest machines will have a hundred or more DRAMs.

more bandwidth per bit of storage. DRAM designers can make this happen by 1) increasing DRAM width, 2) changing the DRAM interface to make it more efficient, or 3) combining the first two strategies. Without significant changes in the DRAM interface, either granularity will grow unacceptably large, or the inadequate bandwidth will increasingly make memory a system-level bottleneck—or both.

Samsung Ex. 1097, p. 30
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

Today's DRAMs have been optimized for low cost and reasonable performance when many DRAMs are used in parallel. They are ineffective when used to build small high-performance memory systems—the very type of memory system needed as processor performance improves and the number of DRAMs per system shrinks.

The changes occurring in the DRAM marketplace result directly from the growing inadequacy of the existing DRAM interface. The following section illustrates how a system's bandwidth and granularity requirements can be combined into a single metric— a metric that enables comparisons of the capabilities of individual DRAMs and selection of viable candidates for a specific application.

2.5 Fill Frequency: A New DRAM Metric[9]

The typical view of the growing inadequacy of the traditional DRAM interface is shown in the left side of Figure 2-11. It shows the *Performance Gap* between the diverging peak bandwidth requirements of CPUs and the peak bandwidth capabilities of a DRAM on unlabeled axes. Though simple and conceptually pleasing, this representation of the current situation fails to explain why the explosive levels of change are occurring now, as opposed to five years ago or five years from now.

A modified and somewhat more realistic representation shows the post-cache CPU bandwidth requirements and notes that a cross over is occurring between this requirement and the bandwidth that is obtainable with a single bank of DRAMs (see the right side of Figure 2-11).

A third approach, presented for the first time in this report, involves quantifying system bandwidth and granularity requirements and comparing them with the capabilities of specific DRAMs. The metric that facilitates this quantitative comparison is called the *fill frequency.*

9. The author's appreciation and thanks go to Subodh Toprani of Rambus Inc. for planting the seeds for this analysis. It is the author's understanding that the idea of combining bandwidth with capacity characteristics of DRAMs originated with Mike Farmwald.

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 31
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

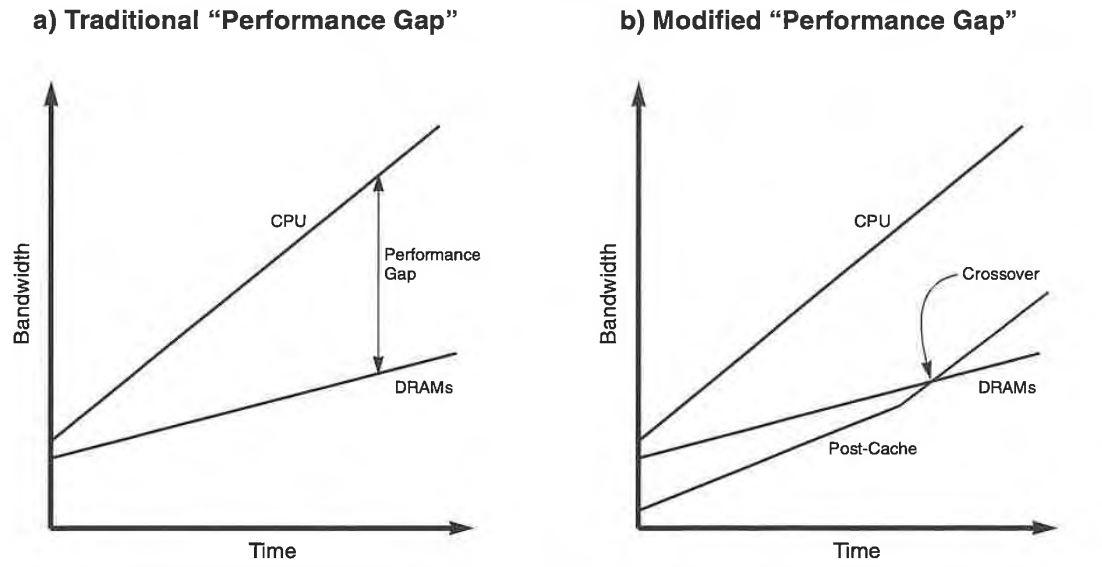


Figure 2-11 Two simplistic views of the performance gap between CPUs and DRAMs inadequately explain the nature of the dilemma facing the DRAM vendor community.

Granularity and bandwidth requirements change with time

As discussed in Section 2.3.2, the granularity of a memory subsystem is equal to the minimum amount of memory that can be present. Often, the granularity of a memory system is also equal to the minimum amount of memory that can be added to a system in the field.[10] The maximum granularity that is practical in a marketplace varies with the marketplace and with time. For example, a PC with an 8-Mbyte granularity was impractical in 1990 because most PC purchasers would not buy a system with a minimum memory size of 8 Mbytes. In 1996, 8 Mbytes is a reasonable granularity for a Pentium-based system. However, it remains an unacceptable granularity for PDAs because of power, physical volume, and cost constraints. A PDA with a minimum memory size of 8 Mbytes would not sell well.

Markets also have specific peak and sustained bandwidth requirements. However, these requirements are more imprecise and more variable within a market segment than the granularity requirements. This is because end users are less aware of

10. Technically, the granularity is the minimum size of a single rank of memory. See Section 4.2.2.

Samsung Ex. 1097, p. 32
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

the amount and cost of their system's bandwidth than they are of the amount and cost of memory. However, a system with too little memory bandwidth is likely to perform poorly—and therefore is unlikely to garner as many Editor's Choice awards. Sales of such a system would suffer in spite of any price reductions that might stem from its lower-than-typical bandwidth memory system.

Trading off bandwidth and granularity

When designing a memory system, developing more bandwidth is always easy, so long as the granularity is allowed to grow proportionately.[11] Difficulties arise when bandwidth has to increase more than granularity or, alternatively, when granularity has to shrink without a commensurate decrease in bandwidth. In summary, it is hard to get a lot of bandwidth out of a small amount of memory. Engineering is the art of trading off market requirements that are naturally at odds with one another: in this case, low memory granularity and high memory bandwidth.

The ratio of the bandwidth to the granularity nicely characterizes the market requirements placed on the memory subsystem design. The ratio, expressed in Mbytes per second per Mbyte (MBps/Mbyte), indicates the degree of relative ease or difficulty in simultaneously meeting a market's bandwidth and granularity requirements. A low ratio means that market requires relatively little bandwidth compared to its large entry-level memory sizes. In contrast, a high ratio means a great deal of bandwidth must be extracted from a small amount of memory. Table 2-2 shows granularity, bandwidth, and the corresponding ratio for a variety of system types. To reiterate, the bandwidth requirements are more fungible than the granularity requirements and represent the approximate center of a range of acceptable bandwidths.

The inverse of the ratio in Table 2-2 is the ratio of the system's memory size—or capacity—to the system's peak memory bandwidth—or fastest rate of data transfer. This ratio is expressed in Mbytes/MBps. This ratio is the amount of time (in seconds) that it takes to completely fill the memory at the maximum fill rate (bandwidth). It follows that the inverse of this ratio (i.e., the ratio of the bandwidth to the granularity) is the frequency with which the memory can be filled, measured in hertz (Hz). Consequently, this ratio is called the fill frequency of the memory system. Because the bandwidth and granularity being divided are

11. A complete discussion of the trade-offs and techniques inherent in memory subsystem design is given in Chapter 2.

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 33
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

System Type	Approximate Minimum Peak Bandwidth (MBps)	Maximum Granularity (Mbyte)	Ratio or Fill Frequency (MBps/Mbyte) (Hz)
3-D Frame Buffer	800	4	200
High-end 2-D PC Frame Buffer	400	4	100
Low-end 2-D PC Frame Buffer	70	1	70
Low-end PC	100	4	25
High-end PC	176	8	22
Workstation	800	64	12.5

Table 2-2 The granularity and bandwidth requirements vary substantially across system types. The ratio of the two is the memory system's required fill frequency.

market requirements, the fill frequency ratio is also a market requirement.

Three important observations follow from Table 2-2:

- Graphics applications are more demanding than main memory applications: the bandwidth requirements are higher and the memory sizes are smaller.
- Among the main memory applications of DRAMs, large computer systems are less taxing of their memory systems than small computer systems. This is because as one climbs up the hierarchy from PCs to workstations to mainframes, the maximum acceptable granularity grows more quickly than the minimum acceptable bandwidth.
- All of the ratios fall within a relatively narrow range. Because bandwidth and granularity requirements generally grow or shrink together, the highest ratio is only one and a half orders of magnitude greater than the smallest.

DRAM sizes and bandwidths also vary

Interestingly, DRAM characteristics include a size and a peak bandwidth deliverable by a single device. For example, a 1M × 16 fast-page-mode DRAM has a size of 2 Mbytes, a width of 2 bytes, and a peak operating frequency of 25 MHz. This width and frequency combine to make a peak bandwidth of 50 MBps and a ratio of peak bandwidth to size of 25 MBps/Mbyte. Table 2-3 shows the peak bandwidth, size, and ratio for a variety of current DRAMs.

Samsung Ex. 1097, p. 34
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

DRAM Type	Width (Bytes)	Maximum Operating Frequency (MHz)	Peak Bandwidth (MBps)	Size (Mbyte)	Ratio or Fill Frequency (MBps/Mbyte) (Hz)
1M×1; -10	1/8	16	2.1	1/8	16
1M×4; -8	1/2	20	10	1/2	20
4M×1; -8	1/8	20	2.5	1/2	5
1M×16; -6	2	25	50	2	25
4M×4; -6	1/2	25	12.5	2	6.25
16M×1; -6	1/8	25	3.1	2	1.6

Table 2-3 The sizes and peak bandwidths of DRAMs vary substantially.

Just as with the ratio of the system requirements of Table 2-2, the ratio in Table 2-3 describes how frequently all the bits in the DRAM can be read or written. Consequently, it too is called the **fill frequency**—this time it is the fill frequency of a particular DRAM. When the units of information (Mbytes) are cancelled from the numerator and denominator of these fill frequency ratios, what remains is "per seconds," which is the same as Hz. In other words, 25 MBps/Mbyte is equivalent to 25 Hz.

In this report, *fill frequency* is frequently abbreviated to *FF*.

DRAM and memory system fill frequencies can be compared

One final crucial observation links the fill frequencies of memory systems and those of DRAMs. Because the fill frequencies of DRAMs represent the absolute maximum rate at which all data in the device can be read or written, no memory system constructed of a particular DRAM can have a fill frequency that exceeds the peak fill frequency of that DRAM.[12]

Interleaving doesn't change peak fill frequencies

In particular, interleaving or other rearrangement of the DRAMs into different logical structures cannot create a memory system that has a greater peak fill frequency than that of its constituent DRAMs: interleaving can create more bandwidth but the granularity necessarily increases proportionately.

Recall that the fill frequency requirements of Table 2-2 were minimum requirements. In order to succeed in the marketplace, a system's memory system had to have a higher fill frequency than the minimum imposed by the marketplace. Now, because

12. The important distinctions between peak, sustained, and sustainable fill frequencies are covered in more detail in Chapter 4.

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 35
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

the fill frequency of a DRAM is an upper bound on the fill frequency of the memory systems that use it, it follows that all memory systems constructed of a DRAM will have fill frequencies less than or equal to the DRAM's. Consequently, if a DRAM's fill frequency is less than a market's required fill frequency, then all memory systems built with that DRAM will have inadequately low fill frequencies: all systems that attempt to use such DRAMs will fail in the marketplace.

Consequently, DRAMs with fill frequencies smaller than a market's minimum required memory system fill frequency **cannot** be used successfully in memory systems for that marketplace. Therefore, the fill frequencies of DRAMs and markets can be directly compared to determine which DRAMs are unacceptable in which markets.

Market requirement fill frequencies are roughly constant

This crucial observation—that DRAM and memory system fill ratios can be directly compared to determine the applicability of particular DRAMs to specific systems—leads to the powerful representation of the dilemma facing the DRAM industry shown in Figure 2-12. This figure shows the fill frequencies of the predominant DRAMs and of several important memory system types from 1980 to the present. During most of the 1980s, memory system fill frequencies were constant or slightly declining. This decline stemmed mostly from steadily growing cache sizes and the resulting reductions in the frequency of references to main memory.

In the latter half of this decade, though, the memory system fill frequencies are likely to increase or at least hold constant. This halt in the decline is likely because of two factors:

- Further increases in cache sizes have decreasing marginal influence on the size of the load on main memory, and
- Introduction of audio, video, and other memory-intensive data streams will add significantly to the processor's load on main memory.

The only scenario in which the system requirement fill frequencies might decline sharply in the future would involve the continued dramatic price erosion caused by a severe oversupply of DRAMs of all types. If the price per bit were reduced by half, the granularity requirements would increase significantly in the short term because end users could buy systems with more memory for the same system cost. Consequently, the system requirement fill frequencies would decrease correspondingly. In the longer term, however, the changing cost equation for PCs

Samsung Ex. 1097, p. 36
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

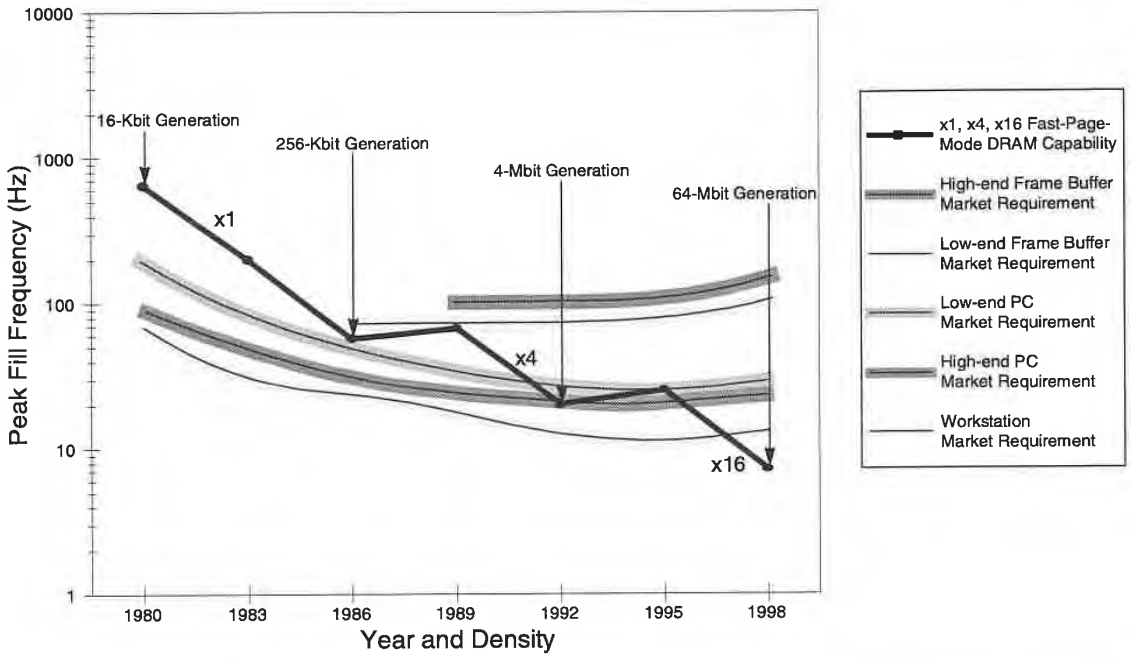


Figure 2-12 Plotting system and DRAM fill frequencies together shows that traditional narrow DRAMs are rapidly becoming inadequate for essentially all applications.

would shift some of the cost savings derived from memory to CPU and other functional units causing a corresponding, and offsetting, increase in bandwidth requirements. Alternatively, a significant drop in the cost of memory would facilitate a drop in the price of comparably configured systems. In this eventuality, the fill frequency requirements would remain unchanged in the face of a reduction in the price of memory. The likely situation is a combination of the two scenarios, in which the decline in memory prices fosters both a reduction in system prices and a continued mild downward pressure on the fill frequency.

But DRAM fill frequencies decline precipitously from generation to generation

In contrast to the modest changes over time in the fill frequency requirements, DRAM fill frequencies have declined steadily. Because the column cycle times have improved slowly relative to the increases in density, the DRAM fill frequency for a constant device width has declined by close to a factor of four every three years. Twice in the last decade, the width of the most common device has quadrupled in order to keep the DRAM fill

© 1996 Steven Przybylski

frequency above the all-important PC memory system fill frequency requirements. At the end of the 1980s, the industry shifted from ×1 to ×4 devices to keep memory system granularity within acceptable bounds. A similar shift to ×16 devices occurred at the 16-Mbit density because PCs were not ready for 4M deep memory systems.

However, at the 64-Mbit density, even ×16 devices fall below the important PC fill frequencies. This indicates that all the standard narrow fast-page-mode devices will be inappropriate for essentially all applications at this density. Consequently, two scenarios are possible:

- A set of new DRAM architectures with fill frequencies above the relevant memory subsystem fill frequencies will be adopted, or
- Expectations will decline and with them the system fill frequencies.

The new DRAM alternatives coming onto the market indicate that the first scenario is already occurring.

Given that new memory architectures will be adopted, the important questions are: 1) which architectures are going to be successful, and 2) when will those architectures be adopted in volume. Each of the many alternatives being proposed by various elements in the DRAM vendor community has its own strengths and weaknesses. This report describes the new devices, exposes their strengths and weaknesses, and highlights the markets to which they are most and least suited.

Concern for 256 Mbits and beyond

The graphs in Figures 2-10 and 2-12 stop with the 64-Mbit generation. However, both graphs hold serious implications for the 256-Mbit and 1-Gbit generations. By the time the 256-Mbit DRAMs are cost-competitive early in the next century, most systems will have one or two DRAMs. Many substantial markets will want less than a full DRAM's worth of memory. Perhaps more disconcerting is that the standard ×16 fast-page-mode DRAM will offer inadequate bandwidth per bit stored by a factor of 50 or so. By the time the 1-Gbit generation is upon us, the average number of DRAMs per system will be at or below 1, and the bandwidth per bit of today's most common organization will be inadequate by a factor of 200. All of the architectures presented in this report as being effective at the 16- or 64-Mbit levels will fail at these higher densities unless further significant changes occur.

Samsung Ex. 1097, p. 38
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

The exponential growth in DRAM density over time has been and will continue to be one of the most powerful forces in electronics. However, the DRAM industry will face unique and daunting challenges in every domain if it is to produce a 1-Gbit device that is at all effective in the systems of the future.

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 39
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

Samsung Ex. 1097, p. 40
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

3 DRAM Basics

3.1 New In The Second Edition

Changes to this chapter from the first editions include minor updates and corrections, and a substantial revision of Section 3.6 on Video RAMs.

3.2 Introduction

All the important traits of memory systems stem from characteristics of DRAMs: their volatility, internal hierarchy, narrow width, and great depth. By illustrating the important characteristics of all DRAMs in general, and of the generic multiplexed-address fast-page-mode DRAM in particular, this chapter sets the stage for discussions of memory systems in Chapter 4 and of new DRAM organizations in Chapters 6 through 8.

The discussion of DRAM basics in this chapter begins with a general description of how a DRAM stores information, how information is presented to a **generic DRAM**, and how it is retrieved again. Section 3.4 then describes the semiconductor technology used to make DRAMs and discusses the evolution of that technology over the past two decades. Section 3.5 contains greater detail about the finer points of DRAM design and their implications for present and future memory systems.

Samsung Ex. 1097, p. 41
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

Computer graphics and multimedia are increasingly important, and **Video RAMs (VRAMs)** play an integral role in the implementation of midrange and high-end graphics subsystems. Section 3.6 highlights the logical and structural differences between DRAMs and VRAMs. The discussion of DRAM and VRAM differences prepares the reader for Section 4.7, on graphics controllers, and Section 10.6.1, on the effectiveness of the new DRAM alternatives in video applications.

3.3 What Is a DRAM?

3.3.1 The Basics

DRAM stands for Dynamic Random Access Memory. *Random access* refers to the ability to access any of the information within the DRAM in random order. *Dynamic* refers to temporary or transient data storage. Data stored in dynamic memories naturally decays over time. In contrast, static memories, such as those in **Static RAMs (SRAMs)** retain data as long as power is applied. Without some mechanism to periodically restore or **refresh** the data, information stored in a DRAM is quickly lost. Traditionally, DRAM designers have included mechanisms to facilitate refreshing the data; the memory system designer is responsible for ensuring that the DRAMs are refreshed frequently enough to prevent data loss.

The heart of the DRAM is the one-transistor storage cell. This storage cell, shown in Figure 3-1, consists of a transistor and a capacitor connected to a bit line and a word line. The word line controls the movement of a bit of information into or out of a cell, and the bit line transmits that bit to the rest of the system.

When the word line is raised, the access transistor is turned on and the bit line is connected to the storage capacitor. When the storage capacitor is charged, the cell is logically storing a 1; when it is discharged, the cell is logically storing a 0. Unfortunately, real capacitors leak charge. Over time, the charge that represents a logic 1 leaks away, causing the logic 1 to change to a logic 0. Before this happens, the cell must be refreshed by reading the value of the cell and restoring a full charge to the capacitor.

Samsung Ex. 1097, p. 42
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

DRAMs are organized in rows and columns

Storage cells are physically and logically oriented in rows and columns, which share word lines and bit lines, respectively. Figure 3-2 shows the internal organization of a ×1 DRAM. The address specifies the individual bit being accessed. Roughly half of the address specifies the row of the addressed bit. This is called the **row address**. The rest of the address is called the **column address** and is used to identify the column of the addressed bit.

A 4M×1 DRAM has 11 row address bits and 11 column address bits. Logically, it has 2,048 rows and 2,048 columns. The cost and size of the package are reduced by decreasing the number of pins on the DRAM IC. This pin reduction is accomplished by sending the two halves of the address on the same address bus, A0..10. The row address is sent first, followed by the column address. When the row address is decoded and a word line is raised, all the bits in a row are transferred to their corresponding bit lines and stored in the sense amplifiers. When the column address arrives, the individual bit or bits accessed are selected from those currently in the sense amps and sent to the I/O pin(s). An important consequence is that once a row access has brought a large number of bits into the sense amplifiers, more than one column access can be performed to retrieve a collection of bits that share the same row address. The collection of bits retrieved from the storage cells and made available by a row access are often referred to as a single page of the DRAM.

Sense amplifiers form a row cache

Having an intermediate storage in the sense amplifiers naturally defines a storage hierarchy within each DRAM. Transfer-

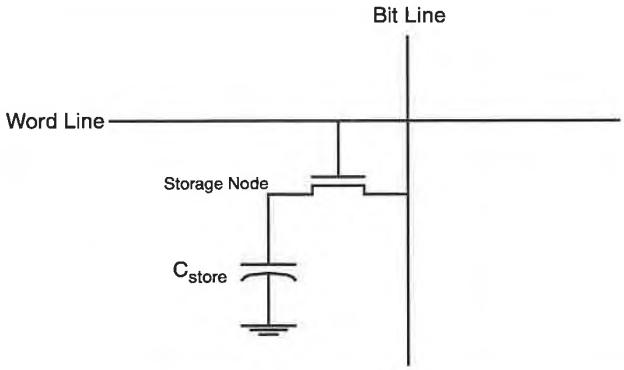


Figure 3-1 The DRAM storage cell consists of a capacitor for storing charge and a transistor for gating that charge onto a bit line.

© 1996 Steven Przybylski

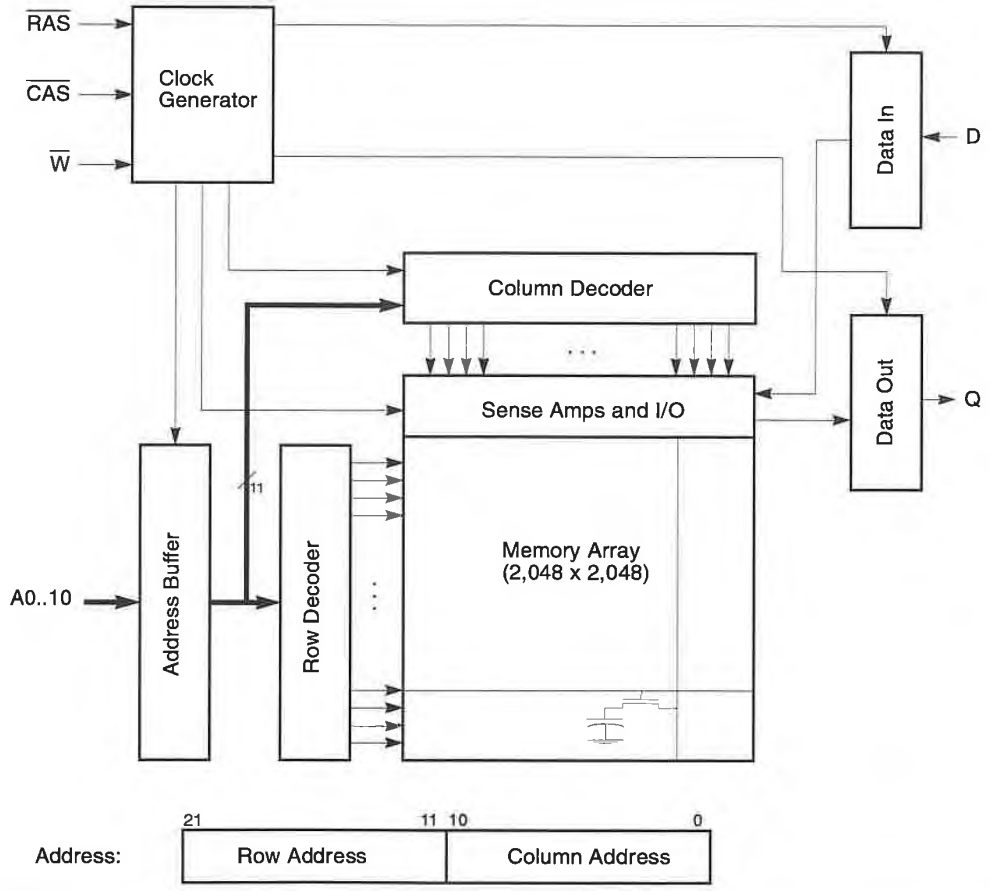


Figure 3-2 The DRAM address is divided into a row address and a column address, which specify the word line and bit line that intersect at the addressed bit.

ring (or reading) the bits from the storage cells into the output pins via the sense amplifiers is much slower than transferring bits from the sense amps directly to the DRAM's pins. The sense amps act like a **cache**: a small high-speed memory that contains a copy of a small fraction of all memory. Since this cache holds a single row of data, it is called a **row cache**. One conceptual view of a DRAM is that a row access loads a row of data from the DRAM array into the row cache, from which the data is accessible via column accesses.

An important aspect of the one-transistor cell is that the read is destructive: the process of reading the state of the cell corrupts

Samsung Ex. 1097, p. 44
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

the value stored. When a word line is raised and the bit line is connected to the storage capacitor, most of the charge stored on the capacitor is transferred to the bit line. Consequently, if the value is not restored prior to the next read operation, the cell will be read incorrectly the second time. To ensure that the bits are read nondestructively, read operations are implemented as a row read immediately followed by a write of the entire row back into all the cells in the accessed row.

A side effect of the need to immediately restore all read bits is that write operations are very similar to read operations. To perform a write of a particular bit in the memory array, the entire row containing the addressed cell is read into the sense amps, where one or more of the bits are replaced with new values. The entire row is then written back into the memory array.

A refresh operation is very similar. Conceptually, it is a row access in which none of the bits accessed in a row are sent off chip. Alternatively, it can be viewed as a write in which none of the bits in the accessed row are replaced with new values. In general, DRAM manufacturers guarantee that if all bits are refreshed every 8–64 ms, then no information will be lost due to a failure to retain charge. Because refreshing the data takes power, devices with longer refresh periods are used in low-power and portable applications. These long-refresh devices are really no different from the regular devices, except that they pass additional tests to verify longer data retention.

The generic fast-page-mode DRAM interface is shown in Figure 3-3. In addition to a multiplexed address bus and a data I/O path, most DRAMs have three or four control pins: $\overline{\text{RAS}}$ **(Row Address Strobe)**, $\overline{\text{CAS}}$ **(Column Address Strobe)**, $\overline{\text{WE}}$ (Write Enable), and sometimes $\overline{\text{OE}}$ (Output Enable). The falling edges of $\overline{\text{RAS}}$ and $\overline{\text{CAS}}$ indicate that the address lines contain the row address and column address, respectively. The Write Enable pin controls whether data is being read from the DRAM or written into it. The write information is provided at or following the falling edge of $\overline{\text{CAS}}$. In DRAMs that use the same pins for data input and output, the Output Enable control input indicates to the DRAMs when to drive output data onto those data pins.

Three styles of generic DRAM

As mentioned above, a row access brings many bits into the sense amps.These bits can be accessed quickly by column accesses without reference to the DRAM array. Three styles of generic DRAM provide access to those bits: **fast-page mode**, **static-column mode**, and **nibble mode**. Figure 3-3 shows accesses to a fast-page-mode DRAM with a total of five column

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 45
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

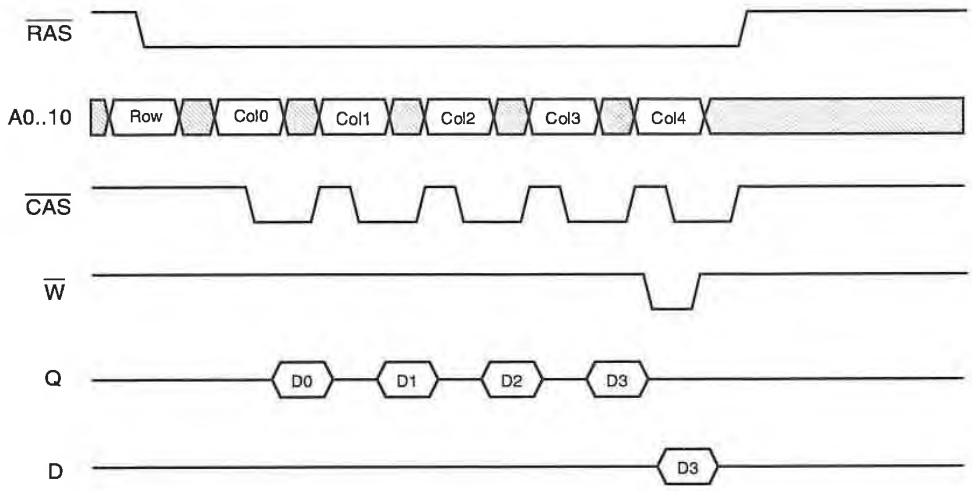


Figure 3-3 Row and column addresses are latched on the falling edges of RAS and CAS, respectively.

accesses (four reads and a write) to one row. With fast-page-mode DRAMs, each new column access is indicated by a falling edge in CAS. In static-column-mode parts, each new column access is indicated by a change in the address bus; CAS need not be toggled each access. Finally, in nibble-mode parts, four bits are serially accessible by toggling CAS while RAS is held low. The DRAM takes care of incrementing the two low-order address bits—the DRAM pays no attention to the address bus after the first column access. Nibble mode is rarely used today due to its inflexibility relative to fast-page mode and static-column mode. Some vendors have even stopped designing nibble mode into their DRAMs. Static-column mode is not as commonly used as fast-page mode, mainly because it is more expensive. It is more expensive because it is not as commonly used, and sales of static-column-mode DRAMs are very much smaller than sales of fast-page-mode DRAMs.

A refresh cycle is generally initiated in one of four ways. A **RAS-only refresh** is like a normal read cycle without any column accesses: RAS is lowered, held low long enough for the DRAM to complete the row access, and then raised without lowering CAS at all. The memory controller provides the address of the row to be refreshed on the falling edge of RAS. An alternative to a RAS-only refresh is a **CAS-before-RAS refresh**. In this case, CAS is

Samsung Ex. 1097, p. 46
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

already low when $\overline{RAS}$ is lowered. The address of the row to be refreshed is taken from an internal row counter; the address bus is ignored. A third alternative, called **hidden refresh**, places a CAS-before-RAS refresh cycle at the end of a regular read cycle, with a small saving in time. Its advantage is that the data from the previous read operation is available on the data pins during the refresh operation. The fourth alternative, found in some DRAMs targeted at low-power operations, is a **self-refresh mode**. By extending a CAS-before-RAS refresh cycle for a long time (more than 64 ms), the part begins to refresh itself autonomously, without any help from the memory controller. System power can be minimized during idle periods by allowing the memory controller to go into an inactive mode.

The refresh requirements of a DRAM are specified by two numbers: 1) the refresh period during which all bits in the DRAM must be refreshed at least once and 2) the number of row addresses that must be refreshed within the refresh period. The ratio of the refresh period to the number of rows is the time between refresh cycles. This ratio is a direct measure of the proportion of total available bandwidth that must be used to refresh the devices. For example, a 1M×4 DRAM that requires 1,024 refresh cycles every 16 ms is busy refreshing data at least 0.8% of the time.

A typical DRAM data sheet specifies 50–100 timing parameters. Together, these parameters describe how the DRAM is to be used and how it behaves. From a system-level perspective, only a handful of these parameters determine the performance of the memory system. These are listed and explained in Table 3-1. The values shown are typical for fast-page-mode 4M×1 DRAMs. DRAMs are identified by their row access times. All DRAMs that are marked -60 have 60-ns row access times. Although the other timing parameters may vary slightly from manufacturer to manufacturer, this variance is generally small enough that comparably marked DRAMs are completely interchangeable in most systems.

The table gives rise to an important observation: a 60-ns DRAM can perform a row access only every 110 ns. More significantly, although the column (or CAS) access time is 15 ns, the absolute maximum rate at which the DRAM performs column accesses is once every 35 ns. In practice, as Chapter 4 demonstrates, the maximum rate is significantly slower: typically every 40–50 ns.

DRAM packages

DRAMs are available in a wide variety of packages. The packages, which are made of plastic or ceramic, have 16 to 50 pins,

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 47
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

Symbol	Name	Typical Values (ns)				Explanation
		-50	-60	-70	-80	
t_{RC}	Random Read/Write Cycle Time	90	110	130	150	The minimum time from the start of one row access to the start of the next
t_{PC}	Page-Mode Cycle Time	30	35	40	45	The minimum time from the start of one column access to the start of the next
t_{RAC}	Access Time from $\overline{RAS}$	50	60	70	80	The minimum time from $\overline{RAS}$ falling to valid data output
t_{CAC}	Access Time from $\overline{CAS}$	13	15	18	20	The minimum time from $\overline{CAS}$ falling to valid data output
t_{AA}	Access Time from Address Valid	25	30	35	40	The minimum time from a valid column address in to valid data out
t_{RCD}	$\overline{RAS}$ to $\overline{CAS}$ delay	18/37	20/45	20/50	20/60	The minimum and maximum time between a falling edge in $\overline{RAS}$ and a falling edge in $\overline{CAS}$

Table 3-1 Only a few of the many DRAM timing parameters determine memory system performance.

depending on the size and the width of the device. A complete description of the packages, their characteristics, and options is given in Section 5.2.1.

3.3.2 DRAM Internals

As with any complex and subtle piece of engineering, the reality of DRAM implementation varies significantly from the straightforward conceptual view presented in Section 3.3.1. This section peels back one more layer of the onion to reveal some of the significant physical and circuit properties of modern DRAMs.[1]

Figure 3-4 shows in more detail the circuit implementation of a dynamic memory array. The important features shown are the differential sense amplifiers, the $V_{cc}/2$ cell plate, the dummy cells, and the segmented and folded bit lines. A real DRAM design may contain all, some, or none of these features.

Small storage cell versus large capacitor

The fundamental challenge of DRAM design is found in the tension between the need for a physically small storage cell and the need for an electrically large storage capacitor (high capacitance). Minimizing the physical size of the storage cell is important because there are so many storage cells per DRAM: one per bit stored. For most DRAMs, 50–60% of the IC's die area consists of storage cells. Larger storage cells make for larger, more

1. For a more thorough description of the history and future of DRAMs and other semiconductor memories, refer to B. Prince, *Semiconductor Memories: A Handbook of Design, Manufacture and Application*, John Wiley, 1991.

Samsung Ex. 1097, p. 48
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

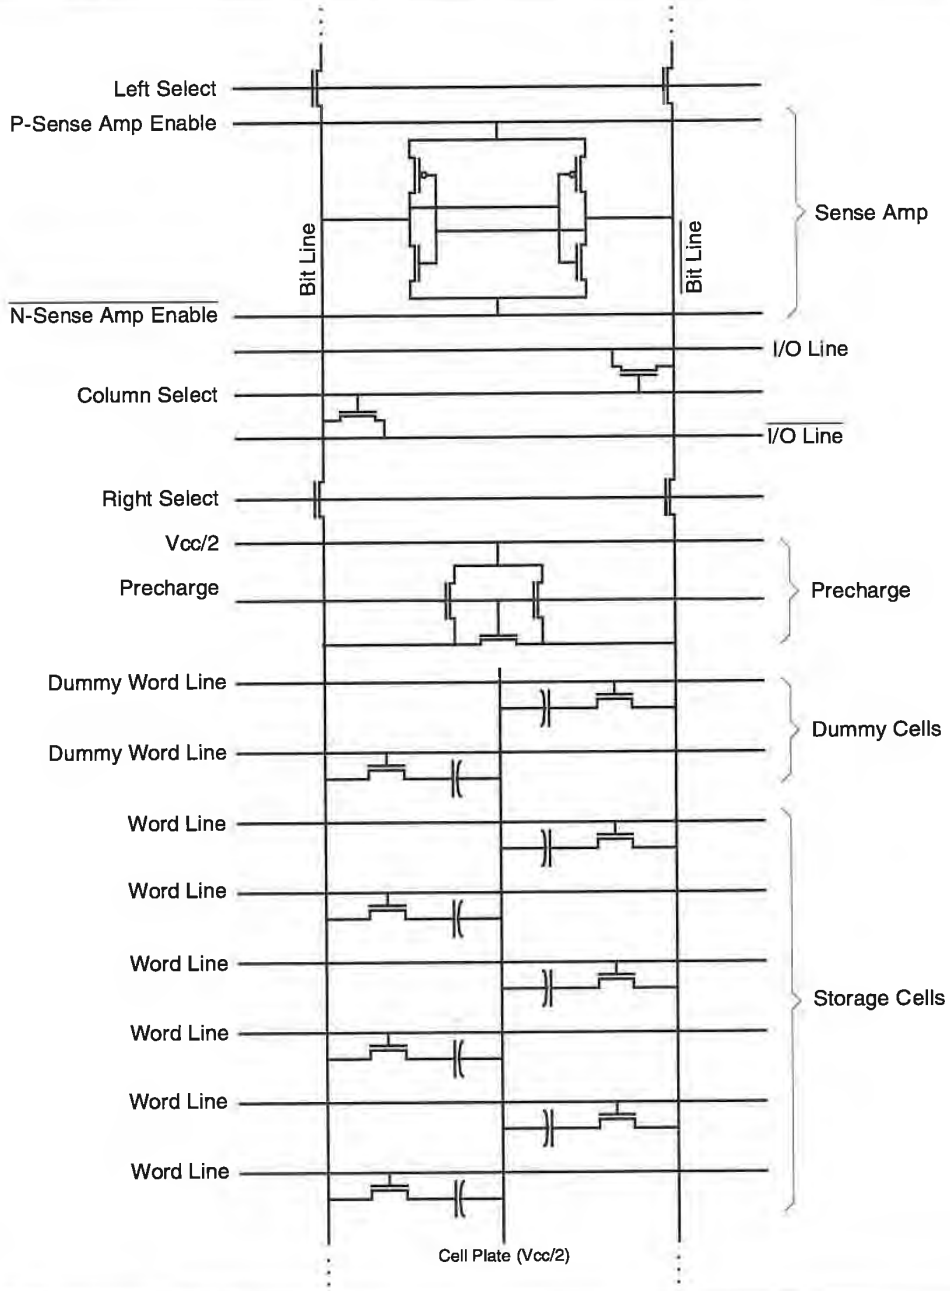


Figure 3-4 A DRAM column contains a differential sense amplifier, column select and precharge circuitry, and a collection of storage cells.

© 1996 Steven Przybylski

expensive, DRAM dice. Maximizing the capacitance of the storage capacitor is important because larger capacitors store more charge at a given voltage. Reading a cell consists of setting up a charge-sharing situation between the parasitic capacitance of the bit line and that of the storage cell. If the ratio of these capacitances is too large, then sensing the storage cell's impact on the bit line will be slower, more difficult, and more error prone. Typically, DRAM designers strive to keep the ratio of the bit line capacitance to the storage cell capacitance to 15:1 or less. Alternatively put, designers ensure that when a charged storage cell is connected to a bit line, the bit line moves by at least 50 mV and preferably by close to 100 mV.

Digital integrated circuits are electrically very noisy, and DRAMs are no exception. To reliably sense such a small change in a bit line, designers have used differential sense amplifiers since the 4-Kbit generation. The internal operation of a read is shown in Figure 3-5. Prior to the start of a row access, each pair of bit lines is equalized (precharged) to a midpoint voltage level. When the row address is decoded, one word line is raised. This connects the addressed storage capacitor to one bit line of every pair. The differential sense amplifier is then primarily amplifying the effect of the charge stored on the addressed storage capacitor. DRAM designers go to great lengths to ensure that the members of the bit line pairs are as nearly identical to each other as possible—both statically (capacitance and resistance) and dynamically (coupling to potential noise sources).

Figure 3-4 shows a row of dummy cells in addition to the many rows of storage cells. In previous generations of DRAMs, whenever a storage cell was activated on one bit line, a dummy cell was activated on the complement bit line. The purposes was to duplicate as much as possible of the storage cell's parasitic capacitance on the complement bit line. This meant that a larger proportion of the voltage difference between the two bit lines resulted from the value stored in the storage cell being accessed rather than from the physical properties of the storage cell itself. However, the setting of the cell plate at $V_{cc}/2$ has all but eliminated the need for dummy cells; they have been used only rarely since the 256-Kbit generation.

The sense amplifier is turned on only after the word line has gone high, the storage capacitor has shared its charge with the bit line, and the voltages of the storage cell and bit line have equalized. A typical sense amplifier is shown in Figure 3-6. The cross-coupled dynamic inverter pair amplifies the 50–150 mV difference between the bit lines and, if the word line is still high,

Samsung Ex. 1097, p. 50
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

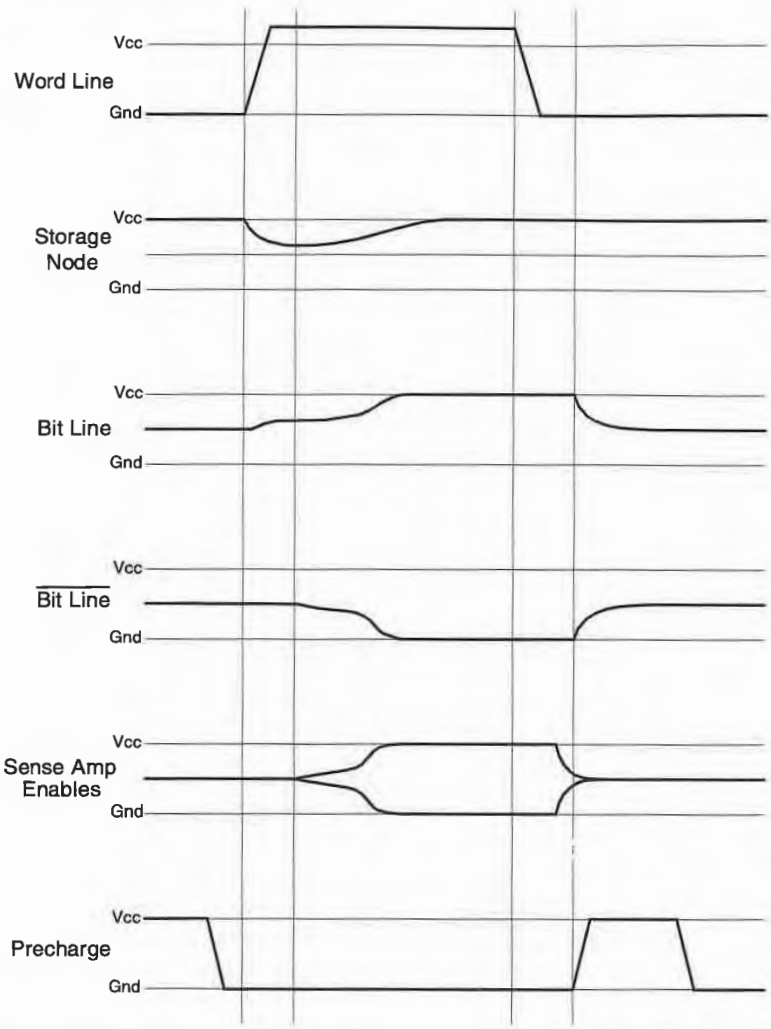


Figure 3-5 Reliable sensing of DRAM cells calls for careful sequencing of word line and sense amps.

automatically restores the charge on the storage capacitor. An important detail is that the coupling capacitances internal to the sense amplifiers are large enough to overwhelm the input signal if the sense amplifier is turned on hard as soon as the bit lines have stabilized. Consequently, the sense amps are generally turned on gently until the bit lines have separated somewhat. When there is no danger of overwhelming the voltage difference on the bit lines, the sense amp is turned on hard, and

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 51
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

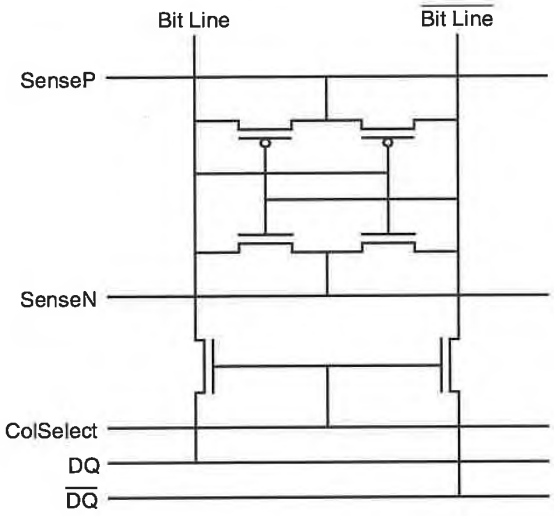


Figure 3-6 A typical sense amplifier consists of a pair of cross-coupled inverters.

the bit lines accelerate away from each other. In general, the need to wait for the bit lines to settle and the delicacy of the data signal at the sense amp are among the main factors limiting a DRAM's row access time.

Many storage bits per bit line

As mentioned above, the ratio of the storage cell capacitance to the bit line parasitic capacitance is crucial to the reliability and speed of bit line sensing. One way to improve the ratio is to limit the length of the bit lines and therefore their parasitic capacitance. However, limiting the number of storage cells on a bit line increases the number of sense amps in a device. Sense amplifiers and their related circuitry can take up significant area and power, so arbitrarily decreasing the number of bits per sense amp is undesirable.

Generally, there are 128 cells per bit line, with two bit lines for each sense amplifier. However, many designs improve the number of storage cells per sense amp by incorporating folded bit lines. In these designs, each of the sense amplifier's inputs are connected to two bit lines through pass transistors—called isolation devices—controlled by SelectLeft and SelectRight, only one of which is asserted during any given row access. This technique doubles the number of storage cells per sense amp from 256 to 512.

Samsung Ex. 1097, p. 52
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

This idea of connecting more than one pair of bit lines to a sense amplifier is taken further with segmented bit lines. In some designs, each primary bit line is connected by pass transistors to several secondary bit lines. The storage cells are on the secondary bit lines. The number of bits per sense amplifier can be increased because total bit line capacitance is divided up among the segments, and one of the segments is active at a time.

Column accesses transfer bits to the I/O pads

A column access consists of shunting one or more bits from the sense amplifiers within the core of the DRAM to the I/O buffers at the periphery and beyond. To get the bits to the I/O buffers quickly, a second set of sense amplifiers is used. The global I/O lines—which connect the primary sense amplifiers in the DRAM array to the secondary sense amplifiers in the I/O pads—are connected to the primary sense amplifiers through pass transistors. Generally, a final output multiplexer/demultiplexer connects the massive I/O drivers with the appropriate secondary sense amps.

Figure 3-2 shows the logical organization of a 4-Mbit DRAM. This representation implies 2K rows and 2K columns, or alternatively, 2K sense amps with 2K storage cells per bit line. Variations exist from design to design, but most 4-Mbit DRAMs have 16K differential sense amps with 128 bits on each bit line. Note that some of the row address bits are used to select banks within the DRAM as opposed to specific rows (see Figure 3-7). The internal structures of ×4 and ×8 DRAMs are similar (see Figure 3-8). Logically, the row and column addresses both decrease in width as the DRAM width increases. Physically, the number of sense amps remains the same independent of DRAM width, so the function of some row and column address bits varies according to the DRAM organization.

Consider two storage cells on adjacent bit lines that are connected to the same sense amplifier. A high voltage level in one cell will trip the sense amplifier one way and be read out as a logic 1. A high voltage level in the other cell will trip the sense amplifier the other way and be read out as a logic 0. In other words, half the bits in the DRAM are stored in true form, and half are stored in complement form. This is one of many design techniques that separate real DRAM implementations from the clean conceptual model of storage cells arranged in a monolithic square array. Other optimizations scatter logically adjacent bits across the DRAM die, making it difficult to deduce a bit's location from its address alone.

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 53
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

Most manufacturers design a single device that, with different bonding and metal options, implements ×1, ×4, and even ×8 versions of fast-page-mode, static-column-mode, and nibble-mode interfaces. If different speed grades, extended refresh periods, supply voltages, and minor variations are counted, several hundred different products can be derived from a single basic design. The number is even larger if all the different packaging options are factored in. The ease of this customization greatly simplifies production planning and management and speeds designs along the inevitable volume-yield curve to lower costs and prices.

3.4 DRAM Fabrication Technology

For many generations, DRAMs have been the IC fabrication technology driver of choice—the proving ground for the latest technology developments and the ever finer geometries at all layers. For example, in mid 1993, 0.6-μm geometries were state-of-the-art for logic processes, yet 16-Mbit DRAMs have been in production at 0.5 μm since early 1992. The reason for this technology leadership is twofold:

1. DRAMs are very regular, easily tested, and quickly diagnosed. Chip failures and poor yields in DRAMs—unlike those in logic-dominated ICs—are relatively easy to correlate to specific processing steps and manufacturing processes.

2. Generic DRAMs are subject to extreme price pressures. From the vendors' perspective, having the most cost-effective technology in place for each generation is imperative. This pushes the vendors to reach for new levels of technology early, so that by the time each generation of DRAM reaches significant volume, the vendors have accumulated substantial experience with the new processes. The accountants, of course, rightfully resist purchasing fabrication equipment until the sales volume to keep it fully utilized is imminent.

The crux of a DRAM technology is its storage capacitor. DRAM designers are forever caught between opposing needs for a physically small cell and a large capacitance. Despite ever decreasing geometries, storage cell capacitors have remained in the 30 to 50 femtofarads (fF) range since the 16-Kbit generation.

Samsung Ex. 1097, p. 54
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

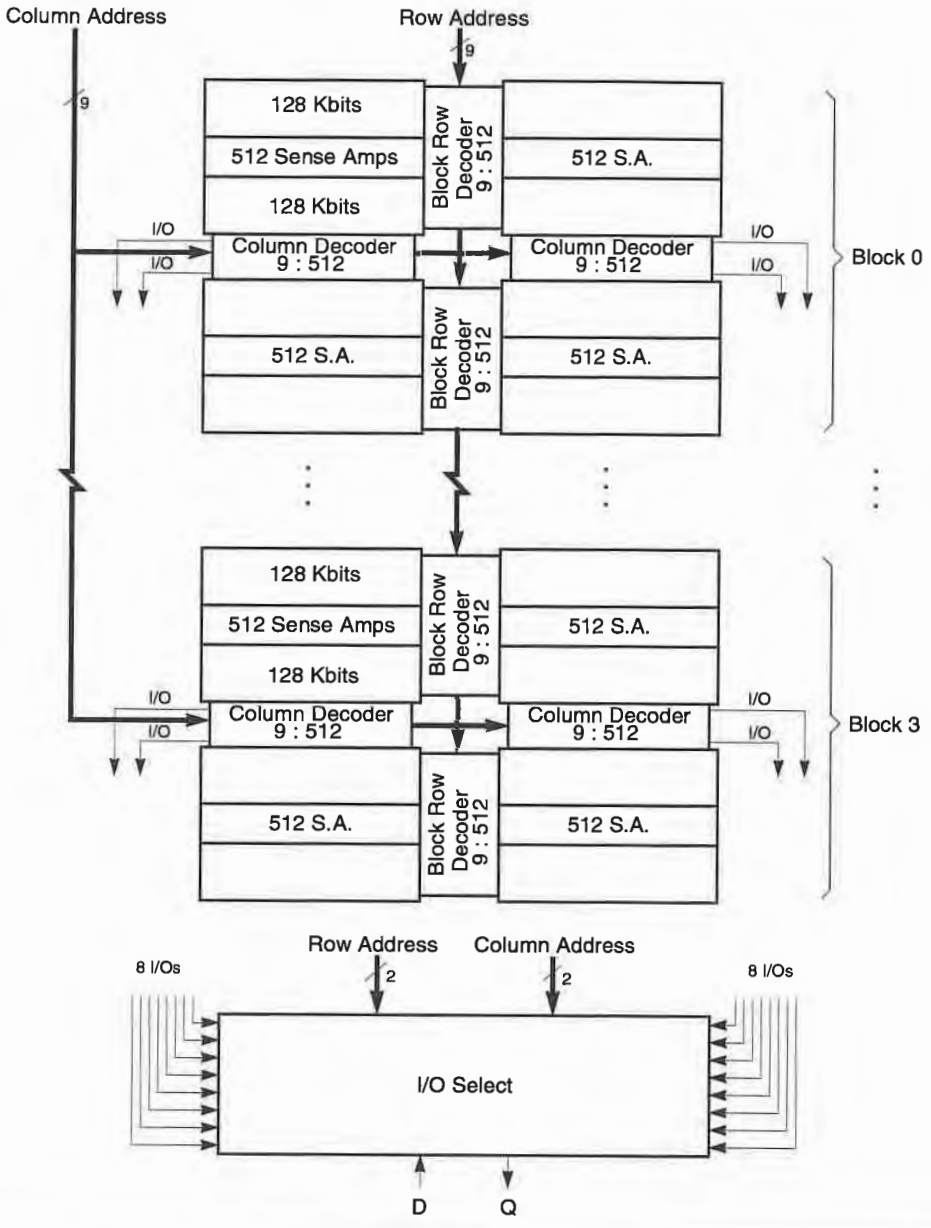


Figure 3-7 A 4M×1 DRAM typically consists of several independent blocks.

© 1996 Steven Przybylski

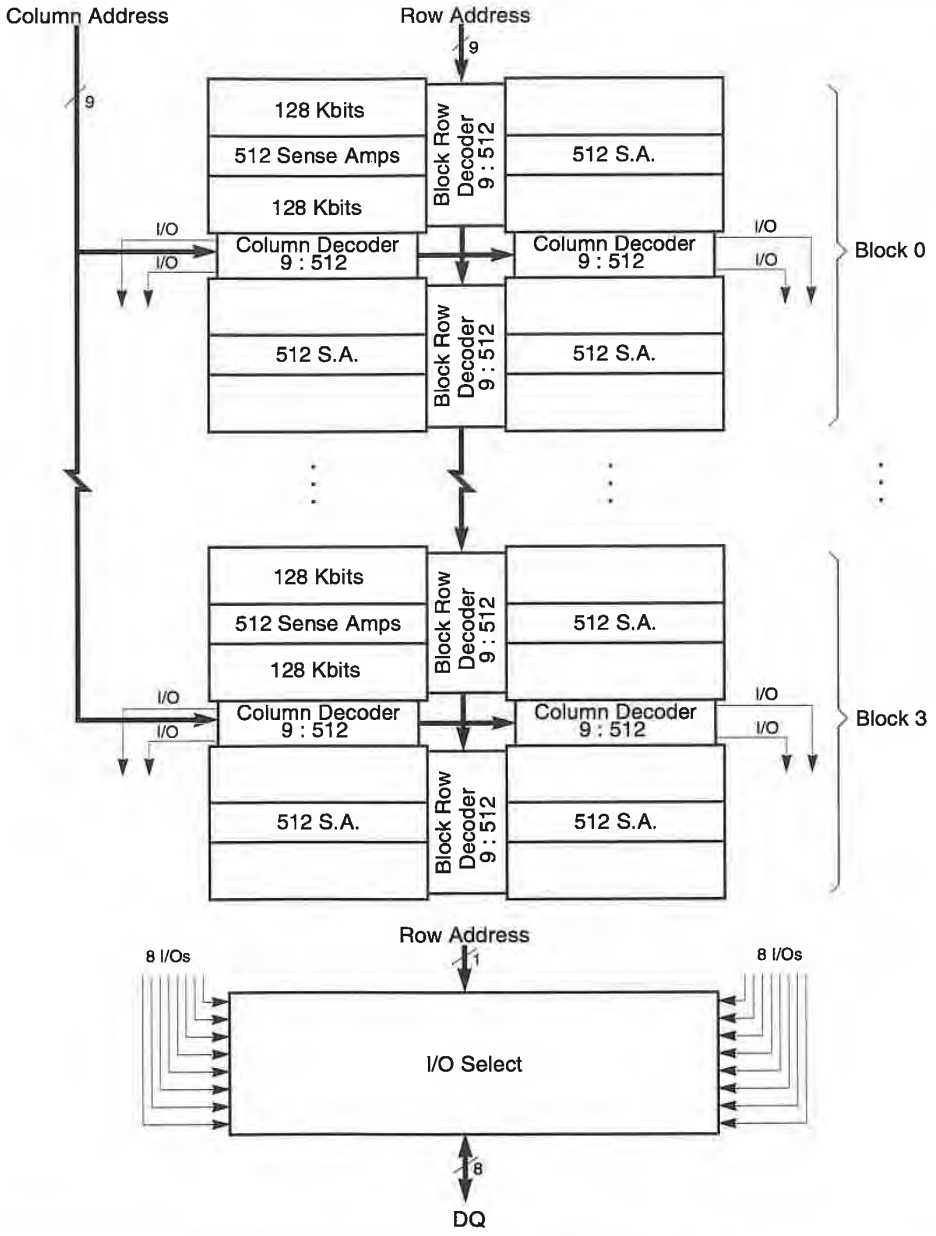


Figure 3-8 A 512K×8 DRAM can be derived from the same design as a 4M x 1 device.

Samsung Ex. 1097, p. 56
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

Although cell capacitance declines with the physical area of the storage capacitor, it has been helped over the years by declining oxide thickness and by the introduction of novel cell types that cram more capacitance into the smaller storage cells.

Three types of storage cells

Over time, three basic storage cells have been used:

1. The planar cell, used from the 4- and 16-Kbit generations through to the 1-Mbit generation
2. The trench-capacitor cell, used in some 1-, 4-, and some 16- and 64-Mbit designs
3. The stacked-capacitor cell, used in 4-, 16-, and most 64-Mbit parts

Planar storage cells

A schematic representation of a simple planar cell is shown in Figure 3-9. The word line runs into the page in a second level of polysilicon that forms the access transistor. In modern devices, the word line also runs in a second level of metal that occasionally makes contact with the polysilicon. This is called word-line strapping. The bit line, typically in metal, runs left to right across the page. The storage capacitor is formed between the diffusion to the right of the access transistor and the cell plate above the diffusion in polysilicon. The cell plate is shared by all the cells in a column. The plate is held at a constant voltage—typically one-half the supply voltage. The magnitude of the storage capacitor is linearly proportional to the area of the overlap of the cell plate and the diffusion underneath it, and inversely proportional to the oxide thickness. Cell area has declined from approximately 450 μm^2 in the 16-Kbit generation to roughly 40 μm^2 in the 1-Mbit generation. In the same period, gate oxides declined from about 1,000 Å to 120 Å.

Trench storage cells

A trench capacitor is shown in Figure 3-10. The trench capacitor is based on a curious characteristic of crystalline silicon. If the crystal lattice is oriented in a particular way, it is possible to etch precise, deep grooves into the surface of a silicon wafer. This exposes a surface area that is greater than the area of the trench opening. If a capacitor is built inside the trench, the increased surface area can be used to build a larger storage capacitor. Typically, as in planar cells, the storage node of the cell is built in diffusion around the outside of the trench. The trench is lined with a thin insulating oxide and filled in with polysilicon to form the cell plate.

Stacked storage cells

Unfortunately, trench capacitors are generally more difficult to build, especially at very fine minimum feature sizes.

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 57
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

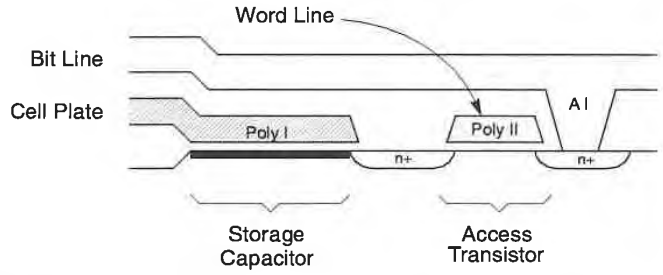


Figure 3-9 A cross-section view of a planar DRAM cell shows the capacitor oriented upward from diffusion to the cell plate in polysilicon.

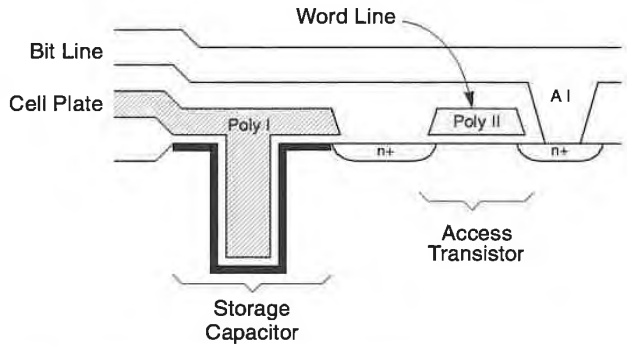


Figure 3-10 In a trench DRAM cell, the capacitor follows the boundary of a trench etched into the substrate.

Consequently, the industry is moving away from trenches and toward stacked cells. Stacked cells are more closely akin to the planar cells used in the earliest generations, except that both plates of the capacitor are built in polysilicon. Furthermore, since both plates are in polysilicon, they can be built on top of the word line and access transistor (see Figure 3-11). Increased area utilization, plus fringe and undulation effects, can yield storage capacitor values that nearly double the values of comparable planar cells.

As feature sizes continue to shrink, the quest for larger cell capacitors will force DRAM process engineers to pursue more

Samsung Ex. 1097, p. 58
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

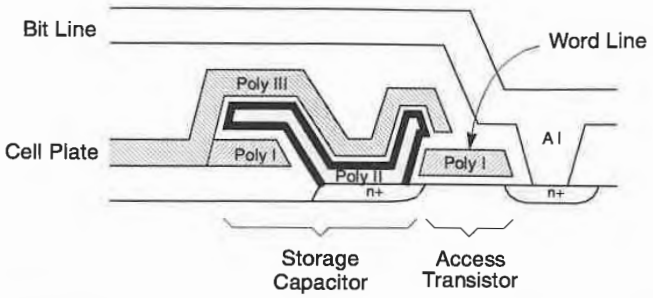


Figure 3-11 In a stacked DRAM cell, the capacitor hugs the top and bottom on a convoluted piece of polysilicon.

elaborate vertical structures, such as overlapped stacked cells, cylindrical cells, and multilayer stacked capacitors. Unfortunately, complex vertical structures require more processing steps to build. This in turn requires more equipment per fab line, which adds to the spiraling costs of DRAM factories and results in a higher cost per DRAM sold (see Section 2.2.2).

With the progression from planar to trench to stacked cells, DRAM processes have become more esoteric and specialized, and less related to the logic-oriented technologies used to build ASICs and processors. These latter devices need more layers of interconnect than are generally available in DRAM processes and have no need for the thin oxide layer used to form the storage capacitor. As DRAM processes become more different from logic processes, DRAMs are in jeopardy of losing their role as the technology leader where new processing ideas first make their presence felt.

Tight logic design also important

The storage capacitor is the primary focus in the design of a DRAM process. However, tight logic design is also important. One difficult aspect of DRAM architecture and layout is implementing the word-line drivers and sense amplifiers on the very fine row and column pitches of the DRAM cells. The area-efficient layout of these structures is important because they take up a significant proportion of the overall die area. Since the 64-Kbit generation, the proportion of the DRAM die taken up by storage cells has hovered in the 50–60% range: close to half the die area is taken up by I/O buffers and pads, wiring, decoders, sense amplifiers, word-line drivers, and other miscellaneous structures.

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 59
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

3.5 Advanced DRAM Concepts

3.5.1 Soft Errors

A significant and inevitable consequence of storing information as charges on capacitors is that the information is rather fragile. In addition to the danger of data loss from leakage or electrical noise, it is possible that an alpha particle[2] impacting the DRAM die will cause a bit to be misread. Because the DRAM is not physically damaged by the event, an error of this kind is called a **soft error**. In the 16- and 64-Kbit generations, soft errors caused by alpha particles generated within the DRAM packages were a significant problem. However, several techniques have reduced the number of alpha particles reaching the storage capacitors and thus the likelihood of an alpha particle's erasing a bit. These techniques include coating the die with a material that absorbs alpha particles, keeping the capacitance of the storage capacitor greater than 50 fF, and using well structures that absorb most of the generated electron-hole pairs. Consequently, soft error rates have declined dramatically in recent generations and no longer pose a significant problem for most system designers.

The standard measure of failure rates is the number of **Failures in Time (FITs)**. Figure 3-12 shows the soft error rate—the expected number of soft errors per bit in the DRAM for every billion operating hours—across six generations of DRAMs. The failure rate for a system is obtained by aggregating the failure rates for all components in the system. Figure 3-13 shows the expected number of errors over 10 years of use for a 2-Mbyte memory system. By the 4-Mbit generation, that number dropped to nearly insignificant levels. System-level approaches to dealing with soft errors are discussed in Section 4.2.5.

3.5.2 Supply Voltages

Another fundamental tension in DRAM design involves the selection of the supply voltage. Designers have increased storage cell capacitance by decreasing oxide thickness as much as possible. Thinner oxides, however, are more susceptible to temporary and permanent damage from high voltages. Electrical

2. An alpha particle is a helium atom stripped of its electrons. The low-level alpha radiation that surrounds us is generally harmless.

Samsung Ex. 1097, p. 60
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

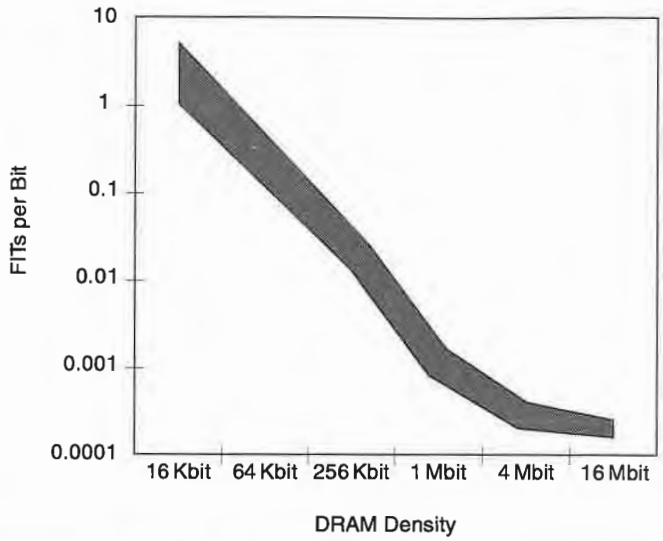


Figure 3-12 The soft error rate declined dramatically over six generations. (Source: *Micron 1993 DRAM Data Book*).

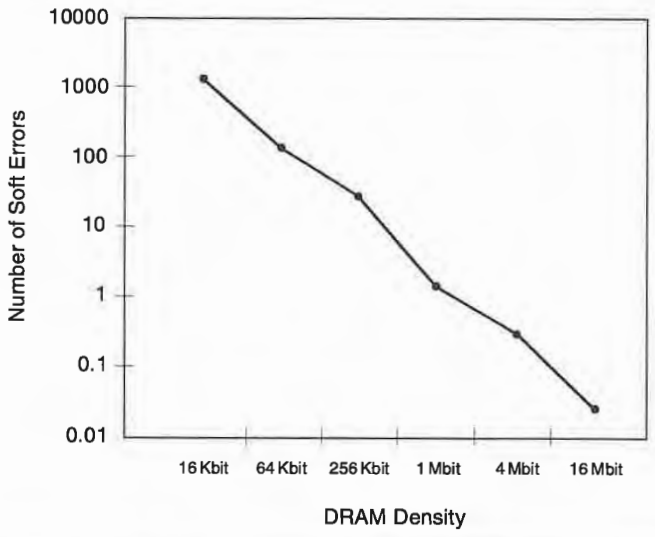


Figure 3-13 The number of soft errors over 10 years of continuous use for a 2-Mbyte memory system has become essentially insignificant. (Source: *Micron 1993 DRAM Data Book*).

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 61
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

stress on the thin oxide and hot electron effects in the access transistor are both minimized by reducing internal voltages. Unfortunately, for many years, 5 V was pervasively used as the system-wide supply voltage.

The net result is that the internal supply voltage, which started at 12 V in the 4- and 16-Kbit levels, settled at 5 V for many generations. At the 4- and 16-Mbit levels, however, designers were forced to reduce the internal supply voltage while retaining the 5 V supply. Practically all modern DRAMs have on-chip voltage regulators. In the 4-Mbit generation, these voltage generators transform the 5 V supply to roughly 4 V internally; in the 16-Mbit generation they transform the supply voltage to 3.3 V. The 16-Mbit DRAMs that are specified with 3.3 V supply voltage are actually made from the same dice as the 5 V parts, but the on-chip voltage converter is bypassed.

As the industry moves to the 64- and 256-Mbit levels, internal supply voltages will drop further, to 2.5–1.5 V. At the chip interface, however, 3.3 V will remain the standard for another generation or two, at which point a shift down to 2.5 V will occur.

3.5.3 Testing and Redundancy in DRAMs

One of the biggest problems associated with the exponential growth in DRAM density is the corresponding exponential growth in testing time. Customers and vendors alike have rightly placed a great deal of emphasis on DRAM quality. All vendors exhaustively test every DRAM prior to shipping it to a customer. These tests not only determine the speed of the device and verify the simple functioning of every bit, but also ensure that data is retained under worst case DC and AC conditions. Typically, a good DRAM die is tested once before assembly and again during or after burn-in and temperature cycling. At times, a third test, after assembly but before burn-in, is also warranted. Only at final test are the DRAMs **speed graded** (also called binned) and marked. In addition, a fraction of all DRAMs are statistically sampled for more exhaustive accelerated life cycle and reliability tests. Full details of each company's test procedures and quality assurance programs can be found in its data books.

As the number of bits per DRAM grows, so does the amount of time needed to test each device. The increased costs of memory testers and the threat of an exponential growth in test time have caused DRAM vendors to implement a variety of parallel

Samsung Ex. 1097, p. 62
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

and automated test modes in their DRAM designs. These modes reduce the test time for each DRAM from minutes or hours to seconds. In general, the idea of these modes is to test more than one bit at a time, or to off load all test generation and checking onto the chip. These test modes are occasionally used by large system vendors to verify the quality of parts coming from their DRAM suppliers.

Most DRAM consumers are surprised to learn that testing accounts for up to 50% of the vendor's production costs. The cost of the testers, handling equipment, and burn-in ovens must be amortized over all the good dice produced. As a DRAM design matures and manufacturing engineers understand its failure modes, the amount of time needed to test each part declines significantly. Reduced test times and higher yields combine to dramatically reduce the cost of producing a working device over the course of a design's lifetime.

Wide DRAMs are more expensive to test than narrow ones. Fixed width testers can simultaneously test more narrow parts than wide parts. This is the primary reason for the cost premium demanded for wide DRAMs. Section 11.3.1 discusses testing and its impact on the production costs in more depth.

DRAM redundancy

An important part of wafer probe testing is the isolation of any stuck or faulty bits in an otherwise good die. All DRAMs have a degree of built-in **redundancy**. Individual rows or columns that are not functional can be replaced by other rows and columns that are redundant by using a laser programmer that, for example, vaporizes polysilicon fuses at predetermined locations on the die. A small amount of redundancy can dramatically improve yield. The numbers of rows and columns that are appropriate depends on:

- The device's raw yield
- The area overhead associated with the extra row and column decoders, storage cells, sense amps, and so forth
- The fraction of failures that can be effectively removed by mapping out one or more offending rows and/or columns

Generally, a small number of extra rows and columns are adequate to multiply the raw yield severalfold. As a DRAM matures and is redesigned onto newer processes, the amount of redundancy is reduced to minimize the die size and to adjust for increased yields.

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 63
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

3.5.4 Speed-Enhancing Techniques

Automatic Transition Detection

One technique commonly used to speed access times from a valid address is Automatic Transition Detection (ATD). Logically, the row or column address is presented a setup time before a RAS or CAS edge. Only after the falling edge of $\overline{RAS}$ or $\overline{CAS}$ is the row or column access started. In fast-page-mode DRAMs, each of the address input pins is connected to a simple circuit that detects when the address pins have changed. When any one of the address pins changes, a column access is begun in advance of the edge in $\overline{CAS}$. To speed the access, the new address is automatically propagated to the column decoders. ATD allows DRAMs to behave as if high-power static circuitry is used through the column address decode path; in reality, lower power dynamic circuitry is used throughout the decoders.

The adoption of automatic transition detection facilitated the transition from page mode to fast-page mode in the late 1980s. ATD is also responsible for the signficant difference between the column access time from the column address and the column access time from the falling of $\overline{CAS}$ that is typical in fast-page-mode DRAMs. The transition from page mode to fast-page mode accounts for much of the rapid decline in column access times and especially column cycle times during that period (see Figure 2-2).

Boosted word lines

Figure 3-5 on page 51 shows internal signals during a row access. In particular, it shows that the word line goes above V_{cc}. Boosting is a common circuit technique that generates voltages greater than the supply voltage. Without boosting the word line, a threshold drop across the access transistor during a write prevents the storage capacitor from getting above about one volt below V_{cc}. A reduced storage cell voltage translates to a smaller stored charge, a longer row access time, a shorter refresh period, and poorer reliability of read operations.

Current sense amplifiers for global I/O lines

As DRAMs have become physically larger, the delay in driving the global I/O lines from the primary sense amplifiers to the I/O pads has been growing as a proportion of the column access time. To combat this problem, many DRAM designers are choosing current sense amplifiers rather than the traditional voltage-sensing variety (see Figure 3-6 on page 52). Current sense amplifiers at the I/O pads can detect changes in the current flowing in the global I/O lines from the DRAM core much more quickly than the traditional sense amplifiers can sense the voltage difference that derives from that current change.

Samsung Ex. 1097, p. 64
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

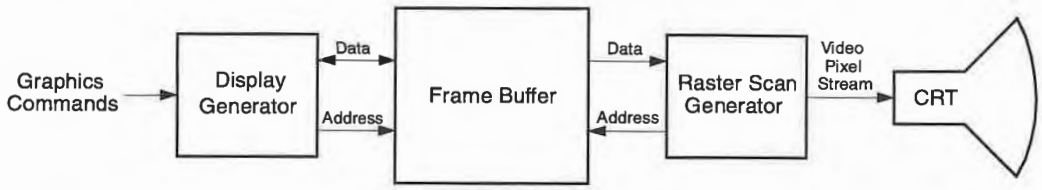


Figure 3-14 This block diagram of a typical graphics controller shows the two sources of references to the frame buffer.

Output drive control Ground bounce is a very serious problem in wide DRAMs. With as many as eight outputs per power and ground pair, the on-chip ground level can be raised substantially by all drivers switching in the same direction simultaneously. The problem is to switch the output voltage as quickly as possible in the face of inductive wire bonds and package leads while keeping ground bounce to tolerable levels. To achieve this, modern output drivers moderate the output drive to limit the rate of change of output current. The output transition time is minimized by turning on the final output transistors in incremental stages.

3.6 Video RAMs (VRAMs)

Figure 3-14 shows a block diagram of a graphics controller. Video commands are received from the processor by the graphics engine, which translates them into a series of pixel operations. These operations are performed on the frame buffer, which holds an exact copy of the screen's current contents. Concurrent with the graphics engine's interpretation of video commands, a separate and relatively independent process reads out the contents of the **frame buffer** and sends the contents to the CRT. Consequently, there are two logically independent ports into the frame buffer RAM: a standard read-write port for access by the graphics engine and a read-only port for dumping the frame buffer to the display.

When generic DRAMs are used in video applications, these two logical ports must be multiplexed over the single random

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 65
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

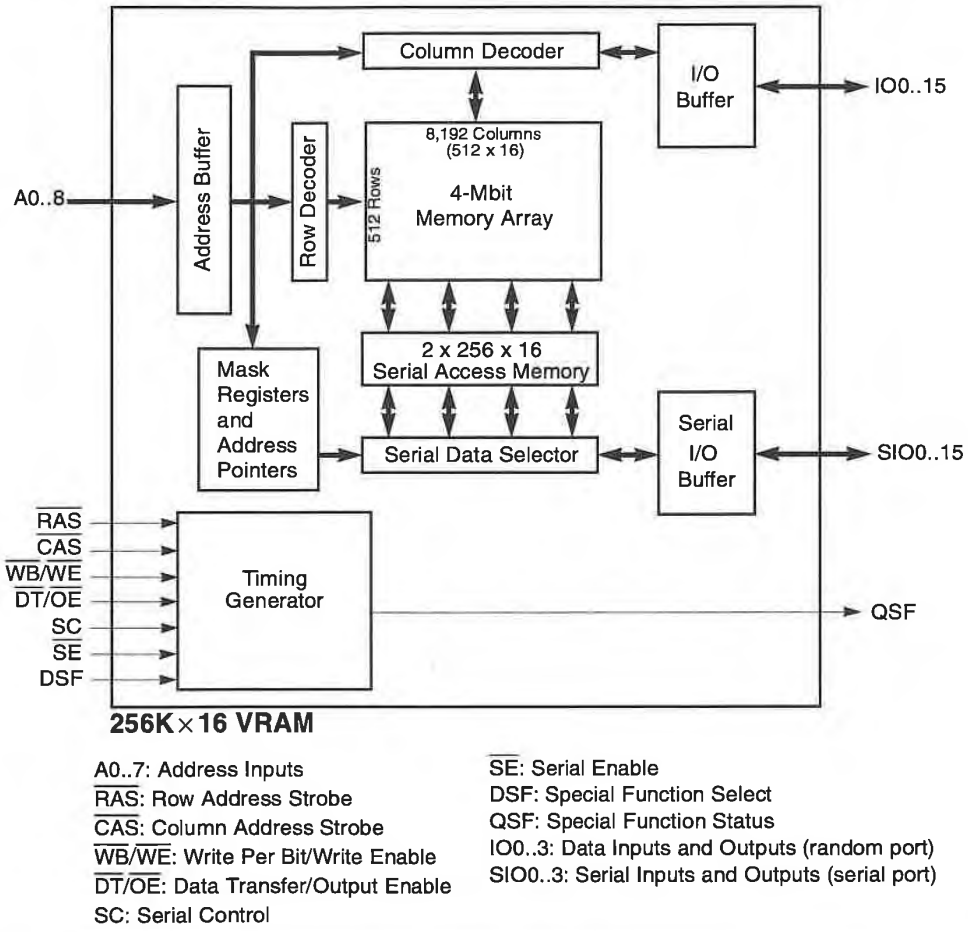


Figure 3-15 This block diagram of a typical 256K×16 VRAM device shows both the random and serial ports.

read-write port available. On the other hand, **Video DRAMs (VRAMs)** are specially suited for frame buffers and other video applications because they have two ports corresponding to the two data streams: a random read-write port that is an augmented version of the standard RAS/CAS interface and a serial read-only or read-write port for displaying the screen. Figure 3-15 shows the internal organization of a 256K×4 VRAM.

The high-speed serial port is connected to a shift register—the Serial Access Memory (SAM)—that is built into the sense amplifiers of the DRAM array. This shift register, which

Samsung Ex. 1097, p. 66
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

typically contains a few kilobits, is loaded in parallel from the DRAM array with a single row access. While the contents of the shift register are streaming out to the display, the random access port is available for read and write accesses by the display generator or CPU.

JEDEC-standardized features

In order to establish a degree of compatibility among the VRAMs produced by different vendors, JEDEC has standardized a set of mandatory and optional features for 256K×4 and 128K×8 VRAMs. Unfortunately, the optional features include most of the useful operations involving the serial port. The optional features are not consistently implemented by all the vendors; in some cases, vendors have omitted some functions, modified others, and introduced still others. The end result is a marketplace even more confusing and treacherous than the generic DRAM marketplace.

Among the JEDEC-standardized features are the normal random port read and write operations, reads and writes between the SAM and DRAM, and the normal selection of DRAM refresh operations. DRAM write options include:

- Per-bit masked writes in which the mask is specified on every write or at the start of a string of writes
- Block writes that cause a 4- or 8-bit color register to be written into four adjacent column addresses

With block writes, the data pins are used to specify which of the four locations are written. Block writes can also have per-bit masks. Of course, operations to load the color and bit-mask registers are included in the optional set. Finally, all recent VRAMs have their SAMs split into two logical halves. This split allows serial transfers to continue out one half of the SAM while the other half is being parallel-loaded from, or parallel-unloaded to, the DRAM.

In addition to the common 256K×4 and 128K×8 organizations, 256K×8 and 256K×16 VRAMs are available or proposed. There also exist a set of triple-ported devices that include one random port and two serial ports. These parts are for use in high-performance GUI accelerators and network interface applications characterized by serial input and output with a need for random access.

VRAM history

When introduced at the 64-Kbit level, VRAMs constituted only a tiny fraction of the overall DRAM market. Sales grew rapidly: at their peak, VRAMs constituted 8–10% of all dynamic RAMs

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 67
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

sold. VRAMs traditionally have been significantly more expensive per bit than DRAMs. Typical price premiums for VRAMs have ranged between 60 and 100% over comparably organized DRAMs. This price per bit premium is primarily due to the large die size penalty, which in turn stems from the incorporation of the shift register into the sense amplifiers. Some VRAMs are nearly twice as large as comparably organized DRAMs, though a more typical area penalty is 40%. Also, as specialty parts with fewer sourcing options and less consistency across vendors, VRAMs have benefited less from the price reduction pressures of direct competition by many manufacturers. To a lesser extent, the price premium also derives from the larger package needed to accommodate both access ports.

The high price of VRAMs, combined with the extreme cost sensitivity of the PC graphics marketplace, drove most PC GUI accelerators to use fast-page-mode DRAMs instead of VRAMs. Since the mid 1980s, the added cost associated with a VRAM frame buffer could only be justified in the high-end PC and workstation graphics systems. In these systems, the improved graphics performance warranted the increased costs.

During 1995, the price premium commanded by 4-Mbit VRAMs dropped substantially, to about 40% over 256K×16 DRAMs. This drop resulted from the actions of a few vendors that wanted to significantly increase their market share in this arena. The price premiums of the 1- and 2-Mbit VRAMs came down as well, but only indirectly and not as dramatically as those of the 4-Mbit VRAMs.

The twilight years

The introduction of several new DRAM organizations targeted specifically at graphics applications has been motivated by several factors: the recent growth in the graphics marketplace, the increasing bandwidth demands that the accompany growing screen resolution and color depth, and higher margins available to the vendors of specialty DRAMs. These new organizations include Synchronous Graphics RAMs (SGRAMs), Multi-bank DRAMs (MDRAMs), and Window RAMs (WRAMs). The attractiveness of the new video and general purpose DRAM organizations has caused several major manufacturers to abandon the development of their 4-Mbit VRAMs. As of 4Q1995, no manufacturer is proposing an 8-Mbit VRAM. During 1996 and 1997, VRAMs, as currently defined, will be substantially replaced by these and other new devices. In the long term, VRAMs are unlikely to survive the newly diversified DRAM marketplace.

Samsung Ex. 1097, p. 68
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

4 Memory System Design

This chapter is concerned with memory systems: how DRAMs are put together in a larger unit that responds to a CPU's read and write requests. The new DRAMs are interesting as individual circuits; more important, however, are the memory systems these new devices can form. The characteristics of these memory systems determine the utility and effectiveness of the new DRAM organizations.

4.1 New In The Second Edition

The changes to this chapter between the first and second editions include minor updates to reflect the changing structure of computer systems. In particular, many of the examples have been changed to reflect Pentium-based systems.

New descriptions of the terminology of latency, bandwidth, ranks, banks, and sub-banks have been added. Section 4.2.3 defined and explains the difference between peak, sustained and sustainable bandwidths and fill frequencies. Descriptions of the different burst transfer orders have also been added.

Section 4.8 also adds a description of Unified Memory Architecture (UMA) PCs and discusses the pros and cons of the different styles of UMA system.

Samsung Ex. 1097, p. 69
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

4.2 Basic Concepts

What memory systems do

The role of a memory system is twofold:

- Reliable and consistent data storage
- Quick responses to the read and write requests of CPUs and I/O devices

Requests come in sizes ranging from a single byte to many words.[1] The definition of *reliable data storage* varies from system to system: in some systems, an occasional dropped bit is acceptable, while other systems require data loss rates at or below once per millennia even in the face of power outages. *Consistent data storage* refers to the memory system's behavior when dealing with reads and writes to the same address, potentially from different sources. For example, the CPU may perform a write to a particular location, followed shortly by a read of the same address. For the memory system to respond with anything other than the written value would present an inconsistent, and potentially fatal, view of memory.

Memory system design is primarily constrained by physical restrictions and bandwidth requirements. However, the DRAMs themselves account for a substantial portion of most systems' cost; this portion of system cost has been substantially beyond the designer's control. Memory system designers can only impact overall system cost somewhat, by selecting one variety of DRAM over another and by trading off the cost of the non-DRAM components with the memory system performance. Frequently, the memory system designer is constrained by the need to facilitate unnaturally small memory systems to minimize the price of the entry-level system.

4.2.1 Single In-Line Memory Modules (SIMMs)

DRAMs come in a variety of packages (see Section 5.2.1). In addition, memory vendors sell a wide variety of **Single In-Line Memory Modules (SIMMs)**. SIMMs, small printed circuit boards with up to 20 DRAMs mounted on them, have 30, 72, or more signals accessible via an edge connector. SIMMs are inserted into small connectors that are soldered onto a larger printed circuit board, or motherboard. Because SIMMs are socketed, they are inherently replaceable and upgradable.

1. The definition of a word varies across the industry. Most commonly, a word is defined as 16 or 32 bits. In this report, unless otherwise specified, a word is 32 bits long.

Samsung Ex. 1097, p. 70
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

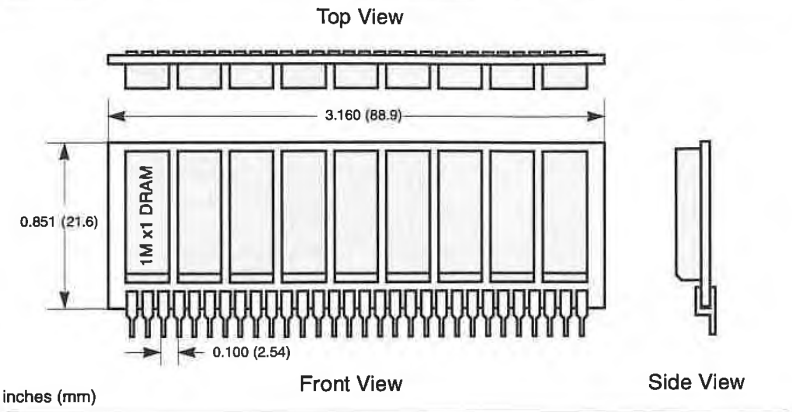


Figure 4-1 Standard 30-pin SIMMs are common in many low-end and midrange memory systems.

The most common SIMMs are ×8 and ×9 organizations on 3.5" long by 0.5 to 1.0" high boards, and ×32 and ×36 organizations on 4.25" long by 1.00–1.45" high boards (see Figure 4-1). Depths of 256K through 16M are widely available. Consequently, SIMMs come in all capacities from 1 to 40 Mbytes and more. They can have DRAMs mounted on one or both sides. The form factors and pinouts of the more common SIMMs are standardized, making them commodity parts.

Figures 4-2 through 4-4 show the internal structures of three SIMMs:

- A small, simple ×8 part
- A ×36 part with two banks and per-byte $\overline{\text{CAS}}$ signals
- A custom ×72 buffered SIMM

The simplest SIMMs just forward the DRAM signals to the SIMM edge. SIMMs of intermediate complexity often logically partition the DRAMs into banks and slices by routing individual $\overline{\text{RAS}}$ and $\overline{\text{CAS}}$ signals to groups of DRAMs. The ×36 part shown in Figure 4-3 has two banks, activated by $\overline{\text{RAS0}}$, $\overline{\text{RAS2}}$ and $\overline{\text{RAS1}}$, $\overline{\text{RAS3}}$. The $\overline{\text{CAS}}$ signals are arranged to enable byte reads and writes. All the standardized 72-pin SIMMs contain speed and size information on four of the pins. This allows the memory controller to adjust its behavior to meet the requirements of the particular SIMM in each socket.

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 71
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

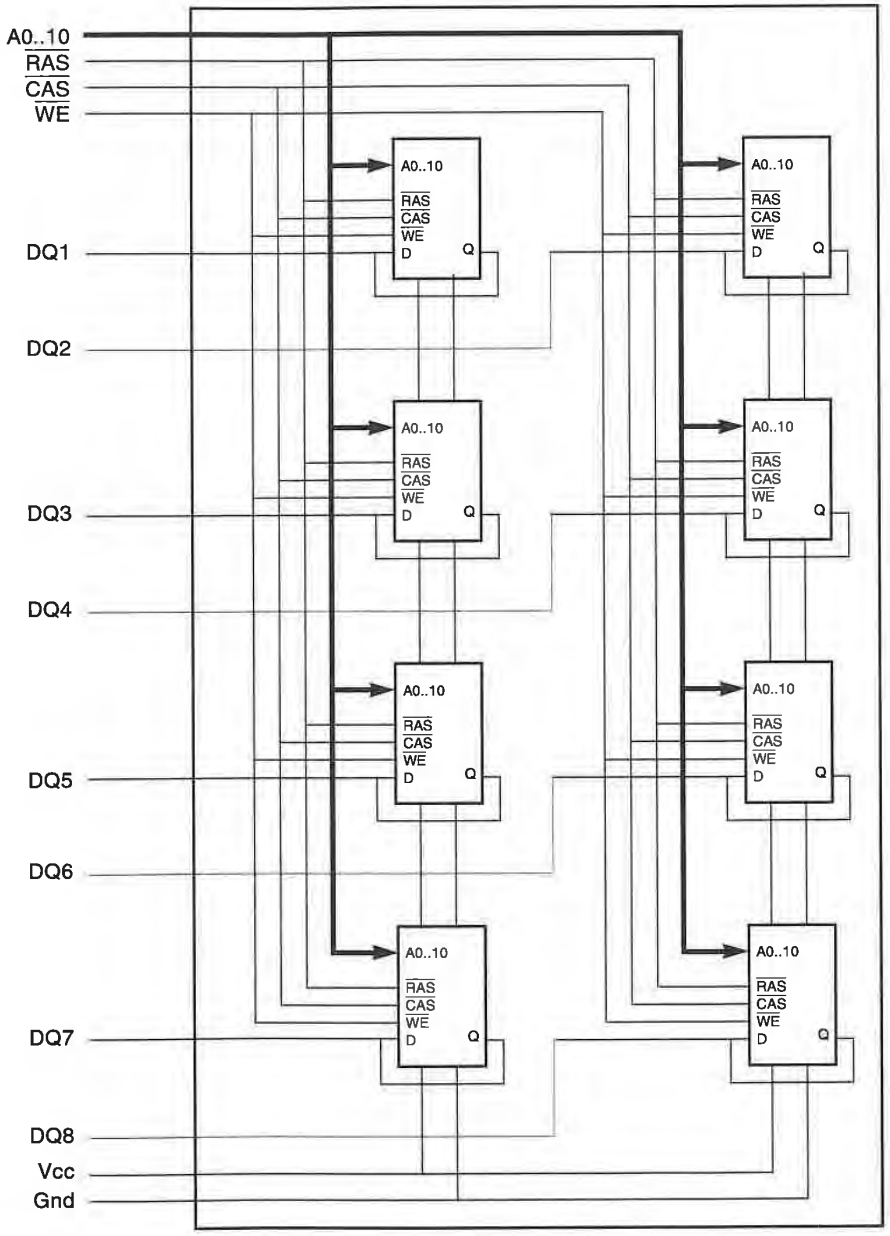


Figure 4-2 A 4-Mbyte byte-wide SIMM can be constructed from eight 4M×1 DRAMs.

Samsung Ex. 1097, p. 72
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

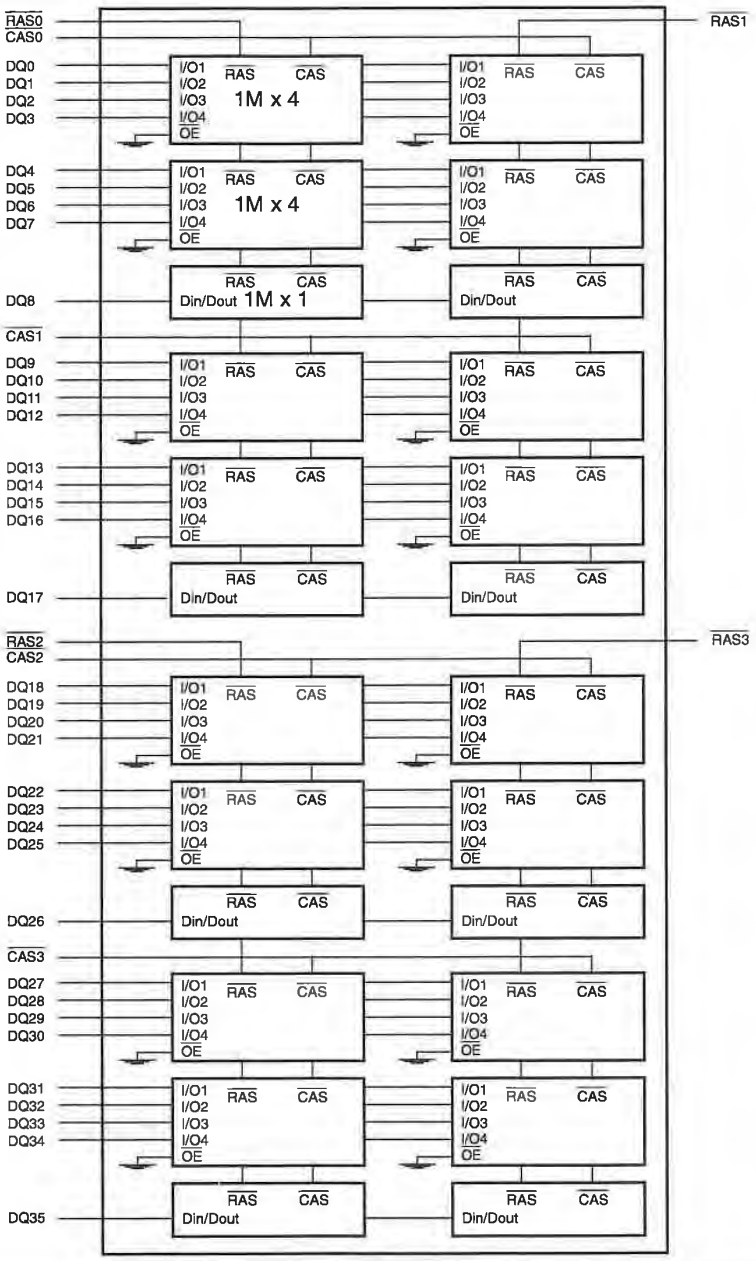


Figure 4-3 A 2M×36 SIMM can be constructed from sixteen 4-Mbit DRAMs and eight 1-Mbit DRAMs, configured in two banks. The 10-bit address bus going to all twenty-four DRAMs is not shown.

© 1996 Steven Przybylski

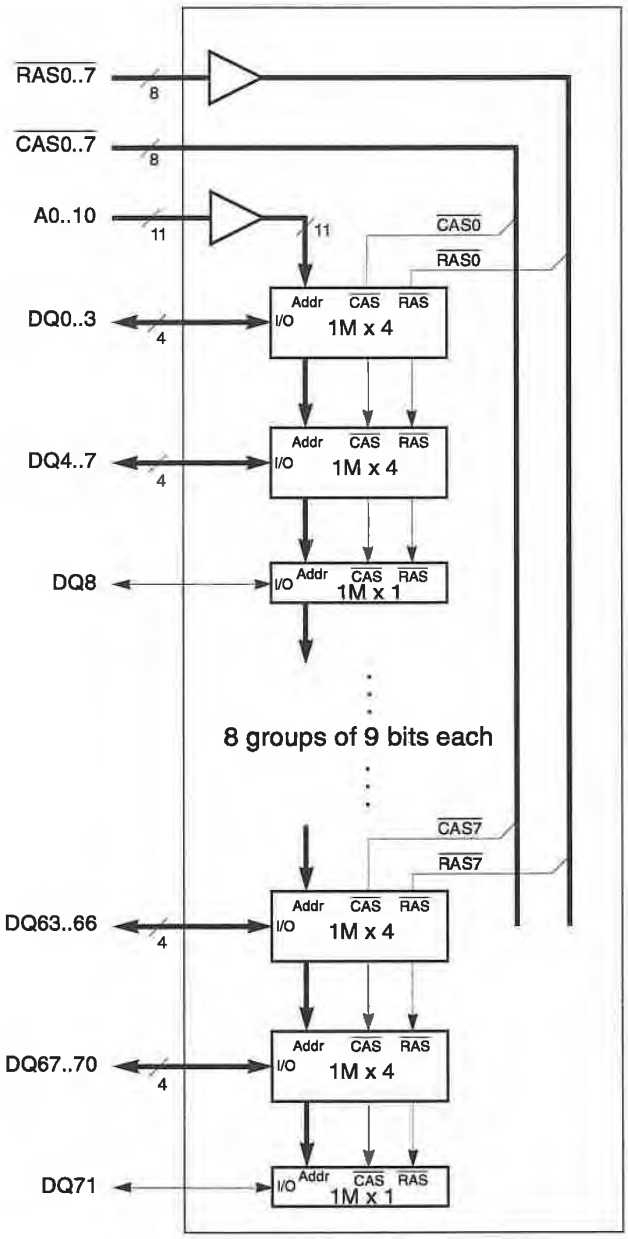


Figure 4-4 On this custom buffered ×72 SIMM, the address lines and most of the control lines are locally regenerated.

Samsung Ex. 1097, p. 74
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

SIMMs were conceived at a time when DRAMs and, consequently, memory systems, were slow. The sole goal of their developers was to facilitate cheap and physically dense memory expansion. The poor electrical characteristics of SIMMs and SIMM sockets are becoming severe impediments to the higher speed operation of memory systems. This problem will be further aggravated as the new generation of DRAMs, with their high frequency interfaces, become common. This problem may be partially alleviated by adding flow-through or registered buffers to SIMMs.

The buffered SIMM in Figure 4-4 has flow-through buffers on the address and control lines to limit the load presented by the SIMM to the memory controller. Limiting the load helps the memory system's overall performance by improving signal quality on the motherboard and decreasing the elapsed time between the memory controller outputs and the DRAM inputs. Adding buffers to the motherboard minimizes variability in the load placed on the primary control signals from the memory controller. Reduced variability in the load improves the signal quality of the motherboard signals and helps increase the frequency of reliable operation.

A number of the newer memory modules come in buffered and unbuffered versions. In many cases, the buffered versions include repeaters on all the control lines except the $\overline{\text{RAS}}$ lines. This is done because of the structure of PC memory systems (see section 4.4). In PCs, the $\overline{\text{RAS}}$ signals are used as bank-enables and so they are naturally low-fanout: if the system has only a few DRAMs per bank, buffers might actually slow the propagation of $\overline{\text{RAS}}$ to the DRAMs. On the other hand, $\overline{\text{CAS}}$, $\overline{\text{WE}}$, and the address lines are sent to all banks, and so they can have large fanouts in multibank systems. Buffering these signals improves the total time to propagate the control signals to the DRAMs.

For some of the new synchronous devices, registered memory modules are of value. These devices allow the total address and control transit time from the memory controller to the DRAMs to be split across two cycles. For large memory systems, this split can facilitate a high clock rate, and therefore a higher peak bandwidth; however, this is done at the expense of an additional cycle of latency for all operations.

In general, for all small- and medium-sized memory systems—including those found in PCs—unbuffered memory modules

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 75
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

typically will yield the highest operating frequency and provide the best overall performance.

New memory modules standardized

Recently, a set of new memory modules has been standardized. Instead of leads along only one side of the bottom edge, these memory modules have leads on both sides of the small printed circuit board. They are therefore called Dual In-Line Memory Modules (DIMMs). Although the widely used memory modules in the ×8 through ×32 widths are SIMMs, at the ×64 width, all modules are DIMMs. Because the pins are on both sides of the module, the 168-pin DIMMs are only about 1 inch wider than the older ×32 72-pin SIMMs currently in use. The 168-pin ×64 DIMMs are poised to become the next widely used memory module for PCs.

Another variety of DIMM achieving wide adoption is called Small-Outline Memory Modules (SOMMs). The density of contacts along the bottom edge of SOMMs is twice that of SIMMs and DIMMs. This greater density gives ×32 and ×36 SOMMs half the width of the traditional 72-pin SIMMs. SOMMs with DRAMs mounted on both sides are well suited to the tight packaging constraints in notebook computers and other portable applications.

4.2.2 Memory System Organization: Granularity and Expandability

Figure 4-5 is a block diagram of a typical midrange memory system. Such a system consists of a **memory controller** that 1) receives commands and addresses from the rest of the system and 2) directly or indirectly drives the address and control lines of the DRAMs. The data flowing to or from the DRAMs usually passes through bidirectional buffers that are also controlled by the memory controller. Alternatively, especially in high-performance systems, the data itself may pass through the controller on the way to or from the DRAMs.

The memory array itself is organized as a number of **banks** with the same widths but varying depths. A read operation begins when the memory controller receives a memory-read command and determines whether the address maps to one of its banks. If the bank being accessed currently has $\overline{\text{RAS}}$ held low, and if the row currently available within the DRAMs is the one that contains the requested address, the data can be quickly retrieved from the DRAMs via a CAS access. If the wrong row is

Samsung Ex. 1097, p. 76
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

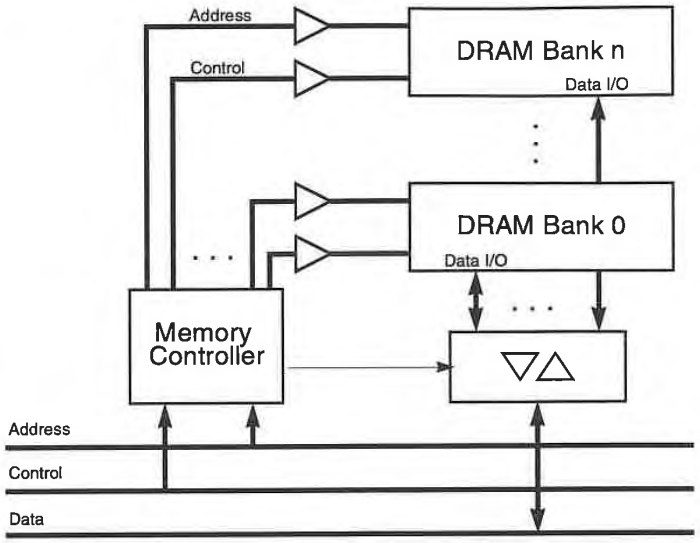


Figure 4-5 Typical memory system for a low-end to midrange computer system.

currently available, then $\overline{\text{RAS}}$ must be raised and kept high for a row precharge time before the access of the desired row can begin. Ultimately, as the data comes out of the DRAMs, the memory controller communicates to the processor that arrival of the data is imminent and directs the buffers to appropriately forward the requested data.

Writes happen similarly, except that the data to be written is usually received by the memory controller simultaneously with, or immediately after, the arrival of the address. This can be in advance of when the DRAMs want to see the data; as a result, the data is often temporarily stored or delayed in the memory controller or transceivers until the DRAMs are ready to perform the appropriate CAS write access.

Of banks and ranks

Until recently, all DRAMs logically consisted of a single homogenous block of memory. Several of the new DRAM architectures are internally divided into banks. This has introduced a degree of confusion: the word *bank* can mean either a physical or a logical grouping of bits: groupings that can be different. To rectify this confusion, a new term, **rank**, has been introduced to refer to the physical structure.

© 1996 Steven Przybylski

Figure 4-6 illustrates the relationships among banks, ranks, DIMMs, and SIMMs. It shows a high-performance 64-bit memory system formed by two-way interleaving[2] DRAMs that have two banks internally. The DRAMs are mounted on DIMMs or SIMMs. In the figure, two different module configurations are shown: a ×32 SIMM and a ×64 DIMM.

A rank is defined as the collection of DRAMs that have to be present or not present together. A rank is made up of a linear array of DRAMs that are accessed and controlled in parallel. All DRAMs within a single rank are the same speed. Usually, though not always, all ranks under the control of a memory controller must be the same speed. In Figure 4-6, the memory system is interleaved, causing a rank to consist of a group of DRAMs on each side of the multiplexer. In this case, a rank is made up of a number of SIMMs or DIMMs.

A bank is the collection of bits that are accessed together to satisfy a block transaction. In Figure 4-6, a block is half the bits in one rank of DRAMs. A **sub-bank** is the collection of bits that are always accessed with the same temporal characteristics. In this case, a sub-bank is half of a bank because the memory system is two-way interleaved. A sub-bank is half the bits in the group of DRAMs—those on one side or the other of the multiplexer. When the banks of a memory system are logically divided into sub-banks, it is typical for the corresponding ranks to be physically divided into **sub-ranks**. In most cases, all of the sub-ranks in a rank must be present or not present together.

What is memory system granularity?

SIMMs and DIMMs come in a variety of interchangeable sizes. For example, a ×32 SIMM constructed from two 1M × 16 DRAMs contains 4 Mbytes, while a ×32 SIMM made with 4M × 4 DRAMs contains 16 Mbytes. Typically, memory controllers are able to control a wide variety of SIMMs and DIMMs. Consequently, a memory system constructed with memory modules can have a wide ranges of sizes depending on the number of ranks populated and the size of the SIMMs or DIMMs used. The **granularity** of a memory system is the amount of memory in a single rank using the smallest density SIMM, DIMM, or DRAM allowed by the memory controller. Since memory systems have an integral number of banks, the granularity of the memory system is also the smallest amount of memory that can be added to the machine.

2. See Section 4.2.4 for an explanation of interleaving.

Samsung Ex. 1097, p. 78
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

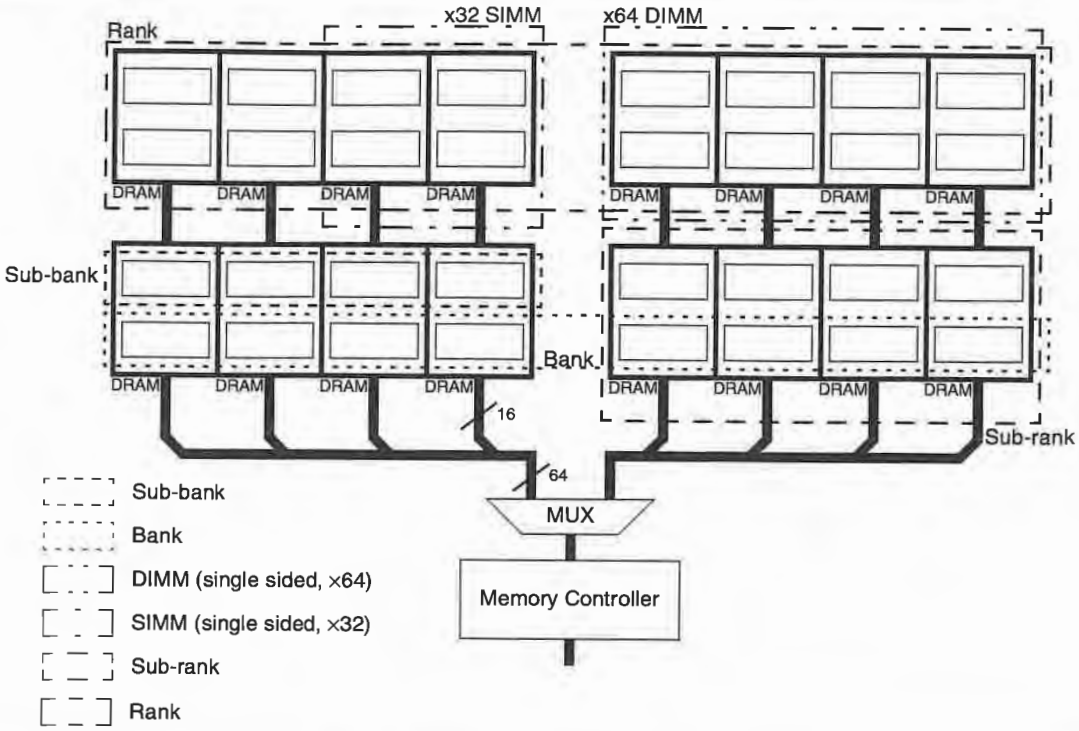


Figure 4-6 Complex memory systems are divided logically into banks and physically into ranks. DIMMs and SIMMs are ganged in parallel to form ranks, each of which may consist of one or more banks.

If the DRAMs used are not internally partitioned into distinct banks—in other words, if the DRAMs consist of a single homogeneous block of memory—then the ranks and banks are one and the same. In this case, the memory banks have the same depth as the DRAMs from which they are constructed. A 32-bit-wide memory bank made of eight 4M×4 DRAMs has a depth of 4M and a width of 4 bytes, for a total of 16 Mbytes. At a given DRAM density, building a bank or rank out of fewer DRAMs that are wider and shallower makes for banks that are shallower and consequently smaller.

Another way to describe the granularity of the memory system is as the product of the memory system's width and the minimum controllable rank depth. For example, a 16-bit memory with a 256K minimum depth has a granularity of 512 Kbytes. In contrast, the granularity of a 64-bit memory built with 4M×4 DRAMs is 32 Mbytes.

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 79
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

Since the depth of a memory rank is generally the same as the depth of the DRAMs it is made of, the granularity of a memory system is:

$$\text{Granularity} = \frac{\text{DRAM Size (bits)} \times \text{Memory Width (bytes)}}{\text{DRAM Width (bits)}}$$

Since DRAM and memory widths both increase slowly, the granularity is primarily related to DRAM density. Consequently, over the past two decades granularity has grown so much that it is often comparable to the minimum desirable memory size. If granularity grows beyond the minimum desirable memory size, then the price of an entry-level system will be inflated by the cost of the unnecessary but unavoidable extra memory.

Slowing the growth in granularity

In the face of ever rising DRAM densities, two ways exist to slow the growth in granularity:

1. Match every increase in DRAM size with a comparable increase in DRAM width.

   The migration path from 4M×4 at the 16-Mbit level would be to 4M×16 at the 64-Mbit level and on to 4M×64 at the 256-Mbit level. This is practical in the short run, but soon becomes impractical as pin counts grow exponentially.

   Although the desirable minimum memory size is also growing, increasing bandwidth requirements indicate a slow but steady growth in memory width. Consequently, the need to increase device width is undiminished by increasing memory system size.

2. Change the memory system organization and implementation so that the memory width is equal to the DRAM width.

   Rambus DRAMs (RDRAMs) and **RamLink DRAMs**, the revolutionary alternatives discussed in Sections 8.2 and 8.3, take this approach. A high-bandwidth path is provided to every DRAM, and memory system granularity is reduced to the size of a single DRAM. In the simplest use of these interfaces, the DRAMs are not ganged in parallel. Conceptually, the memory system width is equal to the DRAM width. Though these solutions dramatically reduce granularity, memory system granularity is still linked to the exponentially growing DRAM size. Unfortunately, short of using an older generation of DRAMs or combining logic and memory on a single die, little can be done to reduce granularity below the size of a single DRAM.

Samsung Ex. 1097, p. 80
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

4.2.3 Bandwidth and Latency: Peak, Sustained, and Sustainable

Improving bandwidth with block transfers

DRAMs share an important feature: once a row has been fetched from the storage cells into the sense amps, the entire row is available via column accesses. The sense amplifiers act as a cache capable of holding a row's worth of bits.

Caches take advantage of the **spatial locality of reference** in programs by fetching and storing a number of contiguous memory locations. For example, when the program executes instruction i, the cache also fetches instructions i+1, i+2, and i+3. This is worthwhile because of the significant likelihood that these instructions will be executed next.

Consequently, in most computer systems today, the vast majority of the read references to main memory are block requests. Typically, 16 or 32 bytes are fetched at a time—sometimes less in older or low-end systems, sometimes more in specialized or very high-performance systems. In Intel 486 systems, 16-byte transfers and 4-byte-wide memory systems are the norm. In such systems, each memory access becomes a row access followed by four column accesses. These accesses deliver the 16 bytes of data to the CPU in four 4-byte chunks in rapid succession (see Figure 4-7). In Pentium® systems, the processor transfers 32 bytes at a time over a 64-bit bus. Again, the transfers occur in four consecutive chunks.

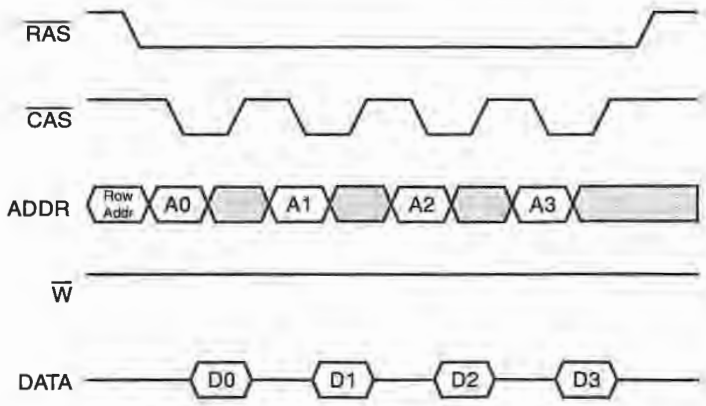


Figure 4-7 A burst of four words consists of four column accesses.

© 1996 Steven Przybylski

Two metrics of performance

Two primary metrics are used to evaluate DRAMs and memory systems: **bandwidth** and **latency**. Latency is the amount of time between the start of a transaction—typically a read transaction—and the return of the first piece of data. From the DRAM's perspective, if the transaction is a column read access, then the latency is the column access time, and if the transaction needs a row access, then the latency is the row access time. From the processor's perspective, though, the latency starts when the address of the request leaves the CPU. From there it has to travel to the memory controller where it is decoded and multiplexed. Then the address and the appropriate control signals are transmitted to the DRAMs. Once the DRAM outputs the first piece of required data, it must flow to and through the memory controller and then on to the CPU. The temporal overhead in getting the address to the DRAMs and the data back to the CPU can add many cycles of latency to the basic DRAM access time.

Bandwidth is the rate at which data is transferred. In fact, a number of different bandwidths together describe the characteristics of a DRAM and a memory system. A DRAM's **peak bandwidth** is the maximum rate at which data can be transferred into or out of a DRAM; constraints imposed by the memory system are not taken into consideration in calculating this metric. For a single DRAM, the peak bandwidth is equal to the width of the DRAM divided by the minimum column cycle time. A memory system's peak bandwidth is the width of the memory bus divided by the actual column cycle time. The actual column cycle time is typically greater than the DRAM's data sheet column cycle time because of design realities: signal propagation delays, synchronization delays, and the overhead of cycle-based controller design.

Actual **sustained bandwidth** in a system is typically much lower than the peak bandwidth. Several factors account for the gap. First, both row and column accesses have latencies during which no data is transferred. Also, in a real system, the overhead associated with transferring information to and from the processor—and often additional overhead within the processor—separates one memory transaction from the next. This time during which no data is being transferred decreases the actual bandwidth of the memory system further below the peak bandwidth. The sustained bandwidth obtained is highly dependent on the workload: the applications being executed. This is because the application determines the mix of row and column accesses and the amount of time between primary and secondary cache misses.

Samsung Ex. 1097, p. 82
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

Sustainable bandwidth defined

Because the actual sustained bandwidth depends on so many factors outside the memory system (such as the application, processor, compiler, operating system, and system architecture), the sustained bandwidth is not a convenient or reliable metric for comparing DRAMs or memory systems. Instead, this report defines a new metric: the **sustainable bandwidth**. Like the sustained bandwidth, the sustainable bandwidth is more representative of DRAM and memory system performance in a system than is the peak bandwidth. However, unlike the sustained bandwidth, the sustainable bandwidth is only dependent on the memory organization, DRAM characteristics, and the relative frequency of row and column accesses.

The sustainable bandwidth of a memory system is the bandwidth deliverable by a memory system assuming an aggressive controller and an infinite supply of memory references queued up for processing. The sustainable bandwidth is calculated assuming that the memory controller can take maximum advantage of any unique features of the DRAM.

The only remaining external parameter is the mix of row and column accesses. This is a function of the application and of the DRAM's row-cache characteristics. The **worst case sustainable bandwidth** is the sustainable bandwidth assuming a pathological reference mix that forces 100% row accesses. The **best case sustainable bandwidth** assumes a trivial reference stream that is satisfied with 100% column accesses to a random mix of banks. For most DRAM organizations, the worst case scenario involves a sequence of references to a single bank. In contrast, the best case scenario typically involves references spread across multiple banks. **Good case** and **bad case sustainable bandwidth** refer to the bandwidths attainable with row column access frequencies that are on the good and bad ends, respectively, of a realistic range of row-cache miss ratios. The range is chosen to be representative of the row-cache behavior that would be encountered in a real system using the specific DRAM and memory organization.

Peak versus sustainable fill frequencies

As defined in Chapter 4, the fill frequency of a DRAM or a memory system is the ratio of an organization's bandwidth to its size. As was just noted, the peak, sustained, and sustainable bandwidths of DRAMs and memory systems vary substantially. Consequently, the peak and sustainable fill frequencies of systems and devices also differ substantially. The sustainable fill frequencies are more meaningful than the peak fill frequencies. However, sustainable fill frequencies are more difficult to calculate properly because they are dependent on all the important

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 83
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

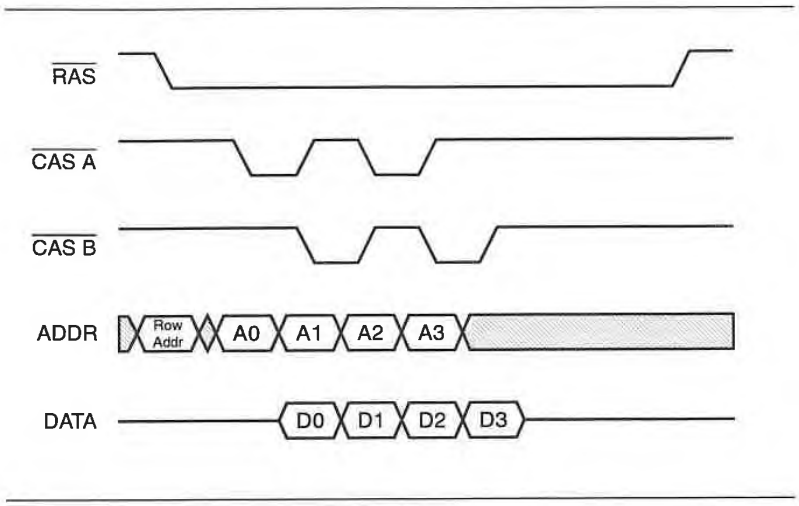


Figure 4-8 Interleaving uses two or more sub-banks to improve bandwidth.

characteristics of the DRAM and memory system—not just the peak column access cycle rate. Note that since the sustainable bandwidth is a range of bandwidths dependent on the mix of row and column accesses (i.e., the row-cache miss ratio), the **sustainable fill frequency** is also a corresponding range of fill frequencies based on the same access type mix.

4.2.4 Performance and Bandwidth Considerations

Interleaving overlaps CAS precharge time

In Figure 4-7, each access is handled entirely within one bank, even if the memory system contains more than one bank. The other banks are idle for the duration of the access. An alternative is to have more than one bank participate in each transaction. Figure 4-8 shows the same 32-byte access handled cooperatively by two banks. Both banks perform RAS accesses, then each bank does two alternating column accesses. This is potentially faster than the design shown in Figure 4-7 because each bank's CAS precharge time is in effect used by the other bank to transfer data. This alternating of two or more banks is called **interleaving**.

Unfortunately, the industry is inconsistent with respect to the terminology of interleaving. For example, Figure 4-9 shows a 32-bit memory system with four SIMMs. Each horizontal pair of SIMMs cooperates to satisfy accesses. This organization can be

Samsung Ex. 1097, p. 84
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

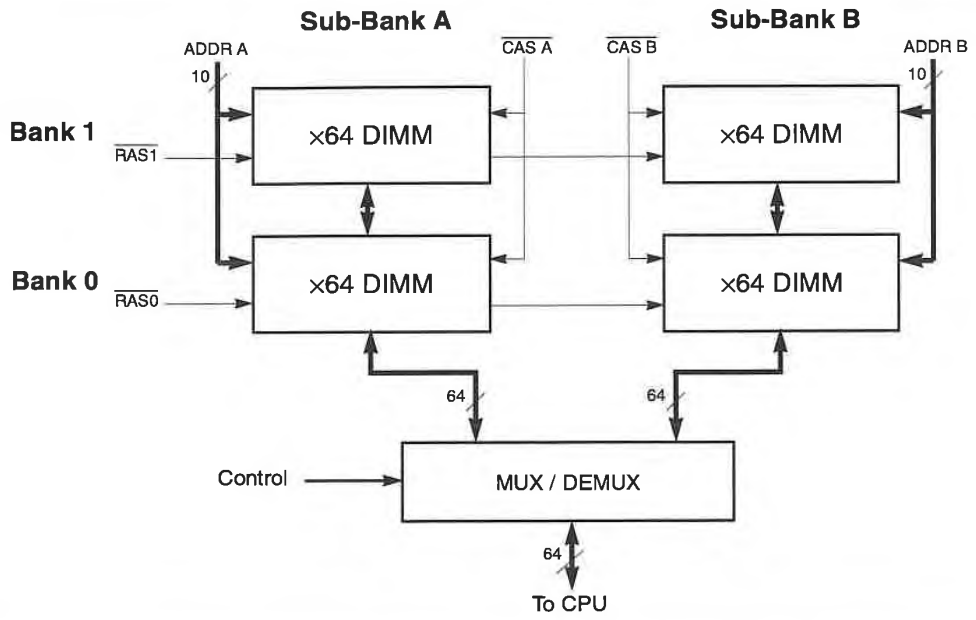


Figure 4-9 This interleaved memory system has two banks. Each bank has two sub-banks.

described as either four banks that are two-way interleaved, or two banks that are each two-way interleaved internally. A two-way interleaved bank is made up of two **sub-banks** that work together to provide data more quickly than either could do individually. The latter description is more consistent with the definition of a bank as the chunk of memory that is accessed for each memory reference; in the remainder of this report, interleaved memory systems are described in terms of banks and sub-banks.

Figure 4-10 shows the partitioning of the physical address bits in a two-way interleaved memory system. The highest order bits specify the bank being accessed; the next highest order bits specify the row address within the bank. Then come the column address bits, which are common to all interleaves. The least significant three bits indicate the interleave and byte lane of the referenced byte. For a 16-byte read access, the low-order four bits of the address would be zero, and the memory controller would have to increment the column address between the first and second column accesses to each half of the bank.

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 85
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

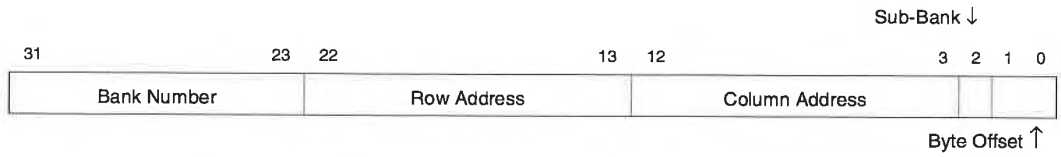


Figure 4-10 Interleaved memory systems use low-order address bits to locate words in sub-banks.

How interleaved and noninterleaved systems differ

Interleaved memory systems are primarily differentiated from noninterleaved systems by their temporally staggered column accesses. Unfortunately, although an n-way interleaved memory system has higher bandwidth than a noninterleaved system, the granularity of a memory system is increased by a factor of n. Another cost of interleaving results from the need to provide a separate address bus and control signals for each sub-bank. In high-speed interleaved memory systems, each sub-bank has its own data bus as well. The net result is that interleaving memory controllers have higher pin-counts than their generic counterparts. In lower speed systems, the sub-banks can share the same data bus. In this case the DRAMs' output-enable signals are used to switch quickly between the sub-banks. However, if the data bus is shared among the sub-banks, an interleaved memory controller can have fewer pins than a wide, noninterleaved controller of comparable bandwidth—if the shared data bus data replaces a wider nonmultiplexed bus.

Aggressive use of page mode

Memory systems can take also advantage of the spatial locality of references by aggressively using page-mode or static-column-mode DRAMs. $\overline{\text{RAS}}$ need not be raised after a memory request is satisfied. If $\overline{\text{RAS}}$ is held low between memory requests, and if the next memory request is to the same row, then the needed data can be retrieved with one or more column accesses.

Unfortunately, if $\overline{\text{RAS}}$ is low when the next memory reference arrives, and the request is to a different row, then the second access is delayed. This delay occurs because $\overline{\text{RAS}}$ must first be raised and held high for the RAS precharge time before the new row access can begin. Using page mode between accesses is therefore a gamble. If the frequency of same-row, back-to-back accesses is high enough, more time is saved than is wasted when the row address is different. The net effect is that the

Samsung Ex. 1097, p. 86
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

mean memory access time is reduced. Generally, aggressive use of page mode is good policy only if there are several banks because instruction and data references are less likely to repeatedly interfere with one another.

Row sense amplifiers as a cache

One way to view the aggressive use of page mode is to think of the sense amps, taken together, as a **cache**: as long as $\overline{\text{RAS}}$ is held low, a portion of main memory is available more quickly than the rest. This **row cache** is a direct-mapped cache in that each location in main memory can be copied into exactly one sense amplifier. The block size of the cache is equal to the size of the memory system row: that is, the row size of each DRAM multiplied by the number of DRAMs accessed in parallel. The number of blocks in this cache is equal to the number of banks held with $\overline{\text{RAS}}$ low. For example, a 16-Mbyte memory system constructed of four 1M×32 SIMMs has a row cache of 16 Kbytes, consisting of four blocks of 4 Kbytes each. In comparison to current CPU caches, this is a fairly small cache with an abnormally large block size.

Row caches exploit the spatial locality of reference in the reference stream reaching main memory. In the absence of a CPU cache, a row cache with a small number of blocks can work very well because there is a great deal of spatial locality within the instruction and data reference streams. The presence of primary and secondary CPU caches reduces the spatial locality of the reference stream reaching main memory. However, in many cases, enough of the reference stream's spatial locality remains to allow moderate, though not dramatic, improvements in system performance if page mode is used aggressively. Interestingly, if the CPU caches are write-through, a great deal more locality of reference exists in the write stream reaching main memory than in the read stream. In these circumstances, it makes sense to keep the DRAMs in page mode after a write, based on the expectation of another write to a similar address soon, but to take them out of page mode after a read.

In many respects, row caches behave similarly to CPU caches. In both CPU and row caches the miss rate is primarily determined by four factors, listed here in descending order of importance:

- Program characteristics
- Cache size
- Block size
- Set associativity

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 87
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

The primary benefit of increased set associativity is not so much a decrease in the miss rate as a reduction in the frequency of pathological conflicts that occasionally cause periods of abnormally high cache miss rates.

Row caches differ from CPU caches in two important ways:

- The block sizes that minimize the miss ratio are significantly higher than in CPU caches.
- The row-cache block size that maximizes performance is the same as the block size that minimizes the row-cache miss ratio.

Because CPU cache blocks are filled serially over several cycles instead of all at once—as row-cache blocks are—the CPU cache block size that maximizes performance is always smaller than the CPU cache block size that minimizes the miss ratio.

Linear and sub-block ordered bursts

There are two different orders in which data is returned to the CPU from the memory system. Unfortunately, each goes by several names. With the first transfer order—called linear or sequential burst—the data returned to the CPU starts with the addressed word, continues linearly to the end of the block, and wraps around to the beginning of the block. With the second transfer order—called sub-block order, interleaved, or Intel order—data transfer begins with the addressed word, and then proceeds with increasingly distant words until the entire block is transferred. The order that data is transferred in for the two types of bursts is shown in Table 4-1.

Start Address	Linear Ordered Bursts				Sub-block ordered Bursts			
	Burst of 1	Burst of 2	Burst of 4	Burst of 8	Burst of 1	Burst of 2	Burst of 4	Burst of 8
0	0	0-1	0-1-2-3	0-1-2-3-4-5-6-7	0	0-1	0-1-2-3	0-1-2-3-4-5-6-7
1	1	1-0	1-2-3-0	1-2-3-4-5-6-7-0	1	1-0	1-0-3-2	1-0-3-2-5-4-7-6
2			2-3-0-1	2-3-4-5-6-7-0-1			2-3-0-1	2-3-0-1-6-7-4-5
3			3-0-1-2	3-4-5-6-7-0-1-2			3-2-1-0	3-2-1-0-7-6-5-4
4				4-5-6-7-0-1-2-3				4-5-6-7-0-1-2-3
5				5-6-7-0-1-2-3-4				5-4-7-6-1-0-3-2
6				6-7-0-1-2-3-4-5				6-7-4-5-2-3-0-1
7				7-0-1-2-3-4-5-6				7-6-5-4-3-2-1-0

Table 4-1 Linear and sub-block ordered burst are identical for short burst and bursts that start at the start of a block.

Samsung Ex. 1097, p. 88
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

The order of the data returned in the sub-block-ordered case seems almost random. However, this data order actually exhibits a number of important properties. First, the address the nth piece of data is obtained by doing the exclusive-OR of n and the starting address of the transfer. Consequently, the order does depend on the length of the burst. For example, bursts of length four are the first part of the bursts of length eight. This makes it easier to abort and subsequently restart long transfers. Second, since the first two pieces of data transferred are always an odd and even pair of words, it is much easier to build DRAMs and memory systems that are half-speed and double-width internally. Sub-block ordering has been used in all Intel processors since the 486 and in some workstations as well.

It is important to note that the two burst orders are identical for short bursts and for all bursts with a zero start address. Overall, the choice of burst type is a matter of convenience and politics, and not a matter of great significance from a performance perspective.

4.2.5 Parity and Error Correction Codes (ECC)

Generally, computer memory systems are supposed to faithfully retain the information entrusted to them. In many applications, even a small chance of losing a single bit is unacceptable. Given that DRAMs occasionally lose information to alpha particles and other transient errors, memory system engineers add some redundancy in order to reduce the probability of bad data being used accidentally.

In addition to the straightforward cost of additional memory bits that must be purchased to provide parity and error protection, hidden costs stem from the additional logic and time in the memory controller and/or processor needed to set and check the redundant bits. More significantly, when the parity or error correction code is formed over 16-, 32-, or 64-bit words, then memory writes that are smaller than the word size must be performed as read-modify-write cycles so that all the redundant bits reflect the changed portion of a memory word. This naturally takes longer than a simple write operation.

Two factors affect the acceptable level of error protection in a computer system: the probability of an error and the likely result of an error. When a computer gets the wrong answer or crashes due to a bad bit, the resulting problems depend greatly on the application. The cost of a flatly wrong answer is

© 1996 Steven Przybylski

potentially much greater in a computer controlling a nuclear reactor, flying a space shuttle, or shuffling money around than in one running a video game. The probability of an error also influences the acceptable level of protection. An expected bit failure rate of one per 10,000 years is probably acceptable in almost any application, while one error per day is probably unacceptable even to the most casual of PC users.

Unprotected memory systems are increasingly common

Three levels of protection from transient errors are commonly used:

No Protection:
: In the extreme, if no redundancy is added to the memory system, any error that does occur is undetectable. The possible consequences of using a bad instruction or datum range from no discernible effect, through an application or system crash, to an unsuspected wrong answer.

Parity Protection:
: The addition of a single parity bit per word of memory allows the detection of any single-bit error within each word. However, correcting the error once it has been detected is impossible since all single-bit errors within a word exhibit the same symptom.

Single-Error Correcting, Double-Error Detecting (SECDED) Codes:
: By adding seven redundant bits per 32-bit word, or eight redundant bits per 64-bit word, it is possible to detect and correct any single-bit error to a word and to detect, but not correct, any two-bit error to a single word.

Given the steady and dramatic decline in soft error rates in DRAMs over time (see Section 3.5.1), unprotected memory systems are becoming more and more acceptable for a wide variety of non–mission critical applications and even those that are mission critical. The limited safeguard afforded by parity protection is becoming irrelevant: either data integrity is sufficiently important to warrant error correction, or the soft error rate is low enough that the system failure rate due to memory errors is well below the failure rate due to other causes. Already, the parity protection that is standard in PCs is not primarily protecting the user from soft errors in the DRAMs. More often, parity protects the user from other sources of intermittent memory failures. These sources include user-installed memory that is fractionally too slow or is poorly seated in its

Samsung Ex. 1097, p. 90
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

socket and overly aggressive motherboard designs that intermittently and inadvertently run DRAMs or other memory system components beyond their rated specifications.

4.2.6 Refresh

DRAMs must have their contents refreshed periodically. Almost all DRAMs must receive a refresh operation on average every 15 μs to ensure that all information is retained under worst case operating conditions. Section 3.3.1 describes the two primary types of refresh cycles offered by generic DRAMs.

In deciding how DRAMs are to be refreshed, the memory system designer addresses the following questions:

- How should available bandwidth be divided between refresh and normal operations?
- Which type of refresh operation should be used?
- Can the processor generate a memory reference while a refresh operation is taking place? If so, how is this handled?

The possible responses to these questions involve choosing one of two basic approaches to DRAM refresh operations:

- Static allocation of a fixed proportion of the memory bandwidth to refresh. Memory references are multiplexed with refresh operations according to the fixed allocation. This strategy results in a fixed but nonminimum memory response time.
- Dynamic allocation of all memory bandwidth to either memory references or refresh operations as required. Such organizations have greater average bandwidth available to the processor at the expense of decreased predictability. The response time for a memory reference will be longer than usual if it arrives at the memory subsystem while a refresh operation is occurring.

To minimize the amount of memory bandwidth used for refresh, it is common to refresh all banks within a memory subsystem simultaneously. Unlike a normal memory reference in which only one bank performs a RAS operation at a time, in a multibank refresh all banks perform RAS operations together. To avoid an unacceptable amount of power supply noise generated by a large number of DRAMs refreshing simultaneously, it is common to slightly stagger the refresh operations of all banks

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 91
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

in the system. The drawback is that when staggered, an entire refresh operation takes somewhat longer than if all banks are refreshed in synchrony. However, the combined benefits of reliable operation and reduced demand on the power supply greatly outweigh this drawback.

4.3 High-Speed Design Issues

An inevitable consequence of the constant improvement in IC fabrication technologies is an increase in the maximum operating frequency of ICs. As intrachip speeds and bandwidths improve, there is constant pressure to increase the speed and bandwidth of the interchip connections. However, the physical laws that govern the maximum rate of signal and information flow come into play much more forcefully in the interchip case because of the longer distances to be traversed. This section highlights some of the impediments and fundamental limitations that come into play when high-speed DRAMs are combined to form high-performance memory systems.[3]

4.3.1 High-Speed Signals and Terminated Transmission Lines

The speed of an electrical signal traveling down a wire is determined by the physical properties of the wire and its surrounding material. For traces on PC boards, the propagation delay[4] of a signal traveling down a wire is between 100 and 500 ps/in, depending on the physical dimensions of the trace and its neighbors, the printed circuit board material, and the number and characteristics of loads on the trace (ICs, connectors, etcetera).

Electrical signals have rise and fall times: the time the signal takes to change from the voltage representing one logic level to that representing the other.[5] When the time necessary for a signal to traverse a wire is much shorter than the rise time of the signal, then the wire and all the components along it act as a

3. A more complete, yet thoroughly approachable, presentation of the complexities of high-speed design can be found in H. Johnson and M. Graham, *High-Speed Digital Design: A Handbook of Black Magic*, Prentice Hall, 1993.

4. It is convenient to express the speed of signal propagation as a delay per unit length. A delay of 100 ps/in is the same as a speed of 256×10^6 m/s.

5. Rise and fall times are generally defined as the elapsed time from the 10 to 90% points: that is, from the point in time that the signal has already made 10% of the total transition from the voltage representing one logic level to the voltage representing the other logic level, to the point in time when it has completed 90% of the transition.

Samsung Ex. 1097, p. 92
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

single lumped (capacitive) load. This is the normal situation in low-speed systems with slowly changing signals. The lumped capacitive load is easily characterized, and signals are well behaved in that transitions are smooth and predictable.

Any of three circumstances can invalidate the basic assumption of long rise times relative to signal propagation delays:

- Traces that are longer than a few inches
- Traces that have unusually slow propagation delays due to heavy loading or unusual construction
- Signals that transition in less than a few nanoseconds

The total propagation delay of a signal from point A to point B is related to both the transition time and the propagation delay on the wire proper. As systems have become faster, there is an increasing need to cut the total propagation delay of interchip communication by pushing rise and fall times to ever lower levels—levels that cause the lumped model of wires to break down.

When the transition time is less than the propagation delay on a wire, the wire behaves as a transmission line. Signals enter the transmission line at a driver and propagate down the wire as a wave front. An appropriate physical analogy is of a long, taut string that is plucked close to one end. The wave caused by the plucking travels down the string; if the far end is held firmly, the wave reflects back to the source. Voltage waves reflect off the ends of wires in an analogous fashion. The magnitude and polarity of the reflected wave depend on the characteristics of the load at the ends of wire relative to the electrical characteristics of the wire itself. In the simplest case, the voltage wave will bounce off the two ends of the wire in ever smaller amounts until a steady state value is reached.

Reflections can originate not only at the ends of the wire but also wherever there is an electrical load on the wire or a change in the wire's physical (and corresponding electrical) properties. Examples of electrical discontinuities that can cause reflections include connectors, ICs, vias from one level of a multilayer PCB to another, and, perhaps most importantly, stubs. Stubs are short sections of wire or trace branching off from a main trace to one or more chips. If the stub is short, it is effectively invisible to the voltage wave, and the wavefront behaves as if the loads were on the main trace. However, if the stub is too long, it becomes a source of muted reflections for signals traveling down the main trace.

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 93
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

Signal reflections cause ringing and delay

Reflections, though not inherently bad, become a problem when they interfere with one another enough to degrade signal quality or to delay the reception of a signal at its intended destination. Figure 5-2 on page 138 illustrates the ringing that can be caused by interference between reflected wavefronts. Ringing not only delays the point at which the signal can be reliably interpreted at the destination, but, when excessive, can physically harm the input electronics.

One way that reflections can be eliminated is by placing an appropriately sized resistor at the end of the trace. This is called parallel, or end, termination of the trace. When the termination resistor matches the characteristic impedance (i.e. resistance) of the trace, then the magnitude of the reflection vanishes to zero. In this case, the effective signal propagation delay is determined by the characteristics of the trace and the length of the wire. Calculation of the delay is not complicated by the reflections. Reflections from all the intermediate discontinuities along the trace are minimized or eliminated by:

- Making the trace's electrical characteristics as uniform as possible along its entire length
- Concentrating the load at the terminated end of the trace or distributing the load evenly along the trace's length

Another form of termination is source termination. In this scheme, a resistor equal to the characteristic impedance of the trace is placed between the driver and the transmission line. The far end of the trace is left unterminated. When the driver switches state, a half-amplitude wave-front propagates to the unterminated far end of the trace. There it reflects back toward the driver and, in the process, the signal attains its full amplitude. When the reflected wave reaches the driver, it is absorbed by the source terminating resistor, and no further reflections occur.

For single-source, single-destination traces, both source and parallel termination are acceptable. In both cases, the receiver at the far end of the trace sees a full amplitude signal in minimum time. Parallel termination causes some additional power to be dissipated in the termination resistor. The magnitude of this dissipation depends on 1) the termination voltage, 2) the voltage levels representing logic 0 and logic 1, and 3) the characteristic impedance of the trace.

Samsung Ex. 1097, p. 94
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

Source termination also involves a complication. Receivers that are not at the far end of the trace from the driver first see a half-amplitude wave traveling toward the far end. Later on, after the incident wave has traveled to the unterminated far end and reflected back, their input voltage completes its transition. The half-voltage state may be insufficient to be recognized as a transition, and so the effective time to communicate with these intermediate receivers is longer than with parallel termination. In some cases, this half-transitioned level can cause excessive power to be dissipated in the receiver's input buffer.

The use of **terminated transmission lines**—uniform and uniformly loaded traces terminated through an appropriate resistor to either the driver or a termination supply—is a fundamental prerequisite to the design of very high-speed systems.

Bidirectional buses pose special challenges

This discussion has focused on the case of a wire driven at one end toward one or more loads at the other end. However, bidirectional buses, with many potential sources and destinations of information, are common in computer systems. These buses pose special challenges because they are typically:

- Physically long
- Heavily and unevenly loaded
- Critical to high-speed operation

Because drivers can be anywhere along the traces in bidirectional buses, signals propagate in both directions. If reflections are to prevented, the bus must be terminated at both ends. Consequently, the DC resistance of the bus is equal to half the characteristic impedance of the trace, and the power needed to drive the bus is doubled. Source termination is inappropriate for multisource buses.

In non–transmission line environments, the total capacitive load on a bus often limits the speed of operation. The bus can sometimes be effectively partitioned through the use of low-resistance NMOS transistors. These transistors selectively connect or disconnect major portions of the overall net when specific communication needs warrant doing so. If the maximum load connected to the driver at one time is dramatically reduced, the time needed to transition the bus from one logic state to the other is also reduced. These NMOS transistors can also be used with high-current TTL drivers to construct 3:1 address

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 95
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

multiplexers that drive the heavy loads found in memory arrays faster than any monolithic TTL multiplexers.

4.3.2 Clock and Data Skew Issues

Unfortunately, getting high-quality electrical signals from point A to point B is only part of the problem. Designers of high-speed systems must also cope with the issue of elapsed time.

Clocks are among a system's most important signals. Clocks define the start and stop of each machine cycle; consequently, the beginning and end of all of interchip communication is defined relative to one or more clock edges. However, like all electrical signals, clocks have rise and fall times and propagation delays. If the clock that defines the start of a cycle does not transition at all ICs at the same time, then the time available within a cycle for communication between ICs is compromised. The maximum difference in time between the arrival of a clock at different parts of a system is called **clock skew**, and various techniques are employed to minimize it:

- Distributing multiple copies of a clock so that one signal does not have to reach all ICs on a board or in a system
- Ensuring that all clock traces are the same length and similarly loaded, thus minimizing differences in the clock propagation delays to all parts of the system
- Using **Phase-Locked Loops (PLLs)** in ICs to minimize the delay from a clock transition at an IC's input pin to the clock transition throughout the chip

Skew can exist between data signals as well. Buses normally carry signals in parallel: the bits traveling down the bus together are temporally aligned—that is, sent and received at the same time. However, if one or more of the traces is longer or more heavily loaded than the rest, the bits on those traces will arrive later than the bits on the other traces. This difference effectively constrains the time at which the bus as a whole can be latched. In the extreme, the difference in time between the arrival of the fastest and slowest bits can impact the minimum cycle time of the system.

Even in a system in which all clock and data skew has been eliminated, signal propagation delays can cause headaches, especially with regard to bidirectional buses. Most commonly, data is set up at the output pins prior to the start of the clock

Samsung Ex. 1097, p. 96
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

cycle in which the data is to be communicated to another IC. At the appropriate edge of the clock, the data flows through the transmitting chip's output buffers. The data then travels down the traces to the receiving IC or ICs. The data must reach the destination devices prior to a setup time before the end of the cycle. At the end of the cycle, the data is latched at the input buffers of the receiving chip(s). Regardless of the direction of information flow, this scheme allows one full cycle (less any clock skew) for an output driver, an input receiver, and propagation along the trace. If the trace is exceptionally long or heavily loaded, or if the output driver is slow to deliver data onto the trace, then the cycle time of the entire system must be increased to compensate.

Coping with slow communication in high-speed systems

One common reason for the output buffer being slow to deliver data is that the data is being generated in the same cycle that it is being communicated. In this case, less than one full cycle is allocated to actual communication between the chips. If the time from the clock to the data being valid on the output pin is too long, and the cycle time of the system is being constrained, then the clock at the receiver can be delayed relative to the clock at the transmitter. One drawback is that the time given to communication in one direction comes directly out of the time available for communication in the other. Also, it can be hard to manage and control many clocks in a system.

The obvious solution to this problem of long propagation times is to orchestrate communication at some lower frequency than the system's primary frequency. This is the idea behind the proliferation of processors with frequency-doubled, -tripled, and -quadrupled clocks internally. To prevent low-speed external communication from restraining high-speed internal computation, the processor's external interface is run at one-half, one-third, or one-quarter the rate of the CPU core. Thus, two, three, or four processor cycles are available to generate, propagate, and receive external signals.

The idea of shifting the send and receive clocks in time to provide more time for signal propagation can be taken further. If the clock at the receiver is shifted one-half cycle out of phase from the clock at the transmitter, symmetry can be restored. If one-and-one-half cycles are used for communication in each direction, then three full cycles, instead of the normal two, are allocated to there-and-back communication. This increased number of cycles naturally increases the latency of operations being performed in the remote IC. This practice is not without its hazards, though. In burst transfers in which a new piece of

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 97
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

data is sent on every cycle, the driver conceptually begins sending the second chunk of information before the first chunk is latched at the destination. If the difference between the driver's data-valid-out time and data-hold time is too great, then the first piece of information may disappear before it can be latched. Under such circumstances, phase shifting the send and receive clocks with respect to each other is of limited utility and the only recourses are to:

- Send data every other cycle, as with a clock-doubled CPU
- Restrict the operating frequency of the system
- Explicitly pipeline the communication by adding latching between the source and destination ICs.

4.3.3 Issues Specific to Memory Systems

As DRAMs move to higher operating frequencies, some situations that commonly arise in the design of high-performance memory systems will begin to confront designers of mainstream memory systems.

The common structure for memory systems today is a wide bus with address, data, and control signals in parallel on a PC board. On this wide bus are SIMM sockets, one or more of which are populated with conventional SIMMs of various configurations. The sockets need not all contain the same type of SIMM. This physical structure is not at all conducive to high-speed operation for the following reasons:

- The motherboard bus is long and nonuniformly loaded. A large maximum load means short worst case transition times and long propagation delays.
- The address and control traces on the SIMMs are heavily loaded and long enough to be reflecting and resonating stubs that contribute to poor signal quality.
- DRAMs, which are quite slow to turn on and off, limit the frequency with which multiple sub-banks sharing the main bus can be interleaved.

One technique that can effectively control the reflections caused by the stubs on the SIMMs is to place small resistors between the main bus and the stubs in series with the connectors. These resistors dampen the effect of the reflections on a motherboard

Samsung Ex. 1097, p. 98
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

trace and help maintain good signal quality throughout the memory system.

In a synchronous memory system, the signal propagation times alone make the use of unbuffered SIMMs untenable at very high frequencies. But the use of buffered SIMMs cause new problems. The buffers add delay, and, more significantly, variability in the delays of address, control, and clock signals: this variability translates directly into clock and data skew. The extraordinary amount of clock skew introduced by standard per-SIMM buffers can be eliminated by using Phase-Locked Loops (PLL) on every SIMM. However, using PLLs adds cost and requires additional start-up time after a power down.

The delay introduced by address and control buffers adds significantly to the total propagation time from the memory controller to the synchronous DRAMs,[6] making single-cycle, end-to-end communication difficult: latching buffers may be necessary for operating frequencies above 66 MHz. The latches may facilitate high operating frequencies, but of course they also increase the latency of all operations by one cycle. A comparable set of buffers may also be needed on the data lines to preserve good signal quality on the motherboard bus and maintain high-frequency operation.

Ideally, the buffers would be on the motherboard instead of on the SIMM in order to minimize the variability of loading on the motherboard bus. However, placing the buffers on the motherboard increases the motherboard's cost and area. The alternative is to place the buffers on the SIMMs, and mandate that all SIMM sockets be filled with either real SIMMs, or dummy SIMMs with loading comparable to that of real SIMMs.

The physical design of memory systems is very important in attaining high-speed operation. Regular placement of DRAMs or SIMMs ensures consistent trace lengths and uniform loading. Also, since the current consumption of DRAMs is very bursty—DRAMs draw a lot of current when doing row and column accesses but little current in between accesses—and because DRAMs are relatively sensitive to power supply noise, careful placement of adequate bypass capacitors throughout the memory array is very important. Bypass capacitors are small capacitors, connected to the power and ground planes of the printed circuit board, that supply the transient current needs of

6. Recall that *synchronous DRAM* refers to any DRAM with a synchronous interface, rather than to a *Synchronous DRAM*, which is a JEDEC-standardized DRAM.

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 99
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

the DRAMs. Inadequate bypassing, or selection of the wrong capacitors, can lead to unreliable operation. Many capacitors designed for low-frequency applications are primarily inductive at high frequencies. Staggering the refresh operations of individual banks and sub-banks is also very important to the maintenance of high-quality supplies.

4.4 Memory Systems for PCs

4.4.1 x86 Machines

Cost is by far the most important aspect of memory system design for PCs. PCs are also sold on the basis of configuration and, increasingly, service. Processor type and frequency are generally the only performance-related design considerations, especially for low-end and mid-range products. Furthermore, in the presence of primary and secondary caches, the memory system's leverage on system-level performance is very small. Memory system design changes that improve performance at the expense of increased cost are seldom worthwhile.

Figure 4-11 shows a block diagram of a typical Intel Pentium®-based desktop PC. The processor bus connects the processor, the secondary cache and two system control ASICs. The System Controller ASIC and the Datapath ASIC separate the processor address and data buses from the memory and I/O buses. The DRAMs are logically and electrically quite close to the processor. Consequently, the overhead in the cache miss penalty, above and beyond the basic DRAM access time, is minimal. The secondary cache control is either incorporated into one of the system control ICs or into a separate dedicated device.

Low-end and midrange systems are often built with sockets for two banks; each bank is constructed of two ×32 SIMMs. Some PCs are now being built with ×64 DIMMs, and many **chip sets** allow a mix of socket types. Commonly, bank depths of 256K–4M are allowed; the wide range of allowed bank depths facilitates a correspondingly wide range of memory sizes: 2–128 Mbytes. With 1M×4 DRAMs, the minimum memory size (granularity) is 8 Mbytes.

To provide greater flexibility to the system designer and end user, chip sets are programmable with respect to several of the

Samsung Ex. 1097, p. 100
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

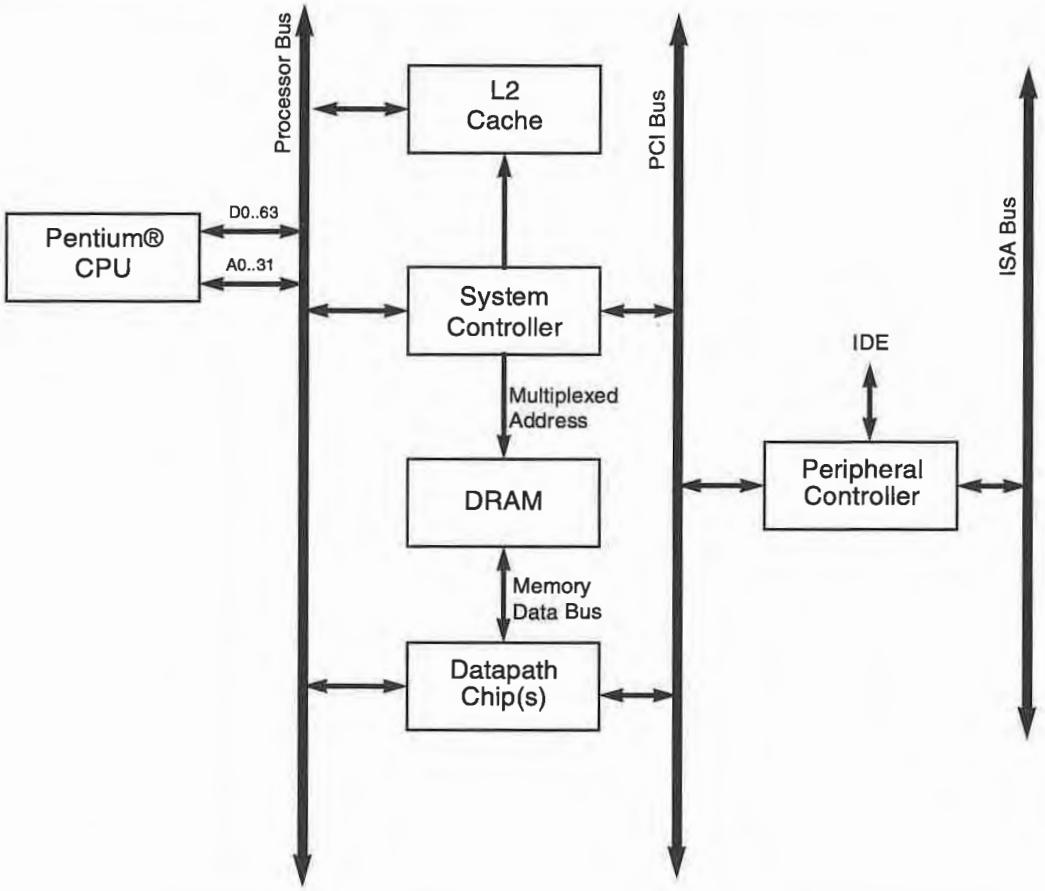


Figure 4-11 Block diagram of a typical Pentium®-based PC.

major DRAM timing parameters. Many PC chip sets support interleaved accesses when two compatible banks are installed. They can also perform hidden and/or staggered refreshes.

Almost all new PCs have at least one level of caching. In Intel 486-based systems, the primary cache is a **write-through cache**: all processor-generated data writes are written to the cache and immediately to main memory. The external secondary cache, if there is one, can be either write-through or write-back. In a **write-back (copy-back) cache**, written data is propagated to main memory only when a cache line that has absorbed a write is replaced. Such a line is called a dirty line. If only write-through caches are present in the system, the

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 101
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

majority of the references reaching main memory are small writes of one to four bytes. In contrast, when there is a write-back cache, almost all references reaching memory are block operations, and between one-third and one-half of these are writes. The Pentium® processor has a separate instruction cache and a write-back data cache on-chip.

Typically, the chip set provides individual $\overline{\text{RAS}}$ signals for each rank supported. Chip sets generally support up to four, six, or eight ranks, though most low- and midrange PCs only provide sockets for two ranks. $\overline{\text{WE}}$ signals are shared by all ranks but are sometimes duplicated on the chip set to reduce the fanout on the motherboard. Four or eight $\overline{\text{CAS}}$ signals are shared by all ranks—one $\overline{\text{CAS}}$ signal for each byte lane in the memory system.[7] This matrix of $\overline{\text{RAS}}$ and $\overline{\text{CAS}}$ signals allows for full-speed byte writes. During a byte write, only one of the four or eight $\overline{\text{CAS}}$ signals is asserted.

A serious consequence of this orthogonal matrix of $\overline{\text{RAS}}$ and $\overline{\text{CAS}}$ lines is that at most one bank can be active (i.e., open and available to column accesses) at a time. Since the DRAM chip-selects are not used in PCs—narrow SIMMs and some early DRAMs do not provide them—if more than one $\overline{\text{RAS}}$ is low at a time, then all open banks would perform the column operation indicated by the shared $\overline{\text{CAS}}$ lines. Because only one bank can be open at a time, the DRAMs' row cache cannot be very effective, since only a fraction of it can hold data at a time. Consequently, most chip-set DRAM controllers precharge the accessed bank immediately after every transaction. In most cases, the precharge is complete by the time the next memory access arrives at, and is recognized by, the chip set. Consequently, the main memory access time is almost always equal to the row access time of the DRAMs plus the inevitable chip-set and bus transfer overhead.

Until recently, most PC chip sets did not try to optimize the L2 cache-miss timing. The chip sets took several cycles to recognize an L2 miss, decode and multiplex the address, and initiate a DRAM row access. This substantial overhead—together with the processor's internal L1 miss overhead—matched the DRAM access time in importance. More recently, though, chip-set designers have been optimizing the main memory controller to minimize the number of wasted cycles in the case of an L2 cache miss.

7. In 486-class machines, the system and memory buses are 32 bits wide, so four $\overline{\text{CAS}}$ lines are needed. In Pentium and Pentium Pro systems, the system memory buses are 64 bits wide, so eight $\overline{\text{CAS}}$ signals are used.

Samsung Ex. 1097, p. 102
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

Fill frequency requirements

In early 1995, before the release of Windows 95, the acceptable granularity of low-end PCs was 4 Mbytes, especially on the 486-class systems. The acceptable peak bandwidth for these systems was around 100 Mbytes/s, making for a market-required fill frequency of around 25 Hz. High-end Pentium systems had 8-Mbyte granularities and peak bandwidths of around 180 Mbytes/s. Consequently, their fill frequency market requirements are around 22 Hz.

However, by the end of 1995, with Windows 95 firmly entrenched as the default operating system on new x86 PCs, granularities jumped by a factor of two: to 8 Mbytes at the low end and 16 Mbytes at the high end. At the low end, the peak bandwidth market requirement also doubled as the 32-bit memory systems of the 486-class machines were replaced by the 64-bit memory systems of Pentium systems—consequently, the fill frequency requirement changed little. At the high end, much of the industry shifted from fast-page-mode DRAMs to EDO DRAMs, effectively raising the peak bandwidth requirement to about 260 Mbytes/s. The two shifts—from fast-page mode to EDO and from 8- to 16-Mbyte granularity—combined to drop the high-end fill frequency requirement to about 16 Hz.

The required fill frequencies for PC memory systems have declined from the range of 30–50 Hz a decade ago to about 15–25Hz recently (see Section 2.5). The ranges are fairly broad: peak bandwidth requirements are very weak because most computer purchasers do not directly care about bandwidth. The primary reason for the decline has been the increasing use of caches in the PCs to reduce the CPU's demands on the memory subsystem. However, with the advent of streamed multimedia data types and significant non-CPU compute requirements, bandwidth and fill frequency requirements are likely to creep up again in the coming years.

4.4.2 The Macintosh Family

Significant differences between the memory systems of Apple Computer's Macintosh family and those of the x86-based PC compatibles reflect Apple's flexibility in design—a flexibility that stems from Apple's control over the Macintosh product line. For example, though most of the recent Macintosh computers have used standard 30-pin SIMMs, the Macintosh IIfx and Macintosh IIsi used unusual 64-pin SIMMs. Apple also routinely solders the basic memory configuration onto the motherboard, providing sockets only for expansion memory. Finally, Apple

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 103
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

normally does not provide parity protection for main memory. Only two of Apple's machines—the Macintosh IIci and Macintosh IIfx—allow the user to attain parity protection through the installation of ×9 SIMMs instead of ×8s.[8]

Some of Apple's latest PowerPC systems were among the first to use the 8-byte 168-pin DIMMs.

4.4.3 Portables

The important design criteria in the portable PC domain are similar to those in the desktop arena. The main difference is the PC domain's emphasis on low power consumption. Most portable systems use special low-power DRAMs with long refresh periods and/or lower operating currents. In addition, to reduce the power consumption per row access, these machines tend to use wide ×16 DRAMs (see Section 6.2). Most notebook and laptop machines take proprietary memory modules instead of generic SIMMs, primarily because of the very tight physical constraints inside the machines. The latest standard memory modules, Small-Outline Memory Modules (SOMMs), give portables the same flexibility that SIMMs provide in the desktop market, but with a higher packing density. Many of the smallest machines use PCMCIA card slots as a mechanism for memory expansion.

Interestingly, some of the early Apple PowerBooks used Pseudo-Static RAMs (PSRAMs) instead of regular DRAMs. PSRAMs are DRAMs with all the refresh control circuitry on chip. They have the important advantages of very low power consumption and entirely autonomous refresh. Consequently, when the system is idle, the memory controller does not have to periodically initiate refresh cycles—it can go to sleep with the rest of the system.

4.5 Memory Systems for Workstations

Workstations and servers are different from PCs in three important respects that impact their respective memory systems.

8. *Designing Cards and Drivers for the Macintosh Family*, Third edition, Apple Computer Inc., 1992.

Samsung Ex. 1097, p. 104
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

First, a typically configured workstation has four to sixteen times more memory than a midrange PC. Second, workstations' greater processing power and higher I/O rates imply substantially higher main memory bandwidth requirements. Finally, performance plays a more significant role in influencing workstation buyers. Therefore, adding some cost to the workstation memory subsystem in order to improve system-level performance by a few percentage points may be justified.

The primary consequence of the larger bandwidth in workstations is that workstation and server memories tend to be wider and are more commonly interleaved to get maximum performance. Both these trends increase the granularity of the system, but fortunately that is generally not a problem because of increased memory sizes.

Workstations and workstation servers span an important organizational boundary. Like PCs, low-end workstations are essentially single-board computers. The processor or processors, the memory, and the bulk of the graphics and I/O capability are on one printed circuit board. Many trade-offs and choices in this domain are fundamentally similar to those in the PC arena. However, at the high end of the workstation market, the systems follow the traditional minicomputer model with processor cards, autonomous memory cards, and I/O cards interconnected by a backplane bus (see Figure 4-12).

To improve performance, workstation designers increasingly use buffered SIMMs. By facilitating shorter column cycle times, buffered SIMMs yield memory systems with higher bandwidths. Since the address buffers are on the SIMMs themselves, the signal traces distributing the address and control signals from the memory controller are shorter and lack long, heavily loaded stubs on the SIMMs. Elimination of these stubs significantly improves signal quality and reduces overall signal propagation time to the DRAMs.

Multiprocessors increase memory traffic

Multiprocessors are rising dramatically in popularity in the workstation arena. The primary impact of multiprocessing on the memory subsystem is increased memory traffic. The increased number of processors means a linear increase in baseline memory traffic. The amount of additional bus and memory traffic can be substantial when processors share data and cooperate on a task.

Aside from increased bandwidth requirements, the impact of multiprocessing on the memory subsystem varies greatly

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 105
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

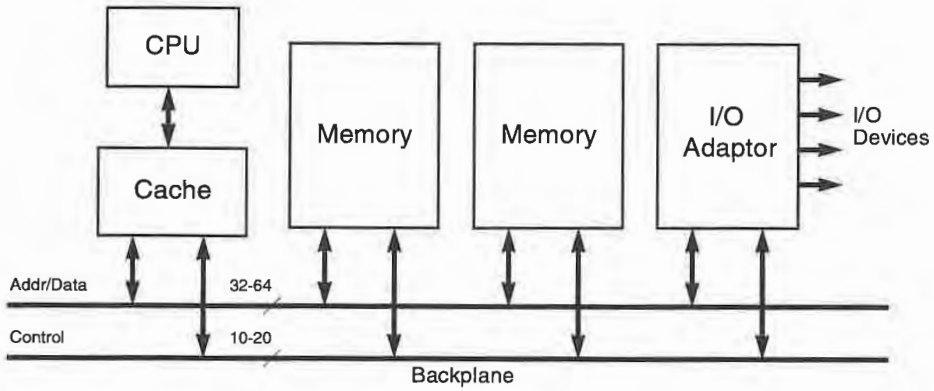


Figure 4-12 Many high-end workstations have CPU(s), memory, and I/O blocks connected on separate cards connected by a backplane.

depending on the type of computer system. The impact is limited in the case of small-scale multiprocessors with snoopy-bus-based **cache-coherence** protocols. The impact is more significant when the cache-coherence problem is solved by keeping directories at main memory. In the latter case, the memory subsystem is charged with keeping track of which memory blocks are currently held in which caches. Memory plays an active role in the sharing of data by mediating between two or more processors bidding for exclusive access to shared data. The need to play this role adds complexity to the memory controller and mandates a secondary bank of RAMs for storing the directory of coherency information. This directory generally adds 5–15% overhead to the size of main memory.

Increased memory traffic in multiprocessors can be combated through the use of a distributed memory architecture. In this architecture, main memory is broken up into smaller chunks and physically distributed throughout the system. Typically, each chunk is associated with a processor. As the system grows, more nodes—each consisting of a processor, some memory, and a network interconnection—are added. With some operating system collusion, most of a processor's memory references will be to its local memory even though there is a shared global address

Samsung Ex. 1097, p. 106
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

space. The total amount of network traffic is therefore limited to shared data and possibly I/O.

Workstations use error protection

Most workstations use error detecting and correcting codes in their memory systems to protect against soft and some hard errors. Use of a single-error correcting/double-error detecting (SECDED) code is an explicit requirement for many workstation market segments.

Finally, workstations are characterized by secondary caches that are larger than those found in PCs. The decreased miss rate of these large caches helps reduce the demands on the memory system. However, large caches also decrease the performance impact of main memory. This reduced performance leverage and the longer cache blocks found in the large secondary caches (typically 32–128 bytes) combine to decrease the emphasis on low memory latency and to increase the importance of high-bandwidth solutions.

Fill frequency requirements

Compared with high-end PCs, workstations have greater bandwidth requirements but dramatically relaxed (i.e., increased) granularity requirements. For example, minimum memory configurations of 64–256 Mbytes are common in this market segment. The net result is that the fill frequency requirements for workstation memory systems are currently in the vicinity of 10–20 Hz. These lower fill frequencies explain why the workstation market adopted the 4M×4 fast-page-mode DRAM, while the PC industry waited for the 1M×16 devices to be cost-competitive.

4.6 Memory Systems for Large Machines

Three classes of large machines are noteworthy: regular mainframes, vector supercomputers, and massively parallel machines. Common to these machines are very large memory sizes and massive amounts of memory bandwidth available throughout the machines. Multigigabyte primary and secondary memories are commonly supported with many gigabytes per second of aggregate bandwidth.

Mainframes

Thus the primary difference between a mainframe and a high-end workstation or server is not so much processing power as

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 107
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

I/O capability. Over the last half-decade, the dramatic performance lead that mainframes once held over microprocessors has vanished. However, no uniprocessor microprocessor-based system exists that can keep a terabyte disk farm meaningfully utilized.[9]

The greater I/O capability of mainframes affects the memory subsystem in two ways: main memory size is increased and memory bandwidth is larger. Mainframe main memories typically start at several hundred megabytes on a minimally configured system and run as large as several gigabytes. Frequently, main memory is logically and physically partitioned into two separate units: a main store and a secondary store. The main store, which is the smaller of the two, is sometimes implemented with SRAMs; the larger, slower secondary store is always constructed of DRAMs. Unlike a cache, which is not visible to software, the operating system manages both the main and the secondary stores.

With 4-Mbit DRAMs, 1 Gbyte of secondary store takes 2,048 ICs. If mounted on 128 SIMMs, this same memory could be placed on a single large PC board. Consequently, though memory sizes have continued to grow over the years, the physical volume of mainframe memory systems has shrunk: the appropriate measure of mainframe memory has changed from cubic meters to cubic feet.

Increased bandwidth requirements are satisfied in three ways:

- Through the use of static RAM where economics permit
- Through greater interleaving of memory banks
- Through wider memory organizations.

Eight- and sixteen-way interleaving is common in machines of this class, as are 512- and 1,024-bit memory buses.

Vector supercomputers

Vector supercomputers make unique demands on their memory systems. As a consequence of the very large problem sizes that these computers work on, even large caches are of little value. Therefore, performance depends heavily on the ability of the memory system to provide multiple streams of vector data simultaneously at very high bandwidths. Each processor is capable of absorbing two double-precision operands and

9. In fact, many recent medium to large machines use many microprocessors, operating in concert with a powerful I/O system, to effectively manage very large data sets.

Samsung Ex. 1097, p. 108
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

generating a double-precision result every cycle. Chaining functional units together increases the number of operands per cycle that the memory unit must provide.

Supercomputer memory systems are very heavily interleaved—up to 512-way interleaved—with multiple paths from the subbanks to the memory controller and on to the processors. The whole subsystem is engineered to minimize the number of resource conflicts between otherwise independent read and write accesses, so that vector processors can operate at full speed the majority of the time. For most vector supercomputers, many cycles of latency are required to start up a vector operation. The emphasis for memory system design is clearly on high bandwidth and low contention for memory. The trade-off between the use of SRAM and DRAM for main memory is driven mainly by economics, performance, and, to a lesser extent, by the physical implications of a lower density SRAM memory system. All midrange to high-end supercomputers have used SRAMs in their main memories. For the lower performance, entry-level machines, designers have occasionally chosen DRAMs to improve affordability.

Massively parallel machines

In many respects, massively parallel processors (MPPs) represent the other end of the spectrum from vector supercomputers. In an MPP, memory is not a single large, complex subsystem. Instead, memory is divided into small chunks and is distributed throughout the machine; it is associated with either individual processors or small groups of processors. The amount of memory under the control of each autonomous memory controller varies significantly according to the system architecture. One class of MPP uses very large numbers of simple processors, small clusters of which share memory systems consisting of a few DRAMs each. With this class of MPP, the problem is not getting massive amounts of bandwidth out of a large number of DRAMs, but squeezing as much bandwidth as possible out of a very small number of RAMs while retaining a relatively simple memory controller.

Fill frequency requirements

The fill frequency requirements of large systems are difficult to categorize because of the wide spectrum of machine architectures and application domains. Furthermore, with the partitioning of main memory into primary and secondary stores, a hierarchy of fill frequency requirements is introduced within a single machine category.

That said, generally, many large systems rely on very large main memories to completely hold the problem or problems

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 109
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

being worked on. In contrast, PCs and—to a much lesser extent—workstations rely on virtual memory to provide the illusion of large physical memory. Consequently, the granularity requirements in large systems are substantially larger than those in other system classes; the granularities are larger even relative to the increased bandwidth requirements. As a result, the fill frequency requirements for the main memories of large general purpose machines are typically lower than those of workstations and PCs.

4.7 Graphics Memory Systems

The visual display of information in many forms is an increasingly important component of what computers do. Graphics controllers are responsible for 1) maintaining an image of the screen's current contents; 2) accepting commands that describe at some level of abstraction changes to the appearance of the screen; and 3) interpreting those commands by changing the controller's internal image of the screen. Also, the contents of the internal image must be continuously translated into the electrical signals needed to drive the physical display. All graphics controllers have a small graphics memory system that encodes the screen's current appearance.

When a one-to-one correspondence exists between locations in this local memory and the appearance of individual pixels on the screen, the memory is said to contain a **frame buffer**. Depending on the number of colors that can be displayed at once, between one and thirty-two bits are used to describe the appearance of each pixel. One bit per pixel implies a monochrome display; eight bits per pixel implies a choice among 256 colors for each pixel; and twenty-four bits per pixel, called *true color*, allows each pixel to take on any of 16.7 million colors.

Part of the graphics controller's job is to regularly scan the entire contents of the frame buffer and drive the video monitor correspondingly. This task takes up a portion of the video memory's total available bandwidth. Table 4-2 lists both the peak and the average amount of bandwidth needed to refresh the screen for a variety of video display configurations. The peak bandwidth is the rate at which data must be fed into the RAMDAC while the CRT's electron beam is drawing the visible portion of the screen. The average bandwidth is computed by multiplying the frame buffer size by the refresh rate. The

Samsung Ex. 1097, p. 110
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

average bandwidth with compression column is discussed in Section 4.8. The average bandwidth and the peak bandwidth differ because pixels are not drawn during the horizontal and vertical retrace periods. Peak bandwidths are 50–100% higher than average bandwidths—though the exact ratio depends on the display configuration.

Low-end graphics controllers use generic DRAMs

Figure 4-13 shows a coarse block diagram of a graphics controller that uses generic DRAMs to hold the frame buffer. The graphics controller uses the single port on the DRAM to make changes to the display contents and to retrieve the frame buffer for display on the screen. Regular generic DRAMs work acceptably as video memories for low-performance graphics with small resolutions and color depths because the total bandwidth required of the frame buffer memory is less than the 80–100 Mbytes/s that can be obtained from a ×32 memory system. Some of the more recent graphics controller chips can make use of ×64 memory systems to generate even more bandwidth. Recently, midrange graphics applications have been using high-speed DRAMs to generate more bandwidth than is possible with regular generic DRAMs, but without the cost and design effort of a VRAM-based solution. Other recent trends are the integration of the RAMDACs onto the same device as the

Name	Resolution	Bits per Pixel	Colors	Frame Buffer Size	Refresh Rate	Peak Display Bandwidth	Average Display Bandwidth	Average Display Bandwidth w/ Compression
VGA	640 ×480	8	256	300 Kbytes	60 Hz	33 Mbytes/s	18 Mbytes/s	1.2-1.8 Mbytes/s
	640 ×480	16	64K	600 Kbytes	60 Hz	66 Mbytes/s	36 Mbytes/s	2.4-3.6 Mbytes/s
SVGA	800 × 600	8	256	480 Kbytes	60 Hz	52 Mbytes/s	29 Mbytes/s	1.9-2.9 Mbytes/s
	800 × 600	16	64K	960 Kbytes	70 Hz	122 Mbytes/s	67 Mbytes/s	4.4-6.7 Mbytes/s
	1,024 × 768	8	256	786 Kbytes	70 Hz	100 Mbytes/s	55 Mbytes/s	3.6-5.5 Mbytes/s
	1,024 × 768	16	64K	1.6 Mbytes	70 Hz	200 Mbytes/s	110 Mbytes/s	7.3-11 Mbytes/s
	1,024 × 768	24	16M	2.4 Mbytes	72 Hz	300 Mbytes/s	170 Mbytes/s	11-17 Mbytes/s
	1,280 × 1,024	8	256	1.3 Mbytes	72 Hz	172 Mbytes/s	94 Mbytes/s	6-10 Mbytes/s
	1,280 × 1,024	16	64K	2.6 Mbytes	72 Hz	344 Mbytes/s	189 Mbytes/s	12-19 Mbytes/s
	1,280 × 1,024	24	16M	3.9 Mbytes	72 Hz	516 Mbytes/s	283 Mbytes/s	19-28 Mbytes/s
	1,600 × 1,200	24	16M	5.8 Mbytes	72 Hz	756 Mbytes/s	415 Mbytes/s	27-41 Mbytes/s
HDTV	1,920 × 1,080	24	16M	6.2 Mbytes	72 Hz	816 Mbytes/s	448 Mbytes/s	N/A

Table 4-2 The display bandwidth grows with the screen resolution, color depth, and refresh rate.

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 111
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

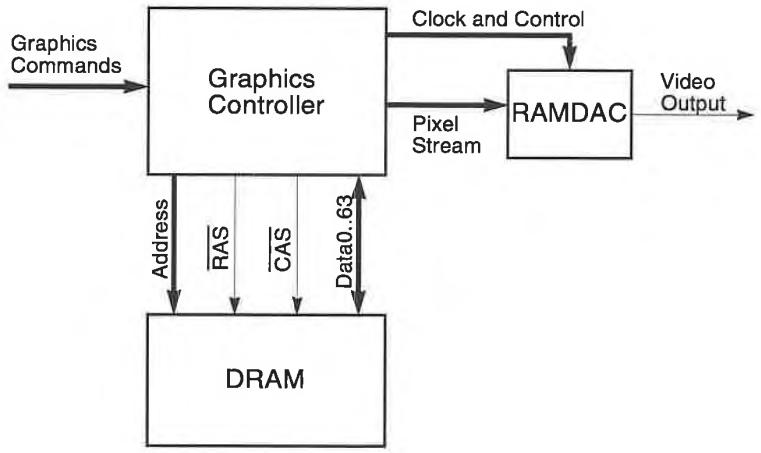


Figure 4-13 Low-end graphics controllers use generic DRAMs.

graphics controllers and the addition of more GUI acceleration into the graphics controller to help relieve the CPU of some of the more mundane display management tasks.

High-end graphics need VRAMs

When higher resolution, more colors, or better video performance is desired, standard DRAMs are no longer adequate for the task. For these systems, **Video RAMs (VRAMs)** are used. As described in Section 3.6, VRAMs are two-port devices designed for use in graphics controllers. In addition to an augmented standard DRAM interface, they have a high-speed serial interface tailored for dumping the frame buffer to the screen. Figure 4-14 shows a block diagram of a graphics controller using VRAMs.

When still higher performance is required, as in the case of real-time 3-D object rendering, even more complex hardware structures are put in place to speed the interpretation of high-level graphics commands from the processor. For example, a Z-buffer can be added to the frame buffer to improve 3-D performance. A Z-buffer holds depth information associated with each pixel to help the controller decide whether some new object being drawn is in front of or behind the object already drawn at that location. Another feature that improves high-performance 3-D is double-buffering. Double-buffering provides one copy of the frame

Samsung Ex. 1097, p. 112
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

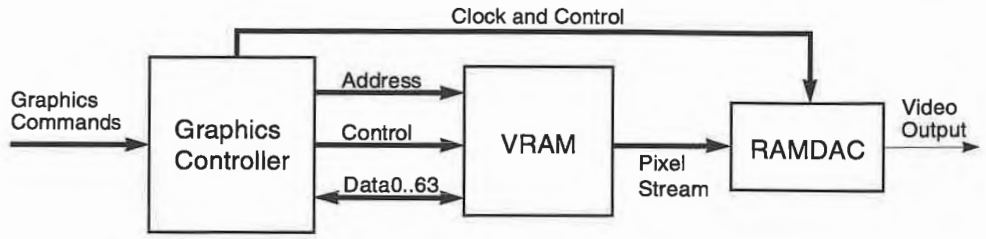


Figure 4-14 Midrange graphics controllers use two-ported VRAMs to increase the frame buffer bandwidth.

buffer for displaying to the screen and a second copy for rendering a scene. Once a scene is fully drawn into the second copy, the two buffers exchange roles, and the display hardware dumps the completed scene to the screen. Two frame buffers can be used in virtual reality applications; the two buffers are used to generate separate images for the left and right eyes, so that the user perceives depth properly.

Tiling improves DRAM utilization

An increasingly important question in the design of frame buffers is the way in which pixels on the screen are mapped onto bytes in the physical memory (i.e., in the frame buffer). One method is natural and intuitive: the screen consists of a rectangular collection of rows (i.e., scan lines), while DRAMs are a rectangular collection of rows of bits. It seems natural then to map scan lines onto DRAM rows. This arrangement, shown in Figure 4-15, was used for many years.

This straightforward mapping is very efficient for screen refresh. A single row access brings in a big chunk or all of a scan line into the row cache, from which it can be dumped to the RAMDAC with column accesses. Unfortunately, though, this arrangement is not very efficient for rendering. All graphical objects except for horizontal lines cross multiple scan lines. Drawing a vertical line requires writing only a few pixels—as few as one—on each scan line. Thus, drawing a vertical line becomes a sequence of row accesses with a single column write access following each. This request sequence has the lowest possible efficiency and yields the lowest sustained bandwidth.

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 113
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

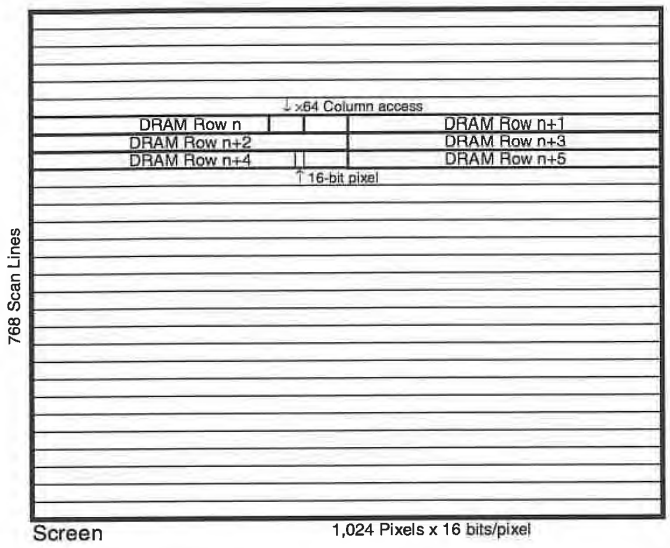


Figure 4-15 In a traditional, untiled frame buffer, each DRAM row is mapped onto a single scan line. In this example, each DRAM row represents a portion of the screen that is one scan line high and 512 pixels across.

Although the traditional, untiled method was used for many years, a tiled frame buffer offers significantly improved performance. In a tiled frame buffer, rectangular regions of pixels are mapped onto DRAM rows. These regions, or tiles, are typically up to eight scan lines tall, and 256 pixels across—plus or minus a factor of two in both directions. The size and best aspect ratio for the tiles depends on the logical layout of the DRAMs being used (i.e., the size of their row cache), the screen resolution, and the color depth.

With a tiled frame buffer, the screen refresh and best case rendering (i.e., when drawing horizontal lines) performance is not noticeably worse than with the untiled mapping. However, the worst case rendering performance is very much better. When drawing a long, thin vertical line into tiles eight scan lines high, each row access is followed by eight column accesses. The worst case sustained bandwidth with a tiled frame buffer is about three times greater than with an untiled one. The average sustained bandwidth is generally improved by 50–100% with the introduction of tiling. The magnitude of the improvement

Samsung Ex. 1097, p. 114
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

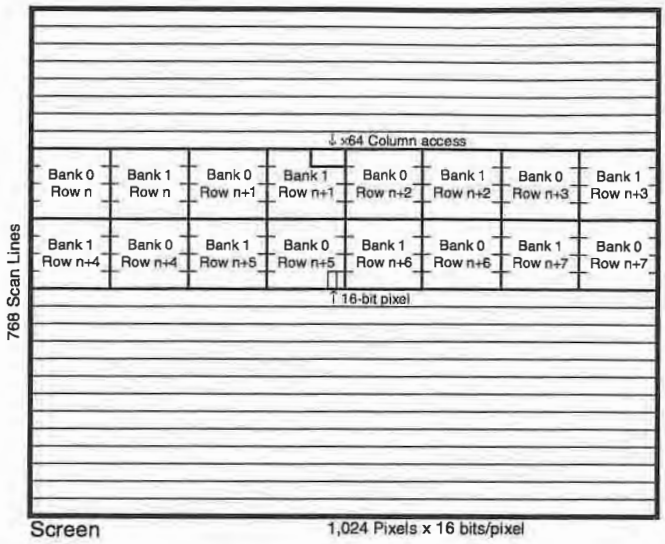


Figure 4-16 A tiled frame buffer has rectangular regions—here 4 pixels by 128 pixels—mapped onto DRAM rows.

depends on the DRAMs, the tiles, the screen resolution, and the objects being rendered.

A final optimization is possible when the frame buffer memory system has several banks. If horizontally and vertically adjacent tiles are placed in different banks, then one bank can perform a row access while another performs column accesses. This technique can hide a substantial amount of the total row access time, further increasing the sustained rendering performance. This tiled frame buffer structure is depicted in Figure 4-16.

Tiling a frame buffer has one drawback: mapping from the x-y screen coordinates to a DRAM row and column address is more complicated than with a untiled frame buffer. However, GUI accelerators need to support a variety of unusual screen resolutions, and the additional complexity in their address generation logic needed to support tiling is insignificant.

Multimedia more demanding of PC graphics

The growth of multimedia applications increases demands on PC graphics controllers. For example, the information content of an NTSC video signal is about 7.5 Mbytes/s, assuming 8-bit pixels. The ability to process and render a video signal onto a portion of the screen requires at least that much free frame

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 115
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

buffer bandwidth plus enough controller complexity to map the NTSC signal into a reduced pixel and color space.

The trend in the design of small portable systems is toward a unified memory system. Rather than having two separate physical memories—one for the main memory and one for the graphics memory—the logical memories are combined into a single physical unit. Unified Memory Architecture (UMA) memory systems are described in detail in Section 4.8.

High fill frequency requirements

Frame buffers—and graphics memory systems in general—have the highest fill frequency requirements of any of the main DRAM application domains. This is because the size and granularity of these memory systems are small—typically 1–4 Mbytes—and the peak and sustained bandwidth requirements are extreme. In 1995, at the low end of the GUI accelerator spectrum, the acceptable granularity was 1 Mbyte, while the peak bandwidth requirement of 100 Mbytes/s derives from a 32-bit memory width and a 25 MHz operating frequency for fast-page-mode DRAMs. These yield a peak fill frequency requirement of 100 Hz. At the high end of the performance spectrum, the minimum granularity is increased to 2 Mbytes, but the peak bandwidth requirement is increased to approximately 400 Mbytes/s: the result is a required peak fill frequency on the order of 200 Hz.

The bandwidth requirements for high-performance 3-D graphics systems are generally substantially higher than the requirements in the 2-D arena. This increase in the peak bandwidth requirement more than offsets the increase in granularity due to double-buffering and the addition of a Z-buffer. Consequently, the ratio of the bandwidth to minimum memory size—i.e., the peak fill frequency requirement—for 3-D graphics is larger by a substantial margin than the ratio for all other significant memory system types.

The high fill frequency requirements of graphics memory systems suggest why the graphics world is still using DRAMs that are 4-Mbit and smaller while the main memory world has shifted to the more cost-effective 16-Mbit devices. As illustrated in Section 2.5, none of the generic 16-Mbit devices' fill frequencies are adequate to meet the graphics market requirements.

Samsung Ex. 1097, p. 116
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

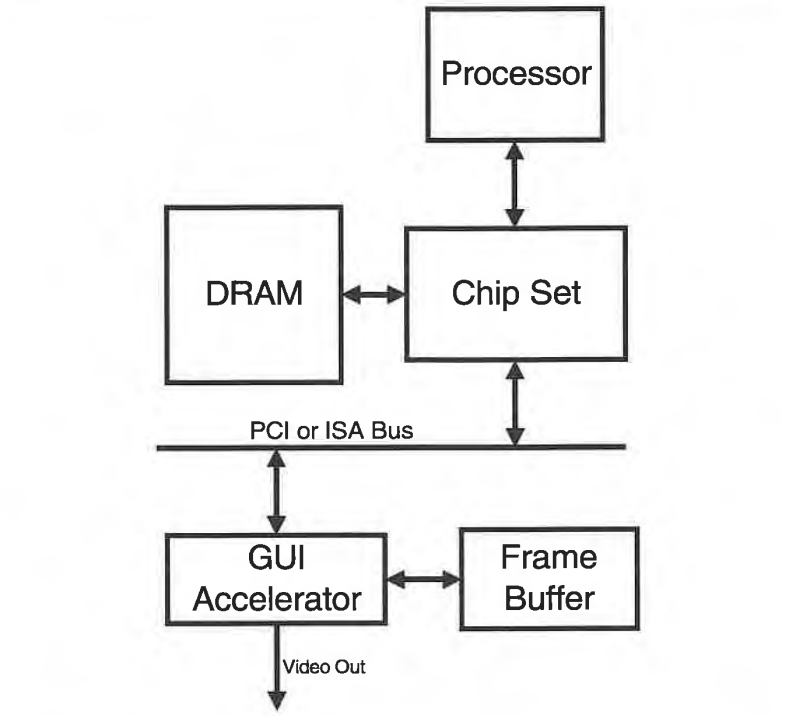


Figure 4-17 The traditional PC graphics architecture.

4.8 Unified Memory Architecture (UMA)

In the traditional PC architecture, the graphics controller is physically and logically separate from the central portion of the system: it sits out on the ISA or PCI bus, on a separate daughtercard inserted into a socket on the motherboard. Consequently, most PCs contain two separate DRAM memories: 1) a main memory array, constructed of SIMMs and controlled by the chip set, and 2) a frame buffer memory, usually soldered down or socketed on the graphics adapter card. Figure 4-17 illustrates this traditional PC system architecture.

The central idea behind **Unified Memory Architecture (UMA)** systems is to house these two logical structures (the main memory and the frame buffer) in a single memory array on the motherboard. The goal is to reduce total cost, increase

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 117
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

functionality, or, preferably, do both. The single memory array is shared and divided between the two logical memories so that the sum of the size of the main memory and the size of the frame buffer is equal to the total amount of memory physically present.

New cost-performance trade-offs

A switch to UMA reduces total cost. Total memory cost is reduced because the graphics memories used to make frame buffers are more expensive per bit than main memory DRAMs. Graphics memories are either lower density or have special functionality (as in VRAMs) to provide the bandwidth and granularity (or equivalently, the fill frequency) needed by graphics applications. Replacing graphics memories with cheaper main memory devices reduces total cost.

Alternatively, the frame buffer bits can be eliminated altogether. Instead of having 16 Mbytes of main memory and 2 Mbytes of frame buffer memory, a system could have 16 Mbytes of total memory, conceptually divided into 14 Mbytes of main memory and 2 Mbytes of frame buffer. The consumer, seeing 16 Mbytes in the advertisement, may not be fully cognizant that some of those bits are not available for general purpose computing.

Reserving some memory for the frame buffer results in some performance loss compared to a system with a separate frame buffer. Obviously, the magnitude of this performance loss depends on both the total amount of memory and the amount of memory taken up by the frame buffer. As the amount of main memory left for the operating system and applications decreases, the number of page faults increases and the user spends more time waiting on the disk. If 1 Mbyte is taken for the frame buffer from a cached 8-Mbyte system, the total performance loss is 5–10%. Stealing 2 of 8 Mbytes of memory for the frame buffer is very noticeable. However, when the total memory size is 16 Mbytes of memory, the performance loss from reserving 2 Mbytes for the frame buffer is again less than 10%. Without a secondary cache, the performance loss associated with the smaller main memory size is increased. In this case, taking 1 Mbyte from an 8-Mbyte system results in a 10–15% total performance degradation.

Of course, the magnitude of the performance loss is highly dependent on the application, the problem size, and the operating system. In some circumstances, the performance loss due to the loss of memory available to the operating system and application will be much greater than the amounts indicated

Samsung Ex. 1097, p. 118
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

here. Similarly, in some situations the loss of main memory will not have a noticeable impact on end-user performance.

In a UMA system, contention over the single memory resource is another source of performance loss. In particular, if the memory array is busy performing graphics operations when a processor cache miss occurs, the processor may stall longer than it would in a traditional architecture machine. Similarly, if the graphics engine is precluded from accessing the memory because the processor is using the memory a lot, then graphics display will not be as snappy. However, queuing theory analyses indicate that so long as the amount of a shared resource—memory bandwidth in this case—is more than marginally greater than the total demand, the average delay from contention should be quite small.

Additional cost savings are effected through eliminating sockets, daughter PC boards, and—in some cases—one or more of the ICs needed to implement the graphics adapter.

One nice aspect of UMA architectures is that only the memory actually used for the frame buffer is purchased. In a traditional graphics daughter card, the graphics memory is binary-sized[10] and must be large enough to hold the frame buffer for the highest supported resolution and color depth. When a lower resolution is used, a portion of the graphics memory system goes unused. For example, a 4-Mbyte frame buffer memory array is required by a graphics card that supports 1,024×768 displays and 24-bit pixels (see Table 4-2). However, whenever the user is in 8-bit VGA mode, only 300 Kbytes of that 4 Mbytes is used; 3.7 Mbytes are unused or "wasted." With a UMA system, the system can be booted in VGA mode with a small frame buffer. This leaves the maximum amount of memory for main memory. Ideally, display resolution could be changed dynamically and the logical boundary between main memory and the frame buffer could be changed on the fly while the system is operating.

In some circumstances, UMA architectures can improve performance over the traditional split architectures. For example:

- When the unified memory system is of a higher performance level than is typical for PC main memories, the average response time to CPU memory requests can improve.
- When some or all of the displayed images originate in main memory—for example, as in a decompressed video signal or

10. The number of bytes is an integral power of 2.

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 119
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

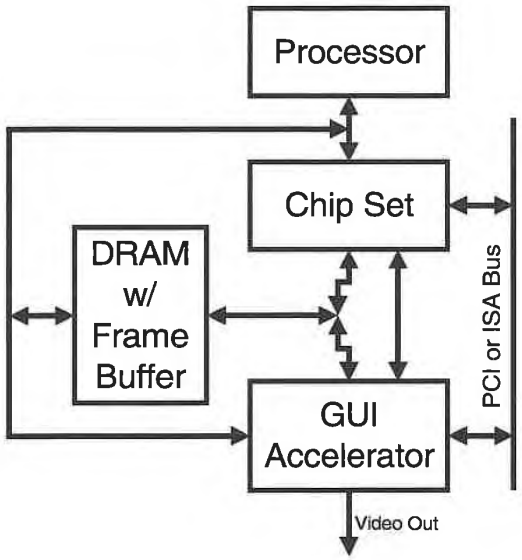


Figure 4-18 A loosely-coupled UMA system.

a texture-mapped scene—then those bits need not traverse the slow and sometimes overcommitted PCI or ISA bus.

Two styles of UMA systems

There are two styles of UMA machines: loosely-coupled (or dual-supplier) and tightly-coupled (or single-supplier). These two styles are shown in Figures 4-18 and 4-19, respectively.

A loosely-coupled (dual-supplier) UMA system contains the same ASICs as a traditional PC contains. However, these ASICs are interconnected differently. A GUI accelerator chip—possibly made by a different vendor than the chip set supplier—shares control of the unified memory array with that chip set.

As of 1Q1996, a three-pin interface connecting the chip set and the GUI accelerator is being standardized by the Video Electronics Standard Association (VESA). The standard provides for the GUI accelerator to request access to the memory array. Having two request priorities reflects the necessity of keeping screen refresh data flowing to the RAMDAC. When a request is granted, the chip set tri-states all the memory array control signals, and the GUI accelerator takes control until it either completes the task or is preempted by the chip set. Several

Samsung Ex. 1097, p. 120
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

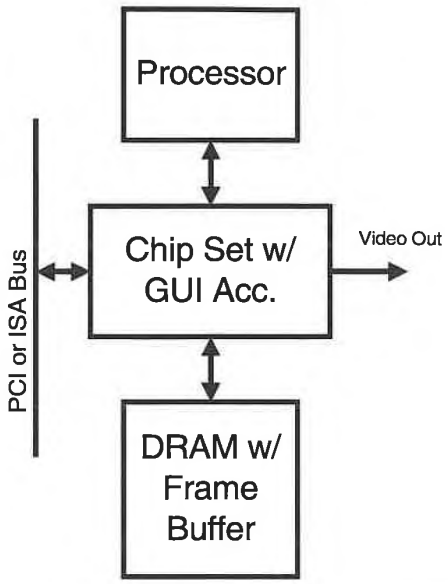


Figure 4-19 A tightly-coupled UMA system.

companies are promoting a slightly extended interface. The addition of another signal from the GUI accelerator to the chip set facilitates three levels of request priority. This allows an urgent need for screen refresh data (i.e., a nearly empty refresh data FIFO) to be differentiated from a nonurgent need (i.e., the refresh data FIFO is still one-half full).

An important part of the VESA standard are the software protocols that allow the system software modules controlling the two logically separate functional units—the BIOS for the GUI accelerator and the BIOS for the chip set—to communicate and share information about the size, type, and speed of the memory and cooperate on the allocation of physical bits to the two logical arrays.

Advantages and disadvantages of loosely-coupled UMA systems

Loosely coupled UMA systems have both advantages and disadvantages relative to the standard PC architecture. A main advantage is that a range of PC products can be developed by inserting different GUI accelerators into a socket on a standard motherboard. This differentiation is similar to that currently common in the product lines of the large PC suppliers: a

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 121
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

low-end system with smaller memory and disk has a cheaper, lower performance graphics subsystem, while the high-end product with the same processor[11] (and motherboard, chassis, and everything else) has more memory, a larger disk, a better CD-ROM, and a better graphics controller and monitor. The ability to efficiently and cheaply create a range of PC products is vital in the current PC marketplace.

The dual-supplier architecture enables a company with a single competency or specialization to offer products for UMA systems. This benefits users by increasing the range of product choice and by increasing competition.

A disadvantage of the loosely-coupled UMA architecture is that graphics data flowing from the CPU to the GUI accelerator must still traverse the PCI bus. Because only one $\overline{\text{RAS}}$ signal is shared between the two units, the GUI accelerator does not necessarily have access to all of memory.

A second disadvantage is the loosely-coupled UMA model's assumption that all banks are closed (i.e., $\overline{\text{RAS}}$ is raised) whenever control is transferred between the two controllers. This assumption is not a problem with traditional PCs and DRAMs: because at most one bank can be kept open at a time (see Section 4.4), little is gained by keeping banks open across the transition of control. However, with several new DRAM architectures, many banks can be kept open simultaneously. If these new DRAMs are to be used effectively, the banks must be kept open over long periods of time.

A third disadvantage of this organization is that there is necessarily additional loading on the DRAM control signals. The additional loading is probably not problematic with current fast-page-mode DRAMs. However, with the higher operating frequencies common among the new DRAM architectures, it is more important to keep the loading on the DRAM control lines to a minimum. Also, some new memory architectures do not easily facilitate having two masters share control of memory.

Perhaps the greatest disadvantage of the loosely-coupled UMA architecture relative to the standard PC architecture is that it is harder to develop a range of chip products that facilitate a range of system-level products.

11. High-end systems often have a higher clock-frequency version of the same processor.

Samsung Ex. 1097, p. 122
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

Tightly-coupled UMA systems

In a tightly-coupled (single-supplier) UMA system, both the chip set and the GUI accelerator are integrated into a single ASIC or group of ASICs: they are no longer two physically separate units. In order to design such a chip set, a company must have competency in both graphics accelerator design and chip-set design—or two companies must cooperate on design and distribution.

One advantage of tightly-coupled UMAs is that as levels of integration and pin counts increase, the total number of chips in the system can decline. This leads to board area and power savings as well as the obvious savings in the bill-of-materials. In addition, a single memory controller enables easier management of high-performance DRAMs to get the most out of multiple banks and pipelined or concurrent transfers. Furthermore, with a single controller, the loading on the DRAM address and control signals is minimized, and so the operating frequency is maximized.

UMA architectures compared

Table 4-3 contrasts the strengths and weaknesses of the three basic PC architectures.

Characteristic	Traditional PC	Loosely-Coupled UMA	Tightly-Coupled UMA
Cost	Two memory arrays Two controllers	One memory array Two controllers	One memory array One larger controller
Performance impact	Baseline	Tolerable	Tolerable
Supported resolutions	Limited by size of frame buffer	Not limited	Not limited
Partitioning difficulties	None	Negligible	Severe in 64-bit system w/o BGA packaging
Ability to choose chip set and GUI accelerator	Completely independent	Independent but linked	One constrained choice
Main memory to frame buffer transfers	Via PCI	Via PCI	Memory to memory
Simultaneous main memory and frame buffer operation	Yes	No	No
Optimizing memory controller(s)	Yes; Independent	No	Yes
DRAM control signal loading	Typical for main memory	Increased	Typical for main memory

Table 4-3 The two forms of UMA architecture for the PC have different strengths and weaknesses.

High peak bandwidth requirements

The average bandwidth requirements placed on unified memory systems are essentially equal to the sum of the average main memory and graphics requirements. The true peak bandwidth requirement is actually the maximum of the individual main memory and graphics peak bandwidth requirements. However, since peak and average bandwidths are related, the realistic

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 123
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

peak bandwidth requirements lie between the maximum of the two individual components and their sum.

As with any system, bandwidth requirements are soft. As bandwidth demands approach the total available bandwidth, memory utilization increases, average access time increases, and system performance degrades. As the performance of the system degrades, bandwidth demands also decrease until the supply and demand for bandwidth are balanced. This degradation of performance occurs smoothly and gradually unless the supply of bandwidth is grossly inadequate.

It is also true that many applications are quite bimodal in their behavior: they go through periods of heavy computation alternating with periods of heavy display activity. For these applications, both the amount of interference between the CPU and the graphics demands on the unified memory system would be minimal because the CPU and GUI accelerators would not need their peak bandwidth requirements simultaneously.

As noted in Section 4.7, the bandwidth requirements from the graphics subsystem vary linearly with display resolution and color depth. The average screen refresh bandwidth—as a function of display resolution and frequency—is shown in Table 4-2. For resolutions above 800×600 with 16 bits of color, the amount of bandwidth needed just to dump the contents of the frame buffer to the RAMDAC is daunting. If screen refresh takes up too much of the total available sustainable bandwidth, then too little is left for the drawing of graphics and for the CPU. The CPU and GUI accelerator then compete for the remaining available bandwidth, and both CPU and graphics performance start to degrade. Unfortunately, screen refresh must have the highest priority when competing for a limited amount of total bandwidth: if the supply of pixels to the RAMDAC falters—if garbage is even occasionally drawn to the screen—user perception of display quality deteriorates rapidly.

Frame buffer compression dramatically reduces bandwidth requirements

Fortunately, the actual information content of the pixel stream to the RAMDAC is dramatically less than its bandwidth. Consequently, the pixel stream and, by extrapolation, the frame buffer, are highly susceptible to compression. Compression of the frame buffer can dramatically reduce the bandwidth needed for screen refresh. The last column of Table 4-2 on page 111 shows the amount of frame buffer bandwidth needed if the frame buffer contents are compressed by a factor of 10 to 15, as is typical in many applications.

Samsung Ex. 1097, p. 124
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

Frame buffers are highly compressible because of two independent phenomena:

1. Groups of horizontally and vertically adjacent pixels commonly have the same or nearly the same color.

2. Most—if not all—of the screen is typically unchanged from one screen refresh to the next.

The first fact means that the information in the frame buffer is highly redundant and can be compressed by a wide variety of general purpose and customized compression algorithms. The second fact means that once a small region of the frame buffer is compressed, the compressed version will most probably be reused many times before that portion of the frame buffer is redrawn.

To illustrate how frame buffer compression can work, one possible algorithm is described below:

1. The frame buffer is conceptually divided into many regions, each most likely being a portion of a scan line. Each region is compressed as a unit. For each region, some storage is needed in addition to the uncompressed frame buffer. This separate storage, for all the regions together, constitutes a compressed version of the frame buffer.

2. If the pixel in the uncompressed frame buffer is changed by the GUI accelerator, then the compressed version of the region to which the pixel belongs is marked as invalid or old.

3. When performing a screen refresh of a region, the compressed version is checked. If it is invalid, then the uncompressed frame buffer is read and sent to the RAMDAC. In the process, it is compressed by one or more compression algorithms. The compressed result is placed in the compressed frame buffer and marked as valid.

4. If, when performing a screen refresh of a region, the compressed version is checked and found valid, then the compressed version is read out from the compressed frame buffer, decompressed, and sent to the RAMDAC.

The overhead associated with such a scheme is the cost of maintaining the compressed frame buffer. This cost is in the form of memory (the size of the compressed memory), the bandwidth needed to maintain the compressed frame buffer, and the gates needed to implement the compression and decompression

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 125
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

algorithms. With appropriately sized regions, all three cost components can be kept small.

One potential problem associated with keeping a compressed copy of the frame buffer is the bandwidth needed to invalidate the compressed representation of the regions every time a pixel is modified. However, since pixel writes are very spatially clustered, caching of recent invalidations works very well to reduce this burden.

As mentioned earlier, frame buffer compression works because screen contents are typically very locally uniform and generally very stable. Business applications in particular have very low rates of pixel change: as the user types in a word processor or spreadsheet, only the pixels representing the words or numbers being modified are changing. But even for applications such as games that involve a lot of dynamic screen activity, much of the screen is essentially static from one scene change to another.

Frame buffer compression less effective with video and 3-D graphics

Frame buffer compression works less well for some applications and application domains. One such application is video. A typical video stream writes every pixel in the video window at a rate of 30 Hz. Even though many of the pixel values are unchanged from one video frame to another, recognizing this fact requires either comparing the new pixel values with the old or computing a hash function on the new pixels in a region and comparing that to the hash value stored in the compressed frame buffer. In the absence of these techniques, a full screen video window would write every pixel every other screen refresh period or so. Such a consistent and high rate of pixel writing greatly reduces the effectiveness of frame buffer compression.

Frame buffer compression is dramatically less effective with 3-D graphics. Unlike 2-D applications, 3-D applications involve smoothly shaded polygons. Furthermore, continuous movement of the background is common. These two characteristics mean: 1) compressing the frame buffer is harder since there are gradients of color across much of the screen, 2) many more pixels are written every screen refresh interval, and 3) many more pixels change value from one screen refresh to the next.

However, the ineffectiveness of frame buffer compression in a 3-D environment is a moot point: the high memory bandwidth demands of 3-D graphics make it unlikely that midrange to high-performance 3-D systems will use a UMA architecture in the first place.

Samsung Ex. 1097, p. 126
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

UMA to enter the market at the low end

As long as total memory bandwidth is constrained by the use of fast-page-mode and EDO DRAMs, UMA architecture will be viable only in low-end PCs. This is because low-end PCs have reduced display resolutions and color depths compared with their midrange and high-end cousins. The low-end marketplace is also the domain where UMA's cost savings are the most significant, the performance penalty is least important to the consumer, and the total bandwidth requirements are lowest. Over time—as higher bandwidth memories become more available and cost-effective—UMA machines are likely to migrate up into the middle of the PC product range. The high end's specialized graphics requirements mean UMA is unlikely to penetrate this market segment in the next few years.

Microsoft actively discourages UMA systems

But the biggest barrier inhibiting the introduction of UMA systems is Microsoft. In 4Q1995 Microsoft began strongly discouraging the PC system vendors from introducing UMA-based systems. The main reason for Microsoft's stance was the performance degradation due to the reduction in main memory size. Windows 95 is tuned for 8 Mbytes of main memory—not 7 or 7.3 Mbytes. But also of significance is Windows 95's ability to adapt to changes of screen resolution on the fly, and its inability to deal with changes in main memory size without a full reboot. Unfortunately, in UMA systems these two go together since the size of main memory and the size of the frame buffer sum to a constant. Microsoft is also very concerned about the success of the PC as a game platform. Games using the DirectDraw and Direct3D APIs make full use of the available off-screen frame buffer store to improve rendering performance. In UMA systems, the tendency is to keep off-screen frame buffer storage to a minimum in order to maximize the size of main memory. The bottom line is that at least in North America, the significant adoption of UMA in the PC marketplace will wait until Microsoft gives its blessing.

Intermediate fill frequencies

UMA memory systems lie between traditional main memory and frame buffer memories in terms of design difficulty. That is because UMA's bandwidth requirements are essentially the sum of the main memory and frame buffer requirements while the granularity requirements are comparable to those of main memories for traditional PCs. Consequently, peak fill frequencies for low-end UMA machines are in the range of 70–100 Hz: higher than the 20–30 Hz of traditional PCs, but less than the 100–300 Hz of typical graphics adapters.

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 127
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

Samsung Ex. 1097, p. 128
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

5 New DRAM Architectures

5.1 New In The Second Edition

Information on these topics has been added since the first edition:

- Packaging types for fast-page-mode 64-Mbit devices
- HSTL and SSTL electrical interfaces
- Burst EDO DRAMs as a new evolutionary option
- SyncLink as a new revolutionary option

5.2 Potential Improvements to the Existing Interface

The changes occurring in the DRAM marketplace are driven by exponential growth in DRAM density. The difference between the growth rate of DRAM density and the growth of main memory size and bandwidth per pin is making the memory system designer's job increasingly difficult (see Chapter 2). Already, the standard DRAM interface is inadequate for several significant market segments. As the industry continues its inevitable evolution to the 16- and 64-Mbit generations, the number and importance of the markets seeking alternatives to the standard narrow interface will grow dramatically.

Samsung Ex. 1097, p. 129
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

Consequently, tremendous energy is being devoted to finding cost-effective alternatives that provide the features, bandwidth and fill frequencies required by these markets.

This chapter discusses the opportunities for improving the existing DRAM interface. Specifically, the current interface is decomposed into five basic elements:

- The physical interface
- The electrical interface
- The pin-level interface
- The temporal interface
- The logical interface

For each of these five elements of the interface, alternatives exist that are more or less appropriate for different market segments. Depending on the type and number of changes in the interfaces, the new DRAM architectures are characterized as either conventional, evolutionary, or revolutionary.

The five aspects of the existing DRAM interface can be varied somewhat independently. Each of the following five sections describes one aspect and discusses the choices associated with it.

5.2.1 The Physical Interface

A DRAM's **physical interface** is the physical and electrical specification of the package: its size, shape, and thickness; the number, type, and spacing of the pins; and the electrical characteristics of the package leads from the solder connections on the board to the bonding pads on the die.

DRAM packages

Four generations of DRAMs—beginning in 1974 with the 4-Kbit generation and ending with the 256-Kbit devices—were packaged in 300 by 800 mil, 16-pin **Dual In-Line Packages (DIPs)** (see Figure 5-1). Because surface-mount assembly replaced through-hole packaging as the manufacturing technology of choice in the mid-1980s, **Small-Outline J-Lead (SOJ)** packages are now more appropriate for applications that use unsocketed DRAMs. SOJ packages are roughly comparable to DIPs, except that the pins are spaced every 50 mils (1.27 mm) along the two long edges instead of every 100 mils (2.54 mm), and the pins are curled up under the package in a J shape instead of

Samsung Ex. 1097, p. 130
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

extending down through the PC board. The small gap created by the J-leads between the package bottom and the board provides a convenient space for low-profile decoupling capacitors.

Recently, the desire for tighter **SIMM** spacing has motivated the use of **Thin Small-Outline Packages (TSOPs)**. TSOPs are small rectangular packages as thin as 44 mils (1.12 mm), with pins every 20 mils (0.55 mm) along the two shorter edges. TSOPs provide reasonable package sizes up to roughly 32 pins. TSOP II packages have pins along the long edge: the pins are spaced every 50 mils (1.27 mm) for pin counts up to 32 and every 31.5 mils (0.8 mm) for larger pin counts. With TSOPs, the SIMM-to-SIMM spacing can be as small as 110 mils (2.79 mm)—much smaller than the 200-mil (5.08 mm) spacing for SOJ-based SIMMs.

Since the 256-Kbit generation, applications needing tightly packed DRAMs have used vertical packages or multi-DRAM memory modules. The **Zigzag In-Line Package (ZIP)**—an older, through-hole vertical package—has leads along one edge every 50 mils (1.27 mm) and adjacent pins offset by 100 mils (2.54 mm) in a zigzag pattern. More recent vertical packages include the new **Surface-Mount Vertical Package (SVP)**. The SVP package, shown in Figure 5-1, has leads every 26 mils (0.65 mm) down one edge and larger surface-mount power leads at either end to provide support during assembly.[1] This package facilitates mounting DRAMs vertically every 100 mils (2.54 mm).

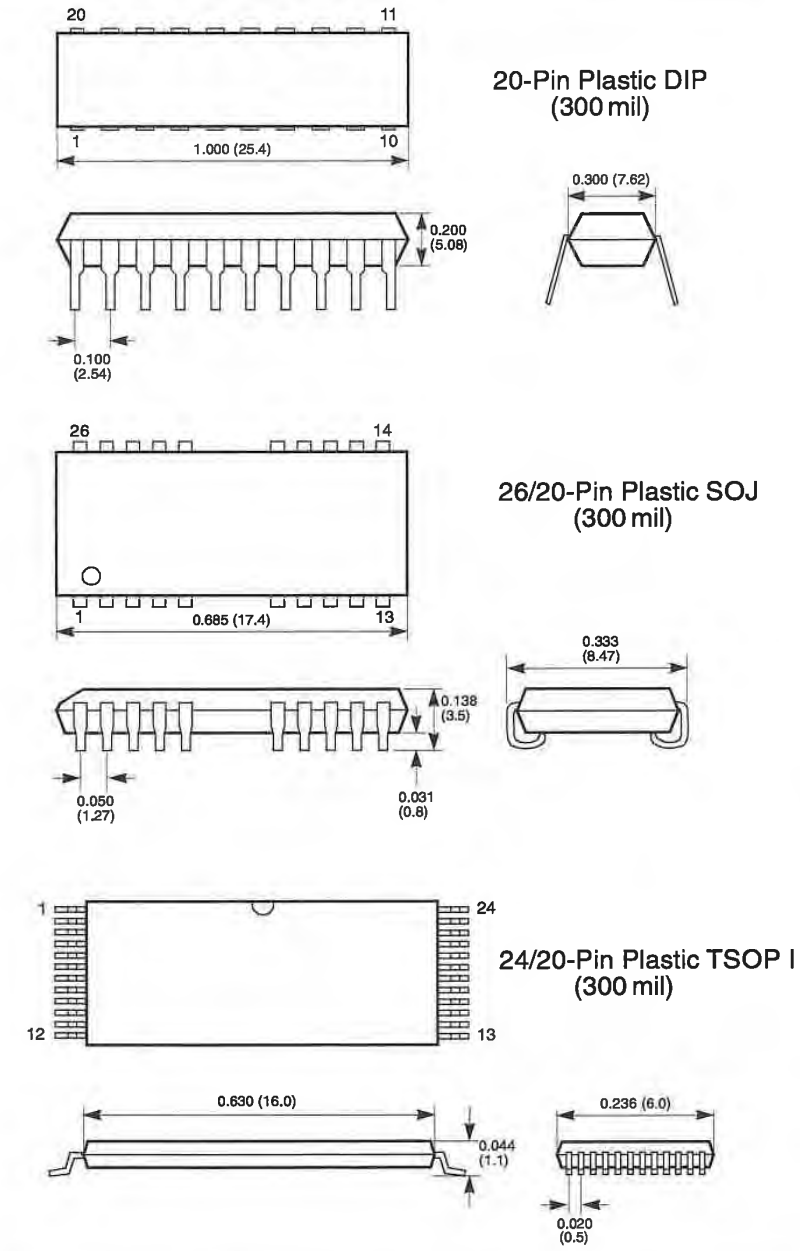
Table 5-1 lists the packaging options and their physical sizes for each of the mainstream DRAM organizations since the 256-Kbit generation. Many of the DRAMs that are available in TSOP packages come in normal and reversed configurations. The reverse pinouts are the mirror image of the normal pinouts along the long axis of the package. These pinouts are offered to facilitate double-sided SIMMs and circuit boards. Reverse-pinout DRAMs on the back side of the board connect to DRAMs on the top surface with a minimum of signal routing.

To date, all DRAM packages have pins only on their perimeters. The area of these packages is proportional to the number of pins and the pin spacing. Memory vendors have kept package sizes reasonable by relying on ever finer pin pitches. Keeping package sizes reasonable has become especially critical as pin counts have climbed from 16 pins for a 4K×1, to 24 for a 4M×4, to 40

1. *NEC Memory Products Data Book*, Volume 1, 1993, pp. 16–48.

© 1996 Steven Przybylski

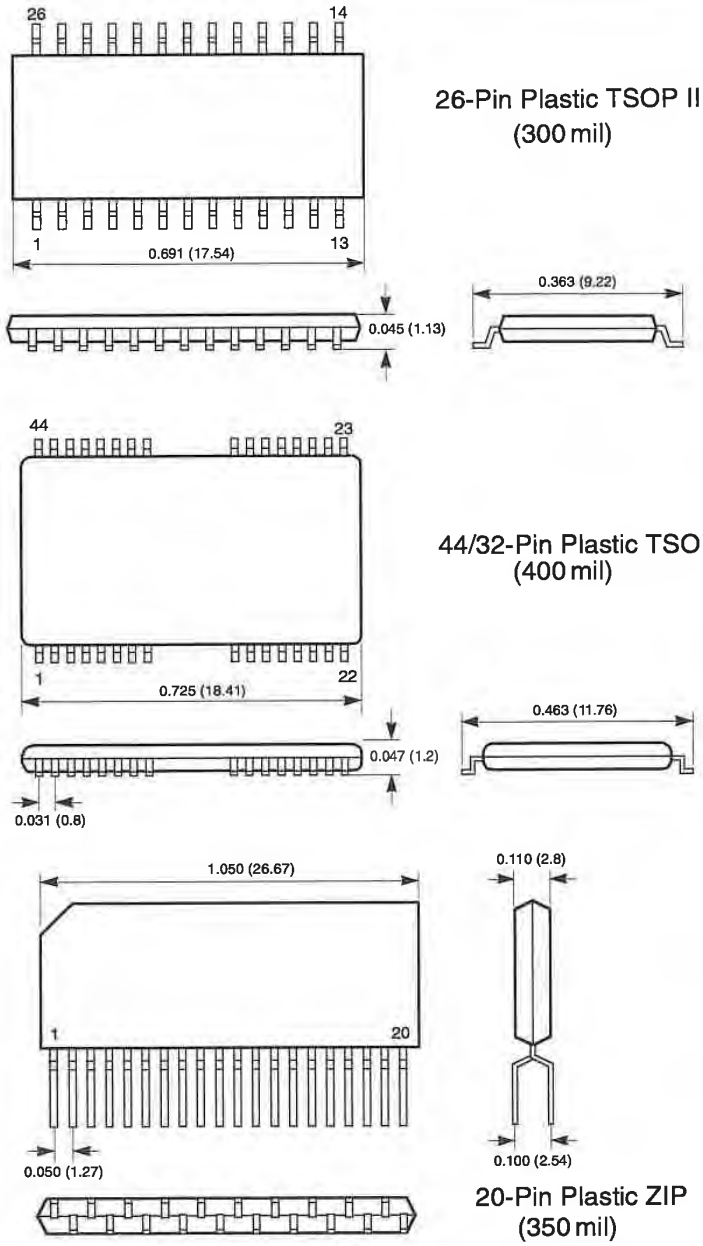
Samsung Ex. 1097, p. 131
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454



Dimensions are in inches (millimeters)

Figure 5-1 DRAMs come in a variety of packages. (*Continues next page*).

Samsung Ex. 1097, p. 132
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454



Dimensions in inches (millimeters)

Figure 5-1 (*Continues next page*).

© 1996 Steven Przybylski

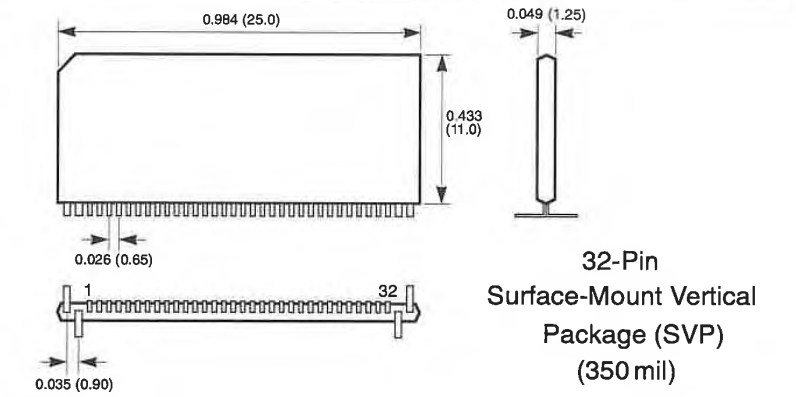


Figure 5-1 (*Continued*).

for the wide ×16 parts. Vertically mounted DRAMs greatly reduce the DRAM's footprint, but often the physical constraints of the system do not allow use of these packages. For example, ZIPs are not usable on SIMMs because the SIMMs themselves must be packed closely together.

Table 5-1 really tells only part of the whole story. Many organizations, especially in the 16- and 64-Mbit generations, are available in more than one package width for each type. This allows manufacturers to move their designs from larger 400- or 500-mil-wide packages to narrower 300- or 400-mil-wide packages as organizations migrate to subsequent technologies with smaller geometries.

Package properties assume greater importance

The electrical properties of DRAM packages have been of little concern until recently. However, DRAM widths and operating frequencies are growing rapidly. As they grow, so does the importance of lead inductance and its contribution to both ground bounce and the degradation of high-speed signals. Fine-pitch packages can have greater inductance and mutual inductance than their bigger cousins. The main advantage of fine-pitch packages is that the smallest possible package can be used. This increases packing density and decreases the propagation delay on the printed circuit board. Fine-pitch packages have naturally shorter lead lengths and consequently lower capacitance. However, package capacitance is much less important than package inductance because, in most cases, the package capacitance is dwarfed by the input capacitance of the IC itself.

Samsung Ex. 1097, p. 134
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

Density	Organization	Pins	DIP	SOJ	TSOP I	TSOP II	ZIP
		Pin Pitch	100 mils (2.54 mm)	50 mils (1.27 mm)	20 mils (0.51 mm)	50, 31, or 26 mils	50 mils (1.27 mm)
256 Kbit	256K×1	16	0.33 × 0.87 **0.29 sq. in.**	—	—	—	0.13 × 0.80 **0.10 sq. in.**
	64K×4	18	0.33 × 0.87 **0.29 sq. in**	—	—	—	20 pins 0.13 × 1.00 **0.13 sq. in.**
1 Mbit	1M×1	18	0.33 × 0.87 **0.29 sq. in**	20 pins 0.34 × 0.68 **0.23 sq. in.**	20 pins 0.24 × 0.63 **0.15 sq. in.**	20 pins 0.37 × 0.68 **0.25 sq. in**	20 pins 0.13 × 1.00 **0.13 sq. in.**
	256K×4	20	0.33 × 0.97 **0.32 sq. in.**	0.34 × 0.68 **0.23 sq. in.**	0.24 × 0.63 **0.15 sq. in.**	0.37 × 0.68 **0.25 sq. in.**	0.13 × 1.00 **0.13 sq. in.**
4 Mbit	4M×1	20	0.33 × 0.97 **0.32 sq. in.**	0.34 × 0.68 **0.23 sq. in.**	0.24 × 0.63 **0.15 sq. in.**	0.37 × 0.68 **0.25 sq. in.**	0.13 × 1.00 **0.13 sq. in.**
	1M×4	20	0.33 × 0.97 **0.32 sq. in.**	0.34 × 0.68 **0.23 sq. in.**	0.24 × 0.63 **0.15 sq. in.**	0.37 × 0.68 **0.25 sq. in.**	0.13 × 1.00 **0.13 sq. in.**
	512K×8	28	—	0.45 × 0.73 **0.33 sq. in.**	—	0.47 × 0.73 **0.34 sq. in.**	0.13 × 1.4 **0.18 sq. in.**
	256K×16	40	—	0.45 × 1.03 **0.46 sq. in.**	—	0.47 × 0.73 **0.34 sq. in.**	0.13 × 2.0 **0.26 sq. in.**
16 Mbit	16M×1	24	—	0.45 × 0.675 **0.30 sq. in.**	—	0.47 × 0.675 **0.30 sq. in.**	—
	4M×4	24	—	0.45 × 0.675 **0.30 sq. in.**	—	0.47 × 0.675 **0.30 sq. in.**	—
	2M×8	28	—	0.45 × 0.73 **0.33 sq. in.**	—	0.47 × 0.73 **0.34 sq. in.**	—
	1M×16	42	—	0.45 × 1.08 **0.49 sq. in.**	—	0.47 × 0.825 **0.39 sq. in.**	—
	512K×32	70	—	0.45 × 1.78 **0.80 sq. in.**	—	0.47 × 1.015 **0.48 sq. in.**	—
64 Mbit	16M×4	32	—	0.45 × 0.825 **0.37 sq. in**		0.47 × 0.84 **0.39 sq. in.**	—
	8M×8	32	—	0.45 × 0.825 **0.37 sq. in**		0.47 × 0.84 **0.39 sq. in.**	—
	2M×16	50	—	0.45 × 1.275 **0.57 sq. in**		0.47 × 0.81 **0.38 sq. in.**	—

Table 5-1 DRAMs are available in a variety of package types with differing areas. Dimensions are in inches.

JEDEC standardizes pinouts

Package dimensions and pinouts—the assignment of signals to pins—are standardized by the Joint Electron Device Engineering Council (JEDEC), one of the standardization bodies run by the Electronics Industries Association (EIA). The JEDEC JC-43 committee and its subcommittees are concerned with all types of semiconductor memories. Members of the committees include

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 135
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

representatives of the major semiconductor vendors and some systems companies. For the most part, these representatives see their charter as being to ensure the physical and, to a lesser extent, electrical compatibility of essentially similar parts across vendors. A company that has developed a new device, package, or module might take the design to JEDEC in an attempt to have it adopted as a standard. Thereafter, any other company wanting to produce a similar part would use the same pinout. This allows customers to use parts manufactured by any vendor, provided the parts are functionally and temporally compatible. The task of ensuring functional and temporal compatibility has almost always fallen on the customer. Furthermore, a configuration adopted by JEDEC is not necessarily free of intellectual property restrictions. JEDEC has been known to reject proposals that were excessively protected by patents. For the most part, though, the body notes but does not act on patent issues related to proposed and adopted standards.[2]

Recent exceptions to this model of the subcommittee's role include **Video RAMs (VRAMs)**, **Synchronous DRAMs (SDRAMs),** and **Synchronous Graphics RAMs (SGRAMs)**. For these devices, JEDEC has undertaken to specify the functionality of the devices as well as the pinouts. This is seen as necessary because as memory devices get more complicated, the likelihood of companies designing incompatible parts with the same pinouts increases. Such a situation would defeat the purpose of having a standardizing body.

Bare die a new "packaging option"

Multichip modules (MCMs) and multichip packages finally began coming of age in 1995. This largely resulted from a great increase in the number of semiconductor companies offering known good bare die. The reliability of these tested and burned-in die is high enough that multichip modules and packages have a high probability of being good. Base DRAM die from several vendors are currently in use in modules for notebook PCs and other high density applications.

5.2.2 The Electrical Interface

The DRAM's electrical interface defines the voltage levels that represent a logic 0 and logic 1, coming and going. Also defined are the maximum current and the average power used by the device during various types of operations.

2. See the list of relevant patents in Appendix D.

Samsung Ex. 1097, p. 136
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

TTL

In the late 1960s, the industry settled on **TTL (Transistor-Transistor Logic)** as the standard voltage levels. Designed as part of the bipolar logic family of the same name, TTL specifies a 5 V supply, a logic 0 of 0.8 V or below, and a logic 1 of 2.0 V or above. To maintain 400-mV noise margins, driving devices must drive at least 2.4 V for a logic 1, and at most 0.4 V for a logic 0. When driving other TTL loads, bipolar TTL devices drive a logic 1 at about 3.5 V and a logic 0 at 0.2 V.

These voltages were developed for bipolar circuits that were used in the first several generations of Small-Scale Integration (SSI) logic parts (7400 series). Consequently, they are not the best choice for NMOS or CMOS devices. For example, CMOS devices usually drive rail-to-rail. Under no-load conditions, a logic 0 will be 0 V and a logic 1 will be V_{cc} (5 V).[3] With typical loads, many CMOS devices are specified to output at least 3.5 V for a logic 1. On the input side, CMOS circuit designers must ensure that under all process and temperature conditions, the input trip point is at around 1.4 V instead of a more natural 2.5 V.

Low-Voltage TTL

Over the past few years, 3.3 V has rapidly gained acceptance as the supply voltage for portable and power-conscious systems. This change brought with it a shift to **Low-Voltage TTL (LVTTL)** signaling levels. The only difference between TTL and LVTTL is the latter's substantially reduced maximum input voltage. Most CMOS circuits are specified with a maximum logic 1 input voltage somewhere between V_{cc}+0.3 V and V_{cc}+1.0 V. The potential incompatibility between TTL and LVTTL occurs when a 5 V TTL CMOS part drives a full-rail logic 1, exceeding the maximum allowed signal input voltage of a 3.3 V LVTTL part.

Large voltage swings between the logic 0 and logic 1 levels make high-speed operation over unterminated traces problematic for both TTL and LVTTL. As either the slew rate or the length of the trace increases, so does the severity of ringing caused by signal reflections. Excessive ringing, shown in Figure 5-2, causes three problems:

- Ringing delays the stability and validity of the voltage signal at the receiver's input.

3. Historically, V_{cc} was used to denote the TTL supply voltage, and V_{dd} denoted the supply voltage for MOS devices. For many years, these two voltages have almost always been the same. Consequently, today V_{cc} and V_{dd} are used quite interchangeably. V_{cc} is preferred in this report because JEDEC uses it to denote the primary supply voltage on memory devices, regardless of the implementation technology.

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 137
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

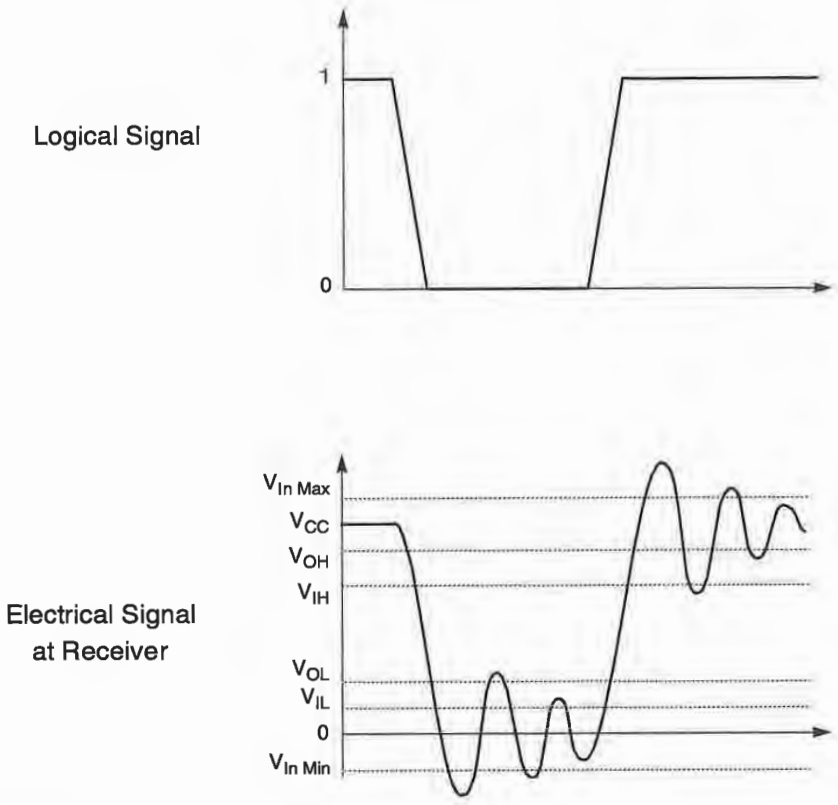


Figure 5-2 Excessive ringing is due to high slew rates and long, unterminated lines.

- Excessive undershoot can violate the input specifications of the receiving device. Typically, CMOS devices specify the lowest allowed input voltage at between -0.3 V and -1.0 V. Exceeding this level injects noise into the device's substrate and can damage the device if extreme.
- If the magnitude of the ringing exceeds the noise margin by a large enough value, very sensitive devices may sense multiple transitions in place of one. This can be especially deadly if occurring on a clock signal.

Another significant problem for TTL and LVTTL is the amount of current needed to drive large capacitive loads quickly. Large

Samsung Ex. 1097, p. 138
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

currents in the output drivers make for high power consumption and large amounts of **ground bounce** both on and off the DRAM IC. Since excessive ground bounce can make a device or system unreliable, the amount of ground bounce must be controlled. This necessity imposes practical limitations on the number and speed of the output drivers. Finally, large signal swings radiate substantial amounts of power. The techniques that can reduce the amount of electromagnetic interference (EMI) generated by a system include:

- Proper shielding in the system's case
- Reduction of trace lengths
- Contiguous ground planes
- Reduction of the swing of high-speed signals

For many years, designers of high-speed systems have avoided these problems with TTL by using ECL with terminated transmission lines. The advantages of ECL include low signal swings for fast transitions and reliable signal quality due to limited and manageable signal reflections.[4] From the perspective of a CMOS system designer, the primary disadvantages of ECL are the unusual power supply levels and the high power dissipation of an output driver.

Recently, however, two proposals targeted at CMOS systems promise ECL's high speed and good signal integrity without its drawbacks. Like ECL, both GTL and CTT are meant to work in terminated-transmission-line environments.

New low-swing electrical interfaces for CMOS systems

Like many of the new signaling schemes, **Gunning Transceiver Logic (GTL)** relies on an externally generated reference voltage, around which the signal swing is centered. In the case of GTL, the nominal logic levels are 400 mV above and below the 0.8-V reference voltage. In addition, all GTL signals are parallel-terminated to a 1.2-V termination supply through a terminating resistor that matches the trace's characteristic impedance. As a result, the driver only has to pull down to a logic 0; the line naturally restores itself to a logic 1 when the driver turns off (see Figure 5-3). Correspondingly, the driver does not dissipate any power when driving a logic 1, but dissipates around 10 mW when driving a logic 0 (about 20 mA into a 50 Ω line). In wider DRAMs, the **Simultaneous Switching**

4. A very pragmatic and approachable treatment of transmission lines and high-speed system design is found in H. Johnson and M. Graham, *High-Speed Digital Design: A Handbook of Black Magic*, Prentice Hall, 1993.

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 139
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

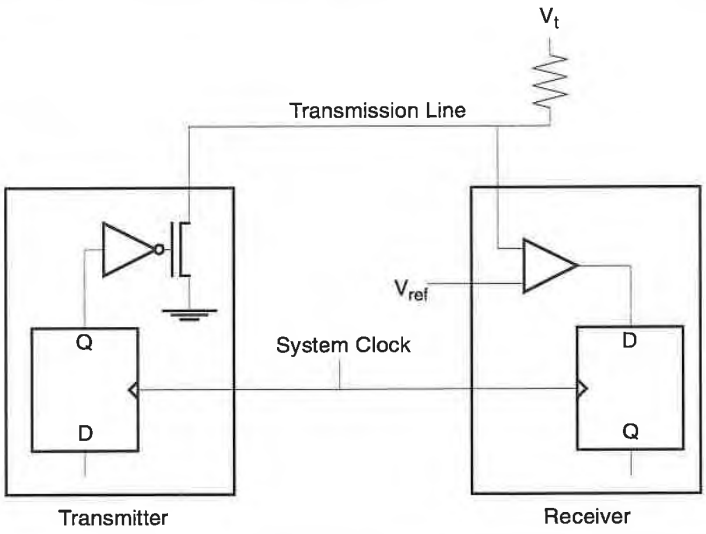


Figure 5-3 Gunning Transceiver Logic (GTL) drivers dissipate power only when driving a logic 0.

Noise (SSN) resulting from 16 or 18 drivers turning on at the same time can cause significant ground bounce on the DRAM unless an adequate number of supply pins are provided on the package.

One alternative to GTL is the **Center-Tap-Terminated (CTT)** interface. Like GTL, CTT relies on an externally generated reference voltage to set the center point of the signal swing. However, unlike GTL and other new electrical interfaces, CTT provides compatibility with existing rail-to-rail CMOS and LVTTL levels. When used with short, unterminated traces, CTT drivers provide regular rail-to-rail swings. However, when used with a longer, terminated transmission line, the driver automatically adjusts to a small swing. In this case, the output swings are 800 mV, and the worst case input levels are tightened to ±200 mV around the reference voltage. The driver monitors the difference between the desired state of the output signal and its actual state, and reduces its drive once a valid level is reached. Both GTL and CTT have been adopted by JEDEC as standards.

The small signal swings of GTL and CTT work well under ideal conditions and in the absence of noise and signal reflections. However, most memory systems are less than ideal. Imperfect

Samsung Ex. 1097, p. 140
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

layout of the memory array and the signal traces can lead to impedance mismatches and trace stubs that cause noticeable signal reflections. These reflections, which can result in significant degradation of signal quality, can be dealt with in one of two ways:

1. Careful design of the board and SIMMs can minimize or eliminate the mismatches and trace stubs that cause the reflections.

2. Use of electrical interfaces tolerant of poor trace characteristics can minimize the impact of the reflections.

New alternatives: HSTL and SSTL

Throughout 1993, GTL was widely hailed as the next standard electrical interface for all forms of chip-to-chip communication. However, it became increasingly clear that GTL lacks robustness in many circumstances. Although GTL will probably continue to be used in backplanes and some point-to-point communication applications, it is unlikely to gain acceptance as a DRAM interface. Instead, JEDEC's JC 16.3 committee has defined another low-swing interface called High Speed Transceiver Logic (HSTL).

HSTL is unique in that it defines four classes of output levels for logic 0 and logic 1. The classes differ according to the load being driven (see Figure 5-4):

Class I:
Push-pull output buffer driving an unterminated capacitive load or a symmetrically-parallel-terminated transmission line[5]

Class II:
Push-pull output buffer driving an externally source-terminated transmission line or a symmetrically-double-parallel-terminated transmission line

Class III:
Push-pull output buffer driving asymmetrically-parallel-terminated transmission line

Class IV:
Push-pull output buffer driving asymmetrically double-parallel-terminated transmission lines

5. Symmetrical parallel termination is termination to Vddq/2; asymmetrical parallel termination is termination to Vddq.

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 141
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

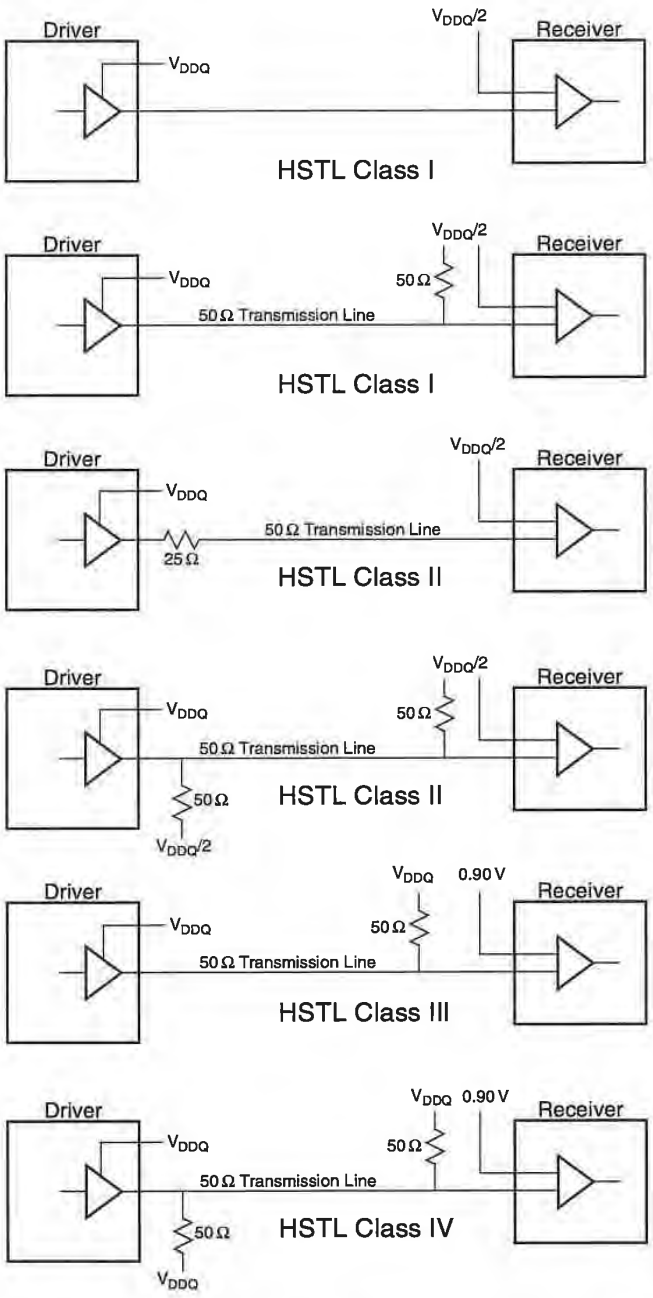


Figure 5-4 The new HSTL JEDEC standard defines 6 electrical environments, each of which is driven by an HSTL driver of one of four classes.

Samsung Ex. 1097, p. 142
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

HSTL is a low-swing interface in an era of higher supply voltages. Therefore, all the levels are defined in terms of VddQ, which is nominally 1.5 V, independent of the IC's supply voltages. Providing for unterminated and series-terminated environments as well as parallel termination facilitates all common interconnection schemes. The asymmetrically terminated schemes of classes III and IV are useful for tri-stated bidirectional buses because with asymmetrical termination, an undriven signal does not float to $V_{DDQ}/2$—the trip point of the receivers in classes I and II, at which they typically burn uncharacteristic amounts of power or even oscillate.

The standard is designed so that any HSTL driver can successfully drive any HSTL receiver. This is accomplished through several techniques. First, output buffers and receivers both input the threshold reference voltage Vref, which is nominally at VddQ/2. Second, although the voltage levels for 0 and 1 are the same except for the asymmetrically terminated cases, the drive currents at which these levels are attained at DC vary by class.

Another variant is **Stub Series (Terminated) Transceiver Logic (SSTL)**. SSTL is very similar to CTT: both have a 3.3 V_{DDQ}, a 1.5 V termination voltage, and a minimum input AC swing of 0.8 V. Like HSTL, SSTL defines two classes of electrical environments. The first, Class I, assumes a single 50-Ω terminating resister at the receiver, much as in HSTL Class I. In order to generate the required signal swing at the receiver, the driver must sink or source 8 mA when driving a logic 0 or logic 1, respectively. SSTL Class II assumes a doubly terminated environment, as in HSTL CLass II. In this case, the output driver must sink or source 16 mA of current. So, unlike most other voltage-mode electrical interfaces, the output driver requirements are specified in terms of current instead of voltage.

However, the important part of SSTL is not actually part of the electrical interface at all, but rather part of the physical environment in which SSTL is expected to be used. Specifically, small series termination resistors are to be placed between stubs on the SIMMs or DIMMs and the main bidirectional transmission line, as in Figure 5-5. The series stub resistors attenuate the impact of reflections coming back down the stubs on the main bus transmission line. This greatly improves signal integrity throughout the shared bus, thereby both helps the effectiveness of the reduced noise margin and reducing the time

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 143
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

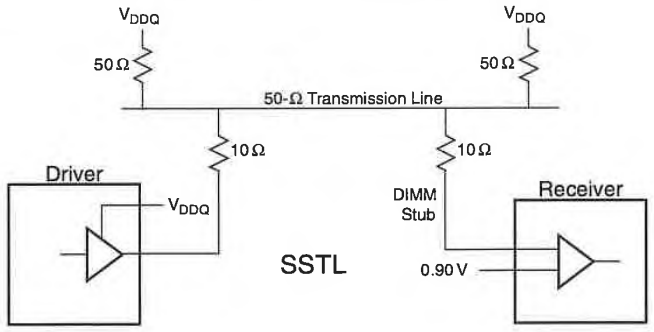


Figure 5-5 Stub series termination attenuates the effects of long stubs and greatly improves the signal integrity on the main transmission line.

needed for the bus to settle to a valid level. SDRAMs with SSTL electrical interfaces should be available by the end of 1996.

Unfortunately, as of 4Q1995, SSTL is still an untested proposed standard. No substantial systems have been built to its specifications, and design experience is equally limited. Until a greater body of experience is developed by the DRAM vendor and consumer communities, SSTL will remain as unproven as GTL was in 1993.

Low-Voltage Differential Signals (LVDS)

Another new signaling interface, **Low-Voltage Differential Signals (LVDS)**, is under development as a potential standard by the IEEE Computer Society P1596.3 Working Group. LVDS will offer an alternative to the ECL levels selected for the Scalable Coherent Interconnect (SCI). The main idea behind the interface is that noise in either the signal conductor or the reference voltage—typically the ground and/or the supply voltage—inhibits reliable use of small signal swings. The use of a differential scheme essentially eliminates noise contribution from the supplies. Using terminated transmission lines facilitates high-speed slew-rate signals over long distances. The resulting interface is straightforward: 250-mV differential swings, centered nominally at 1.2 V above ground. To account for ground shift and noise, receivers should accept input signals centered anywhere between 0 and 2 V. The main disadvantage is that the numbers of both signal wires and output drivers are doubled. Since LVDS is a new standard, it has yet to be used in any systems or components outside the SCI community.

Samsung Ex. 1097, p. 144
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

Rambus Signaling Logic (RSL)

Of course, a host of proprietary solutions are possible. A noteworthy example is the Rambus electrical interface, called Rambus Signaling Logic (RSL) (see Section 8.2). By physically constraining the length of the bus and controlling the electrical properties of the terminated transmission lines, it is possible to construct a bidirectional, high-speed, single-ended bus.

Rambus uses 1 V swings[6] centered around an externally generated reference voltage of approximately 2.2 V. The bidirectionality of the bus is restricted. Data can flow from a bus master to one of the slaves or from a slave to a bus master, but data can flow from one slave to another only under very constrained circumstances. All bus masters must be located at one end of the bus. The bus is terminated only at the non-master end, so half-swing signals originating with slaves reflect off the unterminated end to attain full amplitude.

In general, when terminated transmission lines are used, multiple-source/multiple-destination buses must be terminated at both ends. The restriction of the information flow on the Rambus Channel allows it to be singly terminated. An important side effect is that the current needed to drive a signal is halved in comparison to the more general double-terminated case.

Advantages of a current-mode bus

The Rambus Channel is a current-mode bus. This eliminates the clipping that causes overlapping signals to interfere in voltage-mode buses. Eliminating the interference is especially important given the single-ended termination scheme, since driven signals take two bus propagation times to dissipate off the bus. If Rambus used a voltage-mode electrical interface, the dead time between different devices driving the bus would have to be doubled.

Unlike voltage-mode drivers, current-mode drivers present a high impedance to the bus, regardless of whether they are driving or not. This high impedance means that the drivers do not cause secondary reflections in signals originated by other drivers. Also, current-mode drivers can be built using smaller transistors with smaller parasitic capacitances: when combined with a low capacitance package, these drivers alter the characteristic impedance of the transmission lines less than voltage-mode drivers do.

The Rambus Channel is terminated to 300 mV above the reference voltage: devices sink current when pulling the bus low

6. Early descriptions of RSL specified a 600 mV swing. This was changed in mid 1994.

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 145
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

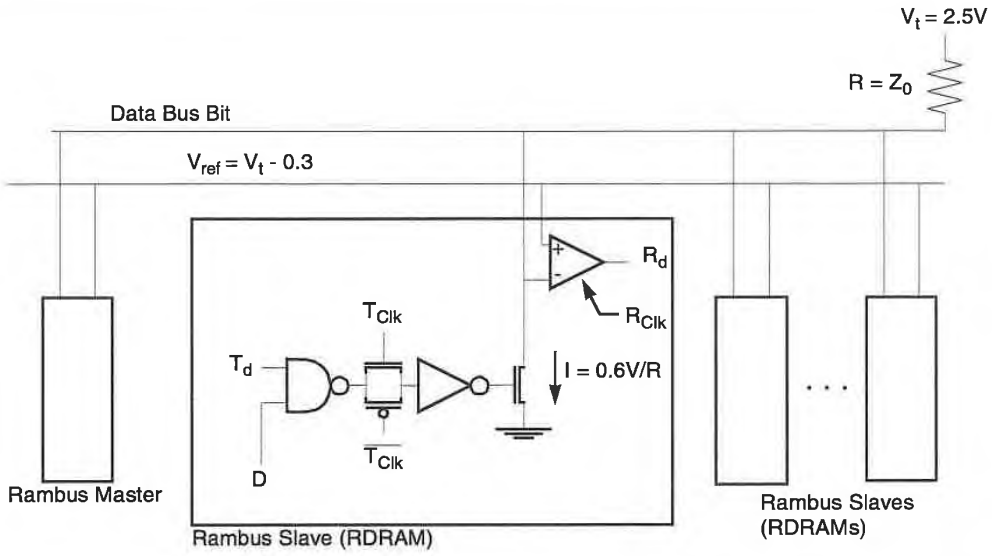


Figure 5-6 The Rambus electrical interface defines a low-signal-swing, current-mode bus.

(logic 1) and release the bus when driving a logic 0 or when not driving the bus (see Figure 5-6). To allow for variation in the characteristic impedance of the transmission lines, the current that a slave must sink is programmable between 10 mA and 35 mA. Although 35 mA multiplied by 10 output drivers is a substantial amount of current, the Rambus DRAM package provides enough ground pins to keep the on-chip ground bounce tolerable. As with other schemes that have externally supplied reference voltages, the input receivers are relatively tolerant of ground bounce because they are essentially differential amplifiers with high common-mode rejection ratios.

Necessary but not sufficient for high-speed operation

Table 5-2 summarizes the important signaling interfaces. Although the new low-signal-swing electrical interfaces facilitate high-speed operation, they address only part of the problem. In particular, large clock skews and on- and off-chip signal propagation delays put short cycle times at odds with large and expandable memory systems. In other words, an effective electrical interface is a necessary, but not sufficient, condition for high-frequency systems.

Samsung Ex. 1097, p. 146
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

Electrical Interface	Supply Voltage	Single-Ended/ Differential	Minimum Output Voltage Swing	Typical Output Voltage Swing	Termination Voltage	External Reference Voltage
TTL (Bipolar)	5 V	Single-Ended	2 V	3.2 V	—	—
TTL (CMOS)	5 V	Single-Ended	2 V	5 V	—	—
LVTTL (CMOS)	3.3 V	Single-Ended	2 V	3.3 V	—	—
ECL (100K)	–4.5 V	Either	590 mV	600 mV	–1.2 V	—
CTT (Large Swing)	3.3 V	Single-Ended	2 V	3.3 V	—	—
CTT (Small Swing)	3.3 V	Single-Ended	800 mV	800 mV	Yes	
GTL	—	Single-Ended	800 mV	800 mV	1.2 V	0.8 V
TLVTTL	3.3 V	Single-Ended	400 m V	—	1.5 V	1.5 V
HSTL (Class I)	1.5 V	Single-Ended	800 mV	—	—/0.75 V	0.75 V
HSTL (Class II)	1.5 V	Single-Ended	800 mV	—	—/0.75 V	0.75 V
HSTL (Class III)	1.5 V	Single-Ended	700 mV	—	0.9 V	1.5 V
HSTL (Class IV)	1.5 V	Single-Ended	700 mV	—	0.9 V	1.5 V
SSTL	3.3 V	Single-Ended	800 mV	8 or 16 mA	1.5 V	1.5 V
LVDS	—	Differential	250 mV	250 mV	None	—
Rambus	—	Single-Ended	1 V	1 V	2.5 V	2.2 V

Table 5-2 The new signaling interfaces facilitate small signal swings and terminated transmission lines.

5.2.3 The Pin-Level Interface

The **pin-level interface** is the overall style of the logical interface; there are four types:

1. A multiplexed address bus with a separate data bus
2. A nonmultiplexed address bus and a separate data bus
3. A single, unified bus for addresses and data
4. A dual- or triple-ported device combining a fairly standard random-access, multiplexed address port with one or two specialized serial ports

Most currently available DRAMs use the first interface type. The second interface type, which includes a non-multiplexed address bus, is found in a few specialty high-speed DRAMs. The third type of pin interface—a unified address/data bus—is found on Rambus DRAMs and in the RamLink proposal. The fourth is found in video RAMs.

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 147
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

5.2.4 The Temporal Interface

The **temporal interface** primarily specifies whether the DRAM interface is synchronous or asynchronous. A synchronous interface has a fixed, free-running clock input; an asynchronous interface uses a combination of edge- and level-sensitive inputs. In addition, the temporal interface includes any special features that specifically address clocking concerns. For example, the Rambus Channel temporal interface specifies that data and the transmission clock always flow together down the bus, thus minimizing the clock skew between the sender and the receiver of the data.

The temporal interface also includes detailed timing information:

- Specific setup and hold times
- Access, cycle, precharge, and refresh times
- Output buffer delays
- Refresh timing

For synchronous parts, many of these times are expressed in cycles; for asynchronous parts, the times are expressed in nanoseconds and, in some cases, milliseconds.

DRAMs are generally offered in three or four speed grades. All timing characteristics are specified separately for each speed grade, and there is a great deal of consistency across vendors within each grade. Consequently, DRAMs are routinely referred to by their row access times (for example, a 70-ns DRAM). Many other important timing numbers can be inferred from the row access time without reference to a manufacturer's data sheet.

5.2.5 The Logical Interface

The **logical interface** is primarily concerned with the sequence of events rather than the details of the timing. It specifies how the bits are organized within the DRAM, what operations are defined, and how those operations are controlled.

For example, the logical interface for a fast-page-mode DRAM describes:

- The number of bits

Samsung Ex. 1097, p. 148
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

- The width of the bus
- The partitioning of the address into row and column addresses
- The sequence of control signals and data transfers necessary to complete each type of read or write operation
- The size, configuration, and functionality of the **row cache** built into the DRAM's sense amplifiers

A large number of logical interfaces exist among the standard asynchronous DRAMs. Not only are there **fast-page-mode**, **static-column-mode**, and **nibble-mode** parts for each DRAM size (see Section 3.3), but there are a variety of widths from ×1 through ×18. In addition, for a given size, width, and type of DRAM, further choices must be made regarding the number of refresh cycles and the relative sizes of the **row address** and **column address**.

Traditionally, row and column addresses have been the same width. However, with the ×16 4-Mbit parts and ×4, ×8, and ×16 16-Mbit parts, the system designer can choose versions that have unequal, or asymmetrical, row and column addresses. One wide part with symmetrical addresses functions very much like four narrower parts from the previous generation. For this reason, wide parts are easily combined with smaller, narrower parts to form ×9 memory arrays and SIMMs. However, parts with wider row addresses consume less power since fewer sense amplifiers are used in row accesses and refresh operations. These technical reasons for choosing one type or the other are often overshadowed by the market dominance of the parts with symmetrical addressing. Asymmetrical parts, with their lower sales volume, are significantly more expensive—especially at the 4-Mbit level.

Many interface combinations from a single design

The number of different combinations of logical interface, speed grade, supply voltage, and package type is staggering. DRAM vendors manage this complexity by using a single IC design to fill many roles. Designs are one-time programmable at the factory through the use of bonding options, metal mask options, and, to a limited extent, laser fuses. For example, a 16-Mbit design might be programmable to any combination of the following features:

- Fast-page mode, static-column mode, or nibble mode[7]

7. Extended-Data-Out (EDO) and Burst EDO DRAMs can be built on the same die as fast-page mode and static-column mode.

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 149
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

- ×1, ×4, ×8, or ×16 wide
- Symmetrical or asymmetrical addresses
- Power supply of 5 V or 3.3 V

Each combination is then offered in three or four speed grades and in two to four packages. As many as 400 different products can be derived from a single IC design. Additional part numbers are generated when the DRAMs are tested to special electrical, timing, environmental, or reliability requirements.

5.3 Generic Narrow DRAMs

Each of the new DRAM alternatives involves a change in one or more aspects of the overall DRAM interface. To clarify these changes and their differences, this section compares them to a generic 1M×4, 5 V fast-page-mode DRAM. This part has the following characteristics:

Physical Interface:
A common package for a 1M×4 is a 26-pin TSOP measuring 0.7 inch (17.5 mm) by 0.36 inch (9.2 mm) with 20 actual pins on a 50-mil (1.27 mm) pitch.

Electrical Interface:
The device drives 5 V CMOS rail-to-rail levels.

Pin-Level Interface:
The standard DRAM interface includes a multiplexed address bus, a bidirectional data bus, and four control signals: $\overline{\text{RAS}}$, $\overline{\text{CAS}}$, $\overline{\text{WE}}$, and $\overline{\text{OE}}$. The address bus is 10 bits wide, and the data bus is 4 bits wide. This package provides for a single power pin and a single ground pin.

Temporal Interface:
The interface to standard DRAMs is fundamentally asynchronous. Though $\overline{\text{RAS}}$ and $\overline{\text{CAS}}$ act as clocks for the purposes of latching addresses and, in some cases, data, they do not display the fixed and regular behavior of clocks.

The reference speed is a row access time of 60 ns (t_{RAC}). The corresponding column access time is 15 ns (t_{CAC}), while the column address access time is 30 ns (t_{AA}). The column cycle time is 40 ns (t_{PC}); the row cycle time

Samsung Ex. 1097, p. 150
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

is 110 ns (t_{RC}). These times are typical of -6 fast-page-mode devices.

Logical Interface:
Section 3.3 details the logical interface of page-mode and fast-page-mode DRAMs. For this generic 1M×4 DRAM, the 20-bit address is evenly split between the row address and the column address. The row cache is therefore one 4-Kbit block, organized as 1K addressable chunks of four bits each. The *status quo* DRAM requires 1,024 refresh cycles every 16 ms.

5.4 Conventional Approaches

Two alternatives to the generic narrow DRAM are conventional in that they change the existing interface as little as possible. Making only limited changes to the logical and temporal interfaces keeps the cost premiums for the improved characteristics to a minimum. System vendors easily accept these two alternatives because of the limited redesign needed to incorporate them into systems, the apparently low risk involved, and the prospect of good pricing throughout a system product's lifetime.

The first of the two conventional alternatives consists of DRAMs that are ×8 or wider (see Section 6.2). These parts have increased bandwidth per bit and decreased DRAM depth for reduced memory system granularity.

Extended Data-Out (EDO) DRAMs (see Section 6.3) are the second alternative. Essentially fast-page-mode DRAMs, they incorporate a small change to the output buffer control. This modification allows greater overlap of successive column accesses within a burst. The end result is a smaller difference between the DRAM's specified minimum column access cycle time and the time actually attainable in a system. EDO DRAMs are called **hyper-page-mode DRAMs** by a few vendors.

5.5 Evolutionary Approaches

The four evolutionary alternatives involve more aggressive changes to the temporal and logical interfaces than the conventional alternatives do. All four evolutionary DRAMs

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 151
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

retain the multiplexed address bus that is the heart of the existing interface.

Burst EDO (BEDO) DRAMs (see Section 7.2) represent the final evolutionary step of the standard asynchronous fast-page-mode interface. As the JEDEC name (Pipelined Nibble Mode DRAMs) implies, Burst EDO devices are like nibble-mode DRAMs in that the DRAMs contain a 2-bit counter to construct a burst of four chunks of data from a single column address. Burst EDO devices are pipelined to facilitate decreasing the column cycle time below the column access time.

The **Enhanced DRAM (EDRAM)** interface is conventional in many respects and evolutionary in others (see Section 7.3). The increased performance of EDRAMs results from two sets of changes. First, the speed of the DRAM core is significantly increased. This increase is reflected in striking improvements in the speed of read and write operations to both the DRAM array and the row cache. In addition, EDRAMs include small but significant changes to the logical interface; these changes greatly improve the handling of data writes.

Synchronous DRAMs (SDRAMs) replace the asynchronous RAS/CAS temporal interface with a synchronous interface (see Section 7.4). In the process, $\overline{\text{RAS}}$ and $\overline{\text{CAS}}$ lose their clock characteristics and become sampled control signals. Some SDRAMs have two internal banks. If the reference stream strictly alternates between the two internal banks, 100% utilization of the synchronous data I/O bus can be attained with moderate-sized block-read transfers—even in the face of continuous row-cache misses.

Essentially all the performance improvement of SDRAMs derives from three changes:

- The new temporal interface
- The logical interface, which facilitates pipelining references to independent banks within a single DRAM
- A future change in the electrical interface, which will facilitate operation at frequencies above 100 MHz

Cached DRAMs (CDRAMs) add a small SRAM onto the same die as the DRAM (see Section 7.5). This SRAM has a dedicated non-multiplexed address bus and separate control signals. The memory controller can use the SRAM as a cache for the DRAM data. The ability to overlap SRAM and DRAM operations allows

Samsung Ex. 1097, p. 152
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

for high internal and external bandwidths and good cache-hit rates. CDRAMs therefore offer significant changes in both the temporal and logical interfaces.

5.6 Revolutionary Approaches

For two decades, the multiplexed address bus has been the crux of the DRAM interface. Many designers now feel that the old standard causes more problems than it solves. Rethinking the entire DRAM interface opens up new opportunities to meet the long term needs of DRAM users and suppliers. The two revolutionary new DRAM interface proposals discard the separate data and multiplexed address buses and replace them with a single byte-wide bus that carries addresses, data, and control.

Rambus Inc. developed the first revolutionary alternative and has licensed it to a variety of DRAM and ASIC vendors. Up to 32 **Rambus DRAMs (RDRAMs)** are vertically mounted on a Rambus Channel: a 9-bit-wide, 250-MHz, synchronous, bidirectional, low-swing current-mode bus. Internally, RDRAMs have two banks; each bank has its own row cache (see Section 8.2).

The Rambus approach involves significant changes to all aspects of the interface. The Rambus Channel specification fully defines new physical, electrical, pin-level, temporal, and logical interfaces. Most significantly, Rambus provides an impressive 500 Mbytes/s peak bandwidth, regardless of whether the channel is populated with one RDRAM or thirty-two. As 16-Mbit and then 64-Mbit DRAMs become common, Rambus's single-DRAM memory system granularity will be more valuable. The first 16- and 18-Mbit RDRAMs are in production; additional 8- and 16-Mbit RDRAMs will become available in 1996.

RamLink and SyncLink

The second revolutionary alternative, **RamLink**, is a recently approved standard developed by the IEEE Computer Society P1596.4 Working Group (see Section 8.3). RamLink is the name given to the protocol developed; RingLink is the name for the physical layer adopted with the protocol. Like the Rambus approach, this alternative replaces the separate address, data, and control buses of the generic DRAM with a single byte-wide, high-speed unified bus that carries all three types of information.

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 153
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

DRAM Type	Alternative	Interface				
		Physical	Electrical	Pin-Level	Temporal	Logical
Reference	Generic					
Conventional	Wide Fast-Page-Mode					
	Extended-Data-Out					
Evolutionary	Burst EDO					
	EDRAM					
	SDRAM					
	CDRAM					
Revolutionary	Rambus					
	RamLink					
	SyncLink					

Table 5-3 Each new DRAM alternative changes different portions of the overall interface. Darker shading indicates more significant changes.

The biggest technical difference between RamLink/RingLink and Rambus is that Rambus facilitates high-speed operation by physically and electrically constraining a traditional bidirectional bus. RingLink, on the other hand, uses a ring topology with point-to-point, low-voltage-swing differential connections.

The advantage of RingLink's approach is that the DRAMs can be physically separated. However, each additional node in the ringlet adds latency to all transactions on that ring. Notably, RamLink's logical interface provides general facilities for both I/O and memory devices, making the technology potentially useful as a general system interconnect.

In 1995, several DRAM vendors took notice of RamLink and decided that with a different physical layer it could become a serious competitor to Rambus. Thus SyncLink is a derivative of the RamLink protocols on a dual-bus physical layer with a total peak bandwidth of 1 Gbyte/s.

SyncLink is in a very preliminary stage of development. Though RamLink is ready to become a formal standard,[8] as of 4Q1995, SyncLink is still in its formative stages. Though eight semiconductor and system companies have banded together to form an industry consortium to support and foster the refinement of SyncLink, SyncLink products are unlikely to exist until the 64-Mbit generation.

8. See Section 8.3.4

Samsung Ex. 1097, p. 154
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

How evolutionary and revolutionary are these designs?

Table 5-3 lists all the conventional, evolutionary, and revolutionary options and shows the degree to which each alternative involves changes to the generic DRAM's physical, electrical, pin-level, temporal, and logical interfaces. The more radical the approach, the greater the number and the severity of the changes to the overall interface. Chapters 6 through 8 explore these changes in greater detail. Chapters 9 and 10 consider whether the changes return commensurate improvements in performance or functionality.

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 155
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

Samsung Ex. 1097, p. 156
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

6 Conventional Alternatives

The conventional alternatives to the generic narrow fast-page-mode DRAMs minimize changes to the interface: primarily, the conventional DRAMs provide improved bandwidth or reduced granularity while taking advantage of the existing designs and cost structure of generic DRAMs. Two DRAM architectures classified as conventional are discussed in this report: wide DRAMs and **Extended Data-Out (EDO) DRAMs**. In conventional wide DRAMs, the logical interface and package are widened from ×4 to ×16 or wider. In Extended Data-Out (EDO) DRAMs, the logical interface is modified slightly to facilitate more efficient burst transfers.

6.1 New In The Second Edition

The second edition of this report includes fill frequencies for all wide and EDO DRAMs, including the new ×32 EDO devices. In addition, more complete and up-to-date design information for EDO devices is provided.

The second edition of this report drops Fast DRAMs as a conventional alternative. Fast generic DRAMs have always found their primary application in graphics systems. Instead, the Burst EDO evolutionary alternative is added to Chapter 7.

Samsung Ex. 1097, p. 157
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

6.2 Wide Conventional DRAMs

Today, wide-word DRAMs are defined informally as DRAMs that have eight or more bidirectional data pins. However, this definition is fluid. At the 256-Kbit level, ×4 DRAMs were considered wide; soon, only parts that are ×16 or more will be considered unusually wide. The unchanging element of this definition is that wide DRAMs are understood to be wider than the current generation's mainstream DRAMs.

Figure 6-1 shows the logical organization of a 1M×16 DRAM. This organization is very much akin to sixteen 1M×1 parts in parallel sharing address and control lines (see Figure 3-2 on page 44). The only real difference is that instead of having one $\overline{\text{CAS}}$ signal, the illustrated part has two. Having two lines facilitates reads and writes to both the upper and lower bytes. Access to the individual bytes is important for a variety of main memory and graphics applications that have to handle byte and subword writes efficiently. If only one $\overline{\text{CAS}}$ were present, a byte write would involve a 16-bit read and a merge of old data with the new in the memory controller, followed by a 16-bit write. In addition to the parts with two $\overline{\text{CAS}}$ pins, some versions have two write-enable pins and one $\overline{\text{CAS}}$ pin. These pins allow read accesses of one byte in parallel with write accesses to the other byte.

Some ×16 parts targeted toward video and graphics applications also offer per-bit writes. A write mask is latched with the falling edge of $\overline{\text{RAS}}$. For portable applications, wide 16-Mbit DRAMs with self-refresh capabilities are available. These DRAMs are able to enter a mode in which they autonomously and indefinitely refresh themselves without any outside intervention.[1] At the 16-Mbit level, designers can also select the number of row and column address bits. Many suppliers are offering three sets of ×16 parts with row address sizes of 10, 11, or 12 bits.

Large packages needed for wide DRAMs

In addition to changes to the logical interface, there are noteworthy changes to the physical and temporal interfaces. Physically, wide devices have fewer address pins and significantly more data pins. Each quadrupling in the number of data pins nominally yields a reduction by one in the number of address pins. The ×1 parts have the fewest pins, and the total number

1. Self-refresh DRAMs are different from Pseudo-Static RAMs (PSRAMs) in that PSRAMs do not have a multiplexed address bus. Despite its DRAM internals, the PSRAM's interface more closely resembles an SRAM's interface than a DRAM's.

Samsung Ex. 1097, p. 158
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

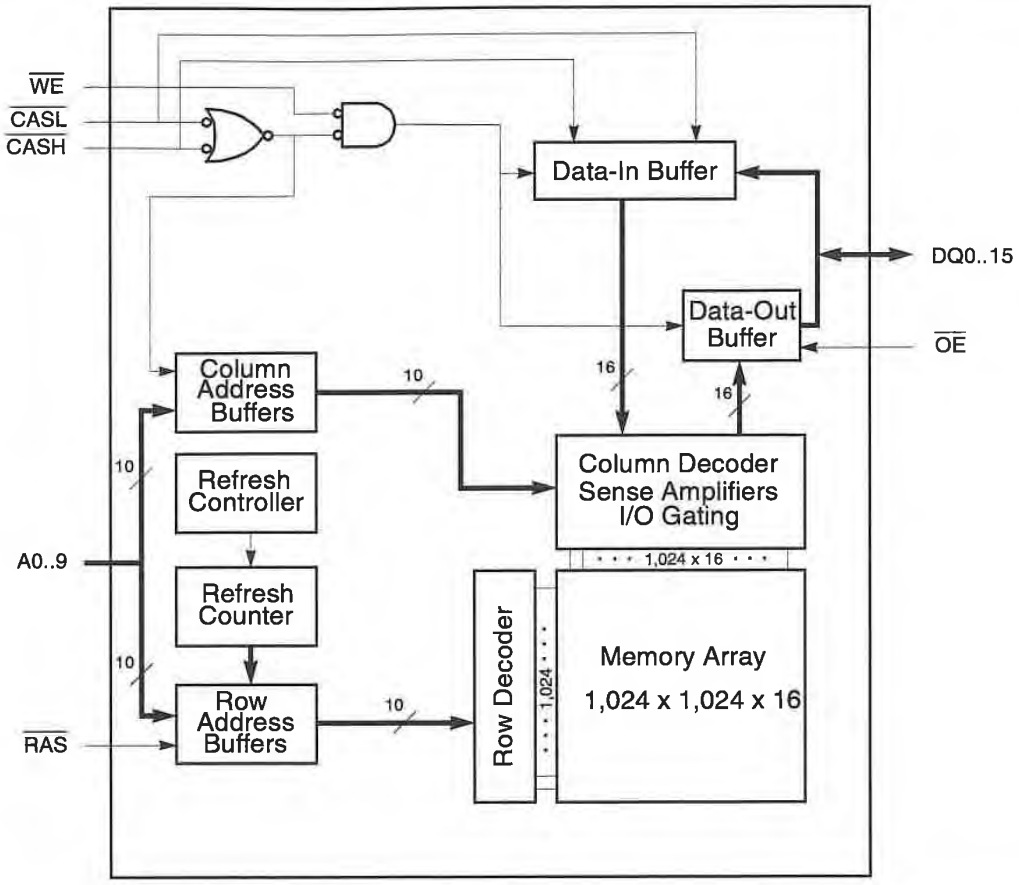


Figure 6-1 Logically, one ×16 DRAM is very much like sixteen ×1 parts integrated in parallel.

grows rapidly for devices wider than ×8. Table 5-1, on page 135, shows the growth in the number of pins associated with increased width. Basically, ×16 devices, which are housed in packages of 40–50 pins, have one power or ground pin for every four data pins (see Figure 6-2). The extra supply pins keep ground bounce and on-chip supply noise to a tolerable level, even when all the outputs are switching in the same direction simultaneously. One supply for every four outputs is consistent with common practice in the microprocessor and ASIC arenas, where typically one supply pair is added for every four to eight outputs, depending on output drive capabilities and package characteristics.

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 159
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

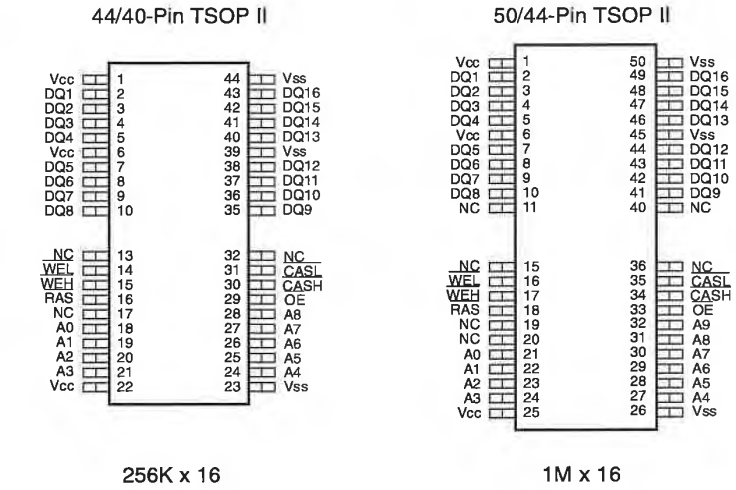


Figure 6-2 The ratios of supply pins to signal pins in ×16 DRAMs are comparable to the ratios in CPUs and ASICs.

A few examples illustrate the relationship between DRAM width and package size. A 16-Mbit ×32 device in a 70-pin fine-pitch TSOP II package covers 0.47 square inches (3.03 cm^2). With this package type, a 16-Mbit ×32 device pays a 38% board area premium over the ×4 part of comparable density. This premium may be sufficiently large to cause a shift in the package type of choice for very wide DRAMs away from SOJ and TSOP II and toward one of the square, fine-pitch packages used for small gate arrays such as Plastic Quad Flat Packs (PQFPs). This trade-off is evident in some of the wide ×32 graphics RAMs.

The primary negative effect on the temporal interface is the slightly reduced speed of the wide parts compared with the speed of the narrower generic ones. The large currents associated with driving many output buffers and large internal nodes make it difficult to drive those signals quite as fast as in a ×1 or a ×4 part. Consequently, wide parts are either unavailable in the fastest speed grade, or are available late and at higher prices.

Higher chip power—but lower system power

On the power consumption front, wide parts burn more power than narrow ones during all types of operations. A ×16 part draws about twice as much power per operation as a ×1 part of

Samsung Ex. 1097, p. 160
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

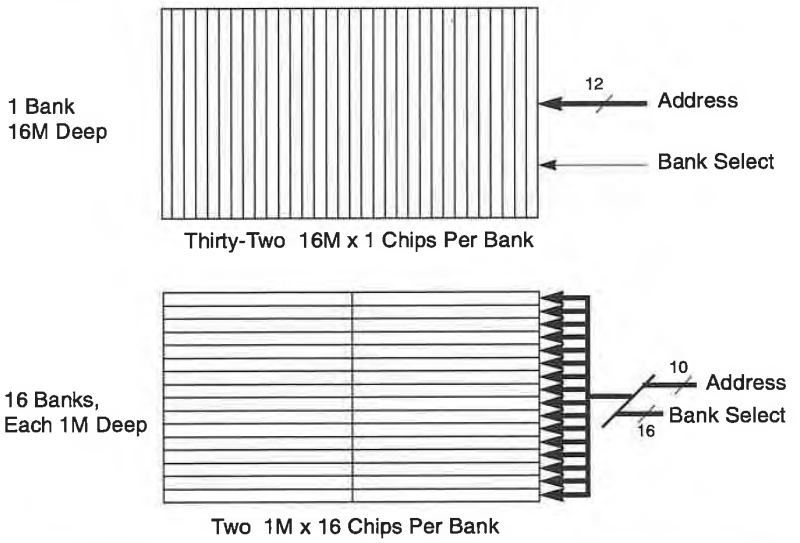


Figure 6-3 In a horizontally organized memory system, only a few chips are active at a time.

the same size and technology. However, designers of portable systems find that the battery life of a machine is maximized through the use of wide parts because fewer devices are active at a time. In addition to the granularity issues often vital in small portable systems, the preference for wider parts is driven by the lower average power needed for reads and writes. Figure 6-3 illustrates that in the extreme case, a 64-Mbyte memory system built with 16M×1 DRAMs uses all the DRAMs during all reads and writes, whereas a memory system built of 1M×16 DRAMs uses just two of the thirty-two total DRAMs. Since a quiescent DRAM draws between 1 and 0.001% of an active DRAM's power, the power saved during row accesses (reads and writes) overwhelms the 50–100% penalty during refresh-only periods.

Cost premiums for wide DRAMs are built in

In addition to taking more board area, wide devices suffer from four cost-related penalties:

- Larger die sizes
- Larger, more expensive packages
- Higher test costs
- Low volumes and, consequently, larger price premiums

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 161
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

This price gap will diminish as the peak demand shifts toward wider DRAMs. Just as ×4 parts were more expensive than ×1 parts at the 256-Kbit level, but achieved cost parity at the 4-Mbit level, so will the price premiums for the ×8s and ×16s decline as their sales increase. However, the intrinsic cost differences are large enough that price parity between ×4 and ×16 devices of the same density will come before the 64-Mbit level. Only with the 64-Mbit generation will the number of ×4 devices shipped be small enough for the narrow DRAMs to become a specialty device.

The cost structure of some wide parts is impacted by the fact that the wide parts are derived from a separate design. Many vendors' 512K× 16 devices are derived from different designs than their 4M×1 and 1M×4 devices. A separate, lower volume design incurs more overhead per DRAM sold than its higher volume cousin running on the same fabrication line. More significantly, these wide designs have slightly larger dice than their narrow counterparts. The additional die area comes from wider buses into each section of the DRAM array and more output buffers. Each 1% increase in area results in about a 2% cost penalty.

Of course, the ×9 and ×18 parts must come from a different design than the ×1, ×4, ×8, and ×16 parts because they have more bits: 4,718,592 for the ×9 and ×18 parts versus 4,194,304 for the others at the 4-Mbit level. These specialty parts are among the lowest volume DRAMs and for this reason demand an additional price premium on a per-bit basis. This price premium strongly motivates systems designers to fabricate a ×9 width by combining two ×4 parts with a ×1 part of a lower density.

The benefits of increased width

Chapter 2 illustrates the two main unfortunate consequences of continued exponential increase in DRAM density:

1. Exponential growth in the granularity of memory systems is overtaking the growth in average memory size.
2. Bandwidth requirements are surpassing the maximum total bandwidth available within many memory systems.

Wide conventional DRAMs primarily address the first of these problems. When the depth of the DRAMs and the corresponding memory banks are reduced, the granularity of the memory system is reduced correspondingly. A ×32 memory system built of 16M-deep DRAMs has a granularity of 64 Mbytes. In contrast, if

Samsung Ex. 1097, p. 162
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

wider, 1M-deep, DRAMs are used, the granularity of expansion is reduced to 4 Mbytes.

First ×32 generic DRAM

In late 1995, Toshiba introduced the first ×32 EDO DRAM. Intended primarily for graphics applications, it is intended as general replacement for the 256K×16 EDO DRAMs that are in common use with 64-bit GUI accelerators. Its performance, shown on Figure 6-6, seems well suited for that marketplace, but in the long run it will have a hard time competing against more specialized ×32 graphics memories such as SGRAM and WRAM. It is also well suited to embedded systems that have ×32 buses but require only 2 Mbytes of memory.

Granularity and fill frequencies

A side effect of reduced granularity is an increase in the total available pin bandwidth in the system. In comparison to a 16M×1 part, a 1M×16 part has ×16 more I/O pins per bit stored and, consequently, ×16 more pin bandwidth per memory bit. In a time when, in many memory systems, all available pin bandwidth is used, increases in memory system bandwidth can come only through increases in 1) memory system size, 2) bandwidth per pin, or 3) bandwidth per bit. Increasing the width of the DRAM is a simple and cost-effective way to implement the third option.

These effects are clearly visible in the devices' fill frequencies. Figure 6-4 shows that at the 16-Mbit level, ×4 fast-page-mode devices have unacceptable fill frequencies for the PC industry. The ×16 devices, while acceptable for PC main memory, are completely unacceptable for graphics applications.[2] For this reason, the 16-Mbit generation only began to penetrate the PC industry when price parity between the 4-Mbit 1M×4 and the 16-Mbit 1M×16 was attained in mid 1995. At the 64-Mbit level, even the ×16 organization has unacceptable fill frequencies for main memory applications. In fact, even a ×32 device would be essentially inadequate, and a 1M×64 would be needed to meet the basic bandwidth and granularity requirements of a PC's main memory. Such a part would be prohibitively expensive in comparison to some of the other alternatives. Given the basic inadequacy of wide, fast-page-mode DRAMs, the industry's search for alternatives at both the 16- and 64-Mbit levels is not surprising.

2. See Section 2.5 for a description of the fill frequency and its utility.

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 163
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

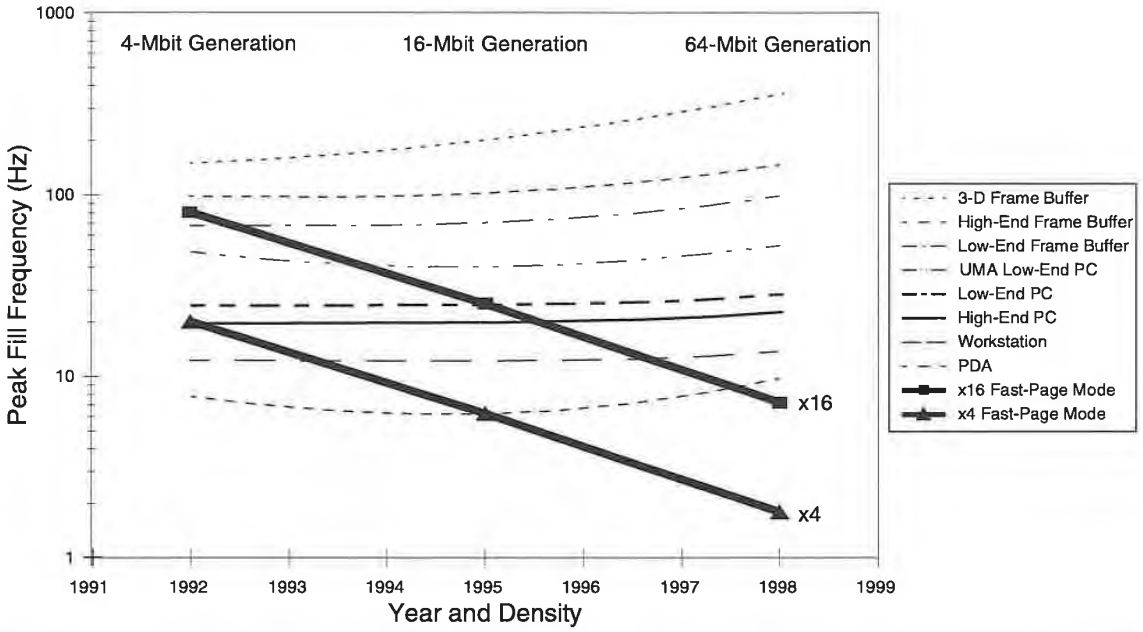


Figure 6-4 Fast-page-mode fill frequencies are barely adequate at the 16-Mbit, and below the needs of all important application domains at the 64-Mbit level—even in ×16 widths.

6.3 Extended Data-Out DRAMs

Extended Data-Out (EDO) refers to a modification of the standard fast-page-mode logical interface. With generic fast-page mode, the DRAM's output buffer is turned off on a read when $\overline{\text{CAS}}$ goes high. This means that $\overline{\text{CAS}}$ must be held low long enough for the data to come out and travel to the memory controller or CPU, where it can be latched. Only then can the column precharge begin. In EDO DRAMs, the output buffer is controlled by the output enable ($\overline{\text{OE}}$) and by $\overline{\text{CAS}}$. Figure 6-5 shows the data output pins for a normal fast-page-mode DRAM and for an EDO DRAM. Because the data remains valid on the output until the next falling edge of $\overline{\text{CAS}}$, the transfer of the data to the latch in the memory controller can be overlapped with the next column precharge. Consequently, burst transfers can realistically be cycled up to 30% more quickly.

Samsung Ex. 1097, p. 164
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

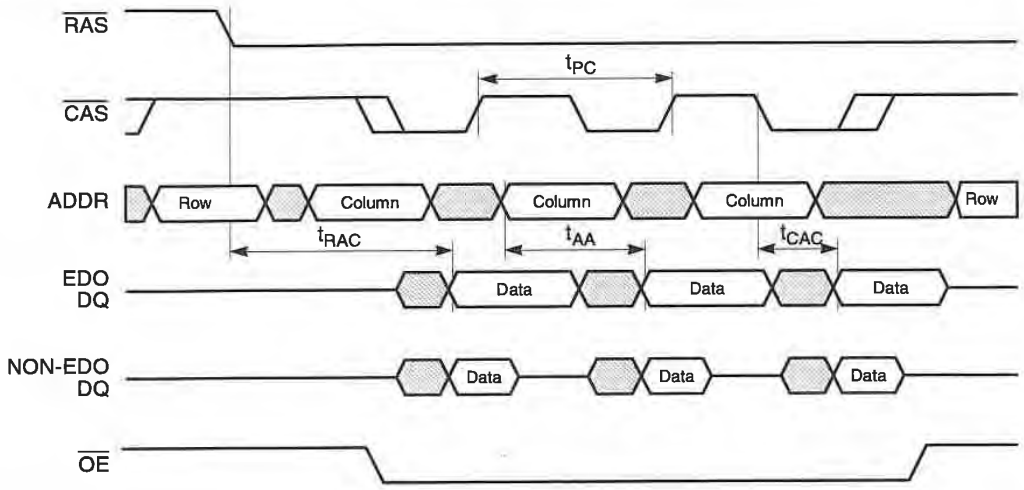


Figure 6-5 Extended Data-Out (EDO) DRAMs preserve data on the output pins even after $\overline{CAS}$ is raised. (Source: *Micron 1993 DRAM Data Book*).

For example, with a 60-ns DRAM, the fastest possible CAS cycle time of 40 ns gives a minimal 5 ns of valid data-out. In contrast, a 60-ns EDO part facilitates a 25-ns CAS cycle time with a more reasonable 10 ns of valid data out. To get 10 ns of data out, a regular fast-page-mode DRAM must be operated with a CAS cycle time of 45 ns.

Table 6-1 compares the realistic column cycle time for fast-page-mode and EDO DRAMs. The most significant differences are for the 60-ns and 70-ns parts. In a 66-MHz system (15-ns cycles) a fast-page-mode memory system is able to return a new piece of data every 3 cycles,[3] while an EDO memory system cuts the data repeat rate to 2 cycles.

EDO parts also implement a slight change in the definition of Write Enable ($\overline{WE}$). Early EDO parts were unaffected by the state of $\overline{WE}$ while $\overline{CAS}$ is high in the middle of a series of column accesses: while $\overline{CAS}$ is high, the output buffer is controlled only by $\overline{OE}$. However, standard SIMMs do not use $\overline{OE}$: all $\overline{OE}$ pins are tied to ground on the SIMM and are not brought out to the

3. Due to the realities of controller design, most Pentium chip sets achieve a 4-cycle repeat rate with -7 fast-page-mode DRAMs at 66 MHz.

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 165
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

Row Access Time	Column Cycle Time	
	Fast-Page-Mode DRAM	EDO DRAM
80 ns	50 ns	35 ns
70 ns	45 ns	30 ns
60 ns	40 ns	25 ns
50 ns	35 ns	20 ns

Table 6-1 EDO cycle times are an improvement over fast-page-mode cycle times.

connector. Consequently, when these EDO DRAMs were placed on SIMMs, the output buffer is always on during a series of column accesses. Unfortunately, having the output buffer continuously on is incompatible with interleaving two or more SIMMs on a shared data bus.

As standardized by JEDEC, interleaving EDO SIMMs on a shared data bus is facilitated by having $\overline{WE}$ or $\overline{OE}$ control the output buffer while $\overline{CAS}$ is high. If $\overline{WE}$ is pulsed low or $\overline{OE}$ pulsed high while $\overline{CAS}$ is high, the output buffer is turned off until the start of the next column read operation. As long as $\overline{WE}$ returns high at least one setup time before the start of the next column access, no ambiguity exists regarding whether $\overline{WE}$ is just controlling the output buffer or whether it is actually initiating a write access.

Hyper-Page Mode

In a characteristic lack of solidarity regarding feature names, the DRAM vendors are split on whether to call these devices Extended Data-Out or Hyper-Page Mode parts. *Hyper-Page Mode* is promoted as the next logical step in the progression from page mode in the mid 1980s to fast-page mode in the late 1980s. At this point, a small number of large companies are committed to *Hyper-Page Mode*; they are unlikely to switch names in the near future—in spite of JEDEC's official adoption of *Extended Data Out Fast Page Mode* as the part's official name.

Advantages of EDO DRAMs

The important advantages of EDO DRAMs include:

1. Only minor changes are required to existing fast-page-mode DRAMs. DRAM vendors can incorporate these changes into their primary designs with little penalty. As a result, EDO

Samsung Ex. 1097, p. 166
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

DRAMs can ride the same yield and cost curves as high-volume generic DRAMs. EDO devices command a price premium of less than 5%.

2. Existing memory controllers need only minor changes to take advantage of EDO DRAMs. In fact, EDO DRAM timings are compatible with those of generic DRAMs. Thus, EDO DRAMs can be used in many existing applications as fast-page-mode DRAMs, though without any performance increase. Given favorable pricing, system manufacturers with a mix of EDO and non-EDO designs could purchase and inventory only EDO parts for all systems. As mid 1995, most PC chip-set vendors are producing or sampling devices that support EDO memories. Intel's Triton chip set, which is currently dominating the PC motherboard market, was among the first chip set in production with EDO support. By late 1995, EDO support became a check-list item without which no chip-set vendor can survive.

3. It is relatively easy for a memory controller to differentiate at power up whether a bank of memory is fast-page mode or EDO:

   - Write a 0 to address A0 and a 1 to address A1.
   - Read A0, latching the data after $\overline{\text{CAS}}$ has risen.
   - Read A1, latching the data after $\overline{\text{CAS}}$ has risen.

   If the two pieces of data are 0 and 1, then the memory is EDO. If the two pieces are both the same, then the memory is fast-page mode. This process is made more reliable by a pull-up resistor on the data line being used for the test.

   The ability to check for EDO versus fast-page mode at power up means that for the end user, EDO and fast-page mode are interchangeable. The same SIMMs and DIMMs can be used for both and freely intermixed on a bank-by-bank basis.

4. Some wide EDO DRAMs for graphics applications have been in volume production since 1993. As of mid 1995, most of the first-tier DRAM vendors were already in volume production with EDO parts. As the remaining suppliers come on-line with mainstream EDO parts by mid 1996, the remaining price premium should entirely disappear. In the longer term, some vendors are predicting a premium for fast-page mode over EDO.

When the price premium for EDO DRAMs drops to zero, then EDO DRAMs will truly provide something for nothing compared with their fast-page-mode counterparts. Such a deal is hard to refuse, especially in light of EDO DRAM's backward

© 1996 Steven Przybylski

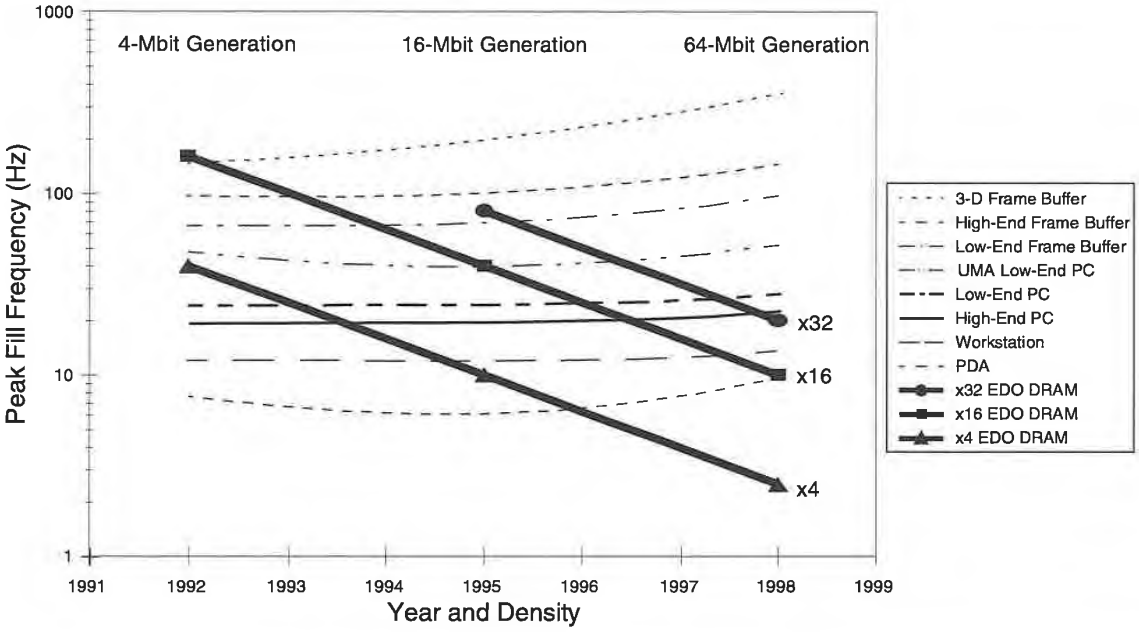


Figure 6-6 EDO peak fill frequencies are somewhat higher than those of fast-page-mode devices.

compatibility. Many suppliers are expecting EDO devices to rapidly win over a substantial portion of the fast-page-mode DRAM market. Over the next few years, EDO DRAMs will fill the growing need for a low-cost, intermediate-performance alternative to generic DRAMs..

EDO DRAMs provide a convenient bridge to **Burst EDO DRAMs**, **Synchronous DRAMs (SDRAMs),** and the other evolutionary and revolutionary alternatives: EDO DRAMs provide less bandwidth per pin than the more aggressive designs, but they do so without the cost premiums that the evolutionary and revolutionary alternatives currently command. When the cost premiums for Burst EDO DRAMs, SDRAMs, and their kin come down, EDO DRAMs may lose their appeal. However, the emergence of EDO DRAMs has clearly delayed the shift to higher performance organizations.

Wide EDO DRAMs have reasonable fill frequencies

The main characteristics that define the conventional alternatives—greater width and modified CAS semantics—are orthogonal. They can be—and are—freely mixed. The first EDO

Samsung Ex. 1097, p. 168
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

devices were in fact the ×16 devices. In the future, all new conventional parts, narrow as well as wide, will be designed as both fast-page-mode and EDO devices. Figure 6-6 shows that the peak fill frequencies of EDO devices are noticeably higher than their fast-page-mode counterparts (see Figure 6-4). In particular, the 1M×16 EDO device comfortably meets the needs of the PC marketplace, while the fast-page-mode devices were adequate but not exceptional in this role

More significant is that at the 64-Mbit density, the ×16 EDO device is just as inadequate for PCs as the equivalent fast-page-mode device. At the 64-Mbit density, even a ×32 EDO will have marginal performance relative to its size, especially in light of the 16-Mbyte granularity that it generates in a ×64 memory system.

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 169
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

Samsung Ex. 1097, p. 170
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

7 Evolutionary Alternatives

Evolutionary DRAM alternatives augment the existing generic DRAM interface. The starting point for these alternatives is the standard separate data bus and a multiplexed address bus controlled by $\overline{RAS}$ and $\overline{CAS}$. The effectiveness of the interface is improved by changing aspects of the logical, temporal, and/or electrical interface. Some of these changes increase bandwidth by increasing the interface's operating frequency and cycle-by-cycle efficiency. Other changes are aimed at improving the efficacy of the DRAM's row cache. Advocates of the evolutionary alternatives argue that limiting the amount of change to the interface eases the transition to higher performance systems and minimizes the associated cost premiums.

The evolutionary approaches are presented in order of increasing degrees of change from the generic DRAM interface; the discussion begins with Burst EDO DRAMs (Section 7.2), Enhanced DRAMs (Section 7.3), Synchronous DRAMs (Section 7.4), and concludes with Cached DRAMs (Section 7.5).

7.1 New In The Second Edition

Since 1994, substantial changes have occurred among the revolutionary alternatives; these changes are reflected in this chapter. A new section, Section 7.2, describes Burst EDO DRAMs. Section 7.4 on SDRAMs has been substantially revised and expanded. The sections on EDRAMs and CDRAMs have also been updated. Of special interest are: fill frequency data for all

Samsung Ex. 1097, p. 171
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

alternatives, the comparisons of Burst EDO DRAMs and SDRAMs, and an updated list of SDRAM cross-vendor incompatibilities.

7.2 Burst EDO (BEDO) DRAMs

Burst EDO DRAMs constitute the latest and one of the most controversial of the new DRAM technologies. They represent the final evolution of the generic asynchronous interface, following the lineage from page-mode through fast-page-mode and EDO DRAMs. The Burst EDO interface also shows great similarity to the nibble-mode interface. Nibble mode was introduced along with page mode and static-column mode in the mid 1980s, but has fallen by the wayside in recent years (see Section 3.3.1).

Like nibble-mode DRAMs, Burst EDO DRAMs are good at providing bursts of four pieces of data at high speed and with limited intervention from the controller. They are particularly targeted at PCs, in which all traffic between the caches and main memory involves bursts of four and processor bus clock frequencies of up to 66 MHz.

7.2.1 Physical, Electrical, Pin-Level Interface

Among the greatest strengths of the Burst EDO interface is its compatibility with both fast-page mode and EDO at a number of levels. Specifically, the physical, electrical, and pin-level interfaces are identical in all three cases: Burst EDO devices are in the same packages with the same signals on the same pins (i.e., the same pinouts), and they use the same LVTTL levels.

Consequently, the same SIMM and DIMM configurations work just as well for Burst EDO as for EDO and fast-page mode. There is no strict requirement for different physical keying of modules, nor is there any difference in the module's presence-detect pins. Just as a memory controller can distinguish between EDO and fast-page mode at power up, it can easily differentiate between fast-page mode, EDO, and Burst EDO. All the controller need do is perform a read burst sequence while changing the address bus. If the data out indicates that the change in the address bus had no impact on the data stream, then the DRAMs are Burst EDO DRAMs. If they are not Burst

Samsung Ex. 1097, p. 172
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

EDO devices, then the algorithm given in Section 6.3 can be used to determine whether they are fast-page mode or EDO.

Burst EDO DRAMs will be available in the same 168-pin ×64 DIMMs that promise to be the next standard SIMM/DIMM for PCs. These DIMMs will come in buffered and unbuffered versions. The unbuffered DIMMs are destined to enjoy more popularity since they are the ones that facilitate 66-MHz operating frequency. As defined by JEDEC, the buffered version of the 168-pin DIMM has non-inverting line buffers on all the address and control signals except $\overline{\text{RAS}}$.

As of 3Q1995, only 16-Mbit, 3.3 V Burst EDO devices are available. Widths of 4, 8, and 16 bits will be available. Micron's 3.3 V parts have 5 V tolerant I/Os, and so are easily clustered with a small 3.3 V regulator to make 5 V SIMMs and DIMMs. The next generation, 64-Mbit, Burst EDO devices have been announced in the ×4, ×8, and ×16 widths.

7.2.2 Temporal and Logical Interface

The differences between Burst EDO DRAMs and EDO and fast-page-mode DRAMs are concentrated in the temporal and logical interfaces. Figure 7-1 shows a timing diagram for a typical read burst of four. The initiation of the row access begins as with the generic and conventional alternatives: $\overline{\text{RAS}}$ is lowered any time after a row address is stable on the address bus. The column address of the start of the burst is then presented and $\overline{\text{CAS}}$ is dropped. However, unlike the generic and conventional alternatives, the read data does not come out until $\overline{\text{CAS}}$ is toggled a second time: column reads and writes are pipelined with a pipeline depth of two.

Subsequent toggles of $\overline{\text{CAS}}$ bring out subsequent pieces of data until four pieces of data have been presented. The address bus is automatically sampled on the fifth falling edge of $\overline{\text{CAS}}$ and every fourth one thereafter for a new column address. In this way, a continuous stream of data can be generated. In fact, to prevent a new burst from starting, the new burst must be explicitly terminated. Bursts can be terminated at any time by changing the state of $\overline{\text{WE}}$ or by raising $\overline{\text{RAS}}$. The latter of course closes the row and initiates a precharge of the memory array. Figure 7-1 shows a second burst being terminated by a change in $\overline{\text{WE}}$ before it really gets going.

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 173
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

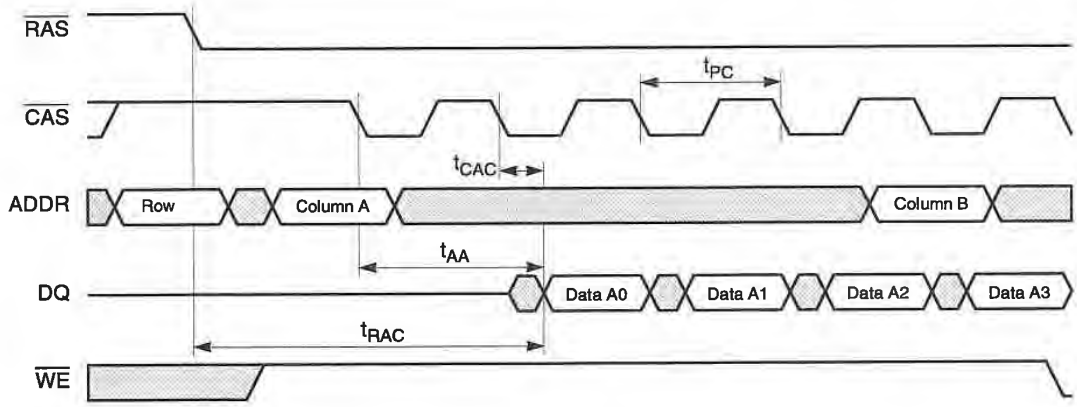


Figure 7-1 In a Burst EDO DRAM, five CAS transitions are needed to read the first four pieces of data.

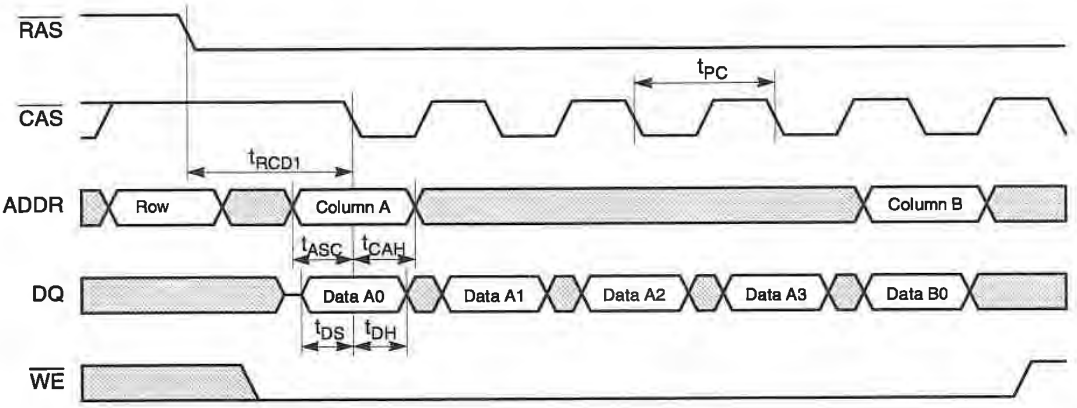


Figure 7-2 Burst EDO write bursts only require four CAS transitions.

Samsung Ex. 1097, p. 174
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

Write bursts are fairly similar to reads except that write bursts are not pipelined. The first piece of data is latched with the column address on the first falling edge of $\overline{CAS}$, and the subsequent three chunks of data are latched on the next three falling edges of $\overline{CAS}$. Writes are indicated by $\overline{WE}$ being lowered before the fall of $\overline{CAS}$ and being held low until $\overline{CAS}$ rises for the last time. Raising $\overline{WE}$ before the end of the burst terminates the burst. Figure 7-2 shows one burst of four, followed immediately by a second burst that is explicitly terminated after one piece of data is transferred.

Because write data is not pipelined, only four CAS pulses are needed for a write burst of four. The conflict between the pipelined read transactions and the non-pipelined write transactions forces a delay of about one cycle between the last CAS pulse of a read burst and the first CAS pulse of an immediately following write transaction. No such delay is needed when a write burst is followed by a read burst.

JEDEC-standardized *Pipelined Nibble Mode*

Burst EDO devices are very much like the older nibble-mode devices, except that read transactions are pipelined. Since the term *burst* has special connotations within JEDEC—i.e. long, linear bursts—the standardizing body's name for Burst EDO is Pipelined Nibble Mode. This circumstance of a popular name which differs from the standardized name is not uncommon in the DRAM world. The situation results from a company or group of companies developing and commercializing devices prior to seeking standardization.

Both burst types are supported

Burst EDO devices support both linear bursts and **sub-block-ordered** bursts.[1] However, in Burst EDO devices the burst type is programmed differently than in all of the other architectures. The burst type for Burst EDO devices is set by performing a Write-and-CAS-before-RAS (WCBR) cycle. This special cycle is initiated by having $\overline{WE}$ and $\overline{CAS}$ low on the falling edge of $\overline{RAS}$. The state of A0 on the falling edge of $\overline{RAS}$ indicates whether the device should perform linear or sub-block ordered bursts. Normal operation is resumed by performing a RAS-only or CAS-before-RAS (CBR) refresh cycle. WCBR cycles exist for fast-page-mode and EDO devices, but only for entrance into special test modes.

Using WCBR cycles to program the burst type preserves maximum compatibility with the conventional alternatives. No new

1. Sub-block ordered bursts are also called Intel-ordered or interleaved bursts. See the glossary entry in Appendix B for details.

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 175
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

pins are required, and the meaning of any of the control signals during normal operations is unaltered except for the minimal changes necessary to implement the burst functionality. The drawback of using WCBR cycles for changing the burst type is the associated temporal overhead. Two full row cycle times (at least 220 ns for a -6 part) are needed to change the burst type and then restore normal operation. If a row is open, it takes close to three row cycle times to close the row, perform the WCBR cycle, exit to normal operation, and then reopen the same or another row.

For most systems, this temporal overhead for changing the burst type is not an issue. For the majority of applications, the burst type will be set one time after power up and not changed again. Since a linear burst and a sub-block-ordered burst are identical if the burst begins at the start of the block, Burst EDO DRAMs in sub-block-ordered mode can satisfy many linear-ordered transfers.

However, in some systems—especially UMA systems with a mix of sub-block-ordered processor requests and linear-ordered graphics requests—the burst type may change quite frequently. In these systems, the overhead from changing the burst type can become noticeable in both the typical and the worst cases. Of secondary importance is the additional burden on the memory controller of remembering which state each bank is in and managing that state appropriately. In general, if the overhead due to burst-type switching becomes too great then it becomes worthwhile to keep the DRAMs in sub-block-ordered mode and rearrange the data words in the memory controller when a linear transfer with a non-zero start address is encountered.

Burst EDO devices are very strongly burst devices. A significant penalty is paid for performing nonburst transactions. Specifically, if a single transfer is desired from each column address, then the maximum repeat rate is one piece of data every four cycles with -5 parts at 66 MHz or with -6 parts at 60 MHz. Pipelining of the read transactions and the autonomous nibble counter hurts because a temporal penalty is paid to stop a burst that is in progress. This characteristic makes Burst EDO a poor choice for graphics applications in which short bursts are common and not clusterable into longer transactions. Burst EDO is also inconvenient for higher performance systems which reach higher peak bandwidths through interleaving. Interleaving decreases the number of chunks in a burst of a fixed total size (see Section 4.2.4).

Samsung Ex. 1097, p. 176
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

Parameter	Symbol	Units	Fast-Page Mode			EDO			Burst EDO		
			-7	-6	-5	-7	-6	-5	-7	-6	-5
Row Access Time	t_{RAC}	ns	70	60	50	70	60	50	70	60	**52**
Row Cycle Time	t_{RC}	ns	130	110	95	130	110	**89**	130	110	**90**
Row Precharge	t_{RP}	ns	50	40	35	50	40	35	50	40	**30**
Column Address Access Time[a]	t_{AA}	ns	35	30	25	35	30	25	35	**28.2**	25
CAS Access Time[b]	t_{CAC}	ns	20	15	13	20	15	13	**15**	**11.6**	**10**
CAS Cycle Time	t_{PC}	ns	45	40	35	**30**	**25**	**20**	**20**	**16.6**	**15**
RAS to CAS Delay	t_{RCD}	ns	20	20	20	**14**	**14**	**14**	**20**	**16.6**	**15**

Table 7-1 Timing parameter differences between fast-page-mode, EDO and Burst EDO DRAMs (Sources: *Micron MT4LC1M16C3, MT4LC1M16E5, and MT4LC1M16H5, and IBM IBM0117805 and IBM0116160 Data Sheets*).

a. In Burst EDO, t_{AA} is from the CAS edge that latches the column address.
b. In Burst EDO, t_{CAC} is from the second through fifth CAS falling edge to the data out.

Important changes to the temporal interface

Burst EDO devices make use of extended data-out times: as in EDO devices, the data is not withdrawn from the data pins as soon as $\overline{CAS}$ is raised.

Table 7-1 compares the important timing parameters of fast-page-mode, EDO, and Burst EDO devices. The access times—both row and column—of the Burst EDO devices are the same as their fast-page-mode and EDO counterparts. Pipelining the read transactions allows a substantial decrease in the column cycle time over that of the other conventional alternatives. Another important change is in t_{AA}—the access time from a column address. In fast-page-mode and EDO DRAMs, t_{AA} is measured from a stable address on the DRAM address pins. However, with Burst EDO devices, t_{AA} is measured from the first falling edge of $\overline{CAS}$—the one that latches in the column address. The other column access time, t_{CAC}, is measured from the second and subsequent falling edges of $\overline{CAS}$. t_{CAC} is carefully chosen to give 5ns of valid data out time in 50-, 60-, and 66-MHz systems with -7, -6, and -5 Burst EDO parts.

Another important detail to note regarding the Burst EDO timing parameters is the row access time of the -5 part. This 52-ns row access time derives from the desire to have DRAMs respond to a row access in four cycles at 66 MHz. In contrast, the -5 fast-page-mode and EDO parts have 50-ns row access times. This 2-ns difference may seem insignificant, but it is not. Most

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 177
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

16-Mbit devices are built on 0.5-µm technologies. These technologies generally do not produce much -5 product. The -5 bin is not expected to become mainstream until the industry moves *en masse* to the 0.4–0.35-µm 64-Mbit technologies. Though some companies are ramping up these lines in 1995, by far most of the capacity is still at the larger geometries. Given that 50-ns row access times are a stretch, the difference between 50 and 52 ns is significant.

7.2.3 System Design Issues

Burst EDO DRAMs are specifically targeted towards Pentium-class PCs with 50-MHz and higher processor buses. The degree to which Burst EDO DRAMs meet the needs of these systems is shown in Table 7-2. The table shows the absolute best case lead-off times and transfer rates in PCs for fast-page-mode, EDO, and Burst EDO memory systems. Lead-off times are given from the DRAM's perspective and assume precharged DRAMs: only the time taken by the row access is accounted for. Data transmission time, chip-set overhead, and processor setup time are not included. These can contribute to both increased lead-off times and slower transfer rates (see Section 10.3).

Table 7-2 corroborates the assertion that the interface changes from fast-page mode to EDO to Burst EDO do not affect the lead-off latencies. Only the transfer rates for the second through fourth data chunks are affected. Table 7-2 also shows that although EDO is adequate to achieve single-cycle transfer rates (x-1-1-1) at up to 50 MHz (best case), Burst EDO extends those rates up to 66 MHz.

The transfer rates in Table 7-2 assume generic full rail 3.3 V LVTTL levels. The cost of maintaining these standard levels is a limitation in the maximum configuration size. Generally speaking, a principal concern in the design of PC memory systems is the total capacitive load on the $\overline{\text{CAS}}$ lines. The other control signals are typically duplicated or go to a single rank of DRAMs; the address lines typically have plenty of time to settle prior to an edge on $\overline{\text{RAS}}$ or $\overline{\text{CAS}}$.[2] To get to 66-MHz timing with -5 Burst EDO DRAMs, at most nine CAS loads can be driven, assuming standard high-current ASIC output drivers. This translates to four DIMMs or SIMMs with two CAS loads per

2. This is different than for fast-page-mode or EDO memory systems, in which the access time from a stable address is important. In Burst EDO devices, the address is latched by $\overline{\text{CAS}}$, and so the arrival time of the address at the DRAM is irrelevant as long as it is at least a setup time ahead of the fall of $\overline{\text{CAS}}$.

Samsung Ex. 1097, p. 178
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

Processor Frequency (MHz)	Bus Frequency (MHz)	Fast-Page Mode			EDO			Burst EDO		
		-7	-6	-5	-7	-6	-5	-7	-6	-5
25	25	2-2-2-2	**2-1-1-1**	2-1-1-1	2-1-1-1	2-1-1-1	2-1-1-1	2-1-1-1	2-1-1-1	2-1-1-1
33	33	3-2-2-2	3-2-2-2	2-1-1-1	**3-1-1-1**	3-1-1-1	2-1-1-1	3-1-1-1	3-1-1-1	2-1-1-1
40	40	3-2-2-2	3-2-2-2	3-2-2-2	3-2-2-2	**3-1-1-1**	3-1-1-1	3-1-1-1	3-1-1-1	3-1-1-1
50	50	4-3-3-3	4-2-2-2	3-2-2-2	4-2-2-2	4-2-2-2	**3-1-1-1**	**4-1-1-1**	4-1-1-1	3-1-1-1
60	60	5-3-3-3	4-3-3-3	4-2-2-2	**5-2-2-2**	4-2-2-2	4-2-2-2	5-2-2-2	4-1-1-1	4-1-1-1
66	66	6-4-4-4	5-3-3-3	4-2-2-2	**6-2-2-2**	5-2-2-2	4-2-2-2	6-2-2-2	5-2-2-2	4-1-1-1
80	40	3-2-2-2	3-2-2-2	3-2-2-2	3-2-2-2	**3-1-1-1**	3-1-1-1	3-1-1-1	3-1-1-1	3-1-1-1
90	60	5-3-3-3	4-3-3-3	4-2-2-2	5-2-2-2	4-2-2-2	4-2-2-2	5-2-2-2	**4-1-1-1**	4-1-1-1
100	66	6-4-4-4	5-3-3-3	4-2-2-2	6-2-2-2	5-2-2-2	4-2-2-2	6-2-2-2	5-2-2-2	**4-1-1-1**
120	60	5-3-3-3	4-3-3-3	4-2-2-2	5-2-2-2	4-2-2-2	4-2-2-2	5-2-2-2	**4-1-1-1**	4-1-1-1
133	66	6-4-4-4	5-3-3-3	4-2-2-2	6-2-2-2	5-2-2-2	4-2-2-2	6-2-2-2	5-2-2-2	**4-1-1-1**
150	60	5-3-3-3	4-3-3-3	4-2-2-2	5-2-2-2	4-2-2-2	4-2-2-2	5-2-2-2	**4-1-1-1**	4-1-1-1
150	75	6-4-4-4	5-4-4-4	5-3-3-3	6-3-3-3	5-2-2-2	5-2-2-2	6-2-2-2	5-2-2-2	5-2-2-2
166	66	6-4-4-4	5-3-3-3	4-2-2-2	6-2-2-2	5-2-2-2	4-2-2-2	6-2-2-2	5-2-2-2	4-1-1-1

Table 7-2 Best case lead-off and transfer rates in PCs from the DRAM's perspective, starting with a precharged DRAM. The shaded entries highlight noteworthy speed bins and frequencies at which Burst EDO or EDO perform better than EDO or fast-page mode, respectively.

module: a total of four ranks.[3] Four ranks is more than adequate for most PC configurations, which typically allow for two or three ranks. Only the larger server-oriented motherboards facilitate four or more ranks of DRAMs.

Burst EDO from the same die as fast-page mode and EDO

One great strength of the Burst EDO architecture is the ease with which an EDO DRAM design can be converted into a Burst EDO DRAM design. In fact, so few additional gates are needed that it is easy to generate a single design that is customized by bond option or metal mask. For example, Micron's second and third generation 16-Mbit database can produce fast-page-mode, EDO, or Burst EDO devices of widths 4, 8, or 16.

This design commonality is tremendously valuable for fab line managers. Because four to six months normally pass from raw wafer to final, salable DRAM, the ability to select the product toward the end of the manufacturing process instead of at the beginning allows for greater flexibility and market responsiveness.

3. Section 4.2.2 defines a *rank* as a physical collection of DRAMs, while reserving the word *bank* for a logical collection of bits.

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 179
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

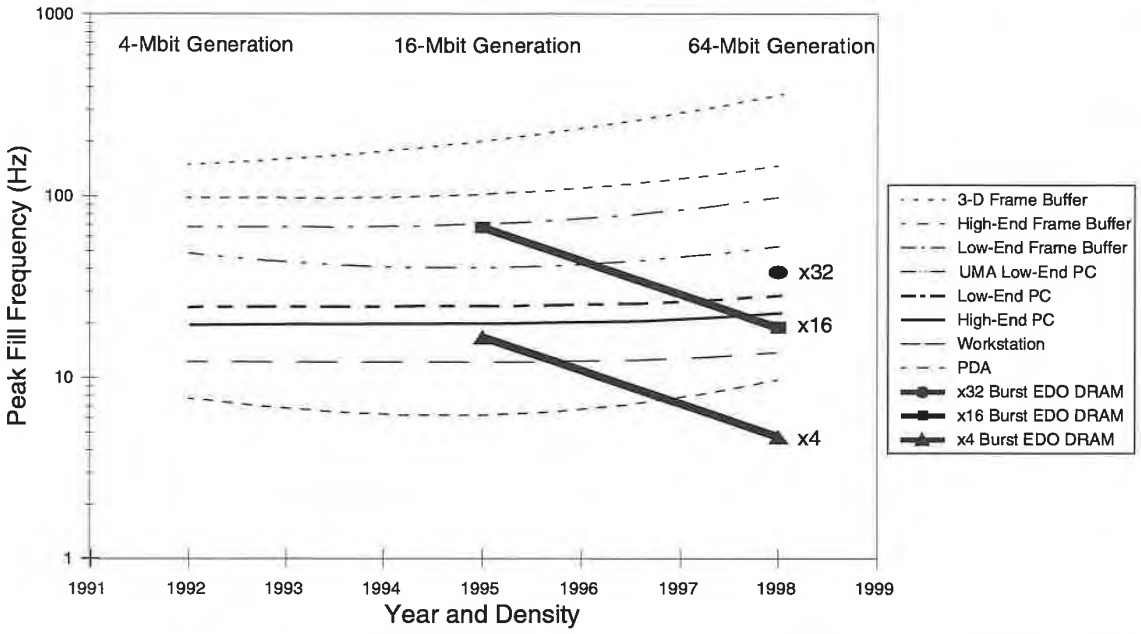


Figure 7-3 Burst EDO devices have quite acceptable peak fill frequencies at the 16-Mbit level, but fall behind at the 64-Mbit level except for ×32 and larger widths.

7.2.4 Burst EDO Fill Frequencies

Because of Burst EDO's higher peak transfer rates, its peak bandwidths are substantially higher than those of EDO or fast-page-mode devices. Consequently, at a given density and width, Burst EDO devices offer better peak fill frequencies. These are shown in Figure 7-3. At the 16-Mbit level, the ×16 device offers better peak performance than is required by most applications. Even the ×4 16-Mbit Burst EDO has pretty respectable fill frequencies for many PC applications. However, the granularity generated by the standard bus widths often makes the latter part undesirable despite its high bandwidth.

At the 64-Mbit level, the ×16 device falls to a level that is probably unacceptable for PCs. Regardless of its operating frequency, the memory system granularity that a 4M × 16 Burst EDO generates in ×64 systems (i.e., 32 Mbytes) will be unacceptable for all but the largest PC configurations, The 4M × 16 Burst EDO will remain a good device for the workstation market. To satisfy

Samsung Ex. 1097, p. 180
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

the larger PC marketplace, the width of the device will have to increase to at least ×32.

7.2.5 History, Current Status, and Future Plans

Burst EDO was developed and quietly promoted behind the scenes throughout 1994. Most of the work was done by Micron Semiconductor, with the help and input of many of its larger customers and partners. In early 1995, with working devices in hand, Micron announced the part and almost simultaneously took it to JEDEC for adoption as a standard. At the time Micron had the public support of three other semiconductor manufacturers and six chip-set vendors. Although most other semiconductor manufacturers were openly neutral to the proposal, a couple actively opposed it. At that time, it was expected that Burst EDO would start showing up in systems toward the very end of 1995 or, more likely, in 1996.

Through the spring and early summer of 1995, Burst EDO gained popularity and momentum. However, many DRAM vendors were concluding that yielding adequate supplies of the -5 Burst EDO DRAMs would be difficult on their existing high-volume fabrication lines. In theory, the demand for Burst EDO would be split between the -6 parts for 60-MHz machines and -5 parts for 66-MHz systems. However, Intel's pricing policy for Pentium processors and general market dynamics has largely shifted the consumer PC industry to the processor speeds (i.e. 100, 133, and 166 MHz) that need 66-MHz processor buses. Consequently, any vendors that could not cheaply produce the -5 parts in high volume would be at a significant competitive disadvantage. Simultaneously, and perhaps not coincidentally, renewed interest arose regarding SDRAMs for PCs. During this period, SDRAM-Lite was proposed as a lower cost, reduced functionality SDRAM, and the projected medium-term price premiums for SDRAMs over EDO DRAMs dropped from more than 15% to between 5 and 10%; and projected long-term SDRAM price premiums dropped to 0% (see Section 7.4).

Despite increased interest in SDRAMs, many of the chip-set and PC system vendors continued to support Burst EDO. By the end of 3Q1995, three chip vendors were sampling sets with Burst EDO support; within a month, five more vendors released their designs. In total, 13 chip-set manufacturers have announced that Burst EDO support will be a part of future designs; 9 OEM or motherboard vendors announced their intent to use the device.

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 181
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

The depth of support among chip-set vendors for Burst EDO is not surprising given the nature of their business and the small effort needed to add Burst EDO support to an existing EDO DRAM controller. Burst EDO support quickly became a checklist item for the chip set. Chip-set vendors need to support Burst EDO to provide system vendors with the option of using Burst EDO devices if they do become mainstream.

However, the lack of support from memory vendors is troubling. As of 4Q1995, only two of the five announced suppliers of Burst EDO DRAM are publicly planning to produce the device in volume within a year. Unfortunately, the number of vendors quietly developing Burst EDO designs is unknown. The inability of system companies to guarantee a large, diverse supply of Burst EDO DRAMs at attractive prices will prevent them from adopting the technology, even if the chip-set capability exists and the modules are physically identical.

Nevertheless, Burst EDO has two compelling arguments in its favor. First, Micron is promising no price premium for Burst EDO or EDO, even at the -5 speed grade.[4] Second, systems can be engineered to accept either EDO or Burst EDO DIMMs. If a system vendor is assured of a secure supply of DRAMs, then it can safely use Burst EDO as a system-level product differentiator. If either the supply or the low premium fails to materialize, the system company can quickly back out of its decision since no new engineering or qualification would be needed. Furthermore, if the aftermarket in Burst EDO DIMMs fails to develop quickly because of a general shortage of supply, then the system vendor could expect to garner a higher than average percentage of the memory upgrades for its machines in the field.

By the end of 1H1996 it should be clear whether Burst EDO is going to make it as a mainstream PC device, become another niche DRAM architecture, or fade away entirely.

7.2.6 Strengths and Weaknesses

Strengths

- 66-MHz cycle rate with substantial backward compatibility
- Same packages and modules as EDO and fast-page-mode DRAMs
- Limited controller redesign required

4. This promise is mitigated by the uncertainty regarding the future price premium for -5 EDO DRAMs over -6 or -7 EDO DRAMs.

Samsung Ex. 1097, p. 182
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

- No price premium over EDO devices
- Easy power-up determination of fast-page mode, EDO, Burst EDO
- Metal mask or bond option for fast-page mode, EDO, Burst EDO
- Shared economies of scale

Weaknesses

- Fast 52-ns devices needed for 66-MHz systems
- Significant resistance from some DRAM vendors
- Memory configuration limited at 66 MHz
- Maximum four dual-CAS-load SIMMs or DIMMs
- No improvement in latency over EDO or fast-page-mode DRAMs
- Poor nonburst performance
- Awkward programming of burst type

7.3 Enhanced DRAMs (EDRAMs)

Enhanced Memory Systems of Colorado Springs, Colo., a subsidiary of Ramtron International Corp., has developed a line of **Enhanced DRAMs (EDRAMs)**. The important features of EDRAMs are a very fast DRAM core and an incrementally improved logical interface for faster handling of memory writes and refresh operations. Other aspects of the interface—physical, electrical, and pin-level—are essentially unchanged from the generic DRAMs, thus reducing the complexity of migrating from standard fast-page-mode DRAMs to EDRAMs.

Though initially offered at the 4-Mbit density in ×1 and ×4 organizations, two 512K×8 configurations were sampled in 1995 and Enhanced Memory Systems has plans for 3.3 V and 16-Mbit versions to be introduced in 1996.

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 183
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

7.3.1 Physical, Electrical, and Pin-Level Interfaces

The physical and electrical aspects of the interface are relatively unchanged from the generic DRAMs. The parts require a 5 V supply and use standard CMOS TTL levels with rail-to-rail output levels.

The pin-level interface is modified slightly from the generic interface:

1. $\overline{\text{RAS}}$ is called $\overline{\text{RE}}$ (Row Enable).
2. $\overline{\text{CAS}}$ is called $\overline{\text{CAL}}$ (Column Address Latch).
3. $\overline{\text{OE}}$ is called $\overline{\text{G}}$ (Output Enable).
4. $\overline{\text{W}}$ is called W/$\overline{\text{R}}$ (Write/Read Mode).
5. A second write-control pin, $\overline{\text{WE}}$, is added.
6. A chip select, $\overline{\text{S}}$, is added.
7. A refresh-control pin, $\overline{\text{F}}$, is added.
8. A hit output pin, $\overline{\text{Hit}}$, is added to one of the ×8 parts.
9. Four burst-type and length-control pins are added to the other ×8 part.
10. Unlike generic 1M×4 DRAMs, which have 10-bit address buses, the 1M×4 EDRAMs have the same 11-bit address buses as the 4M×1 EDRAMs.

As a result of the extra pins and the wider address bus, ×1 and ×4 EDRAMs have 22 and 24 signal pins, respectively, instead of the 16 pins of generic 4M×1 DRAMs and the 18 of generic 1M×4 DRAMs. In order to support the higher speed, the usual 2 supply pins are replaced by 6. The end result is that both the ×1 and ×4 EDRAM parts are in 300-mil 28-pin **SOJ packages** (see Figure 7-4). However, the 28-pin SOJ is nominally only 50 mils (1.27 mm) longer than the 26/20-pin SOJ used for the generic counterparts. The extra supply and signal pins therefore translate into at most a 7% board area penalty. By keeping the overhang past the last pins to a minimum, the actual area penalty ranges from 2 to 5% over the actual 20-pin SOJ packages used by generic DRAM vendors. Unfortunately, the board area penalty would be larger with any other package type; with SOJ packages, the EDRAMs fit comfortably on standard **SIMMs**.

Samsung Ex. 1097, p. 184
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

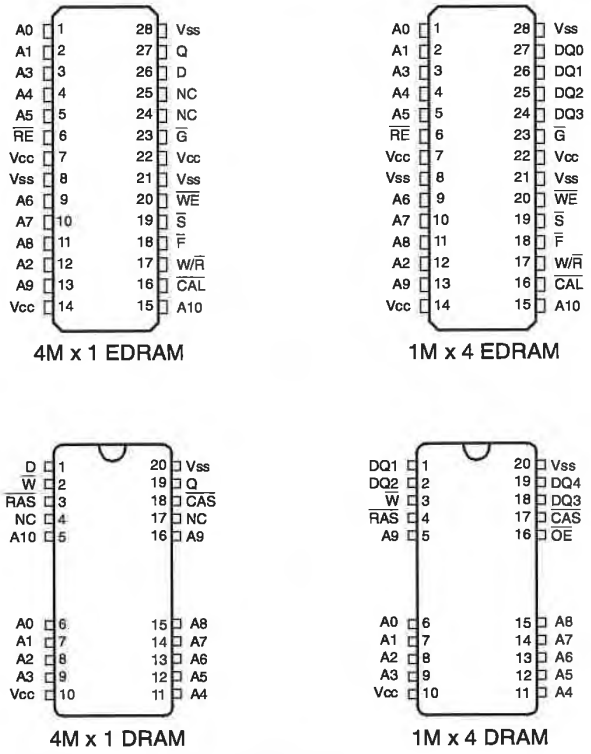


Figure 7-4 EDRAMs are available in 28-pin SOJ packages.

Enhanced Memory Systems offers standard 72-pin ×32 and ×36 EDRAM SIMMs. Both single-sided and double-sided versions are available. Though currently available only in SOJ packages, the ×1 and ×4 EDRAMs will also be available in 44-pin 300-mil TSOP II packages in 1996.

The ×8 EDRAMs will be in 44-pin TSOP II packages instead of the 28-pin TSOP II that is normal for 512K×8 DRAMs. The larger premium is due to the total of 13 power and ground pins needed to support the high clock rates. The same 44-pin package will be used for the 1M×4 and 1M×4 TSOP II versions, as well as for the future 2M×8 EDRAM.

Despite their higher speed, EDRAMs draw no more power than generic DRAMs. Their normal operating power dissipation is 660 mW, compared with standard DRAMs' 500–600 mW.

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 185
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

However, the improvements to the logical interface described in Section 7.3.3 allow EDRAMs to spend less time than generic DRAMs in the highest power states.

Memory controllers and modules have to include the four new control pins. If accommodating these new pins poses a serious problem, some of the new functionality in the EDRAM logical interface can be forgone by tying some of the control signals together. Grounding $\overline{S}$ and tying $\overline{CAL}$, $\overline{F}$, and $\overline{WE}$ together provides most of the benefit of the EDRAM organization with effectively the same set of control pins found on generic DRAMs.

Enhanced Memory Systems sells a line of ×32 and ×36 72-pin SIMMs. These SIMMs bring out all of the new control signals by using a combination of the Presence Detect (PD) and some of the No Connect (NC) pins on the standard SIMM. All of the remaining NC pins in the standard SIMM definition are used for additional power and ground pins. Compatibility with standard SIMMs is jeopardized through the use of the predefined A11 pins as a ground. Additionally, the ×36 SIMMs are different from normal ×36 SIMMs: instead of presenting four 9-bit bytes to the outside world, they present four 8-bit bytes and one 4-bit write-per-bit nibble. This means that if used as parity protection, a parity bit mask must be placed on the data bus when $\overline{RE}$ is lowered to enable the parity bits to be written.

7.3.2 Temporal Interface

The EDRAM's essential feature is its speed. While most manufacturers offer 4-Mbit devices down to 60-ns row access times, EDRAMs come in 30-, 35- and 45-ns speed grades.[5] Unlike most DRAMs, EDRAMs are graded by their **row-cache**[6] access time, which is comparable to the fast-page-mode column address access time. EDRAMs are offered with 12-, 15- and 20-ns row-cache access times; most -60 generic fast-page-mode DRAMs have 30-ns row-cache access times. Table 7-3 summarizes the main EDRAM timing parameters and contrasts them to those of a generic DRAM. Figure 7-5 illustrates the timing parameters.

The EDRAM's temporal interface has two important characteristics. First, all parameters are dramatically faster than the generic's corresponding values: some of the parameters are

5. As this report goes to press, Enhanced Memory Systems is about to introduce its newest speed grade: a -12 part with a 30-ns row access time.

6. Enhanced Memory Systems calls the EDRAM row cache a *page cache* as well as a *row register*.

Samsung Ex. 1097, p. 186
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

	EDRAM					Fast-Page-Mode DRAM			
		-12		-15			-60		
Parameter	Symbol	Min	Max	Min	Max	Symbol	Min	Max	Units
RAS[a] Access Time	t_{RAC}		30		35	t_{RAC}		60	ns
Row Address Setup Time	t_{ASR}	5		5		t_{ASR}	0		ns
Row Address Hold Time	t_{RAH}	1.5		1.5		t_{RAH}	10		ns
RAS Cycle Time	t_C		55		65	t_{RC}		110	ns
Row Precharge Time	t_{RP}		20		25	t_{RP}		40	ns
Column Address Access Time	t_{AC}		12		15	t_{AA}		30	ns
CAS Access Time						t_{CAC}		15	ns
CAS Cycle Time	t_{PC}		12		15	t_{PC}		40	ns
Column Address Setup Time	t_{ASC}	5		5		t_{ASC}	0		ns
Column Address Hold Time	t_{CAH}	0		0		t_{CAH}	15		ns
CAS Pulse Width	t_{CAE}		5		6	t_{CAS}		15	ns
OE Access Time	t_{GQV}		5		5	t_{OEA}		15	ns
OE Turn-Off Time	t_{GQZ}	0	5	0	5	t_{OEZ}	0	15	ns
CAS High to Data Change	t_{CQX}	5		5		t_{OFF}	0		ns
CAS High to Data Valid	t_{CQV}		15		17	t_{CPA}		35	ns
Write Pulse Width	t_{WP}		5		5	t_{WP}		10	ns
Write Pulse to CAS High	t_{CWL}		5		5	t_{CWL}		15	ns

Table 7-3 The 1M×4 EDRAMs are significantly faster than the fastest common 1M×4 fast-page-mode DRAMs. All times are in nanoseconds. (Sources: *DM2202 EDRAM Data Sheet* and *Samsung KM44C1000B Data Sheet*).

a. Standard DRAM signal names are used for the sake of comparison.

reduced by as much as two-thirds. Second, Enhanced Memory Systems' designers have shifted from a 0-setup-time environment to a 0-hold-time environment. This shift saves some time during a burst of column accesses by overlapping each column precharge time with the transmission and setup of the subsequent column address.

An unfortunate reality of the generic DRAM temporal interface is the unavoidable gap between the CAS cycle time listed in the data sheet and the minimum practical page-mode cycle time. With EDRAMs, this gap is narrowed but not altogether eliminated. Table 7-3 lists the CAS cycle time for the -60 generic DRAM as 40 ns. However, making a nontrivial, expandable memory system with a page-mode cycle time less than 50–60 ns

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 187
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

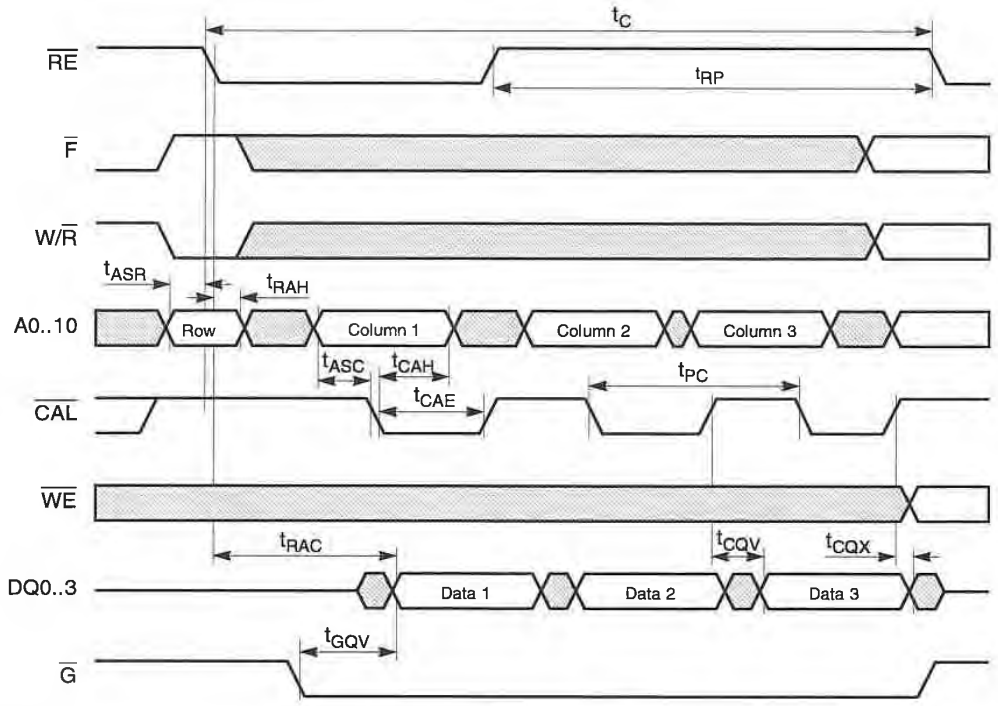


Figure 7-5 EDRAM burst transfers are very similar to generic DRAM burst transfers.

is quite difficult. As specified, -15 EDRAMs cycle at 66 MHz (15-ns cycles), but when operated at this speed, read data is guaranteed to be valid for only 5 ns.[7] To increase the **valid data-out time**, a -15 EDRAM would have to be operated at less than 66 MHz or be used in an interleaved memory system in which the valid data-out timing was determined by the faster $\overline{G}$ timing. Realistically, a bank of -15 EDRAMs can be cycled at about 50 MHz; two-way interleaved memory systems can get to a 100-MHz data transfer rate on a shared data bus—thanks to a very fast 5-ns data output time from $\overline{G}$ (Output Enable). As with all interleaved systems, the consequence is doubled memory system granularity.

7. The valid data-out time of a static-column access is determined by the cycle time, the access time from an address change (t_{AC}=12 ns for the -12 parts and 15 ns for the -15 parts), and the data-invalid time after an address change (t_{AQX}=5 ns for both speed bins).

Samsung Ex. 1097, p. 188
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

Alternatively, -12 EDRAMs could be used. These theoretically cycle at 83 MHz. In a 66-MHz system they provide 8 ns of valid data-out time.

In 1996, a set of EDRAMs with EDO outputs will become available. The -15 EDO parts should be able to operate in systems at 66 MHz without the need for interleaving, while the -12 EDO EDRAMs should function in systems up to 83 MHz.

SRAM performance from a DRAM

As a consequence of these high speeds, an EDRAM memory system responds as quickly as an SRAM-based secondary cache most of the time. Table 7-4 shows the burst fill rates for several hypothetical PC memory systems. The first of the four numbers in each table location is the number of cycles needed to return the first double-word (32 bits) or quad-word (64 bits) from memory; the next three numbers represent the additional cycles needed to return each subsequent double- or quad-word. The total number of cycles for the whole burst refill is given in parentheses. These numbers include controller overhead and processor-related delays (such as the setup time and the initial delay in presenting the address to be fetched).

For these small memory systems, the EDRAM's raw speed matches the speed of an SRAM cache when they both hit and substantially betters the speed of a fast-page-mode DRAM when they both miss. The system with an SRAM cache can show higher performance than the EDRAM system if the SRAM secondary cache miss rate is substantially lower than the EDRAM's row-cache miss rate. In practice, the EDRAM memory system tends to outperform the combination of an SRAM secondary cache and fast-page-mode main memory. On the cost front, the savings from eliminating the SRAM are somewhat offset by EDRAM's cost per bit premium. For small- to medium-size memory systems, the EDRAM solution is both faster and cheaper than the traditional combination of an SRAM and a generic DRAM.

7.3.3 Logical Interface

The logical interface of the first-generation EDRAMs is very similar to that of the standard fast-page-mode DRAMs, with the advantage that EDRAMs can be used as static-column-mode DRAMs as well. This dual functionality is accomplished by slightly changing the semantics of the Column Address Latch ($\overline{\text{CAL}}$). The other changes are in the functionality of the row cache and in the handling of writes and refreshes.

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 189
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

DRAM Type	System Speed	Burst Read		Burst Write	
		Cache Hit	Cache Miss	Cache Hit	Cache Miss[a]
-12 EDRAM	33 MHz	2:1:1:1 (5)	3:1:1:1 (6)	2:1:1:1 (5)	3:1:1:1 (6)
	50 MHz	2:1:1:1 (5)	4:1:1:1 (7)	2:1:1:1 (5)	3:1:1:1 (6)
	66 MHz	3:1:1:1 (6)	5:1:1:1 (8)	2:1:1:1 (5)	3:1:1:1 (6)
SRAM Cache and Fast-Page-Mode DRAM	33 MHz	2:1:1:1 (5)	4:2:2:2 (10)	2:1:1:1 (5)	4:2:2:2 (10)
	50 MHz	2:1:1:1 (5)	6:2:2:2 (12)	2:1:1:1 (5)	6:2:2:2 (12)
	66 MHz	3:1:1:1 (6)	8:3:3:3 (17)	2:1:1:1 (5)	8:3:3:3 (17)

Table 7-4 An EDRAM-based PC memory system can respond as quickly as an SRAM secondary cache backed by a fast-page-mode DRAM.

a. A cache miss in the case of a write-through row cache; a write-back in the case of an SRAM cache.

In fast-page-mode DRAMs, the row cache is loaded whenever a row access occurs with $\overline{RAS}$ falling. When $\overline{RAS}$ is raised again, its contents are lost: fast column accesses to the cache can be made only while $\overline{RAS}$ is held low. In EDRAMs, the row cache is loaded with the read row access initiated by the falling edge of the row enable ($\overline{RE}$); the row cache remains available until the start of the next read row access. In other words, reads to the row cache can continue even after $\overline{RE}$ is raised and the DRAM core has been precharged.

The EDRAM's row cache has the same size and organization as the fast-page-mode DRAM's: the cache consists of a single block containing all the data logically fetched from the DRAM array during a row access. For the 4M×1 EDRAM, the row cache is a single, direct-mapped block of 2,048 bits. For the 1M×4 EDRAMs, the block is structured as 512 words of 4 bits each. For the 512K×8 EDRAMs, the block is structured as 256 words of 8 bits each. This means that in the 1M×4 and 512K×8 parts, the row address is 11 bits wide, while the column address is 9 bits wide. In contrast, all generic 1M×4 parts have symmetrical 10-bit row and column addresses.

On-chip comparator checks for row-cache hits

Figure 7-6 shows a block diagram of the ×4 part. In addition to the usual elements, a latch holds the row address of the last row access, and a comparator checks whether the current row address is the same as the previous one. When the comparator detects a match, the data can be read directly out of the row cache. The block diagram also shows that the read path from the row cache to the output port is independent of the write path from the I/O block to the sense amps. Separation of these

Samsung Ex. 1097, p. 190
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

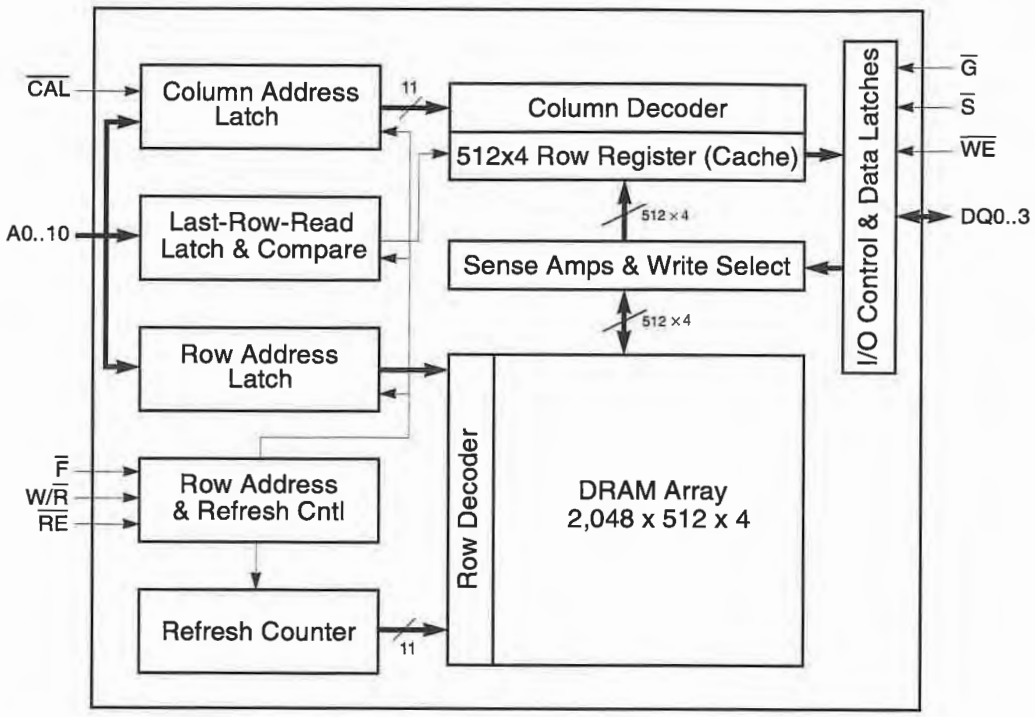


Figure 7-6 The EDRAM block diagram shows the row cache and the separate read and write paths to the DRAM array.

paths facilitates writes into the DRAM array that do not interfere with the row cache.

Burst reads are overlapped with row precharge and refresh

Figure 7-5, on page 189, shows a row access followed by a sequence of row-cache read operations. The timing diagram is the same regardless of whether the row address is to the same row as the last row access. If the last-row-read comparator indicates a match, then the row access is not performed, and the data is available from the row cache immediately. Otherwise, the data is not available until a row access is completed, a full 35 ns after the falling edge of $\overline{RE}$.

Figure 7-5 also depicts column accesses continuing after $\overline{RE}$ is raised; EDRAMs, unlike generic DRAMs, allow row-cache accesses to overlap with row precharges. For burst transfers of

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 191
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

three or more words, if $\overline{RE}$ is raised as soon as possible, the row precharge time is completely hidden underneath the cache accesses, and the next row access can begin as soon as the previous burst access is complete.

Like row precharges, refresh operations can be hidden underneath cache accesses. If the refresh mode pin, $\overline{F}$, is asserted when $\overline{RE}$ drops, then instead of a cache-loading DRAM read, a cache-preserving refresh cycle begins. Unlike the hidden refresh cycles of generic DRAMs, these refresh operations are truly hidden: they can be entirely hidden underneath row-cache read operations. As in the generic case, the refresh address is taken from an on-chip refresh counter. Of course, refreshes can also be done by $\overline{RE}$-only refresh cycles, in which case the address of the row to be refreshed is supplied by the memory controller. $\overline{RE}$-only refresh cycles are performed like write operations with no data transferred in order to avoid changes to the row cache. One way or another, each row must be refreshed every 64 ms—four times less frequently than refresh operations in generic DRAMs. This greatly reduces standby power consumption.

Hidden precharge leads to lower average power

Generic DRAMs draw almost as much power doing column accesses as row accesses: since $\overline{RAS}$ is low during column accesses, substantial power is being burned in the DRAM array even though no row accesses are being performed. EDRAMs draw substantially less power performing row-cache-hit accesses than generic DRAMs do because the DRAM array can be quickly precharged and brought into its lower power state. Averaged over a collection of row and column accesses, EDRAM's power consumption is substantially less than the generic DRAMs' even though their peak power consumptions are comparable.

The column address latch is a flow-through latch that captures the column address on the falling edge of $\overline{CAL}$. However, the read transaction proceeds normally even if the memory controller does not lower $\overline{CAL}$. If $\overline{CAL}$ is held high, the EDRAM behaves much like a static-column-mode DRAM on reads, with data flowing out after changes in the column address. The disadvantage of this mode is that the column address must be held until the row-cache access is complete, sometimes causing a slower system cycle time.

Fast writes can bypass the row cache

Figure 7-7 shows a write operation to an address currently in the cache. The memory controller initiates a write by holding W/$\overline{R}$ high as $\overline{RE}$ is lowered. Thereafter, a write is performed by strobing $\overline{WE}$ while $\overline{CAL}$ is low. If the row address is the same as

Samsung Ex. 1097, p. 192
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

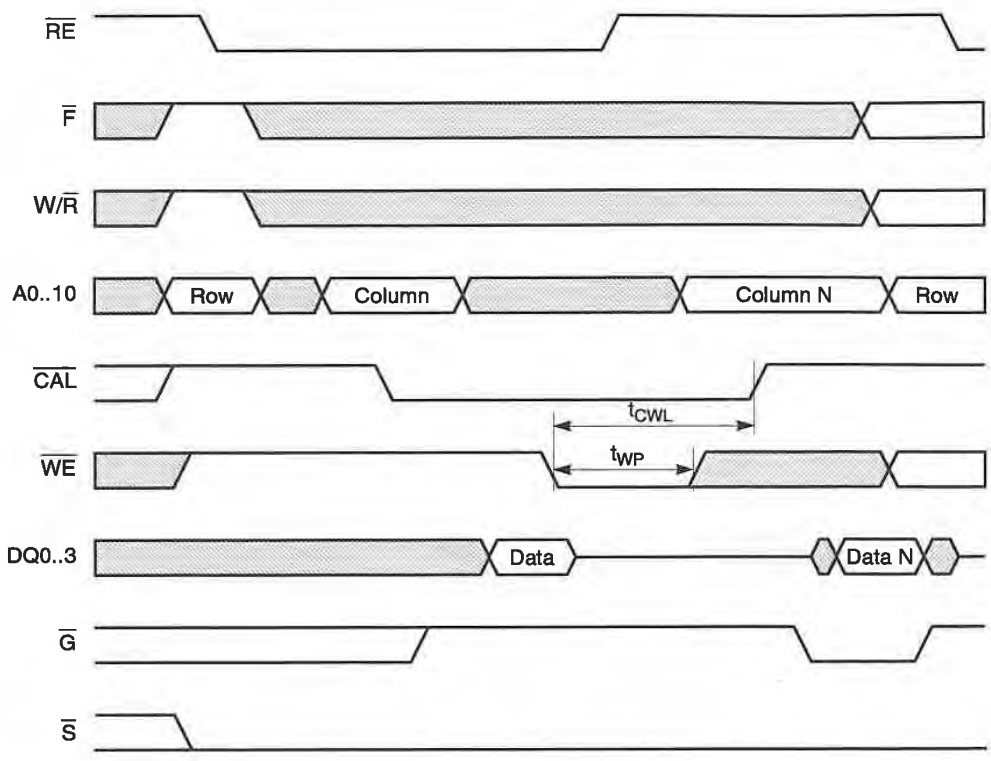


Figure 7-7 This figure shows an EDRAM with both a row-cache write and a row-cache read being completed underneath a DRAM array precharge.

the row address of the last row read, then the write data goes into both the row cache and the memory array. If the row addresses are different, then the write data goes into the memory array only, leaving the row cache untouched.

This improvement to write operations is one of EDRAM's significant improvements to the logical interface. Not only are writes to the DRAM array retired very quickly, but they do not necessarily disturb the row cache. In systems with a write-through primary cache, such as the Intel 486, writes to the memory system significantly outnumber reads, so retiring them quickly is very important. Furthermore, if the row cache has a row of instructions or read data, then writes can be retired without

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 193
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

causing subsequent row-cache read misses. This feature is also useful if the primary cache is a write-back cache. In this case, writes to main memory result from the replacement of dirty blocks in the primary cache. These writes are almost always to different row addresses than the memory reads needed to satisfy the original primary cache miss. Handling the write-backs without disturbing the row-cache contents significantly improves the row-cache miss ratio. Systems lacking a primary cache also benefit from EDRAM's fast writes around the cache.

In all cases, the address bus and $\overline{CAL}$ are freed up as soon as $\overline{CAL}$ rises and is held high for the minimum column precharge time (5 ns). This means that new operations, whether writes to the same row while $\overline{RE}$ is held low or reads to the row cache, can begin immediately, before the memory-array write is actually completed. This early release of the interface in the case of writes further reduces the effective memory-write time. Writes are effectively dispatched in only 15 ns, except in the cases of a write followed by a read or a write to a row that is not in the row cache.

Other write-oriented features include read-modify-write cycles and a version of the 1M×4 EDRAM with write-per-bit capabilities. This DM2212 part is appropriate for holding the parity bits in a ×36 memory system or SIMM. All three narrow 4-Mbit EDRAMs (4M×1, 1M×4, and 1M×4 with write-per-bit) are derived from the same mask set.

The main difference between the EDRAM's and generic DRAM's write interfaces is the separation of the standard $\overline{W}$ into two pins: $\overline{WE}$ and W/$\overline{R}$. Most of the functionality of $\overline{W}$ is carried by $\overline{WE}$. The main purpose of W/$\overline{R}$, and $\overline{F}$ as well, is to indicate that the current row access should not corrupt the current row-cache contents.

Memory controller considerations

The EDRAM interface preserves enough of the generic DRAM interface to facilitate a single memory controller design able to control both EDRAMs and generic fast-page-mode DRAMs. Such memory controllers facilitate pairs of products with different prices and performance levels from a single motherboard. Having a controller that can control either type of DRAM also affords a degree of security of supply: a system vendor can ship a lower performance system with generic DRAMs if EDRAMs are in short supply. National Semiconductor and V3 have announced a line of embedded system memory controllers for the Intel 960 series of CPUs. These controllers can control any combination of generic DRAM and EDRAM SIMMs.

Samsung Ex. 1097, p. 194
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

To take full advantage of all the performance enhancements in the EDRAM design, the memory controller has to keep a duplicate copy of the EDRAM's last-row-read latch and its associated comparator. Only by duplicating this latch can the memory controller understand how quickly a write can happen, and whether a row access is needed to satisfy a CPU read request.

Wider EDRAMs to add more row cache and burst transfers

Sampling of two 512K×8 EDRAMs began in 1995. The first of these EDRAMs maintains continuity with the 4M×1 and 1M×4 parts by maintaining all the existing signals and timing relationships. It significantly improves the row cache by exposing all of the sense amps in the device as row-cache bits. The size of the row cache is increased by logically dividing the one existing bank into four banks that have one-fourth the depth. The result is a very well-organized four-row 8-Kbit row cache. To manage these four row registers, the design has four address comparators and four last-row-read latches (see Figure 7-8). Address bits A_8 and A_9 specify the bank to be used by both row and column operations. This organization allows all four banks to be "open" simultaneously: column access reads can occur from any of the row registers, and background memory array writes can occur to any bank without disturbing the row cache. In general, this bank and row-cache organization greatly improves the effectiveness of the row cache, especially in small memory systems with only a few **ranks**.

To relieve the memory controller from the need to keep a copy of all last-row-read registers, a cache-hit output is provided. However, the timing of this signal is such that it is difficult to use at full speed: the controller is best off duplicating the comparators and latches locally.

Another important improvement in the ×8 devices is the addition of a new control signal, QLE. When QLE is held low, the output buffers behave like those in the 4M×1 and 1M×4 devices. However, when QLE is held high, a latch in the output buffer is enabled and the EDRAM behaves like an EDO DRAM (see Section 6.3). Extending the data output time allows 66-MHz operation in a system with the -15 parts or 83-MHz operation with the -12 parts.

It is interesting that the designers have chosen to bring QLE to the outside world even though in most systems it is unlikely to change state. Normally, the selection between fast-page mode and EDO happens by bonding out different versions of $\overline{\text{CAS}}$ and creating two different part numbers. Providing a separate pin to control the EDO functionality allows EDO behavior for static-

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 195
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

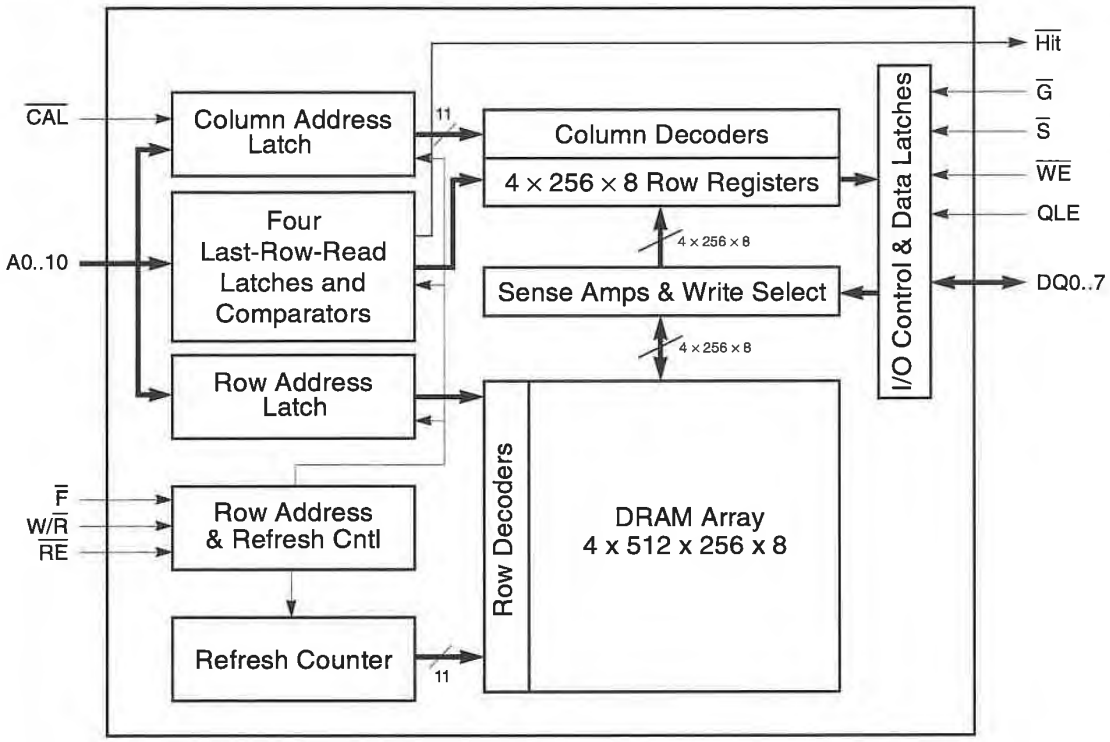


Figure 7-8 The ×8 EDRAMs have four banks internally.

column-mode accesses as well as for page-mode accesses. When used as a static-column-mode device, addresses change arbitrarily and $\overline{CAL}$ is not toggled. In this case, QLE can be used to latch the output data and to eliminate the temporal penalty normally due to the large difference between the access time from address (t_{AC}) and data-invalid time from address (t_{AQX}).

Explicit burst control signals

The second ×8 version adds four dedicated pins to control burst transfers. In conjunction with an on-chip burst counter, these pins provide bursts of 2, 4, 8, or 256 bytes in either linear or sub-block (Intel) ordering. Specifying the size on a reference-by-reference basis simplifies changing the burst size between CPU-, I/O-, and graphics-originated references. In contrast, SDRAMs require reprogramming of the command register to

Samsung Ex. 1097, p. 196
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

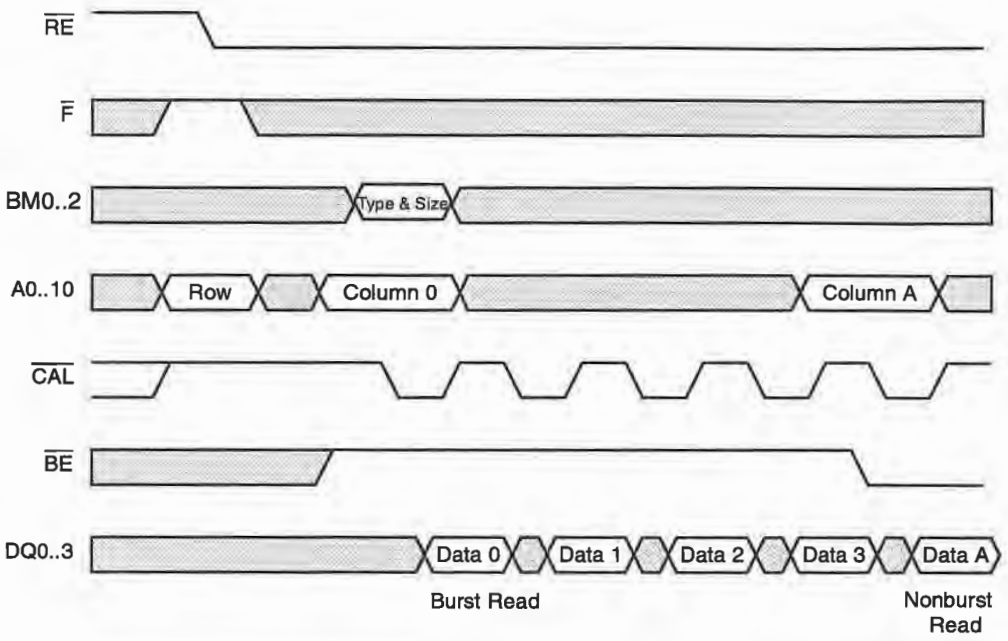


Figure 7-9 The EDRAM burst mode is different than that of either Burst EDO DRAMs or SDRAMs.

change the burst length, while Burst EDO DRAMs only support a single burst length.

The EDRAM's burst mode differs greatly from any other burst mode. To initiate a burst, the Burst Enable signal, BE, is raised prior to the falling edge of CAL. On that first falling edge of CAL, the start address and burst type and size, as indicated by the three Burst Mode pins, BM0..2, are latched by the EDRAM. Subsequent rising edges of CAL for which BE is asserted cause subsequent addresses to be generated internally, and the corresponding data to be accessed from the row cache. Bursts can optionally involve the DRAM array, depending on whether RE is asserted (see Figure 7-9).

As with first-generation EDRAMs, these second-generation EDRAMs will facilitate reads and writes to the row cache in parallel with DRAM writes. With four blocks in the row cache, the opportunity for parallelism will be greater, especially for video and DSP applications with multiple simultaneous source

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 197
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

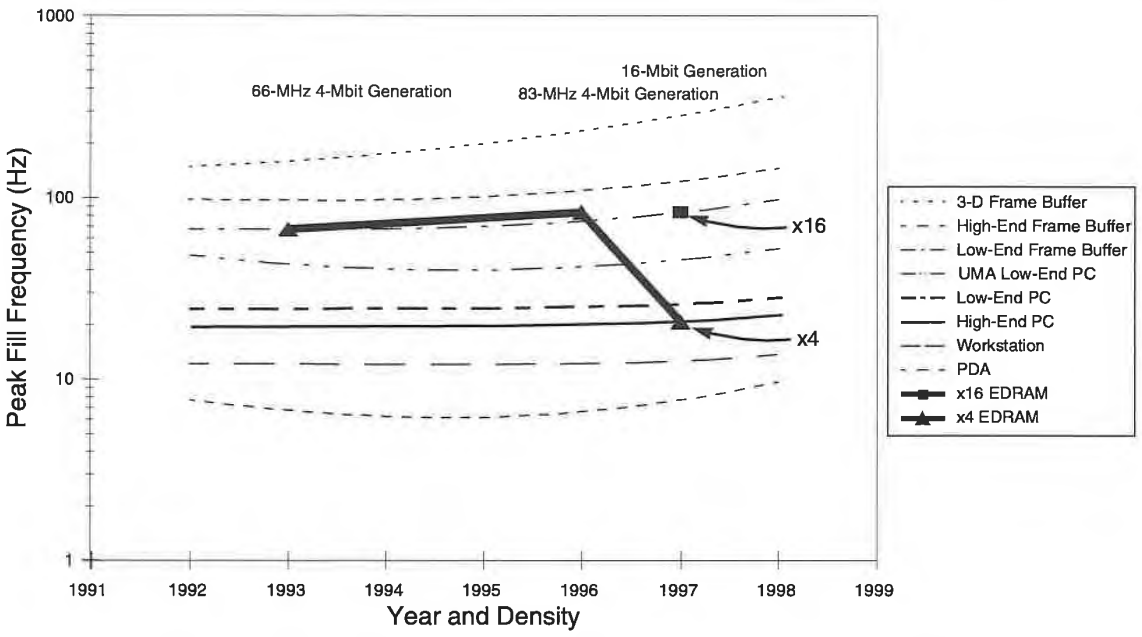


Figure 7-10 The EDRAMs' peak fill frequencies support their use in UMA and graphics systems.

and destination streams to different areas in memory. Both the number of banks and the number of row-cache bits per DRAM bit are quadrupled. Together, these features further improve the effectiveness of EDRAM memory systems as replacements for SRAM-based secondary caches.

7.3.4 EDRAM Fill Frequencies

The peak operating frequency of the current -15 EDRAMs is 66 MHZ. Consequently, the peak bandwidth per DRAM is 33 Mbytes/s for the ×4 EDRAMs, and 66 Mbytes/s for the ×8 devices. With a 1/2-Mbyte size (4 Mbits), these bandwidths translate to peak fill frequencies of 66 Hz for the ×4 devices and 132 Hz for the ×8 devices. The corresponding fill frequencies for the -12 parts are 83 Hz and 166 MHz, respectively. These fill frequencies are adequate for all but the most demanding graphics applications (see Figure 7-10).

Samsung Ex. 1097, p. 198
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

Most of the EDRAM features impact the sustained and sustainable fill frequencies but not the peak fill frequencies. Changes to the logical interface improve the effectiveness of the row cache and significantly increase the percentage of references that can be handled with column accesses instead of row accesses. Just as significant for the calculation of sustainable fill frequencies is the short row-access time. The shorter column access time has less impact on the sustainable bandwidth and fill frequency because with 32-byte transfer in four chunks most of the total access time is spent doing the data transfer, essentially regardless of which DRAM architecture is being used.

Given EDRAM's single-cycle column access time, the maximum sustainable bandwidth is equal to the peak bandwidth. The worst case sustainable bandwidth—assuming continuous row accesses and bursts of four—is 25 Mbytes/s for the -15 ×4 device and 50 Mbytes/s for the -15 ×8 part.[8] Consequently, the ranges of sustainable fill frequencies are 55–66 Hz for the -15 ×4 and 100–132 Hz for the -15 ×8 devices. For the -12 parts, worst case bandwidths are 30.3 Mbytes/s and 60.6 Mbytes/s for the ×4 and ×8 widths respectively. Therefore, this fastest speed grade develops sustainable fill frequencies of 60.6–83 Hz for the ×4 parts and 121.2–166 Hz for the ×8 parts. These are the best sustainable fill frequencies at this density and these widths. Even the worst case sustainable fill frequencies are very high. This implies very consistent performance across the full range of reference patterns. Further comparative analysis of sustainable fill frequencies is given in Chapter 10.

All fill frequencies for ×1 EDRAMs are one-fourth those of the ×4 devices, or 13.75–16.5 Hz for the -15 parts, and 15.2–20.8 Hz for the -12 parts.

7.3.5 Miscellanea

Corporate relationships

Ramtron International has designed high-speed DRAMs for NMB Semiconductor since the 256-Kbit generation. Though Ramtron licensed the products exclusively to NMB, Ramtron retained ownership of these products and the processes that largely made them possible. In mid 1993, NMB sold its semiconductor operations to NPNX, a subsidiary of Nippon Steel.

8. With bursts of four, the precharge time is completely hidden underneath the transfers, facilitating a total access time of 35 ns + 3 × 15 ns = 80 ns for the -15 parts. The total is 66 ns for the -12 parts.

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 199
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

In mid 1995, Ramtron International spun off the EDRAM business into a wholly owned subsidiary called Enhanced Memory Systems. Ramtron proper now deals solely with its nonvolatile Ferroelectric RAMs (FRAMs). Enhanced Memory Systems continues to use NPNX as its only source of EDRAMs. That makes Enhanced Memory Systems essentially the only fabless DRAM vendor.

NMB and Ramtron jointly created United Memories, also of Colorado Springs, Colo., as the design center for EDRAMs, high-speed standard DRAMs, and the processes used to fabricate them. In 1995, Enhanced Memory Systems sold its stake in United Memories to NPNX.

In May 1995, IBM and Ramtron announced that IBM had signed a nonexclusive licensing agreement. Under this agreement, IBM will provide EDRAMs to Enhanced Memory Systems as well as optionally sell EDRAMs under the IBM label. Enhanced will begin offering IBM-produced plug-compatible 5 V parts 4M×1 and 1M×4 parts in 1Q1996. IBM will also produce parts with EDO outputs, four banks, a 3.3 V supply, and low-power (i.e., self-refresh) consumption. The EDO and four-bank parts will have a similar interface to the ×8 parts. The IBM parts will be the first ×1 and ×4 EDRAMs in TSOP packages.

Process and physical layout

Currently, EDRAMs are fabricated in the same 0.6-μm CMOS process used for NPNX's high-speed DRAMs. The EDRAMs' speed is quite impressive given that most 4-Mbit DRAMs are built with comparable geometries. The speed derives from a number of factors. Two important factors are field-shield isolation and a low-resistance interconnect to help speed the settling and precharging of the bit lines. IBM's EDRAMs will be made on IBM's standard 4-Mbit technology. This means that IBM's EDRAMs will enjoy the same learning curve and cost structure as their high-volume generic DRAMs.

Enhanced Memory Systems hopes that IBM will also provide them with a route to 16-Mbit devices in 1996, as NPNX currently does not have a high-volume 16-Mbit process.

The EDRAM is physically laid out as 16 blocks of 256 Kbits each. Each block includes 1K sense amps on one end and a 512-bit row register on the other. Each sense amplifier has 128 storage cells on each of its four bit lines. Of the total of 8K row-register bits on the chip, either 2K or 8K act as the row cache at any one time, depending on the device organization. Because the row registers are built into naturally occurring column-

Samsung Ex. 1097, p. 200
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

select circuitry, the area penalty for the extra functionality is relatively small. Enhanced Memory Systems routinely calls the row registers an SRAM cache. However, it is largely a marketing term because no discrete SRAM block, with bit lines and sense amplifiers, exists.

SRAM L2 replacement

As suggested earlier, Enhanced Memory Systems markets EDRAMs primarily as a cost-effective alternative to the combination of SRAM-based secondary caches and fast-page-mode or EDO DRAMs. Simulations and a demonstration prototype show that a PC with an Intel 486DX2-66 processor and an EDRAM memory system consistently outperforms a comparable machine with a 256-Kbyte cache and a fast-page-mode memory system. Also, a DX4100 motherboard without any SRAM L2 cache performs comparably to a 100-MHz Pentium® with a 256-Kbyte L2 cache and fast-page-mode DRAM.

A major question is whether the public will buy a PC with a special high-performance DRAM instead of a large secondary cache. Naive end users and unsophisticated salespeople may compare a secondary cache size with a row-cache size and assume that more is better. This danger is not present in the embedded-system marketplace, which is simultaneously less averse to innovation and more technically driven. It is in this domain that EDRAMs have met with the greatest success.

The performance advantage and attractiveness of EDRAMs should increase as processor bus speeds reach and move beyond 66 MHz. At these speeds, fast-page-mode and even EDO DRAMs mandate a large number of wait states on every secondary cache miss.

EDRAMs are also attractive in the video arena as a replacement for VRAMs. EDRAMs' high bandwidth and effective cache structure facilitate performing both screen updates and display scanning through a single port. In fact, several of the announced EDRAM design wins are graphics- and virtual reality-oriented.

Diverse and growing design wins

The most significant design win for EDRAMs is a DX4-100 motherboard produced by Ocean Information Systems. Other announced and unannounced design wins are graphics accelerators, several communications bridges/routers, and other embedded systems. For most design wins, EDRAMs enable one or more of the following:

- Replacement of a secondary cache and generic DRAM

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 201
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

- High performance with moderate power consumption
- Combining a video memory and a main memory in a UMA environment
- Replacement of an expensive SRAM main memory

In comparison to the other DRAM alternatives and system architectures, EDRAMs offer either:

- Lower system cost at comparable or slightly better performance
- Higher performance at comparable or greater system cost

As of 4Q1995, there were over 140 EDRAM design wins and over 1.4 million units have been shipped. Though tiny by generic DRAM standards, this number of design wins is significant given Enhanced Memory Systems' small size and limited market presence. As of May 1995, Enhanced Memory Systems had over $39 million in orders on its books for 1995: a threefold increase over 1994 sales. NPNX is severely constraining the supply of EDRAMs and the lead time for new orders has reached 16 weeks. Enhanced Memory Systems is counting on IBM coming on line in early 1996 to dramatically increase supply and ease the current critical shortage. Until that time, EDRAMs will remain hard to get in substantial volumes.

EDRAM market barriers

Enhanced Memory Systems's biggest hurdles are convincing potential clients that 1) it is a significant DRAM vendor; 2) its leverage with IBM and Nippon Steel is sufficient to ensure a solid supply of EDRAMs in the face of a continuing industry-wide shortage of DRAMs; 3) the source of supply is sufficiently diversified to survive a natural disaster or yield crash; and 4) the price premium over generic DRAMs—currently 25%—will continue to provide system-level cost advantages. These business issues will ultimately impact the success of EDRAMs more than the device's significant technical merits. This is especially true in light of the delays and delivery problems that Enhanced Memory Systems' customers have faced over the past year—problems that result from Enhanced Memory System's single-source relationship with NPNX.

7.3.6 Future Developments

If IBM extends its licensing agreement with Enhanced Memory Systems, IBM should become a source for high-speed 16-Mbit

Samsung Ex. 1097, p. 202
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

EDRAMs in 1996. Enhanced Memory Systems plans to offer 2M×8 and 1M×16 devices, each with 8 banks and a total of 32 Kbits of row cache. The same -12 speed bin is expected to be the mainstay part. EDO, burst-transfer modes, and TSOP packaging will carry forward to the new generation.

Given no change in the column-cycle time, these devices' fill frequencies will be one-fourth of the fill frequencies of the corresponding 4-Mbit devices. However, given the introduction of a ×16 device, the sustained and sustainable fill frequencies will be slightly more than one-fourth to one-half of current levels (see Figure 7-10).

The peak fill frequencies for the -12 2M×8 and 1M×16 should be 41.6 Hz and 83 Hz, respectively.

7.3.7 Strengths and Weaknesses

Strengths

- Very fast row and column accesses
  - 30-ns row accesses; 12-ns cache accesses
  - Precharge time hidden underneath burst transfers
- Efficient write interface
  - Writes do not necessarily disturb the row cache
  - Interface freed up before DRAM write is complete
- Refresh cycles hidden under column accesses
- Lower average power than generic DRAMs

Weaknesses

- TTL signal levels
- Asynchronous interface
- Late with 16-Mbit devices
- Proprietary part
  - No second source until 1996

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 203
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

7.4 Synchronous DRAMs (SDRAMs)

Recognizing the limitations of the generic DRAM interface, and desiring to develop a higher performance standard interface for a new generation of commodity DRAMs, a number of DRAM vendors and users joined together under the auspices of the EIA/JEDEC JC42.3 DRAM Standards Committee to develop a synchronous DRAM interface. Although most of the evolutionary and revolutionary DRAMs are actually synchronous—that is, they include a free running clock input—only DRAMs that conform or nearly conform to the proposed standard are referred to as **Synchronous DRAMs (SDRAMs)**.

The SDRAM interface retains a multiplexed address bus and adds a high-speed clock to 1) synchronize the flow of addresses, data, and control on and off the DRAM and 2) facilitate the pipelining of operations. A two-bank internal organization provides for maximum utilization of the high-speed interface. Consequently, the main changes are in the temporal and logical interfaces.

7.4.1 Physical, Pin-Level, and Temporal Interfaces

Figure 7-11 shows the package and pin definitions for a 16-Mbit ×8 SDRAM that conforms to the JEDEC standard. Also included in the standard are specifications for 16-Mbit ×4, ×9, ×16, and ×18 parts, as well as 4-Mbit ×16 devices. The ×16 parts of both generations and the 16-Mbit ×18 parts are put into 50-pin TSOP packages by reducing the number of supply pins per data I/O pin from 1:1 to 1:2. In the case of the ×18 parts, the ratio is even higher. The ×4 SDRAMs are generally placed in the same 400-mil 44-pin TSOP II packages as the ×8 and ×9 parts, but with four of the data I/O pins changed to No-Connects (NC).

At the 64-Mbit level, the predominant package will be a 54-pin TSOP II. This package is being used for the ×4, ×8, and ×16 devices. A new ×32 organization will be in an 84-pin TSOP II. Although most vendors will begin by using packages that are 500-mil wide, package widths will ultimately stabilize at 400 mil.

Samsung Ex. 1097, p. 204
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

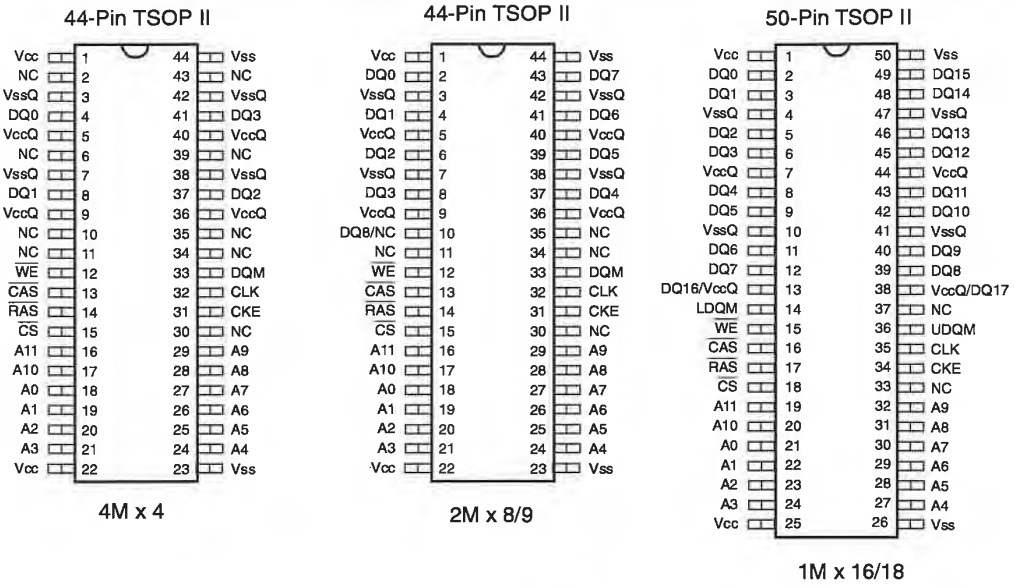


Figure 7-11 Synchronous DRAMs pinouts have been defined for a variety of device widths.

Two sets of DIMMs defined

JEDEC has also defined two sets of 8-byte DIMMs for use with SDRAMs. The first set of DIMMs have 200 pins while the second set have 168 pins. The former set have more power and ground pins and so should be usable at a higher clock frequency, while the second set are physically compatible with the 168-pin 8-byte DIMMs defined for fast-page-mode and EDO DRAMs. Both sets of DIMMs come in buffered and unbuffered versions. In the buffered versions, the free-running clock is regenerated on the DIMM by a phase-locked loop (PLL) and an on-DIMM register captures all the address and control signals and drives them to the individual SDRAMs. Both varieties are defined in ×64, ×72, and ×80 versions. However, the ×72 configuration only

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 205
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

supports word writes—not byte writes. Unlike some ×72 fast-page-mode and EDO DIMMs, the ×72 SDRAM DIMM is not logically eight 9-bit bytes wide.

The 200-pin DIMM employs a conventional parallel presence-detect that encodes:

- the depth and width of the DIMM
- the data sheet minimum cycle time of the SDRAMs used
- whether the DIMM is buffered and latched, or unbuffered
- whether the SDRAMs depend on the 2n-rule[9]
- whether the SDRAMs support early precharge commands after a read burst transfer[1]
- whether low-power SDRAMs are used

The 168-pin DIMMs encode the same information, but do so with a more pin-efficient serial presence-detect scheme. The 200-pin DIMMs are physically keyed to differentiate between current 3.3 V LVTTL parts and future parts with different supply voltages and electrical interfaces.

Of all the possible combinations provided for, the unbuffered ×64 168-pin DIMMs are likely to be most common.

SDRAM features

In addition to the standard $\overline{\text{RAS}}$, $\overline{\text{CAS}}$, and $\overline{\text{W}}$ pins, SDRAMs include a clock (CLK), a clock enable (CKE), a Chip Select ($\overline{\text{CS}}$), and an Input Mask/Output Enable (DQM). The ×16 and ×18 parts have upper and lower output enables to control the two bytes separately. In most generic DRAMs, all the address pins are used for both row and column addresses. SDRAMs are different: some of the upper address bits have specific meanings. For example, in the 16-Mbit ×8 parts, A11 is used exclusively as a bank-select. These semantic definitions are detailed in Section 7.4.2.

The new temporal interface specifies that all address, data, and control inputs are latched on the rising edge of the clock. Outputs change after the rising edge of the clock. Figure 7-12 shows the nine parameters that fully specify the detailed timing aspect of the temporal interface. Most other temporal parameters can be effectively expressed in terms of cycles instead of

9. See Section 7.4.2 for explanations of these terms.

Samsung Ex. 1097, p. 206
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

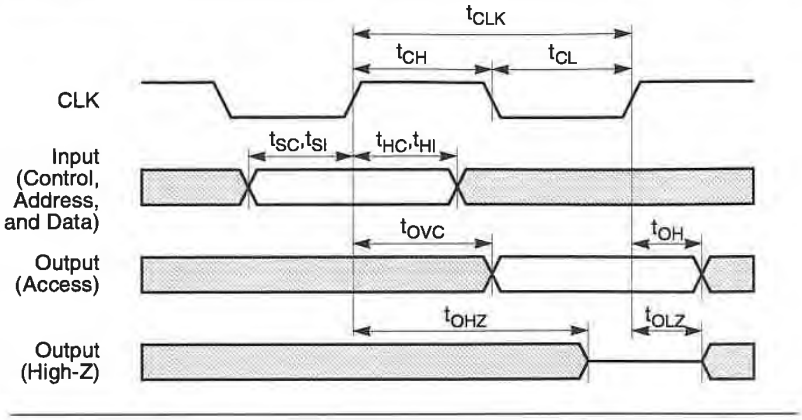


Figure 7-12 The SDRAM's temporal interface, unlike the generic DRAM's, is defined by a small number of parameters.

Parameter	Symbol	-10		-12		-15		Units
		Min	Max	Min	Max	Min	Max	
Clock Period	t_{CLK}	10		12		15		ns
Clock High Time	t_{CH}	3–4		3–4		4–5		ns
Clock Low Time	t_{CL}	3–4		3–4		4–5		ns
Input Setup Time	t_{SI}	2–3		2–3		3–3.5		ns
Input Hold Time	t_{HI}	1–2		1.5–3		1.5–4		ns
Output Valid from Clock, Col. Latency = 1	t_{OVC1}		25–30		30–35		35–40	ns
Output Valid from Clock, Col. Latency = 2	t_{OVC2}		10–15		13–16		16–19	ns
Output Valid from Clock, Col. Latency = 3	t_{OVC3}		8–9		9–11		12–14	ns
Output in Low-Z	t_{OLZ}	0		0		0		ns
Output in High-Z	t_{OHZ}		8–10		8–10		9–11	ns
Output Hold Time	t_{OH}	3–4		3–4		3–4		ns

Table 7-5 The detailed timing parameters that specify an SDRAM behavior are similar to those of generic DRAMs. (Sources: Vendors' preliminary and advanced data sheets).

nanoseconds. Table 7-5 shows typical values of these parameters for several common speed grades.

Timing parameters becoming more consistent

One early criticism of the SDRAM standard was that it did not specify any of the temporal parameters. Consequently, when the DRAM vendors produced their first designs, the setup and hold

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 207
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

parameters, among others, varied widely. In the latest set of data sheets, however, there is much more consistency across vendors. For example, most vendors are working towards setup times of 3–4 ns and hold times of 1–2 ns. Vendors still opting for more hold time and less setup time are becoming the minority. Similarly, the time from rising clock to valid data-out have become much more consistent across vendors.

In spite of this progress, the detailed timing parameters are still much less consistent than those of generic DRAMs. And complete interchangeability of SDRAMs from the perspective of the temporal interface is far from assured. Furthermore, substantial differences still exist in the specification of the parts and even in the names used for many of the secondary timing parameters.

At-speed operation difficult

The typical timing parameters leave little room for systems designers to maneuver. For example, when operated at 100 MHz, the -10 parts nominally present valid data for 4–6 ns.[10] However, if the DRAM clock is phase-aligned with the memory controller clock, then the data has an inadequate 1–2 ns to propagate from the DRAM into the memory controller. Skewing the DRAM clock earlier than the memory controller clock helps with this problem but effectively increases the SDRAM setup time. Of course, clock skew throughout the memory system subtracts directly from the valid data-out time, making it even harder to operate the SDRAMs at speed.

In medium to large memory systems constructed with DIMMs, DRAM timing parameters are too tight to facilitate at-speed operation: -10 parts will not facilitate 100-MHz operation, -12 parts will not facilitate 83-MHz operation, and so on. Generally, for the largest memory systems, speed grades that are at least one bin faster than would nominally be expected are required to ensure reliable operation. For example, most vendors' -12 parts are needed to facilitate 66-MHz operation in typical PC configurations.

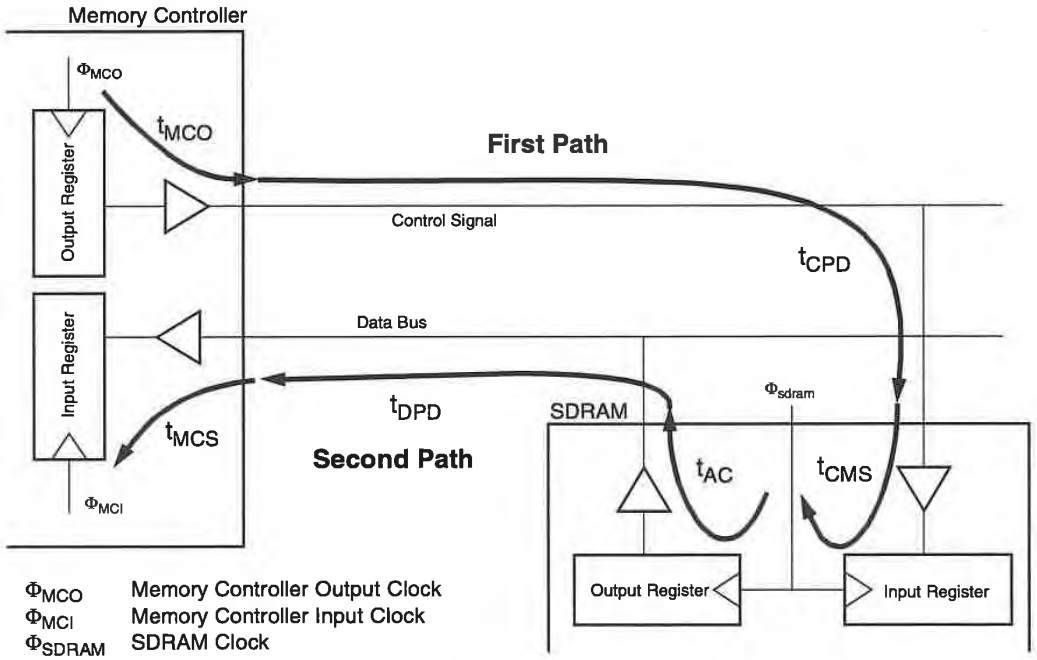
Figure 7-13 illustrates why 100-MHz SDRAMs are needed in 66-MHz systems. Figure 7-13 shows a pair of critical timing paths in a typical SDRAM memory system. The first path originates with a clock on the memory controller, follows the flow of a control signal to the SDRAM, and terminates in the input register in the SDRAM. The three components of this path are the

10. The data-valid time is the time that output is guaranteed to be valid: that is, the cycle time (t_{CLK}) plus the output hold time (t_{OH}), minus the data access time from the clock (t_{OVC}).

Samsung Ex. 1097, p. 208
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

clock-to-output delay of the memory controller (t_{MCO}), the propagation delay of the control signal (t_{CPD}), and the SDRAM's control signal setup time (t_{CMS}). The sum of these three must be less than the sum of the cycle time (t_{CYC}) plus the skew between the memory controller output clock and the SDRAM clock ($t_{MCO-SDRAM}$).

The second path starts with the SDRAM clock, follows the output of a data bit back to the memory controller, and terminates in the memory controller's input register. The three components of this path are the SDRAM's access time from clock (t_{AC}), the propagation delay of the data bit (t_{DPD}), and the memory



Φ_{MCO}	Memory Controller Output Clock
Φ_{MCI}	Memory Controller Input Clock
Φ_{SDRAM}	SDRAM Clock
t_{MCO}	Memory Controller Clock-to-Output Time
t_{CPD}	Control Signal Propagation Delay
t_{CMS}	SDRAM Control Setup Time
t_{AC}	SDRAM Access Time from Clock
t_{DPD}	Data Signal Propagation Delay
t_{MCS}	Memory Controller Setup Time

Figure 7-13 The difficulty in achieving at-speed operation derives from the SDRAMs' long access time from clock

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 209
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

Timing Parameter	Symbol	Typical Value	Notes
Memory Controller Output From Clock	t_{MCO}	5 ns	Fast ASIC
Control Signal Propagation Delay	t_{CPD}	5 ns	LVTTL DIMMs
SDRAM Control Signal Setup Time	t_{CMS}	2 ns	83-MHz part
SDRAM Access Time From Clock	t_{AC}	11 ns	83-MHz part
Data Propagation Delay	t_{DPD}	6 ns	LVTTL DIMMs
Memory Controller Setup Time	t_{MCS}	3 ns	Fast ASIC
Sum of the Above		32 ns	
Two Clock Periods	$2 \times t_{CYC}$	30 ns	
Memory Controller Clock Skew Needed		4 ns	

Table 7-6 With LVTTL DIMMs, 66-MHz operation is hard with 83-MHz SDRAMs.

controller data setup time (t_{MCS}). The sum of all three components must be less than the cycle time (t_{CYC}) plus the skew between the memory controller input clock and the SDRAM clock ($t_{SDRAM-MCI}$).

An examination of both paths leads to the conclusion that the sum of all six components must be less than two cycle times plus the skew between the memory controller's input and output clocks. Otherwise, the data cannot get back into the memory controller's input register in time.

In a simple memory controller, the input and output clocks are copies of a master clock and therefore very little skew exists between them. If skew is deliberately added between the input and output memory controller clocks, then the maximum amount of skew that can be added is determined by the memory controller's and SDRAM's data-hold times: the designer must ensure that the correct piece of data is always latched.

PC memory configurations are characterized by two to four ranks of memory constructed of DIMMs. With LVTTL levels, it can easily take 6–8 ns for a signal to propagate the length of the memory data bus and stabilize at the far end. Table 7-6 shows typical timing parameters and demonstrates that 66 MHz is not easy even with 83-MHz parts. This mismatch—between 1) the sum of the delays and 2) two times the clock period—indicates the problem: with these timings, data from the SDRAM will arrive in the memory controller late.

Samsung Ex. 1097, p. 210
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

There are three ways to solve this problem:

- System vendors can buy 100-MHz SDRAMs that have 9-ns clock-to-data-out times.
- DRAM vendors can improve the design of their SDRAMs to reduce the clock-to-data-out times in their slower—and presumably cheaper—speed bins.
- The SDRAM clock and the memory controller or CPU input register clock can be phase-shifted with respect to the memory controller output register clocks to give more time to the flow of data coming back to the memory controller.

Of these three options, the second is obviously preferred. However, without on-chip PLLs and tighter clock requirements into the SDRAM, it is difficult for the DRAM vendors to dramatically improve these data-out times.

Buffered DIMMs offer little help

The DIMMs standardized by JEDEC for SDRAMs come in buffered and unbuffered versions. The buffered varieties include a phase-locked loop (PLL) for generating a phase-coherent local copy of the clock, and a register for latching and redistributing the address and control signals. Though the on-DIMM PLL helps temporally place the SDRAM clock exactly where it is desired, the register does not substantially increase the maximum operating frequency of the system. Since the register clock and the SDRAM clock are necessarily the same, it is impossible to move them with respect to one another. Consequently, the most time-critical part of the cycle to the SDRAM and back to the memory controller—that being from the SDRAM clock back to the memory controller via the data bus—is not impacted by the addition of the register. The register's primary contribution is to add a full cycle of latency to all memory operations.

The SDRAM DIMMs encode the data sheet minimum cycle time in their presence-detect pins. However, since the temporal interface currently varies significantly across manufacturers, even knowledge of the minimum cycle time is not adequate to ensure proper and reliable operation at any specific system operating frequency.

Small systems can operate at speed

In small memory systems and in those with soldered-down parts, signal integrity and data transit times are substantially improved, often facilitating at-speed operation with a minimum of clock shifting.

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 211
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

As SDRAMs mature and migrate to better processing technologies, new and faster speed bins will be introduced. However, since a large part of the limitation on the memory system's operating frequencies is derived from memory controller characteristics and from the physical dimensions of the memory system, system operating frequencies will not increase commensurately.

Other timing inconsistencies

SDRAM vendors are also inconsistent in their definitions of the cycle-based timing parameters. For instance, vendors use two different but related definitions of the access times: 1) vendors may specify the row access time as precise numbers of cycles at several different operating frequencies, or 2) they may specify it as a base value, measured in nanoseconds. With the second definition, the number of cycles is obtained by taking the base value, dividing by the cycle time, and rounding up to the nearest integer. Though the first type of specification can be used for all higher level timing parameters, the latter cannot be used for those parameters that are inherently a fixed number of cycles independent of the cycle time.

Currently, some vendors use cycle counts (the first method) to specify the behavior of their SDRAMs, while others use a combination of the two approaches. DRAM vendors with the flexibility to choose how to specify their parts can describe how their parts work as accurately as possible. This inconsistency in specification styles does not really interfere with compatibility within a system. However, the system designer's task of ensuring compatibility is made more complicated.

7.4.2 Logical Interface

The internal structure of a 16-Mbit ×8 SDRAM is shown in Figure 7-14. The primary elements are two 1-Mbyte memory arrays with their own row decoders and sense amplifiers. Common to both banks are the column address latch, counter, and decoders; the row address latch and refresh counter; the data I/O buffers; the control block; and the mode register. The sense amps form a row cache that consists of two blocks of 4 Kbits each. Unlike some other DRAM architectures, SDRAMs have no row-cache storage apart from the sense amplifiers. Therefore, the contents of a row-cache block are lost whenever the corresponding DRAM bank is precharged.

SDRAMs differ from most other DRAM families in how the addressing varies across members of the family. Figure 7-15 shows that at a given density, the number of row address bits is

Samsung Ex. 1097, p. 212
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

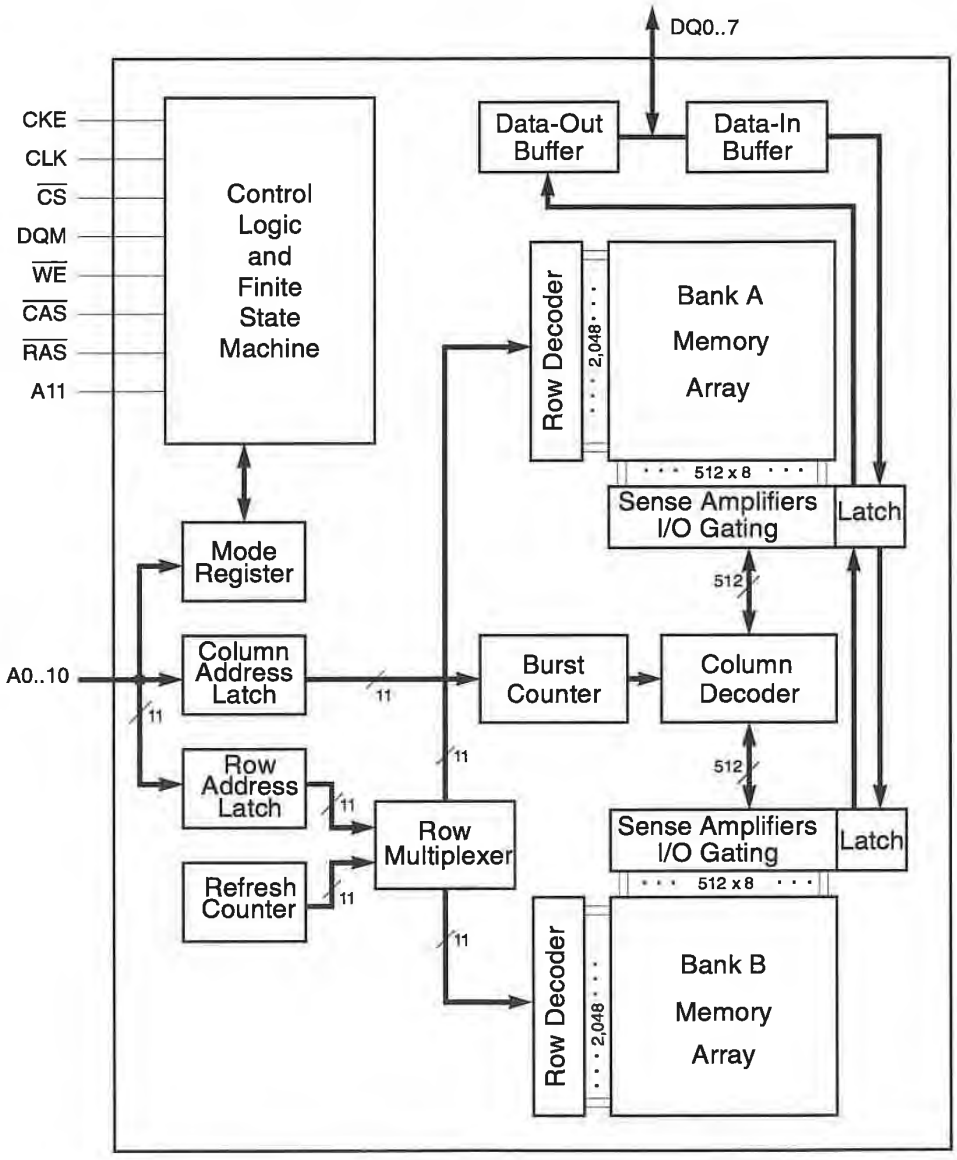


Figure 7-14 An SDRAM has two internal banks and a mode register that customizes the SDRAM's behavior.

© 1996 Steven Przybylski

constant, and all of the trade-off between width and address bits happens in the column address. All 4- and 16-Mbit devices are dual-bank devices; at the 64-Mbit level, both dual-bank and quad-bank parts will probably exist. Because the row-cache size is a function of the number of row address bits, keeping the number of row address bits constant means that all 16-Mbit SDRAMs have the same cache size. This makes it easier to design a single die that can be bonded-out in any of two or three widths.

SDRAMs are unlike the other DRAM families in another way: the number of column address bits is strictly a function of the width of the device and not its density. In the transition from 4-Mbit to 16-Mbit to 64-Mbit at a constant width, the number of row address bits increases, while the number of column address bits stays the same.

In byte-parity SIMMs constructed of fast-page-mode DRAMs, each 9-bit byte is typically formed by combining one lower density ×1 device with two ×4 devices of the current generation. However, it is not possible to mix SDRAMs of different densities in a single **rank**[11] because SDRAMs of different densities do not have the same number of row and column address bits. In place of this capability, JEDEC included ×9 and ×18 devices in the standard. Since no vendor is currently planning on these latter wider parts, it is essentially impossible to build a byte-oriented parity memory system. However, this deficiency of the SDRAM architecture is increasingly irrelevant due to the general trend to ECC and word-oriented memory systems when memory protection is needed. The need to mix devices of different densities will resurface when it becomes desirable to combine two ×32 64-Mbit devices with a ×8 16-Mbit device to form a ×72 memory system. Given current plans for the 64-Mbit generation, the ×8 16-Mbit and ×32 64-Mbit devices will not be combinable in a single rank because each will have different numbers of row and column address bits.

Mode register facilitates programmable characteristics

The mode register is a key feature differentiating SDRAMs from generic DRAMs. The mode register provides a mechanism for customizing the behavior of SDRAMs to meet specific system needs. Figure 7-16 shows the format of the mode register, which is loaded with the contents of the address bus when all of $\overline{CS}$, $\overline{RAS}$, $\overline{CAS}$, and $\overline{WE}$ are low on the rising edge of the clock. The most important fields are the burst length (BL), the burst type (BT), and the latency mode (LTMode). The burst length encodes

11. See Section 4.2.2 for the distinction between bank and rank.

Samsung Ex. 1097, p. 214
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

Density	Width	RAS/ CAS	A13	A12	A11	A10	A9	A8	A7	A6	A5	A4	A3	A2	A1	A0
4 Mb	×16	RAS	NP	NP	NP[a]	NP	BS[b]	Row Address								
		CAS	NP	NP	NP	NP	BS	AP	Column Address							
16 Mb	×4	RAS	NP	NP	BS	Row Address										
		CAS	NP	NP	BS	AP[c]	Column Address									
	×8	RAS	NP	NP	BS	Row Address										
		CAS	NP	NP	BS	AP	X[d]	Column Address								
	×16	RAS	NP	NP	BS	Row Address										
		CAS	NP	NP	BS	AP	X	X	Column Address							
64 Mb	×4	RAS	BS	BS/R12[e]	Row Address											
		CAS	BS	BS/X[f]	X	AP	Column Address									
	×8	RAS	BS	BS/R12	Row Address											
		CAS	BS	BS/X	X	AP	X	Column Address								
	×16	RAS	BS	BS/R12	Row Address											
		CAS	BS	BS/X	X	AP	X	X	Column Address							
	×32	RAS	BS	BS/R12	Row Address											
		CAS	BS	BS/X	X	AP	X	X	X	Column Address						

a. NP = Not Present.
b. BS = Bank Select.
c. AP = Auto-precharge.
d. DC = Don't care.
e. BS/R12 = Bank select or row address bit, depending on whether two or four banks.
f. BS/X = Bank select or Don't Care, depending on whether two or four banks.

Figure 7-15 Within each family of SDRAMs, the number of column address bits is constant.

the $\log_2$ of the burst length. Most vendors are implementing bursts of one through eight chunks, and some are implementing full page transfers as well. Full page transfers facilitate moving a complete row into or out of the DRAM in a single sustained burst. The length of a full page burst varies with the number of

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 215
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

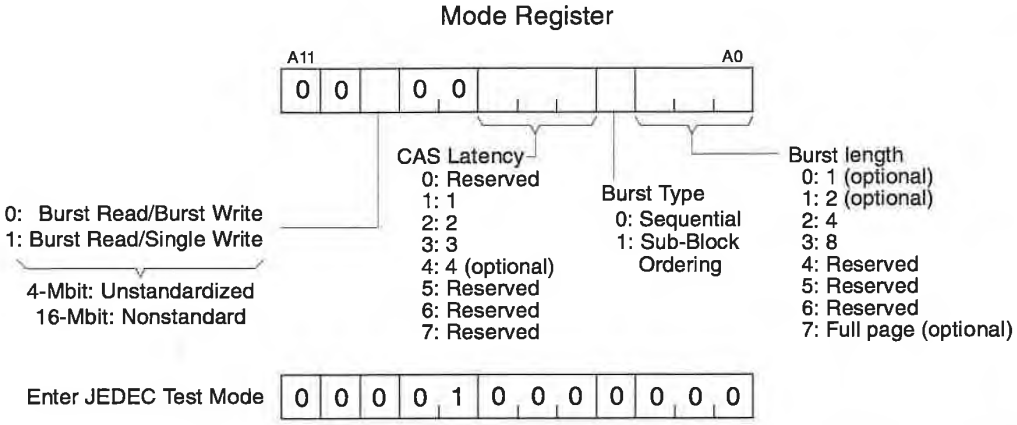


Figure 7-16 The SDRAM mode register controls the burst characteristics, the column access latency, and the test mode.

column address bits, and so is a function of the width of the device. The burst type determines whether the bursts are sequential or **sub-block ordered.**[12] Sequential bursts begin with the addressed data and continue to the end of the binary-aligned block containing the addressed data, at which point the transfer wraps around to the beginning of the block. Some vendors restrict bursts of one—and in some cases, bursts of one and two—to sequential bursts.

The latency mode field encodes the CAS latency—the latency in cycles from a column address to the data appearing on the data bus. The CAS latency is programmable: as the cycle time increases, the number of cycles needed by the SDRAM to perform a column access decreases. The same principle applies to a variety of operations. However, the CAS latency is the only timing parameter that causes SDRAM behavior to change on a cycle-by-cycle basis, depending on the operating frequency. Since the SDRAMs cannot know the current operating frequency or whether the memory controller expects the data to be delayed by one cycle or two, the latency field is provided to instruct the SDRAM in which cycle to present read data. Table 7-7 shows the values of a variety of temporal parameters as a function of **speed grade** and operating frequency.

12. Sub-block ordering is also called interleaved or Intel ordering. See the entry for sub-block ordering in the glossary (Appendix B).

Samsung Ex. 1097, p. 216
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

Speed Grade			-10				-20			
Frequency	f_{OP}	MHz	100	66	50	33	100	66	50	33
Cycle Time	t_{CYC}	ns	10	15	20	30	10	15	20	30
CAS Latency	t_{CAC}	cycles	3	2	2	1	—	3	2	2
RAS to CAS Delay	t_{RCD}	cycles	3	2	2	1	—	3	2	2
RAS Latency	t_{RAC}	cycles	6	4	4	2	—	6	4	4
RAS Cycle Time	t_{RC}	cycles	10	7	6	4	—	10	7	5
RAS Active Time	t_{RAS}	cycles	7	5	4	3	—	6	4	3
RAS Precharge Time	t_{RP}	cycles	3	2	2	1	—	4	3	2

Table 7-7 The timing parameters specified in cycles are a function of the speed grade and the operating frequency. (Source: *NEC uPD42116 Advance Data Sheet*).

For the 16-Mbit generation, the JEDEC standard defines the high-order three bits of the mode register to be zero. The JEDEC standard does not specify the format of the 4-Mbit device's mode register. This is primarily because the 4-Mbit device was initially offered by Hitachi alone prior to the device's adoption by JEDEC. Hitachi decided to put some additional functionality—specifically, a burst-write enable—in bit 9 of the register. When the 4-Mbit pinout was added to the SDRAM standard, the new mode-register format was not included—presumably because the 4-Mbit mode register differed from the 16-Mbit definition.

Since the introduction of the 4-Mbit SDRAMs, most vendors have added the burst-write enable bit to A9 of the mode-register of their 16-Mbit devices. Whether the use of this bit in this way violates the standard or whether it is an allowed extension of it is open to debate. However, JEDEC could decide to use this bit for something else. In the extremely unlikely event of this occurring, 16-Mbit parts with such a burst-write enable would become incompatible with the standard.

Mode register writes can be costly

An unfortunate aspect of the SDRAM definition is that the mode register must be loaded to change the burst length. In a system in which instruction cache refills, data cache refills, and I/O block transfers are all of different lengths, the controller will have to either closely manage the mode register or break up longer transfers into smaller chunks. Both of these alternatives involve extra overhead.

An SDRAM must be completely idle before the mode register can be programmed. Once both banks are closed and precharged, a single-cycle mode-register set command is issued,

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 217
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

followed by three idle cycles. Only then can a new command can be issued. If the mode register has to be changed on a regular basis, substantial portions of the total available peak and sustainable bandwidth are consumed performing these mode-register changes. In the very worst case, the SDRAM's sustainable bandwidth is reduced by a factor of eight.

Breaking up long transfers into shorter requests may not involve a reduction in the sustained bandwidth provided by the SDRAM. Although splitting long transfers into sequences of short transfers does consume additional command bandwidth, command bandwidth is generally in copious supply. Splitting long requests up is only moderately complicated for the memory controller. The problem with this strategy is that it only works with sub-block ordered requests and linear requests with a start address of zero. Unfortunately, for sequential bursts with non-zero starting offsets, the wrapping at the end of the block varies according to burst size.

Of the burst lengths defined, only bursts of four and eight are not optional. A full page burst is only ever defined for sequential bursts. Most vendors have chosen to implement bursts of one, two, four, and eight. However, some vendors have made a burst of one—or in some cases, one and two—valid only for sequential bursts. Most, but not all, have chosen to implement full page sequential bursts.

Figure 7-17 illustrates the operation of a Synchronous DRAM. The figure shows a simple write operation, followed first by a read operation to the same row, and then by an explicit precharge of the active bank. The burst length has been previously programmed to be two. In this example, both the CAS latency and the RAS-to-CAS delay are three cycles. Note that:

- A column access command initiates a burst of data.
- Precharge can begin before the data-read burst is complete.

The limited overlap of row precharge with data transfer helps improve the SDRAM's bus utilization. With a burst length of four, most SDRAMs sustain about 33% of peak bandwidth when satisfying alternating row accesses and column accesses to the same bank. However, the interface's real strength shows when the two banks operate simultaneously. Unfortunately, with only two banks per device, this cannot not happen as often as would be ideal in a single-rank memory system.

Samsung Ex. 1097, p. 218
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

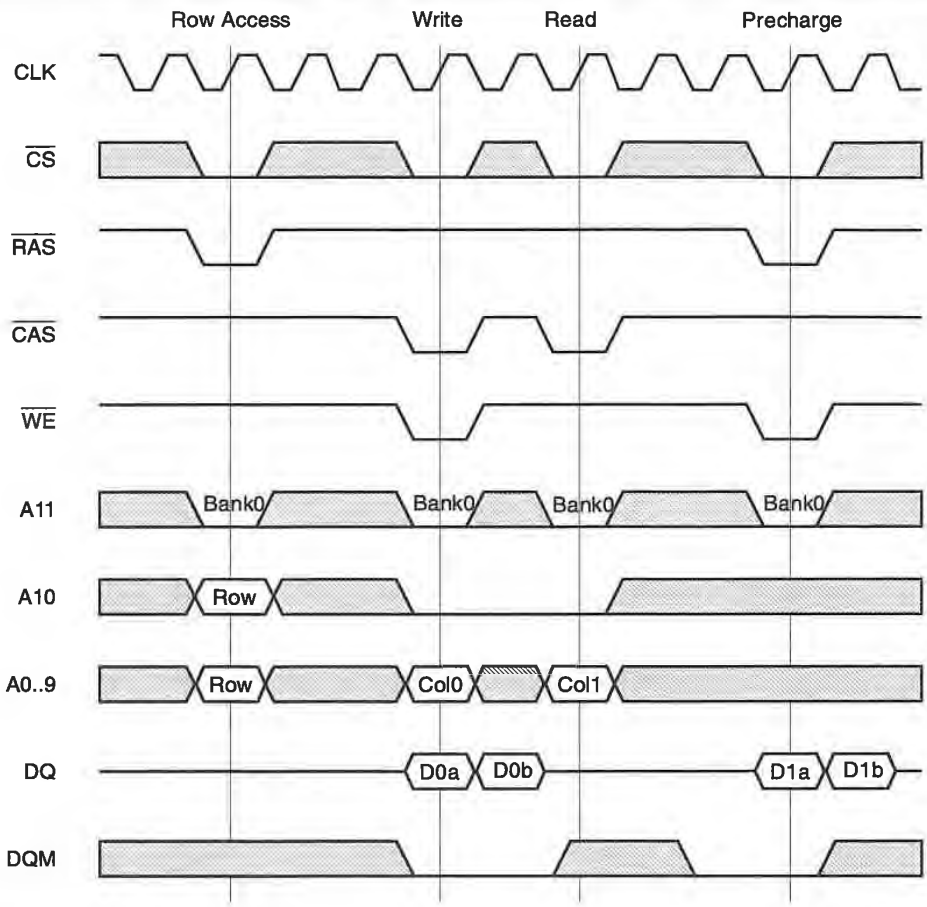


Figure 7-17 Commands encoded on the control pins specify the SDRAM's operations.

Advantages of two banks

The dual-bank organization can be used to improve the sustained bandwidth in two ways. Figure 7-18 illustrates the first way. When short bursts are appropriate, after both banks have completed row accesses, successive column accesses can alternate between the two banks. Because column accesses can be pipelined, short random accesses to one or both of the rows are sustainable at the full speed of the data bus. Column read and write operations to the two banks can be scheduled so that the gaps between the end of one transfer and the start of the next are minimized or eliminated. Some SDRAMs allow new column accesses every cycle; with these parts, the data bus can be utilized fully even for bursts of length one.

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 219
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

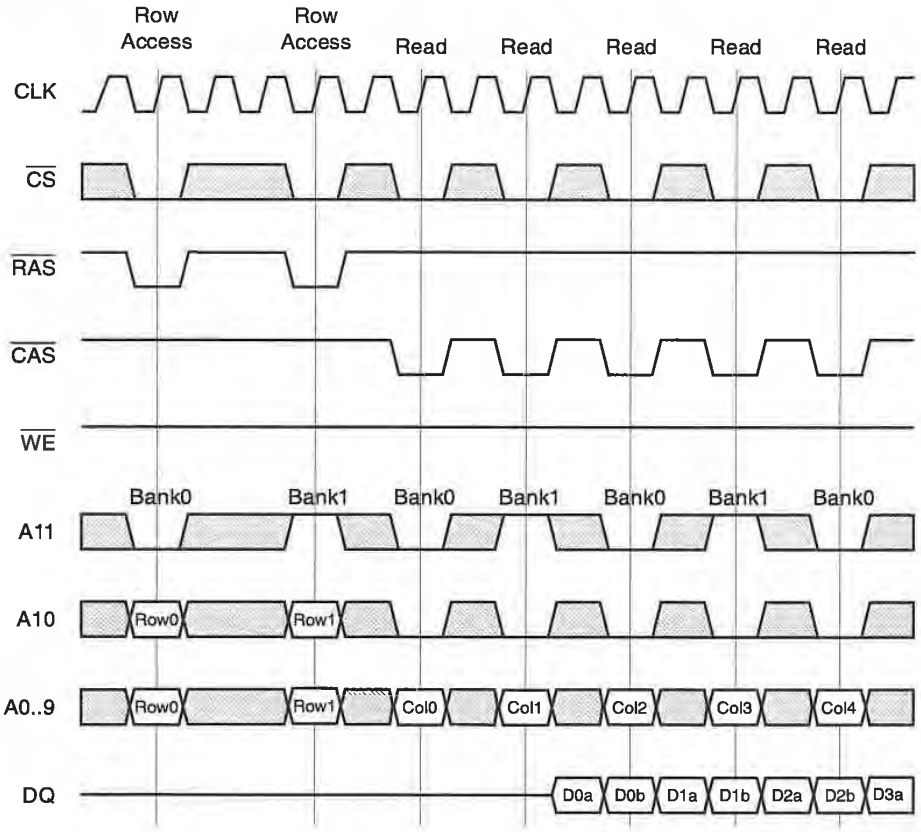


Figure 7-18 Alternate column accesses can be to different SDRAM banks.

Figure 7-19 illustrates the second way in which two banks can improve sustained bandwidth. With longer bursts, a row access to one bank can be partially or completely hidden under column accesses to the other bank. In Figure 7-19, it takes nine cycles from the end of a burst for a bank precharge, a row access, and a column access to occur. As shown, this time can be almost completely filled with a single burst access of length eight to the other bank. Thus, by alternating between the two banks and by using longer bursts, the data bus can be completely utilized even in the face of continual row-cache misses. Of course, this technique relies on a reference stream that strictly alternates between the banks. In general, the more banks in the memory

Samsung Ex. 1097, p. 220
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

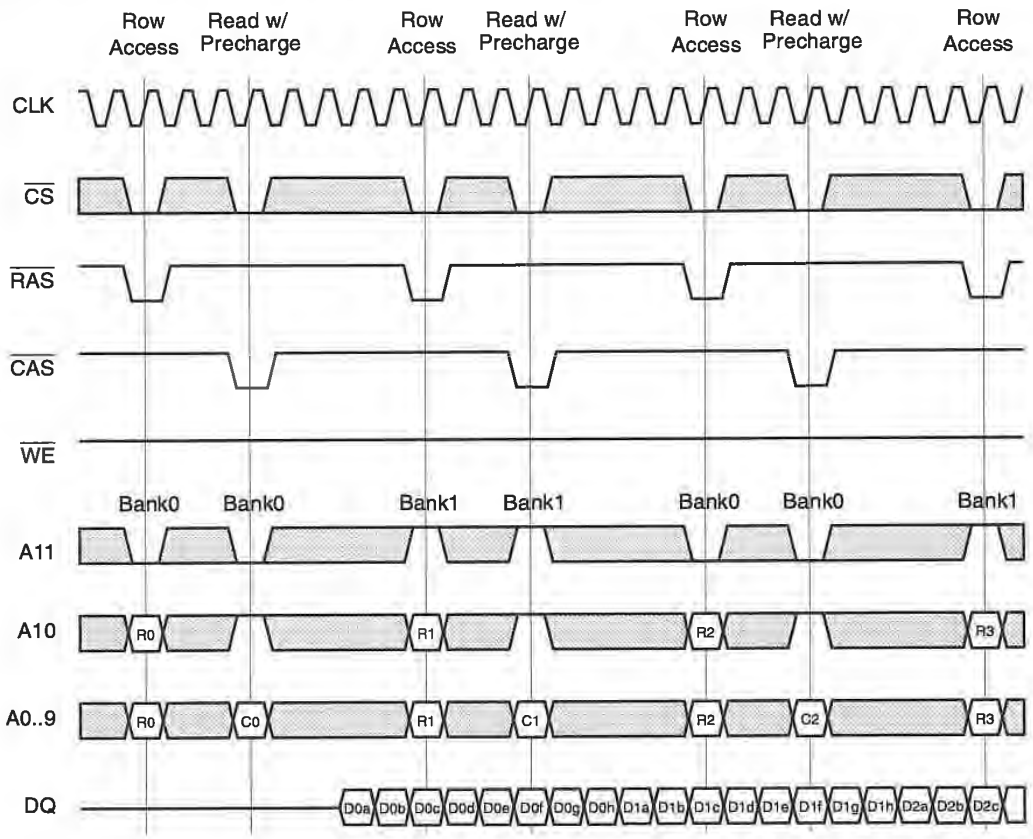


Figure 7-19 Dual-bank SDRAMS can achieve 100% bus utilization—even in the face of continuous row-cache misses—by using long bursts and alternating banks.

system, the higher the likelihood that adjacent references will not access the same bank.

Other significant features include the ability to abort burst transfers; optional automatic precharge at the end of a burst; autonomous self-refresh mode; power-down mode; data masking on input and output; and the ability to precharge both banks with one command. Table 7-8 lists all the commands.

Figure 7-20 shows the state diagram of a single bank of the SDRAM, including all the allowed transitions. This figure demonstrates the symmetry in the command set between reads and writes, and between normal bursts and those with an automatic

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 221
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

Command	Conditions					Notes
	$\overline{CS}$	$\overline{RAS}$	$\overline{CAS}$	$\overline{WE}$	Other	
NOP	L	H	H	H		No operation
Mode-Register Set	L	L	L	L		Load the mode register from A0..7,A9..11[a]
Row Access	L	L	H	H	A11 = Bank	Load a row-cache block with a row from one bank
Read	L	H	L	H	A11 = Bank A10 = L	Column (row-cache) burst read
Read w/Auto-Precharge	L	H	L	H	A11 = Bank A10 = H	Column read with automatic precharge of the bank at the end of the transfer
Write	L	H	L	L	A11 = Bank A10 = L	Burst write into the row cache and the DRAM array
Write w/Auto-Precharge	L	H	L	L	A11 = Bank A10 = H	Burst write into the row cache and the DRAM array followed by a row precharge
Precharge	L	L	H	L	A11 = Bank A10 = L	Precharge one bank
Precharge All Banks	L	L	H	L	A10 = H	Precharge both banks
Burst Terminate[b]	L	H	H	L	A11 = Bank	Terminate a burst only
Auto-Refresh	L	L	L	H		Refresh both banks from on-chip refresh counter
Self-Refresh	L	L	L	H	CKE = L	Refreshing autonomously while CKE = L
Power Down	X	X	X	X	CKE = L	Power down if both banks precharged when CKE is lowered

Table 7-8 Control information is encoded into the control signal and the upper address bits.

a. Use of mode-register bits A9..11 is nonstandard. The JEDEC specification defines them to be zero.
b. Not part of the JEDEC standard, though the standard reserves this command for this function.

precharge at their conclusion. The only asymmetry of substance in the state diagram is that the regular read and write bursts can be arbitrarily terminated by a precharge command or another read or write command. However, the bursts with an automatic precharge can be suspended but never terminated: they must eventually complete with the aforementioned precharge operation.

Dual-bank row cache

The SDRAM model of a row cache is very similar to that of a generic DRAM, except that there are two blocks of row cache per DRAM instead of just one. As a row access is completed, a block of the row cache is loaded with a page from a DRAM array. Read and write column accesses can be made freely to that block of the cache so long as the row is active. As soon as the DRAM array is precharged in preparation for the next row

Samsung Ex. 1097, p. 222
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

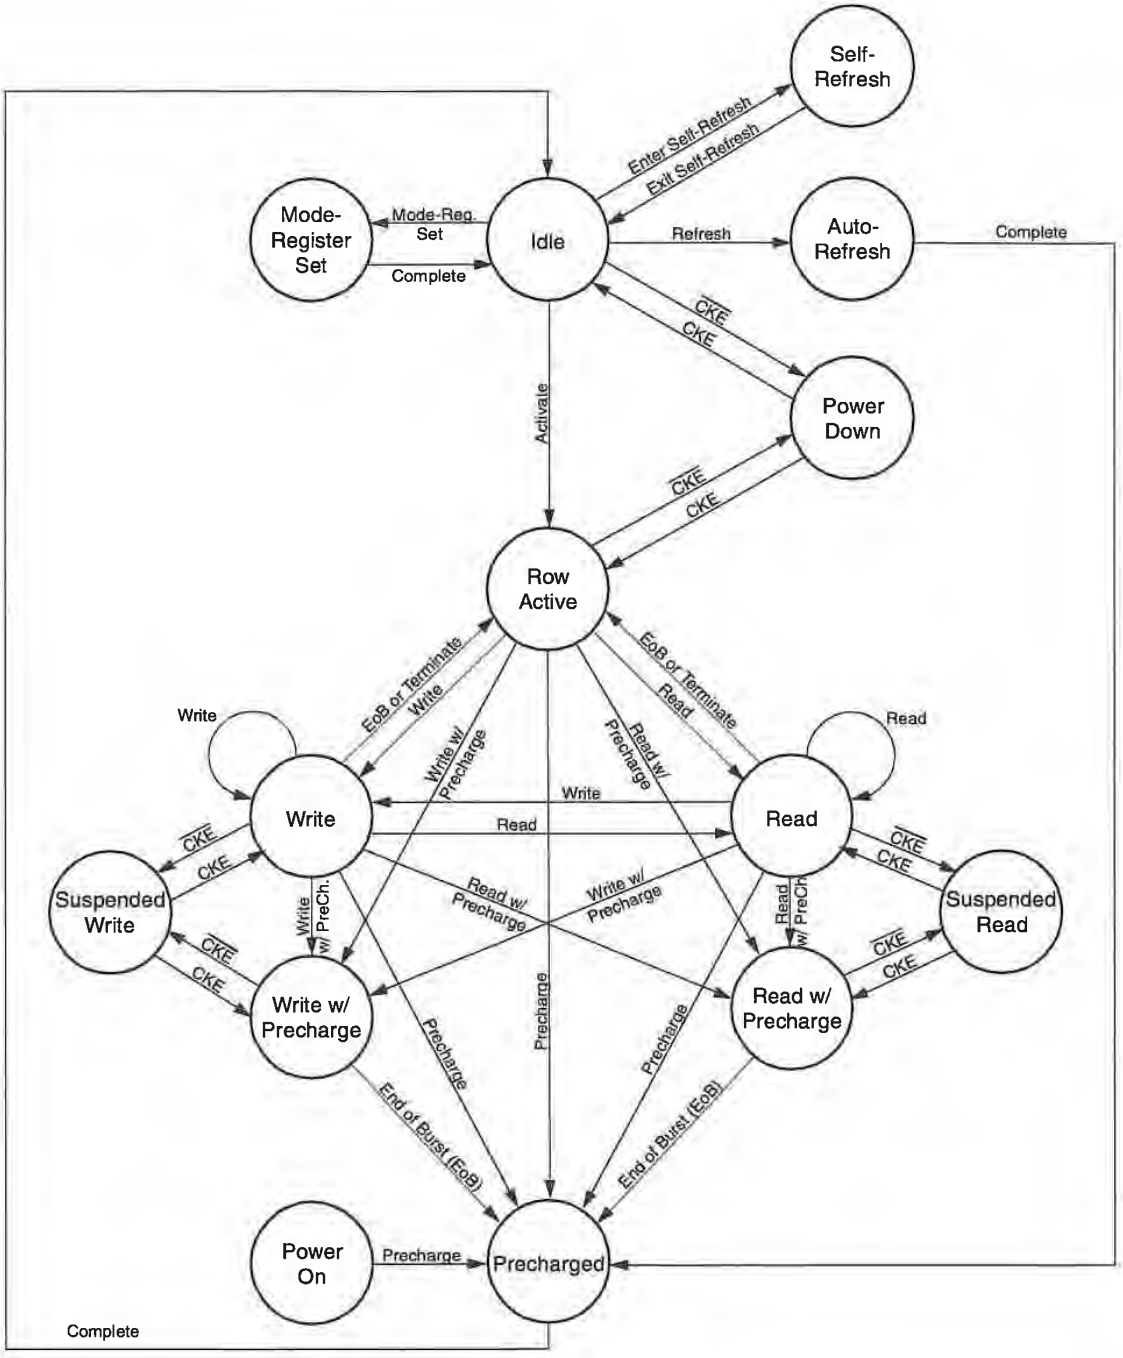


Figure 7-20 The state diagram for an SDRAM bank has 14 states. (Source: *NEC Spring 1995 Data Book*).

© 1996 Steven Przybylski

access, the corresponding block of the row cache becomes unavailable until the next row access completes.

The pulsed-RAS model of the interface has a subtle positive impact on the row-cache organization. In most memory systems constructed of fast-page-mode or EDO devices, the $\overline{\text{CAS}}$ lines are shared among all the ranks. This means that only one rank can have $\overline{\text{RAS}}$ active at a time—otherwise more than one rank would respond to column accesses. Consequently, in a multirank memory system, only one block in the row cache can be active—i.e., hold data—at a time. In the case of SDRAMs, all banks in all ranks of a memory system can be active simultaneously: all the bits in the row cache can hold data at the same time. In larger memory systems with many ranks, this improvement in the row-cache effectiveness can significantly change the row-cache miss ratio or, alternatively, the mix of row and column accesses.

Flexible precharge operations

Precharge operations can be initiated in one of two ways. Explicit commands can be used to precharge one or both banks. Alternatively, read and write operations can be annotated to cause a precharge of the referenced bank at the end of the data transfer. In general, a cycle or two elapses while the data flows from the sense amplifiers to the output pins. To account for this latency, explicit precharge commands are allowed to occur one to two cycles before the first cycle after the end of a read transfer. The precharge command can be presented at the same time as the last chunk of data if the column access latency is one cycle. If the column access latency is two or more, then the precharge command can be presented concurrently with the next-to-last piece of output data.

The minimum time between the end of a burst write and a precharge command to the same bank is vendor-specified: many of a vendor's first parts allow a precharge command in the cycle immediately following the last write chunk of data. However, several vendors require one extra cycle of delay between the end of the write data and the precharge command. At least one vendor requires two extra cycles of delay in some circumstances.

An important characteristic of the SDRAM truth table is that chip select ($\overline{\text{CS}}$) must be asserted for all commands: if $\overline{\text{CS}}$ is not asserted, then all other control signals are ignored. The implication is that only $\overline{\text{CS}}$ needs to operate at full speed. The other high-fanout control signals (A0..11, $\overline{\text{RAS}}$, $\overline{\text{CAS}}$, and $\overline{\text{WE}}$) can change more slowly, provided they have reached valid levels at the clock edges when $\overline{\text{CS}}$ is asserted. This optimization can

Samsung Ex. 1097, p. 224
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

reduce the number of high-speed buffers needed in the system to four (for CLK, CKE, DQM, and $\overline{CS}$) at the expense of an extra cycle or two of latency for some operations.

Two SDRAM pins are used to control the flow of data onto and off of the DRAM: Clock Enable (CKE) and Data Mask (DQM). Asserting DQM during a read burst tri-states the data bus two cycles hence: the data that was to be driven then is never driven. Asserting it during a burst write causes individual pieces of data to be ignored and not written into the row cache and DRAM array. Unlike in a read, there is no temporal offset when using DQM in a write burst: if DQM is asserted during the third cycle of a burst, the third datum is not written.

Read and write bursts can be spread out in time by using CKE. Asserting CKE during a cycle negates all effect of the next clock within the DRAM: latency and burst counters are not incremented, write data is ignored, and read data that would normally be output in that cycle is output in both that cycle and the next.

Power-down mode

The SDRAM power-down mode is entered by deasserting clock enable (CKE) when both banks are precharged. Once power-down mode is entered, the clock can be withdrawn altogether to further reduce power consumption. If a Self-Refresh command is given in the cycle that CKE is lowered, then the DRAM will autonomously refresh itself for the duration of the power down, regardless of whether the clock is withdrawn or not. Power-down mode is exited by restoring the clock and reasserting CKE. Most parts allow a new command to be delivered to the SDRAM in either the same cycle that CKE is reasserted or the next cycle. The JEDEC standard actually allows the vendor to specify the number of cycles that must elapse between exiting power down and issuing the first new command, so this is a potential source of incompatibility among the different vendors' parts.

If the SDRAM is in self-refresh mode while powered down, then it could be just beginning a refresh operation when power-down mode is exited. In this case, the first new command must be postponed for a full row access time after the restoration of CKE. The length of this delay is another vendor-specified time.

The intent that SDRAMs come out of a clockless power-down mode essentially instantaneously precludes the use of a **Phase-Locked Loop (PLL)** in SDRAMs. PLLs are further precluded in SDRAMs by the absence of a statement in the JEDEC specification that the clock is a single frequency: both variable length

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 225
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

and variable duty cycle clocks are technically allowed. The penalty for banning Phase-Locked Loops is that SDRAMs have larger than necessary clock-to-data and clock-to-high-Z access times. Ironically, some SDRAMs will end up on buffered SIMMs with a PLL in the clock distribution ASIC. This burdens the system with the power-down constraints of a PLL without providing the benefits of a PLL to the SDRAM itself.

Incompatibilities frustrate designers

Even among JEDEC-compliant SDRAMs, significant differences in timing and functionality make it awkward or difficult for a memory system designer to ensure compatibility among multiple sources of SDRAMs. Though some differences have been reduced over the past two years, both substantial and subtle incompatibilities remain. Differences in supported burst length, mode-register configuration, column access repeat rate, and timing details have been noted. In addition, some manufacturers plan to offer unusual speed grades and subtly nonstandard functionality. Most significant is the lack of consistency for the coarse timing parameters: the number of cycles an SDRAM needs to perform specific functions at specific clock rates.

SDRAMs also show inconsistency about the termination of burst transfers. A burst transfer can be terminated in its middle by sending a new read, write, or precharge command to the same bank. In the case of a burst write, the standard is unclear on which datum is the last one of the terminated burst to be written into the DRAM array: it could be the datum presented in the cycle before the new command, the datum presented in the cycle with the new command, or even the datum presented in the cycle after the new command. However, any incompatibility can be masked by using the data mask pin (DQM) to explicitly indicate where the write data stops.

The infamous 2n-rule

Two types of SDRAM implementations exist: pipelined and interleaved. Internally, interleaved designs are double-wide and half-frequency. Pipelined implementations run at the external frequency throughout the design. In order to facilitate interleaved designs, the JEDEC SDRAM standard defines a *2n-rule*. The rule states that a burst may only be terminated by a new read, write, or precharge command presented on a cycle that is an even number of cycles into the burst. Pipelined implementations do not require controllers to observe this rule. This difference is important for graphics applications that terminate full page transfers of video-refresh data. It is unlikely that main memory controllers will make use of burst termination, largely because terminated bursts cannot be restarted where they left

Samsung Ex. 1097, p. 226
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

off: they must be repeated in their entirety or the remainder must be broken down into smaller, binary-sized, chunks.

However, the 2n-rule is important to main memory controllers that want to pipeline column read or write accesses by issuing a new column access before the previous one has completed, especially if bursts of one are used. SDRAMs that rely on the 2n-rule cannot accept back-to-back column addresses—regardless of the banks involved—whereas completely pipelined SDRAMs are able to accept back-to-back column accesses.

Both pipelined and interleaved SDRAMs continue to be offered. Data sheets are not always clear about the applicability of the 2n-rule. Indeed, one device relies on the 2n-rule only when the fastest speed grade is operated at its maximum allowed frequency.

Other incompatibilities

Other differences among the various SDRAMs include the initilization sequence after power on[13] and the number of cycles of inactivity needed for a mode-register write to take effect. Sometimes very careful reading of a data sheet is necessary to uncover these subtle incompatibilities. For example, one current data sheet states in a footnote that when the CAS latency is three cycles and the burst length is a full page, then bursts cannot be terminated by a write command: only a Burst Terminate command or a Precharge Command will work (reliably).

Many of the ongoing incompatibilities can be traced to the early history of SDRAMs, when several vendors were developing their first designs in parallel with the standardization process. As potential users examined the initial SDRAMs, it became clearer which features and characteristics were problematic or of little value. The differences among the vendors' offerings have significantly diminished over the past two years. However, the process is still in progress; the SDRAMs will not be as homogeneous across vendors as the generic DRAMs for some time. In the short term, most systems designers should be able to find several vendors with products that overlap sufficiently in their logical and temporal characteristics to be effectively interchangeable.

The current plethora of minor incompatibilities creates a significant danger to the end user in the aftermarket. If an SDRAM inserted in a system in the field is subtly incompatible with the

13. The JEDEC standard specifies a specific initialization sequence. However, some vendors also allow more aggressive or slightly different sequences.

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 227
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

assumptions made by the memory controller designer, then the system may not work or, worse, may not work reliably. This may be the case even though the new SDRAMs are of the right organization and speed. This problem is aggravated by the fact that there is no way for a memory controller to determine the specific timing requirements of an SDRAM that it is controlling. Furthermore, as with all memories, the behavior of an SDRAM that is miscontrolled is undefined: if the timing is close to be being acceptable, the SDRAM may function properly some, most, or all of the time. System reliability and system returns due to memory upgrade and speed problems are a significant problem in the PC industry. If this industry were to adopt SDRAMs in their current state, the return rate could only get significantly worse.

Clearly, a memory controller designer must carefully examine many characteristics of the SDRAMs to be used in conjunction with his or her design. A thorough examination of the data sheets and careful formulation of the baseline assumptions will ensure that several vendors' SDRAMs will be compatible with the controller and will minimize problems in the field. Questions that the memory controller designer must answer include:

1. Will the SDRAM be available in the desired speed grades?
2. Does the SDRAM implement all combinations of column access latency and burst transfer orders and lengths that the controller can generate?
3. If the controller can ever violate the 2n-rule, does the SDRAM require adherence to the 2n-rule at any of the possible operating frequencies?
4. Does the SDRAM implement the burst-write enable bit in the mode register? If not, does the controller need single-write capability? If the SDRAM does have the burst-write enable bit, can it be accidentally enabled or disabled?
5. Does the memory controller's power-on initialization sequence properly initialize the SDRAM, in particular with regard to the placement of the mode-register write?
6. Are the SDRAM's setup and hold times compatible with the controller and system configuration?
7. Are the SDRAM column latency and other cycle-based timing parameters consistent with the memory controller's expectations?

Samsung Ex. 1097, p. 228
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

8. If the controller generates Terminate Burst commands, does the SDRAM support them for the burst lengths being used?

9. If read or write bursts are to be terminated by any means, does the SDRAM deliver or write the expected number of chunks of data in the vicinity of the terminating event?

10. If read or write with automatic precharge is to be used, then does the SDRAM complete the precharge before the controller can generate the next row activate command?

11. If a write command can be followed by a precharge command, is the SDRAM's write recovery time consistent with the memory controller's expectations for all expected operating frequencies?

12. If self-refresh will possibly be needed, are the SDRAM's requirements for exiting self-refresh met, especially with respect to the behavior of CKE?

13. Are the memory controller's assumptions regarding the period of inactivity after a mode-register write consistent with the SDRAM's requirements?

14. Does the SDRAM implement any nonstandard feature that could conflict with the controller's expected mode of operation?

7.4.3 Electrical Interface

A serious limitation of the existing generic DRAM interface is its electrical interface. The large voltage swings and heavily loaded, yet unterminated, traces of regular TTL environments make high-speed operation problematic (see Section 5.2.2). In recognition of this limitation, most SDRAM vendors will eventually offer a choice of electrical interfaces. In addition to LVTTL for lower speed operation, all vendors will eventually offer some other interface for use in higher speed systems. The frequency at which the use of LVTTL becomes untenable depends strongly on the memory system's logical and physical structure. But for most large memory systems, a switch away from LVTTL levels is required around 80 MHz. However, engineers will likely be able to make smaller, fixed-size memory systems work reliably up to 100 MHz and maybe slightly beyond.

Unfortunately, SDRAM vendors have had difficulty agreeing on the appropriate electrical interface for high-speed memory systems. Initially, GTL was widely espoused as the electrical

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 229
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

interface of choice, with CTT being offered as a reasonable alternative. However, GTL was found to have inadequate noise margin, and CTT lacked wide support. In 1994 and early 1995 JEDEC's JC 16.3 committee worked to develop a new high speed, low-swing electrical interface for terminated transmission line environments. HSTL, approved as a standard in mid 1995, defines four classes of interconnect according to the placement and type of termination (see Section 5.2.2). However, even HSTL was found inadequate when faced with the poor quality and stub-laden transmission lines found in DIMM-based memory systems. A derivative of CTT, called **Stub Series Terminated Transceiver Logic (SSTL),** has been developed and should become a standard in early 1996. As illustrated in Figure 5-5 in Section 5.2.2, SSTL introduces small resistors between the primary bus on the motherboard and the stubs on the DIMMs. These resistors dampen the reflections off of the far end of the DIMM stubs and significantly improve the signal settling time and general signal integrity on the motherboard bus.

With the advent of SSTL, the industry may finally have an electrical interface that will take traditional memory modules above 100 MHz. Just as the upper frequency limit for LVTTL is strongly dependent on the physical and electrical characteristics of the memory system, so too SSTL has upper limits. Given standard DIMMs and DIMM sockets, Spice simulations indicate that with current clock-to-data-out times, SSTL will yield 125-MHz memory systems in most cases, and up to 150 MHz in more physically constrained environments. These system speeds will probably be attained with parts binned at up to 166 MHz on a tester. In order to reach memory bus speeds in excess of 150 MHz, both the physical size of the memory system and the clock-to-data-out times will have to be substantially reduced.

With SSTL being standardized in early 1996, SSTL-capable SDRAMs will probably first appear in late 1996. Widescale use is unlikely before the end of 1997. Until then, LVTTL is the only mainstream electrical interface for SDRAMs and the other evolutionary alternatives. Consequently, it will be a few more years before conventional memory system operating frequencies rise above the 66–100 MHz range.

7.4.4 SDRAM Fill Frequency

A DRAM's peak fill frequency is defined as the ratio of the DRAM's peak bandwidth to its size (see Section 2.5). Since the

Samsung Ex. 1097, p. 230
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

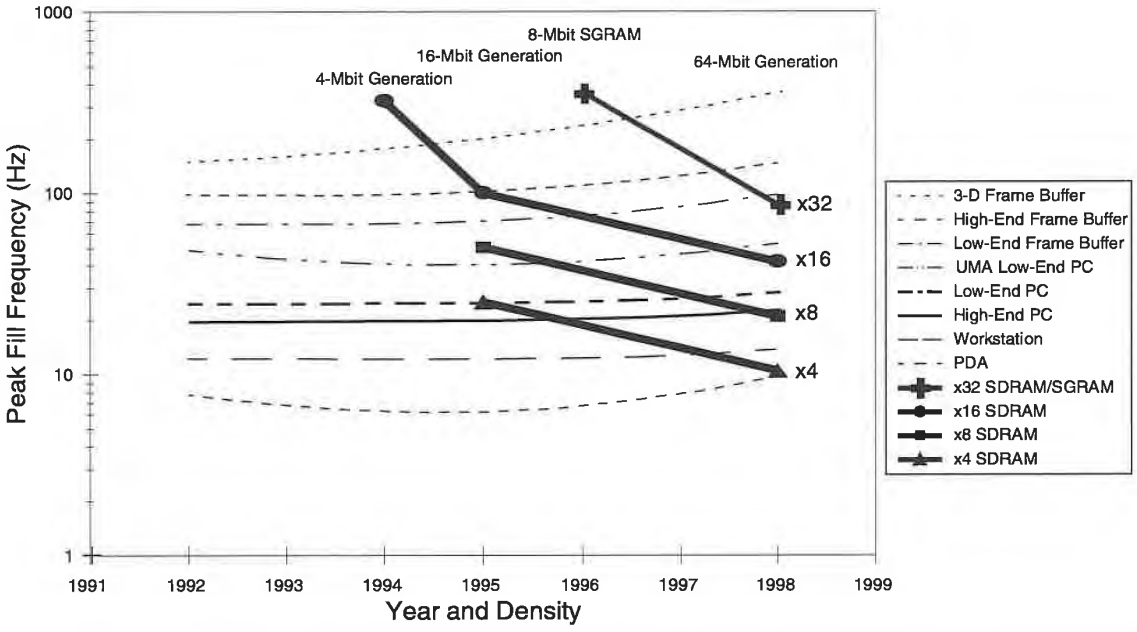


Figure 7-21 SDRAM peak fill frequencies are substantially higher than those of fast-page-mode or EDO devices.

peak fill frequency is independent of system constraints, it is defined strictly in terms of data sheet values. Consequently, the peak fill frequencies of the SDRAMs shown in Figure 7-21 are based on peak frequencies of 80, 100, and 166 MHz for the 4-, 16-, and 64-Mbit generations, respectively. These speed grades are the highest expected to be available; they are not necessarily the frequencies at which these devices would operate in a system. The frequency of operation in a real system plays a fundamental role in determining the sustainable fill frequencies discussed in Chapter 10. The dates used in Figure 7-21 represent the dates on which the various devices were widely available rather than the dates on which they became cost-competitive. The 8-Mbit SGRAM is included on this chart because of its close kinship to standard SDRAMs.

The peak fill frequencies of SDRAMs are substantially higher than those of the conventional DRAM alternatives. Superficially, a number of positive matches exist between target markets and SDRAM capabilities. However, some of apparent matches do not stand up to a closer examination.

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 231
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

A necessary, not sufficient, condition

It is important to remember that meeting a system's peak fill frequency requirement is a necessary but not sufficient condition to ensure a DRAM's acceptability in that system marketplace. A system's fill frequency requirement is a combination of the market's peak bandwidth and memory granularity requirements. When a DRAM's peak fill frequency is lower than the system's peak fill frequency requirement, the DRAM is unacceptable in that system. When a DRAM's peak fill frequency is higher than the system's fill frequency requirement, the DRAM may still be acceptable: the individual bandwidth and granularity requirements must still be checked to verify acceptability of the DRAM. In other words, the fill frequency directly indicates unacceptability, but not acceptability.

For example, based on an examination of the peak fill frequencies of Figure 7-21, it appears that the 16-Mbit ×4 SDRAMs adequately meet the PC marketplace's needs. In contrast, Figures 6-4 and 6-6 indicated that the 4M×4 fast-page-mode and EDO devices have peak fill frequencies below the PC market requirements. However, a closer examination shows that this particular SDRAM satisfies the bandwidth requirement by a large margin but fails to meet the granularity requirement: in a 64-bit memory system, the minimum memory size with any 4M×4 device—be it fast-page-mode, EDO, or SDRAM—is 32 Mbytes: more than that market will currently tolerate. To reiterate: although the fill frequency of the 4M×4 SDRAM does not indicate unacceptability, this organization cannot satisfy the PC marketplace's granularity requirement.

Similarly, a 166-MHz 64-Mbit ×16 SDRAM has a peak fill frequency of 41.5 Hz.[14] This peak fill frequency appears adequate for the PC main memory marketplace. However, this SDRAM's granularity in 64-bit memory systems will also be 32 Mbytes. Unless DRAM prices fall below $10 per megabyte by 1998, 32 Mbytes will still be an unacceptably large minimum memory size. Thus, if SDRAMs are being used in entry-level PCs in 1998, then the 2M×32 organization will be used in those systems shipping with the minimum memory size. The narrower 4M×16 SDRAMs would be used only in the larger and more expensive system configurations.

14. Its bandwidth is 166 MHz×16 bits/cycle, and its size is 64 Mbits. The ratio of the two is 41.5 Mbits/s per Mbits, or 41.5 Hz.

Samsung Ex. 1097, p. 232
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

7.4.5 SDRAM-Lite

In early 1995, Burst EDO DRAMs were announced with much fanfare. For a while it looked as if Burst EDO would be adopted as the next major DRAM standard for PCs. At the time, SDRAMs were languishing with high price premiums and few volume design wins. Several major DRAM vendors were concerned about the prospect of Burst EDO becoming dominant in the short term because of the DRAM speeds needed to reach 66-MHz system speeds (see Section 7.2). Consequently, a significant effort was made to develop a 16-Mbit derivative of SDRAM that could be more cost-competitive and an effective alternative to Burst EDO.

During the summer of 1995, a proposal for SDRAM-Lite was developed. The proposal was to strip the JEDEC standard SDRAM definition of everything that would not be used in a Pentium-class PC. Consequently, the target was 66-MHz bus operation with a CAS latency of three and burst lengths of one and four. Housed in the same packages as regular SDRAMs, SDRAM-Lites are to be available in ×4, ×8, and ×16 widths, and they are to be binned and marked according to common detailed timing parameters needed to reach specific system clock frequencies: the -50 parts would work in 50-MHz systems; the -67 parts would work in 6-MHz systems; and the -75 parts would work in 75-MHz systems. SDRAM-Lites would be 3.3 V parts and use LVTTL I/O levels. Table 7-9 illustrates the differences between the full JEDEC SDRAM standard, the common SDRAM implementations, and the proposed SDRAM-Lite definition.

Several vendors promised 5% price premiums over comparable speed and width EDO DRAMs in 1996. The 5% price premium represented a significant reduction from the premiums for SDRAMs that were then being projected, and was low enough to force an industry-wide reconsideration of Burst EDO.

The expectation was that by eliminating many superfluous features, the overall cost of producing an SDRAM-Lite device could be reduced significantly below the cost of producing a JEDEC-compliant SDRAM. However, this reasoning failed to take into account the often tenuous relationship between cost and price as well as the relationship between features and cost of DRAMs (see Section 11.3). In particular, since the features removed from the SDRAM specification to make SDRAM-Lite—most notably, the variable column access latency and the long burst lengths—have no impact on the die size, the cost savings for an

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 233
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

Characteristic		JEDEC Standard	Common SDRAMs	SDRAM-Lite
CAS Latency	1	Yes	Yes	No
	2	Yes	Yes	No
	3	Yes	Yes	Yes
	4	Optional	No	No
Burst Length	1	Optional	Yes	Yes
	2	Optional	Yes	No
	4	Yes	Yes	Yes
	8	Yes	Yes	No
	Full Page	Optional[a]	Most	No
Burst Type	Linear	Yes	Yes	Yes
	Sub-Block	Yes	Yes	Yes
Single Write		No	Most	No
Burst terminate		Yes	Yes	No
2n-Rule		Yes	Mostly not	No
Consistent Timing Parameters		No	No	Yes

Table 7-9 SDRAM-Lite parts will only support those SDRAM features needed in PCs.

a. A full-page burst is only defined for linear (sequential) bursts.

SDRAM-Lite over SDRAM will stem only from improved yield and reduced test times.

The removal of these features would have essentially no impact on the yield but could noticeably decrease the test time and, therefore, the test cost. Since the test cost represents a substantial portion of the overall cost-of-goods, improving the testability of the device can significantly reduce a device's total cost.

In opposition to these cost-of-goods savings are the largely intangible costs associated with a proliferation of part numbers. Though there is some marketing benefit to having a full-featured premium line and a lower-price entry-level offering, there are distribution, inventory, and management costs associated with having another set of parts. The DRAM industry is already struggling to manage the proliferation of DRAM types and variations. Even in this case, in which the distinction between an SDRAM and an SDRAM-Lite is only made very late in the manufacturing process, the cost of maintaining an extra product line is justified only if the lower cost of goods derived from a

Samsung Ex. 1097, p. 234
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

reduction in the test time directly facilitates penetration of the PC market.

SDRAM-Lite losing momentum

In the last quarter of 1995, as the DRAM vendors recognized that the cost savings associated with the reduction in test time would not live up to initial expectations, their initial enthusiasm for SDRAM-Lite waned substantially. Though some vendors still plan to introduce an SDRAM-Lite line, it is unclear whether SDRAM-Lite will actually come to fruition, or whether SDRAMs will drop in price sufficiently quickly to accomplish what SDRAM-Lites were meant to do: become the next major DRAM of choice for PCs.

The peak and sustainable fill frequencies of SDRAM-Lite parts are the same as the fill frequencies of the comparable density, width, and speed SDRAMs. However, since the SDRAM-Lite parts are to be marked by system clock frequency—not the maximum frequency on a tester—the peak fill frequencies for 75-MHz 16-Mbit SDRAM-Lite parts are somewhat lower than those for the 100-MHz SDRAMs shown in Figure 7-21. As for the sustainable fill frequencies (used as a metric for comparison in Chapter 10), SDRAM-Lites perform exactly as the SDRAMs do with only a few insignificant exceptions.

7.4.6 History and Barriers to Acceptance

Early in 1993, interest in SDRAMs grew as NEC and Samsung began sampling their initial offerings: NEC's first JEDEC-compliant SDRAM and Samsung's single-bank, level-RAS design. Despite the initial fanfare and a number of other DRAM vendors announcing JEDEC-complaint parts, interest in SDRAMs waned substantially in 1994 and 1995. In fact, through 1995 the most successful SDRAM was the 4-Mbit ×16 SDRAM spearheaded by Hitachi. Despite the lack of JEDEC recognition, this part was very successful in the video game market because of its small memory granularity and substantial peak bandwidth.

Despite strong technical interest, the systems companies did not adopt the initial SDRAMs for several reasons:

1. High price premiums

   Despite repeated calls from system vendors for at most a 5% price premium, SDRAM price premiums remained at about 50% through 1994, declining to about 25% by the end of 1995. Only in the second half of 1995 did several vendors start to aggressively predict price premiums of 10% or less

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 235
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

in the medium term, starting as early as mid 1996. Several are now predicting no price premium over EDO DRAMs by mid 1997 or early 1998.

2. Significant incompatibilities among the early vendors

   A great many incompatibilities in both the logical and temporal interfaces existed among the early entrants into the SDRAM market. Initially, finding a group of vendors close enough in their specifications that a single controller could be designed to work at speed with more than one of them was difficult regardless of availability and pricing. This difficulty in ensuring multiple sources of supply discouraged several significant potential early adopters.

3. No DIMM standard at the outset

   Without a DIMM standard, early adopters would have had to design and produce their own DIMMs, only to be faced with a redesign once a DIMM standard was announced. It was mid 1994 before JEDEC ratified an 8-byte 200-pin SDRAM DIMM standard and early 1995 before it became clear that a 8-byte 168-pin DIMM standard would also be needed for compatibility with the fast-page-mode and EDO modules.

4. No high-speed electrical interface available

   As was discussed earlier, the plan at the outset was to use LVTTL for the lower clock frequencies, and a different electrical interface for higher clock frequencies. Without a credible high-speed interface emerging until mid 1995, all the performance-oriented early adopters would have been forced to use LVTTL and lower clock frequencies. Under these circumstances, much of the performance advantage over fast-page-mode and EDO alternatives was eroded.

5. No 64-Mbit SDRAM defined at the outset

   First previews of the 64-Mbit SDRAMs did not emerge until late 1994. Prior to this point, a memory controller designer had to either guess at the future SDRAMs' organization or forgo any future in-the-field memory upgrades.

6. Initial confusion between Samsung's level-RAS and JEDEC's pulsed-RAS parts

   This confusion was relatively short-lived. However, when both Samsung and NEC made their initial product and second-source announcements in March of 1993, it was unclear to many which, if either, the systems industry would adopt.

7. Lackluster support from the DRAM vendor community

   By the end of 1994, faced with an absence of orders, several DRAM vendors declined to move their initial SDRAM

Samsung Ex. 1097, p. 236
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

prototypes from sampling into production. This was a strong indication to the systems companies that SDRAMs were not yet ready for prime time.

Many problems have been solved

By the end of 1995, over two-and-one-half years after the first product announcements, many of these initial barriers to acceptance have been eliminated or significantly reduced:

1. Price premiums have declined, and several vendors are promising 5% price premiums over EDO by mid 1996. Many are predicting no price premium over EDO by 1997 or 1998.

2. Though it has yet to migrate from paper into silicon, a reasonably credible high-speed electrical interface (SSTL) is now in place.

3. The SDRAM DIMM standards are in place, though there are in fact two 8-byte standards, each with myriad variations, but no 4- or 2-byte standards.

4. The 64-Mbit SDRAMs are firming up, though both 2- and 4-bank versions are likely to be available, at least at the outset.

5. Samsung's level-RAS SDRAM has faded away.

Barriers to SDRAM acceptance remain

What remains is a shorter but still significant list of barriers to SDRAM acceptance:

1. The list of incompatibilities across vendors at the temporal level has been decreasing, but it remains significant. Most cycle-count differences have been eliminated, but some remain.

2. The standard does not adequately constrain the SDRAM design space. Thus, a controller that works with a collection of SDRAMs today may not necessarily work with future SDRAMs from these or other vendors.

3. Though Burst EDO and SDRAM-Lite are fading from contention, some confusion and uncertainty currently exists regarding the choice between SDRAM, SDRAM-Lite, and Burst EDO. Uncertainty and risk makes it difficult to commit to any alternative.

4. There is no central repository for information regarding incompatibilities and errata across the various SDRAM vendors. Each controller designer must cope individually with the various obstacles.

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 237
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

Several of these "barriers" existed in the generic DRAM marketplace for years. But with the additional complexity, higher frequencies, and greater incompatibilities among SDRAMs, the need to systematically address the issues that affect interoperability and reliable in-the-field memory upgrades have become more significant.

Samsung and NEC to produce compatible parts

Perhaps the most encouraging piece of recent news regarding SDRAMs was the January 1996 announcement that Samsung and NEC would work together to create completely plug-compatible SDRAMs with essentially identical temporal and logical interfaces. NEC and Samsung are the world's largest and third largest DRAM vendors. Their combined worldwide DRAM market share was 22% in 1994.[15] Together they probably have enough clout to define a de facto standard in those areas that JEDEC left unspecified. The major PC chip-set and system companies have also been working with the DRAM vendors to define a set of timing parameters that will work reliably and economically in PCs. Together, these two efforts will likely force the other DRAM vendors to follow suit and conform to a common temporal and logical specification. If this happens, then the road to a commodity SDRAM marketplace will have been substantially shortened.

7.4.7 Prices, Costs, and Design Wins

Early in 1993, interest in SDRAMs grew as NEC and Samsung began sampling their initial offerings. However, interest cooled during 1993 as memory system designers saw the initial price premiums. As discussed in depth in Chapter 11, system companies are very sensitive to DRAM prices. Many express the sentiment that even price per bit premiums of 5% are too high. The first SDRAMs began sampling with 50–100% price premiums. With premiums sitting at 25% or above, the volume ramp initially expected for 1995 failed to materialize. Consequently, several vendors declined to move their initial SDRAM designs into manufacturing. Now, with second and third generation 16-Mbit SDRAMs being completed, some vendors are aggressively promoting SDRAMs as a low premium (sub-10%) alternative for late 1996, and as a no-premium device in 1997 or 1998. Once a 5% price premium is reached, the market will overcome any remaining obstacles and begin to drive a substantial volume ramp. With a volume ramp beginning in earnest in 1997, it will likely be 1998 before the number of SDRAMs shipped is a

15. Source: Dataquest.

Samsung Ex. 1097, p. 238
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

reasonable fraction of the number of fast-page-mode and EDO DRAMs shipped. At that point, the 64-Mbit 2M×32 and 4M×16 devices are likely to be the most popular.

Diminishing die size penalties

The current SDRAM price premium stems from three sources: larger die sizes, lower cumulative volumes, and amortized engineering costs. The die size penalties for Synchronous DRAMs are about 15% at the 4-Mbit density. However, at the 16-Mbit density, vendors are claiming area penalties between 0 and 5% for Synchronous DRAMs compared with the generic counterparts. Generally speaking, the cost of a working die goes up with the square of its area, so these area penalties will translate into 0–10% cost penalties. Since the cost of an untested die is less than 50% of the overall cost of a salable DRAM, the area penalty for Synchronous DRAMs should not contribute significantly to the overall cost of manufacture or, ultimately, to the prices paid by end users. At the 64-Mbit level, the difference between the die sizes of comparable SDRAMs and EDO DRAMs will have vanished.

Reasonable test times

As mentioned earlier, test time is actually a more important component of the overall DRAM cost structure than the die size. Initially, many industry insiders believed that SDRAMs would incur a significant test time penalty relative to generic DRAMs. It was thought that dual-bank designs would pay an especially heavy penalty because of the need to ensure that operations in one bank did not affect the behavior of the other bank. Estimates of additional test time were used to justify the prediction of a long-term floor of 7–15% in the price premium for SDRAMs over generic DRAMs.

However, for several reasons, test times for SDRAMs can be comparable to those for comparable generic DRAMs:

- Test time is strongly dependent on the bandwidth at the chip interface. High-bandwidth interfaces yield shorter test times.
- Fewer asynchronous timings need to be tested.
- Waiting for the charge on the storage capacitors to decay constitutes a large portion of a DRAM test program. With a dual-bank architecture, one bank can be tested while waiting on the other—the tests of the two banks can be overlapped. The total test time can be less than the sum of the test times of the individual banks.

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 239
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

- The logic and memory portions of an SDRAM can be effectively tested separately.
- The run times of some of the classic DRAM tests grow with the square of the number of rows in the DRAM array. Splitting the DRAM into two banks can halve the time needed for these tests.

There is no strong technical reason for a dramatic cost difference between SDRAMs and generic DRAMs, especially at the 64-Mbit level and beyond. As with all DRAMs, the relative supply and demand will eventually determine the SDRAM price premium.

Now that the earliest SDRAM implementations are several years old, the first set of patents relating to SDRAMs are being issued. A list of recent patents relating to the new DRAM architectures is given in Appendix D. However, it is still too early to tell whether any company will develop a strong enough portfolio to extract significant royalties from the other players in this market, and, if so, whether the consumer's net price will be noticeably affected.

Design wins starting to ship

As of the end of 1995, the largest SDRAM design win shipping in volume is Sega's Saturn video game. This popular system, which began shipping in May 1995, uses four 256K×16 SDRAMs for part of the main CPU's memory, and another five SDRAMs in the graphics subsystem. In the same market domain, the reference design for 3DO's new game system is also SDRAM-based.

Sony has been shipping midrange MIPS-based workstations with SDRAM main memories since mid 1995. The NWS-4000 workstations operate 66-MHz (-15) SDRAMs at 44–50 MHz. Another pair of design wins are Fujitsu's SuperSPARC-based DS/90 7800 and DS/90 7900 servers. These systems use 66-MHz SDRAMs at 60 MHz. The memory system is two-way interleaved. In addition, Silicon Graphics Inc. has indicated that it intends to use SDRAMs in some of its future workstation products.

In November of 1995, Toshiba announced that their new high-end notebook computer—the Tecra 700—would offer an SDRAM main memory and no secondary cache. Toshiba previously used EDO memory in its high-end cacheless notebook computers. The Tecra started shipping late in 1995. Interestingly, all of the early advertisements for the Tecra 700 by systems retail stores

Samsung Ex. 1097, p. 240
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

completely ignore the novelty of its main memory. The ads note the size of main memory and generally obscure the absence of the secondary cache. However, none of them mention SDRAM at all. Initially at least, in the marketing of systems, SDRAMs are seen more as a liability than as an asset. Undoubtedly, part of the problem is the cost of memory upgrades. In early 1996, as the cost of memory for desktop systems was approaching $20 per megabyte and custom fast-page-mode SIMMs for popular notebook computers were available for about $35–40 per megabyte, memory upgrades for Tecra systems were selling for $70-80 per megabyte.

Several chip-set vendors in both the PowerPC and x86 camps have introduced SDRAM support in their latest offerings. A few of these are in production already, though as of 1Q1996, most are still sampling. However, since all of these parts are also capable of controlling EDO—and often Burst EDO—DRAMs, the SDRAM capability is not extraordinary. The advent of SDRAM-capable chip sets is a prerequisite to the adoption of SDRAMs by the PC system vendors.

7.4.8 Current Offerings

As mentioned previously, Samsung and NEC were first out of the chute—NEC with its JEDEC-compliant dual-bank design, and Samsung with its single-bank, level-RAS part. Samsung has since dropped its single-bank part. Since 1Q1993 most major DRAM vendors have developed SDRAMs; several are developing their first or second complete redesigns. Most of the initial designs were not productized.

Several of these companies are already sampling 64-Mbit SDRAMs. Most of the industry expects the price crossover from the 16-Mbit generation to the 64-Mbit devices to happen in 1998. However, some vendors are trying to drive the crossover as soon as possible—possibly as soon as mid 1997.

Hitachi, Hyundai, and Fujitsu are currently offering 4-Mbit SDRAMs with speed grades of 50–80 MHz, with 66 MHz being the most important. Hitachi is also developing a 2-Mbit SDRAM to further meet the high bandwidth and small memory size needs of the game and consumer arenas.

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 241
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

Synchronous Graphics RAMs

Synchronous Graphics RAMs (SGRAMs) are related to SDRAMs. SGRAMs are wide (×32) 8-Mbit SDRAMs targeted at video applications. They include block write and masked-write features to speed graphics rendering. As of 4Q1995, at least eight vendors plan to offer SGRAMs; several vendors are already sampling. By mid 1996, a full complement of SGRAMs should be available. However, since SGRAMs are fundamentally SDRAMs at heart, many of the interoperability issues found in SDRAMs are also evident in SGRAMs.

7.4.9 Strengths and Weaknesses

Strengths

- JEDEC standard
  - Likely to become a commodity part
  - Eventually many vendors with (nearly) interchangeable parts
- High bandwidths possible
  - Dual-bank organizations facilitate high bus utilization
- DIMM, 64-Mbit, and electrical interface plans in place
- All banks can be kept open simultaneously

Weaknesses

- Widespread minor differences among the vendors hamper multiple sourcing and reliable third-party in-the-field upgrades
- Row-cache contents do not persist after a bank precharge
- Does not inherently address memory system granularity
- Does not inherently address memory latency issues
- Confusion regarding SDRAM-Lite and Burst EDO

Samsung Ex. 1097, p. 242
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

7.5 Cached DRAMs (CDRAMs)

The third and most aggressive of the evolutionary DRAMs is the **Cached DRAM (CDRAM)**. The Cached DRAM design grew out of two observations:

- Small caches with relatively large block sizes have the same miss ratios as large caches with small block sizes.
- The bandwidth between the DRAM array and the I/O pads is greater than the pin bandwidth by several orders of magnitude.

Mitsubishi engineers concluded that by migrating a secondary cache onto the DRAM ICs, comparable performance could be obtained with a much smaller secondary cache—provided the block size was increased. The result is the Cached DRAM: a DRAM with a small embedded SRAM that functions as a secondary or tertiary cache.

CDRAM history

In 1992, Mitsubishi introduced the first CDRAM, a 1M×4 part; the part was ready for volume production in February 1993. In August 1993, Mitsubishi and Samsung announced plans to build interchangeable 256K×16 CDRAMs and to jointly specify parts at the 16-Mbit level. As of 4Q1993, Mitsubishi had de-emphasized the 1M×4 and a completed 4M×4 CDRAM. Mitsubishi is instead focusing on the 256K×16 at the 4-Mbit level and on a 1M×16 at the 16-Mbit level. Because the 1M×4 failed to gain any design wins, and has been withdrawn, the following discussion focuses on the newer 4- and 16-Mbit×16 parts.

Figure 7-22 shows a block diagram of the 256K×16 device. Conceptually, the CDRAM is organized very much like two separate devices—a 256K×16 DRAM and a 1K×16 SRAM—connected via a read path and a write path. Unlike other DRAM organizations, in which the sense amplifiers form a single level of caching in the form of a row cache, these CDRAMs have two levels of caching. The SRAM is typically used as a cache and the sense amps form yet another level of caching as a typical row cache. To disambiguate these, in this report the two caches are referred to as the SRAM cache and the row cache, respectively.

Clearly, the degree of change from the generic DRAM is substantial. Additional address and datapaths mean new physical, pin-level, and logical interfaces; the clock input implies a new temporal interface.

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 243
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

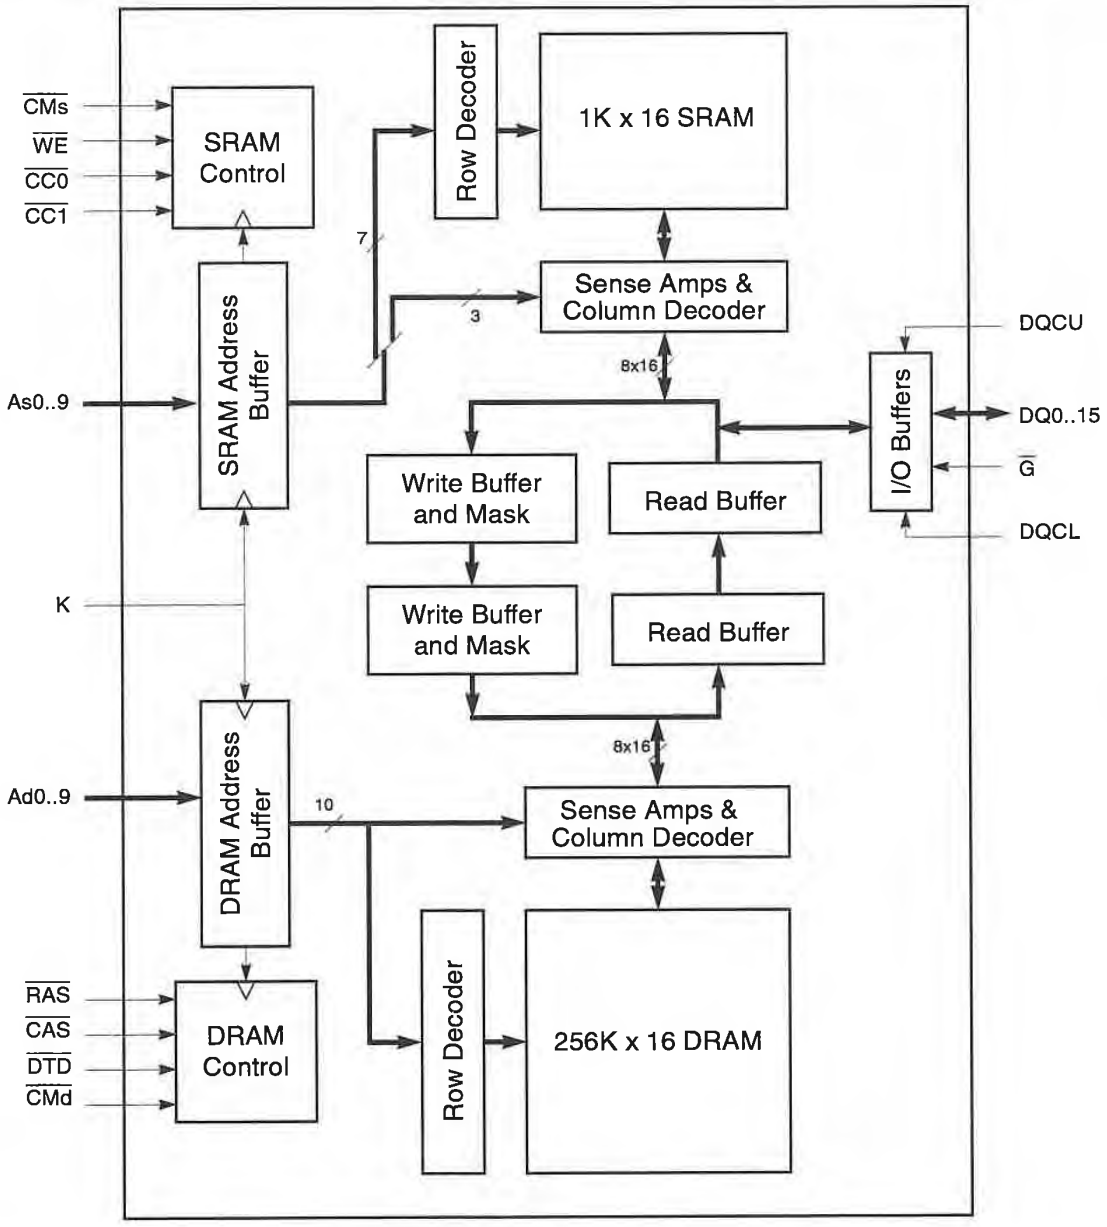


Figure 7-22 The second generation 256K×16 CDRAM has two write buffers and two read buffers between the SRAM and DRAM blocks.

Samsung Ex. 1097, p. 244
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

7.5.1 Physical, Pin-Level, Temporal, and Electrical Interfaces

The pin count for both the 256K×16 and the 1M×16 is 70 pins, but the part is in a 400-mil TSOP package with a 0.65-mm lead pitch. As a result, the total board area used per part is 0.43 square inches (2.77 cm^2), compared with 0.33 square inches (2.13 cm^2) for a generic 256K×16 DRAM in a 40-pin TSOP II package. Despite having 75% more pins, the CDRAM takes up only 26% more area than a generic 256K×16 DRAM. Furthermore, the wide CDRAM suffers no area penalty compared with the ×16 VRAMs and the other graphics-oriented devices.

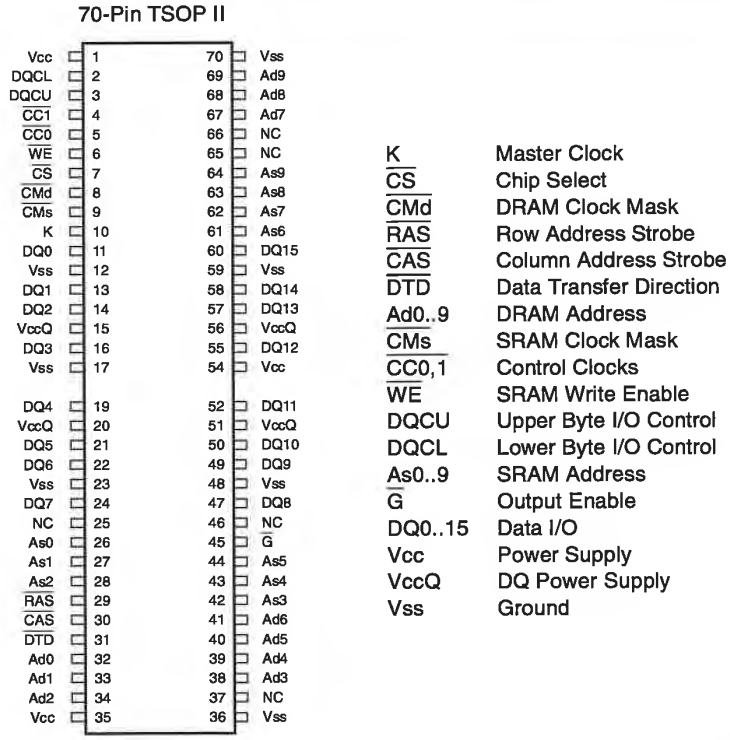


Figure 7-23 The ×16 CDRAM has a fairly large package despite the fine 0.65-mm lead pitch.

© 1996 Steven Przybylski

The new pins include a separate 10-bit non-multiplexed address bus for the SRAM, a total of 11 control signals, a clock input, and 14 supply pins.

Three output modes

The clock input is a fundamental part of the changes to the pin-level interface. As with Synchronous DRAMs, the input clock provides the reference for both the output of data and the input of addresses, commands, and data. The CDRAMs support three separate output configurations or modes, referred to as Transparent, Latched, and Registered (see Figure 7-24). The I/O configuration is selected by programming a control register within the CDRAM. Table 7-10 shows values for the setup, hold, and access times that define most of the temporal interface.

Parameter	Symbol	Minimum or Maximum			Units
		-12	-15	-20	
Clock Time	t_K	12	15	20	ns
Clock High Time	t_{KH}	4	4	4	ns
Clock Low TIme	t_{KL}	4	4	4	ns
Data, Address, and Control Signal Setup Time	t_{xxS}	4	4	4	ns
Data, Address, and Control Signal Hold Time	t_{xxH}	2	2	2	ns
Chip Select Setup Time	t_{CSS}	4	4	4	ns
Chip Select Hold Time	t_{CSH}	2	2	2	ns
Access Time from Clock High (All modes)	t_{KHA}	12	15	20	ns
Access Time from Clock Low (Latched mode)	t_{KLA}	9.5	12	16	ns
Access Time from Clock High (Registered mode)	t_{KHAR}	9.5	12	16	ns

Table 7-10 The second generation 4-Mbit CDRAM support clock supports speeds up to 83MHz. (Source: *M5M4V4169TP Preliminary Specification, Rev 2.0*).

The three output modes behave differently on a read. In Transparent Mode, the data flows out immediately as soon as it is available after the command that generated it. In Latched Mode, the event that triggers the output data is the falling edge of clock after the rising edge that latched the command. In Registered Mode, the next rising edge of clock triggers the driving of output data. In Latched and Registered Modes, the data comes out a little sooner relative to the triggering clock edge. As with the other evolutionary alternatives, the access time from the appropriate clock edge is equal to or only slightly less than the minimum (and expected) clock cycle time. At full speed, data arrives in the next clock cycle and is latched into the controller

Samsung Ex. 1097, p. 246
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

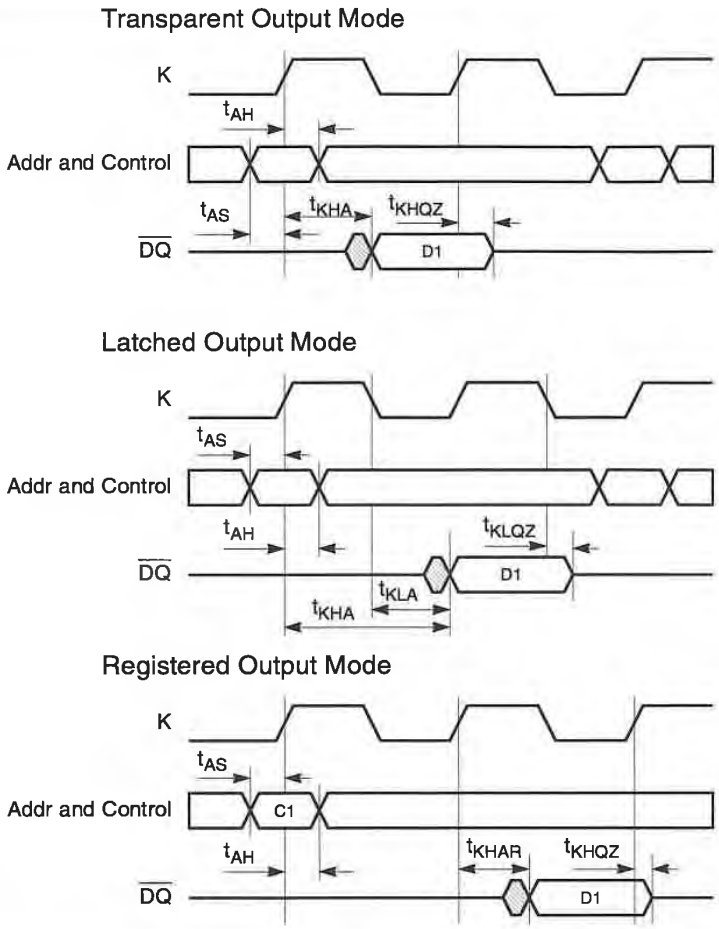


Figure 7-24 CDRAMs have three output modes that differ according to which clock edges control the data input and output.

either 1, 1.5, or 2 cycles after the edge that initiated the data transfer. The choice of the appropriate output mode depends on: operating frequency; device speed grade; loading on the address, control, and data signals; relative skew between the CDRAM's clock and the controller's; and clock edges being available within the controller for latching the data. Each of the three modes is appropriate in different circumstances. The registered output mode is closest to what is available in SDRAMs.

The CDRAMs are 3.3 V devices with full rail-to-rail LVTTL levels. This electrical interface is quite acceptable at operating

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 247
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

frequencies of 50- and 66-MHz, and remains marginally acceptable up to the 100-MHz level that the 16-Mbit CDRAM is capable of. After the DRAM industry finally settles on a high-speed interface for SDRAMs, the CDRAMs are likely to migrate to the new electrical interface as well.

The DRAM array at the heart of Mitsubishi's CDRAM is a standard 3.3V, 75-ns DRAM with a 135-ns cycle time. The SRAM access time varies from 12 ns in the fastest speed grade to 20 ns in the slowest. The DRAM access time varies correspondingly from 75 to 85 ns. The 16-Mbit CDRAMs, described in Section 7.5.3, are offered in -10, -12, and -15 speed grades with 10-ns, 12-ns, and 15-ns SRAM access and cycle times, respectively.

7.5.2 Logical Interface

The logical interface of first generation CDRAMs was highly constrained in some respects and very flexible in others. The second generation ×16 parts have a new logical interface that is very clean and orthogonal. Consequently, most of the constraints evident in the unsuccessful first generation are not present in the current devices.

The strength of the CDRAM architecture derives from its wide (128-bit) internal datapaths and ability to control the SRAM and DRAM independently. These capabilities allow for high bandwidth transfers between the two arrays and create opportunities for parallelism within the memory system. Table 7-11 lists the operations allowed. Essentially any SRAM operation can be paired with any DRAM operation.

Parallel operation of the SRAM and DRAM

Figure 7-25 shows some of the power of this interface: the figure shows a DRAM row access followed by a series of column accesses into the DRAM. These all occur in parallel with a series of otherwise independent SRAM reads and writes from the external interface. It also shows a write of the entire contents of the Read Buffer (128 bits) into the SRAM, and reads the Read Buffer directly to the I/O interface.

SRAM operations can occur on every cycle, and column operations can occur up to every other cycle when operating at full speed. The only potential conflict that might cause the two streams to interfere with each other is the need to write into the single-ported SRAM array. Overlapping SRAM and DRAM operations is especially handy if the SRAM is managed as a

Samsung Ex. 1097, p. 248
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

Memory	Opcode	Function	
None	NOP	No Operation	
	SCR	Set Command Register	
SRAM	SPD	SRAM Power Down	
	DES	Deselect SRAM	
	SR	SRAM Read	SRAM ⇒Dout
	SW	SRAM Write	Din ⇒ SRAM
	BRT	Buffer Read Transfer	RB ⇒ SRAM
	BWT	Buffer Write Transfer	SRAM ⇒ WB1
	BRTR	Buffer Read Transfer and Read	RB ⇒ SRAM ⇒ Dout
	BWTW	Buffer Write Transfer and Write	Din ⇒ SRAM ⇒ WB1
	BR	Buffer Read	RB ⇒ Dout
	BW	Buffer Write	Din ⇒ WB1
DRAM	DPD	DRAM Power Down	
	SRF	Self-Refresh	
	DRT	DRAM Read Transfer	DRAM ⇒ RB
	DWT1	DRAM Write Transfer 1	WB1 ⇒ WB2 ⇒ DRAM
	DWT1R	DRAM Write Transfer 1 and Read	WB1 ⇒ WB2 ⇒ DRAM ⇒ RB
	DWT2	DRAM Write Transfer 2	WB2 ⇒ DRAM
	DWT2R	DRAM Write Transfer 2 and Read	WB2 ⇒ DRAM ⇒ RB
	ACT	DRAM Activate (Row access)	
	PCG	DRAM Precharge	
	ARF	DRAM Auto-Refresh	
	DNOP	DRAM NOP	

Table 7-11 The second generation CDRAM has orthogonal SRAM and DRAM operations. (Source: *M5M4V4169TP Target Specification, Rev. 2.0*).

write-back cache: once a dirty line is transferred into the write buffers, the DRAM write can occur without impeding SRAM operations.

Since the DRAM is fundamentally an asynchronous device, the latency and repeat rates, when measured in cycles, vary according to the actual operating frequency. Just as in SDRAMs, the latency in the command register dictates when the read buffer is to be loaded with column read access data.

Not only can data transfer operations be overlapped, but DRAM precharges and refresh operations can also be hidden

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 249
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

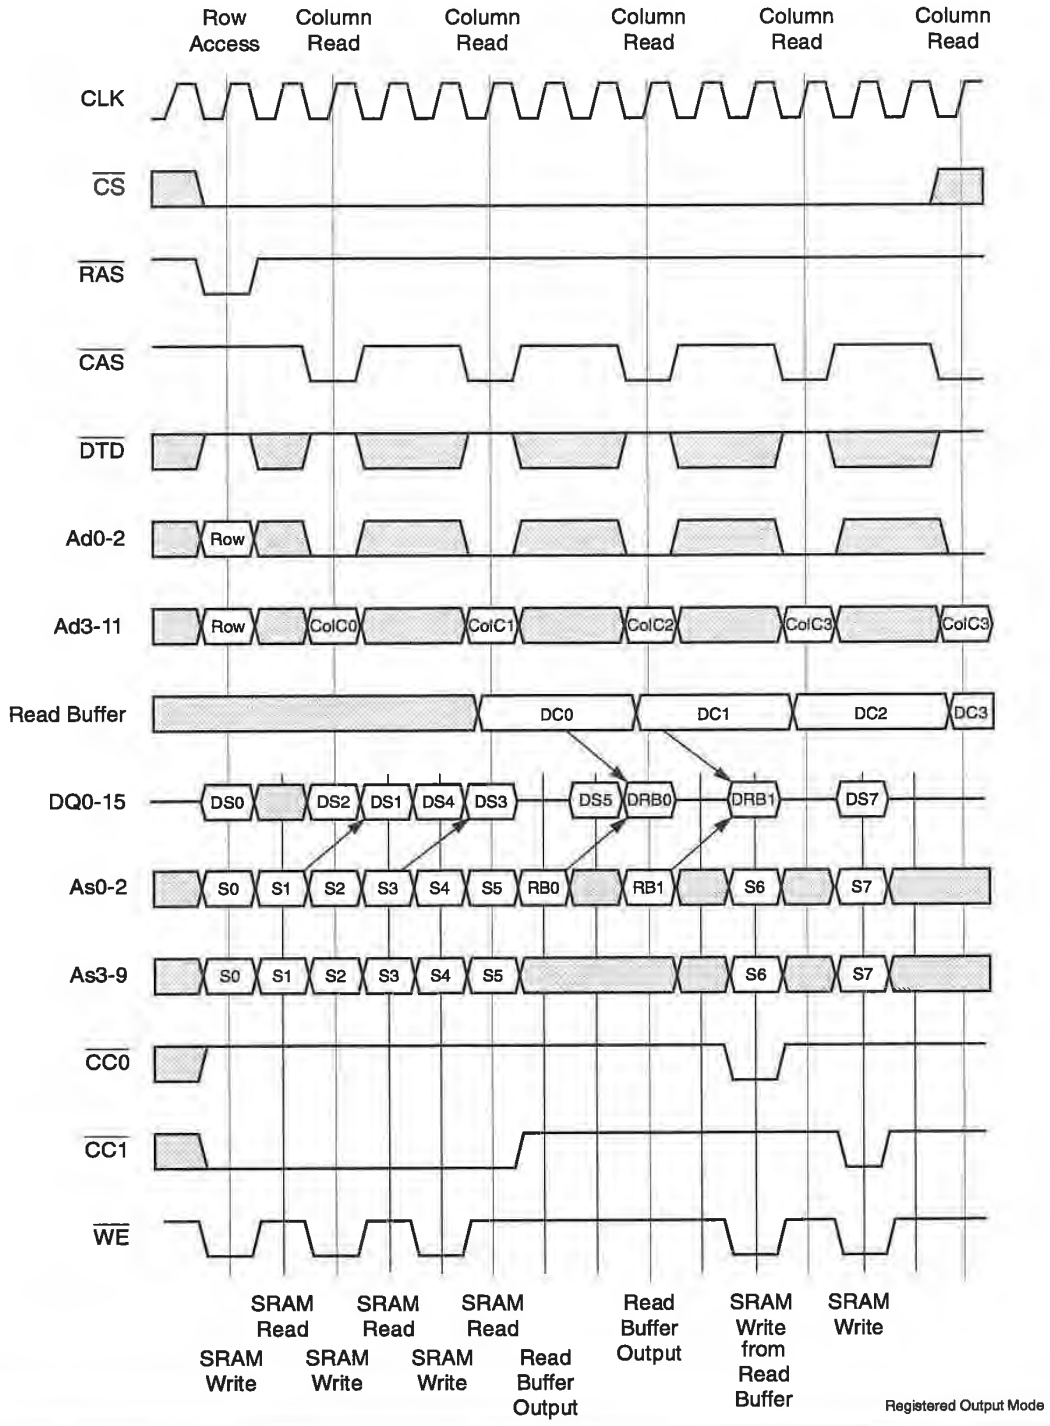


Figure 7-25 CDRAMs allow DRAM and SRAM operations to occur concurrently.

Samsung Ex. 1097, p. 250
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

underneath SRAM data transfers. Furthermore, each array can be independently powered down to minimize power consumption according to resources utilization. In addition, the DRAM has self-refresh capability when powered down.

The internal datapaths are organized as eight words of 16 bits each. Write masks associated with the write buffers allow for 16 bit writes from the I/O pins into the DRAM array. Writes from the SRAM into the DRAM are always in 128-bit chunks. The read buffers buffer DRAM data on its way to the SRAM or the I/O interface. The second level of the read buffer is an architecturally invisible staging register that allows the DRAM to start a new operation before the full latency, as determined by the command register, elapses.

The data-out portion of the logical interface is tailored through programming command registers. As mentioned before, the command register controls the nature and behavior of the data bus and selects the placement and behavior of registers and latches in the data-out path. The three choices are: Transparent-Output Mode; Latched-Output Mode; and Registered-Output Mode (see Figure 7-24). Also controlled by the command register are the SRAM read latency and the operation of the write masks associated with the write buffers (see Figure 7-22).

In systems with a 64-bit bus, four ×16 CDRAMs would typically be ganged together. In this case, the aggregate transfer size between the SRAM and DRAM would be a comfortable 64 bytes. The size of the row cache would be 1 Kbyte and the total size of the SRAM cache would be 8 Kbytes. Though very small compared to the typical external L2 caches constructed of discrete SRAMs, the CDRAM's SRAM cache can be quite effective due to the high bandwidth and fairly low latency transfers between it and the DRAM.

SRAM as a local memory or cache

The SRAM in the CDRAM need not be used as a traditional cache. Because all control of this traditionally organized SRAM array is left up to the memory controller, the array can be used as a cache or local memory. As a cache, it can be direct-mapped, set-associative, or fully associative. The block and fetch size of the SRAM cache can be any multiple of 16 bytes per CDRAM.

This flexibility in the SRAM cache organization constitutes both a burden and a blessing. On the burden side, the SRAM constitutes only the data part of a cache. The missing tag storage and cache control must be housed in the memory controller. Fairly complicated control logic is needed to make the most of the

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 251
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

CDRAM's features, and the controller must contain enough tags for the largest possible number of memory banks. For small, fixed-sized memories such as frame buffers, tag storage is quite tolerable: 128 tags of one byte each. However, for large, expandable main memory applications, the amount of tag storage needed in the controller could become significant, especially if the memory controller is to be implemented in an FPGA or a small ASIC. However, if the CDRAM is being used to eliminate an SRAM-based secondary cache, the cost reductions make the cost increment in the controller insignificant.

On the positive side, all small-associativity caches are equally realizable, so the memory designer is able to tailor the design to the market and the expected workload. On the negative side, the memory controller must include the tags and control.

7.5.3 The 16-Mbit CDRAM

The 16-Mbit CDRAM has many similarities and a few significant differences. Perhaps most significantly, the organization of the device is a 1M×16 device with a 16-Kbit SRAM: compared to the 4-Mbit CDRAM, the device width and SRAM size stayed the same while the DRAM size quadrupled. This was undoubtedly done to keep the die size penalty as low as possible.

The package and electrical interface remained the same. The three output modes were retained, as was the 70-pin TSOP II package. The pinouts of the two devices are identical except for three additions to pins that were No Connects at the 4-Mbit level: additions were two new DRAM address bits and one new control signal, $\overline{\text{BTE}}$, for enabling burst transfers into and out of the SRAM without requiring a new SRAM address each cycle.

All commands listed in Table 7-11 remain the same and no new ones were added. The power-down modes and self-refresh remain.

It is important to note that the new DRAM address bits are both row addresses. The number of column addresses (seven) remains the same. This maintains maximum compatibility with the 4-Mbit CDRAM. The drawback is that the row cache remains the same 256 bytes per device. Also, if four 4-Mbit CDRAMs are to be replaced by one 16-Mbit CDRAM, the total SRAM size is cut by a factor of four.

Samsung Ex. 1097, p. 252
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

Faster speed grades

New to the 16-Mbit generation is the -10 speed bin, which supports maximum clock rates of 100 MHz. The temporal characteristics of the DRAM array have been improved to be comparable with most 16-Mbit DRAMs: the -10 part corresponds to a -6 fast-page-mode device, while the -12, -15, and -20 correspond to -7, -7, and -8 parts, respectively. Consequently, given the same speed bin and operating frequency, some DRAM operations complete one cycle sooner than with the 4-Mbit CDRAM. The data output times for the -10 part are 10 ns in Transparent Mode and 8 ns in Latched and Register Modes. The very tight data output times for all speed grades mean that it will be difficult to make CDRAMs to run at their maximum rated frequencies in anything but small, nonsocketed memory systems.

The new burst mode allows for automatic transfers of bursts of length one, two, four, or eight in either linear or sub-block order.[16]

7.5.4 CDRAM Fill Frequencies

The CDRAM's peak fill frequencies are calculated differently from those of the other DRAMs. This is because a couple of different memory sizes and peak bandwidths must be considered. For example, the SRAM is 2 Kbytes, and the DRAM is either 0.5 or 2 Mbytes. The peak bandwidth into the device is 166 Mbytes/s (2 bytes times 83 MHz) for the 4-Mbit device and 200 Mbytes/s (2 bytes times 100 MHz) for the 16-Mbit CDRAM. The peak bandwidth for the path between the SRAM and the DRAM is 666 Mbytes/s (16 bytes times 41.5 MHz) for the 4-Mbit device, and 800 Mbytes (16 bytes times 50 MHz) for the 16-Mbit device.

Based on the DRAM array size, these bandwidths translate into the fill frequencies shown in Figure 7-26.[17] The peak bandwidths between the SRAM and DRAM are very high though unsustainable over long periods because they rely on data being in the SRAM to start with. More realistic peak fill frequencies are those based on the bandwidth through the I/O interface and the DRAM size: 333 Hz for the -12 4-Mbit device[18] and 100 Hz for the -10 16-Mbit device. Both of these peak fill frequencies

16. Sub-block ordered transfers are also called Intel-ordered or interleaved transfers.

17. Figure 7-26 shows the 4-Mbit CDRAM becoming available in 1994. However, the faster speed grade that was initially available was the -15 part, not the -12 on which the fill frequencies are based. Consequently, the peak fill frequencies shown for the 4-Mbit CDRAM are 133 Hz and 533 Hz.

18. Or, as shown in Figure 7-26, 266 Hz for the -15 4-Mbit CDRAM.

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 253
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

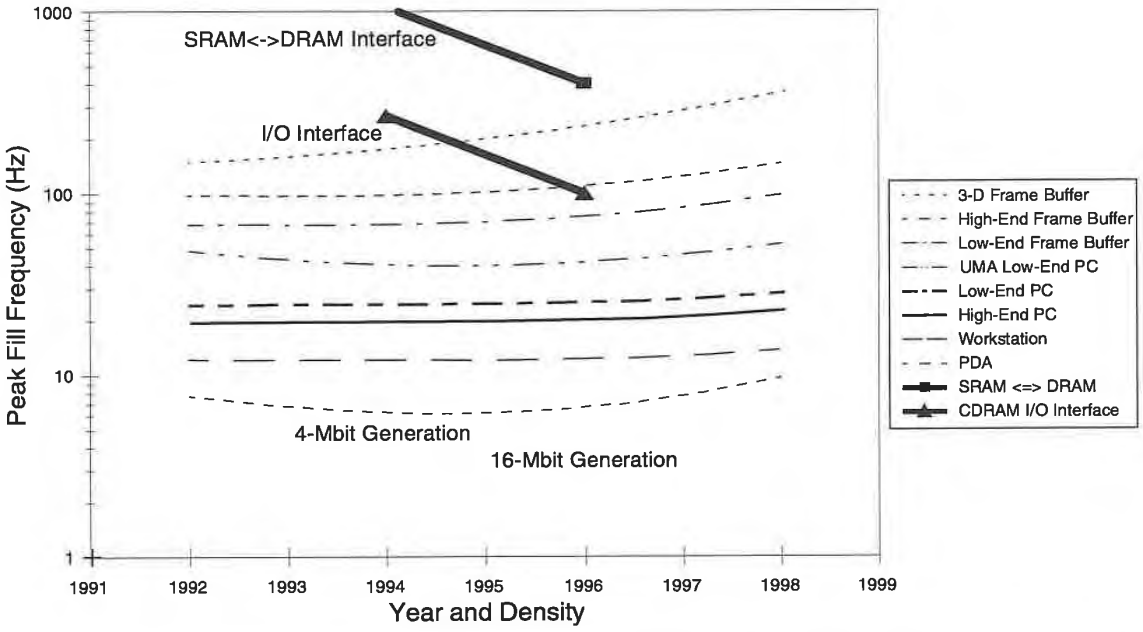


Figure 7-26 The ×16 CDRAMs have high peak fill rates at both the DRAM and pin-level interfaces.

are substantially higher than those evident among any of the conventional alternatives.

From the perspective of performance and fill frequencies, the main benefit of the split SRAM/DRAM interface is to increase the sustainable and sustained fill frequencies. The fact that the external data bus can be 100% utilized transferring data even in the face of a steady stream of row-cache misses means that for most memory organizations and burst sizes the worst case sustainable fill frequency is very high: it is essentially equal to the best case sustainable fill frequency. Both the best case and worst case sustainable fill frequencies are equal to the peak I/O interface fill frequency.[19]

For many high-performance embedded, video, and DSP applications, the sustained bandwidth delivered by conventional DRAMs is significantly below the peak bandwidth because of poor row-cache miss rates. The flexibility inherent in the

19. Please see Section 4.2.3 for an explanation of the subtle differences between peak, best and worst case sustainable and sustained bandwidth and fill frequency.

Samsung Ex. 1097, p. 254
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

CDRAM interface facilitates very high memory bus utilizations, and consequently high sustained bandwidth, even when continually performing row accesses. These applications are characterized by high bandwidth requirements and multiple data streams with good locality of reference. When these characteristics are present, most references can be satisfied out of the SRAM instead of having to continually wait on DRAM row accesses.

7.5.5 Miscellanea

Power consumption

CDRAMs conceptually and physically contain both an SRAM and a DRAM. Thus, the power they draw is equivalent to the sum of the power consumption of a conventional SRAM and a DRAM. The 256K×16 part typically draws 540 mW in the SRAM and 300 mW in the DRAM: these amounts are fairly typical of the discrete devices they replace.

Prototype system demonstrated

In order to demonstrate the performance of CDRAM systems, Mitsubishi constructed a prototype PC with a single-bank memory system using first generation CDRAMs and an Intel 486DX2-66 CPU with no secondary cache. When compared with a variety of cached and uncached systems with the same processor, the CDRAM system matched or bettered the performance of systems with 128-Kbyte and 256-Kbyte secondary caches.[20]

Fifteen percent price premium

In the first generation CDRAM, the die size penalty associated with adding the 16-Kbit SRAM to a 4-Mbit DRAM is roughly 7%. The ongoing size penalty, the extra testing time, and the low volumes translate into a cost premium to manufacture CDRAMs. As of 4Q1995, Mitsubishi was quoting CDRAMs at a fixed 15% price per bit premium over fast-page-mode DRAMs. For many memory consumers, particularly in cost-sensitive markets, this premium disqualifies CDRAMs from consideration. In other applications, the need for a specialized controller is a serious problem.

For other applications, a 15% price premium is tolerable if adopting CDRAMs eliminates some expensive RAMs, such as VRAM in a frame buffer or SRAM in a secondary cache. Mitsubishi strongly advocates a unified memory organization for a wide range of PCs and workstations. Such a memory system, built in whole or in part from CDRAMs, contains both main

20. C. Hart, "CDRAM in a Unified Memory Architecture," in *Proceedings of Spring 1994 Compcon*, IEEE Computer Society Press, 1994.

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 255
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

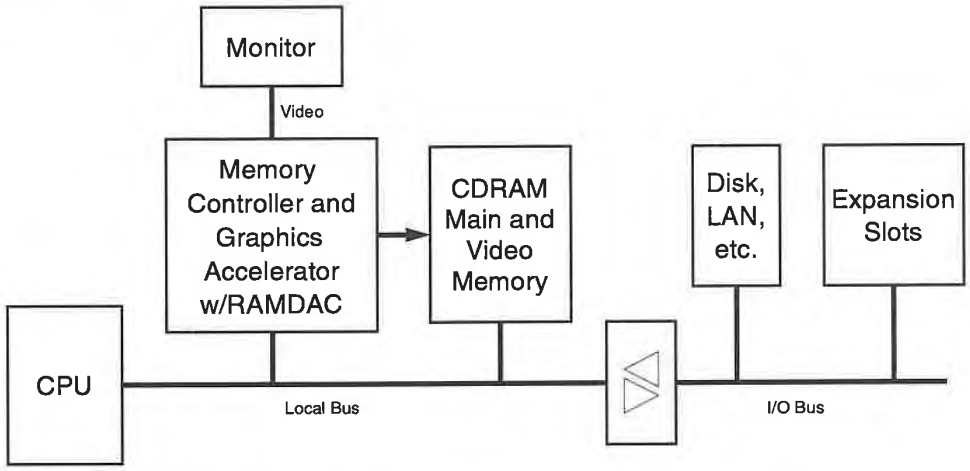


Figure 7-27 With a unified memory architecture, the main and video memories are combined.

memory and a video frame buffer (see Figure 7-27 and Section 4.8). CDRAMs provide enough bandwidth to efficiently satisfy all the CPU, graphics, video, and I/O references.

Multiple sources are available

Mitsubishi's second-source agreement with Samsung focuses on the second generation 256K×16 CDRAM. Samsung began sampling this device in early 1995, while Mitsubishi has had it in production since mid 1994. Both vendors are moving forward on the design of the 1M×16 organizations. Mitsubishi sampled this device in 1Q1995. Unlike the first generation CDRAM, the improved second generation devices have met with some success in the marketplace. Mitsubishi notes several design wins in both the arcade video game arena and as a hard disk drive (HDD) disk cache.

Though these design wins are undoubtedly adequate to justify the migration of the existing CDRAMs into production, they do not move CDRAMs out of the niche device category. Unless the DRAM marketplace changes substantially and CDRAMs find their way into some significant PC, consumer, or set-top applications, total market share will remain less than a few percent indefinitely.

Samsung Ex. 1097, p. 256
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

7.5.6 Strengths and Weaknesses

Strengths

- Large SRAM for low cache miss rates
- Flexible and efficient dual interface
- High peak bandwidths and fill frequencies
  - Good opportunities for overlapping SRAM and DRAM operations
  - High data bus utilization
- Low latency for high performance without a discrete SRAM secondary cache

Weaknesses

- Significant cost-per-bit premium
- Memory controller must implement cache tags and cache control
- Larger package results from two address buses
- Only two sources of supply

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 258
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

8 Revolutionary Alternatives

Since 1974, the essential characteristics of the DRAM interface have been the separate data bus and the multiplexed address bus. All of the conventional and evolutionary alternatives described in Chapters 6 and 7 share these elements. However, three recent architectures—Rambus, RamLink, and SyncLink—discard the conventional pin-level interface for a byte-serial bus that transfers control, addresses, and data. SyncLink, a derivative of RamLink, includes one high-speed bus for address, control, and write data, and a second bus for read data. Since these new interfaces represent such a dramatic break from the traditional interface, these three organizations are revolutionary in nature.

8.1 New In The Second Edition

Additions to Section 8.2 on Rambus DRAMs include increased coverage of the 8- and 16-Mbit RDRAMs, a description of the characteristics of the primary initial Rambus design wins, and a discussion of changes to the Rambus organization and protocol that are to be included in the 64-Mbit generation. Additional material includes discussion of the Y Channel and Rambus's recent design wins. The focus of the discussion is now on the 8- and 16-Mbit RDRAMs, rather than on the original 4.5-Mbit device.

Section 8.3 discusses the final form of the pending RamLink standard. Section 8.4 describes SyncLink, a derivative of

Samsung Ex. 1097, p. 259
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

RamLink, which is still being defined as of 4Q1995. SyncLink's technical characteristics and the structure of the industry consortium promoting it are discussed in detail.

Peak fill frequencies have been added as appropriate.

8.2 Rambus DRAMs (RDRAMs)

As the number of DRAMs per system declines (see Figure 2-10), there is a growing need for high-bandwidth memory systems with fewer—even just a single—DRAM. Generic and conventional DRAMs fail to meet the needs of these systems because the maximum bandwidth that can be generated by a single device is too limited. Adequate bandwidth can be generated only by ganging many DRAMs in parallel, but this is done at the expense of a correspondingly larger minimum memory size.

Today, many graphics subsystems fall into this category. A single 16-Mbit DRAM can meet the capacity needs of most PC GUI accelerators. In the not-too-distant future, most PCs' main memories will be constructed of four or fewer 64-Mbit DRAMs. Increasingly, the need is for new memory architectures that provide high bandwidth from a small number of DRAMs.

Rambus Inc. is a company founded in 1990 to develop and promote such a new high-bandwidth DRAM interface. The heart of its interface's organization is the Rambus Channel: a byte-wide, 500- or 533-Mbytes/s bidirectional bus connecting a memory controller with a collection of **Rambus DRAMs (RDRAMs)**.

Rambus Inc. has fully specified all aspects of a Rambus-based memory system: the temporal and logical behavior of RDRAMs; the physical and electrical characteristics of the Rambus Channel; and all aspects of its SIMM equivalent (RModules and RSockets). Rethinking all aspects of the DRAM interface, Rambus engineers (REngineers?) developed a unique solution that facilitates high-bandwidth memory systems with as few as one RDRAM. This solution involves significant changes to each of the DRAM interface's five aspects (see Section 5.2). The main fetaures of the Rambus Channel and of RDRAM's:

- Peak bandwidth of 500 or 533 Mbytes/s per channel, regardless of the number of DRAMs[1]
- Byte-wide organization with 8- or 9-bit bytes

Samsung Ex. 1097, p. 260
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

- Interoperable 8-, 16-, 18-, and 64-Mbit RDRAMs
- Vertically or horizontally mounted DRAMs
- Physically constrained bidirectional bus using terminated surface-trace transmission lines
- Expandability via transceivered or untransceivered modules
- Up to 22 RDRAMs or six modules per Rambus Channel at 533 Mbytes/s[2]
- Reduced voltage-swing, current-mode signaling
- Synchronous, burst-oriented protocol
- One bank per Mbyte in 8- and 16-Mbit RDRAMs[3]
- Large row caches

RDRAMs were first introduced at the 4-Mbit level with a 4.5-Mbit device that was sampled by Toshiba in 3Q1993. The second generation 16- and 18-Mbit devices were first sampled by NEC in 2Q1994 and Toshiba in 3Q1994. These devices and their derivative, the 8-Mbit RDRAM, have become the mainstream Rambus devices—the 4-Mbit device failed to generate any design wins, and was not moved into production. Consequently, this section focuses on the characteristics of the 16-Mbit RDRAMs; a historical footnote in Section 8.2.6 describes the original 4.5-Mbit devices. Changes that Rambus is introducing at the 64-Mbit level are presented in Section 8.2.7.

8.2.1 Pin-Level and Physical Interfaces

Typically, the pin-level and physical interfaces define the form and package of an integrated circuit without reference to the systems in which the IC is used. Rambus is unique: it is both a DRAM specification and a channel specification. The physical and electrical properties of both the RDRAMs and the bus on which they are placed are rigorously defined because high-frequency operation relies as much on careful physical design of the printed circuit board as it does on high-speed circuits.

1. RDRAMs capable of 533 Mbytes/s should become available starting in 1H1996. Higher bandwidth speed bins will be available in the future.
2. At 533 Mbytes/s, the maximum number of RDRAMs on the channel drops to about 24.
3. Future 64-Mbit RDRAMs will have one bank per 2 Mbytes.

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 261
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

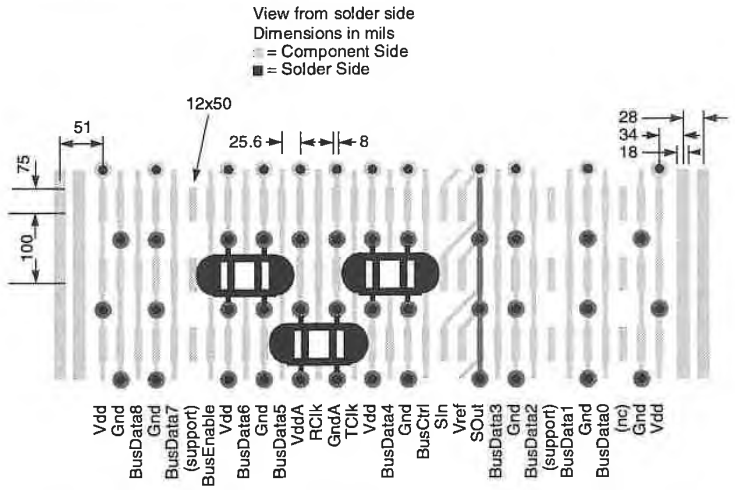


Figure 8-2 The physical layout of the Rambus Channel is well defined to ensure good signal quality.

The Rambus Channel

The Rambus Channel consists of 15 signals—13 high-speed signals and 2 low-speed signals. With five power pins, eight ground pins, and a reference voltage, RDRAMs fit comfortably into a 32-pin Surface Horizontal Package (SHP) or a standard 32-pin Surface Vertical Package (SVP) (see Figures 5-1 and 8-1). Since the pin-level interface is entirely independent of DRAM density, the same packages are used for all 8-, 16-, and 64-Mbit RDRAMs.[4]

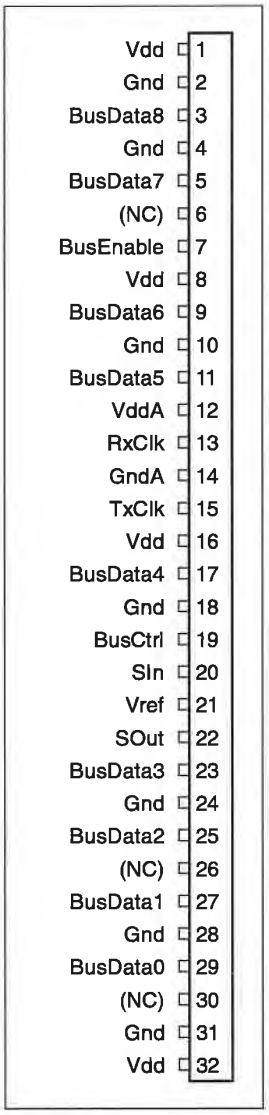
A Rambus Channel is physically limited to about 10 cm by the propagation delay of signals and the 2-ns bit time on the bus. To reach 533 Mbytes/s, the bit time on the bus is reduced to 1.87 ns, and the maximum channel length is reduced to 9.38 cm. The Rambus Channel master is the memory controller at one end of the bus. The other end of the bus is terminated. The RDRAMs are Rambus Channel slaves, responding to requests from the memory controller.

The Rambus physical specification dictates the placement of the RDRAMs and decoupling capacitors on the board (see Figure 8-2). In order to maintain a consistent signal propagation speed across the channel, the individual traces must be as

4. Technically, a master need not drive the Sin and Sout chain to initialize the DRAMs on a channel. Consequently, the number of pins needed on a master is 31 pins per channel.

Samsung Ex. 1097, p. 262
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

32-Pin SVP Pinout



Pin Descriptions

Signal	I/O	Description
BusData[8-0]	I/O	Signal lines for request, write data, and read data packets. A Request Packet contains the address, operation codes, and count of the bytes to be transferred. This is a low-swing, active-low signal referenced to Vref.
RxClk	I	Receive clock. Incoming Request Packets and write Data Packets are aligned to this clock. This is a low-swing, active-low signal referenced to Vref.
TxClk	I	Transmit clock. Outgoing Acknowledge and Read Data Packets are aligned with this clock. This is a low-swing, active-low signal referenced to Vref.
Vref	I	Logic threshold voltage for low-swing signals.
BusCtrl	I/O	Control signal for framing packets, transmitting part of the operation code, and acknowledging requests. Low-swing, active-low signal referenced to Vref.
BusEnable	I	Control signal to enable the bus. Long assertions of this signal will reset all devices on the channel. It is also used for Serial Address Packets. This is a low-swing, active-low signal referenced to Vref.
Vdd, VddA		+5V power supply. VddA is a separate analog supply.
Gnd, GndA		Circuit ground. GndA is a separate analog ground.
SIn	I	Initialization daisy chain input. TTL levels. Active high.
SOut	O	Initialization daisy chain output. TTL levels. Active high.

Figure 8-1 All RDRAMs are packaged in 32-pin SVP packages. The Rambus Channel includes 13 high-speed signals.

© 1996 Steven Przybylski

identical as possible. This is easily achieved though since the channel is confined to a single layer of a board and the pin assignment and pitch are identical in all RDRAMs. Though the variation across the Rambus Channel needs to be well controlled, the physical characteristics of the channel as a whole need only be controlled to ±20%. Though the Rambus Channel was designed with printed circuit boards of four or more layers in mind, two-level boards can be used, though with a severely reduced channel length.

The main packaging option for RDRAMs is the **Surface Horizontal Package (SHP)**. The SHP package is basically a 70-pin 0.65-mm pin-pitch TSOP II package with pins on only one edge. The opposite edge has some support pins but only at its extremes. The SHP package takes up about 13 mm of channel, allowing for at most 11 such devices per Rambus Channel at 500 Mbytes/s. At 533 Mbytes/s, 10 SHP packages can be accommodated. The SHP package is especially useful for applications that are vertically constrained and those that require only a few RDRAMs per channel. As of 4Q1995, not all Rambus vendor literature includes the SHP package, though eventually all vendors are offering RDRAMs in both packages.

The Surface Vertical Package (SVP) packages are mounted vertically, as close as every 100 mils (2.54 mm). Consequently, up to 24 SVP-packaged RDRAMs can be accommodated on a single Rambus Channel at 500 Mbytes/s, and up to 22 SVP-packages RDRAMs at 533 Mbytes/s.

Although the SHP package is more wasteful of channel length, most early applications of Rambus's technology involve only a small total memory size per channel, with a few packages soldered down and a single socket for memory expansion. As of 1Q1996, all announced design wins use SHP packages.

Channel and RDRAM width

Note that the Rambus Channel is actually nine bits wide. This width reflects the ongoing need to facilitate parity or ECC protection. For lower cost systems and for systems that can tolerate today's low soft error rates, the channel protocol allows ×8 RDRAMs to coexist on a channel with the standard ×9 parts. In recent years, parity protection has begun to disappear in PCs due to both cost pressures and the recognition that software failures are much more common than soft memory failures (see Section 4.2.5). Consequently, the ×8 RDRAMs will ultimately be more popular than the ×9 parts. Though frame buffers have never been parity protected, some early users of RDRAMs in graphics applications are using the ×9 parts instead of the ×8.

Samsung Ex. 1097, p. 264
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

The extra width facilitates either 18-bit pixels, an easy way to differentiate between RGB and YUV encoded pixels, or to otherwise tag pixels as special.

Socketed modules for memory expansion

Rambus Inc. has defined two sets of memory expansion modules (RModules) and matching sockets (RSockets). Vertical RModules extend vertically above the channel. Since vertical RSockets each occupy 12.3 mm of channel, the maximum memory system per Rambus Channel consists of eight modules of 24 RDRAMs each, or 192 memory chips. At 533 Mbytes/s, the maximum memory system consists of six Vertical RSockets populated with RModules with 22 SVP RDRAMs each, for a total of 132 devices. In either case, even with 2-Mbyte RDRAMs, the maximum memory size is substantially more than anyone is likely to put on a single channel.

A second type of expansion module is the Horizontal RModule and its matching Horizontal RSocket. As would be expected, in this case, the module is parallel to and on top of the main Rambus Channel. The main advantage of the horizontal RSocket is that a normal PC board with a horizontally mounted RModule fits within the 0.8" board-to-board spacing of PC add-in cards. Its disadvantage is that a much smaller number of sockets can be placed on a channel. Up to four Horizontal RSockets can fit on a channel, and Horizontal RModules can have up to eight SHP RDRAMs on them. So in a low-profile configuration, the maximum memory system consists of 32 RDRAMs at either 500 or 533 Mbytes/s. Again, since most of the early adopters of the Rambus technology at the 16-Mbit level are developing graphics and multimedia solutions, a maximum memory size of 32 RDRAMs is much greater than the one to four RDRAMs typically present on a channel in these systems.

Figure 8-3 shows an example of a memory system that has four RDRAMs and a socket with a four-RDRAM Vertical RModule. All sockets, horizontal and vertical, are keyed to prevent insertion of future low-voltage RModules into higher voltage sockets. A channel populated only with modules is less heavily loaded than a channel populated with RDRAMs. The lower capacitance translates into a higher signal propagation speed and consequently a slightly longer maximum length.

RTransceivers

The Rambus Channel that the memory controller is on is dubbed the Primary Channel. Some RModules include a Rambus Transceiver (RTransceiver) to isolate the Primary Channel on the motherboard from the Secondary Channel on the RModule. An RTransceiver which acts as a slave on the primary

© 1996 Steven Przybylski

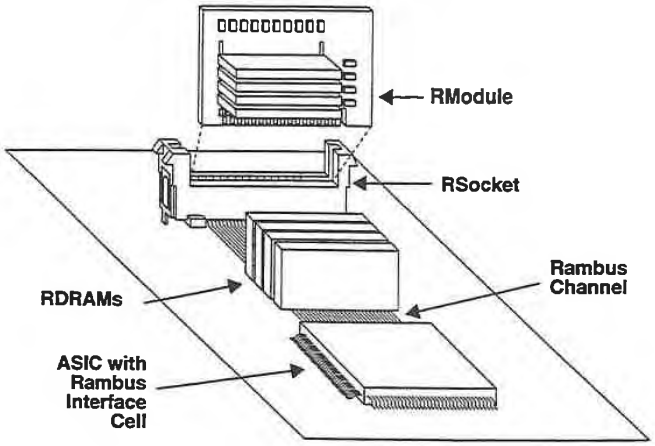


Figure 8-3 RDRAMs can be soldered directly onto the Rambus Channel or placed on expansion RModules.

channel and as a master on the secondary channel (see Figure 8-4). The transceivers will be in 42-pin SVP packages; this organization minimizes wasted space on the RModules. Transceivers are not allowed as slaves on the secondary channels.

RModules with transceivers and a secondary channel are called Secondary RModules. To maintain the electrical properties of the primary channel when no secondary RModule is installed, the socket must be filled with an unpopulated Null RModule. RModules without transceivers can be placed only at the end of the primary channel and must include the bus termination. These are called Primary RModules because they extend the primary channel. When a Primary RModule is not installed, the socket must be filled with a Terminator, which provides bus termination and clock routing.

A third option for expansion involves the use of a Y Channel, as shown in Figure 8-4. A Y Channel involves having two complete channels joined at their master ends. This joining affects the impedance of the channel and complicates clock distribution (discussed in Sections 8.2.2) but facilitates a one-time upgrade without the need to have a socket populated with a Null RMod-

Samsung Ex. 1097, p. 266
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

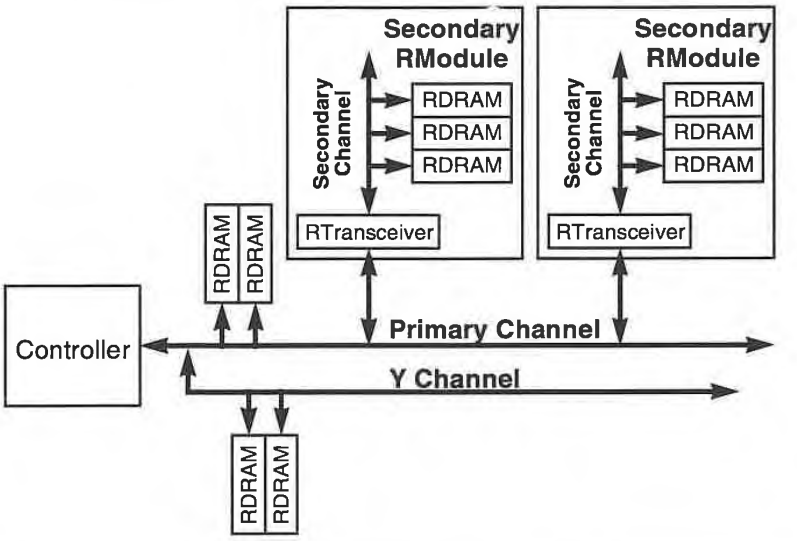


Figure 8-4 RTransceivers bridge the module's secondary channels to the PC board's primary channel. Y Channels add additional flexibility.

ule prior to the upgrade. An empty Horizontal RSocket close to the master does not create a significant impedance discontinuity that would impact signal propagation. When the RSocket is populated with a Primary RModule, the Y Channel is formed and the combined channels' operating characteristics are changed. All initial Rambus design wins use either a Y Channel or a Primary RModule as the expansion memory to a soldered-down base memory constructed of SHP RDRAMs.

8.2.2 Electrical Interface

High-frequency operation requires not only careful physical design but a robust electrical interface. In the case of Rambus, the electrical interface is called Rambus Signaling Logic (RSL). RSL defines the voltage levels on the Rambus Channel and ensures the interoperability of RDRAMs with different supply voltages. RSL differs from most high-speed signaling standards in that, under normal circumstances, the system designer can ignore the details of the electrical interface. If only RSL-compliant parts are placed on a properly laid-out Rambus Channel, then signal quality and reliable operation are guaranteed. RSL-compliant parts include RDRAMs, ASICs incorporating a

© 1996 Steven Przybylski

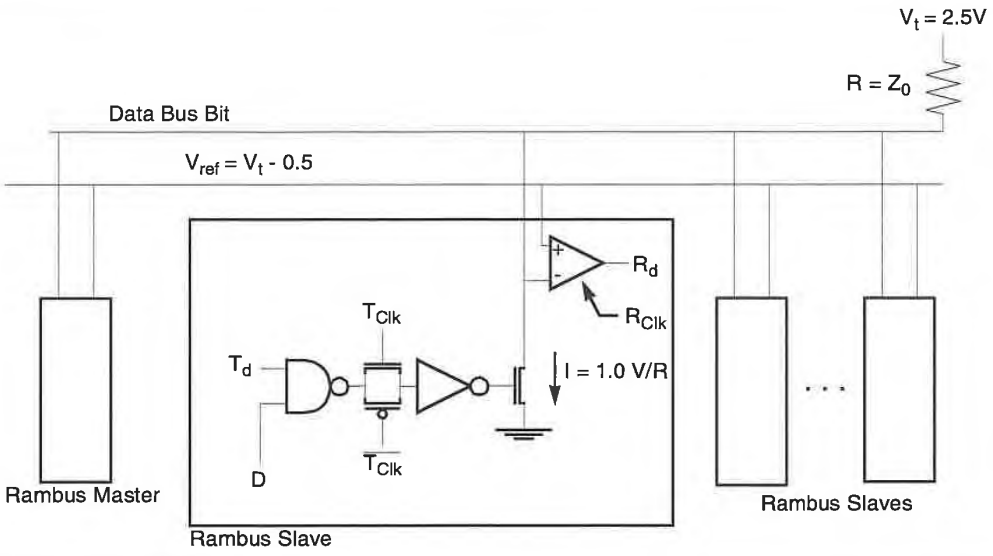


Figure 8-5 Rambus Signaling Logic relies on low swings and terminated transmission lines.

Rambus ASIC Cell (RAC), and processors or other devices with a Rambus interface built in.

RSL depends on low signal swings and uniform, terminated transmission lines to maintain exceptional signal quality at high speeds. Figure 8-5 shows the basic circuit structure of the bus. Signals are terminated to approximately 2.5 V at one end of the bus. The other end, with the bus master, is left unterminated. A reference voltage approximately 300 mV below the termination voltage is generated by a simple resistor network on the board.

A logic 0 is represented by the termination voltage of about 2.5 V. A logic 1 is represented by about 1.5 V. When a master or a slave wants to drive a logic 1, it sinks enough current from the bus to pull the bus down to 1.5 V. When a device is "driving" a logic 0, or when no device is driving, the bus is pulled up to a logic 0 by the terminating resistor. The total signal swing is roughly 1 V, centered around an externally generated reference voltage. The minimum allowed swing, as defined by $V_{IL,Max}$ and $V_{IH,Min}$, is 700 mV. The maximum allowed swing is 1.6 V. The variability in the levels and swing allows for imprecise setting of the termination and reference voltages and for variability across the channel.

Samsung Ex. 1097, p. 268
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

A singly terminated, bidirectional bus

The bus is terminated only at the non-master end. Slaves initiate half-swing signals that propagate down the bus in both directions. When the signal reaches the unterminated master end, it reflects, reaches full amplitude, and propagates back to the terminated end. Because the Rambus Channel is a **current-mode bus**, one device can start driving before the previous driver's signal has completely propagated off the end of the bus. In contrast, with **voltage-mode buses**, overlapping signals clip and interfere with one other. With voltage-mode buses, the transition from one bus driver to another must be separated in time by at least two bus propagation times. With current-mode buses, this requirement does not exist.

To allow for variation in the characteristic impedance of the transmission lines, the current that a slave must sink to drive a logic 1 is programmable at system initialization to between 4 and 36 mA.[5] Although 36 mA multiplied by 10 output drivers is a substantial amount of current to sink per DRAM, the Rambus DRAM package provides enough ground pins to keep the on-chip ground bounce tolerable. Also, differential input receivers that compare the input signal to an externally generated reference voltage are less susceptible to ground bounce than are the single-ended receivers used in generic DRAMs.

The output drive current is set independently for each slave and master device on the channel at power up. For each device in turn, the master programs greater and greater drive currents until the slave is able to pull down the bus to V_{REF}. By doubling the value that generates a swing from V_T to V_{REF}, a symmetrical waveform is generated around V_{REF}.

The variability in the output drive capabilities of the slaves and masters means that loaded trace impedances between 20 – 80 Ω are supported. This range is crucial for a number of reasons. First, it facilitates poor tolerance in the trace characteristic impedance from board to board. Second, it allows the use of two-layer PC boards with their higher characteristic impedances. Third, it permits the use of Y Channels (see Section 8.2.1 and Figure 8-4). With a Y Channel in place, the master and all the slaves see a characteristic impedance that is half the trace impedance and the termination impedance. If the nominal trace impedance is 50 Ω, then the characteristic impedance when a Y Channel is present is 25 Ω—still within the allowed range of 20–80 Ω.

5. The maximum value of I_{OL} at the minimum programmed setting is 4 mA, while the minimum value for I_{OL} at the maximum programmed current setting is 36 mA.

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 269
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

As the channel's characteristic impedance decreases, the current needed to drive a valid level increases. Consequently, power consumption while driving data increases somewhat for all devices in a Y-Channel configuration.

Because voltage levels are set relative to the reference voltage, which is allowed to vary between 1.9 V and V_{DD}-0.8 V, the supply voltages of Rambus masters and slaves have no impact on the interoperability of 5 V and 3.3 V parts. In fact, the signaling voltage levels will support a further decline in the RDRAM supply voltage to 2.5 V: this shift is expected at the 64- or 256-Mbit level. RSockets and RModules are physically keyed to prevent the accidental insertion of a 3.3 V module into a socket supplying 5 V to Vdd. For now, however, all RDRAMs in production are 3.3 V devices and no 5 V devices are planned, so interoperability is not an important feature until the 2.5 V parts are introduced.

The very high signal quality needed for 250- or 266-MHz operation is possible because each trace has a very uniform characteristic impedance. Generally, uniformity of the traces' impedance is ensured by uniform and compact placement of the RDRAMs. In cases where the loading is less than perfectly uniform or the RDRAMs are more widely separated than is desirable, a uniform characteristic impedance can be created by manipulating the trace widths appropriately.

It is also important to make sure that the terminating resistors match the impedance of the transmission line. A mismatch between the two causes reflections and reduces the noise margin on the bus. The physical layer specification of multilayer PC board layout requires that each termination resistor matches its transmission line's characteristic impedance to within 10%.

Low EMI generation

The electromagnetic interference (EMI) generated by an electronics system is related to the rise times of the signals, their voltage swings, and their traces' effectiveness as antennas. This latter influence is determined by the length and layout of the traces and their coupling to a ground plane. As ringing inherently has substantial high-frequency components, good signal integrity is fundamental to reducing EMI radiation (see Section 4.3.1).

Despite its very high operating frequency, a Rambus Channel's short channel length, low voltage swing, good signal integrity, and effective coupling to a ground plane make it a relatively low-intensity source of electromagnetic interference. A Rambus memory system typically generates significantly less EMI noise

Samsung Ex. 1097, p. 270
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

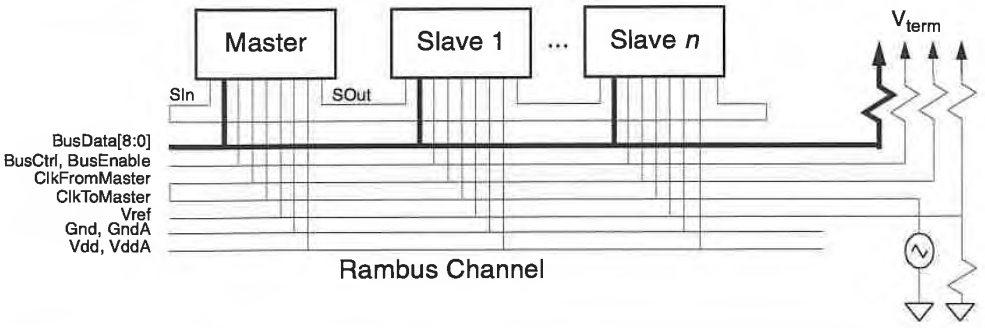


Figure 8-6 Data always flows with a clock, either toward or away from the master at one end of the channel.

than a comparable bandwidth memory system constructed of a conventional or evolutionary alternative.

8.2.3 Temporal Interface

The Rambus Channel is a synchronous bus running at 250 or 266 MHz, with data transferred on each edge of the clock. Each bit exists on the bus for a scant 2 ns. One of the most difficult aspects of high-frequency digital system design is the management of clock skew—the differing phase relationship between a global clock and data at different places in the system. Clock skew interacts with time-of-flight delays to complicate the flow of data traveling in one direction or the other on a bus.

Eliminating clock skew

The Rambus Channel elegantly solves the problem of clock skew: data and a clock always travel together down the bus. If a Rambus device transmits the data with a fixed-phase relationship to a clock that travels in parallel with the data, the phase relationship is maintained at the receiving device. The only clock-to-data skew at the receiver comes from the fractionally different propagation speeds of the individual bus traces. By specifying the maximum deviation of both the bus's physical characteristics and the RDRAM's load on the bus, the Rambus specification minimizes skew across the entire length of the bus.

This solution is facilitated by the fact that all transactions are either to or from the master at one end of the bus. ClkFromMaster is used as the reference clock for transmissions from the master to a slave; ClkToMaster is used for responses from a slave to the master. Figure 8-6 illustrates how the two clocks are really the

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 271
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

same signal. It starts at the terminated end of the bus, travels to the master end, then doubles back to the far end of the bus, where it too is terminated. This configuration yields the convenient side effect that the two clocks are identical for the master.

Rise and fall times on the channel are nominally 400 ps, though they are allowed to be as fast as 300 ps. The designers of RDRAMs and RACs (see Section 8.2.5) must take great care in order to meet the 350-ps data setup and hold times.[6] Even more challenging are the transmit windows. At 250 MHz, data bits have to be driven on the bus for 2 ns, with the leading and trailing edge of each bit positioned to within 300 ps. Sophisticated digital Delay-Locked Loops or Phase-Locked Loops[7] are required to maintain these tight timing constraints with a minimum of phase jitter.[8]

Because data is transmitted on both edges of the clock, bit positions are labeled either even or odd. Packets and data bursts always begin on the falling edge of the appropriate clock with an even bit. Groups of four clock periods are used to transmit 8 bytes of data, called an **octbyte**.

Fortunately, all of the difficult engineering associated with the 266-MHz operating frequency is encapsulated in the design of the RDRAMs and RAC ASIC megacells: most system designers never need to concern themselves with these aspects provided they follow the physical Rambus Channel layout rules.

Y Channel poses special challenges

The Y Channel shown in Figure 8-4 illustrates the important role that clock skew has high-speed system design. A challenge in designing the Y Channel was arranging for the concurrent flow of clock and data down either primary Rambus Channel. One possible alternative was for the master to have two clock inputs, one for tapping the clock loop on each channel. This alternative significantly complicates the design of the master since the master needs to send and receive data in synchrony with two different, out-of-phase, clocks.

The other basic alternative, and the one chosen by Rambus, is to have separate clock loops for the two channels, but to

6. The setup and hold times are specified to be at least 0.35 times the actual clock period, which is nominally 4 ns for 500 Mbytes/s and 3.75 ns for 533 Mbytes/s

7. The original 4-Mbit RDRAM design used a phase-locked loop. All the 8-, 16-, and 18-Mbit devices use a delay-locked loop.

8. M.Horowitz *et al.*, "PLL Design for a 500 MB/s Interface," *ISSCC Digest of Technical Papers*, February 1993, pp. 160–161; and T. Lee *et al.*, "A 2.5 V Delay-Locked Loop for an 18Mb 500MB/s DRAM," *ISSCC Digest of Technical Papers*, February 1994, pp. 300–301.

Samsung Ex. 1097, p. 272
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

fine-tune them so they are in phase with each other at their master ends. If this phase relationship is maintained, then the master sends and receives data identically to and from slaves on either the primary or the Y Channel.

Maintaining a phase relationship at the master means ensuring that trace lengths of the two clock paths from a common source are nearly identical. Any phase mismatch between the two clocks eats directly into the setup and hold times and begins to limit the operating frequency of the channels. Since the total setup and hold specifications are 35% of the quarter-cycle time (or 350 ps each at 500 Mbytes/s), it is important to keep the skew between the two clock loops to less than 50 ps. Since signals travel at around 160 ps/inch on unloaded surface traces of FR4 PC boards, each tenth of an inch difference in the path lengths of the two channels' clocks translates into about 16 ps of skew. These lengths have to be measured from a common source to their respective master ends. If the loading of either channel is variable, then maintaining the skew is more difficult because changes in loading of the clock (and other) traces changes both the propagation delay of the traces and the effective skew. Consequently, systems designed to accept a Y Channel are designed to accept Primary RModules with a specific clock travel time from the connector to the far end of the channel and down the channel back to the connector.

Consequences of speed binning

The original Rambus concept was to have a single-speed device. All Rambus Channels were to operate at 500 Mbytes/s and all RDRAMs of each generation would have the same access times. This strategy had both important benefits and serious consequences.

The foremost advantage was that any memory controller could know the detailed speed characteristics of the device simply by querying the DeviceType register to determine the RDRAM's size and indexing through a small, fixed, look-up table. The unfortunate disadvantage is that the standard speed had to be set at a relatively slow level so that there would be essentially no yield loss due to speed. The need to eliminate yield drop-out due to low-speed parts is responsible for much of the long latencies described in Section 8.2.4.

The way things have developed, some of the detailed speed characteristics of the devices vary across manufacturers. And now, with the introduction of both 500 and 533 MBytes/s devices, the clean specification of the temporal interface is further complicated. The master is unable to determine whether a

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 273
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

particular RDRAM came from one frequency bin or another. Furthermore, if the RDRAM vendors begin binning based on latency as well as frequency, a variety of RDRAMs with identical logical characteristics but slightly different temporal characteristics will appear. If this occurs, the memory controller will have few resources to figure out the actual temporal characteristics of the devices it is controlling. Unfortunately, this represents no real change from the current confused and error-prone situation with generic, conventional, and evolutionary DRAMs.

Physical and electrical interface independent of logical interface

The physical and electrical specifications of the Rambus Channel are very closely coupled to produce an easy and reliable high-speed interconnect. However, little interdependence exists between these two aspects of the complete specification and the logical interface that defines the protocol by which the masters communicate with the RDRAMs. In fact, the physical and electrical layers can support a wide variety of point-to-point and single-master communication protocols. This capability has been used in Silicon Graphics Inc.'s IMPACT workstation. Two ASICs in this workstation's graphics subsystem communicate over a dedicated Rambus Channel using a custom protocol.

8.2.4 Logical Interface

The Rambus logical interface is defined in terms of the protocol on the Rambus Channel and the response of the RDRAMs to the various message types. The basic model for an RDRAM is shown in Figure 8-7. It illustrates that RDRAMs are effectively a ×64 or ×72 DRAM with a high-speed 8:1 multiplexor/demultiplexor on the front end. All operations and data transfers occur in groups of eight bytes temporally aligned for easy shunting onto the correct byte lane of the internal ×72 organization.

RDRAMs divided into two banks

The main functional blocks are the Rambus interface with the associated Delay-Locked Loops; the control unit with its configuration registers; and the two DRAM banks, each with its own **row cache**.

Bigger and better row cache

In the case of the 18-Mbit RDRAMs, each of the DRAM banks consists of 9 Mbits and is buffered by 18 Kbits of row cache built into the sense amplifiers. As in the generic DRAMs, the row cache is direct-mapped with a very long block size. However, in RDRAMs, there are two blocks per DRAM instead of one block as in the generic designs. Also, RDRAMs offer substantially more row cache than generic DRAMs. RDRAMs have 1 row-cache bit for every 512 DRAM bits instead of the 1 row-cache bit

Samsung Ex. 1097, p. 274
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

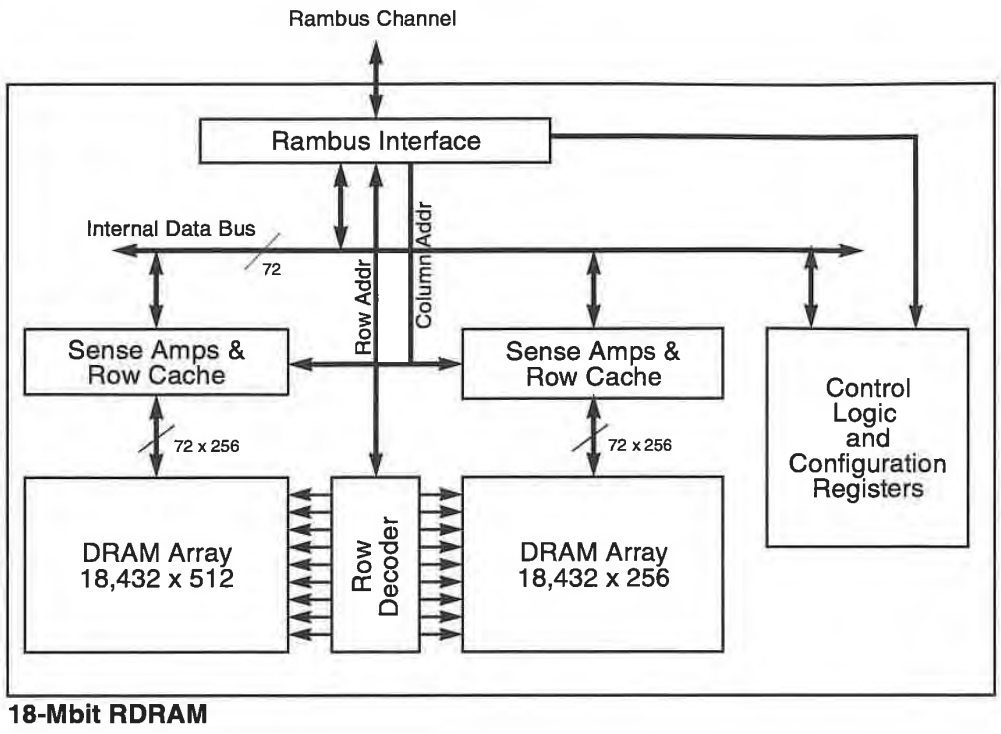


Figure 8-7 The 18-Mbit Rambus DRAMs have two independent banks of 9 Mbits. Each bank has a row cache of 18 Kbits.

for every 2,048 bits offered by a generic 4M×4 DRAM and the 1,024 bits offered by a generic 1M×16 DRAM. In the case of the 16-Mbit RDRAMs, each of the two 8-Mbit banks has a 16-Kbit row cache.

Not only do RDRAMs have proportionately more row cache, but that row cache is better organized than the row caches of most conventional and evolutionary alternatives. For example, a ×64 memory system built of 1M×16 EDO DRAMs has one bank per 8 Mbytes. In contrast, a memory system of 8- or 16-Mbit RDRAMs has eight banks per 8 Mbytes. The beneficial impact of having more banks—or alternatively, more blocks in the row cache—is increased in larger memory systems. The memory controllers in PC chip sets in particular can only keep one bank of EDO memory active at a time, whereas in an RDRAM memory system, all banks can be active simultaneously. All the bits in the row cache can hold data at the same time.

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 275
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

An 8-Mbit RDRAMs is logically half of a 16-Mbit device. The 8-Mbit RDRAMs include a single bank and a single line of row cache. Any future 9-Mbit devices would be single-bank devices logically equivalent to half of an 18-Mbit RDRAM.

Table 8-1 lists the many RDRAM configuration registers and identifies their functions. To summarize, registers identify the RDRAM's size; type; manufacturer; base address; delay and access intervals; and refresh and operational modes.

Register	Description
DeviceType	Read-only register that provides type, size, and version information
DeviceID	Specifies the base address of the RDRAM
Delay	Specifies the column access delays
Mode	Includes device-enable, refresh mode, and pull-down current control
RefRow	Self-refresh row and bank information
RASInterval	Specifies the row access timing parameters
MinInterval	Read-only register providing the minimum timing parameters
AddressSelect	Defines the address bit swapping in use
DeviceManufacturer	Contains a manufacturer code and part type code
Row	Contains the row address of the row-cache contents.

Table 8-1 RDRAMs contain 10 configuration registers that provide information about the device and determine much of its behavior.

A backplane bus on a board

In many ways, the Rambus Channel is more akin to a backplane bus than to the generic DRAM interface. The master and the slaves communicate via groups of bytes or packets that contain address and control information (Request Packets) or data (Data Packets). Rambus slaves monitor the bus for Request Packets. When a slave recognizes a request that addresses its memory space, the slave responds at a prespecified time according to the function specified in the request. The time at which it is to respond is specified in the RDRAM's configuration registers.

Figure 8-8 shows the format of a Request Packet. A Request packet contains 36 bits of byte address, 8 bits of byte count, and 6 bits of opcode. The 6 bytes of information occupy the bus for 12 ns on 500 Mbytes/s channels. RDRAMs support block reads and writes to the row cache, simple register reads and writes, and a broadcast register write.

Samsung Ex. 1097, p. 276
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

	Bus Data								
Bus Cntl	**8**	**7**	**6**	**5**	**4**	**3**	**2**	**1**	**0**
Start	Op0	Addr9..2							
Op1	Op3	Addr17..10							
OpX1	Addr26..18								
Op2	Addr35..27								
OpX2	Rsrv		Count6,4,2			Rsrv			
Close/Rsrv	Rsrv		Count7,5,3			Count1,0		Addr1,0	

Figure 8-8 A Rambus Request Packet takes 12 ns on the bus.

The most noteworthy aspect of the Request Packet is the unusual placement of the bits within the various bytes. Several reasons exist for this particular placement of the fields. First, it is important to keep related information physically adjacent within the DRAMs. For example, by interleaving count bits the way they are, a full 8-bit count field can be constructed within the RDRAM without having to bus any bits across a large distance. A second reason for the unusual field placement is history: first-generation 4.5-Mbit RDRAMs did not support OpX[2..1], Addr[1..0], or Count[1..0].[9]

Once a slave recognizes a read or write to its address space, its first job is to acknowledge receipt of the request. This is done with a 2-bit Acknowledge Packet. The slave places this packet on the BusCntl line at a fixed time after the end of the Request Packet. As with most of the fixed delays inherent in the protocol, the duration of the delay is specified in the Delay register. Some of the delay values must be programmed identically in all slaves, while others can vary within established windows from device to device on one Rambus Channel.

The responding slave has two choices with respect to the Acknowledge Packet: it can positively (Ack) or negatively (Nack) acknowledge the request. An Ack tells the master that the slave is currently able to handle the request. In the case of a read, the data follows forthwith. A Nack tells the master that the slave is currently unable to complete the request, and that the master should reissue the request later.

9. See Section 8.2.6 for further discussion of 4-Mbit RDRAMs.

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 277
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

Two reasons to Nack a request

A slave may Nack a request for two reasons:

- If an RDRAM is currently performing some other operation, then it is unable to handle the new request. In some circumstances this would be indicative of a controller error.
- If the row cache does not currently contain the addressed data, then the RDRAM must perform a row access before the request can be completed.

The RDRAM is responsible for bringing the referenced data into the appropriate row cache as soon as possible. This is done with the expectation that the Nacked request will be repeated shortly. Notice that the data read and write operations are implicitly to row caches only. Data to be read or written must be moved to the row cache via a row access before the read or write successfully completes.

When the master receives a Nack to a request, it can either continue trying the request until it is eventually Acked or wait the expected time for the RDRAM to perform a row access before retrying the request. Any pending requests to other RDRAMs on the channel can be performed while waiting for the first RDRAM to complete the necessary row access.

After a block-read Request packet is Acked, a Data Packet is sent from the RDRAM back to the master. Data Packets contain an integral number of **octbytes**,[10] ranging in total from 8 to 256 bytes. Here, a byte is defined as eight or nine bits. Since all Data Packets are an integral number of octbytes long, the low-order three bits of the address and count fields must be zero for all memory read requests.

For write commands, byte addressability and byte granularity of transfers are provided for. The slaves can infer from the low-order three bits of the address and count fields which bytes in the first and last octbytes are real data. The bytes in those octbytes that are not data are empty placeholders needed to maintain a consistent temporal framework. This scheme simplifies the design of a Rambus interface by ensuring that bytes with comparable low-order address bits share the same position within their respective octbytes. Block transfers are not allowed to cross row-cache boundaries, which occur at 2-Kbyte intervals in the 8-, 16-, 18-, and forthcoming 64-Mbit devices.

10. An *octbyte* is eight contiguous bytes.

Samsung Ex. 1097, p. 278
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

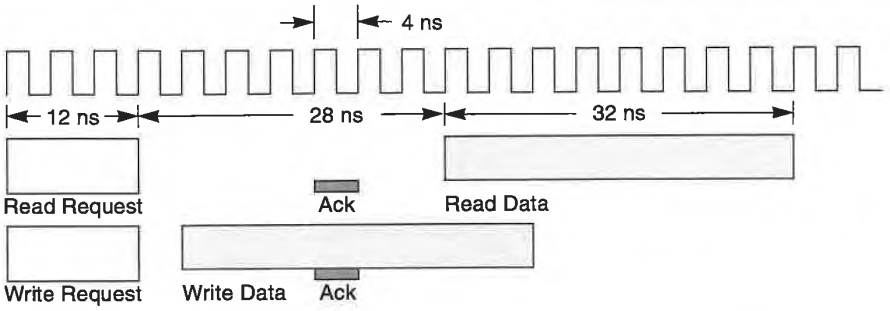


Figure 8-9 With 16-Mbit RDRAMs, row-cache hit data starts 28 ns after the end of the request.

The use of the ninth bit is left unspecified by the protocol. Though intended primarily for parity or ECC, the ninth bit could also be used for data, tags, directories, or prefetch hints. The Rambus protocol does not define any additional facilities for error checking or correcting. If the system architect feels it necessary, the integrity of the memory system as a whole can be ensured by generating and checking parity or ECC as data comes into the memory controller.

High bandwidth but poor read miss latency

Figure 8-9 shows the relative placement of the Request, Acknowledge, and Data Packets for typical 16-Mbit generation devices operating at 500 Mbytes/s. For simple block write transactions, the data follows shortly after the Request Packet, before the acknowledge is received. In the case of a read request that hits in the row cache, the data begins coming back to the master quite quickly—28 ns after the end of the Request Packet. Given that data transfers at 500 Mbytes/s, a 16-byte read takes a total of 72 ns. A 32-byte transfer takes 104 ns. At 533 Mbytes/s, the cycle counts remain the same, yielding a latency to the start of data transfer of 26 ns, and total access times of 67.5 ns and 97.6 ns respectively for 16- and 32-byte transfers, respectively.

In the case of a row miss (see Figure 8-10), the first read request is Nacked. If the referenced bank has a row in the row cache (i.e., has an open row), then the RDRAM begins a precharge and a row access to bring the referenced address into the row cache. A repeat request will be Acked any time after 76 ns from the end of the first Request Packet. The data begins flowing 28 ns after the end of the second Request Packet, so the overall access time

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 279
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

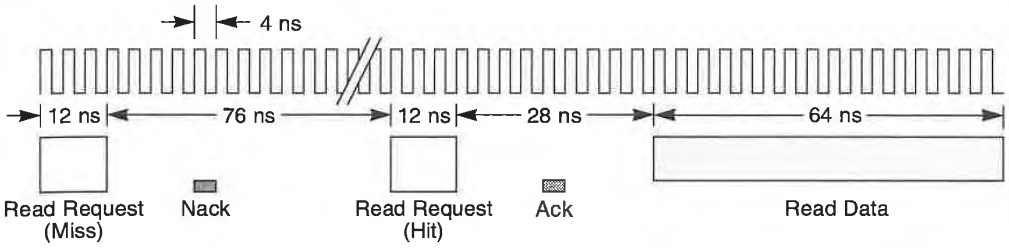


Figure 8-10 In the case of a row-cache miss, two read transactions are needed to retrieve the data.

is 160 ns for a 16-byte access, or 192 ns for a 32-byte transfer. Again, when operating at 533 Mbytes/s, the cycle counts remain the same, yielding total access times of 150 ns and 180 ns for 16- and 32-byte transfers, respectively.

In generic DRAMs, the column access is allowed to begin just before the row access is complete. In other words, the column address decode begins before the sense amplifiers have completely stabilized at their final voltage levels. In RDRAMs, the same thing happens despite the two separate Request packets. If the second Request packet arrives as soon as is allowed, then it is accepted and column decode begins even though the row access is not technically complete.

Long row-cache miss latency appears to be the most troublesome characteristic of RDRAMs. However, comparisons at the memory-system level show that RDRAM row-cache miss latencies are comparable to those of the other alternatives (see Chapter 10).

Furthermore, RDRAMs benefit from low row-cache miss rates: low miss rates reduce the average memory latency to the first word of data. Although the RDRAM's row cache is better organized than most of the alternative DRAMs' row caches, in some systems and applications the RDRAM's row cache will not work well. In these cases the average memory access time and the sustained bandwidth can fall significantly if the average transfer size is small. Examples include: 1) systems without any form of caching between the CPU's primary cache and the memory system and 2) systems with a single RDRAM. The RDRAM row

Samsung Ex. 1097, p. 280
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

cache is ineffective in the second system type because references to one RDRAM bank are not allowed at the same time that the other bank is performing a row or column access. This exclusion makes it impossible to overlap row-cache misses with other activity. In any case, Rambus DRAMs are most effective when transfer sizes are large enough that overall access time is dominated by the transfer period and not by the initial latency.

Eleven timing parameters set in registers

In total, eleven timing parameters are specified in three of the configuration registers. Together, these timing parameters completely specify both how the master will talk to the RDRAM, and how the RDRAM will talk back. All fields are preprogrammed with reset values that are guaranteed to be functional at the particular device's fastest specified clock rate. However, optimal performance is typically achieved when some of the parameters are reduced from their preset values. These reductions are done by the master via register write requests.

It is not always easy for the master to know what these delays should be set at—or even what they are set at. Several fields are write-only, so the master cannot query them. Furthermore, in some cases the reset values vary by manufacturer. The only way for a master to know what the reset value is or to figure out the optimum value is to query the DeviceManufacturer register for the vendor and part ID so that it can index through a look-up table of known optimal values. Unfortunately, this method fails for any vendors or RDRAMs introduced after the memory control BIOS was set.

As long as the right speed device is placed on a Rambus Channel, correct functionality is guaranteed. However, with the introduction of a variety of speed grades, correct functionality is no longer guaranteed. There is no way for the controller to detect that a slow part (i.e., 500 MBytes/s) is placed on a fast channel (i.e., 533 Mbytes/s). The device may work most of the time, but fail occasionally under extreme environmental conditions. This, of course, is a problem that has plagued the DRAM and computer industry since the first DRAMs were binned according to access time.

Two flavors of block read

Block data reads come in two flavors: sequential and nonsequential (or random). As one might expect, sequential block reads begin at the requested address and proceed linearly to successively larger addresses. Nonsequential block read requests combine a Request Packet with a Serial Address Packet. The Serial Address Packet is a stream of 8-bit column addresses serialized onto the BusEnable line. Each column

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 281
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

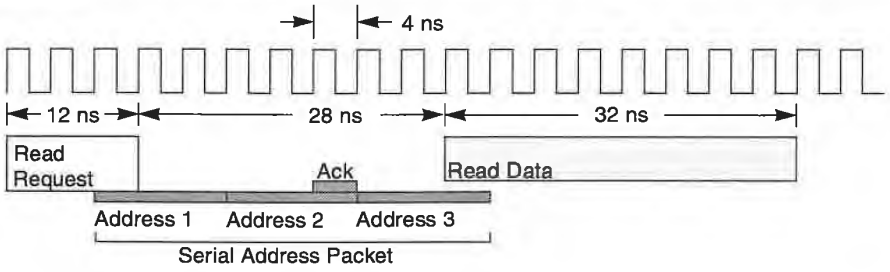


Figure 8-11 Serial address packets on the BusEnable line convey octbyte addresses.

address in the Serial Address Packet specifies the column address of one octbyte in the Data Packet to be sent back to the master. Since both the address of an octbyte and the resulting data consume the same amount of time on the channel, the addresses for octbytes toward the end of a long Data Packet can be sent after the first data bytes have already been transferred.

Figure 8-11 shows the temporal relationship between the Serial Address Packet and the data octbytes that it controls. Note that the address of the first octbyte is contained in the Read Request Packet. Since the Data Packet contains four octbytes, the Serial Address Packet contains the column addresses of the last three octbytes in the transaction.

Nonsequential burst reads are very useful in a number of circumstances. Because they facilitate a full speed burst of eight-byte chunks selected arbitrarily from a 2-Kbyte region of memory (i.e., a row),[11] they can be used to effect a variety of fetch strategies. Most notably, desired-word-first[12] and sub-block-ordered (Intel interleaved) transfers are straightforward. The constraint that each octbyte is linearly ordered and starts with an 8-byte aligned byte is not a concern: only in cases of a processor bus narrower than 64 bits does the memory controller have to perform a small amount of data reordering to attain the desired result.

11. Because the Serial Address Packet is limited to 8 bits of address per octbyte, only the low-order 11 bits of the full byte address can be specified in this way. If subsequent generations of RDRAMs have rows that are larger than 2 Kbytes, nonsequential read and write bursts will be limited to portions of a row.

12. Also called critical-word-first.

Samsung Ex. 1097, p. 282
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

Nonsequential burst reads are also extremely useful in graphics application in which the frame buffer has been tiled (see Section 4.7). In these circumstances, vertically adjacent sections of sequential scan lines are mapped onto different portions of the same DRAM row. A drawing operation that manipulates more than one scan line touches several noncontiguous locations in a single row. Such complex references to a single tile can be arbitrarily combined into a single nonsequential burst read (or write) with no loss of bandwidth.

Eleven flavors of writes

The Rambus protocol defines three types of write transactions and four types of write masking. Eleven of the twelve possible combinations are supported. All writes involve sending a Data Packet very shortly after the Request Packet, as shown in Figure 8-9. The Data Packet is initiated before the Slave has time to Ack or Nack the request. If the Data Packet is long, it can be terminated by the master if the request is Nacked. Some write requests also involve a Serial Address Packet being sent on the BusEnable line. The three types of write are:

Wseq Sequential Write

Sequential writes corresponds to sequential reads: the Data Packet contains a linearly sequential stream of octbytes. Byte granularity of the write is accomplished by allowing the low order 3 bits of the count and start address to be non-zero. If they are non-zero, some of the first and last octbytes in the Data Packet are not written.

Wnsq Nonsequential Write

Nonsequential writes correspond to nonsequential reads: a Serial Address Packet is sent in parallel with the Request and Data Packets. Complete octbytes are transferred (subject to the optional masking described below), so the low order three bits of the address and count fields of the Request Packet must be zero.

Wbns Nonsequential Write with Byte Masking

This class of write includes eight bytes of mask (one octbyte) into the Data Packet before every eight octbytes of data. Each byte of mask corresponds to one of the subsequent octbytes of data; each bit in a byte of mask enables or disables the writing of one byte in the corresponding octbyte of data.

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 283
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

Each of these write types (except the last) comes in four flavors. Three of them make use of a 72-bit register[13] called MDReg. The four flavors are:

Npb — No Per-Bit Mask

The write is a generic write with no masking on a per-bit basis.

Dpb — Data-Per-Bit

With this flavor of write, the data to be written is provided in the Data Packet, and the MDReg is used as a per-bit mask that enables or disables the writing of each bit in every octbyte in the transfer.

Mpb — Mask-Per-Bit

In mask-per-bit writes, the data to be written comes from the MDReg, and the contents of each octbyte in the Data Packet are interpreted as 72 bits of per-bit mask. This is exactly the same as Data-Per-Bit writes with the roles of the Data Packet and the MDReg reversed.

Bpb — Both-Per-Bit

In this type of write, every octbyte of data in the Data Packet is preceded by an octbyte of mask. Each bit in the mask octbyte enables or disables the writing of the corresponding data bit in the following octbyte of data. As a side effect, the MDReg is loaded with the last octbyte of mask for use in subsequent Mask-Per-Bit or Data-Per-Bit writes.

Essentially all combinations of byte masking and data sequentiality with the per-bit masking are allowed. A complete specification of a write combines a type with a flavor. For example, WseqNpb is a simple write of linear data with no byte of bit masking. WbnsMpb involves writing the contents of the MDreg to a random series of octbyte addresses (as specified in the Serial Address Packet) with both per-byte and per-bit masks provided in the Data Packet.

The only disallowed combination is Nonsequential Writes with Byte Mask (Wbns) with Both-Per-Bit (Bpb) masking. This combination does not make sense because it includes a redundant

13. In ×8 RDRAMs, such as the 8- and 16-Mbit parts, the MDReg is 64 bits long. The MDReg is not listed in Table 8-1 because it is not accessible via register read and write requests.

Samsung Ex. 1097, p. 284
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

degree of masking and an unorthogonal mixing of mask and data octbytes in the Data Packet.

Of the eleven allowed combination, only one or two would be used in main memory applications. Most others are potentially useful in graphics systems—depending on the layout of pixels in memory and the types of pixel operations that the GUI accelerator needs to perform. It is unlikely that any single GUI accelerator would make use of all eleven write combinations.

Explicit closing of a row is coming

Though not all 16-Mbit generation RDRAMs support it, the Request Packet includes a "Close" indicator (see Figure 8-8). When set, the master is indicating that this is likely to be the last access to this row of this bank, and that the RDRAM should close the row and precharge the bank after the request is completed. If the addressed bank is precharged, the latency of the next row access is significantly reduced. The Close bit is one of the few instances in which the RDRAMs of a single generation do not all exhibit the same functionality. For the RDRAMs that support the close bit, the row access time given a closed bank is 64 ns from the end of the Request packet to the start of the Data Packet. When operating at 533 Mbytes/s, this row access time is reduced to 60 ns.

Long transfers can be terminated

Since the maximum burst length is 256 bytes, Data Packets can contain up to 32 octbytes.[14] A maximum-length packet resides on the bus for 0.5 µs. This is a long time for the memory controller to be locked off of the channel. To eliminate potential problems associated with the duration of a maximum-size packet, the master can terminate read or write data transfers by sending a serial control packet on the BusCntl line. The Data Packet in progress is terminated after the next one or two octbytes have been sent, depending on whether a read or write transaction is being terminated.

Address mapping improves row-cache hit rate

An interesting feature built into all RDRAMs is an address mapping function that reallocates the row, column, and RDRAM select bits, scattering logically adjacent DRAM rows across the RDRAMs. Figure 8-12 shows how eight row address bits are swapped with eight bank base address bits to interleave rows across 128 RDRAMs (256 banks). For systems with fewer RDRAMs, from zero to eight row address bits can be swapped with the corresponding bank select bits by programming the AddressSelect register.

14. Since Wnbs transactions add an octbyte of mask for every 8 octbytes of data, the actual maximum size of the Data Packet is 288 bytes.

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 285
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

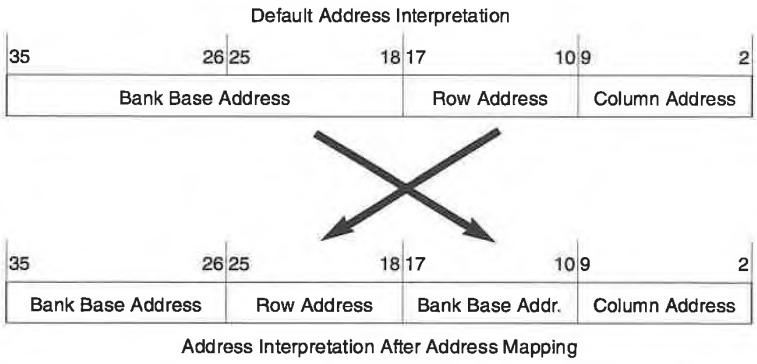


Figure 8-12 Address mapping shifts the location of the row address in the 36-bit physical address to distribute logically adjacent rows across DRAMs.

Two refresh options

There are two basic ways in which RDRAMs are refreshed during normal operation. The first way is to use Touch commands to force at least one row access of every row in every bank in all RDRAMs within the refresh interval. At the 16-Mbit density, the refresh interval is set at 17 ms. In other words, the memory controller must send a Touch command to each 16-Mbit RDRAM every 16.6 µs and to each 8-Mbit RDRAM every 33.2 µs. A Touch command is a read or write access of zero bytes: there are actually four such commands with slightly different characteristics.

The second way to keep RDRAMs refreshed is to perform a special write to the MinInterval register that initiates a burst of four row accesses. In this case, the RDRAM keeps track of the row number and automatically cycles though all the rows in both banks. Because four row accesses are done as a result of one command, the frequency with which the register writes need to be done is one-fourth the frequency needed if touch commands are used. If it is desirable or acceptable to have all RDRAMs refresh simultaneously, then the writing of the MinInterval register of all RDRAMs on a channel can be accomplished with a single Write Register Broadcast request.

Power reduction features

Rambus is potentially very attractive for high-performance portable applications: it provides very good bandwidth out of as little as a single DRAM. Consequently, it provides the minimum possible memory system granularity. Power consumption is potentially very good in that the large row cache should decrease the row-cache miss ratio and decrease the number of

Samsung Ex. 1097, p. 286
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

row accesses. Furthermore, when a row-cache miss does occur, only a single DRAM bank is accessed. In contrast, if narrow generic DRAMs are used, many DRAMs are active on a row-cache miss.

The main obstacle to low-power operation is the Rambus Channel's high operating frequency. Given that dynamic power is equal to CV^2f, the 250-MHz operating frequency of the Rambus interface, though small, is relatively power hungry. The maximum power ratings for all 16-Mbit RDRAMs are given in Table 8-2. On the bright side, at most a single RDRAM actively sends or receives data at any given time. However, the maximum inactive power of 220 mW is much greater than the 10–20 mW that a generic DRAM draws while idly self-refreshing itself.

Mode	Description	Maximum	Units
PowerDown	Device shut down	0.33	mW
Standby	Device inactive	220	mW
Active	Evaluating a Request Packet	550	mW
Read	Sending data to master	760	mW
Write	Receiving data from master	860	mW

Table 8-2 Individual RDRAMs draw substantial power due to high operating frequencies.

These power consumption numbers are somewhat deceptive because they are the absolute maximum values. Specific devices may have lower maximum power ratings than those listed here. Actual power consumption depends on the specific device and the frequency and mix of references. Rambus Inc. has a proprietary spreadsheet that calculates the expected power consumption of a memory system based on these parameters. This makes it difficult to construct realistic power consumption comparisons without their help on a system-by-system basis.

RDRAM's Standby Mode reduces power consumption. One nice feature allows most of the RDRAMs in a memory system to be in Standby Mode while a few devices—containing, for instance, a frame buffer—stay active. Although nine of the ten high-speed receivers are disabled in Standby Mode, the high-speed Delay-Locked Loops remain active, so that power consumption is still significant.

PowerDown Mode for lowest power level

RDRAMs also include a PowerDown Mode that disables the internal clock generators. In this mode, the vast majority of

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 287
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

power consumption is due to refreshes, just as with conventional DRAMs. PowerDown Mode is entered by setting a bit in the Mode register and is exited by counting the number of Serial Mode Packets seen on the BusEnable line. When in PowerDown Mode, a refresh of a single row in one bank is initiated by a positive edge of SIn line. The pulse on SIn is transferred to SOut, so that a 60-kHz square wave applied to the start of the SIn-SOut chain keeps all the RDRAMs on the channel refreshed. This unique scheme eliminates the need for on-chip timing loops that vary from device to device and are sensitive to voltage and temperature.

For low-power systems, four issues are relevant:

- Power dissipated while actively transmitting or receiving data
- Power dissipated while awake but not active
- Power dissipated while in low-power mode
- The ease with which some or all of the DRAMs enter or leave the lowest power mode

The PowerDown Mode dissipation is as low as with any self-refresh conventional DRAM. The power dissipated while awake but not transmitting is somewhat higher than the power dissipated by other DRAM solutions because of the high-speed clocks. However, due to the very high data-transfer rate, the energy consumed per byte transferred is lower than with any other DRAM alternative. This characteristic, combined with the small number of output drivers in the memory controller and the ability to power down portions of the memory system, makes the 16-Mbit RDRAM potentially attractive for portable applications with substantial bandwidth requirements.

Device initialization

The configuration registers determine so much of the RDRAMs' behavior that the RDRAM initialization procedure is a vital element of the overall interface. A low-speed TTL serial chain loops through all the devices on the Rambus Channel via their select-in (SIn) and select-out (SOut) pins (see Figure 8-6). The Rambus master uses this loop in conjunction with the BusEnable line to reset and initialize all the RDRAMs on the primary and secondary channels before high-speed operation begins.

Secondary Rambus Channels

Secondary channels, whether on RModules or not, are connected to the primary channel via RTransceivers (see Figure 8-4 on page 267). The secondary channels have their own clock

Samsung Ex. 1097, p. 288
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

domain that is synchronous with the transceiver's transmit clock. The transceiver acts as a bridge device, forwarding packets from one channel to the other with minimum delay. In fact, the delay from the start of a Request Packet to the start of the response is limited to four cycles: for either a row-cache hit or miss, the access time for an RDRAM on a secondary channel is 16 ns longer than the access time of the same device would be if it were on the primary channel.

8.2.5 Miscellanea

DRAM cores come in two basic flavors:

- Some keep the word line raised continuously from the start of a row access until the start of the next array precharge.
- Other DRAM cores pulse the word line once near the start of a row access.

For RDRAMs of the first variety, data writes propagate all the way into the DRAM cells. Consequently, after a write transaction, delay is necessary before the next request can be considered. When a row miss is encountered, the word line must be lowered and the array precharged before the needed row access begins. The need to lower the word line and precharge the array significantly increases the row-miss latency. These delays are also present in most of the other new DRAM architectures.

RDRAMs based on pulsed word-line arrays can have their row caches in the sense amps decoupled from the DRAM array. A data write propagates as far as the sense amps and stays there. On the next row miss to the bank, a row access must first be done to write the dirty data from the row cache to the array, followed by a second row access to retrieve the new row. However, if the row cache is clean (not written to since the previous row access), the row access to write the row-cache contents into the array is unnecessary.

All current 8-, 16-, and 18-Mbit RDRAMs behave like the second variety—the type that pulse their word lines. All suffer an eight-clock penalty on row-cache misses when the replaced row-cache line is dirty. On 500 Mbytes/s channels, this adds 32 ns to the baseline 116 ns row-cache-miss access time. At 533 ns, a dirty row cache adds 30 ns to the 109 ns row-cache-miss access time. For the RDRAM that support the Close bit, a dirty row cache impacts how long it takes to close an "open" bank and to completely precharge it in preparation for the next row access.

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 289
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

The frequency with which these penalties come into play, and their effect on average latency and sustained bandwidth, are highly dependent on the particulars of the application and the memory system.

Although the focus at Rambus is currently on DRAMs, the specification of the Rambus Channel allows other forms of memory. In particular, Rambus ROMs and flash memories are likely to be available in the very long term if Rambus becomes widely accepted as a main memory device. Furthermore, other forms of Rambus slaves are possible, including RAMDACs, graphics controllers, and polled I/O devices.

Designing ASIC Rambus masters

In an ideal world—at least from the perspective of the people at Rambus—processors would have Rambus Channel interfaces integrated alongside their traditional wide, slow I/O interfaces. If Rambus becomes widely adopted, this will probably happen, because much of Rambus's potential performance is lost when an intermediate device, or bridge, is used. In the interim—and indefinitely for some applications—a Rambus interface ASIC megacell, such as the Rambus ASIC Cell (RAC), will be the only means of creating a Rambus master.

The RAC is a diffused megacell designed and certified by Rambus Inc. As of 4Q1995, the RAC is available on Toshiba, NEC, LSI Logic, and Hitachi gate arrays in .5 and .35 µm technologies. TSMC and LG Semicon will have RAC cells for their ASIC technologies in 1Q1996 and 2Q1996, respectively. IBM—the latest Rambus ASIC licensee—will have its RAC available by mid 1996.

The RAC provides an 8:1 multiplexing of the Rambus Channel, so that the rest of the memory controller ASIC need only deal with 72 bits of data every 16 ns. The RAC, shown in Figure 8-13, also provides some clocking support, but all generation and interpretation of the Rambus protocol is up to the memory system controller. Because the RAC implements a fixed 8:1 multiplexer, if the Rambus Channel is running at 250 MHz, then the ASIC side of the RAC runs at 62.5 MHz. If this frequency is awkward, then the designer has four options:

- Add buffers to interface the two temporal domains.
- Slow down the entire Rambus subsystem to facilitate a more convenient RAC interface frequency.
- Speed up the Rambus Channel to 533 Mbytes/s, which translates to a RAC interface frequency of 66 MHz.

Samsung Ex. 1097, p. 290
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

RDRAMs and RACs capable of 533 Mbytes/s are expected to be available in 1H1996.

- Run the Rambus Channel and the processor domains synchronously with each other but at different—but related—frequencies. For example, if the ASIC runs at 66 MHz and the Rambus Channel operates at 8×3/4 speed (i.e., 50 MHz), then there is always a clock edge on the Rambus Channel that lines up with each clock edge in the ASIC temporal domain. Consequently, in this case the mismatch between the two frequencies does not introduce any additional latency. With other convenient ratios of operating frequencies, a small amount of latency is added to some fraction of requests as an appropriate clock edge is waited for.

The RAC adds latency. In the case of a row-cache hit to a 16-Mbit RDRAM, the first byte is returned to the master no sooner than 40 ns after the start of the Request Packet. With an 8:1 demultiplexing of the receive data, the first byte is unavailable within the memory controller until one cycle after the seventh byte has arrived. This adds another 20 ns to the raw RDRAM latency to the first usable byte. However, the entire transaction is increased in length by only a single clock at its start and another at its conclusion—a total of 4 ns. For many applications, a delay of 20 ns in the delivery of the first chunk of octbyte is not very important because both the utilization of the Rambus Channel and the sustained bandwidth are unaffected. For most implementations, though, the multiplexing/demultiplexing function of the RAC will directly affect system-level performance.

Rambus Protocol Unit simplifies controller design

Though the RAC deals with all the high-speed signals, it does not aid at all with protocol generation and interpretation. Rambus provides a synthesizable Verilog Rambus Protocol Unit (RPU) to help with controller design. With about 3,000 gates, the RPU manages most of the Rambus protocol generation tasks and provides a much simplified memory read/write interface for memory controller designers to build on or modify.

Multiple masters are possible

The Rambus protocol is a general purpose single-master, multiple-slave protocol that would work well for any memory device. Though Rambus Inc. does not emphasize this feature, the protocol does facilitate pseudomasters. Pseudomasters do not initiate slave reads and writes, but they do supply or accept the data being transferred between the master and a true slave. These pseudomasters must be placed between the master and the slaves on the channel. Technically, by having the masters

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 291
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

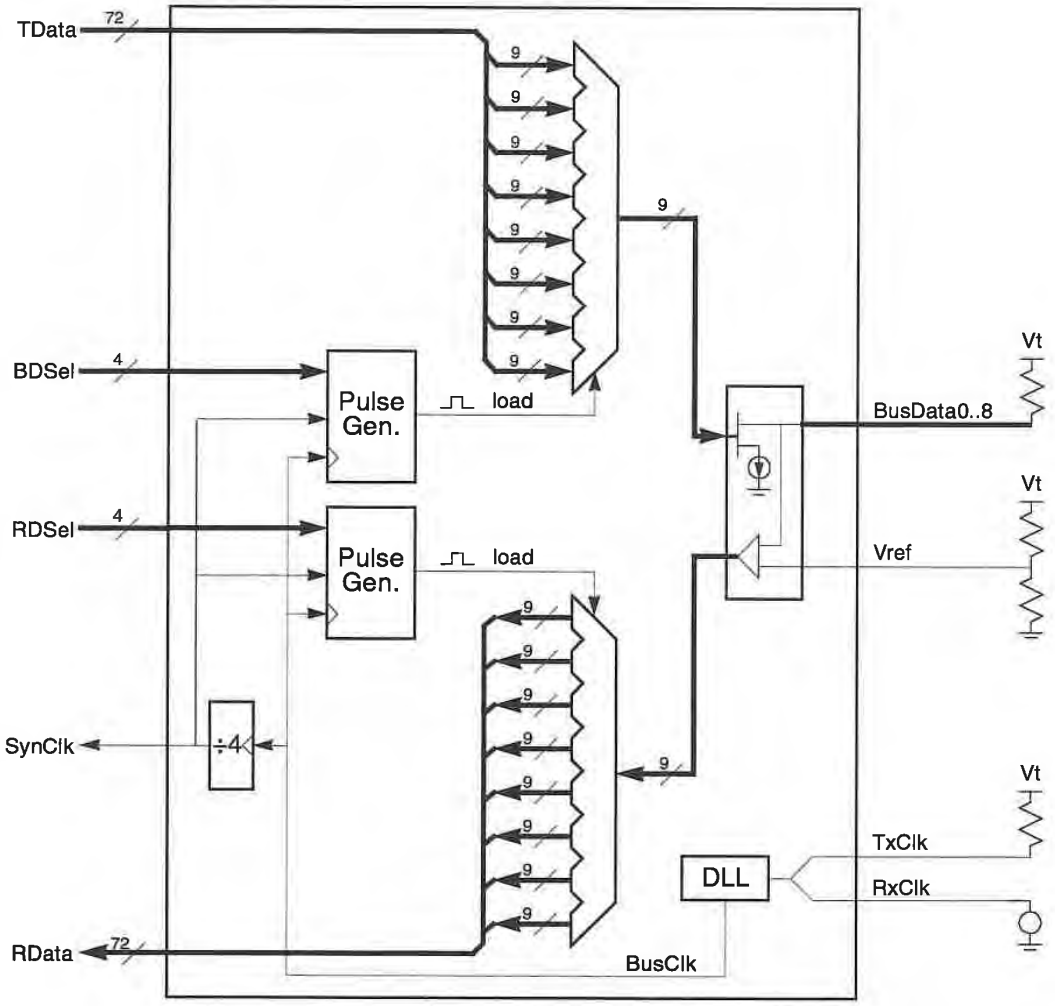


Figure 8-13 The RAC is essentially an 8:1 multiplexer/demultiplexor with some clock-generation circuitry.

Samsung Ex. 1097, p. 292
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

together at the unterminated end of the channel, it should be possible to have two true masters on a channel. If the masters are physically close enough, then the small phase difference between their clocks will not affect reliable operation. So far, no dual master system has been demonstrated. With no provision for interrupts or traditional DMA, the Rambus Channel would not make an effective general purpose memory and I/O interconnect without further extension or the use of nonstandardized peripheral signals.

As noted earlier, with 9-bit RDRAMs, the ninth bit is uncommitted in the protocol. It can be used for parity, ECC codes, prefetch hints, or anything else the system architect dreams up. If a 9-bit RDRAM is used, and if the ninth bit is used entirely for parity or ECC, then the communication of data to and from the DRAM is protected from error as a side effect of the protection against errors in the memory array. However, the Request and Acknowledge Packets cannot be protected against transmission errors. Similarly, the Data Packets cannot be protected against transmission errors if 1) 8-bit RDRAMs are used, or 2) the ninth bit is used for something other than error protection. Of course, all of these problems, particularly the difficulty in checking address and control information, are shared by the other DRAM architectures and by nonredundant electronics in general.

Granularity, performance, and design migration

One of the greatest problems with generic DRAMs is that many DRAMs must be used in parallel in order to build a memory system with reasonable bandwidth. Consequences of parallel-ganging generic DRAMs include increased memory system granularity, wider buses, more pins on the memory controller, and greater system-level power consumption.

The greatest strength of the Rambus system is that high-bandwidth memory systems can be constructed out of a single DRAM. Thus the granularity of the memory system is a single DRAM. Increasingly, many applications demand memory systems that are small but capable of delivering high performance. As densities increase past 16 Mbits, the number of applications that use one to four DRAMs is becoming very significant.

Rambus Inc. claims that row-cache miss rates as low as 5% are possible for general purpose code. The subject of row-cache miss rates is, however, rather speculative. These rates depend a great deal on the application, the size of main memory, and the allocation of the application to the memory banks. However, the RDRAM's relatively large row-cache size and its address-swapping feature mean that the row-cache miss rates should be

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 293
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

lower than for the other alternatives. The actual row-cache miss ratio will depend strongly on the primary and secondary cache organizations and the characteristics of the applications being run.

Rambus's peak bandwidth of 500 or 533 Mbytes/s is deceptive: the maximum sustained bandwidth is substantially lower. For example, with 32-byte transfers, the maximum sustained bandwidth for the 16-Mbit RDRAMs ranges between 166 Mbytes/s and 308 Mbytes/s, depending on the row-cache miss ratio. The maximum sustained bandwidth declines further with shorter bursts but increases with longer bursts.

Although the sustained bandwidth of a Rambus Channel is significantly lower than 500 Mbytes/s, it is nevertheless substantially higher than the bandwidth of any comparable generic DRAM or VRAM memory system. A generic -60 ×4 DRAM can sustain 4–12 Mbytes/s. A high-performance -70 ×8 VRAM[15] can sustain 50–70 Mbytes/s on its two ports. Furthermore, just as DRAMs and VRAMs are ganged in parallel to increase their net bandwidth, Rambus Channels can be used in parallel in either synchronous or asynchronous configurations to further improve total memory system bandwidth.

Ganged multiple-channel memory systems

When a memory system consists of more than one Rambus Channel, logically adjacent bytes can be either placed together within one RDRAM or scattered across the channels in several RDRAMs. In the first case, the channels are relatively independent of each other. Individual memory references are handled on one channel, and each of the channels can simultaneously handle different memory references. In this case, the peak bandwidth used on any one transfer remains 500 Mbytes/s but the maximum sustainable bandwidth for a system with n channels is n times that of a single-channel system. This requires n independent channel controllers.

When adjacent bytes are interleaved across all the channels, all channels must be used in parallel on each reference. The channels are operated in lockstep and must be populated identically: effectively, a single Rambus Channel that is n bytes wide is formed. In this case, the transfer rate for all memory references is n times 500 Mbytes/s. However, because the Request Packets must be duplicated on all the channels, the maximum sustained bandwidth increases much more slowly. The memory controller for this configuration is essentially the same as for a single

15. Based on NEC's μPD482234 2-Mbit VRAM.

Samsung Ex. 1097, p. 294
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

Rambus Channel, except that the datapaths are proportionately wider.

Secondary cache replacement

Like the vendors of several new DRAM architectures, Rambus Inc. indicates that in some cases a Rambus memory system can be used to replace both the generic memory system and the external secondary cache. The actual cost and performance trade-off depends greatly on the speed of the processor interface, the memory size, and the applications. For many configurations, a Rambus-based memory system offers cost savings and performance comparable to that of medium to large secondary cache systems. These performance levels are especially attainable if the application's spatial locality of reference is sufficient to allow a prefetching Rambus memory controller can take advantage of excess Rambus bandwidth to decrease mean memory access time.

System-level cost savings

In many ways, an RDRAM is a 64-bit DRAM with a high-speed 8-bit interface. Transfers are in multiples of 64 bits, and the 533 Mbytes/s peak bandwidth is equivalent to that of a ×64 SDRAM operating at 66 MHz. In comparison to a ×64 SDRAM—or any of the wider, lower frequency DRAMs—a main benefit of the Rambus architecture is that the total number of pins needed to implement the memory system is reduced. The number of pins needed on the memory controller is 31 pins per channel instead of about 110 for a conventional ×64 memory system. This savings facilitates a smaller controller package or leaves more controller pins for additional channels. It also greatly eases the controller pin-count constraints placed on tightly-coupled UMA architectures (see Section 4.8).

Reducing the number of pins on the DRAMs and DRAM modules also means that the memory system can be physically smaller, and that memory expansion can more easily be made externally accessible. Secondarily, since many fewer signals are transitioning simultaneously, less electromagnetic interference (EMI) is generated.

Overall, if the per-bit price premium is low, then in many circumstances, a narrow, high-speed interface can yield a lower overall system cost than the more conventional alternative.

8.2.6 First Generation 4-Mbit RDRAMs

The second generation 18-, 16-, and 8-Mbit RDRAMs represent a significant improvement over the 4.5-Mbit devices. Toshiba

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 295
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

began sampling a 4.5-Mbit RDRAM in 3Q1993. However, with no design wins at this density, the part was never moved into production. The second generation devices have since become the only RDRAMs available, and the 4.5-Mbit device and architecture have been relegated to a historical footnote.

Rambus was first conceptualized as a main memory solution. Consequently, the first Rambus protocol—as implemented at the 4.5-Mbit level—lacked the features that make its 16-Mbit devices good graphics-oriented devices: only linear, sequential read and write bursts were supported, with no bit or byte masking on writes. The MDReg register did not exist and only ×9 devices were planned. The ability to terminate a transfer and PowerDown mode are also new in the second generation.

In the first generation architecture, Refresh was primarily intended to be handled autonomously by the DRAM. Though manual refresh was possible, it was much more awkward and inefficient than with the 16-Mbit devices. Recognizing that memory controllers inherently want to know the state of the DRAMs they control, Rambus dropped self-refresh in the second generation. Simultaneously, one of the configuration register fields that controlled self-refresh was eliminated.

Throughout the industry, DRAM row and column access times decreased noticeably from the 4-Mbit to the 16-Mbit generations. The same was true for RDRAMs. Though both generations of RDRAMs transferred data at 500 Mbytes/s, the 4.5-Mbit RDRAM row-cache hit and miss latencies were substantially longer than the current 16-Mbit devices' latencies.

8.2.7 The 64-Mbit RDRAMs

In mid 1995, Rambus began disclosing details of its architecture for the 64-Mbit generation. As of this writing, not all the technical details have been disclosed. However, enough has been released to get a flavor of the changes in store.

The rationale for the changes being introduced is rooted in four important trends:

- Superscalar, out-of-order processors capable of having several outstanding memory references will be widely available by 1997.
- Processor, multimedia, and I/O bandwidth demands on main and graphics memories will continue to increase.

Samsung Ex. 1097, p. 296
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

- UMA and integrated multimedia architectures imply that memory systems will increasingly have to simultaneously meet the needs of several logical sources of references.
- Most memory systems populated with 64-Mbit devices will have a minimum size (or granularity) of four or fewer DRAMs.

The metagoal for the 64-Mbit RDRAM was to create a mainstream device that is well suited to the PC main memory market. Specific goals set for the 64-Mbit RDRAMs included the ability to propagate concurrency of operation down into the DRAMs, and to increase the sustainable bandwidth on a single channel with short transfers.

These goals are being reached through changes in the RDRAM organization, the RDRAM capabilities, and the Rambus Channel protocol:

- The 64-Mbit RDRAMs will be 4-bank devices. Each bank will have a 2-Kbyte row cache, making for a 1,024:1 ratio of row cache to memory size.
- Concurrent row and column operations will be allowed in a single RDRAM. Separate row and column access resources will be present in RDRAMs, so that one access of each type can be active simultaneously. Specifically, if an access is in progress, then a new access to the same RDRAM can be initiated in all but two cases. If the second access is a row access to a bank that is being used by the first access—regardless of whether the first access is a row or column access—then although the second access can be accepted and queued in the RDRAM, it cannot be begun until the first access is substantially complete. This restriction makes sense: if the second access is a row access, one of the first things it wants to do is to clear the row cache by precharging the sense amplifiers.

  The ability to overlap transactions within an RDRAM is an improvement over the 16-Mbit architecture, which allowed at most one active operation per device.
- Row accesses will no longer need to have two request packets. In the 16-Mbit generation, the first request is NACKed, and the second is ACKed (see Figure 8-10). In the 64-Mbit generation, the time from a request that generates a row-cache miss to the data will be controlled by the memory controller without need for a second request packet.
- Some of the access latencies will be reduced to reflect the general speed improvements derived from the 64-Mbit

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 297
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

processing technologies (typically 0.35 μm and below). The row-cache-hit access latency and the closed-bank row access will be the same as in the 533 Mbytes/s 16-Mbit generation parts. When operating at 533 Mbytes/s, the row-cache-miss access latency from the end of the Request packet to the start of the Data Packet will be reduced to 90 ns.

Consistent with the general migration to 533 Mbytes/s, all 64-Mbit RDRAMs will be capable of this speed or greater. Speeds to 600 MBps are projected for use on constrained Rambus Channels. This is not to say that the DRAM vendors will not sell slower parts as 500 Mbytes/s parts.

Also, to remove the timing differences between a row-cache miss to a clean row cache and to a dirty row cache, all 64-Mbit RDRAMs will use level word-line DRAM arrays. This will make it easier for the memory controller to predict when data will be available in the case of a row-cache miss.

Despite all these changes, 64-Mbit devices will be able to cohabitate with the 16-Mbit devices on the same channel. Issues regarding the number of banks and bank organization aside, controllers designed for the 16-Mbit generation will be able to control 64-Mbit devices as well, though at a lower performance level.

The end result of these changes is that a single 533 Mbytes/s RDRAM will be able to sustain 425 Mbytes/s bandwidth with 100% row-cache misses and 32-byte transfers—provided adjacent accesses are to different banks. This bandwidth is the maximum that the channel can sustain, given the overhead associated with the sending of the Request Packets and with turning the bus around. This same bandwidth is attainable with row-cache hits even if all the references are to the same bank. With consistent support for the Close bit, the closing and precharging of banks will potentially be hidden under data transfers from other banks. This rate of 425 Mbytes/s contrasts very favorably with the maximum bandwidth deliverable by a single 16-Mbit RDRAM performing 32-byte transfers: 190 Mbytes/s with 100% row cache misses and 266 Mbytes/s for 100% row cache hits.

The die size penalty for the Rambus-specific circuitry will continue to decline as a percentage of the total die area. Rambus's target for the 64-Mbit generation is a die size less than or equal to that of ×32 EDO DRAM.

Samsung Ex. 1097, p. 298
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

The changes in the 64-Mbit RDRAMs apparently address all the deficiencies in the 16-Mbit generation that limit the sustained bandwidth to be substantially less than the 500 Mbytes/s peak value. Though the 64-Mbit generation will first be available in late 1996, it is unlikely that this new generation will be cost-competitive until late 1997 or 1998. Although there have been no product announcements by the RDRAM vendors, Rambus has hinted strongly that efforts to retrofit these 64-Mbit enhancements back into the 16-Mbit generation are underway. The specific timing and circumstances of the functionality upgrade to the current generation are as yet unknown.

8.2.8 RDRAM Fill Frequencies

Figure 8-14 shows the peak fill frequencies for all previous and proposed RDRAM devices. The figure illustrates clearly the reason for the 4.5-Mbit device's lack of marketplace success: 1) the device had superior peak bandwidth but it lacked the features required by the graphics marketplace; and 2) the device provided more bandwidth than was needed by the main memory market. The main memory market was unwilling to pay a premium for performance it did not need.

At the 16-Mbit level, RDRAM peak fill frequencies are in line with the requirements of the graphics market. These peak fill frequencies are greater than those of any other main memory-oriented devices at this density. At this density, RDRAMs compete directly with a host of graphics-oriented RAMs: video RAMs (VRAMs), Synchronous Graphics RAMs (SGRAMs), Window RAMs (WRAMs), and Multibank DRAMs (MDRAMs).

In graphics applications, the random access read and write modes work well in conjunction with tiled frame buffers to increase the average transfer size substantially above what would be common with most of the other DRAM architectures. This has the effect of reducing the gap between the peak and sustained bandwidths and, by inference, fill frequencies.

In main memory applications, the typical transfer size is less than that found in graphics application: 32 bytes in PCs and up to 128 bytes in workstations. At the low end of this range, the 8-, 16-, and 18-Mbit DRAMs' sustainable bandwidth is substantially less than their peak of 500 or 533 Mbytes/s. For these devices, the gap between the peak and the sustained fill frequencies is larger than for most evolutionary alternatives. This artifact illustrates that peak fill frequencies are not the best

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 299
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

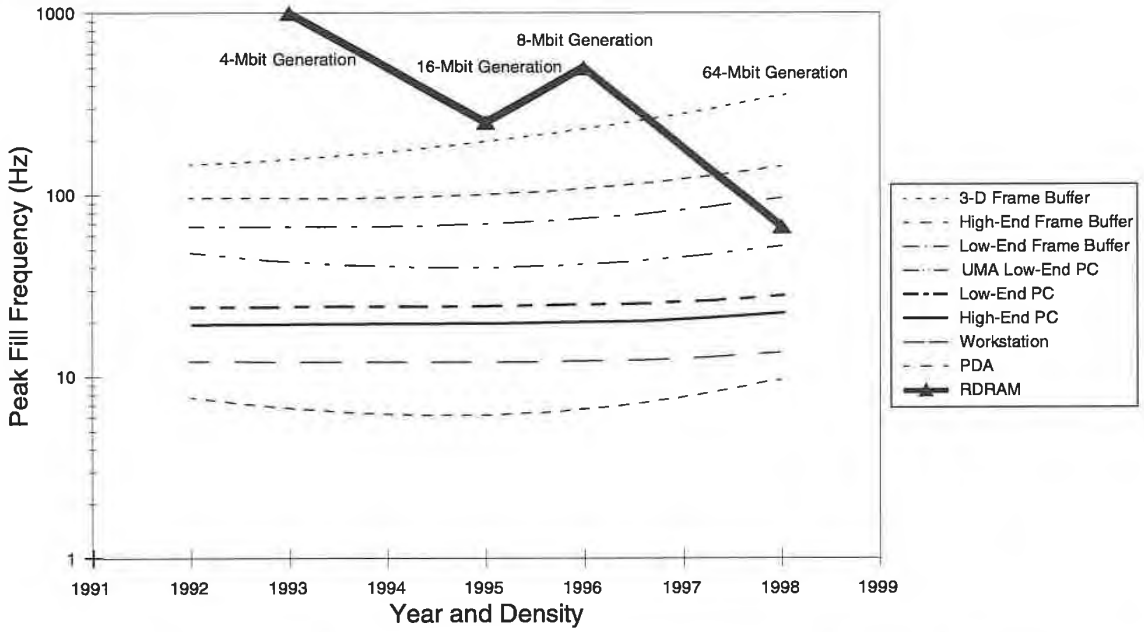


Figure 8-14 RDRAMs' peak fill frequencies are very high at the 4- and 16-Mbit generations. This makes them more appropriate for graphics and UMA systems. The 64-Mbit devices will be more appropriate for main memory applications.

metric for comparing DRAMs to each other or to system requirements. The sustainable fill frequencies used in Chapter 10 provide a more realistic metric.

At the 64-Mbit level, RDRAMs are one of the few organizations that have enough bandwidth to satisfy the UMA and main memory marketplaces with granularities as low as 8 Mbytes. All other potential candidates at this density are wide—×32 or wider. The Rambus camp has the opportunity to capitalize on the narrower interface in order to provide the main memory market with a more cost-effective solution than any of the evolutionary alternatives.

8.2.9 Current Status and Business Issues

Licensing strategy

Rambus's strategy is to make RDRAMs a standard. Rambus plans to accomplish this goal by licensing its technology to a variety of large DRAM manufacturers and other semiconductor

Samsung Ex. 1097, p. 300
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

vendors. Rambus itself is an intellectual property company: it does not sell DRAMs directly. In total, 15 companies have taken Rambus licenses. Those requiring a license include companies wishing to build chips with a Rambus interface: either controllers or slaves. Customers of the DRAM vendors and ASIC houses need not purchase a license.

Seven RDRAM vendors

Rambus initially licensed Toshiba, NEC, and Fujitsu—all major players in the DRAM marketplace. Toshiba sampled its 4.5-Mbit RDRAMs in 3Q1993. NEC and Toshiba began sampling 16- and 18-Mbit parts in 3Q1994 and 1Q1995, respectively, and both have moved into production as of 4Q1995. Both companies are also sampling the single-bank 8-Mbit RDRAMs, with production to begin in 1Q1996. Fujitsu had working 4.5-Mbit RDRAMs and RACs but decided not to commercialize the parts.

Since March 1993, Rambus has licensed four new DRAM vendors: OKI, LG Semicon (previously called Goldstar), Samsung, and Hitachi. NEC, Toshiba, OKI, and LG Semicon should all be in production with 8-, 16-, and 18-Mbit devices by mid 1996, all at both 500 and 533 Mbytes/s. Hitachi has announced that its first RDRAMs will be at the 64-Mbit level. As of 4Q1995 Samsung is silent on the topic of its RDRAM development and production schedule.

Seven ASIC licensees

Initially, Toshiba, NEC, and Hitachi were licensed to produce RACs for their ASIC families. As of 4Q1995, the original three ASIC vendors have been joined by LSI Logic, and all four have incorporated the RAC into their most recent 3.3 V 0.5-μm and 0.35-μm ASIC families. Other announced ASIC licensees are LG Semicon, TSMC, and, most recently, IBM. TSMC is expected to have RACs available in 1H1996, while LG Semicon has not announced general availability of its RAC capabilities.

No DRAM or ASIC vendor has yet announced availability of RTransceivers.

For the design and manufacture of sockets for RModules, Rambus has licensed both Augat and Molex. Horizontal and vertical RSockets are both in production.

Other licensees

Another publicly announced licensee is Cirrus Logic. Cirrus is working on a family of Rambus-based GUI accelerators to be available in 1996. The product line will ultimately include devices with one and two Rambus Channels. A separate UMA chip set/GUI accelerator is also planned.

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 301
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

Rambus is also working in conjunction with its ASIC licensees to provide the core logic of a UMA memory controller for Pentium systems. Individual chip-set or GUI accelerator companies could augment the design to provide differentiation and value added.

Other licensees include Chromatics Research, whose Mpact Media Engine relies on an RDRAM for local memory; Oxford Computer, which is developing a Rambus-based DSP chip; and Brooktree. Brooktree's case is interesting because no specific products have been announced. However, the encoding of the DeviceManufacturer register has a value reserved for Brooktree-designed slaves. This implies that Brooktree has or had plans to offer some form of Rambus-based RAMDAC, GUI accelerator, or multimedia engine.

Despite all this activity and longstanding predictions of the arrival of general-purpose processors that include a Rambus interface, no such processors have been announced. However, NEC does offer a Rambus-processor bridge chip set to facilitate R4000-based systems.

Significant design wins emerging

Licensee arrangements are important to Rambus, but volume sales of RDRAMs will be driven by customers of the licensees—companies that do not necessarily have any relationship with Rambus. To date, only a few companies have announced the development of a Rambus-based end product, and as of 3Q1995, only one is actually shipping.

That product is SGI's Indigo2 IMPACT 3-D workstation. Each workstation has between 6 and 18 16-Mbit RDRAMs, depending on the configuration. The RDRAMs store the frame buffer. In addition, two ASICs, each with a RAC, communicate at 500 Mbytes/s via a custom protocol over a dedicated Rambus Channel. SGI expects to sell about 30,000 of these systems between mid 1995 and mid 1996.

Another very significant design win that will be in production soon is Nintendo's new Ultra64 game machine. This new generation device basically consists of a MIPS architecture microprocessor, a graphics and I/O chip, and a 16-Mbit RDRAM on a two-layer PC board measuring about 2 inches by 3 inches. It also supports user expansion of the Rambus-based main memory via a Primary Horizontal RModule directly accessible by the user. The game player will be introduced in Japan in 1H1996 and in North America in 2H1996. This will undoubtedly be the first device to consume a million RDRAMs—probably within a few

Samsung Ex. 1097, p. 302
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

months of its introduction. Nintendo has predicted that the unit will sell three million units in its first year in Japan alone. If successful, tens of millions of these game machines could be sold over several years.

Other technology demonstrations

Rambus has built a single-channel eight-RDRAM graphics demonstration system with three double-buffered, full-motion 640×480 NTSC video sources displayed onto the four quadrants of a 1,280×1,024×18 frame buffer. The peak bandwidths on the Rambus Channel into and out of this frame buffer sum to 344 Mbytes/s.

Price premium will decline at 16-Mbit level

Significant concern surrounds the price premium for RDRAMs. Though the die size penalty for adding a Rambus interface to Toshiba's 4.5-Mbit DRAM was about 15%, the die size penalty at the 16-Mbit generation will be close to 10% by the end of 1996 compared with 1M×16 EDO DRAMs. Though initial price premiums have been substantial, NEC expects the cost of a 16-Mbit RDRAM to fall below that of four 256K×16 EDO devices by the end of 1995. Generally, the current generation of RDRAMs are expected to track the pricing of the other graphics-oriented 8-Mbit DRAMs.

Concern regarding Rambus's price premium centers on the additional test time and expense needed to verify the high-speed interface. Rambus Inc. counters this concern by pointing out that the high-speed tester needed to test the interface is cheap in comparison to the massive production memory testers. The final high-speed test of the interface is relatively short; in the future it may also be used to verify complete post-burn-in functionality. Furthermore, as with the dual-bank SDRAMs, substantial opportunity exists to reduce overall test time by testing one bank of the RDRAM while waiting for the bits in the other bank to decay.

Royalties on RDRAMs and memory controllers

In the long term, Rambus Inc. intends to earn the majority of its income from royalties paid by its licensees: a small percentage on RDRAMs and a larger percentage on processors and logic products incorporating a Rambus interface. The royalties apply to the ASIC memory controller's RACs on the device and are paid by the semiconductor vendor. These costs are presumably but not necessarily passed on to the ASIC vendor's customer. The customer need not have a relationship with Rambus and would not typically pay any royalties directly to Rambus. In general, the royalty arrangements apply to the last sale of the chip incorporating the Rambus technology as a separate chip.

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 303
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

Some potential users of the technology are concerned that royalties may keep the overall price premium higher than it would otherwise be. Rambus Inc. seems to appreciate the risk of stifling adoption of its technology by overpricing. However, since RDRAM pricing is between the licensees and their customers, Rambus Inc. has relatively little influence on the day-to-day pricing of RDRAMs. Regardless, philosophical and pragmatic concern persists regarding this "taxed technology," especially among the DRAM and processor vendors.

Much of the concern about Rambus's royalty arrangement stems from the current situation with Texas Instruments. As a result of an extensive patent portfolio, TI derives significant royalties from other DRAM manufacturers. However, with respect to Rambus, this concern seems somewhat misplaced. Rambus Inc. is, in many respects, more akin to Dolby Labs than to Texas Instruments. The royalties and up-front licensing fees paid by the DRAM and processor vendors to a large extent pay for specific goods and services—that is, for engineering work that the vendors would otherwise have to perform and for access to a standardized marketplace that might not otherwise exist.

If RDRAMs do become so wildly successful that they dominate the future DRAM marketplace, then Rambus Inc. would be in a good position to demand substantial royalties from latecomers to the market. However, a more likely scenario involves the fragmentation of the overall DRAM marketplace, without any single dominant architecture (see Section 11.5). In this sort of environment, price premiums and, consequently, royalties, will have to be kept reasonably low to capture and protect market share. Ultimately, markets that derive significant benefit from Rambus's bandwidth and granularity will be willing to pay for it, and those that do not will seek cheaper alternatives.

Summary

The 16-Mbit RDRAMs are essentially graphics devices. They are well placed to participate in the coming replacement of standard VRAMs and fast-page-mode DRAMs across a wide range of graphics solutions. Rambus's primary short-term competition includes SDRAMs, WRAMs, and MDRAMs. In this domain the 8- and 16-Mbit RDRAMs' higher latency does not pose a performance problem. At the 64-Mbit level, RDRAMs become much more suited to UMA and main memory applications, with favorable characteristics compared with the wide ×32 EDO DRAMs or SDRAMs needed at that density.

Samsung Ex. 1097, p. 304
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

1995 marked a significant milestone for Rambus: the first shipment of a product including RDRAMs. In 1996, the number of such products should increase, causing a signficant volume ramp to develop. However, it will likely be 1998 before RDRAMs represent more than a few percent of the total DRAM marketplace.

8.2.10 Strengths and Weaknesses

Strengths

- High-performance memory systems with just one DRAM
- Minimum possible memory system granularity
- Cookbook solution for high-speed design
  - Very standardized across vendors
- In some cases, the large row cache facilitates elimination of a second-level cache
- Small number of pins on the memory controller
- Good match for frame buffers
- Broad industry licensing
  - Seven major DRAM vendors
  - Seven ASIC vendors
- Much improved 64-Mbit architecture
  - More banks
  - Concurrency
  - Much better sustained bandwidth

Weaknesses

- High latency, especially in 16-Mbit generation
- Limited parallelism within a single 16-Mbit RDRAM
- Negative industry perception of royalty structure
- Byte-serial to word-parallel delays
- Die size penalty at 8- and 16-Mbit levels
- Not evolutionary from EDO

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 305
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

8.3 RamLink DRAMs

The second revolutionary interface is the RamLink standard developed by the IEEE Computer Society P1596.4 Working Group.[16] Though completed in the summer of 1995, the standard is still pending ratification by the IEEE. When ratified, it will become IEEE Standard 1596.4-1996. This working group applied the techniques and general philosophy adopted in the Scalable Coherent Interface (SCI) to develop a ring-topology interconnect for memory and I/O devices. RamLink is revolutionary because, like Rambus, it replaces the RAS/CAS multiplexed address interface with a single byte-wide interface that carries addresses, control, and data. However, RamLink differs from the Rambus Channel in that its primary physical-layer specification is of a ring topology constructed of high-speed point-to-point links.

As with Rambus, the entire generic DRAM interface has been rethought. As a result, all aspects of the interface are new. An important contribution of the RamLink standard was the differentiation of the logical specification from the electrical and physical specification. Specifically, the protocol that defines how the master communicates with the slave DRAMs is called RamLink; the primary physical layer, being the ring topology, is called RingLink. The RamLink standard defines an alternate physical layer specification called SyncLink. This specification was the starting point for, but is different from, the final SyncLink interface described in Section 8.4.

Though the explicit partitioning of the logical and physical portions of the interface is very convenient, it requires some special nomenclature. In this report, *RamLink* is used when discussing something specific to the logical level. *RingLink* is used in reference to aspects of the physical interface, and *RamLink/RingLink* is used when describing the memory system formed when the two are used together.

A RamLink/RingLink memory system is made up of a collection of ringlets. Each ringlet consists of a single master and up to 62 RamLink slaves (see Figure 8-15). Slaves can bridge to other ringlets. Each slave has a 500-Mbytes/s input port from the previous slave and a similar link to the next one.

16. The primary author of this report was a member of the P1596.4 Working Group and of the Working Group's Ballot Review Committee.

Samsung Ex. 1097, p. 306
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

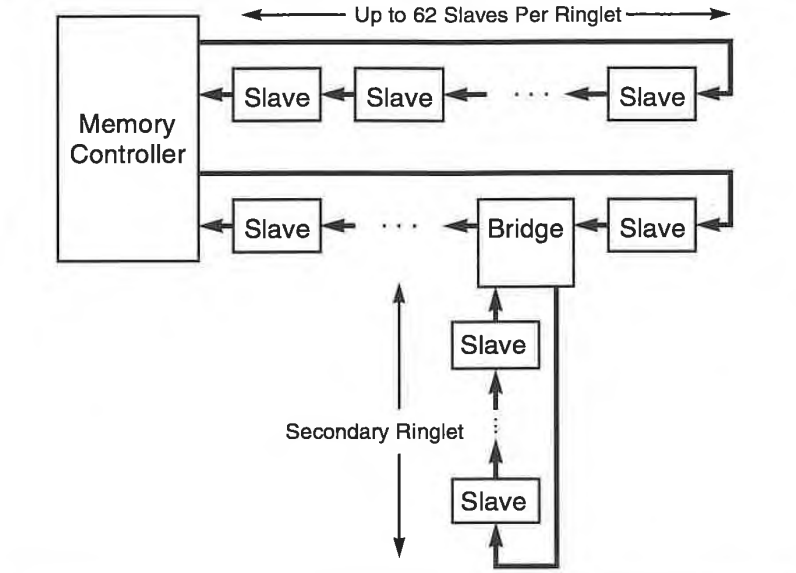


Figure 8-15 A RamLink memory system consists of a collection of ringlets, each with up to 60 slaves or secondary ringlets.

8.3.1 Pin-Level Interface

RingLink devices have two 8-bit or 9-bit ports, one for the incoming stream and one for the outgoing. In addition, each port carries a Flag signal for delimiting the start and end of packets, a 250-MHz clock that is phase-aligned to the data, and a quarter-speed reference frequency clock (see Table 8-3). Transmitting data on both edges of the clock doubles the bandwidth at a given clock frequency. Alternatively, transmitting data on both clock edges means that the clock signal is not running at twice the frequency of the data signals. At 250 MHz, each port transmits data at 500 Mbytes/s. RingLink devices also have a low-speed LinkOn signal in and out for controlling power-down modes (see Section 8.3.3). In total, RingLink has eleven to thirteen high-speed receivers, eleven to thirteen high-speed drivers, and one low-speed signal in. All signals are differential.

Figure 8-16 shows the presumed internal structure of a RingLink DRAM. The clock reference input is used as a stable frequency source, and the strobe input is used to provide the sample point for the incoming data bits. An on-chip Phase-Locked Loop generates the local clock from the reference

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 307
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

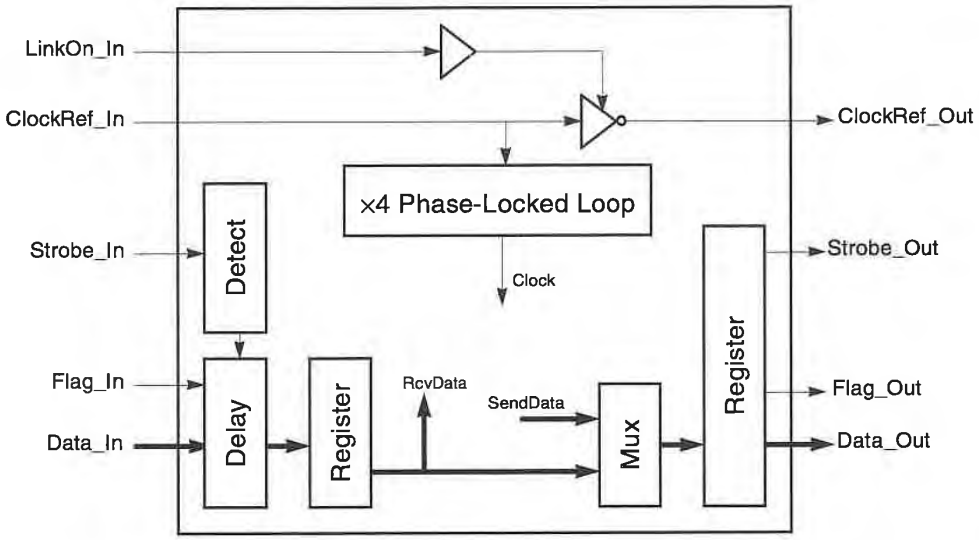


Figure 8-16 A RingLink slave will generate its local clock from the input clock reference.

Signal Name	Pins	Description
$D\_In_{0..8}$	18	A 500-Mbytes/s data stream from the previous RamLink device
Flag_In	2	Signals the start and end of incoming packets
Strobe_In	2	Incoming 250-MHz clock for D_In and Flag_In
ClockRef_In	2	Incoming quarter-speed reference frequency clock
Select_In	2	Optional chip select (1/2 speed)
LinkOn_In	2	A low-speed power-down control signal
$D\_Out_{0..8}$	18	A 500-Mbytes/s data stream to the previous RamLink device
Flag_Out	2	Signals the start and end of the outgoing packets
Strobe_Out	2	A 250-MHz clock synchronized with D_Out and Flag_Out
ClockRef_Out	2	Optional daisy-chained chip select
Select_Out	2	A low-speed power-down control signal (1/2 speed)
jTag	4	Optional standardized test port

Table 8-3 Each RingLink port has one low-speed and ten high-speed signals.

Samsung Ex. 1097, p. 308
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

clock, and delay blocks line up the incoming data to the internal clock.

8.3.2 Physical, Electrical, and Temporal Interfaces

The RingLink definition is not a physical-level specification in the traditional sense because it does not define a package or the specific pinouts. The working group felt that physical packages are properly set by JEDEC and the marketplace. For similar reasons, many details of the temporal interface, such as the setup and hold times, access times, and minimum response times, are being left to the semiconductor vendors that could eventually implement RamLink DRAMs.

For the electrical interface, the group chose the **Low-Voltage Differential Signals (LVDS)** standard developed by the IEEE P1596.3 Working Group as an alternative to ECL levels for the Scalable Coherent Interface (SCI) (see Section 5.2.2). To summarize, this signaling scheme uses full differential signals for all high-speed communication. Therefore, each high-speed signal takes two pins on a package, making 54 nonsupply pins the minimum on a RingLink device. However, differential drivers require fewer supply pins than their single-ended counterparts.

LVDS uses terminated transmission lines and 250-mV differential swings nominally centered around 1.2 V above ground. To account for ground shift and noise, receivers must accept input signals centered anywhere between 0 and 2 V. This proposed signaling standard can be used over relatively long distances at high speed and with modest power dissipation in the drivers. The implication is that a RingLink ringlet can meander around a motherboard to all the memory and I/O components without difficulty. With appropriate ribbon cables, a ringlet could even leave the motherboard and return.

With the exception of the LinkOn signal, all of the RamLink differential signals are parallel-terminated at the destination.

The SyncLink alternate pin-level and electrical interface

The primary RamLink topology is a ring. Rings suffer from increased access latency as the number of elements on the ring increases. For small systems that are very sensitive to long access latencies, the RamLink standard includes an alternate physical and electrical interface called SyncLink. The SyncLink topology is akin to two parallel Rambus Channels: the first link

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 309
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

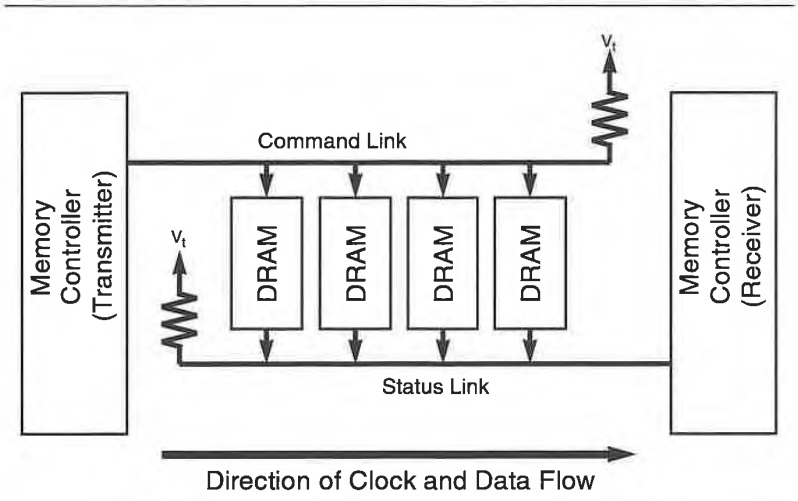


Figure 8-17 Data and clock flow in the same direction on both SyncLink buses.

is used for controller-to-slave communications, and the second one is used for slave-to-controller replies. However, SyncLink is unlike Rambus in its use of a single clock for data flowing to and from the controller. This is accomplished by conceptually having the controller at both ends of the bus. Having two buses simplifies the scheduling of returned data and doubles the total available bandwidth. SyncLink buses use small-signal-swing drivers and parallel termination at one end of the bus (see Figure 8-17). The command link is narrow (8 bits) and fast, while the status link is wider (16 bits) and slower (one-half frequency). The RamLink standard does not specify the important physical and electrical details that ensure interchangeability and interoperability at high frequencies. To reiterate, the SyncLink defined in the RamLink standard is similar to but different from the proposed SyncLink standard described in Section 8.4.

8.3.3 Logical Interface

RamLink is primarily a logical interface. Its purpose is to enable a group of RAM and I/O devices to work together at a logical level. It leaves the detailed physical, temporal, and electrical details to be worked out in the marketplace. The logical interface and bus protocols are designed to work essentially unchanged, regardless of whether a RingLink, SyncLink, or other topology is used.

Samsung Ex. 1097, p. 310
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

As with most protocols that are consistent with ring network interconnects, RamLink's logical interface is based on the concept of a packet. A packet is a group of adjacent bytes transmitted on a link. Each byte has an implicit meaning. RamLink defines six types of packets:

Request Packets:
: Read and write Request Packets are sent from the master to a slave to initiate a data transfer. Write Request Packets include the data to be written.

Response Packets:
: In response to a Request Packet from the master, DRAM slaves send back a Response Packet with a positive acknowledgment (Ack) and possibly some data.

Retry Packets:
: Slaves return a Retry Packet (Nack) instead of a Response Packet when they are unable to satisfy a request within the time allotted by the master.

Idle Packets:
: The controller uses Idle Packets to fill up otherwise idle time on the ringlet. Slaves use the Idle Packets flowing past them on the ringlet to transmit interrupts back to the master.

Event Packets:
: Like Idle Packets, Event Packets initiate events in slaves that do not need responses.

Copy Packets:
: Copy packets—an optional extension to the protocol—can be used in conjunction with memory-mapped I/O.

The packet formats as transmitted on RingLink are shown in Figure 8-18. All packets are an even number of bytes and start on the first half of a clock with a logic 1 on the Flag signal. The fields in the packets are defined as follows:

slaveID: A 6-bit slave identifier that is zero for the first slave in the ring and increments for each successive slave.

transID: Up to four transactions can be in progress per slave at any one time. The transaction identifier, transID, disambiguates the transactions and their responses.

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 311
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

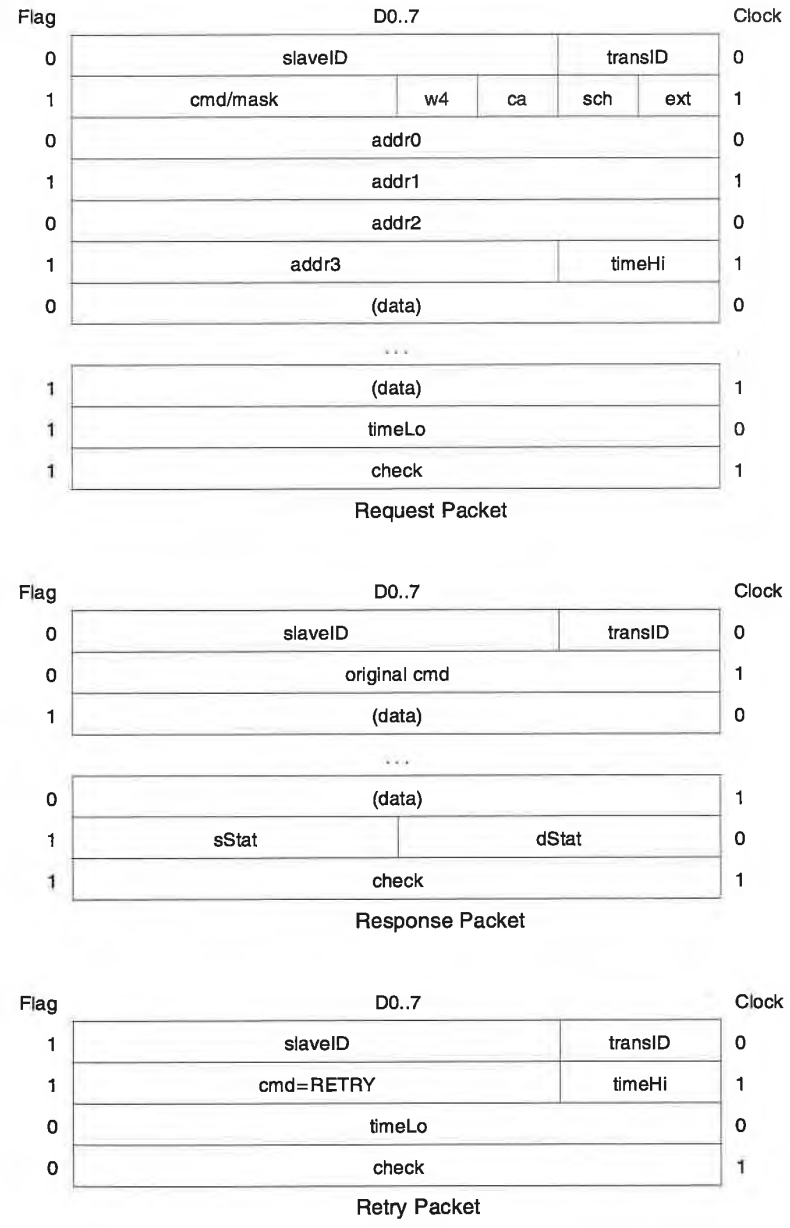
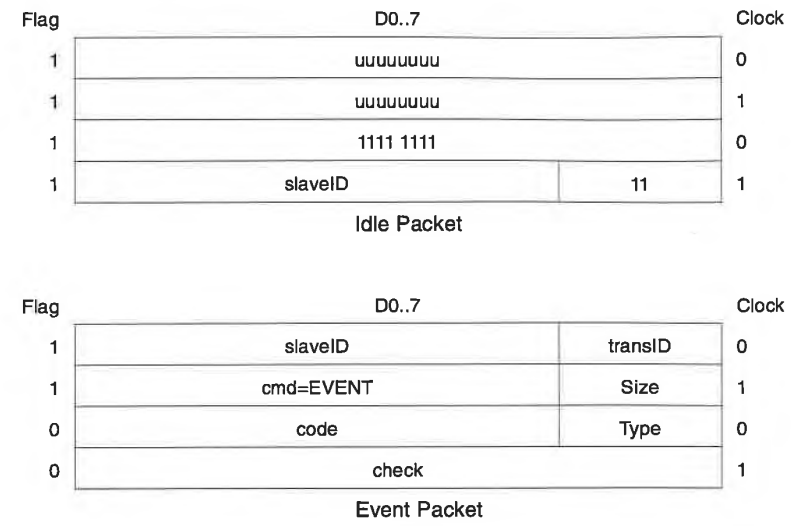


Figure 8-18 The five RamLink packet types range in length from 4 to 12 control and address bytes plus up to 64 data bytes. (*Continues next page*).

Samsung Ex. 1097, p. 312
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454



Flag	D0..7		Clock
1	uuuuuuuu		0
1	uuuuuuuu		1
1	1111 1111		0
1	slaveID	11	1

Idle Packet

Flag	D0..7		Clock
1	slaveID	transID	0
1	cmd=EVENT	Size	1
0	code	Type	0
0	check		1

Event Packet

Figure 8-18 (*Continued*).

w4: Selects whether the cmd/mask field contains a command or a byte mask for a masked-write command. Table 8-4 lists all the commands.

ca: This single-bit, nonbinding cache hint indicates whether the next access to the addressed DRAM is likely to be to the same row. If set, the bit indicates a spatially nonsequential access is likely, so the DRAM array should be precharged.

ext: When set, the ext bit indicates that an extended Request Packet format with 48 bits of address is being used.

addr0..3: Normal Request Packets contain 30 bits of per-slave quad-byte address. The low-order two bits of the byte address are always zero.

sch,timeHi..Lo:
The 10-bit time field indicates precisely when the slave is supposed to return a Response or Retry Packet. If the sch bit is set, the time field specifies the maximum time that the slave has to respond. In this case, the master will keep the ringlet clear for the slave to respond any time up until the specified period has elapsed. In Retry Packets, the time field indicates the minimum amount

© 1996 Steven Przybylski

of additional time required to complete the request successfully.

data: The data portion of Request and Response Packets contains 4, 8, 16, 32, or 64 bytes of data, as determined by the request command code. Request Packets only contain data if they are write requests. Slaves do not generate Response or Retry Packets in response to write requests.

sStat: The sStat field in Response Packets contains a standard specified status code that indicates normal or abnormal completion of the request and the type of error in the case of an abnormal completion.

dStat: The dStat field is a vendor-specified status field that varies from device to device.

code: In Event Packets, the code contains the event or command that the addressed slave is supposed to act upon.

check: All packets except Idle Packets are trailed by a check byte that verifies the integrity of the address, command, and data information in the packet.

cmd/mask				w4	ca	Command	Description
0	0	0	0	0	x	Read0	Precharge or refresh
0	0	0	1	0	x		Reserved
0	Size			0	x	Readxx	Read xx bytes, $xx = 2^{Size}$, $4 \le xx \le 64$
0	1	1	1	0	x	Response	Response Packet with data or status
1	0	0	0	0	0	Retry	Retry Packet with extra time needed
1	0	0	0	0	1	Event	Responseless Request Packet
1	0	0	1	0	x		Reserved
1	0	1	0	0	x	Copy64	Read64, generating a Write64
1	Size			0	x	Writexx	Write xx bytes, $xx = 2^{Size}$, $8 \le xx \le 64$
1	1	1	1	0	x		Reserved
Mask				1		WriteSb	A 4-byte write-under-mask.

Table 8-4 The RamLink command field encodes the packet type as well as the operation type and size.

The RingLink ringlet is a finely scheduled resource. In order to reduce the latency through the slaves, the protocol is constructed so that slaves pass on incoming bytes unconditionally, without interpreting them. This means they only drive the link

Samsung Ex. 1097, p. 314
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

when specifically directed to by the master. The link master is responsible for scheduling all the packets on the link so that a slave does not blindly overwrite a valid packet when driving a Response or Retry Packet. Consequently, read Request Packets indicate specifically when the Response or Retry Packets are expected. If the slave is to respond as soon as possible, then the ring must be kept clear of other valid packets until the slave responds.

An advantage of this strict RingLink, or SyncLink, scheduling is that the total bandwidth of a ring is more than 500 Mbytes/s. By arranging to have downstream slaves overwrite the Request Packets to upstream slaves with Response Packets, the peak ring bandwidth can be doubled to 1 Gbyte/s. However, taking advantage of this possibility calls for some very careful scheduling. Some memory controllers would be unable to perform such detailed scheduling.

Per-slave latency minimized

One exception exists to the rule that slaves need not interpret packets as they flow past: the case of interrupts in Idle Packets. Since RingLink is meant as a general purpose interconnect, slaves can piggyback interrupt requests on Idle Packets. In order to send an interrupt to the master, a slave watches for the first byte of an Idle Packet. When the fourth byte comes by, instead of forwarding it unchanged, the slave inserts its slave ID. With this scheme, a downstream slave could overwrite an upstream slave's interrupt request with one of its own. In this respect, the scheme is logically equivalent to a daisy-chained interrupt line. The interrupt handlers must be cognizant of this possibility and behave appropriately.

Since packets do not flow through the slaves in a SyncLink system, the Idle Packets cannot be used to transmit interrupt information back to the host. Instead, separate point-to-point interrupt lines between the slaves and the memory controller are anticipated.

RamLink supports reads and writes of up to 64 bytes. Writes of fewer than 4 bytes are accomplished by sending 4 bytes of data and a byte mask in the Request Packet's cmd/mask field. Though nominally defined as an 8-bit byte, the proposed standard provides for 9-bit devices as well. The 9-bit devices can be used on an 8-bit ringlet, but the ninth bit of the wide device becomes inaccessible. The 8-bit devices cannot be used on a 9-bit ringlet since they would not pass through the ninth data bit with the same timing as the others. The 8-bit devices would be used in systems with no ECC or in systems in which ECC is handled on

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 315
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

the DRAMs; the 9-bit devices would be used when ECC is implemented in the memory controller. ECC will probably not be implemented on the DRAMs for some time.

Although 64 slave identifiers are available, the protocol limits the number of slaves to 62. The other two slave IDs are used for addressing the master and for broadcasts. Regular Request Packets are able to address 4 Gbytes per slave. In compliance with the IEEE Standard 1212-1991 Control and Status Register (CSR) Architecture, a full 256 Mbytes at the top of each slave's address space is reserved for control and status registers. However, simple RamLink slaves need not implement the CSRs. This contrasts with the Rambus approach, which uses separate read and write commands to get access to the configuration registers.

In addition to the regular Request Packet shown in Figure 8-18, an extended format request adds two additional bytes of per-slave address and two reserved bytes for future expansion. Using extended Request Packets, the per-slave addressability is increased to 281 terabytes, with a total of 512 Mbytes of CSR space.

Explicitly addressing slaves with a slave identifier has two consequences. First, it is very easy for a slave to determine that it is being addressed. Second, if the memory controller is to present a contiguous memory space to the system processor, it must create a mapping from the processor's visible address space onto the address spaces of the individual RamLink DRAMs under its control. Most modern operating systems do not require a contiguous main memory, but many older operating systems, including all versions of MS-DOS, do require a contiguous main memory.

Event Packets

Event Packets reset or initialize the slaves, control power-down and auto-refresh modes, and diagnose the ringlet or bus. The RamLink specification defines two low-power modes: Standby and Shutdown. In Standby Mode, a slave feeds the flag and clock signals to the next slave, but turns off its data receivers and transmitters. Standby Mode is exited when a logic 1 is propagated down the flag wire. In Shutdown Mode, all the high-speed receivers and drivers are turned off. Shutdown slaves stay in this lowest power mode until the LinkOn pin is asserted. DRAMs must refresh themselves when in either reduced power state. An unfortunate consequence of a ring topology is that all the devices on a ringlet must be fully turned on for accesses to be made to any device. The ability to selectively shut down

Samsung Ex. 1097, p. 316
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

parts of the memory system is possible only with a multiple-RingLink system or with a SyncLink topology.

Under normal circumstances, RamLink DRAMs autonomously refresh themselves. However, in order to ensure a shorter maximum response time, DRAMs can be manually refreshed via the Read0 command or via a Refresh-Now event broadcast to all DRAMs.

A common complaint about the ring topology of the RamLink pin-level interface is that each slave in the ring adds additional latency to all operations on the ring. Early versions of the proposed RamLink standard had a minimum amount of logic between a slave's input port and its output port. This minimized per-slave latency in the ringlet. Unfortunately, in order to deal with potential clock skew and jitter problems, more registers and explicit delays were added to the flow-through path. The unpleasant consequence is a greater per-device ring latency. Even with a relatively optimistic 4- to 8-ns per-slave delay, the total ring latency can be significant if more than a few devices are in the primary ring. This latency can be dealt with in three ways:

- Limit all rings to only a few devices
- Use the SyncLink pin-level interface
- Ignore it

The last of these three solutions relies on the observation that ring latency is not a problem in heavily used memory systems because it is small in comparison to the wait in the memory controller for access to the ring. This solution unfortunately assumes that the memory system is a performance-bottleneck in the system—a situation the industry is trying to avoid as it moves away from generic DRAMs.

8.3.4 Current Status and Future Developments

Though several DRAM vendors showed some interest in RamLink, most of RamLink's early development was accomplished by volunteers from system vendors. The RamLink proposal was essentially complete by 4Q1993, and it was submitted for formal ballot in 2Q1994. The ballot committee[17] revised the standard in late 1994, and it was resubmitted to the working group

17. Including the primary author of this report.

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 317
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

for a second vote in mid 1995. RamLink passed the second vote unanimously.

Pending intellectual property issues

At that point the proposed RamLink standard was passed to the IEEE for ratification. However, the IEEE has delayed formal ratification pending intellectual property concerns, especially—but not exclusively—with respect to Rambus. As of this writing, this issue continues to hang over RamLink, and by inference, SyncLink.

Even if RamLink is ratified by the IEEE, such standardization does not mean that RamLink is free of intellectual property constraints. It may be that any implementation would infringe on one or more preexisting or currently pending patents. If this proves to be the case, then one of RamLink's chief advantages—its nonproprietary nature—would be compromised.

Physical and electrical interfaces to be set by the market

The heart of the RamLink proposal is the logical interface. Other physical, electrical, and temporal interfaces may be adapted to the RamLink logical interface as market conditions dictate. This lack of complete specification may result in even more confusion than is currently evident in the SDRAM arena. Failure of the DRAM vendors to agree on these important details could significantly delay adoption of the standard.

Because the RamLink proposal is still so new, no one has yet designed a RamLink interface. Though some experience from the SCI case can be carried forward to RamLink, uncertainty exists regarding the RamLink interface's overhead and its impact on the price premium.

RamLink eclipsed by SyncLink

In late 1994, several elements of the DRAM community took notice of RamLink and expressed a desire to refine it and promote it as a commercial product. However, rather than modify the already-complete RamLink, it was decided to create a new and separate standard based on, but not identical to, the SyncLink physical layer defined in the RamLink annex. Thus was born the SyncLink standard described in Section 8.4. With all the interest and effort currently focussed on the new SyncLink, RamLink is at this point destined to obscurity. There is no significant probability that any RamLink/RingLink or RamLink/SyncLink DRAMs will ever be designed or sold. More likely, if RamLink is implemented, it would be as a general board-level interconnect mostly handling I/O traffic.

Samsung Ex. 1097, p. 318
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

Interestingly, a close cousin of RamLink is MicroUnity's MediaChannel. MediaChannel is very much akin to a minimalistic RamLink protocol on a reduced RingLink physical layer. The 8-bit wide ring can have up to 4 chips in it. Headers are 1 byte of command, 0 or 4 bytes of address and 0 or 8 bytes of data followed by a CRC check byte. The channel also uses low-swing (600 mV) differential signaling and a flow-through architecture. Each ring can generate up to 1 Gbyte/s of bandwidth.

8.3.5 Strengths and Weaknesses

Strengths

- High bandwidth possible from just a few DRAMs
- Minimum memory system granularity
- Nonproprietary standard
- Generalized interconnect
- Not physically constrained
- Elegant sequential/nonsequential hint mechanism
- Dual RamLink/SyncLink physical and electrical levels

Weaknesses

- Overtaken by SyncLink
- Long latency
- Latency increases with RamLink ringlet size
- Physical and electrical interfaces incompletely specified
- To access any device in a RingLink ring, all devices must be active

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 319
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

8.4 SyncLink DRAMs (SLDRAMs)

Section 8.3 describes the IEEE Ramlink standard, which defines 1) a logical protocol (RamLink) for memory controllers to talk to DRAMs and other slaves, and 2) two physical layers—RingLink and SyncLink. In late 1994, just as RamLink was becoming a standard, discussions were begun within the IEEE Computer Society P1596.4 working group on a potential successor. These discussions, which for the first time attracted significant interest from the DRAM community, soon focussed on RamLink's SyncLink physical layer as a natural starting point. Rather than delay ratification of RamLink, a separate working group, P1596.7, was formed to develop a new, though derivative, DRAM interface. This new interface, also called SyncLink, has evolved well beyond its RamLink/SyncLink heritage. SyncLink is a work in progress. The description presented here is accurate as of version 0.77 of draft standard, dated January 3, 1996.

At the root of many of RamLink's problems was its development by a working group of volunteers. The reliance on volunteers slowed the effort and limited the resources available for research. Many design decisions were based on intuition and design experience rather than on experimental analysis.

SyncLink Consortium and Trade Association

In order to save SyncLink from a similar fate, several companies interested in ensuring the existence of a nonproprietary high-bandwidth DRAM interface formed a SyncLink Trade Association to help drive the development of the standard. The SyncLink Trade Association intends eventually to be a broadly-based and open Association that promotes SyncLink and supports its users and manufacturers, providing a forum for rapid industry standardization.

The SyncLink Trade Association in turn formed a SyncLink Consortium in October 1995. The consortium's initial set of charter membership consisted of Hyundai, Mitsubishi, and Texas Instruments. These have since been joined by Samsung, Micron, Fujitsu, MOSAID, and Apple. The consortium is open to additional members on an equal and nondiscriminatory basis, and at this writing, several other large systems vendors are seriously considering membership. However, the price of membership to new members is not fixed. It is expected to rise over

Samsung Ex. 1097, p. 320
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

time as SyncLink moves through the standardization process into the commercial world.

The SyncLink Consortium was formed to provide a mechanism for holding closed meetings in order to handle confidential proposals and to develop intellectual property. That intellectual property would be shared by the consortium members. The consortium could choose to share that intellectual property with the Trade Association members, or even the general public, if that would speed the adoption of SyncLink. The consortium is technically a nonprofit organization. Any income that it should derive in future from its intellectual property will go to support the SyncLink Trade Association.

Initially, the consortium's priority is to complete the design studies needed to set the IEEE Working Group in the most promising direction. Much of this work is being carried out by MOSAID, an experienced and highly-respected DRAM design house based in Kanata, Ontario, Canada. At present, little information about the Consortium proceedings and plans has been made available to those outside the consortium

The SyncLink Trade Association is a part of SCIzzL, the Association of Scalable Coherent Interface (SCI) Local-Area MultiProcessor Users, Developers and Manufacturers. SCIzzL's charter includes providing start-up assistance for all SCI-related groups. SCIzzL is a nonprofit organization operating under the corporate umbrella of Santa Clara University, of Santa Clara, Calif. SCIzzL has been administrating all P1596.x Working Group meetings, which include those on SCI (now ANSI/IEEE Standard 1596-1992) and its derivatives, including P1596.3, Low Voltage Differential Signals (LVDS, see page 135), P1596.4, RamLink, and now P1596.7, SyncLink.

The SCIzzL membership fee is currently $25,000 per year for Sponsors, and $50,000 per year for Executive Sponsors. As the early SCIzzL membership is dominated by companies with an interest in SyncLink, a large part of the fees garnered to date have gone to start up the SyncLink Consortium. In the future, however, companies with other SCI-related interests may SCIzzL. In this eventuality, the SyncLink-related activities will be accounted separately from the SCI-related activities.

The SyncLink Consortium and SyncLink Trade Association meetings are closed, while IEEE Working Group meetings are by definition open to the public. Initially, it was intended that the consortium and trade association would exist merely to

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 321
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

support the working group and to speed its efforts. However, as the effort has gathered steam, the consortium is playing a more active role in deciding the shape of the final result; the working group is becoming more of a forum for public ratification of the consortium's proposals.

Broad project goals

The SyncLink Consortium and IEEE Working Group are trying to design a DRAM and interconnect architecture that will define a commodity DRAM at the 64-Mbit and larger density levels. In order of importance, the three objectives are:

- Develop an interface that is applicable to a wide variety of future DRAM applications
- Facilitate multiple concurrently active requests
- Support other RAM-like devices, such as ROM, Flash memory, memory-mapped I/O, and bridges to other interconnects

Overall, the first objective is most important and the third is very much less important than the first two. Consequently, it will be acceptable if SyncLink ends up being suboptimal for connecting non-DRAM devices. The first goal is being achieved by migrating as much complexity as possible out of the SLDRAMs into the memory controller and by keeping the SLDRAMs as simple as possible, but not simpler.

8.4.1 Bus Structure and Pin-Level Interface

SyncLink's bus structure derives from the SyncLink physical layer of the RamLink standard. Two buses connect a collection of DRAMs with a memory controller. One bus, dubbed the SendLink, is dedicated to communication from the memory controller to the DRAMs. It is a narrow (8 or 9 bit), high-speed (500 MHz or greater) bus that is parallel-terminated at the end further from the controller. The SendLink carries command, address, and write data. The second bus, called the SinkLink,[18] is dedicated to traffic from the DRAMs to the controller. Since it is a multi-source/single-destination bus, the SinkLink is fundamentally harder to operate at high speeds than the SendLink. Consequently, this bus is operated at one-half the speed and

18. The bus (SinkLink) and the standard (SyncLink) are homophones—they are pronounced alike but are spelled differently. SinkLink (the bus) is spelled with an i and a k, while SyncLink (the standard) is spelled with a y and a c. This makes for interesting sentences when talking about the SyncLink SinkLink.

Samsung Ex. 1097, p. 322
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

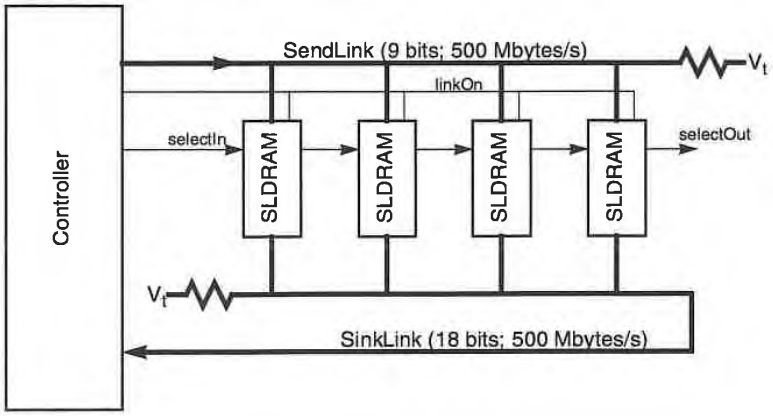


Figure 8-19 SyncLink defines a narrow, high-speed SendLink and a wider, slower SinkLink to connect the memory controller to the SLDRAMs.

twice the width of the SendLink: 250 MHz and 16 or 18 bits wide (see Figure 8-19).

Figure 8-20 shows the conceptual pinout of an SLDRAM. In addition to the two buses, each DRAM receives a high-speed clock, a Flag signal, a Select signal, and a Link-On signal. The input clock runs at the SinkLink rate and in phase with the data arriving on the SendLink. The Select signals are daisy-chained through all the SDRAMs on the link and are used during initialization and synchronization of the system. A JTAG interface is optional at this level of the specification.

The SinkLink is defined as two bytes wide: an A byte and a B byte. Labeling the byte lanes helps clarify the address of bytes on the SinkLink and facilitates the definition of both 8/16- and 9/18-bit SLDRAMs. SyncLink allows the mixing of 8/16 and 9/18-bit SLDRAMs on the same link. A single package and module is to be used by both widths to avoid any impedance mismatches on the 9th bit of the SendLink or the 9th and 18th bits of the SinkLink.

That the SinkLink has to double back in order to return to the memory controller is a cause for concern. This is hard to do without any noticeable impedance discontinuity or length variation across the bus. The difficulty of this task is eased somewhat by the lower speed of the SinkLink, but aggravated by its greater width.

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 323
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

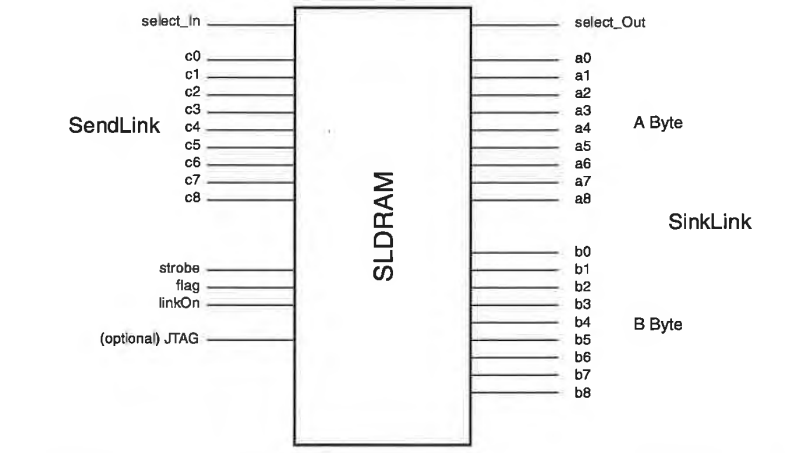


Figure 8-20 SyncLink defines two buses and a couple of additional clock signals per DRAM.

8.4.2 Physical, Electrical, and Temporal Interfaces

Traditionally, DRAM packages and electrical interfaces have been the domain of two JEDEC subcommittees. The SyncLink Working Group and the Consortium intend to leave these aspects out of the interface in the IEEE SyncLink standard. Instead, the members of the SyncLink Consortium will at some point take a proposed SLDRAM package and module definition to JEDEC for standardization. As of this writing, no specific packaging or module plans have been developed. The only general statement regarding modules is that it is expected that the standard memory expansion mechanism will be a small module with one SLDRAM and an EEPROM on it. The EEPROM could provide details of the SLDRAM's organization and temporal characteristics to the memory controller during system initialization. Details of this scheme have yet to be proposed.

The reason that the number of SLDRAMs per module has to be limited to one or two is that the data flows synchronously on the two links. Any physical structure that provides for multiple SLDRAMs on an expansion module will make it very difficult to

Samsung Ex. 1097, p. 324
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

maintain the phase relationship of data on the two links on the motherboard.

This limitation of the number of devices per module has a significant impact on the system design. Specifically, a socket must be provided for every SLDRAM in the maximally configured system. Without the use of transceivers—of which there is no mention yet—the ratio of minimum memory size to maximum memory size will be lower than is currently available to memory system designers. Furthermore, at the 64-Mbit level, a significant number of devices will be needed in a large memory system. Providing the sockets and space for the maximum memory system will burden the minimum memory system.

As of 4Q1995, the proposed electrical interface for SyncLink is SSTL, a derivative of CTT, as defined by JEDEC (see Section 5.2.2). This low-swing voltage-mode interface facilitates high speed and reasonably good signal integrity in a single-ended parallel-terminated transmission line environment. Both buses are terminated to V_T through a termination resistor matched to the characteristic impedance of the bus.

Similarly, the detailed timing parameters of DRAMs have traditionally been set by a loose consensus among the DRAM vendors and major consumers. The Consortium's intent is to let the same loose consensus ultimately determine the timing parameters of SyncLink DRAMs. The process is very similar to what was followed for SDRAMs: initial SDRAM designs varied significantly across vendors; subsequent iterations migrated closer and closer to a common timing specification.

8.4.3 Logical Interface

The logical definition of an SLDRAM includes a description of both the structure of the RAM and the format and semantics of the protocol used to communicate with the RAM.

SLDRAMs are logically divided into a number of distinct and independent sub-RAMs (see Figure 8-21). The number of sub-RAMs per device is not predetermined and can vary across generations and across products. The sub-RAMs are not cognizant of the other sub-RAMs on the device or link: it is the controller's responsibility to make sure that they do not conflict with one other. Specifically, the controller must ensure that no two sub-RAMs try to drive the shared SinkLink simultaneously.

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 325
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

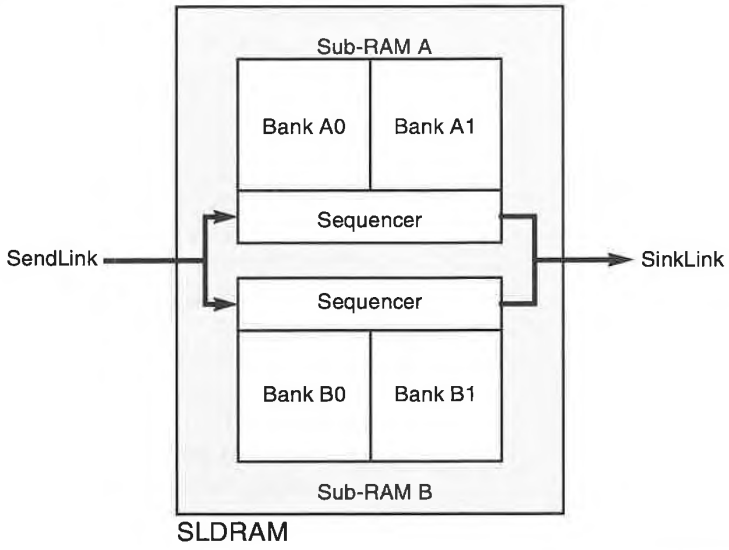


Figure 8-21 SLDRAMs contain one or more sub-RAMs, each of which consists of a sequencer and one or more banks of memory.

Each sub-RAM contains a sequencer and one or more banks. The multiple banks within a sub-RAM are allowed to be open simultaneously, but concurrent read or write transfers to more than one bank in a sub-RAM are not permitted. For example, a single SLDRAM could accept write data to a bank in one sub-RAM while driving read data from another bank in a different sub-RAM, but these two transactions could not be concurrent if both banks were in the same sub-RAM, regardless of whether the sub-RAMs are in the same SLDRAM.

Memory controllers communicate with SLDRAMs via packets on the SendLink. The start of a packet is indicated by a change in the Flag signal. Packets start with two bytes of header that identify the sub-RAM being accessed and the command to be performed. Packets specifying read and write commands contain two bytes of column address; packets specifying OpenRow commands include four bytes of address. Other command types have no need for any address bytes. Figure 8-22 shows the general packet format, while Table 8-5 shows all the command encodings.

Samsung Ex. 1097, p. 326
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

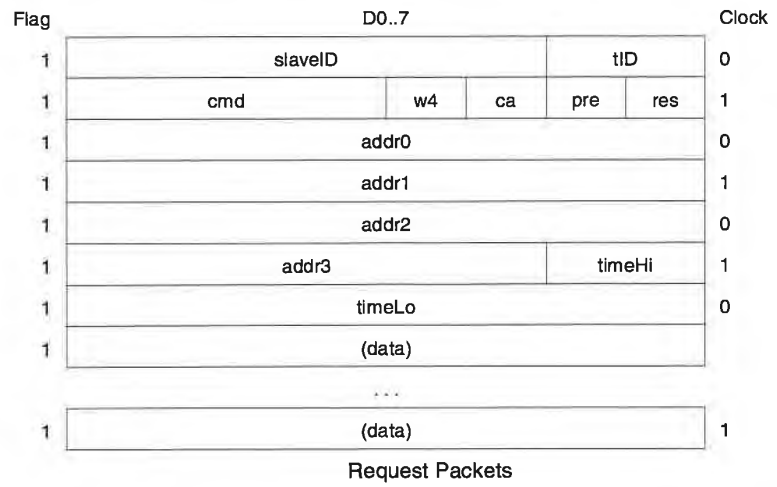


Figure 8-22 SyncLink request packets have two bytes of header, two or four bytes of address, and zero or one byte of timing information.

cmd/mask				w4	ca	Command	Description
0	0	0	0	0	x	Read0	Open row/refresh
0	0	0	1	0	x		Reserved
0	Size			0	x	Readxx	Read xx bytes, $xx = 2^{Size}$, $4 \le xx \le 64$
0	1	1	1	0	x		Reserved
1	0	0	0	0	0		Reserved
1	0	0	0	0	1	Event	Responseless request packet
1	0	0	1	0	x		Reserved
1	0	1	0	0	x		Reserved
1	Size			0	x	Writexx	Write xx bytes, $xx = 2^{Size}$, $8 \le xx \le 64$
1	1	1	1	0	x	Retry	Reserved
Mask				1	x	WriteSb	A 4-byte write-under-mask.

Table 8-5 The SyncLink command field encodes the packet type as well as the operation type and size.

The similarity of Figures 8-16 and 8-22 and of Tables 8-4 and 8-5 illustrate SyncLink's RamLink heritage. Many of the request packet fields are similar to those found in RamLink:

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 327
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

slaveID: A 6-bit slave identifier indicates which SLDRAM on the link is being addressed. The one closest to the controller on the SendLink is device 0, with device numbers increasing with each successive SLDRAM.

tID: The tID field identifies which sub-RAM on the addressed device is being accessed. At this writing, the proposed standard calls for up to four sub-RAMs per device, but the boundary between the slaveID and tID fields is quite arbitrary. At most one active command per sub-RAM is allowed at a time.

w4: Selects whether the cmd/mask field contains a command or a byte mask for a masked-write command.

ca: This single-bit cache hint indicates whether the next access to the addressed DRAM is likely to be to the same row. If set, the bit indicates that a spatially non-sequential access is likely, so that the DRAM array should be precharged.

pre: When set, the pre bit indicates the command is operating on an open bank, and so only two bytes of address are included in the header.

addr0..3: Normal Request Packets contain 30 bits of per-slave quad-byte address. The low order two bits of the byte address are always zero and not transmitted.

timeHi..Lo:
The 10-bit time field indicates precisely when the slave is supposed to return read data to the controller on the SinkLink.

Five request types In all, there are five different request types:

OpenRow:
The addressed bank is opened to the addressed row. If the bank is open at another row at the outset, that row is closed and the bank precharged first. The full four bytes of address are supplied. However, no time information is needed, so the TimeLo byte is not sent. The net result is a six-byte request packet.

ReadofOpen:
Perform a column read access of a block of data from an open bank. Only the low order two bytes of address are transmitted, making for a five-byte request packet. The device and sub-RAM are identified in the first byte of

Samsung Ex. 1097, p. 328
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

the request, and the bank and column address of referenced data are contained in the two bytes of address.

ReadofClosed:
The ReadOfClosed request type allows read operations to be initiated to banks that are currently not open to the addressed row. If the bank is precharged or open to a different row, then the appropriate row access is automatically performed by the SLDRAM. Since both the full address and time information are needed, this request has a seven-byte Request Packet.

Write: Write requests are only allowed to open banks. The write data immediately follows the packet header. No time delay information is needed, so none is provided. Since writes are only allowed to open rows, only the low order two bytes of address (including the bank and column addresses) are provided in the packet header. No timing information is needed either, so the total packet size is four bytes for the header plus the data to be written.

Event: Event packets convey commands to SLDRAMs that do not involve data movement. The control slaveID initialization, a reduced power state called StandBy, and auto-refresh. No address or time information is sent; a single-byte event code and a single byte of data bring the event packet size to four bytes. No data is ever sent back to the memory controller as a result of an Event Request.

Since the entire protocol is quad-byte oriented, the address of bytes with quad-bytes is irrelevant and not transmitted. The only important relationship is the order in which the bytes of a quad-byte go into an SLDRAM on the SendLink relative to the order in which they come out on the SinkLink. This relationship is shown in Table 8-6. The first and third bytes going in on the SendLink always come out on the A byte of the SinkLink, and the second and fourth bytes going in on the Sendlink always come out on the B byte.

Effectively a ×32 DRAM

Because the protocol is quad-byte addressed, an SLDRAM can be thought of as a ×32 DRAM with a total of 24 address and data pins.

Data transfers are 4 to 64 bytes long. All data transfers are linear, and sub-block order (i.e., interleaved) transfers are not supported. All SLDRAMs will support block transfers that start at the beginning of a block and proceed to the end without need for

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 329
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

Clock	Clock Phase	SendLink Order
n	0	First Byte
	1	Second Byte
n+1	0	Third Byte
	1	Fourth Byte

Clock	Clock Phase	SinkLink Order	
		A Byte	B Byte
n	Both	First Byte	Second Byte
n+1	Both	Third Byte	Fourth Byte

Table 8-6 SyncLink defines the relationship between the order of the bytes on the SendLink and that of the bytes on the SinkLink.

wrapping around.[19] For transfers that are eight bytes or longer, SLDRAMs may optionally support transfers that start in the middle of a block, wrap from the end of the logical block, and complete with the block's beginning. For example, a 16-byte read with an initial address of 12 would transmit bytes from addresses (12,13), (14,15), (0,1), (2,3), (4,5), (6,7), (8,9), (10,11). The pairs of addresses in parentheses represent those bytes transmitted simultaneously on the A and B bytes of the SinkLink.[20]

Because all read requests contain the relative time at which the SLDRAM is to drive the requested data onto the SinkLink, it is fairly easy for the controller to intermix read requests and write requests on the SendLink so as to keep the SinkLink heavily utilized—but without adding unnecessary latency.

Partitioning address bits into bank address bits, row address bits, and column address bits is a matter of convention. As the proposal now stands, different SLDRAMs of the same generation could have different partitions. However, the controller must understand the partitioning because it has to keep track of whether a request will be completed with a row access or a column access. If the controller fails to comprehend the

19. This type of transfer is characterized by a start address which has at least the *n* least-significant bits being zero, where *n* is log2 of the transfer size.

20. Again, because the low order two address bits are irrelevant, the addresses of the bytes within the quad-words are arbitrary.

Samsung Ex. 1097, p. 330
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

partitioning of address bits, and loses track of which rows are open, then erroneous behavior results.

Broadcast writes and events

Currently, up to 64 slaveIDs are possible. However, at most 62 SLDRAMs can be present on a link. SlaveID 62 is reserved and used during the initialization procedure. SlaveID 63 is reserved to indicate a broadcast of a request (most commonly an Event Request) to all the sub-RAMs in all the SLDRAMs on the link.

SlaveID initialization

The mechanism by which SLDRAMs are given their slaveIDs is borrowed essentially unchanged from RamLink. A hard reset event is broadcast to all devices, forcing them to set their slaveIDs to 62 and to clear their Select_Out pins. Since the Select_In and Select_Out pins of the SLDRAMs are connected in a daisy chain, they can be used to sequentially initialize all of the devices on the link. A second event, configHard, sets the slaveID of an SLDRAM if the SLDRAM's slaveID is 62 and its Select_In is high. If these two conditions are met, then the slaveID is set to the value specified in the Event Request, and Select_Out is raised. By broadcasting a configHard, raising the Select_In to the first device, and then broadcasting a series of configHard events with increasing slaveID values in the data fields, all SLDRAMs on the link are given unique IDs.

Low and lowest power modes

In addition to the normal operating mode in which an SLDRAM is ready to accept commands, SyncLink defines two other states: StandBy and ShutDown.

The entire memory system is put into a low-power StandBy state by lowering the LinkOn signal. StandBy is intended for short periods of inactivity. It facilitates turning off input receivers and output drivers, but the SLDRAMs remain essentially active. Standby is exited by raising the LinkOn signal a few cycles before the next request on the SendLink. By providing separate LinkOn signals to each SLDRAM, portions of a memory system can be powered down independently.

The lowest power mode is called ShutDown. In this mode, on-chip PLLs are turned off. ShutDown is intended for longer durations of inactivity, so a minimum shutdown time is enforced to ensure clean shutdown entry and exit. Unlike some other architectures, the SLDRAMs do not self-refresh themselves during ShutDown. This means that all the SLDRAMs must be periodically brought out of ShutDown for the controller to perform refresh operations. ShutDown is entered by sending a ShutDown event prior to removing the LinkOn signal. After LinkOn has been low for a time, the clock can be removed. ShutDown is

© 1996 Steven Przybylski

exited by reapplying clock, raising LinkOn, then waiting a predetermined length of time for the SLDRAM's clock circuitry to lock onto the input clock.

Manual- and auto-refresh only

The draft SyncLink standard allows for manual and auto-refresh only. Like all DRAMs, SLDRAMs can be refreshed by manually opening each row in every bank at least once per refresh period. This can be efficiently done with reads of zero bytes. Auto-refresh, in which the DRAM keeps track of the address of the row to be refreshed next, is initiated by sending refreshNow events to the sub-RAMs.

There are actually four flavors of refreshNow events, named refreshNow0 through refreshNow3. Memory controllers use exactly one flavor of them at a time. If refreshNow0 events are used, then the controller will broadcast at least 4,096 of these events every 64-ms interval. This translates to one every 15.6 µs on average–the typical refresh rate for today's DRAMs. In the future, though, the rate at which banks will have to be refreshed is expected to increase. This is because the data retention time divided by the number of rows per bank is declining and will certainly decline below 15.6 µs by the 256-Mbit generation. In the future, refreshNow events will have to be broadcast more frequently. That is where refreshNow1 through refreshNow3 come in. If refreshNow1 events are used, then the memory controller sends at least 4,096 of the events every 32 ms. The corresponding durations for refreshNow2 and refreshNow3 are 16 ms and 8 ms respectively: refreshNow3 events are sent on average once every 1.95 µs.

Devices that only need a refresh event every 16 µs must also work if the higher rate refresh events are used. They do this by ignoring the appropriate percentage of the events and acting on one event in every two, four, or eight events, as the case may be.

8.4.4 Changes from RamLink

SyncLink stands on its own as a complete specification. However, a comparison of the SyncLink that is part of the RamLink definition and the draft SyncLink standard is interesting. Most of these changes resulted directly from the significant input of the DRAM community—input that was very limited during the development of RAMLink. As of draft 0.77 of the proposed standard:

Samsung Ex. 1097, p. 332
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

1. All non-data information coming back to the memory controller has been eliminated. The SinkLink transmits only data bytes.

2. The check bytes on all packets to and from the memory controller have been eliminated.

3. The imprecise scheduling option has been eliminated. The time that data is to be driven on the SinkLink is precisely specified in the Request Packets.

4. There is now no possibility of error and retry. If the data is not available when the memory controller wants it, then the memory controller has made an error, and the behavior is undefined.

5. Unlike RamLink, SyncLink does not provide for any prespecified registers on the DRAM. All information needed to respond to read and write commands is encoded in the commands. In fact, there are no addressable entities on-chip other than memory. The CSRs that optionally specified the DRAM's characteristics have been eliminated. An on-module PROM must be added to communicate device capabilities to the master.

6. Data transfers to or from open rows now only transmit two bytes of address, instead of the full four bytes.

7. The extended addressing model has been eliminated.

8. Self-refresh is eliminated other than in ShutDown mode.

9. The ca bit is no longer a nonbinding hint; it is now a binding precharge indicator.

10. Sub-RAMs and banks have been added.

Overall, these changes make the DRAMs simpler and more predictable. Consequently, the changes mandate that the controller be much more aware of the DRAM's internal organization and precise state. This is a return to the normal situation that exists for the conventional and evolutionary alternatives. Fundamentally, the memory controller needs to be cognizant of the DRAMs' internal state whenever the memory controller needs to make the maximum use of the DRAMs' capabilities.

8.4.5 SLDRAM Fill Frequency

A great deal of uncertainty still exists about the final configuration and speed of any commercial SyncLink parts. However, if

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 333
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

current plans to run both the SendLink and the SinkLink at 500 MBps come to fruition, the total peak bandwidth of an SLDRAM comes to 1 GBps. At the 64-Mbit density, this would translate to a peak fill frequency of 125 Hz—higher than that of any currently available DRAM. As with all the DRAM architectures, the fill frequency will decline by nearly a factor of four—to 32 Hz—at the 256-Mbit level.

Discussion of the sustainable fill frequency will have to wait until more information is available about the internal characteristics of real SLDRAMs.

The characteristics that determine an SLDRAM's sustainable fill frequency—the number of sub-RAMs, the number of banks, their latency, and so forth—are still undetermined. Consequently, any discussion of the gap between a future SLDRAM's peak and sustainable fill frequency would be generally premature. It is safe to say, though, that with the introduction of sub-RAMs and of two independent, unidirectional buses, the sustainable fill frequency could be very close to the peak fill frequency under many circumstances.

8.4.6 Current Status and Business Issues

As of 4Q1995, the SyncLink Consortium has grown to seven members, and the SyncLink Consortium and—to a lesser extent—the IEEE Working Group have been meeting regularly to resolve the outstanding architectural issues. Many such issues remain essentially in all aspects of the interface. Though most remaining issues are fairly minor, the SyncLink described here may still change substantially.

Despite the Consortium's financial resources, the working group continues to be essentially a volunteer effort. A relatively small amount of effort is being put in on a month-to-month basis. Given the pace of development, a ballottable draft will probably be completed in 1996. Separately, much work is still needed on the physical, temporal, and electrical interfaces.

Once the SyncLink architecture is more solidified, work can begin on implementations. Given the normal design and debug cycles of the DRAM industry, it will be late 1997 at the very earliest before any samples are available. If the Consortium pursues the SDRAM model of allowing the market to determine the acceptable temporal specification, another 18 to 24 months are likely to pass before a degree of uniformity and cost-competi-

Samsung Ex. 1097, p. 334
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

tiveness is developed. Realistically, the first SLDRAMs will be 64-Mbit devices; widescale adoption may have to wait until the 256-Mbit level.

The list of companies "supporting" SyncLink is impressive–Hyundai, Mitsubishi, Texas Instruments, Samsung, Micron, Fujitsu, MOSAID, and Apple have all paid to join the Consortium. However, this support has yet to translate into a significant increase in activity. Although, these companies are motivated by the prospect of a high-bandwidth, royalty-free, industry standard for future applications, their participation at the working group meetings has been spotty.

The desire for a royalty-free high-bandwidth solution is largely motivated by Rambus's business model. That model divides the DRAM vendors into two camps: those with a Rambus license and those without. Regardless of the ultimate success of RDRAMs as alternatives to SDRAMs and the other existing alternatives, the non-Rambus camp does not currently have a stylistically comparable offering. That places them in jeopardy of losing market share if RDRAMs do in fact become widely accepted. Furthermore, it has yet to be determined to what degree any SyncLink-like solution would be out from under Rambus's or other companies' intellectual property umbrellas.

SyncLink is still a very young technology. The architecture is still in development; actual SLDRAMs are several years away. However, SyncLink exhibits enough promise and backing that it may potentially become a significant market force in the future.

8.4.7 Strengths and Weaknesses

Strengths

- High bandwidth possible from just a few DRAMs
- Minimum memory system granularity
- Possibly a nonproprietary standard
- Simplified RAM model with sub-RAMs and banks
- Growing industry backing

Weaknesses

- Very preliminary–many unknowns remaining
- Memory expansion mechanism limited.

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 335
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

- Memory controller can't directly determine device characteristics
- Physical and electrical interfaces incompletely specified

Samsung Ex. 1097, p. 336
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

9 Chip-Level Comparison

This chapter summarizes the characteristics of the DRAM alternatives discussed in the previous chapters. The characteristics highlighted here are those evident from the examination of a single DRAM, independent of the surrounding memory system. Chapter 10 continues the comparative analysis by exploring in detail the features of memory systems built with the alternative DRAMs.

Table 9-1 compares the important characteristics of the alternative DRAMs. As one moves across the table from the **conventional** to the **evolutionary** and **revolutionary** alternatives, the amount of change from the existing generic DRAM interface increases.

9.1 New In The Second Edition

Table 9-1 has been significantly enhanced in this edition:

- Burst EDO and SyncLink alternatives have been added.
- Data on 64-Mbit devices has been added where it is already known. However, most 64-Mbit specifications are still subject to change.
- Peak fill frequencies have been added.
- The ratios of row-cache sizes to DRAM sizes have been added.

Samsung Ex. 1097, p. 337
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

- The conventional alternatives have been summarized in a single EDO column.
- The data for all alternatives has been updated to reflect new devices and new processing technologies

The text of this section has been updated to reflect the new and updated data.

9.2 Pin-Level and Physical Interfaces

The DRAM industry is moving away from the standard asynchronous interface. The inefficiencies of this old interface significantly constrain latency and maximum sustainable bandwidth. The four basic strategies for improving the interface involve changing the interface incrementally (wide DRAMs and **Extended Data-Out (EDO) DRAMs**), significantly (**Burst EDO DRAMs**, **EDRAMs**, **SDRAMs,** and **CDRAMs**), or radically (**Rambus DRAMs** and **SyncLink DRAMs**).

The incremental strategy used by EDO and EDRAM preserves the asynchronous nature of the interface but changes the semantics of $\overline{\text{RAS}}$ and $\overline{\text{CAS}}$ to eliminate the most egregious bottlenecks. These devices maintain a high degree of compatibility with generic DRAMs, facilitating memory controllers capable of controlling both generic DRAMs and the new devices.

The other devices that tinker with the **pin-level interface** make it synchronous, with a free-running clock (or two in the case of Rambus) entering each device. In the Burst EDO case, there is no clock input. Instead, $\overline{\text{CAS}}$ is given many characteristics of a clock, but only while a burst data transfer is in progress. For this reason, Burst EDO is considered pseudo-synchronous.

The pin-level interface is changed most significantly in CDRAMs, Rambus, and SyncLink. In addition to the standard separate data and multiplexed DRAM address bus, CDRAMs have a separate address bus and associated control signals. Rambus discards the multiplexed address bus altogether for a single byte-wide interface that carries addresses, data, and control information. The Rambus Channel is a high-speed bidirectional bus. In contrast, SyncLink currently relies on two high-

Samsung Ex. 1097, p. 338
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

speed buses: one from the controller to carry commands and write data, and a second to carry read data back to the controller. Rambus currently facilitates 533 Mbytes/s per channel, regardless of the number of DRAMs on the bus. If SyncLink devices become available as currently planned, they will support a 1 Gbyte/s peak bandwidth.

SOJ, TSOP II, and SVP packages for the new alternatives

DRAMs currently come in six basic package types (see Figure 5-1), although the **SOJ** is by far the most common. Table 5-1 lists the various packaging options and their sizes for the generic and wide DRAMs. The same packages and pinouts are used by the fast generic and Extended Data-Out (EDO) parts. EDRAMs are currently available only in the slightly larger SOJ format, though in 1996 they will be available in 300-mil 44-pin TSOP II packages. Because of a few extra control pins, SDRAMs are in packages that are slightly longer than those used for the generic or wide DRAMs.

The extra address bus and control pins on the CDRAM add significantly to the total pin count and consequently to the package size.

RDRAMs are available in both **Surface-Mount Vertical Packages (SVP)** that maximize the number of DRAMs on the physically short channel and **Surface-Mount Horizontal Packages (SHP)** that maintain a low system profile with or without memory expansion. The SHP package is essentially a TSOP II package with signal pins along only one side. An important consequence of adopting a bus-style interface for DRAMs is that the interface and package do not change from one generation to the next. RDRAMs of different generations can be intermixed freely because they have the same physical and electrical interfaces.

9.3 Logical Interface

The **logical interfaces** of the alternatives vary in two main areas: 1) the size and organization of the row cache and 2) the semantics of the read and write operations.

The **row cache** of generic DRAMs consists of a single line that is the size of a row in the DRAM array. The cache is physically built into the sense amplifiers. Access to the row cache varies according to whether the DRAM has a fast-page-mode,

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 339
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

			Generic Fast-Page-Mode DRAM[a]	Conventional	Evolutionary				Revolutionary	
				EDO DRAM	Burst EDO	EDRAM	SDRAM	CDRAM	RDRAM	SyncLink
Vendors	**Proponent**		Market inertia	Various	Micron	Enhanced Memory Systems	JEDEC Committee	Mitsubishi	Rambus	SyncLink Consortium
	Vendor/ Manufacturer		Many	Many	Micron initially, others later	NPNX; IBM starting 1Q1996	Most	Mitsubishi, Samsung	Toshiba, NEC, OKI, LG Semicon, Samsung, Hitachi	None committed yet
Pin-Level Interface			Asynchronous RAS/CAS	Asynchronous RAS/CAS	Asynchronous RAS; Pseudo-synchronous CAS	Enhanced asynchronous RAS/CAS	Synchronous RAS/CAS	Synchronous RAS/CAS and SRAM address	High-speed synchronous 9-bit bus	High-speed synchronous 9- &18-bit buses
Logical Interface	**Current Sizes and Organizations**	**4 Mbit**	×1, ×4	×8, ×9[b], ×16, ×18[b]		×1, ×4	×16	×16	8-Mbit ×8	None
		16 Mbit	×1, ×4	×8, ×16, ×32[c]	×4, ×8, ×16		×4, ×8, ×16	×16	×8, ×9	
		64 Mbit	×4	×8, ×16[c]						
	Future Organizations	**4 Mbit**				×8				64-Mbit soonest availability
		16 Mbit				×8, ×16[d]				
		64 Mbit		×32	×4, ×8, ×16		×4, ×8, ×16, ×32		×8	
	Row-Cache Organizations	**4 Mbit**	1 row of 4 Kbits[e]	1 row of 4 Kbits[a,e] 1 row of 8 Kbits[f,e]	Same as EDO	1 row of 2 Kbits; 4 rows of 2 Kbits in new versions	2 rows of 4 Kbits each (8 Kbits)	128 × 128 (16 Kbits) + DRAM row cache	1 row of 16 Kbits in 8-Mbit RDRAM	Vendor-specific
		16 Mbit	1 row of 8 Kbits[e]	1 row of 8 Kbits[a,e] 1 row of 16 Kbits[f,e]		8 rows of 4 Kbits each (32 Kbits)[d]	2 rows of 4 Kbits each (8 Kbits)	128 × 128 (16 Kbits) + DRAM row cache	2 rows of 16 Kbits or 18 Kbits each (32 or 36 Kbits)	
		64 Mbit	1 row of 16 Kbits[e]	1 row of 16 Kbits[a,f,e]			2 or 4 rows of 4 Kbits each (8 or 16 Kbits)		4 rows of 16 Kbits each (64 Kbits)	
	Ratio of Row-Cache Size to DRAM Size	**4 Mbit**	1:1,024[e]	1:1,024[a,e] 1:512[f,e]		1:2,048 or 1:512 in new parts	1:512	1:256	1:512 in 8-Mbit RDRAM	Vendor-specific
		16 Mbit	1:2,048[e]	1:2,048[a,e] 1:1048[f,e]	1:2,048[a,e] 1:1,048[f,e]	1:512[d]	1:2,048	1:1,024	1:512	
		64 Mbit	1:4,096[e]	1:4,096[a,f,e]	1:4,096[a,f,e]		1:4,096 or 1:8,192[g]		1:1,024	

Table 9-1 Chip-level comparison of all the alternative organizations. (*Continues next page*).

a. For ×4 organization.
b. Fast-page mode only. EDO parts are defined but not widely available.
c. ×9 and ×18 parts are defined but not widely available.
d. To be sampled in late 1996.
e. For symmetric organizations with more equal numbers of row and column address bits.
f. For ×16 organization.
g. For two-bank 64-Mbit SDRAMs.

Samsung Ex. 1097, p. 340
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

			Generic Fast-Page-Mode DRAM[a]	Conventional	Evolutionary				Revolutionary	
				EDO DRAM	Burst EDO	EDRAM	SDRAM	CDRAM	RDRAM	SyncLink[b]
Temporal Interface	Best Common Row Access Latency[c]	4 Mbit	60 ns	60 ns[a,d]		30 ns	70 ns	60 ns		
		16 Mbit	60 ns	60 ns	52 ns	30 ns[e]	60 ns	60 ns	60 ns[f,g]	
		64 Mbit	50 ns	50 ns	52 ns		50 ns[m]		TBD[h]	TBD
	Row-Cache Access Latency	4 Mbit	30 ns[i]	30 ns[i]		12 ns	35 ns	12 ns		
		16 Mbit	30 ns[i]	30 ns[i]	25 ns[i]	12 ns[e]	30 ns	10 ns	26 ns[f,j]	
		64 Mbit	25 ns[i]	25 ns[i]	25 ns[i]		28 ns		TBD	TBD
	Peak Per-Pin Bandwidth	4 Mbit	28 Mbits/s[k]	40 Mbits/s[k]		83 Mbits/s	80 Mbits/s	83 Mbits/s		
		16 Mbit	28 Mbits/s[k]	40 Mbits/s[k]	67 Mbits/s[k]	83 Mbits/s[e]	100 Mbits/s	100 Mbits/s	533 Mbits/s[f,l]	
		64 Mbit	33 Mbits/s[k]	50 Mbits/s[k]	67 Mbits/s[k]		150 Mbits/s[m]			333 Mbits/s[n]
	Peak Device Bandwidth	4 Mbit	112 Mbits/s	160 Mbits/s[a] 640 Mbits/s[d]		333 Mbits/s[a] 666 Mbits/s[o]	1.1 Gbits/s	1.3Gbit/s		
		16 Mbit	112 Mbits/s	160 Mbits/s[a] 640 Mbits/s[d]	533 Mbits/s[o] 1.1 Gbits/s[d]	666 Mbits/s[o] 1.3 Gbits/s[d,e]	800 Mbits/s[o] 1.6 Gbits/s[d]	1.6 Gbits/s	4.8 Gbits/s[f,l]	
		64 Mbit	133 Mbits/s	200 Mbits/s[a] 800 Mbits/s[d]	533 Mbits/s[o] 1.1 Gbits/s[d]		1.2 Gbits/s[o,m] 2.4 Gbits/s[d,m]			8–9 Gbits/s
	Peak Fill Frequency	4 Mbit	28 Hz	40 Hz[a] 160 Hz[d]		83 Hz[a] 167 Hz[o]	320 Hz	332 Hz	533 Hz for 8-Mbit RDRAM[l]	
		16 Mbit	7 Hz	10 Hz[a] 40 Hz[d]	33.3 Hz[o] 67 Hz[d]	42 Hz[o] 83 Hz[d,e]	50 Hz[o] 100 Hz[d]	100 Hz	267 Hz[l]	
		64 Mbit	2.1 Hz	2.5 Hz[a] 10 Hz[d]	8.3 Hz[o] 16.7 Hz[d]		18.75 Hz[o,m] 37.5 Hz[d,m]		67 Hz[l]	128 Hz

Table 9-1 (*Continued*).

a. For a ×4 part.
b. No SLDRAM product plans have been announced. This analysis is based on a hypothetical 64-Mbit SLDRAM with a 1 Gbyte/s peak bandwidth.
c. From the start of row-cache miss access to start of data return, from the DRAM's perspective.
d. For a ×16 part.
e. To be sampled in late 1996.
f. For 8-Mbit RDRAMs too.
g. Includes row access and row-cache hit access but not the initial Request Packet or data transfer times. Assumes that the RDRAM supports the Close bit and that the bank is closed
h. To Be Determined
i. The t_{AA} access time from column address.
j. Does not include Request Packet time or data transfer time.
k. Based on minimum CAS cycle time.
l. Higher frequency RDRAMs to become available in 1996.
m. Expected maximum common speed grade.
n. Averaged over both SendLink and SinkLink, assuming 500 Mbytes/s per bus.
o. For a ×8 part.

			Generic Fast-Page-Mode DRAM[a]	Conventional	Evolutionary				Revolutionary	
				EDO DRAM	Burst EDO	EDRAM	SDRAM	CDRAM	RDRAM	SyncLink
Physical Interface	**Package[b]**	**4 Mbit**	26/20 TSOP II	26/20 TSOP II[a] 44/40 TSOP II[c]	Same as EDO	28 SOJ[a] 44 TSOP II[a,d]	50 TSOP II	70/68 TSOP II		Unspecified
		16 Mbit	26/24 TSOP II	26/24 TSOP II[a] 50/44 TSOP II[c]		44 TSOP II	44 TSOP II[a,d] 50 TSOP II[c]	70/68 TSOP II	32 SVP or SHP[e]	
		64 Mbit	32 TSOP II	32 TSOP II[a] 50 TSOP II[c]			54 TSOP II			
Electrical Interface	**Supply Voltage**	**4 Mbit**	5 V or 3.3 V	5 V or 3.3 V		5 V or 3.3 V	3.3 V	3.3 V		
		16 Mbit	5 V or 3.3 V	5 V or 3.3 V	3.3 V	5 V or 3.3 V	3.3 V	3.3 V	3.3 V[e]	
		64 Mbit	3.3 V	3.3 V	3.3 V		3.3 V			3.3 V
	Signaling Interface	**4 Mbit**	TTL or LVTTL	TTL or LVTTL		TTL or LVTTL	LVTTL	LVTTL		
		16 Mbit	TTL or LVTTL	TTL or LVTTL	LVTTL	TTL or LVTTL	LVTTL, SSTL	LVTTL	RSL[e]	
		64 Mbit	LVTTL	LVTTL	LVTTL		LVTTL, SSTL			SSTL
Price	**Approximate Long-Term Per-Bit Price Premium[f,g]**	**4 Mbit**	Reference (1.00×)	1.0×		1.25×	1.05–1.20×; 1.0× by 1998	1.15×		
		16 Mbit			1.0×			1.15×	1.1× to 1.2×[h]	
		64 Mbit			1.0×					Unknown

Table 9-1 (*Continued*).

a. For a ×4 part.
b. See Table 5-1, on page 135, for a complete list of packaging options and characteristics.
c. For ×16 organization.
d. For a ×8 part.
e. For 8-Mbit RDRAMs too.
f. Subject to market dynamics.
g. All relative to the same speed and width fast-page-mode DRAM.
h. Relative to 1M×16 EDO DRAM.

static-column-mode, or nibble-mode interface. The size of the row cache typically doubles with each generation as the column addresses grow. The row cache also grows as the width increases. This means that most generic DRAMs have more sense amplifiers than logical row-cache bits. With the exception of the CDRAM, which has a separate SRAM structure, the alternatives with more row cache than the generic DRAMs do not add sense amplifiers to the design, but instead use more of the existing sense amplifiers. A few of the organizations—most notably the newer EDRAMs and the 16-Mbit RDRAMs—expose all of the existing sense amplifiers as row-cache bits.

Most of the evolutionary (EDRAM, SDRAM, and CDRAM) and revolutionary (RDRAM and SyncLink) alternatives partition the row cache into more than one block—usually by dividing the DRAM into multiple banks.[1] Increasing the number of lines reduces the size of each line and increases the memory system's ability to cache a variety of addresses simultaneously. The CDRAM takes this partitioning strategy to an extreme by dividing its 16 Kbits of row cache into 64-bit blocks. The draft SyncLink standard introduces another layer into the memory hierarchy by dividing a single DRAM into a number of autonomous sub-RAMs, each of which contains a number of banks.

The CDRAM is also unique in the extent to which the system designer is allowed or required (depending on one's point of view) to manage the use of the row-cache SRAM. The system designer can choose to implement a set-associative cache or to use the SRAM as a fast, local memory. In all other DRAMs, the row caches are direct-mapped and managed completely autonomously.

EDRAMs have one, four, or eight banks, depending on the device and generation. SDRAMs have two banks at the 4- and 16-Mbit densities, and nominally have four banks at the 64-Mbit level. The number of banks per RDRAM is one, two, and four for the 8-, 16-, and 64-Mbit devices, respectively. The 18-Mbit parts are logically identical to the 16-Mbit devices except that they implement 9-bit bytes. Because RDRAMs on a single Rambus Channel are not ganged together in parallel, the number of banks per memory system can be many times higher than in any of the conventional or evolutionary cases. The proposed and draft RamLink and SyncLink standards do not specify the number of banks per device, and no specific product plans have been announced.

1. First-generation EDRAMs have single-block row caches.

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 343
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

Read and write operations vary

The other side of the logical interface—the details of performing reads and writes—also varies significantly across the options. At one extreme, the amount of change from fast-page mode to EDO is relatively minor. At the other extreme, the RamLink protocol bears little resemblance to the fast-page-mode interface.

The EDO logical interface differs from the generic fast-page-mode DRAM's in the semantics of the $\overline{\text{CAS}}$ signal and the control of the DRAM's output buffer. Burst EDO makes a further evolutionary change to the semantics of $\overline{\text{CAS}}$ by turning $\overline{\text{CAS}}$ into a clock for the duration of the burst. Burst EDO also includes pipelining of the read access path: five pulses of $\overline{\text{CAS}}$ are needed to transfer four pieces of data. EDRAMs incorporate somewhat larger changes to the logical interface; these changes improve the efficiency of writes and facilitate row-cache operations underneath the refresh and precharge of the DRAM array. Writes are retired as quickly as row-cache reads, and the write-through to the DRAM array happens after the fact, in parallel with subsequent row-cache read accesses.

SDRAMs differ from generic and conventional DRAMs in that all address and control signals are latched on the falling edge of clock. Instead of bringing $\overline{\text{RAS}}$ low and holding it low to cause a row access, it is pulsed low for one cycle. The two banks of the 4- and 16-Mbit SDRAMs and the four banks of the 64-Mbit devices add a whole new dimension to the DRAM interface by facilitating simultaneous row and column operations in different banks. This allows the data bus to be utilized completely, even in the face of row-cache misses if the reference pattern permits. Though all the Synchronous DRAMs are similar, there are ongoing incompatibilities in the temporal specification and around the periphery of the logical specification. As a result, system designs that make use of certain features will be unable to use some vendors' SDRAMs. The number and significance of these incompatibilities has decreased significantly since SDRAMs were first introduced in 1993. The SDRAMs will continue to migrate towards a common specification as they become more accepted.

Mode registers new to DRAMs

Unlike generic or conventional RAMs, both SDRAMs and CDRAMs have mode registers that customize their behavior. In SDRAMs, the mode register controls the column access time and both the burst length and style. In CDRAMs, the mode register enables output latches or registers and also controls the column latency. The row-cache access latency is programmable to facilitate minimum access times across a wide range of

Samsung Ex. 1097, p. 344
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

memory system frequencies. RDRAMs contain 10 registers that control many aspects of their behavior and report their status. RamLink also specifies a large number of registers—most of them optional—and adopts the IEEE Standard 1212-1991 Control and Status Register (CSR) Architecture as a uniform mechanism for accessing them.

Both Rambus and RamLink define packet-based protocols. Adjacent groups of bytes on the bus or ring encode read or write commands. Commands include address, size, and type information, and, in the case of writes, the new data as well. Rambus defines bursts up to 256 bytes, while RamLink allows up to 64-byte bursts. Both protocols also define a form of back-off to deal with row-cache misses. If a read or write access is to an address not currently in a row cache, the DRAM Nacks the request and proceeds to do a row access to bring the desired data into the cache. A Rambus bus master either knows how long to wait for the row access to be complete or continually polls the DRAM. If the master polls the DRAM, the DRAM repeatedly Nacks the request until the row access is complete. This scheme is becoming optional in the 64-Mbit generation in order to help increase the device's sustainable bandwidth.

In contrast, when a RamLink DRAM tells the memory controller that the DRAM is unable to satisfy a request, it also indicates when the memory controller can ask for the data with some confidence of success.

The SyncLink logical interface is much simpler than that of RamLink.[2] It is much more like that of conventional DRAMs in that the controller is expected to know everything about the state of the DRAM and when it is going to respond to a request.

9.4 Electrical and Temporal Interfaces

Designers have long recognized that TTL signaling levels and unterminated transmission lines are not conducive to high-speed operation. The developers of high-speed circuits and systems have been torn between the advantages of compatibility with existing devices and the performance improvements offered by a change in signaling levels.

2. Based on version 0.77 of the draft SyncLink standard.

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 345
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

TTL and **LVTTL** levels are so pervasive in electronic systems that all the conventional and evolutionary alternatives continue to run at these levels. Of the proponents of these alternatives, only the SDRAM camp vocally endorses the use of a small-signal-swing **electrical interface** as a necessary but not sufficient condition for operation above 80–100 MHz in a large or expandable memory system. However, SDRAM vendors—after having rejected **GTL, CTT,** and **HSTL** as candidates for the interface of choice for high-speed memory systems—are moving to introduce **SSTL** as an industry-wide standard for frequencies above 100 MHz.

Two approaches to high-frequency operation

Each of the three revolutionary alternatives approaches high-speed operation differently. Rambus, RamLink, and SyncLink have adopted highly constrained physical and/or electrical specifications that facilitate guaranteed high-speed operations. All three use low-swing signals on controlled-impedance terminated transmission lines. The differences between the three revolutionary alternatives stem from different conceptions about the size of the system. The Rambus solution is based on a physically short bus where the allowed length is limited by the bit time and the signal propagation velocity. In contrast, the **Low-Voltage Differential Signals (LVDS)** interface adopted by RamLink uses differential unidirectional signals to ensure high-speed operation over longer distances. Like future high-speed SDRAMs, SyncLink DRAMs will use SSTL as their electrical interface. SyncLink defines two buses: a full speed, byte-wide outgoing bus for commands and for write data flowing from the controller to the DRAMs; and a half-speed, 18-bit bus for read data flowing to the memory controller.

A key aspect of the Rambus electrical interface is the constrained bidirectionality of the bus: data flows either to or from a bus master at one end of the bus. This constraint facilitates 1) singly terminated transmission lines that yield lower power consumption and 2) the unambiguous use of two clocks, one of which is synchronous with data flowing in each direction. This novel solution effectively eliminates clock skew as a system design problem.

Capacitive loading varies

A frequently overlooked aspect of electrical and physical interfaces is the capacitive loading that the DRAMs present. In non-transmission-line environments, the total loading directly impacts the rise and fall times; in transmission-line environments, the magnitude of the capacitive loading impacts both the speed of signal propagation and, in the case of nonuniform loading, the magnitude of any reflection generated at the load.

Samsung Ex. 1097, p. 346
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

Generic DRAMs typically have input capacitances of 5 pf on the address and data pins and a capacitance of 7 pf on the control pins. In memory systems or SIMMs where many narrow DRAMs are ganged in parallel, the total capacitive loading on the address and control pins—which typically are connected to all DRAMs in a bank—can become very large. For example, the memory controller for a ×64 memory system constructed of ×4 DRAMs needs to drive in excess of 100 pf. A significant fraction of this total is due to the capacitive load presented by the SIMM or DIMM sockets and board trace.

Most alternative DRAM architectures have comparable input capacitances. The wide, conventional EDO DRAMs follow this pattern. Most SDRAMs and CDRAMs offer 5 pf on the address and control pins and 7 pf on the data and clock pins. Some fast conventional devices have somewhat lower input capacitances as a means of improving operating frequency.

The EDRAM parts are unusual in that their input capacitance on the address, control, and data pins is larger than usual—7 pf, 10 pf, and 6 pf, respectively—but the input capacitance on the output enable pin ($\overline{G}$) is only 2 pf. The effort directed toward accomplishing this very low input capacitance stems from the designers' recognition that EDRAMs' high speed will be useful in interleaved memory systems only if the access time from output enable is very fast, even in comparison to its fast column access times. By minimizing the load on this signal, the two sub-ranks of two-way interleaved memory systems operating at 66 MHz can share the same data bus.

In regard to input capacitances, RDRAMs are again in a class by themselves. All high-speed pins on an RDRAM are specified to have input capacitances of less than 2.0 pf. Furthermore, the maximum variation between any pair of pins on one chip is 0.08 pf. These low values and tight tolerances are crucial elements in maintaining good signal quality, high signal propagation speed, and very low skew on the Rambus Channel. The low values are achieved mostly through careful design of the package and IC. The input capacitance is lowered by the use of current-mode output drivers because current-mode drivers can be built with smaller output transistors than comparable voltage-mode drivers.

Low latency and/or high bandwidth

In the **temporal interface**, all the alternatives except EDRAMs start with DRAM cores of essentially the same speed. Row access times of 60–70 ns are common for second-generation 4-Mbit DRAMs and mainstream 16-Mbit devices. At the

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 347
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

64-Mbit level, 50-ns row access times are expected to be common. In contrast, EDRAMs are built around an impressively fast DRAM core capable of 30-ns row accesses.[3] This fast row access time and the correspondingly fast column access time explain EDRAMs' good performance despite their asynchronous interface and TTL or LVTTL signal levels. In contrast, the designers of the CDRAM concentrated only on reducing the latency of row-cache accesses. Consequently, CDRAMs have row access times comparable to those of generic DRAMs.

The SDRAM and RDRAM designs offer high and very high bandwidth, respectively, but with only average latency.

At the chip level, EDO DRAMs have little advantage over generic DRAMs. On paper, generic DRAM column access times are quite fast. However, when $\overline{CAS}$ is cycled close to its limit, the data-valid-out time disappears to zero. This does not happen to the same extent with EDO DRAMs. Consequently, EDO DRAMs show a more significant performance improvement over generic DRAMs in real memory systems (see Chapter 10).

9.5 Peak Fill Frequencies

Table 9-1 includes the peak fill frequencies of all the devices (see Section 2.5). The fill frequencies of the DRAMs represented in Table 9-1 span two-and-a-half orders of magnitude. At the upper end, the 8-Mbit RDRAMs have a peak fill frequency of 533 Hz; at the lower extreme, the ×4 64-Mbit fast-page-mode DRAMs have a peak fill frequency of 2.1 Hz. At one extreme, there are devices capable of meeting the most demanding memory system requirements, while the DRAMs at the other end of the spectrum are not very useful at all.

As expected, the device's fill frequencies increase substantially in the migration from conventional to generic to revolutionary alternatives. Within each DRAM architecture, the peak fill frequencies increase with device widths. Most significant is the consistent decline in the fill frequencies with each new generation of DRAMs.

Section 2.5 explains how peak fill frequencies can be used to definitively rule out the effective use of some devices in some markets. This comparative analysis is found in Section 11.2.5.

3. For -12 parts.

Samsung Ex. 1097, p. 348
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

9.6 Power Consumption

As more portable devices include significant memory and computational power on board, DRAM power consumption becomes more important. A measure of dissipated power is not included in Table 9-1 because power consumption varies considerably across vendors and part numbers. Low-power devices with reduced refresh requirements as well as battery backup and self-refresh modes are available in all the conventional widths. These features will eventually make their way into the new architectures. SDRAMs and the ×16 CDRAMs have built-in low-power standby modes that enable self-refresh and dramatically reduce power consumption.

RDRAMs also offer a number of features that reduce the average power consumption in mobile applications. The most important is a power-down mode that turns off all the power-hungry Phase-Locked Loops and clock drivers. The power-down mode facilitates powering down most of the memory system while keeping a frame buffer active to do display refreshes.

Interesting, the portable systems community is also anxiously awaiting a new low-signal-swing electrical interface. However, this community's concern is less with increased performance than with reduced power consumption. The high-speed low-swing interfaces use less energy per bit transferred than LVTTL because:

- The energy used is proportional to the square of the signal swing. The high-speed electrical interfaces use no more than 1 V of signal swing instead of the 3.3 V used by TTL and LVTTL.
- High-speed interfaces facilitate higher clock frequencies. Since energy consumed per bit is equal to the power consumption multiplied by the time each bit is on the bus, higher clock frequencies translate into lower energy consumed per bit transferred.

This topic is considered in greater depth in Section 10.4.3.

© 1996 Steven Przybylski

9.7 Price and Supply

Chapter 2 shows that the price per DRAM bit has fallen incredibly over the past 20 years. However, as the abrupt 60% rise in DRAM street prices in the summer of 1993 so clearly showed, predicting future DRAM prices is a gamble. Fortunately, the relative prices of different organizations, speeds, and package types are less subject to extreme fluctuations than the absolute prices of DRAMs as a group: prices of DRAMs with different speeds or widths tend to vary together.

Specialized devices with few sources of supply and well-established demand in the form of design wins in production command higher relative prices than widely available generic devices. When making decisions about which DRAM to use, system designers place greater emphasis on comparing DRAMs' **price premiums** than on comparing their absolute prices. A DRAM's price premium is the gap between its price and the comparable generic's price, expressed as a percentage of the generic's price. As the number of suppliers of a product increases, or as demand softens due to the availability of technically feasible alternatives, the price premium declines.

Factors that set lower limits on a device's relative price are die size, test costs, and package costs. For the evolutionary and revolutionary choices, test cost is the most significant contribution to the long-term price premium. This is because the die size penalties for the added functionality are already small. As DRAM densities increase, the fraction of the new DRAMs' total area attributable to new features will decrease further. For example, the die size penalty for a Rambus interface is 15% for the 4-Mbit device, but the penalty shrinks to less than 10% at the 16-Mbit level. And as engineers begin working on designs, they discover clever ways to decrease the area overhead. For example, initial speculations estimated the die size penalty for a 16-Mbit Synchronous DRAM at 5–10%. Now that several vendors have completed SDRAM designs, actual area penalties of 0–5% are more typical. More than one manufacturer has indicated that its SDRAM devices suffer no die size penalty over its generic DRAMs of comparable width.

Testing time is an important contributor to DRAM cost

Test costs can contribute up to 50% of the DRAM vendor's manufacturing costs. If not done carefully, the additional test time needed to verify the proper operation of an SRAM block, multiple on-chip banks, and high-speed interfaces contributes significantly to a device's production costs. However, the new features

Samsung Ex. 1097, p. 350
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

can be used to actually reduce test time and cost. High-speed interfaces reduce the duration of the data I/O portions of tests, and dual-bank architectures can test one bank while waiting for the bit in the other bank to decay. The total time to test two sub-banks can be less than to test a single bank that is the size of the two sub-banks combined.

Superficially, many of the price premiums for the new DRAMs look unacceptably high. However, in many applications, the new DRAMs are replacing VRAMs. VRAMs currently command a price per bit premium of about 50%; the new DRAMs' 0–25% premiums look very attractive in comparison. The new DRAMs' price premiums also appear less imposing when compared with the nontrivial premiums for wide, fast, or low-power generic DRAMs.

Considerable care must be taken when playing the relative price per bit game. In a perfect world, all vendors would quote their projected price per bit premiums relative to a well-defined standard, such as the ×1 or ×4 part of the same density. However, while some vendors quote their price premiums this way, other quotes price premiums relative to the generic DRAMs of comparable width.

The most important determinant of the new DRAMs' future prices is the volume shipped. This in turn depends on the numbers and sizes of the initial design wins and the slope of the volume ramp that results. Chapter 11 explores which market segments are most likely to adopt the new alternatives; Section 11.3 looks specifically at the expected pricing of the new DRAM organizations.

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 351
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

Samsung Ex. 1097, p. 352
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

10 System-Level Comparison

Chapter 9 compares the alternative DRAM architectures at the chip-level, one DRAM at a time. More significant, however, are the characteristics of the overall memory system. The memory system characteristics facilitated by the DRAMs vary as widely and as unexpectedly as those of the DRAMs themselves. This chapter presents a system-level comparison of the new DRAMs. Focusing on a whole memory system highlights the criteria of most interest to system architects and product planners.

Unfortunately, no single memory system is representative of all DRAM applications. Each DRAM alternative is more or less suited to particular system types. Consequently, this chapter contains five main sections. Each section focuses on a single broad class of memory system. The first four sections highlight memory systems in four major classes of computer systems: workstations, low-end desk-top PCs, portable devices, and high-performance embedded systems. Because several of the new DRAMs are especially suited to graphics applications, the final section focuses in part on a typical frame buffer configuration.

These five types of memory systems represent five distinct points in a multidimensional design space characterized by:

- The minimum and maximum size of the memory system
- The absolute and relative importance of bandwidth and latency
- The physical and electrical constraints
- The total cost

Samsung Ex. 1097, p. 353
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

Other DRAM applications include mainframes, vector supercomputers, massively parallel computer systems, and consumer electronics. Sections 10.5.2 and 10.6.2 discuss these applications.

10.1 New In The Second Edition

There have been substantial changes to this chapter in the second edition.

- Two new alternatives were added to the analysis: Burst EDO and SyncLink. In addition, the conventional alternatives have been changed to focus solely on narrow and wide EDO devices.
- All data and analysis has been updated to reflect new devices and speed bins within each alternative.
- The widths of the memory systems in the workstation and PC system analyses have been doubled to reflect changes in common system architecture. In particular, the PC analysis is now Pentium-centric instead of 486-centric.
- The basis of comparison for the PC memory system has been changed to include 1) the time spent in transit to and from the processor and 2) a more accurate accounting of the overhead cycles spent in the chip set. The latencies and data access patterns are now those seen by the CPU and not those seen by the memory controller.
- An analysis of UMA memory systems has been added to Section 10.3.
- A discussion of the overall performance impact of changes to the PC memory system has been added to Section 10.3.
- The aggressiveness of the workstation memory system design has been increased. The operating frequencies of the SDRAM and CDRAM memory systems have been increased to reflect what is possible with detailed tuning of clocks throughout the system.
- Available data on the 64-Mbit generation has been added where it is consistently available. Analysis of the impact of the next generation on the workstation and PC memory systems has been added.

Samsung Ex. 1097, p. 354
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

- Analysis of sustainable fill frequencies has been added to Sections 10.2 and 10.3.
- Section 10.5 on high-performance embedded systems has been added.

10.2 High-Performance Memory System

The model for the high-performance memory system is a high-end graphics workstation or server. Both uniprocessor and multiprocessor configurations require a combined processor and I/O bandwidth of several hundred megabytes per second. Entry-level memory sizes of 64 to 128 Mbytes are appropriate, and maximum memory size should be in the 1-Gbyte range. In its maximum configuration, this is a multiboard computer with either a conventional backplane or a plethora of daughter cards to accommodate the full spectrum of processor, memory, and I/O expansion capabilities. Memory expansion occurs by either expanding the memory under the umbrella of the primary memory controller, adding new memory subsystems, or combining these two strategies. The main memory system is protected with ECC, and refresh and ECC control are located in each memory controller.

The processors all have large secondary or even tertiary write-back caches, and the I/O boards have buffers that group I/O transfers into blocks of at least 32 bytes. Consequently, very few word transfers occur into or out of the memory system. If the computer system emphasizes graphics or video processing, then the graphics subprocessor contains a separate frame buffer, and low-level graphics and video data do not coexist in main memory. Section 10.6.1 discusses graphics memory systems.

The most important design criteria for this memory system are cost and sustained bandwidth. The minimum configuration contains enough memory that small percentage changes in the price per bit correspond to significant changes in end-user cost and in memory upgrade price. Provided enough bandwidth exists to support all the activity, memory system performance is not a primary factor in determining system-level performance. The large secondary cache mostly decouples system-level performance from memory system characteristics: dramatic

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 355
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

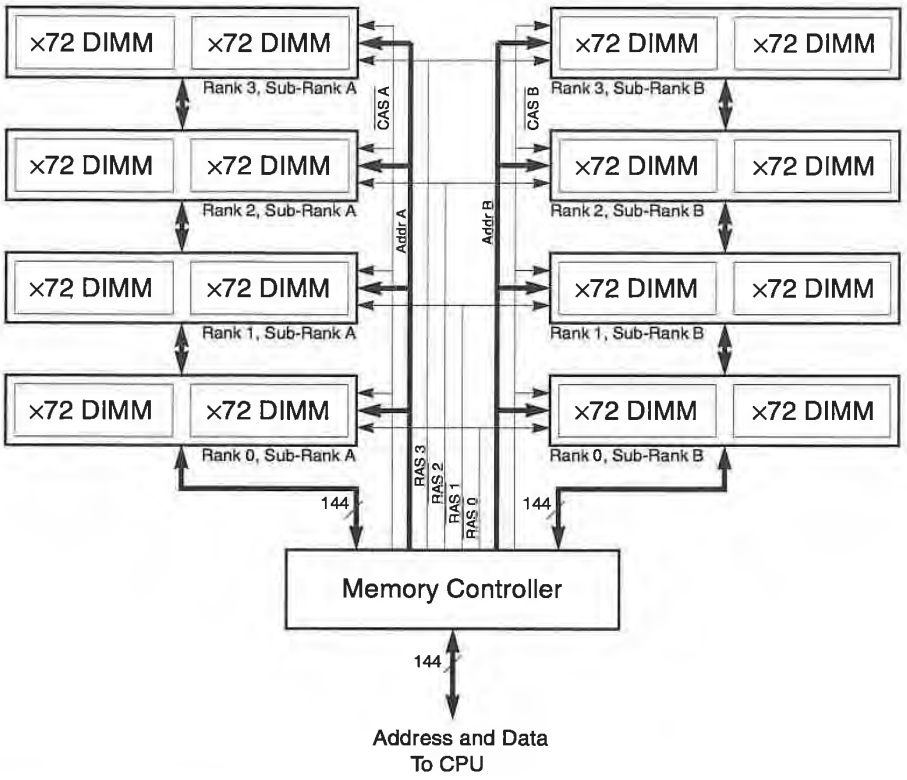


Figure 10-1 This high-performance memory system has four banks, each interleaved with two sub-banks.

changes in the speed of memory typically affect program execution time by only a few percent.

Interleaving is common in high-performance systems

Figure 10-1 shows a memory system built with generic DRAM SIMMs. Each memory controller controls several ranks; each rank is two-way interleaved. The primary datapaths, which are 144 bits wide, consist of 128 bits of data and 16 check bits. Each sub-rank is physically constructed of either four ×36 SIMMs or two ×36 DIMMs, making eight SIMMs or four DIMMs per rank.[1]

1. Section 4.2.2 defines the relationship between banks, ranks, sub-banks, and sub-ranks. When generic and conventional DRAMs are used, each rank contains exactly one bank, so the terms are effectively equivalent.

Samsung Ex. 1097, p. 356
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

The minimum memory size with a single controller and 4-Mbit ×4 DRAMs is 32 Mbytes. If each controller can manage up to four ranks populated with either 4- or 16-Mbit ×4 DRAMs, the maximum memory size per controller is 512 Mbytes. If SIMMs populated with ×1 DRAMs are used, the maximum memory size per controller becomes 2 Gbytes.

The main reason to interleave a memory system, rather than just increase the bank width, is that interleaving relieves the memory controller of the need to juggle the two halves of a wider bank. In an interleaved system, accesses to the sub-banks are staggered in time so that the data flows through the controller with minimal delay, first from one bank and then from the other. In long transfers, the sub-banks alternate, each providing data on alternate cycles. Thus, the memory controller's cycle time is half the column cycle time of the DRAMs or, alternatively, the column cycle time of the DRAMs is equal to two memory controller cycles.

Staggering the sub-bank accesses means that demands on the power supply are spread out. This reduces power supply noise.

In Figure 10-1, each set of sub-ranks has its own data bus from the memory controller. At lower frequencies, memory sub-ranks can share the same data bus. Access to the shared memory data bus can be controlled by use of the DRAM's Output Enable pin. Sharing the data bus saves a substantial number of pins on the memory controller: either zero or one memory data buses exit the memory controller instead of two. When no data buses exit the memory side of the controller, external buffers connect the memory data bus to the processor bus or backplane (as in Section 10.3). This technique is limited to frequencies that allow enough time for a DRAM to drive a heavily loaded memory data bus in less than one cycle. The time needed to drive the data bus is a function of the capacitive load presented by the data bus, the strength of the output drivers in the DRAM, and the DRAM's intrinsic output-enable access time. For example, EDRAMs can cycle faster than generic DRAMs in this configuration, partly because the EDRAM output-enable access time is 5 ns, much less than the 15 ns of generic DRAMs. With synchronous interfaces, especially at higher frequencies, interleaving two sub-ranks offers no advantage over having a single rank that is twice as wide.

Hypothetical memory systems using the alternatives

Comparable memory systems are hypothesized using each of the new alternatives. Conventional and EDRAM alternatives have asynchronous interfaces like the generic fast-page-mode

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 357
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

DRAMs:' the comparison memory systems are presumed to have the same structure as that shown in Figure 10-1. Burst EDO DRAMs, SDRAMs, and CDRAMs have a synchronous or pseudo-synchronous interface that is not helped by external interleaving. Instead, their memory systems are assumed to have noninterleaved banks that are twice as wide: 288 bits in total. These memory systems have comparable numbers of pins on their memory controllers and comparable numbers of DRAMs active at any one time. To match the increased bandwidth on the DRAM side, the bus on the processor side of the bus would also have to be double-width or double frequency to avoid the necessity of buffering data within the memory controller. The presumed operating frequency of the evolutionary memory systems is derived from a combination of factors:

- The temporal characteristics of the fastest parts that are generally available of each type
- The delays associated with signal propagation from the memory controller to the DRAMs on SIMMs or DIMMs
- The presumed temporal characteristics of a single- or multiple-chip ASIC-based memory controller
- The timing margin needed for reliable operation across variations of temperature and voltage

These factors combine to introduce at least 5 ns of slack between the minimum specified cycle time of the DRAM and the presumed system cycle time. In addition, extra cycles of access time are added as appropriate to account for the cycles spent propagating addresses and data to and from the DRAMs. However, in the case of Burst EDO DRAMs and SDRAMs, unbuffered SIMMs and DIMMs are assumed—the latency is not affected by the address and control register on buffered DIMMs available with these devices.

Assumed operating frequencies

The bottom line is that the presumed operating frequencies of the memory systems used in this section to compare the evolutionary alternatives are substantially below the maximum frequencies of the parts used. Specifically, all 4-Mbit evolutionary memory systems are assumed to operate at 66 MHz (15-ns cycles), even though theoretical minimum cycle times for the listed parts are 12 ns. At the 16-Mbit level, the memory systems are presumed to operate at 75 MHz when using 10-ns parts. There are two exceptions. The first is the 16-Mbit Burst EDO DRAM, which is limited to 66-MHz operation. The second exception is the 16-Mbit EDRAM, which presumably can run in

Samsung Ex. 1097, p. 358
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

a system at 75 MHz despite its -12 specification. This is because of its fast clock-to-data-out times and unique EDO output driver functionality.

These system operating frequencies are higher than many believe practical or reliable. However, since this section specifically focuses on the workstation marketplace, a high level of engineering sophistication can be assumed. Consequently, the listed operating frequencies rely on the precision clock alignment necessary to extract the highest possible operating frequency out of the chosen speed grades.

Deciding on a revolutionary memory system that is comparable to the generic memory system in Figure 10-1 is quite difficult. The difficulty results from the lack of consistency from one application to another in the metrics used to determine the characteristics of the memory system. The organization of the revolutionary memory systems most comparable to the system in Figure 10-1 depends on whether board area, controller complexity, overall cost, or power dissipation is to be equalized across the comparison.

Multiple parallel channels

The revolutionary memory systems used in this comparison are consistent with the cost and performance goals typical of a high-performance workstation design. A Rambus memory system consisting of eight channels operated in parallel was chosen because this number provides high performance with substantially fewer pins on the memory controller than are found on memory controllers in the conventional and evolutionary systems. Compared with having eight independent Rambus Channels, operating eight channels in parallel maximizes the peak transfer rate for each transaction, but hurts the sustained bandwidth.

A four-link[2] SyncLink system was chosen because it uses approximately the same number of signal pins as the Rambus system. As a fair comparison to Rambus, the four parallel SyncLinks are operated synchronously in parallel. A bus-based SyncLink system was chosen over a ring-topology RamLink system because RamLink is unlikely to be commercialized as a DRAM interface, while SyncLink has substantially more industrial support currently. Although no SLDRAMs are likely prior to the 64-Mbit level, 16-Mbit SLDRAMs are hypothesized for the purposes of a performance comparison.

2. A SyncLink Link is the pair of buses and associated control signals on which a collection of SLDRAMs reside.

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 359
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

		Generic Fast-Page-Mode DRAM	Conventional		Evolutionary				Revolutionary	
			EDO DRAM	Wide EDO	Burst EDO	EDRAM	SDRAM	CDRAM	RDRAM	SyncLink[a]
DRAM Organization	4 Mbit	1M×4	1M×4	256K×16		1M×4	256K×16	256K×16		
	16 Mbit	4M×4	4M×4	1M×16	2M×8	2M×8[b]	2M×8	1M×16	1M or 2M×8	16 Mbit
	64 Mbit	16M×4	16M×4	4M×16	4M×16		4M×16		8M×8	64 Mbit
DRAM Marking and Speed	4 Mbit	-6 (60 ns)	-6 (60 ns)	-6 (60 ns)		-12 (30 ns)	-12 (70ns)	-12 (60 ns)		
	16 Mbit	-6 (60 ns)	-6 (60 ns)	-6 (60 ns)	-52 (52 ns)	-12 (30 ns)	-10 (60 ns)	-10 (60 ns)	533 Mbytes/s	60 ns
	64 Mbit	-5 (50 ns)	-5 (50 ns)	-5 (50 ns)	-52 (52 ns)		-6 (50 ns)			
Physical Memory Organization		1–4 ranks of 2 sub-ranks of 36 chips each	1–4 ranks of 2 sub-ranks of 36 chips each	1–4 ranks of 2 sub-ranks of 9 ×16 chips each	1–4 ranks of 36 ×8 chips or 18×16 chips each	1–4 ranks of 72×4 chips or 18×16 chips each	1–4 ranks of 36 ×8 chips or 18×16 chips each	1–4 ranks of 9 chips each	8 channels with 1–6 RModules, each with 1–22 RDRAMs[c]	4 links with 1–12 SLDRAMs[c]
Memory Bus Width		2×144 bits	2×144 bits	2×144 bits	1×288 bits	1×288 bits	1×288 bits	1×288 bits	8×9 bits	4×(9 + 18) bits
Expansion Mechanism		DIMMs	DIMMs	DIMMs	DIMMs	SIMMs	DIMMs	SIMMs[d]	1–22 RDRAM RModules	single SLDRAM DIMMs
Granularity/ Minimum Memory Size	4 Mbit	32 Mbytes	32 Mbytes	8 Mbytes		32 Mbytes	8 Mbytes	8 Mbytes		
	16 Mbit	128Mbytes	128Mbytes	32Mbytes	64Mbytes	64 Mbytes	64 Mbytes	32 Mbytes	8 Mbytes[e,c]	8 Mbytes[c]
	64 Mbit	512 Mbytes	512 Mbytes	128 Mbytes	128 Mbytes		128 Mbytes		64 Mbytes[c]	32 Mbytes[c]
Maximum Memory Size	4 Mbit	128 Mbytes	128 Mbytes	32 Mbytes		128 Mbytes	32 Mbytes	32 Mbytes		
	16 Mbit	512 Mbytes	512 Mbytes	128 Mbytes	256 Mbytes	256 Mbytes	256 Mbytes	128 Mbytes	264 Mbytes	96 Mbytes
	64 Mbit	2 Gbytes	2 Gbytes	512 Mbytes	512 Mbytes		512 Mbytes		1.06 Gbytes	384 Mbytes
Controller Signal Pins		476	476	476	332	332	332	352	232	272

Table 10-1 Large high-performance memory systems use many DRAMs to generate high bandwidth. (*Continues next page*).

a. Based on hypothetical SyncLink devices. SyncLink DRAMs are unlikely to be available until the 64-Mbit generation.
b. To be sampled in 1996.
c. The eight Rambus Channels and four SyncLink links are ganged together. This multiplies the effective block size and granularity of the memory system by eight and four, respectively.
d. CDRAM's large package size implies a nonstandard SIMM design.
e. Using 8-Mbit RDRAMs.

Samsung Ex. 1097, p. 360
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

		Generic Fast-Page-Mode DRAM	Conventional		Evolutionary				Revolutionary	
			EDO DRAM	Wide EDO	Burst EDO	EDRAM	SDRAM	CDRAM	RDRAM	SyncLink[a]
DRAM Organization	**4 Mbit**	1M×4	1M×4	256K×16		1M×4	256K×16	256K×16		
	16 Mbit	4M×4	4M×4	1M×16	2M×8	2M×8[b]	2M×8	1M×16	1M or 2M×8	16 Mbit
	64 Mbit	16M×4	16M×4	4M×16	4M×16		4M×16		8M×8	64 Mbit
Physical Memory Organization		1–4 ranks of 2 sub-ranks of 36 chips each	1–4 ranks of 2 sub-ranks of 36 chips each	1–4 ranks of 2 sub-ranks of 9×16 chips each	1–4 ranks of 36 ×8 chips or 18×16 chips each	1–4 ranks of 72×4 chips or 18×16 chips each	1–4 ranks of 36 ×8 chips or 18×16 chips each	1–4 ranks of 9 chips each	8 channels with 1–6 RModules, each with 1–22 RDRAMs	4 links with 1–12 SLDRAMs
Number of Banks	**4 Mbit**	1–4	1–4	1–4		1–4[c]	2–8	1–4		Vendor-specific
	16 Mbit	1–4	1–4	1–4	1–4	8–32	2–8	1–4	1[d]–264	
	64 Mbit	1–4	1–4	1–4	1–4		4–16		4–528	
Row-Cache Organization	**4 Mbit**	1–4 blocks of 32 Kbytes (32–128 Kbytes)	Same as in page-mode DRAMs	1–4 blocks of 16 Kbytes[e] (16–64 Kbytes)		1–4 blocks of 16 Kbytes[f] (16–64 Kbytes)	2–8 blocks of 8 Kbytes (16–64 Kbytes)	128–512 blocks of 256 bytes (32–128 Kbytes)		Vendor-specific
	16 Mbit	1–4 blocks of 64 Kbytes[g] (64–256 Kbytes)		1–4 blocks of 32 Kbytes (32–128 Kbytes)	1–4 blocks of 32 Kbytes (32–128 Kbytes)	8–32 blocks of 16 Kbytes (128–512 Kbytes	2–8 blocks of 16 Kbytes (32–128 Kbytes)	128–512 blocks of 256 bytes (32–128 Kbytes)	1–264 blocks of 16 Kbytes[h] (16–4224 Kbytes)	
	64 Mbit	1–4 blocks of 128 Kbytes (128–512 Kbytes)		1–4 blocks of 32 Kbytes (32–128 Kbytes)	1–4 blocks of 32 Kbytes (32–128 Kbytes)		4–16 blocks of 8 Kbytes (32–128 Kbytes)		4–528 blocks of 16 Kbytes[i] (64-8448 Kbytes)	
Cache Size: Memory Size	**4 Mbit**	1:1,024	1:1,024	1:512		1:2,048[i]	1:512	1:256		Vendor-specific
	16 Mbit	1:2,048	1:2,048	1:1,024	1:2,048	1:512	1:2,048	1:1,024	1:512	
	64 Mbit	1:4,096	1:4,096	1:4,096	1:4,096		1:4,096		1:1,024	

Table 10-1 (*Continues next page*).

a. Based on hypothetical SyncLink devices. SyncLink DRAMs are unlikely to be available until the 64-Mbit generation.
b. To be sampled in 1996.
c. New 4-Mbit EDRAMs have four times as many banks.
d. Using 8-Mbit RDRAMs
e. Using DRAMs with 9-bit row addresses.
f. With ×8 and new ×4 4-Mbit EDRAMs, the row cache would be 4–16 rows of 16 Kbytes each.
g. Using DRAMs with 11-bit row addresses.
h. The eight Rambus Channels and four SyncLink links are ganged together. This multiplies the effective block size and granularity of the memory system by eight and four, respectively.
i. With ×8 and newest ×4 EDRAMs, the ratio of DRAM bits to row-cache bits is 1:512.

Samsung Ex. 1097, p. 361
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

		Generic Fast-Page-Mode DRAM[a]	Conventional		Evolutionary				Revolutionary	
			EDO DRAM[b]	Wide EDO[b]	Burst EDO[b]	EDRAM[b]	SDRAM[b]	CDRAM[b]	RDRAM[c,d]	SyncLink[c,e,f]
DRAM Organization	4 Mbit	1M×4	1M×4	256K×16		1M×4	256K×16	256K×16		
	16 Mbit	4M×4	4M×4	1M×16	2M×8	2M×8	2M×8	1M×16	2M×8	16 Mbit
Physical Memory Organization		1–4 ranks of 2 sub-ranks of 36 chips each	1–4 ranks of 2 sub-ranks of 36 chips each	1–4 ranks of 2 sub-ranks of 9×16 chips each	1–4 ranks of 36 ×8 chips or 18x16 chips each	1–4 ranks of 72×4 chips or 18x16 chips each	1–4 ranks of 36 ×8 chips or 18x16 chips each	1–4 ranks of 9 chips each	8 channels with 1–6 RModules, each with 1–22 RDRAMs	4 links with 1–8 SLDRAMs
Memory Bus Width and Speed	4 Mbit	2×144 bits @ 25 MHz	2×144 bits @ 33 MHz	2×144 bits @ 33 MHz		1×288 bits @ 66 MHz	1×288 bits @ 66 MHz	1×288 bits @ 66 MHz		
	16 Mbit				1×288 bits @ 66 MHz	1×288 bits @ 75 MHz[g]	1×288 bits @ 75 MHz	1×288 bits @ 75 MHz	8×9 bits @ 533 Mbytes/s	4×(9 + 18) bits @500 Mbytes/s
Peak Bandwidth	4 Mbit	128 bits @ 50 MHz = 800 Mbytes/s	128 bits @ 66 MHz = 1.1 Gbytes/s	128 bits @ 66 MHz = 1.1 Gbytes/s		256 bits @ 66 MHz = 2.1 Gbytes/s	256 bits @ 66 MHz = 2.1 Gbytes/s	256 bits @ 66 MHz = 2.1 Gbytes/s	8 channels × 8 bits @ 533 Mbits/s = 4.3 Gbyte/s	4 links × 2 buses/link × 500 Mbytes/s = 4 Gbyte/s
	16 Mbit	128 bits @ 50 MHz = 800 Mbytes/s	128 bits @ 66 MHz = 1.1 Mbytes/s	128 bits @ 66 MHz = 1.1 Gbytes/s	288 bits @ 66 MHz = 2.1 Gbytes/s	256 bits @ 75 MHz = 2.4 Gbytes/s	256 bits @ 75 MHz = 2.4 Gbytes/s	256 bits @ 75 MHz = 2.4 Gbytes/s		
Row-Cache-Hit Access Time (32-byte access)	4 Mbit	(1 + 2 + 1 + 1) × 20 ns = 100 ns	(1 + 3 + 1 + 1) × 15 ns = 90 ns	(1 + 3 + 1 + 1) × 15 ns = 90 ns		(1 + 1 + 0 + 1) × 15 ns = 45 ns	(1 + 3 + 0 + 1) × 15 ns = 75 ns	(1 + 2 + 0 + 1) × 15 ns = 60 ns		
	16 Mbit	(1 + 2 + 1 + 1) × 20 ns = 100 ns	(1 + 3 + 1 + 1) × 15 ns = 90 ns	(1 + 3 + 1 + 1) × 15 ns = 90 ns	(1 + 2 + 0 + 1) × 15 ns = 60 ns	(1 + 1 + 0 + 1) × 13.3 ns = 40 ns	(1 + 3 + 0 + 1) × 13.3 ns = 67 ns	(1 + 2 + 0 + 1) × 13.3 ns = 53 ns	15 ns + 11.3ns + 26 ns + 15 ns + 7.5 ns + 15 ns = 90 ns[h]	16 ns + 8 ns + 30 ns + 4 × 4 ns + 16 ns = 86 ns
Row-Cache-Hit Access Time (128-byte access)	4 Mbit	(1 + 2 + 7 + 1) × 20 ns = 220 ns	(1 + 3 + 7 + 1) × 15 ns = 180 ns	(1 + 3 + 7 + 1) × 15 ns = 180 ns		(1 + 1 + 3 + 1) × 15 ns = 90 ns	(1 + 3 + 3 + 1) × 15 ns = 120 ns	(1 + 2 + 3 + 1) × 15 ns = 105 ns		
	16 Mbit	(1 + 2 + 7 + 1) × 20 ns = 220 ns	(1 + 3 + 7 + 1) × 15 ns = 180 ns	(1 + 3 + 7 + 1) × 15 ns = 180 ns	(1 + 2 + 3 + 1) × 15 ns = 105 ns	(1 + 1 + 3 + 1) × 13.3 ns = 80 ns	(1 + 3 + 3 + 1) × 13.3 ns = 106 ns	(1 + 2 + 3 + 1) × 13.3 ns = 93 ns	15 ns + 11.3 ns + 26 ns + 15 ns + 30 ns + 15 ns = 112.6 ns	16 ns + 8 ns + 30 ns + 16 × 4 ns + 16 ns = 134 ns

Table 10-1 (*Continues next page*).

a. Total access time is 1 cycle of address decode and drive, 2–3 cycles of initial column access, 1 cycle each for subsequent 16-byte chunks, plus 1 cycle additional data latency through the controller.
b. Total access time is 1 cycle of address decode and drive, 2–3 cycles of initial column access, 1 cycle each for subsequent 32-byte chunks, plus 1 cycle additional data latency through the controller.
c. Total access time is memory controller overhead, Request Packet plus column access delay, transceiver overhead, transfer time on eight channels in parallel, plus controller overhead.
d. The eight Rambus Channels and four SyncLink links are ganged together. This multiplies the effective block size and granularity of the memory system by eight and four, respectively.
e. Based on hypothetical SyncLink devices that have 1 Gbyte/s peak bandwidth and access times comparable to those of a -60 fast-page-mode DRAM. No SyncLink DRAMs smaller than 64 Mbits are expected.
f. Total access time is memory controller overhead, Read Packet time plus column access delay, transfer time on four links in parallel, plus controller overhead.
g. Assumes use of EDO outputs in 16-Mbit EDRAMs.
h. Due to the 8-byte transfer granularity on the Rambus Channel, if a 32-byte access is not 64-byte aligned, the access time is 8 ns longer than shown.

		Generic Fast-Page-Mode DRAM[a]	Evolutionary		Evolutionary				Revolutionary	
			EDO DRAM[b]	Wide EDO[b]	Burst EDO[b]	EDRAM[b]	SDRAM[b]	CDRAM[b]	RDRAM[c,d]	SyncLink[c,e,f]
Physical Memory Organization		1–4 ranks of 2 sub-ranks of 36 chips each	1–4 ranks of 2 sub-ranks of 36 chips each	1–4 ranks of 2 sub-ranks of 9 ×16 chips each	1–4 ranks of 36 ×8 chips or 18×16 chips each	1–4 ranks of 72×4 chips or 18×16 chips each	1–4 ranks of 36 ×8 chips or 18×16 chips each	1–4 ranks of 9 chips each	8 channels with 1–6 RModules, each with 1–22 RDRAMs	4 links with 1–8 SLDRAMs
Row Access Time (32-byte access) (Array precharged)	**4 Mbit**	(1 + 4 + 1 + 1) × 20 ns = 140 ns	(1 + 6 + 1 + 1) × 15 ns = 135 ns	(1 + 6 + 1 + 1) × 15 ns = 135 ns		(1 + 3 + 0 + 1) × 15 ns = 75 ns	(1 + 5 + 0 + 1) × 15 ns = 105 ns	(1 + 6 + 0 + 1) × 15 ns = 120 ns		
	16 Mbit	(1 + 4 + 1 + 1) × 20 ns = 140 ns	(1 + 6 + 1 + 1) × 15 ns = 135 ns	(1 + 6 + 1 + 1) × 15 ns = 135 ns	(1 + 4 + 0 + 1) × 15 ns = 90 ns	(1 + 3 + 0 + 1) × 13.3 ns = 67 ns[g]	(1 + 6 + 0 + 1) × 13.3 ns = 106 ns	(1 + 6 + 0 + 1) × 13.3 ns = 106 ns	15 ns + 11.3 ns + 60 ns + 15 ns + 7.5 ns + 15 ns = 123.5 ns[h,i]	16 ns + 8 ns + 60 ns + 4 × 4 ns + 16 ns = 116 ns
Row Access Time (128-byte access) (Array precharged)	**4 Mbit**	(1 + 4 + 7 + 1) × 20 ns = 260 ns	(1 + 6 + 7 + 1) × 15 ns = 225 ns	(1 + 6 + 7 + 1) × 15 ns = 225 ns		(1 + 3 + 3 + 1) × 15 ns = 120 ns	(1 + 5 + 3 + 1) × 15 ns = 150 ns	(1 + 6 + 3 + 1) × 15 ns = 165 ns		
	16 Mbit	(1 + 4 + 7 + 1) × 20 ns = 260 ns	(1 + 6 + 7 + 1) × 15 ns = 225 ns	(1 + 6 + 7 + 1) × 15 ns = 225 ns	(1 + 4 + 3 + 1) × 15 ns = 135 ns	(1 + 3 + 3 + 1) × 13.3 ns = 106 ns[g]	(1 + 6 + 3 + 1) × 13.3 ns = 146 ns	(1 + 6 + 3 + 1) × 13.3 ns = 146 ns	15 ns + 11.3 ns + 60 ns + 15 ns + 30 ns + 15 ns = 146.5 ns[i]	16 ns + 8 ns + 60 ns + 16 × 4 ns + 16 ns = 164 ns
Row-Cache-Miss Access Time (32-byte access) (Array not precharged)	**4 Mbit**	(1 + 7 + 1 + 1) × 20 ns = 200 ns	(1 + 10 + 1 + 1) × 15 ns = 195 ns	(1 + 10 + 1 + 1) × 15 ns = 195 ns		(1 + 5 + 0 + 1) × 15 ns = 105 ns	(1 + 7 + 0 + 1) × 15 ns = 135 ns	(1 + 9 + 0 + 1) × 15 ns = 165 ns		
	16 Mbit	(1 + 7 + 1 + 1) × 20 ns = 200 ns	(1 + 10 + 1 + 1) × 15 ns = 195 ns	(1 + 10 + 1 + 1) × 15 ns = 195 ns	(1 + 6 + 0 + 1) × 15 ns = 120 ns	(1 + 5 + 0 + 1) × 13.3 ns = 93 ns[g]	(1 + 9 + 0 + 1) × 13.3 ns = 146 ns	(1 + 9 + 0 + 1) × 13.3 ns = 146 ns	15 ns + 11.3 ns + 109 ns + 15 ns + 7.5 ns + 15 ns = 173 ns[h,i,j]	16 ns + 8 ns + 100 ns + 4 × 4 ns + 16 ns = 156 ns
Row-Cache-Miss Access Time (128-byte access) (Array not precharged)	**4 Mbit**	(1 + 7 + 7 + 1) × 20 ns = 320 ns	(1 + 10 + 7 + 1) × 15 ns = 285 ns	(1 + 10 + 7 + 1) × 15 ns = 285 ns		(1 + 5 + 3 + 1) × 15 ns = 150 ns	(1 + 7 + 3 + 1) × 15 ns = 180 ns	(1 + 9 + 3 +1) × 15 ns = 210 ns		
	16 Mbit	(1 + 7 + 7 + 1) × 20 ns = 320 ns	(1 + 10 + 7 + 1) × 15 ns = 285 ns	(1 + 10 + 7 + 1) × 15 ns = 285 ns	(1 + 6 + 3 + 1) × 15 ns = 165 ns	(1 + 5 + 3 + 1) × 13.3 ns = 133 ns[g]	(1 + 9 + 3 + 1) × 13.3 ns = 186 ns	(1 + 9 + 3 + 1) × 13.3 ns = 186 ns	15 ns + 11.3 ns + 109 ns + 15 ns + 30 ns + 15 ns = 195 ns[i,j]	16 ns + 8 ns + 100 ns + 16 × 4 ns + 16 ns = 204 ns
Time to First Data at CPU[k]	**4 Mbit**	80/120/180 ns	75/120/180 ns	75/120/180 ns		45/75/105 ns	75/105/135 ns	60/120/165 ns		
	16 Mbit	80/120/180 ns	75/120/180 ns	75/120/180 ns	45/75/105 ns	40/67/93 ns	67/106/146 ns	53/106/146 ns	86/120/169 ns	78/108/148 ns

Table 10-1 (*Continues next page*).

a. Total access time is 1 cycle of address decode and drive; 4–10 cycles of precharge, row access, and column access; 1 cycle for each 16 bytes; plus 1 cycle of additional latency through the memory controller.
b. Total access time is 1 cycle of address decode and drive; 3–9 cycles of precharge, row access, and column access; 1 cycle for each 32 bytes; plus 1 cycle of additional latency through the controller.
c. Total access time is memory controller overhead, Request Packet, row access delay, transceiver overhead, transfer time on four channels in parallel, plus additional memory controller overhead.
d. The eight Rambus Channels and four SyncLink links are ganged together. This multiplies the effective block size and granularity of the memory system by eight and four, respectively.
e. Based on hypothetical SyncLink devices that have 1 Gbyte/s peak bandwidth and access times comparable to those of a -60 fast-page-mode DRAM. No SyncLink DRAMs smaller than 64 Mbits are expected.
f. Total access time is memory controller overhead, Read Packet time plus column access delay, transfer time on four links in parallel, plus controller overhead.
g. Assumes use of EDO outputs in 16-Mbit EDRAMs.
h. Due to the 8-byte transfer granularity on the Rambus Channel, if a 16-byte access is not 32-byte aligned, the access time is 8 ns longer than shown.
i. Assumes a clean row-cache replacement.
j. RDRAMs with shorter row-cache-miss latencies will be available in 2H1996.
k. Each triplet includes the time, including controller overhead, to the first 8-byte chunk at the CPU in the case of a 1) row-cache hit, 2) row access (DRAM already precharged), and 3) row-cache miss including precharge, respectively.

		Generic Fast-Page-Mode DRAM	Conventional		Evolutionary				Revolutionary	
			EDO DRAM	Wide EDO	Burst EDO	EDRAM	SDRAM	CDRAM	RDRAM[a]	SyncLink[b]
Physical Memory Organization		1–4 ranks of 2 sub-ranks of 36 chips each	1–4 ranks of 2 sub-ranks of 36 chips each	1–4 ranks of 2 sub-ranks of 9 ×16 chips each	1–4 ranks of 36 ×8 chips or 18×16 chips each	1–4 ranks of 72×4 chips or 18×16 chips each	1–4 ranks of 36 ×8 chips or 18×16 chips each	1–4 ranks of 9 chips each	8 channels with 1–6 RModules, each with 1–22 RDRAMs	4 links with 1–8 SLDRAMs
Worst Case Sustainable Bandwidth[c] (32-byte reads) (100% row misses)	**4 Mbit**	32 bytes / 8 × 20 ns = 200 Mbytes/s	32 bytes / 10 ×15 ns = 213 Mbytes/s	32 bytes / 10 ×15 ns = 213 Mbytes/s		32 bytes / 4 ×15 ns = 533 Mbytes/s	32 bytes / 8 × 15 ns = 267 Mbytes/s[d]	32 bytes / 8 × 15 ns = 267 Mbytes/s		
	16 Mbit	32 bytes / 8 × 20 ns = 200 Mbytes/s	32 bytes / 10 ×15 ns = 213 Mbytes/s	32 bytes / 10 ×15 ns = 213 Mbytes/s	32 bytes / 9×15 ns = 237 Mbytes/s[e]	32 bytes / 5×13.3 ns = 480 Mbytes/s	32 bytes / 8×13.3 ns = 301 Mbytes/s	32 bytes / 8 × 13.3 ns = 301 Mbytes/s	32 bytes / 143 ns = 223 Mbytes/s[f]	32 bytes / 100 ns + 14 × 2 ns = 250 Mbytes/s[g,h]
Best Case Sustainable Bandwidth[c] (32-byte reads) (100% row hits)	**4 Mbit**	32 bytes / 3 × 20 ns = 533 Mbytes/s	32 bytes / 4 ×15 ns = 533 Mbytes/s	32 bytes / 4 ×15 ns = 533 Mbytes/s		32 bytes / 1 × 15 ns = 2.13 Gbytes/s	32 bytes / 1 × 15 ns = 2.13 Gbytes/s	32 bytes / 1 × 15 ns = 2.13 Gbytes/s		
	16 Mbit	32 bytes / 3 × 20 ns = 533 Mbytes/s	32 bytes / 4 ×15 ns = 533 Mbytes/s	32 bytes / 4 ×15 ns = 533 Mbytes/s	32 bytes / 4 ×15 ns = 533 Mbytes/s[e]	32 bytes / 1×13.3 ns = 2.4 Gbytes/s	32 bytes / 2 × 13.3 ns = 2.4 Gbytes/s	32 bytes / 1 × 13.3 ns = 2.4 Gbytes/s	32 bytes / 60 ns = 533 Mbytes/s[f]	2×32 bytes / 17 × 2 ns = 1.88 Gbytes/s[g,h]
Worst Case Sustainable Bandwidth[c] (128-byte reads) (100% row misses)	**4 Mbit**	128 bytes / 14 × 20 ns = 457 Mbytes/s	128 bytes / 16 ×15 ns = 533 Mbytes/s	128 bytes / 16 ×15 ns = 533 Mbytes/s		128 bytes / 6×15 ns = 1.42 Gbytes/s	128 bytes / 11 × 15 ns = 776 Mbytes/s	128 bytes / 8 × 15 ns = 1.07 Gbytes/s		
	16 Mbit	128 bytes / 14 × 20 ns = 457 Mbytes/s	128 bytes / 16 ×15 ns = 533 Mbytes/s	128 bytes / 16 ×15 ns = 533 Mbytes/s	128 bytes / 9×15 ns = 948 Mbytes/s	128 bytes / 7×13.3 ns = 1.37 Gbytes/s	128 bytes / 10× 13.3 ns = 962 Mbytes/s	128 bytes / 8 × 13.3 ns = 1.2 Gbytes/s	128 bytes / 165 ns= 776 Mbytes/s[i]	128 bytes / 100 ns + 38 × 2 ns = 727 Mbytes/s[g,j]
Best Case Sustainable Bandwidth[c] (128-byte reads) (100% row hits)	**4 Mbit**	128 bytes / 9 × 20 ns = 711 Mbytes/s	128 bytes / 10 ×15 ns = 853 Mbytes/s	128 bytes / 10 ×15 ns = 853 Mbytes/s		128 bytes / 4 × 15 ns = 2.13 Gbytes/s	128 bytes / 4 × 15 ns = 2.13 Gbytes/s	128 bytes / 4 × 15 ns = 2.13 Gbytes/s		
	16 Mbit	128 bytes / 9 × 20 ns = 711 Mbytes/s	128 bytes / 10 ×15 ns = 853 Mbytes/s	128 bytes / 10 ×15 ns = 853 Mbytes/s	128 bytes / 4 ×15 ns = 2.13 Gbytes/s	128 bytes / 4 × 13.3 ns = 2.4 Gbytes/s	128 bytes / 4 × 13.3 ns = 2.4 Gbytes/s	128 bytes / 4 × 13.3 ns = 2.4 Gbytes/s	128 bytes / 82 ns = 1.6 Gbytes/s[i]	2×128 bytes / 41 × 2 ns = 3.1 Gbytes/s[g,j]

Table 10-1 (*Continued*).

a. RDRAMs with higher worst case sustainable bandwidths will be available in 2H1996. The 64-Mbit RDRAMs will be architecturally capable of substantially higher sustainable bandwidths.
b. Based on hypothetical SyncLink devices that have 1 Gbyte/s peak bandwidth and access times comparable to those of a -60 fast-page-mode DRAM. No SyncLink DRAMs smaller than 64 Mbits are expected.
c. Assuming 32- or 128-byte references from multiple independent sources handled in parallel or pipelined when possible.
d. Each bank can deliver one transaction every five cycles.
e. Burst EDO DRAMs suffer a 3-cycle penalty for bursts of length less than 4.
f. If eight Rambus Channels were used independently instead of being ganged, the sustained bandwidths for 32-byte transfers would be 1.63 Gbytes/s, and 2.3 Gbytes/s, respectively.
g. Assumes an equal number of read and write requests. With only one type of request or the other, the sustainable bandwidth is about one-half of the sustainable bandwidth with a mix of transaction types.
h. If four SyncLink links were used independently instead of being ganged, the sustained bandwidths for 32-byte transfers would be 753 Mbytes/s, and 3.1 Gbytes/s, respectively.
i. If eight Rambus Channels were used independently instead of being ganged, the sustained bandwidths for 128-byte transfers would be 3.0 Gbytes/s, and 3.5 Gbytes/s, respectively.
j. If four SyncLink links were used independently instead of being ganged, the sustained bandwidths for 128-byte transfers would be 3.5 Gbytes/s, and 3.7 Gbytes/s, respectively.

Samsung Ex. 1097, p. 364
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

Table 10-1 characterizes a comparable memory system built out of the alternative DRAMs. Ideally, this table would have been generated only after nine complete designs were completed, built, and exhaustively tested. Instead, Table 10-1 and the figures that follow it are based on a careful and pragmatic reading of the appropriate DRAM data sheets. Errors and inaccuracies—both optimistic and pessimistic—are inevitable. The same caution applies to the analyses of PC and portable memory systems in Sections 10.3 through 10.6.

Memory controller pin counts vary significantly

The memory controller has a 144-bit combined address and data port to the backplane and 288 bits of data bus to the DRAM arrays. Another one to two dozen control signals exist on the processor side. Depending on design choices and the replication of signals to the DRAMs, there are at least 44 address and control pins on the DRAM side. This makes for a total count of 476 signal pins and upwards of 260 supply pins. Consequently, under most circumstances the memory controller is partitioned into a number of ICs.

The Burst EDO and SDRAM systems have somewhat fewer signal pins than the generic interleaved system because without interleaving there is a reduced need to duplicate the DRAM address bus and control signals. The signal pin count is approximately 332. Because of the higher operating frequencies on the DRAM side, the total number of supply pins would likely remain the same despite the reduction in the number of signals on the DRAM side.

In the CDRAM system, the DRAM address and control signals are matched by SRAM address and control signals. These additional signals increase the memory controller signal pin count to around 352. The Burst EDO, CDRAM, and SDRAM pin-count estimates assume that the backplane bus is 144 bits wide. Of course, the 288-bit backplane needed to keep up with these DRAMs' peak bandwidth would add even more pins.

A controller for an eight-channel Rambus system has 248 Rambus Channel pins (including powers and grounds), 144 backplane signals, and about 70 supply pins, for a total of about 462 pins. However, 8 channels operating in parallel generate 4.2 Gbytes/s peak bandwidth during data transfers. A backplane substantially wider than 144 bits would be needed to absorb this peak bandwidth.

The quad-link SyncLink controller has approximately 128 signal pins and about 112 supply pins on the DRAM side.

© 1996 Steven Przybylski

Samsung Ex. 1097, p. 365
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

Combined with the backplane side, the total pin count comes to 272 signal pins and about 200 supply pins. This memory system also generates more peak bandwidth than a 144-bit backplane can reasonably deliver—a well-balanced system would also probably include a 256- or 512-bit backplane.

10.2.1 Memory Size and Granularity

The granularity of a memory system is related to the number and width of the sub-ranks and the width of the DRAMs:[3]

$$\text{Granularity} = \frac{\text{DRAM Size} \times \text{Number of Sub-Ranks} \times \text{Sub-Rank Width}}{\text{DRAM Width}}$$

Given a constant memory organization, the granularity (minimum memory size) increases as DRAM width decreases and as DRAM density increases.

In the case of high-performance memory systems, large minimum memory sizes are quite tolerable. Granularity becomes a problem at the 16-Mbit density and beyond only if ×1 DRAMs are used. Most machines with 0.25 Gbyte of memory require more bandwidth than memory systems constructed of these narrow devices can provide. The situation becomes even more dramatic with 64M×1 parts,[4] which yield a minimum memory size of 2 Gbytes in the memory system depicted in Figure 10-1.

In contrast, the total memory size possible with wide parts is probably too small for most high-performance machines. The acceptable minimum memory sizes are larger than the 32 Mbytes generated by 1M×16 devices. Economics and physical constraints will likely push system architects of such machines to use more than 64 Mbytes per memory controller. Consequently, workstation designers are still very comfortable with the ×4 and ×8 parts at the 16-Mbit generation. At the 64-Mbit generation, there will likely be a movement toward the ×16 devices for the minimum configurations. The flexibility that SIMMs and DIMMs provide to configure memory systems late in the manufacturing process—and without impacting any soldered components—is very valuable to systems vendors.

For Rambus and SyncLink, memory size and granularity depend on the size of the RModules or the SyncLink equiva-

3. See Section 4.2.2 for the definitions of *granularity*, *rank*, and *bank*.
4. Many industry insiders expect the ×1 organization to be dropped at the 64-Mbit level.

Samsung Ex. 1097, p. 366
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

A List of Acronyms

ASIC	Application-Specific Integrated Circuit
ASP	Actual Sales Price
ATD	Automatic Transition Detection
BiCMOS	Bipolar Complementary Metal-Oxide-Semiconductor
BIOS	Basic Input/Output System
CAL	Column Address Latch (EDRAM only)
CAS	Column Address Strobe
CBR	CAS Before RAS
CDRAM	Cached Dynamic Random Access Memory
CMOS	Complementary Metal-Oxide-Semiconductor
CPU	Central Processing Unit
CRT	Cathode Ray Tube
CS	Chip Select
CSR	Control and Status Register
CTT	Center-Tap-Terminated Interface

Samsung Ex. 1097, p. 479
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

DIMM	Dual In-Line Memory Module
DIP	Dual In-Line Package
DLL	Delay-Locked Loop
dpi	dots per inch
DRAM	Dynamic Random Access Memory
DSP	Digital Signal Processor/Processing
DTB	Data Transfer Buffer (CDRAM only)
DTV	Digital Television
ECL	Emitter-Coupled Logic
EDO	Extended Data-Out
EDRAM	Enhanced Dynamic Random Access Memory (a trademark of Enhanced Memory Systems, Inc.)
EEPROM	Electrically Erasable, Programmable, Read-Only Memory
EIA	Electronic Industries Association
EMI	Electromagnetic Interference
F	Refresh Mode (EDRAM only)
fF	Femtofarad (unit of capacitance)
FF	Fill Frequency
FIT	Failures in Time
FPGA	Field Programmable Gate Array
FPM	Fast-Page Mode
FRAM	Ferroelectric Random Access Memory
G	Output Enable
GTL	Gunning Transceiver Logic
GUI	Graphical User Interface

Samsung Ex. 1097, p. 480
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

HDD	Hard Disk Drive
HDTV	High Definition Television
HSTL	High Speed Transceiver Logic
IC	Integrated Circuit
IEEE	Institute of Electrical and Electronic Engineers
I/O	Input/Output
JEDEC	Joint Electron Devices Engineering Council
LVDS	Low-Voltage Differential Signals
LVTTL	Low-Voltage TTL
MCM	Multi-Chip Module
MDRAM	Multibank Dynamic Random Access Memory
MPP	Massively Parallel Processor
NC	No Connect
NMOS	N-Channel Metal-Oxide-Semiconductor
NRE	Non-Recurring Expenses
NTSC	National Television System Committee
OE	Output Enable
OS	Operating System
PAL	Programmable Array Logic
PC	An x86-based Personal Computer
PCB	Printed Circuit Board
PCI	Peripheral Component Interconnect
PCMCIA	Personal Computer Memory Card International Association
PD	Presence Detect

© 1995 Steven Przybylski

Samsung Ex. 1097, p. 481
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

PDA	Personal Digital Assistant
PLL	Phase-Locked Loop
PROM	Programmable Read-Only Memory
PSRAM	Pseudo-Static Random Access Memory
QA	Quality Assurance
R&D	Research and Development
RAC	Rambus ASIC Cell
RAM	Random Access Memory
RAMDAC	Random Access Memory and Digital-to-Analog Converter
RAS	Row Address Strobe
RDRAM	Rambus Dynamic Random Access Memory (a trademark of Rambus Inc.)
RE	Row Enable (EDRAM only)
RGB	Red Green Blue
RPU	Rambus Protocol Unit
RSL	Rambus Signaling Logic
ROM	Read-Only Memory
R/W	Read/Write Control
S	Chip Select (EDRAM only)
SAM	Serial Access Memory
SCI	Scalable Coherent Interface
SSI	Small-Scale Integration
SDRAM	Synchronous Dynamic Random Access Memory
SECDED	Single-Error Correction, Double-Error Detection Code

Samsung Ex. 1097, p. 482
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

SGRAM	Synchronous Graphics Random Access Memory
SHP	Surface-Mount Horizontal Package
SIMM	Single In-Line Memory Module (a trademark of Wang Laboratories)
SLDRAM	SyncLink Dynamic Random Access Memory
SOJ	Small-Outline J-Lead Package
SOMM	Small-Outline Memory Module
SRAM	Static Random Access Memory
SSTL	Series Stub Terminated Logic
SVP	Surface-Mount Vertical Package
TLVTTL	Terminated Low-Voltage TTL
TSMC	Taiwan Semiconductor Manufacturing Corporation
TSOP	Thin Small-Outline Package
TTL	Transistor-Transistor Logic
UMA	Unified Memory Architecture
VESA	Video Electronics Standards Association
VRAM	Video Random Access Memory
W	Read/Write Control
WCBR	Write and CAS Before RAS
WE	Write Enable (EDRAM and others)
WRAM	Window Random Access Memory
YUV	Representative of the luminance, chrominance color space
ZIP	Zigzag In-Line Package

© 1995 Steven Przybylski

Samsung Ex. 1097, p. 483
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

Samsung Ex. 1097, p. 484
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

B Glossary

Auto-Refresh

In auto-refresh operations, the address of the row to be refreshed is taken from a row address counter internal to the DRAM instead of from the address bus. CAS-before-RAS refresh and hidden refresh are auto-refresh operations. *See also* **CAS-before-RAS Refresh; Hidden Refresh; RAS-Only Refresh; Refresh; Self-Refresh Mode.**

Bandwidth

Bandwidth is the rate of flow of information. It is typically measured in units of either megabits or megabytes per second, depending on the application. Peak bandwidth and sustained bandwidth are analogous to maximum speed and average speed on an extended trip. *See also* **Peak Bandwidth, Sustained Bandwidth and Sustainable Bandwidth**

Bank

See **Memory Bank.**

Cache

A cache is a small high-speed, architecturally invisible buffer that holds copies of some of the locations in main memory. A cache hit occurs when a memory reference can be satisfied from the cache, and a cache miss occurs when the main memory must be accessed in order to satisfy the request. *See also* **Cache Miss Ratio; Row Cache; Set Associativity; Write-Through Cache.**

Samsung Ex. 1097, p. 485
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

Cache Coherence

In multiprocessor systems with caches, one processor can write to a memory location currently held in another processor's cache. A hardware or software mechanism is needed to ensure that any processor reading the location after the write has completed, sees the new value and not the old. In other words, the caches must be kept coherent; the mechanism used to keep them coherent is called the cache coherency protocol or algorithm.

Cache DRAM (CDRAM)

CDRAMs offer an evolutionary alternative to generic DRAMs. The concept and interface was developed by Mitsubishi, and Samsung has adopted it as well. In addition to a DRAM array, CDRAMs have a small separate SRAM that can be used as a cache for the DRAM data. Through the use of separate address bus and control signals, SRAM transactions can occur in parallel with DRAM operations, thereby facilitating robust, high-bandwidth memory systems.

Cache Miss Ratio

A cache miss is a reference to a cache that cannot be satisfied by the cache. Cache misses cause references to main memory. The cache miss ratio is the fraction of all references to the cache that miss in the cache. (Also called *cache miss rate*).

CAS

See **Column Address Strobe.**

CAS-before-RAS Refresh

If $\overline{\text{CAS}}$ is low when $\overline{\text{RAS}}$ is lowered, then most generic DRAMs will perform a CAS-before-RAS refresh cycle on the row pointed to by an internal refresh counter. The refresh counter is subsequently automatically incremented.

Center-Tap-Terminated Interface (CTT)

CTT is a new electrical interface standard that provides high speed and low signal swings with terminated lines while at the same time being backward compatible with rail-to-rail low-voltage (3.3 V) CMOS in unterminated environments.

Chip Set

Most PCs are built around chip sets. A chip set is a collection of ICs, generally designed by either a fabless chip set vendor or a semiconductor manufacturer and

Samsung Ex. 1097, p. 486
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

sold to the PC system vendor. A chip set mates with a processor, DRAMs, SRAMs (if a secondary cache is included), and the usual assortment of I/O devices to form the majority of the ICs on the motherboard.

Clock Skew

Clock skew is any difference between the actual arrival time of a clock edge and its logical, or conceptual, arrival time. Most commonly, this is the difference in time between the arrival of the clock at two different chips or circuits due to time-of-flight and loading delays on the clock traces. The term is also used to express the difference in arrival time of a clock and some data that is logically synchronous with the clock. In a synchronous system, the cycle time minus the clock skew is the amount of time available for useful work.

Column Access

A column access is a high-speed access to information in the currently accessible row. The currently accessible row is sometimes thought of as residing in the row cache. In generic DRAMs, a column access is performed by presenting a column address and asserting the column address strobe. The column access time is the time from the start of the column access (as indicated by the time the column address is present at the DRAM) to the time when the data is available to the DRAM's output pins. *See also* **Column Address; Column Address Strobe.**

Column Address

The column address is the portion of the overall address that specifies the columns in the DRAM containing the bit or bits to be accessed. In generic, conventional, and evolutionary DRAMs, the row address is presented first to the DRAM; the column address is presented next. The same address pins are used for both.

Column Address Strobe ($\overline{\text{CAS}}$)

$\overline{\text{CAS}}$ is one of the control signals on the generic DRAM interface. An initial falling edge of $\overline{\text{CAS}}$ indicates the start of a read or write column access and the presence of the column address on the address bus. For the significance of subsequent transitions in $\overline{\text{CAS}}$, *see* **Nibble Mode; Page Mode; and Static-Column Mode;**

Column Cycle Time

The column cycle time in a system is the time from the start of one column access to the start of the next.

© 1995 Steven Przybylski

Samsung Ex. 1097, p. 487
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

Commodity DRAM
A commodity part is a standard part that is widely available from several interchangeable sources. For example, the differences between the various vendors' page-mode DRAMs are relatively insignificant: DRAMs for a product can generally be purchased from any of the vendors. This interchangeability promotes several market characteristics, including aggressive price competition and lower profit margins as vendors seek to differentiate their products on the basis of price, quality, dependability of supply, reliability, and, to a lesser extent, speed and specialized features. Other examples of commodity parts are SSI TTL, standard SRAMs, ROMs, PROMs, and PALs.

Copy-Back Cache
See **Write-Back Cache.**

Current-Mode Bus
In a bus connecting more than one source of information, the bus drivers can 1) drive the individual signals that make up the bus to specific voltage levels, which determine whether a logic 0 or a logic 1 is being transmitted (a voltage-mode bus) or 2) source or sink a specific amount of current into the bus (a current-mode bus). In the latter case, the receiver of the information can detect either the current itself or the voltage change that it produces.

Daughter Board
Typically, a daughter board is a small printed circuit board (PCB) that is inserted vertically into a socket on the system's main PCB, or **motherboard**.

Delay-Locked Loop (DLL)
See **Phase-Locked Loop (PLL).**

Design Win
Semiconductor vendors compete to have their ICs designed into system products. Success in this competition is based primarily on technical merit and price. When a system vendor decides to use a particular part in a system product, it is called a design win for the semiconductor vendor. Semiconductor vendors with new products count and advertise design wins as a measure of market acceptance. Unfortunately, one high-volume design win from a major system vendor is more significant in every way than one hundred minor design wins, so the sheer number of design-ins is not necessarily meaningful.

Samsung Ex. 1097, p. 488
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

Differential Signaling
Information can be communicated from one place to another through either differential or single-ended signals. Single-ended signals transmit either the information or its complement. Differential signals transmit both. Differential signaling can be more reliable and somewhat faster, but requires twice as many pins and board traces.

Direct-Mapped Cache
See **Set Associativity.**

Dual In-Line Package (DIP)
A long, narrow package with pins every 100 mils along the two long sides. It is a through-hole, as opposed to surface-mount, package: if not socketed, the pins are soldered into holes through the printed circuit board (PCB). The DIPs used for DRAMs are mostly 30 mils wide and from 800 to 1,000 mils long, depending on the number of pins. DIPs lost ground to SOJ packages at the 4-Mbit level.

Dynamic Memory
Dynamic memories, such as DRAMs, require regular refreshing to retain data while power is applied. In contrast, static memories, such as SRAMs and Flash memories, do not require any special activity to retain data while power is applied.

Electrical Interface
A DRAM's electrical interface defines the voltages, or currents, used to communicate a logic 0 and a logic 1 to or from a DRAM. *See also* **Interface; Logical Interface; Physical Interface; Pin-Level Interface; Temporal Interface.**

Embedded DRAM
An embedded DRAM is a DRAM array and its associated control circuit on the same integrated circuit as a substantial amount of logic.

Enhanced DRAM (EDRAM)
EDRAMs are an evolutionary alternative to the existing DRAM interface; they are sold by Ramtron. EDRAMs offer a high-speed DRAM core and some small but significant changes to the logical interface that improve the handling of writes and refreshes. EDRAM is a trademark of Ramtron International Corp.

© 1995 Steven Przybylski

Samsung Ex. 1097, p. 489
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

Extended Data-Out (EDO) DRAM

An EDO DRAM is a page-mode DRAM in which the output buffer is not turned off with the rising edge of $\overline{\text{CAS}}$. This allows the data-out portion of one column access to be overlapped with the address input of the next column access in a burst. This pipelining of column accesses decreases the column cycle time and increases the bandwidth of the part.

Failures in Time (FIT)

The failure rates of components are measured in FITs: the expected number of failures in one billion operating hours. The failure rate of a system can be calculated from the individual failure rate of all its constituent components.

Fast-Page Mode (FPM)

Fast-page mode is the variant of page mode that has been in use since the mid 1980s. With fast-page mode, the DRAMs have automatic transition detectors on the address pins that detect the arrival of a new row or column address. This facilitates the beginning of a new row or column address before the corresponding transition on $\overline{\text{RAS}}$ or $\overline{\text{CAS}}$, thereby apparently decreasing access times.

Fill Frequency (FF)

The fill frequency of a DRAM or a memory system is the ratio of the bandwidth it can generate—measured in Mbytes/s—divided by its capacity—measured in Mbytes. The result—with units of Hz—indicates how frequently every stored bit can be touched. *See also* **Peak Fill Frequency, Sustained Fill Frequency and Sustainable Fill Frequency.**

Frame Buffer

A frame buffer is a portion of memory that contains a digital representation of an image being actively displayed on a physical display device, such as a CRT. A specific location in the frame buffer is reserved for each pixel on the screen. For most display devices, the frame buffer contents are read out continuously, translated into a video signal, and sent to the display. The frequency at which the frame buffer is dumped to the screen is called the frame rate. *See also* **Pixel; Video Memory; Video RAM.**

Ganged Memory Devices

In memory systems built of conventional or evolutionary DRAMs, multiple DRAMs are operated in parallel

Samsung Ex. 1097, p. 490
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

in a bit-sliced manner in order to construct a logical and physical unit that is wider than the individual devices. These devices are said to be ganged together. The same technique and terminology applies to almost all memory technology

Generic DRAM

In this report, generic DRAM refers to a narrow (×1 or ×4) page-mode DRAM with the standard interface.

Ground Bounce

Integrated circuits have two supply voltages: 1) Ground, with a potential of 0 V and 2) Vcc, typically with a potential of either 5 V or 3.3 V. When internal and external signals change from a logic 0 to a logic 1, or vice versa, current is drawn from the supply pins. Drawing greater than normal current from the supply pins causes the supply voltages to locally change from their nominal values: Ground rises above 0 V, and Vcc drops below 3.3 V or 5 V. These transient changes in the supply voltages are called ground bounce and supply droop, respectively. Too much ground bounce or supply droop can cause other circuits in the vicinity to operate more slowly or unreliably.

Granularity

The granularity of a memory system is the smallest amount of memory that can be added. It is also the smallest realizable memory system, given the memory controller and chosen DRAMs. Another way to describe the granularity of a memory system is as the shallowest depth DRAM that the memory controller can control multiplied by the memory system width.

Gunning Transceiver Logic (GTL)

GTL is a small-signal-swing, single-ended electrical standard being proposed as a replacement for TTL in digital high-speed CMOS systems.

Hidden Refresh

For most generic DRAMs, a hidden refresh cycle is a CAS-before-RAS refresh cycle on the end of a normal read operation. Since $\overline{CAS}$ is held low for the duration of the refresh cycle, the data read during the read cycle is held on the output pins while the row refresh occurs.

Interface

The DRAM interface is the specification of every aspect of a DRAM's behavior. In this report, the interface is conceptually broken down into five components: the

© 1995 Steven Przybylski

Samsung Ex. 1097, p. 491
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

physical interface, the electrical interface, the pin-level interface, the temporal interface, and the logical interface. Each component is discussed in depth in Chapter 5. *See also* **Electrical Interface; Logical Interface; Physical Interface; Pin-Level Interface; Temporal Interface.**

Interleaving

An interleaved memory system has two or more sub-banks within each bank. The sub-banks are column accessed in a staggered fashion to improve the peak bandwidth.

Known Good Die (KGD)

A known good die is a bare, unpackaged integrated circuit that has been, to some degree, tested and burned-in by a semiconductor vendor prior to sale. The known good die is then combined with other bare die in a **multi-chip module** by the purchaser.

Latency

Latency is the length of time from the start of a memory request to the arrival of the first piece of data. In some contexts, the latency of a transfer also includes the data transfer period; in these cases, the latency extends from the start of the request to the arrival of the last piece of data.

Lead-Off Time

In the PC world, the latency to, and including the transfer of, the first of four chunks of data in a burst is called the lead-off time.

Locality of Reference

All programs exhibit a locality of reference. That is, there are patterns in the addresses of referenced instructions and data that increase the likelihood of near-term use and reuse of objects that are physically close to the datum currently being accessed. Some degree of locality of reference in programs is necessary in order for caches and block transfers from memory to be effective. *See also* **Spatial Locality of Reference; Temporal Locality of Reference.**

Logical Interface

The logical interface of a DRAM specifies how the control and address signals are used to transfer data into and out of a DRAM. It includes the relative ordering of events or signal transitions and command encodings. The logical interface also includes details of the exter-

Samsung Ex. 1097, p. 492
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

nally visible internal organization of the DRAM, including the row-cache organization and a definition of associated operations. From another perspective, it includes all details of the interface as a whole that are not specifically part of the physical, electrical, pin-level and temporal interfaces. *See also* **Electrical Interface; Physical Interface; Pin-Level Interface; Temporal Interface.**

Low-Voltage Differential Signals (LVDS)

LVDS is a new electrical standard developed by the IEEE Computer Society P1596.3 working group. It is meant as a replacement for ECL levels in a low-cost, high-performance CMOS environment. Differential signaling provides for high-speed communications with very good noise immunity.

Low-Voltage TTL (LVTTL)

As far as the minimum logic 1 level and the maximum logic 0 level, LVTTL is identical to TTL: at an input, logic 0 is anything below 0.8 V, and logic 1 is anything above 2.4 V. The difference between TTL and LVTTL stems from LVTTL's maximum logic 1 input voltage, which generally cannot exceed the supply voltage by more than 0.3–1.0 V. Thus, a 3.3 V LVTTL part has a maximum logic 1 input level that is below the typical logic 1 output level of most 5 V CMOS TTL parts. *See also* **Rail-to-Rail Signaling; Transistor-Transistor Logic (TTL).**

Memory Bank

Traditionally, a bank of memory consists of one or more DRAMs that are ganged in parallel and accessed together as a wider organizational unit. In this report, a bank refers to the logical collection of bits that are accessed simultaneously. For generic and monolithic DRAMs, this definition is functionally the same as the traditional one. However, some DRAMs are internally logically partitioned into more than one bank. When such DRAMs are aggregated to form a memory system, each rank has a corresponding number of banks.

Memory Controller

The memory controller is a logic unit that oversees the memory system. It receives addresses and commands from the processor and orchestrates the retrieval and forwarding of the requested data from the DRAMs.

© 1995 Steven Przybylski

Samsung Ex. 1097, p. 493
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

Minimum Feature Size
The machines used to fabricate integrated circuits have physical limitations. Consequently, each fab line has a smallest transistor or device that can be reliably and economically manufactured. The size of the smallest feature is naturally called the minimum feature size. It is typically the drawn, actual, or effective width of a minimum-sized transistor.

Motherboard
The motherboard of a small- to medium-sized computer system is the main printed circuit board (PCB) on which the CPU and associated circuits are mounted. In PCs, memory and many I/O devices are also on the motherboard. Other, often optional, I/O and memory components are mounted on small **daughter boards** and inserted vertically into sockets on the motherboard.

Multi-Chip Module (MCM)
An MCM is a package in which more than one bare integrated circuit is mounted. Having more than one IC in a package improves the overall circuit density of the system and as well as the speed of communication between the ICs that are packaged together.

Nibble Mode
Nibble-mode DRAMs allow up to four column accesses per row access. Incrementing the two low-order address bits is handled by the DRAM. For this reason, the address bus is ignored after the first transition in $\overline{\text{CAS}}$.

Nonvolatile Memory
A nonvolatile memory retains stored data even when no power is supplied to the device. Floppy disks, hard disks, ROMs, EEPROMs, and Flash memories are all nonvolatile. *See also* **Volatile Memory.**

Octbyte
Eight sequential bytes in Rambus parlance.

Octet
An 8-bit byte in IEEE parlance.

Octlet
Eight sequential bytes in IEEE parlance.

Page Mode
With a page-mode generic DRAM, all column accesses are initiated by a falling edge in $\overline{\text{CAS}}$. The column

Samsung Ex. 1097, p. 494
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

address is always latched from the address bus on this edge. *See also* **Fast-Page Mode.**

Peak Bandwidth
The peak bandwidth of a DRAM or a memory system is the maximum rate at which data can be transferred into or out of the unit. For a DRAM, it is typically the width of the device divided by the minimum column cycle time listed on the data sheet. For a memory system, the peak bandwidth is the width of the memory system divided by the actual column cycle time of the memory system. *See also* **Bandwidth, Sustained Bandwidth and Sustainable Bandwidth.**

Peak Fill Frequency
The peak fill frequency of a DRAM or a memory system is its peak bandwidth divided by its size. *See also* **Fill Frequency, Sustained Fill Frequency and Sustainable Fill Frequency.**

Performance Gap:
The performance gap is the perceived difference between the peak bandwidth requirements of modern CPUs and the peak bandwidth provided by current DRAMs.

Phase-Locked Loop (PLL)
Phase-Locked Loops are circuits that extract the primary frequency component from an incoming signal. In digital systems, Phase-Locked Loops and their cousins, Delay-Locked Loops (DLL), are used to generate multiple clocks with fixed phase and frequency relationships. For example, a PLL can be used to generate a clock that is twice the speed of, but in phase with, an incoming clock. PLLs and DLLs are increasingly used to avoid input and output buffer delays. By looping a PLL around the clock driver, one or more internal or external clocks can be generated that are exactly in phase with a reference clock. This is crucial to reducing the amount of clock skew in a system. *See also* **Clock Skew.**

Physical Interface
The physical interface of a DRAM is the specification of its package: the size, shape, thickness, and material of the package proper; the number, type, and spacing of the pins; and the electrical characteristics of the package leads from the solder connection on the board to the bonding pad on the die. *See also* **Electrical Interface;**

© 1995 Steven Przybylski

Samsung Ex. 1097, p. 495
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

Interface; Logical Interface; Pin-Level Interface; Temporal Interface.

Pin-Level Interface

The pin-level interface of a DRAM specifies 1) the number and width of the buses entering or leaving the DRAM and 2) the number, meaning, and sense of any individual control signals. *See also* **Electrical Interface; Interface; Logical Interface; Physical Interface; Temporal Interface.**

Pixel

Pixel is derived from *PICture ELement*. Each addressable or definable spot on a video or graphics image is called a pixel. The sizes of screens and images are defined in terms of pixels. Pixels have depth measured in bits. The more bits per pixel, the more different shades of gray or different colors each pixel can be. *See also* **Frame Buffer; Video Memory.**

Presence Detect

Most SIMMs and DIMMs reserve some pins to encode information about the logical organization of the module and the speed of the DRAMs on the module. The pins can also be used to detect the presence or absence of a module in a particular socket. These pins are called the module's presence-detect pins.

Price Premium

The price premium for a DRAM alternative is the price per bit difference between the alternative and the generic DRAM of comparable width. It is commonly expressed as a percent of the generic DRAM price.

Rail-to-Rail Signaling

The electrical interface of an IC determines the voltages that are used to represent a logic 0 and a logic 1. When one of these voltages is at or near one supply voltage (i.e. Ground or Vcc), and the other logic value is near the other supply rail, then the electrical signaling interface uses rail-to-rail signaling. The alternative is a small signal swing, in which logic 0 and logic 1 are represented by voltages that are substantially closer together than Vcc and Ground. *See also* **Electrical Interface; Small Signal Swing.**

Rambus DRAM (RDRAM)

An RDRAM is a specially designed DRAM that conforms to the physical, electrical, temporal, and logical interfaces that together define the Rambus Channel.

Samsung Ex. 1097, p. 496
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

The Rambus Channel provides 500 Mbytes/s peak bandwidth over a byte-wide address, data, and control bus to between one and thirty-two RDRAMs or modules. Rambus and RDRAM are trademarks of Rambus Inc.

RamLink

RamLink is a logical and temporal description of a high-speed DRAM interface developed by the IEEE Computer Society P1596.4 Working Group.

Rank

A rank of DRAMs, SIMMs, or DIMMs is the physical collection of devices that have to be present or not present as a group for the memory system to operate properly. A rank of memory typically consists of one or more DRAMs that are ganged in parallel and accessed together as a wider organizational unit. A rank may include one or more logical banks, depending on the internal logical structure of the constituent DRAMs or modules. The granularity of a memory system is the smallest supported capacity of one rank of memory.

RAS

See **Row Address Strobe.**

RAS-Only Refresh

For most generic DRAMs, a specified row can be refreshed by lowering $\overline{\text{RAS}}$ and subsequently raising it without toggling $\overline{\text{CAS}}$. The row address specifies the row to be refreshed.

Read Buffer

A read buffer is a buffer or latch between two levels in the cache or memory hierarchy of a system. The buffer temporarily holds a small amount of data fetched from one level of the hierarchy until such time as the other memory is ready to receive the data. Read buffers are frequently needed when the data path widths are different in the two memories.

Redundancy

DRAMs contain extra rows and columns of storage cells that are used to selectively replace flawed rows and columns. During initial wafer probe, each die is tested to see which rows and columns are good. The wafer is then placed on a laser fuse programmer, and fuses are blown to cause redundant rows and columns to replace the defective ones.

© 1995 Steven Przybylski

Samsung Ex. 1097, p. 497
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

Reference, memory

Processors make requests of the memory hierarchy to read (fetch) or write data. These requests are called memory references. Sometimes the caches are able to satisfy the requests, and sometimes main memory becomes involved. Load and store instructions generate read and write references, respectively. Instruction execution also implicitly generates read references to bring the instructions to be executed into the processor.

Refresh

DRAMs naturally lose their data over time if they are left idle for long periods of time (i.e. hundreds of milliseconds). To prevent unwanted data loss, all the bits in the DRAM must be periodically refreshed. Typically, a refresh operation must be performed on each row of the DRAM within a 16-ms interval to ensure that no data is lost. *See also* **Auto-Refresh; CAS-before-RAS Refresh; Hidden Refresh; RAS-Only Refresh; Self-Refresh Mode.**

Ringing

Ideally, when a signal transitions from a logic 0 to a logic 1, or vice versa, its voltage changes instantaneously from the voltage representing one logic level to the voltage representing the other. In reality, a period of transition occurs while the voltage moves from one value to the other. When the transitions are rapid, the signal traces long, and the traces are improperly terminated, the voltage of the signal can overshoot the target and then bounce around, settling only slowly to the final voltage level. This is called ringing. In extreme cases, ringing can be physically harmful to ICs. See Figure 5-2 on page 138 for an example.

RingLink

RingLink is the primary physical layer specified in the **RamLink** IEEE Standard. RingLink defines a ring of up to 62 DRAMs or other slaves with 500 Mbyte/s, point-to-point connections.

Row

A DRAM is conceptually and physically organized as a matrix of rows and columns. The rows are accessed as a unit by a row access, after which the individual bits in the row are accessible with column accesses. *See also* **Column Access; Row Access.**

Samsung Ex. 1097, p. 498
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

Row Access

A row access makes the bits of one of the many rows in a DRAM accessible via subsequent column accesses. Conceptually, one row of the DRAM is brought into the row cache from which it can be directly accessed. In generic DRAMs, a row access is performed by presenting a row address and asserting the row address strobe. *See also* **Column Access; Row; Row Access; Row Cache.**

Row Address

The row address is the portion of the overall address specifying the row in the DRAM that contains the bit or bits to be accessed. In generic, conventional, and evolutionary DRAMs, first the row address and then the column address are presented to the DRAM. The same address pins are used for both.

Row Address Strobe ($\overline{\text{RAS}}$)

$\overline{\text{RAS}}$ is one of the control signals on the generic DRAM interface. A falling edge of $\overline{\text{RAS}}$ indicates the start of a read, write, or refresh access as well as the presence of the row address on the address bus.

Row Cache

In DRAM interfaces, once a row of bits is read into the sense amps, some or all of those bits can be accessed more quickly than other bits that are only resident in the DRAM core. Consequently, the sense amplifiers can be viewed as a small cache of the DRAM proper. The typical cache consists of one block that is loaded in parallel on a row access and potentially accessed many times before the next row access. (Also called *page; row buffer,* or *Last-Row-Read cache*).

Self-Refresh Mode

When in self-refresh mode, a DRAM refreshes itself completely autonomously, without any outside intervention. *See also* **Auto-Refresh; Refresh.**

Set Associativity

When a piece of data is loaded into a cache, it can be placed in one of several places. The number of distinct locations in the cache in which piece of data can be placed is the set associativity of the cache. If the set associativity is one—that is, if there is exactly one location in the cache which can hold a copy of each location in main memory—then the cache is said to be direct mapped. If each location in main memory can be copied into any location in the cache, the cache is said to be

© 1995 Steven Przybylski

Samsung Ex. 1097, p. 499
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

fully associative. For small caches with small numbers of blocks, increasing the set associativity from one to two or four can significantly reduce the number of cache misses and the cache miss ratio. *See also* **Cache; Cache Miss ratio.**

Simultaneous Switching Noise (SSN)

When an output driver changes state, it momentarily drives or sinks substantially more current than it does in a quiescent state. Simultaneous switching noise is the ground bounce, supply droop, and other power supply noise that comes when many drivers simultaneously change states from a logic 0 to a logic 1, or *vice versa*, thus placing extraordinary demands on the supply voltages. The effects are felt both on- and off-chip. *See also* **Ground Bounce.**

Single-Ended Signaling

See **Differential Signaling**.

Single In-Line Memory Module (SIMM)

A small elongated printed circuit board with a DRAM mounted on it. SIMMs are available in a variety of standard and special configurations from the memory vendors. They are very common in PCs because of their compact size and easy end-user installabilty. SIMM is a trademark of Wang Laboratories.

Small-Outline J-Lead Package (SOJ)

The SOJ is a long, thin surface-mount package with pins every 50 mils (1.27 mm) along the two long sides. The leads are bend up and under the package in the shape of the letter J.

Small Signal Swing

The electrical interface of an IC determines the voltages that are used to represent a logic 0 and a logic 1. When the two voltages are less than about 1.5 V apart, then the electrical signaling interface uses small signal swings. *See also* **Electrical Interface; Rail-to-Rail Signaling.**

Soft Error

A soft error is a failure to retain information. A hard error is due to a physical flaw or damage in the DRAM. A soft error, which occurs without damaging the DRAM, can be caused by alpha particles, operation of the DRAM beyond rated specifications, or just plain chance.

Samsung Ex. 1097, p. 500
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

Spatial Locality of Reference

Spatial locality of reference is a property of programs: data and instructions that are physically close to data and instructions currently in use are likely to be used in the near future. Physically close pieces of data have similar addresses. The spatial locality of reference of programs promotes the use of larger blocks and transfers in caches and memory subsystems. *See also* **Locality of Reference; Temporal Locality of Reference.**

Speed Grade

After good DRAMs are packaged and burned in, they are tested a final time. One purpose of this final test is to determine the speed of each DRAM. Based on its speed, each device is put into a different speed bin, or speed grade. For example, 4-Mbit devices typically come in 60-ns, 70-ns, and 80-ns speed grades. Speed-graded parts are also referred to as *binned* parts.

Static-Column Mode

A static-column-mode DRAM does not require that $\overline{\text{CAS}}$ be toggled to initiate subsequent column accesses following a row access and a first column access. Transitions in the address bus are enough to indicate the start of a new read or write column access to the DRAM.

Static Memory

Static memories, such as **SRAMs** and Flash memories, do not require any special activity to retain data while power is applied.

Static RAM (SRAM)

A Static RAM uses four or six transistors to store one bit of information. In contrast, a DRAM uses one transistor and a capacitor. SRAM storage cells are correspondingly larger than DRAM cells, so fewer cells can fit on a die. Consequently, SRAMs cost more per bit. However, SRAMs are faster than DRAMs and need not be refreshed to retain data.

Storage Capacitor

Each memory cell in a DRAM contains an access transistor and a storage capacitor. The value of the bit stored in the cell depends on the voltage stored on the capacitor.

© 1995 Steven Przybylski

Samsung Ex. 1097, p. 501
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

Sub-Bank

Two or more sub-banks are interleaved to form a memory bank. *See also* **Interleaving; Memory Bank.**

Sub-Block Ordering

Sub-block ordering is a well-defined order for returning words or bytes in a multiword burst transfer. Linear ordering starts at the addressed word, proceeds to the end of the addressed block, and then continues from the beginning of the addressed block. In contrast, sub-block ordering successively fills larger and larger sub-blocks, including the addressed byte or word, until the entire block is filled. The important features of sub-block ordering are: 1) the first half of the transfer of n bytes is identical to the complete transfer of n/2 bytes and 2) the sub-block ordering is unambiguously defined for all bus widths. Addresses for sub-block ordering are the exclusive OR of the addressed byte and binary counter counting from 0 to the block size. Because it was first widely used on the Intel 486 processor, sub-block ordering is also called *Intel mode* or *Intel ordering*.

Surface-Mount Vertical Package (SVP)

This relatively new package type has pins along one long edge so that the ICs are mounted vertically on the board. Unlike ZIP packages, this surface-mount package is compatible with modern board fabrication techniques. *See also* **Zigzag In-Line Package (ZIP).**

Sustainable Bandwidth

The sustainable bandwidth of a memory system is the expected bandwidth that can be extracted from the memory system assuming an aggressive memory controller and maximally stressful CPU and application. It is a function of the DRAMs used, the mix or row to column accesses, the distribution of transfer sizes. The sustainable bandwidth is a better metric for comparing DRAMs than peak bandwidth, but easier to calculate and less dependent on non-DRAM parameters than the sustained bandwidth, The sustainable bandwidth should be reported as a quadruple for the **best, good, bad and worst cases**, corresponding to 100% column accesses spread across multiple banks, a realistic but favorable mix of row and column accesses, a realistic by unfavorable mix of row and column accesses, and 100% row accesses to a single bank. *See also* **Bandwidth, Peak Bandwidth and Sustained Bandwidth.**

Samsung Ex. 1097, p. 502
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

Sustainable Fill Frequency

The sustainable fill frequency of a memory system is the ratio of its sustainable bandwidth to its capacity.*See also* **Sustainable Bandwidth, Fill Frequency, Peak Fill Frequency and Sustained Fill Frequency**

Sustained Bandwidth

The sustained bandwidth of a memory system is the bandwidth under a specific set of conditions. Its is a function of the processor, the total system workload, the system architecture, the memory controller, and the phase of the moon. *See also* **Bandwidth, Peak Bandwidth and Sustainable Bandwidth.**

Sustained Fill Frequency

The sustained fill frequency of a memory system is its sustained bandwidth divided by its size. *See also* **Fill Frequency, Peak Fill Frequency and Sustainable Fill Frequency.**

Synchronous DRAM (SDRAM)

SDRAM refers to a group of JEDEC-compliant and noncompliant DRAMs with evolutionary, synchronous interfaces instead of the traditional asynchronous interface. Two-bank versions provide internal interleaving and the possibility of 100% data bus utilization even for random accesses. *See also* **Interleaving.**

SyncLink

Originally, SyncLink was just an alternate physical topology defined in the annex of the RamLink specification. Instead of being ring-based, SyncLink consisted of two high-speed unidirectional buses for communication from the master and from the slaves, respectively. *See also* **RamLink.** Now, however, SyncLink is a draft IEEE standard on its own, with no reliance on, and decreasing compatibility with, RamLink.

Temporal Interface

The temporal interface of a DRAM specifies details relating to the timing of the device. It includes access times, data hold times for outputs, and setup and hold times for inputs. The temporal interface, which specifies timing between events or signal transitions, is closely related to the logical interface, which specifies the relative order of events or signal transitions. *See also* **Electrical Interface; Interface; Logical Interface; Physical Interface; Pin-Level Interface.**

© 1995 Steven Przybylski

Samsung Ex. 1097, p. 503
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

Temporal Locality of Reference

Temporal locality of reference is the property of programs that dramatically increases the likelihood of near-term reuse of the instruction or data currently being used. Temporal locality of reference is what makes caches of all sorts, including DRAM row buffers, work well. Without temporal locality, there would be no advantage to storing data in high-speed buffers on the chance of reuse sometime in the near future. *See also* **Cache; Locality of Reference; Spatial Locality of Reference**.

Temporal Overhead

Temporal overhead is time spent not doing useful work. For example, the time needed to get a signal from one part of a circuit board to another is sometimes considered overhead. Because of temporal overhead, access and cycle times in a system are greater than the minimum time specified in DRAM data sheets.

Thin Small-Outline Package (TSOP)

The TSOP package is a surface-mount package with gull-bend leads that make contact with the board next to, rather than underneath, the package. This allows for a lower profile than is possible with an SOJ package. Pins are on a 50-mil (1.27 mm) pitch in TSOP I packages and on a 31-mil (0.8 mm) pitch in the larger TSOP II parts.

Transistor-Transistor Logic (TTL)

Transistor-Transistor Logic (TTL) is the most widely used electrical interface standard. Developed in the late 1960s, its large voltage swings and unterminated signal traces make it appropriate for operating frequencies up to 75–100 MHz. *See also* **Low-Voltage TTL.**

UMA

See **Unified Memory Architecture.**

Unified Memory Architecture (UMA)

In a Unified Memory Architecture computer system, the main memory and any graphics or multimedia memory are physically stored in one memory system, regardless of whether the two are logically separate memory structures.

Samsung Ex. 1097, p. 504
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

Valid Data-Out Time

The valid data-out time of a particular DRAM or IC is the amount of time that data is valid on its output pins. This is the time available for the data to travel to its destination and be latched. The valid data-out time is a function of the frequency at which the part is operated.

Video Memory

Video memory is a portion of a computer's memory that is reserved for use by the video or graphics subsystems. Frequently, it is logically and physically separate from the main memory of the system. Often it is used exclusively to store a frame buffer. *See also* **Frame Buffer; Video RAM.**

Video RAM (VRAM)

A VRAM is a two-port DRAM tailored to the needs of graphics controllers. In addition to the random read-write port of a generic DRAM, a VRAM contains a serial port designed for dumping a frame buffer to a screen. *See also* **Frame Buffer; Video Memory.**

Volatile Memory

A memory is volatile when it loses stored information when power is removed. A memory that retains its data even when no power is applied is nonvolatile. SRAMs and DRAMs are volatile memories. *See also* **Nonvolatile Memory.**

Voltage-Mode Bus

See **Current-Mode Bus.**

Wafer Probe

All of the dice on a wafer are preliminarily tested prior to the wafers being sawed into individual dice. The wafer probe test is used to eliminate grossly bad dice before assembly.

Wait State

When a cache miss occurs, the processor generally has to wait until the needed data is fetched from memory. Wait states are the idle cycles during which the processor is doing little or no useful work.

Write-Back Cache

In a write-back cache, processor writes are written to the cache but not to main memory. Newly written data must be marked as *dirty* in the cache, so that when the data is being replaced from the cache, the new value

© 1995 Steven Przybylski

Samsung Ex. 1097, p. 505
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

can be propagated to main memory. Also called a *copy-back cache. See also* **Write-Through Cache.**

Write Buffer
A write buffer is a small buffer or group of buffers that is logically between two levels in the cache or memory hierarchy. Its task is to hold write data until the larger cache or memory is ready to receive it. Write buffers allow the smaller cache or CPU to proceed with other work instead of waiting on the slower memory.

Write Per Bit
Many DRAMs targeted for graphics applications offer a write per bit feature. If $\overline{WE}$ is low when $\overline{RAS}$ is lowered, then a write mask is read in on the data lines. Any column write operations performed while $\overline{RAS}$ is low affect only those bits specified in the write mask. This is useful if multiple pixels are represented across the width of the DRAM.

Write-Through Cache
In a write-through cache, all processor-generated writes, in addition to being written into the cache, flow though to main memory. Thus main memory is always kept current with the latest values generated by the processor. *See also* **Write-Back Cache.**

Zigzag In-Line Package (ZIP)
ZIP packages are vertical packages with through-hole pins along one edge. Pins are every 50 mils (1.27 mm), with alternating pins offset 100 mils (2.54 mm) in a zigzag pattern.

Samsung Ex. 1097, p. 506
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

C List of Vendors

Fujitsu Microelectronics, Inc.
Integrated Circuits Division
3545 North First Street
San Jose, CA 95134-1804
408.922.9000

LG Semicon, Inc.
3003 North First Street
San Jose, CA 95134
408.432.1331

Hitachi America, Ltd.
Hitachi Plaza
2000 Sierra Point Parkway
Brisbane, CA 94005-1819
415.589.8300

Hyundai Electronics America
Semiconductor Division
166 Baypointe Parkway
San Jose, CA 95134
408.473.9200

IBM Technology Products
P.O. Box 7932
Mount Prospect, IL 60056-7932
800.426.0181

Matsushita Electric Corp. of America
1600 McCandless Drive
Milpitas, CA 95035
408.946.4311

Micron Semiconductor, Inc.
2805 East Columbia Road
Boise, ID 83706-9698
208.368.3900

Mitsubishi Electric Corp.
Electronic Devices Group
Mitsubishi Electronics America, Inc.
1050 East Arques Avenue
Sunnyvale, CA 94086
408.730.5900

Mosel-Vitelic Corp.
3910 North First Street
San Jose, CA 95134-1501
408.433.6000

Motorola Semiconductor
5005 East McDowell Road
Phoenix, AZ 85008
602.244.6900

Samsung Ex. 1097, p. 508
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

NEC Electronics, Inc.
475 Ellis Street
Mountain View, CA 94039
415.960.6000

NPNX
1580 Yamamoto
Takayama-shi
Shiga Prefecture
Japan 294
81.470.23.3121

OKI Semiconductor, Inc.
785 North Mary Street
Sunnyvale, CA 94086
408.720.1900

Enhanced Memory Systems, Inc.
1850 Ramtron Drive
Colorado Springs, CO 80921
719.481.7000

Samsung Semiconductor, Inc.
3655 North First Street
San Jose, CA 95134-1713
408.954.7000

Siemens Components, Inc.
IC Division
10950 Tantau Avenue
Cupertino, CA 95014-0716
408.777.4500

© 1995 Steven Przybylski

Samsung Ex. 1097, p. 509
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

Texas Instruments, Inc.
Semiconductor Group
P.O. Box 1443
Houston, TX 77251-1443
713.274.2000

Toshiba America Electronics Components, Inc.
9775 Toledo Way
Irvine, CA 92718
714.455.2000

Samsung Ex. 1097, p. 510
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

E.8 A 16-Mbit and 18-Mbit Rambus DRAM Data Sheet

Reprinted with permission of Rambus Inc.

Samsung Ex. 1097, p. 809
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

Samsung Ex. 1097, p. 810
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

16-Megabit RDRAM (2M x 8) and 18-Megabit RDRAM (2M x 9)

Description

The 16-Megabit Rambus™ DRAM (RDRAM™) is an extremely high-speed CMOS DRAM organized as 2M words by 8 or 9 bits and capable of bursting up to 256 bytes of data at less than 2 nanoseconds per byte. The use of Rambus Signaling Logic (RSL) technology makes transfer rates greater than 500 MHz achievable while using conventional system and board design methodologies. Lower effecttive latency is attained by operating the dual 2KByte sense amplifiers as high speed caches, or by using Random Access mode to facilitate large block transters .

RDRAMs are general purpose high-performance memory devices suitable for use in a broad range of applications including PC and consumer main memory, graphics, video, and any other application where high-performance at low cost are required.

For detailed description of the Rambus RDRAM and Rambus Signaling Logic (RSL) technology, see the *RDRAM Design Guide*.

Features

- Rambus Interface:
  - Over 500 MB/sec peak transfer rate per RDRAM
  - Rambus Signaling Logic (RSL) interface
  - Synchronous protocol for fast block-oriented transfers
  - Direct connection to Rambus ASICs, MPUs, and Peripherals
  - Only 15 active signals that require just 31 pins total on the controller interface
  - 3.3 volt operation
  - Additional/multiple Rambus Channels provide an additional 500 MB/second bandwidth each
- Dual 2KByte sense amplifiers may be operated as caches for low latency access
- Random Access mode enables any burst order at full bandwidth
- Features for graphics include random-access mode, write-per-bit and mask-per-bit operations
- Control and refresh logic entirely self-contained
- On-chip registers for flexible addressing and timing
- Available in horizontal and vertical surface mount plastic packages
- Higher performance than VRAM at DRAM cost

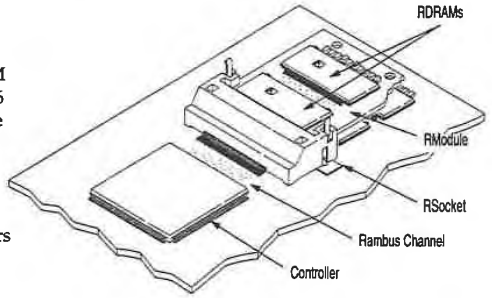


Figure 1: Rambus Subsystem

System Benefits

- Superior performance at a lower system cost than alternative multi-DRAM chip solutions
- Fully engineered solution includes clock chips, memory expansion sockets, simple layout
- For graphics subsystems addressing display resolutions of 1024x768x8 or above, provides lowest cost frame buffer, greater performance, fewest controller pins, ease of memory expansion
- For Pentium® processor class main memory, provides faster memory subsystem, fewer components, and 2MB granularity
- Sufficient performance for unified memory system architectures in cost sensitive consumer products

Part Numbers

The 16-Megabit and 18-Megabit RDRAMs are available in both horizontal surface mount (SHP) and vertical surface mount (SVP) packages, and with a 500 MHz clock rate and a 533 MHz clock rate. The part numbers for the various options are shown in Table 1.

Table 1: Part Numbers by Option

Options	500 MHz	533 MHz
16-Megabit SHP		
16-Megabit SVP		
18-Megabit SHP		
16-Megabit SVP		

Samsung Ex. 1097, p. 811
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

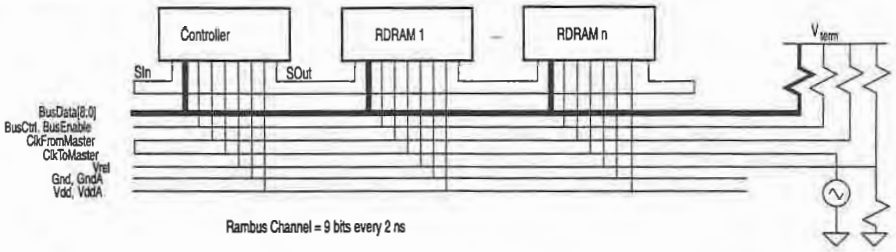


Figure 2: Controller and RDRAMs Connect to Terminated Transmission Lines

Rambus System Overview

A typical Rambus memory system has three main elements: the Rambus Controller, the Channel, and the RDRAMs. The logical representation of this is shown in Figure 2.

The Rambus Channel is a high-speed, byte-wide, synchronous bus used to connect Rambus devices together. The Channel carries all address, data, and control information to and from devices. Transfer of data on the Rambus Channel is managed through the use of a high level block-oriented protocol.

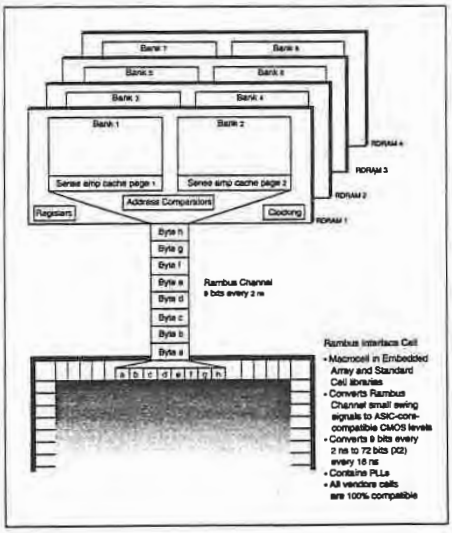


Figure 3: Data Transfer on the Rambus Channel

The Rambus Channel has thirteen high-speed Rambus Signaling Logic (RSL) I/O signals that are used to transfer information at 2 nanosecond intervals. These signals use low voltage swings (logic 0 is about 2.6V and logic 1 is about 1.6V when Vref is 2.1V) to achieve high bus speeds. Two TTL level signals are used for initialization and powerdown operation. Fourteen signals supply power and DC voltage references to the RDRAM, and the remaining pins are No Connects (reserved for future expansion).

The Rambus Interface is implemented on both master and slave devices. Rambus masters (ASIC devices, memory controllers, graphics engines, peripheral chips, or microprocessors) are the only devices capable of generating transaction requests. RDRAMs are slave devices and can only respond to requests from master devices.

Figure 1 shows a typical physical implementation of a Rambus system. It includes a controller ASIC that acts as the Channel master and a base set of RDRAMs soldered directly to the board. An RSocket™ is included on the Channel for memory upgrade using RModule™ expansion cards.

RDRAM Packages and Pinouts

RDRAMs are available in either an EIAJ standard (#ED-7424) 32-pin vertical surface mount package (SVP) or an EIAJ standard 32-pin TSOP-style horizontal surface mount package (SHP).

Each package has 32 signal pins and four mechanical pins that provide support for the device. The mechanical pins are located on the same side as the signal leads in the SVP, and on the opposite side from the signal leads in the SHP. The pins are arranged so that a printed circuit board can be designed to support either package.

Samsung Ex. 1097, p. 812
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

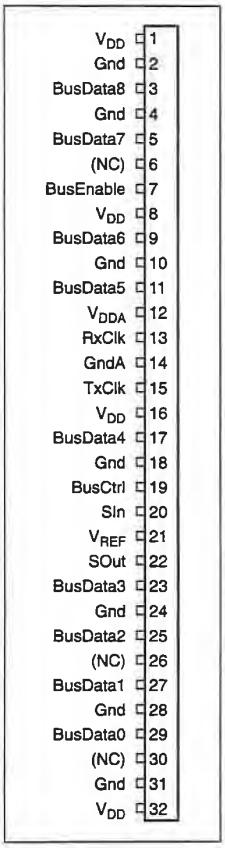


Figure 4: SHP and SVP Pin Numbering

Table 2: Pin Descriptions

Signal	I/O	Description
BusData[8-0]	I/O	Signal lines for request, write data, and read data packets. The request packet contains the address, operation codes, and the count of the bytes to be transferred. These are low-swing, active-low signals referenced to Vref.
RxClk	I	Receive clock. Incoming request and write data packets are aligned to this clock. This is a low-swing, active-low signal referenced to Vref.
TxClk	I	Transmit clock. Outgoing acknowledge and read data packets are aligned with this clock. This is a low-swing, active-low signal referenced to Vref.
Vref	I	Logic threshold reference voltage for low swing signals.
BusCtrl	I/O	Control signal to frame packets, transmit part of the operation code, to acknowledge requests, and to interrupt (terminate) pending transactions. This is a low-swing, active-low signal referenced to Vref.
BusEnable	I	Control signal to manage the operating modes of the RDRAMs and to transfer column addresses for random–access (non-sequential) transactions. This is a low-swing, active-low signal referenced to Vref.
Vdd, VddA		+3.3V power supply. VddA is a separate analog supply for clock recovery in the RDRAM.
Gnd, GndA		Circuit ground. GndA is a separate analog ground for clock generation in the RDRAM.
SIn	I	Initialization daisy chain input. TTL levels. Active high.
SOut	O	Initialization daisy chain output. TTL levels. Active high.

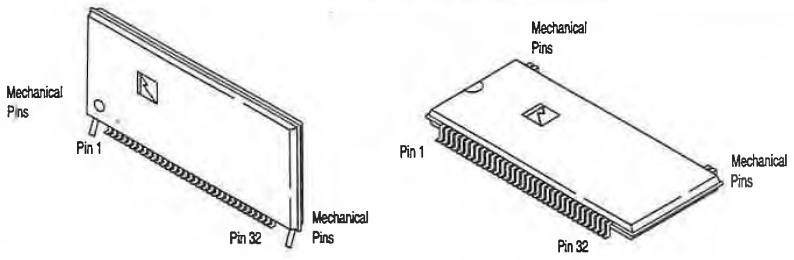


Figure 5: Vertical SVP and Horizontal SHP Packages

Samsung Ex. 1097, p. 813
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

Protocol

The transaction protocol used in Rambus systems is built from several types of information packets. These include the request, acknowledge, serial mode, and data packets.

Request Packet: A master device initiates a transaction by generating a six-byte request packet containing address, control, and byte count information as shown in Figure 6. The Op and OpX fields in the Request packet contain a command that is used to instruct the RDRAM which operation is being requested. A summary of these commands is shown in Table 3.

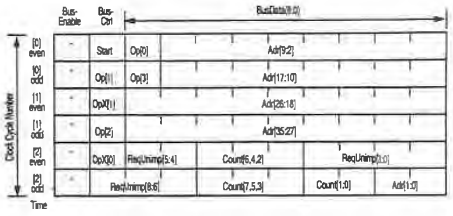


Figure 6: Request Packet[1]

Table 3: Command Summary

Op[3:0]	OpX[1:0]	Name	Description
0000	00	Rseq	Read sequential data from memory space.
0000	01	Rnsq	Read random-access (non-sequential) data from memory space.
0100	00	WseqNpb	Write sequential data to memory space with no per-bit mask application.
0100	01	WseqDpb	Write sequential data to memory space with data-per-bit masking. Static bit masks are supplied by the MDReg while write data is supplied in the data packet.
0100	10	WseqBpb	Write sequential data to memory space with mask-per-bit masking. Both write data and dynamic bit masks are supplied in the data packet.
0100	11	WseqMpb	Write sequential data to memory space with mask-per-bit masking. Static write data is supplied by the MDReg while dynamic bit masks are supplied in the data packet.
0110	00	Rreg	Read data from register space.
0111	00	Wreg	Write data to register space.
1000	00	WnsqNpb	Write random-access (non-sequential) data to memory space with no per-bit mask application.
1000	10	WnsqBpb	Write random-access (non-sequential) data to memory space with mask-per-bit masking. Both write data and dynamic bit masks are supplied in the data packet.
1000	11	WnsqMpb	Write random-access (non-sequential) data to memory space with mask-per-bit masking. Static write data is supplied by the MDReg while dynamic bit masks are supplied in the data packet.
1100	00	WbnsNpb	Write random-access (non-sequential) data to memory space with byte masking and no per-bit mask application. Both byte masks and write data are supplied in the data packet.
1100	01	WbnsDpb	Write random-access (non-sequential) data to memory space with byte masking and data-per-bit masking. Static bit masks are supplied by the MDReg while byte masks and write data are supplied in the data packets.
1100	11	WbnsMpb	Write random-access (non-sequential) data to memory space with byte masking and mask-per-bit masking. Static write data is supplied by the MDReg while byte masks and dynamic bit masks are supplied in the data packets.
1111		WregB	Broadcast write to register space of all responding devices with no acknowledge permitted.

1. A ▭ in this diagram signifies that this pin is not used by this packet. If it is not used by another packet, it is pulled to a logic zero value.

Samsung Ex. 1097, p. 814
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

Acknowledge Packet: Each RDRAM monitors the Channel for a request to access its assigned memory range. The device matching the address range requested then drives an acknowledge packet back to the master.

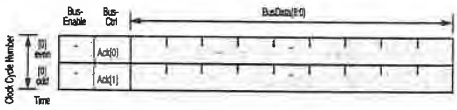


Figure 7: Acknowledge Packet[1]

The Ack[1:0] field in the Acknowledge packet carries the RDRAMs response to the request. If the RDRAM is able to complete the operation as requested, it returns an Okay response. If the RDRAM is unable to complete the operation as requested, it returns a negative acknowledge response (Nack). The encoding of the Ack[1:0] bits is shown in Table 4.

Table 4: Ack[1:0] Encodings

Ack [1:0]	Name	Description
00	Non-existent	Indicates passive acceptance of the request (WregB), or indicates that the addressed device did not respond (all other commands).
01	Okay	Indicates that the request was accepted by the addressed (responding) device.
10	Nack	Indicates that the request could not be accepted because the state of the responding device prevented an access at the fixed timing slot.
11	Ack3	This should not be returned by this responding device. Initiating devices will, when presented with this combination, have an undefined response.

In response to an acknowledged command, the RDRAM either drives a data packet back to the master in the case of a read, or accepts a data packet from the master in the case of a write. Figure 8 shows example 16 byte read and write transactions. The actual timing from the end of a request packet to data and acknowledge packets is adjustable through RDRAM register settings.

1. A [ ] in this diagram signifies that this pin is not used by this packet. If it is not used by another packet, it is pulled to a logic zero value.

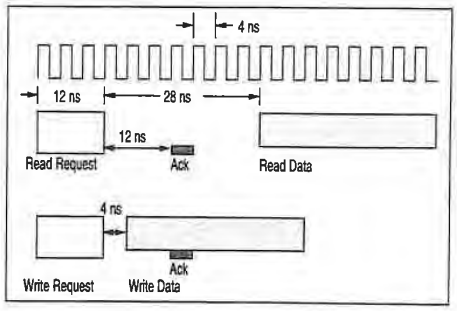


Figure 8: Sample 16-Byte Read and Write Transactions

Serial Address packet: The non-sequential (Random Access) commands specify the eight column address bits needed to access random octbytes within the open page. These address bits are provided using a Serial Address packet.

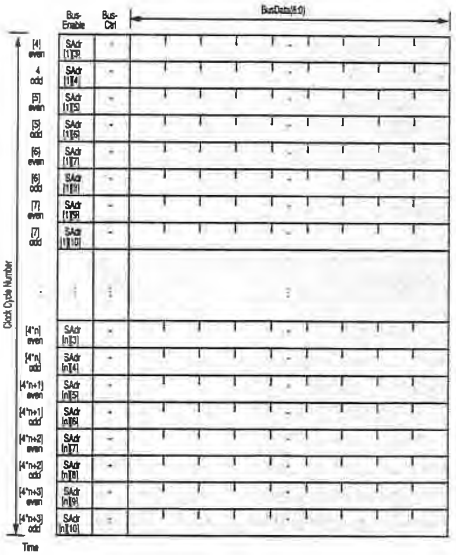


Figure 9: Serial Address Packet Format[1]

Samsung Ex. 1097, p. 815
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

Serial Control Packet: The protocol also allows the Channel master to issue an early termination instruction for a memory read or write transaction. This is done using a Serial Control packet.

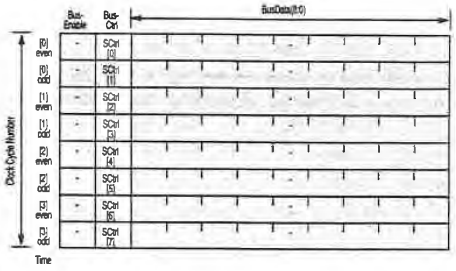


Figure 10: Serial Control Packet Format[1]

The value of the Serial Control field specifies whether the instruction should continue or terminate. This is shown in the table below.

Table 5: Serial Control Field

Serial Control Field	Description	Value
SCtrl[7:0]	Continue	00000000
SCtrl[7:0]	Terminate	00100000

Serial Mode Packet: Serial Mode Packets are used to instruct the RDRAM to perform a operating mode change. A Serial Mode Packet is simply a pulse on the BusEnable line.

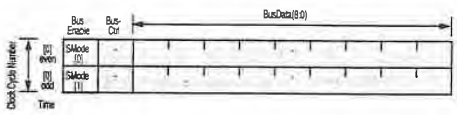


Figure 12: Serial Mode Packet[1]

RDRAM Overview

The figure on the next page is a block diagram of the RDRAM device. The Rambus Channel interface consists of a clock generator, a receiver, and a transmitter. The clock generator uses the external clock signals RxClk and TxClk (tapped off the Channel traces ClockFromMaster and ClockToMaster) and creates the internal signals RClk and TClk. These are used by the receiver and transmitter, respectively, to transfer a bit every 2 nanoseconds on each wire between the RDRAM and the master device. The receiver and transmitter blocks also contain multiplexing and storage hardware to permit the internal RDRAM data paths to operate at the slower clock rate (but equivalent bandwidth) of eight bytes transferred every 15 or 16 nanoseconds (four clock cycles).

The RDRAM also contains control logic and configuration registers. The registers are read and written using special register space commands and control various aspects of RDRAM operation as described on page 11.

The remaining logic consists of a standard DRAM memory core and row sense amplifier caches.

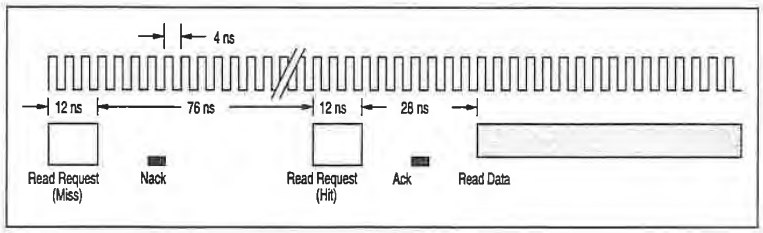


Figure 11: Sample 32-Byte Read Miss and Read Hit Transactions

1. A [-] in this diagram signifies that this pin is not used by this packet. If it is not used by another packet, it is pulled to a logic zero value.

Samsung Ex. 1097, p. 816
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

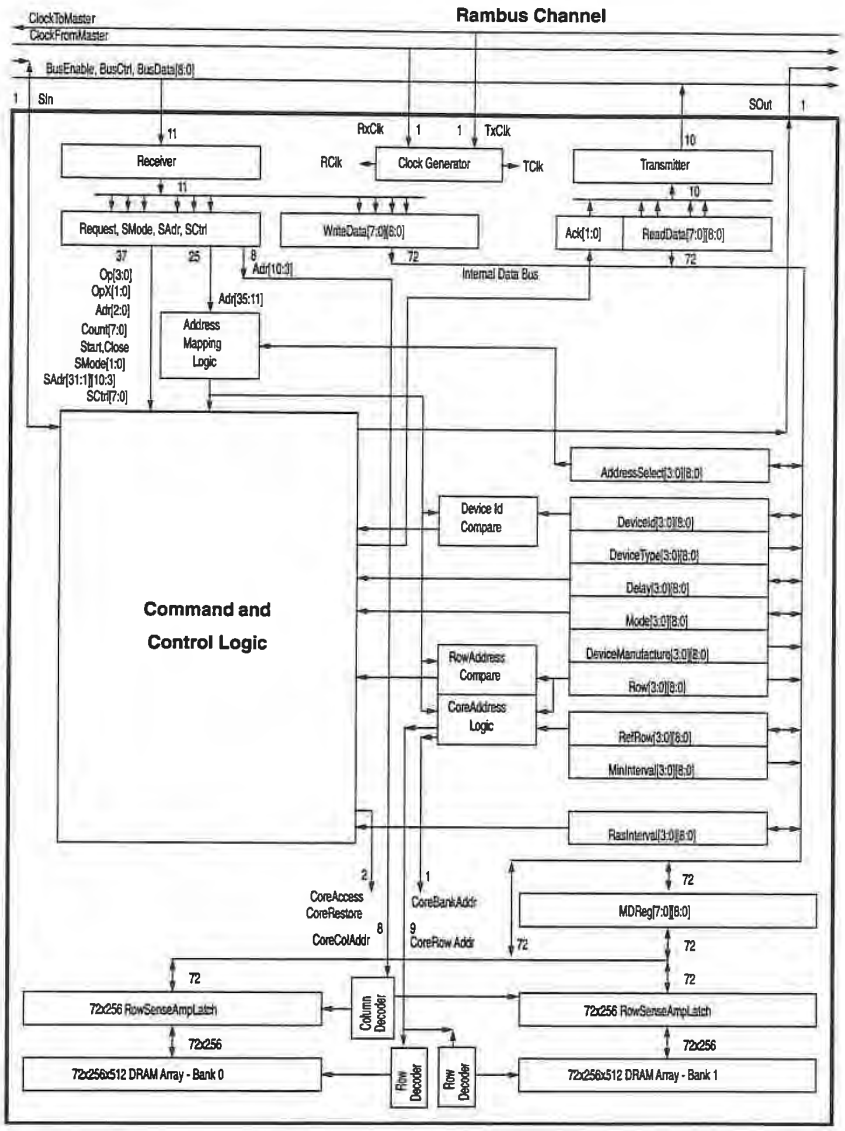


Figure 13: RDRAM Block Diagram

Samsung Ex. 1097, p. 817
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

RDRAM Operation

The RDRAM is broken down into two independent banks of memory, each bank storing a full 1 Mbyte of data (see Figure 13). Each of these banks has a 2KByte open page associated with it that is built out of sense amplifier arrays. These sense amplifiers hold the last accessed row of their associated bank in the sense amplifiers. This allows further accesses to the same row of memory to result in page hits. With the row already stored in the sense amplifiers, subsequent data can be accessed with very low latency. Each RDRAM added to a system adds two open pages to the memory system, helping to increase hit rates.

A page miss results when a row is accessed that is not currently stored as one of the open pages. When this happens, the requesting master is sent a NACK Acknowledge packet indicating the requested row is not yet available. The RDRAM then loads the requested row into the sense amplifiers and waits for the master to submit a retry of the previous request. Figure 11 shows an example of a read miss followed by a read hit for a 32 byte memory read operation.

The amount of time that is needed before the retry can be serviced depends on whether the data in the open page is clean or dirty. The sense amplifiers act as a "write back" cache in that data written to the open page is not written into the actual DRAM cells until the page is closed. If the data in an open page is clean (not previously written) when a new page is requested, the open page does not need to be written back into the DRAM. If the data in an open page is dirty, then additional time must be added to the miss retry delay to account for the writeback operation.

Address Mapping

Address mapping hardware is provided to increase page hit rates by allowing system designers to easily perform n-way RDRAM interleaving. In a non-interleaved memory system, contiguous blocks of addresses follow each other in sequence in one RDRAM, which is then followed by the next RDRAM.

Using address mapping, adjacent blocks of data (2K or greater) can be separated across several RDRAMs, and therefore across several open pages. This allows a more optimal mapping of the pages as caches and creates higher effective page hit rates. In a typical system containing, for example, eight RDRAMs, hit rates could be expected to be as high as 95%. Address mapping is easily adjusted by writing a control register in each RDRAM.

Transaction Concurrency

Concurrent transactions can be used to optimize RDRAM utilization in high performance applications by taking advantage of available Channel bandwidth during page miss latency periods. When a miss in one RDRAM takes place, that device will be busy loading a new row into one of its sense amp caches. The Channel and all other RDRAMs will still be available for use. While waiting for the first RDRAM to finish loading its open page, a transaction to another RDRAM can be initiated. In systems where memory accesses can be queued, a transaction can take place for any pending access residing in a different RDRAM.

Pretouching can be used in systems where certain memory accesses are predictable, such as video applications. This is done when an application is finished with a particular RDRAM and about to access a different one. If the next access to an RDRAM is known in advance, a transaction can be first generated that will cause a row miss and prepare the RDRAM for its next access. When the device is next accessed, the required row of data will already be loaded in the open page and a page hit will take place.

Random Access Mode

Non-contiguous blocks of memory can be accessed through the use of the read and write non-sequential (Random Access) operations. With these commands, multiple eight-byte blocks (octbytes) of data within an open page can be accessed in any order. To do this, the master device sends a request packet specifying a non-sequential operation along with the address of the first octbyte to be accessed. The master device also generates a serial address packet (see Figure 9) on the BusEnable signal that specifies the address of the next octbyte. Successive serial address packets continue to specify new addresses within the open page while data is continuously transferred until the access is complete.

Random Access mode can be used to satisfy the burst order of processors like the Intel™ Pentium ™ at full RDRAM bandwidth.

Samsung Ex. 1097, p. 818
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454



Bit Masking

Three forms of bit masking are available for memory write operations. These operations are referred to as data-per-bit (Dpb), mask-per-bit (Mpb), and both-per-bit (Bpb) masking. An eight-byte Mask Data register (MDReg) within the RDRAM is used to hold the static value of either mask or data information for these operations. The Mask Data register is an RDRAM internal register that is written by the Bpb commands, and is used by the Mpb and Dpb commands.

With the Dpb operation, the MDReg is used to hold a static mask that is applied to all octbytes of data written to the RDRAM core. With the Mpb operation, the MDReg is used to hold an octbyte of static data that is masked by dynamic bit masks supplied in the data packets before being written to the RDRAM core.

The Bpb operation requires data packets to alternate between mask and data octbytes. The even data packets (starting with data packet 0) carry bit masking information which is placed in the MDReg while the odd data packets carry the data to be masked by the latest contents on the MDReg. This type of operation is also used to set the MDReg for later use in Dpb and Mpb operations.

Table 6 shows the source of the mask and data for each of the write commands. The first eight columns show the Wseq and Wnsq sequential and non-sequential (random-access) write commands. Each has four bit mask sub-commands: Npb, Dpb, Mpb, and Bpb.

A write command consists of writing from one to 32 octbytes of data W[31:0]. Each octbyte of data is masked by an octbyte of bitmask Bit[31:0]; that is, if the bitmask bit is set, the corresponding bit of write data is written, if the bitmask bit is clear, the bit in memory is left unchanged.

Each of the first eight columns of Table 6 shows the source of the up-to-32 octbytes of write data and bitmask for the eight write commands. The Npb commands use no bitmask at all - effectively the bitmask is all ones. The Dpb command takes a single octbyte of bitmask from the MDReg and applies it to all data octbytes that are written. The Dpb command does the reverse and takes a single octbyte of write data from the MDReg and writes it to each octbyte of memory using a different bitmask from the data packet. The Bpb commands take an octbype of bitmask and an octbyte of write data alternately, and writes them to a single octbye of memory.

Byte Masking

Contiguous byte masking is supported by the WseqNpb command. This command uses the Adr[2:0] and Count[2:0] fields of the request packet (Figure 6) to specify the byte masks of the first and last octbytes of a data packet. Refer to the *RDRAM Design Guide* for a more complete description of contiguous byte masking. Table 6 does not show the contiguous byte masking option.

The RDRAM also supports a more general form of byte masking called non-contiguous byte masking. This is available with the Wbns write commands.

As can be seen in column nine of Table 6 for the WbnsNpb command, the first data packet, and every ninth thereafter, contain byte masking information that is applied to the eight data packets that follow. This means data packets 0, 9, 18, and 27 (gray boxes in the table) are not written to memory, but are instead used as byte masks for the eight octbytes of data that follow. This means that the data packet Data[35:0] may consist of up to 36 octbytes of information: 4 octbytes of byte mask information and 32 octbytes of data that is actually written to memory.

Each bit of the 64-bit byte mask (one octbyte) controls whether a byte of the following 64 bytes of data (eight octbytes) is written (one) or not written (zero) to memory.

Static bit masking is also available with non-contiguous byte masking. This is shown in the last two columns of Table 6 labeled WbnsDpb and WbnsMpb. The WbnsDpb command applies a static bit held in the MDReg to each octbyte of data W[31:0]. Each bit of the data octbyte is written if the corresponding bits of both the bitmask and bytemask are a one, and is not written if either is a zero.The WbnsMpb command applies a different bitmask Bit[31:0] to a static data octbyte held in the MDReg. Each bit of the data octbyte is written if the corresponding bits of both the bitmask and bytemask are a one, and is not written if either is a zero.

Samsung Ex. 1097, p. 819
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

Table 6: WriteData, BitMask, and ByteMask Sources[2]

Data Octbyt	Wseq Npb	Wseq Dpb	Wseq Mpb	Wseq Bpb	Wnsq Npb	Wnsq Dpb	Wnsq Mpb	Wnsq Bpb	Wbns Npb	Wbns Dpb	Wbns Mpb
MDReg	-	Bit [35:0]	W [35:0]	-	-	Bit [35:0]	W [35:0]	-	-	Bit [35:0]	W [35:0]
Data[0]	W[0]	W[0]	Bit[0]	Bit[0]	W[0]	W[0]	Bit[0]	Bit[0]	Byte [7:0]	Byte [7:0]	Byte [7:0]
Data[1]	W[1]	W[1]	Bit[1]	W[0]	W[1]	W[1]	Bit[1]	W[0]	W[0]	W[0]	Bit[0]
Data[2]	W[2]	W[2]	Bit[2]	Bit[1]	W[2]	W[2]	Bit[2]	Bit[1]	W[1]	W[1]	Bit[1]
Data[3]	W[3]	W[3]	Bit[3]	W[1]	W[3]	W[3]	Bit[3]	W[1]	W[2]	W[2]	Bit[2]
Data[4]	W[4]	W[4]	Bit[4]	Bit[2]	W[4]	W[4]	Bit[4]	Bit[2]	W[3]	W[3]	Bit[3]
Data[5]	W[5]	W[5]	Bit[5]	W[2]	W[5]	W[5]	Bit[5]	W[2]	W[4]	W[4]	Bit[4]
Data[6]	W[6]	W[6]	Bit[6]	Bit[3]	W[6]	W[6]	Bit[6]	Bit[3]	W[5]	W[5]	Bit[5]
Data[7]	W[7]	W[7]	Bit[7]	W[3]	W[7]	W[7]	Bit[7]	W[3]	W[6]	W[6]	Bit[6]
Data[8]	W[8]	W[8]	Bit[8]	Bit[4]	W[8]	W[8]	Bit[8]	Bit[4]	W[7]	W[7]	Bit[7]
Data[9]	W[9]	W[9]	Bit[9]	W[4]	W[9]	W[9]	Bit[9]	W[4]	Byte [15:8]	Byte [15:8]	Byte [15:8]
Data[10]	W[10]	W[10]	Bit[10]	Bit[5]	W[10]	W[10]	Bit[10]	Bit[5]	W[8]	W[8]	Bit[8]
Data[11]	W[11]	W[11]	Bit[11]	W[5]	W[11]	W[11]	Bit[11]	W[5]	W[9]	W[9]	Bit[9]
Data[12]	W[12]	W[12]	Bit[12]	Bit[6]	W[12]	W[12]	Bit[12]	Bit[6]	W[10]	W[10]	Bit[10]
Data[13]	W[13]	W[13]	Bit[13]	W[6]	W[13]	W[13]	Bit[13]	W[6]	W[11]	W[11]	Bit[11]
Data[14]	W[14]	W[14]	Bit[14]	Bit[7]	W[14]	W[14]	Bit[14]	Bit[7]	W[12]	W[12]	Bit[12]
Data[15]	W[15]	W[15]	Bit[15]	W[7]	W[15]	W[15]	Bit[15]	W[7]	W[13]	W[13]	Bit[13]
Data[16]	W[16]	W[16]	Bit[16]	Bit[8]	W[16]	W[16]	Bit[16]	Bit[8]	W[14]	W[14]	Bit[14]
Data[17]	W[17]	W[17]	Bit[17]	W[8]	W[17]	W[17]	Bit[17]	W[8]	W[15]	W[15]	Bit[15]
Data[18]	W[18]	W[18]	Bit[18]	Bit[9]	W[18]	W[18]	Bit[18]	Bit[9]	Byte [23:16]	Byte [23:16]	Byte [23:16]
Data[19]	W[19]	W[19]	Bit[19]	W[9]	W[19]	W[19]	Bit[19]	W[9]	W[16]	W[16]	Bit[16]
Data[20]	W[20]	W[20]	Bit[20]	Bit[10]	W[20]	W[20]	Bit[20]	Bit[10]	W[17]	W[17]	Bit[17]
Data[21]	W[21]	W[21]	Bit[21]	W[10]	W[21]	W[21]	Bit[21]	W[10]	W[18]	W[18]	Bit[18]
Data[22]	W[22]	W[22]	Bit[22]	Bit[11]	W[22]	W[22]	Bit[22]	Bit[11]	W[19]	W[19]	Bit[19]
Data[23]	W[23]	W[23]	Bit[23]	W[11]	W[23]	W[23]	Bit[23]	W[11]	W[20]	W[20]	Bit[20]
Data[24]	W[24]	W[24]	Bit[24]	Bit[12]	W[24]	W[24]	Bit[24]	Bit[12]	W[21]	W[21]	Bit[21]
Data[25]	W[25]	W[25]	Bit[25]	W[12]	W[25]	W[25]	Bit[25]	W[12]	W[22]	W[22]	Bit[22]
Data[26]	W[26]	W[26]	Bit[26]	Bit[13]	W[26]	W[26]	Bit[26]	Bit[13]	W[23]	W[23]	Bit[23]
Data[27]	W[27]	W[27]	Bit[27]	W[13]	W[27]	W[27]	Bit[27]	W[13]	Byte [31:24]	Byte [31:24]	Byte [31:24]
Data[28]	W[28]	W[28]	Bit[28]	Bit[14]	W[28]	W[28]	Bit[28]	Bit[14]	W[24]	W[24]	Bit[24]
Data[29]	W[29]	W[29]	Bit[29]	W[14]	W[29]	W[29]	Bit[29]	W[14]	W[25]	W[25]	Bit[25]
Data[30]	W[30]	W[30]	Bit[30]	Bit[15]	W[30]	W[30]	Bit[30]	Bit[15]	W[26]	W[26]	Bit[26]
Data[31]	W[31]	W[31]	Bit[31]	W[15]	W[31]	W[31]	Bit[31]	W[15]	W[27]	W[27]	Bit[27]
Data[32]	-	-	-	-	-	-	-	-	W[28]	W[28]	Bit[28]
Data[33]	-	-	-	-	-	-	-	-	W[29]	W[29]	Bit[29]
Data[34]	-	-	-	-	-	-	-	-	W[30]	W[30]	Bit[30]
Data[35]	-	-	-	-	-	-	-	-	W[31]	W[31]	Bit[31]

2. The shaded data packet contains byte masking information that is applied to the eight data packets that follow.

Samsung Ex. 1097, p. 820
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

RDRAM Registers

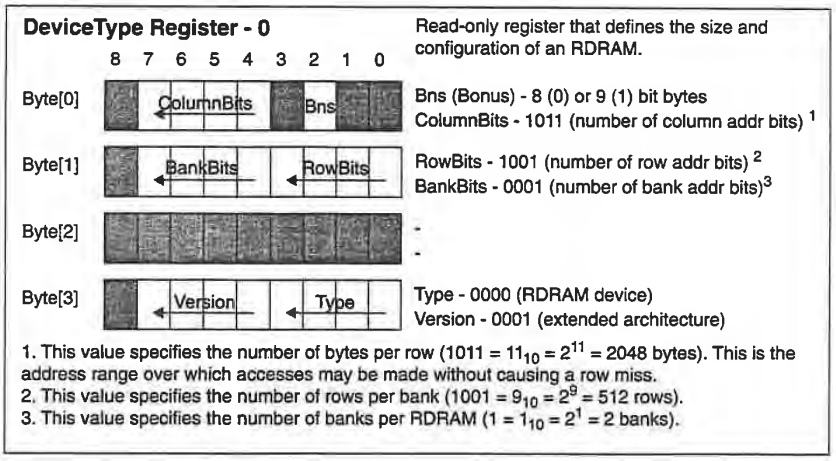
The 16M RDRAM contains ten registers. These are read and written with the Rreg , Wreg, and WregB commands. They are used to provide configuration information to the RDRAM controller (DeviceType, MinInterval, and DeviceManufacture), to control device, bank, and row addressing (DeviceId, AddressSelect, and Row), to control refresh (RefRow), to control RDRAM timing (Delay and RasInterval), and to control RDRAM operation (Mode). The following table summarizes these functions:

Table 7: Register Summary

Register Name	Reg. #	Description
DeviceType[3:0][8:0]	0	Read-only register that defines the size and configuration of the RDRAM.
DeviceId[3:0][8:0]	1	Used to specify the base address for the RDRAM.
Delay[3:0][8:0]	2	Used to specify CAS timing parameters.
Mode[3:0][8:0]	3	Used to initialize the RDRAM and set the I_{OL} output current.
RefRow[3:0][8:0]	5	Used to specify the next row and bank of the RDRAM to be refreshed.
RasInterval[3:0][8:0]	6	Used to specify RAS timing parameters.
MinInterval[3:0][8:0]	7	Read-only register defining minimum timing parameters for CAS accesses.
AddressSelect[3:0][8:0]	8	Used to specify address bit swapping to maximize RDRAM cache hit rate.
DeviceManufacturer[3:0][8:0]	9	Read-only register containing a manufacturer code.
Row[3:0][8:0]	128	Used to specify the currently sensed row in each bank.

The following diagrams show the individual fields of the RDRAM registers. The color of a field denotes its usage: dark-gray is unimplemented, light-gray is read-only, and white is read-write. The arrow within each multi-bit field points from least-significant bit to most-significant bit. Bit and byte numbering use little-endien notation.

Figure 14: Registers

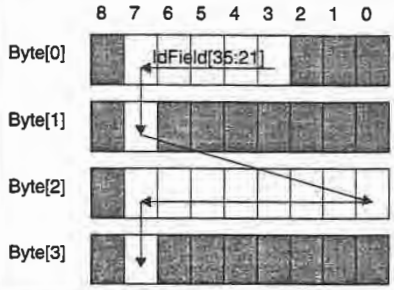


Samsung Ex. 1097, p. 821
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454



DeviceId Register - 1

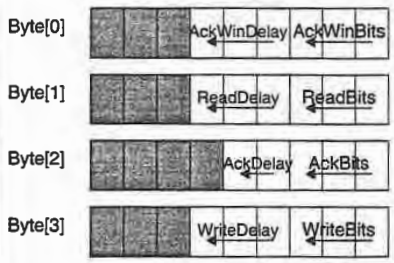
Read-write register for setting the base address of the RDRAM.



IdField - Compared to Adr[35:21] of request packet to select RDRAM.

Delay Register - 2

Read-write and read-only register to program the CAS access delays.



AckWinBits - 011 (Number of AckWinDelay bits)
AckWinDelay - Normally set to 101 (5 busclocks)

ReadBits - 011 ((Number of ReadDelay bits)
ReadDelay - Normally set to 111 (7 busclocks)

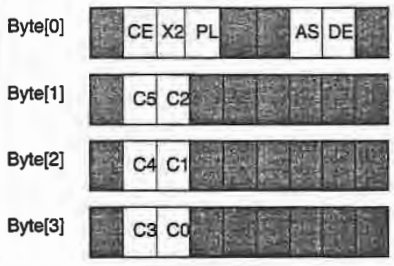
AckBits - 010 ((Number of AckDelay bits)
AckDelay - Normally set to 011 (3 busclocks)

WriteBits - 011 ((Number of WriteDelay bits)
WriteDelay - Normally set to 001(1 busclocks)

AckWinDelay adjusts the size of the acknowledge window. ReadDelay, WriteDelay, and AckDelay adjust the time from the end of the packet to the start of read data, write data, and the acknowledge packets, respectively.

Mode Register - 3

Read-write register for initializing the RDRAM and for controllling operating modes.



DE(DevEn) - Selects RDRAM at initialization
AS - Set to one

PL - Selects powerdown wake-up time
X2(CCMult) - Set to one. Specifies I_{OL} current

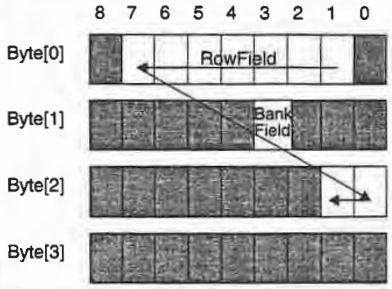
CE - Set to one. Specifies I_{OL} output current

C[5:0] -Set to {63...31}. Specifies I_{OL} output current

Samsung Ex. 1097, p. 822
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

RefRow Register - 5

Read-write register for setting the next row refreshed by SetRR burst refresh.

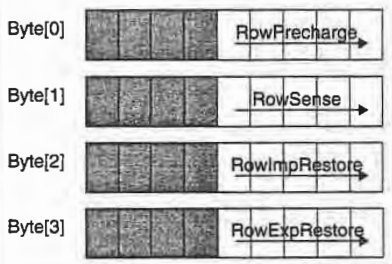


RowField - Next row to be refreshed by SetRR

BankField - Next bank to be refreshed by SetRR

RasInterval Register - 6

Read-write register to program the RAS access intervals. Note that fields are in bit-reversed order.



RowPrecharge - Set the precharge interval (normal value is 10000 = 1 busclock cycles)

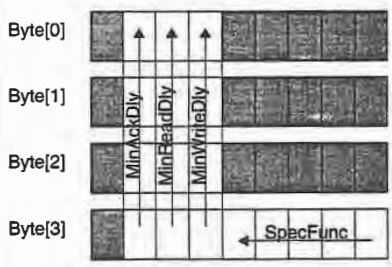
RowSense - Set the sense interval (normal value is 01100 = 6 busclock cycles)

RowImpRestore - Set the implicit restore interval (normal value is 10010 = 9 busclock cycles)

RowExpRestore - Set the explicit restore interval (normal value is 01000 = 2 busclock cycles)

MinInterval Register - 7

Read-only register (configuration info) and write-only register (special control).



MinWriteDelay - 0001 (minimum WriteDelay of RDRAM)

MinReadDelay - 0111 (minimum ReadDelay of RDRAM)

MinAckDelay - 0011 (minimum AckDelay of RDRAM)

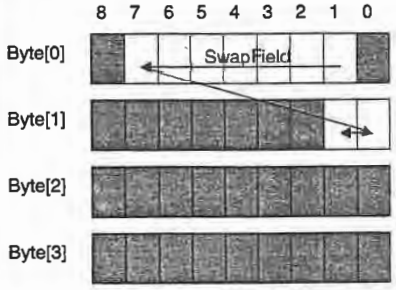
SpecFunc - Performs SetRR burst refresh and SetPD powerdown entry

MinWriteDelay, MinReadDelay, and MinAckDelay specify the minimum number of cycles allowed between a request packet and a write data, read data, and acknowledge packet, respectively. SpecFunc is a write-only field that is used to initiate a SetRR burst refresh or to place the RDRAM into powerdown mode.

Samsung Ex. 1097, p. 823
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

AddressSelect Register - 8

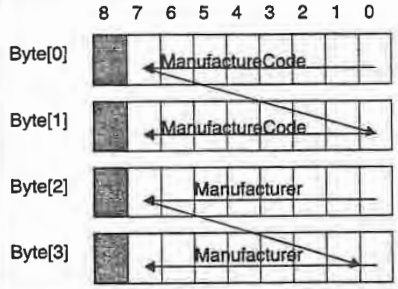
Read-write register for swapping subfields of the Adr field of request packet. This maximizes the rowhit rate for many applications.



SwapField - Each bit swaps a pair of bits from Adr[28:20] and Adr[19:11].

DeviceManufacture Register - 9

Read-only register with configuration information for RDRAM.

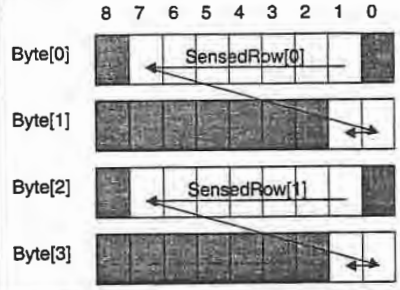


ManufactureCode - Manufacturer specific information.

Manufacturer - Contains code specifying the manufacturing company.

Row Register - 128

Read-write register with addresse of currently-sensed row in each bank of RDRAM.



SensedRow[0] - Address of currently sensed row of bank 0.

SensedRow[1] - Address of currently sensed row of bank 1.

Samsung Ex. 1097, p. 824
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

Refresh

The RDRAM is a dynamic device, and the memory array must be refreshed every 17 ms (t_{REF}). To support this need, the RDRAM includes all of the logic necessary to support three refresh modes built in. These refresh modes are:

- Manual Refresh: The Rambus Channel master uses a register write transaction (SetRR) to initiate a single burst refresh of four rows.
- Touching: A single row is refreshed each time that a read or write request is made to that row in the RDRAM.
- Powerdown Mode Refresh: A single row is refreshed with each pulse on the SIn/SOut pins.

When the RDRAM is in PowerDown mode, it can be refreshed by passing a periodic pulse at a frequency of 60.2 kHz or greater through the SIn/SOut pins. This minimum frequency is equal to $1024/t_{REF}$.

Operating Modes

The RDRAM has three operating modes; Active, Standby and PowerDown. The three modes are distinguished by two factors, their power consumption, and the time that it will take the RDRAM to service a request from that mode.

The control logic within the RDRAM includes a counter that counts Serial Mode packets. It takes a specific number of packets to cause the RDRAM to transition from a low-power mode to the Active state. This counter is active in all three operating modes.

In Active mode, the RDRAM is active and ready to immediately service a request packet. Power consumption is also highest in Active mode.

An RDRAM automatically transitions to Standby mode at the end of a transaction. While in this low power state, each RDRAM monitors the BusEnable signal for a serial mode packet while ignoring other activity on the remaining Channel signals. The Channel master sends a serial mode packet bring all RDRAMs temporarily out of Standby and into Active mode so they can respond to a request packet. Once the request packet is acknowledged, all of the RDRAMs return to Standby mode with the exception of the one responding to the request. That device returns to Standby mode once the read or write operation is complete.

Unlike conventional DRAM memory systems where each device in an entire bank of memory must be kept active and consuming power through an entire access, Rambus memory systems use only one active device while all others remain in a lower power state.

Power consumption may be greatly reduced by using the PowerDown mode. This mode is entered manually by setting the Special Function bit *SetPD* in the MinInterval register. Entering this mode causes the device to write back and precharge both cache lines, disable the internal clock generator, and disable most DC current sources. The BusEnable receiver is kept active to detect serial mode packets used to exit powerdown mode. The only significant power consumption in powerdown mode is due to refresh.

Since the RDRAM's internal clocks are disabled while in powerdown mode, refresh must be maintained manually by the master device. This is done by supplying a low frequency square wave on the SIn TTL signal. This propagates through each RDRAM and is used to initiate asynchronous refresh operations in each device.

Each RDRAM may be placed in either low or high threshold powerdown mode. Threshold refers to the number of serial mode packets required to wake up the RDRAM. A low threshold requires relatively few serial mode packets while a high threshold requires a larger number. The actual power dissipation is identical in both modes.

An example of where these modes are used is in a portable computer application. In this example, sleep mode is implemented by placing a majority of the RDRAMs in high threshold powerdown while the RDRAM that contains the frame buffer is placed in low threshold powerdown. This permits screen refresh to take place without powering up the entire memory system.

Samsung Ex. 1097, p. 825
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454



Table 8: RDRAM State Transitions

Current State	Operation	Next State
ResetState	Count < $t_{MODEARMAX}$[a]	ActiveState
ActiveState	Count >= $t_{MODEARMAX}$	ResetState
ActiveState	Count >= $t_{MODEDELAYMAX}$[b]	StandbyState
ActiveState	Write *SetPD* in the *MinInterval* Register	PowerDownState
ActiveState	$BusCtrl_{EVEN}$ = 1 (Start Bit)	Transaction Framing States
StandbyState	Count >= $t_{MODESAMIN}$[c]	ActiveState
PowerDownState	Count >= $t_{MODEPA[PL]MIN}$[d]	ActiveState

a. $t_{MODEARMAX}$ is the maximum number of SMode packets necessary to cause a transition from ActiveState to ResetState.
b. $t_{MODEDELAYMAX}$ is maximum delay (in clock cycles) after a transaction is complete for the RDRAM to enter the StandbyState.
c. $t_{MODESAMIN}$ in the minimum number of SMode packets necessary to cause a transition from StandbyState to ActiveState.
d. $t_{MODEPAMIN}$ is the minimum number of SMode packets necessary to cause a transition from PowerDownState to ActiveState.

Initialization

The Channel master resets the RDRAM devices on the Channel by asserting the BusEnable signal for $t_{MODEAR,MIN}$ clock cycles. After the RDRAM has been reset, the base address and register space registers revert to their default values. Because the default address of all devices is zero until initialization is complete, individual devices cannot be addressed from the Channel, although the devices can and will respond to broadcast commands.

In a typical system application, the SIn pin of the first RDRAM is connected to the SOut pin of the Channel master (refer Figure 2). The SIn pin of the next RDRAM is connected to the SOut pin of the first RDRAM, and so on. SOut of the last device is then connected to SIn of the Channel master to complete the daisy chain. An RDRAM will not respond to a write command (other than a broadcast write) until the SIn pin is set to 1. If power-down mode is not being used, two pins can be saved on this controller by connecting the SIn pin of the first RDRAM to a logic high, and leaving the SOut pin of the last RDRAM unconnected.

To start the initialization sequence, SIn is asserted on the first RDRAM in the chain to enable it. The Channel master then writes the desired device address to the AddressSelection register and enables it using the *DevEn* bit. Enabling the device sets the SOut pin (and the SIn pin on the next device in the chain) to 1. This process continues until all of the RDRAMs have been initialized.

Initializing an individual RDRAM consists of these steps:

1. Broadcast write any shared register parameters to the RDRAM registers. Shared parameters are values that are common to all of the devices in the system, such as:
   - $t_{ACKWINDELAY}$
   - $t_{READDELAY}$
   - $t_{ACKDELAY}$
   - $t_{WRITEDELAY}$
2. Load the DeviceID register of the first (next) device using a register write transaction. The four bytes of the DeviceID should be unique to that device.
3. Set the *DevEn* (Device Enable) bit in the Mode register to enable the RDRAM. Setting DevEn will cause the assertion of the SIn pin of the next device in the SIn/SOut chain.
4. Loop through steps 2 and 3 until all of the RDRAMS have been initialized.

When these steps have been completed for every device in the chain, all of the slave devices should have unique, contiguous DeviceID values, and will have their *DevEn* bits set. Any further configuration of the devices can be performed across the Channel.

Samsung Ex. 1097, p. 826
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

Absolute Maximum Ratings

The following table represents stress ratings only, and functional operation at the maximums is not guaranteed. Extended exposure to the maximum ratings may affect device reliability. Furthermore, although devices contain protective circuitry to resist damage from static electric discharge, always take precautions to avoid high static voltages or electric fields.

Symbol	Parameter	Min	Max	Unit
$V_{I,ABS}$	Voltage applied to any RSL pin with respect to Gnd	- 0.5	$V_{DD,MAX}$+0.5	V
$V_{I,TTL,ABS}$	Voltage applied to any TTL pin with respect to Gnd	- 0.5	V_{DD}+0.5	V
$V_{DD,ABS}$	Voltage on VDD with respect to Gnd	- 0.5	$V_{DD,MAX}$+1.0	V
$T_{J,ABS}$	Junction temperature under bias	- 55	125	°C
T_{STORE}	Storage temperature	- 55	125	°C

Thermal Parameters

Symbol	Parameter and Conditions	Min	Max	Unit
T_J	Junction operating temperature	0	100	°C
Θ_{JC}	Junction-to-Case thermal resistance		5	°C/Watt

Capacitance

Symbol	Parameter and Conditions	Min	Max	Unit
C_I	Low-swing input parasitic capacitance		2	pF
$C_{I,TTL}$	TTL input parasitic capacitance		8	pF

Power Consumption

Mode	Description	Min	Max	Unit
Powerdown	Device shut down		0.33[a]	mW
Standby	Device inactive		220[a]	mW
Active	Device evaluating request packet		550[a]	mW
Read	Data being transferred from device		760[a]	mW
Write	Data being transferred to device		860[a]	mW

a. These power numbers are manufacturer-dependent; the numbers shown represent "typical" maximum power levels.

Samsung Ex. 1097, p. 827
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

Recommended Electrical Conditions

Symbol	Parameter and Conditions	Min	Max	Unit
$V_{DD, VDDA}$	Supply voltage — 3.3-volt version	3.0	3.6	V
V_{REF}	Reference voltage	1.9	V_{DD} - 0.8	V
V_{IL}	Input low voltage	V_{REF} - 0.35	V_{REF} - 0.8	V
V_{IH}	Input high voltage	V_{REF} + 0.35	V_{REF} + 0.8	V
$V_{IL,TTL}$	TTL input low voltage	- 0.5	0.8	V
$V_{IH,TTL}$	TTL input high voltage	1.8	V_{DD}+0.5	V

Electrical Characteristics

Symbol	Parameter and Conditions	Min	Max	Unit
I_{REF}	V_{REF} current @ $V_{REF,MAX}$	-10	10	µA
I_{OH}	Output high current @ ($0 \le V_{OUT} \le VDD$)	-10	10	µA
I_0	I_{OL} Output current @ V_{OUT} = 1.6V @ C[5:0] = 111111 (63_{10})[a]	0.0	4.0	mA
I_{20}	I_{OL} Output current @ V_{OUT} = 1.6V @ C[5:0] = 110001(49_{10})[a]	18.0	22.0	mA
I_{40}	I_{OL} Output current @ V_{OUT} = 1.6V @ C[5:0] = 011111 (31_{10})[a]	36.0	44.0	mA
$I_{I,TTL}$	TTL input leakage current @ ($0 \le V_{I,TTL} \le VDD$)	-10.0	10.0	µA
$V_{OL,TTL}$	TTL output voltage @ $I_{OL,TTL}$= 1.0mA	0.0	0.4	V
$V_{OH,TTL}$	TTL output high voltage @ $I_{OH,TTL}$= -0.25mA	2.0	V_{DD}	V

a. This is the value written into the C[5:0] field of the Mode register. Values of I_{OL} in between the I_0, I_{20}, and I_{40} values are produced by interpolatingC[5:0] to intermediate values. For example, C[5:0] = 101000 (40_{10}) produces an I_{OL} in the range of 27.0 to 33.0 mA.

Recommended Timing Conditions

Symbol	Parameter	Min	Max	Unit
t_{CR}, t_{CF}	TxClk and RxClk input rise and fall times	0.3	0.8	ns
t_{CYCLE}	TxClk and RxClk cycle times	3.75	4.5	ns
t_{TICK}	Transport time per bit per pin (this timing interval is synthesized by the RDRAM's internal clock generator)	0.5 (2ns @ t_{CYCLE} = 4ns)	0.5 (2.5ns @ t_{CYCLE} = 5ns)	t_{CYCLE}
t_{CH}, t_{CL}	TxClk and RxClk high and low times	45%	55%	t_{CYCLE}
t_{TR}	TxClk-RxClk differential	0	0.7	ns
t_{DR}, t_{DF}	Data/Control input rise and fall times	0.3	0.6	ns

Samsung Ex. 1097, p. 828
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

Symbol	Parameter	Min	Max	Unit
t_S	Data/Control-to-RxClk setup time	0.35		$t_{CYCLE}/4$
t_H	RxClk-to-Data/Control hold time	0.35		$t_{CYCLE}/4$
t_{REF}	Refresh interval		17	ms
$t_{LOCK,RESET}$	RDRAM internal clock generator lock time from Reset mode		750 (3µs @ t_{CYCLE} = 4ns)	t_{CYCLE}
$t_{LOCK,POWERUP}$	RDRAM internal clock generator lock time from PowerUp mode		750 (3µs @ t_{CYCLE} = 4ns)	t_{CYCLE}

Timing Characteristics

Symbol	Parameter	Min	Max	Unit
t_{PIS}	SIn-to-SOut propagation delay @ $C_{LOAD,TTL=40pF}$		25	ns
t_Q	TClk-to-Data/Control output time	0.6	1.4	$t_{CYCLE}/4$
t_{QR}, t_{QF}	Data/Control output rise and fall times	0.3	0.5	ns

Rambus Channel Timing

The next table shows important timings on the Rambus Channel for common operations. Please refer to the *RDRAM Design Guide* for all possible interactions that could occur on the Rambus Channel. All timings are from the point of view of the Channel master, and thus have the bus overhead delay of 4ns per bus transversal included where appropriate.

Symbol	Parameter	Min	Max	Unit
t_{CYCLE}	TxClk and RxClk cycle times	3.75	4.5	ns
$t_{RESPONSE}$	Start of request packet to start of acknowledge packet	6 [a]	9 [a]	t_{CYCLE}
$t_{READHIT}$	Start of request packet to start of read data packet for row hit (Okay)	10 [a]	17 [a]	t_{CYCLE}
$t_{WRITEHIT}$	Start of request packet to start of write data packet for row hit (Okay)	4 [a]	11 [a]	t_{CYCLE}
$t_{RETRYSENSED}$ CLEAN (no restore)	Start of request packet for row miss (Nack) to start of request packet for row hit (Okay). The previous row had not been written	22 [b]		t_{CYCLE}
$t_{RETRYSENSED}$ DIRTY (restore)	Start of request packet for row miss (Nack) to start of request packet for row hit (Okay). The previous row had been written	30 [b]		t_{CYCLE}
$t_{RETRYREFRESH}$ CLEAN (no restore)	Start of request that performs a burst refresh (SetRR) until the start of a request that will not have a Nack acknowledge due to the pending refresh. The previously sensed row had not been written	209		t_{CYCLE}
$t_{RETRYREFRESH}$ DIRTY (restore)	Start of request that performs a burst refresh (SetRR) until the start of a request that will not have a Nack acknowledge due to the pending refresh. The previously sensed row had been written	217		t_{CYCLE}

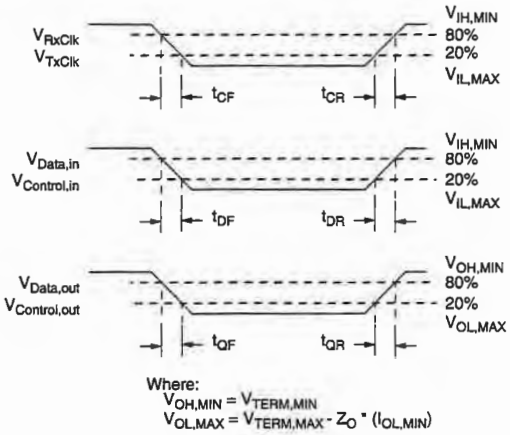
Samsung Ex. 1097, p. 829
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

Symbol	Parameter	Min	Max	Unit
$t_{READBURST32}$	Start of request packet to end of 32 byte read data packet for row hit (Okay)	26 [c]		t_{CYCLE}
$t_{READBURST256}$	Start of request packet to end of 256 byte read data packet for row hit (Okay)	138		t_{CYCLE}
$t_{WRITEBURST32}$	Start of request packet to end of 32 byte write data packet for row hit (Okay)	20 [d]		t_{CYCLE}
$t_{WRITEBURST256}$	Start of request packet to end of 256 byte write data packet for row hit (Okay)	132		t_{CYCLE}
$t_{READDELAY}$	End of request packet to beginning of read data packet	7 [a]	14	t_{CYCLE}
$t_{WRITEDELAY}$	End of request packet to beginning of write data packet	1 [a]	8	t_{CYCLE}
$t_{ACKDELAY}$	End of request packet to beginning of acknowledge packet	3 [a]	6	t_{CYCLE}
$t_{ACKWINDELAY}$	Window in which an acknowledge packet will be sent	5 [a]	12	t_{CYCLE}
$t_{SERIALREAD-OFFSET}$	Delay from the beginning of a serial address packet or serial control packet to the beginning of the corresponding read data subpacket	13	13	t_{CYCLE}
$t_{SERIALWRITE-OFFSET}$	Delay from the beginning of a serial address packet or serial control packet to the beginning of the corresponding write data subpacket	5	5	t_{CYCLE}
$t_{POSTMEM-WRITEDELAY}$	Delay from the end of the current memory space transaction to the beginning of the next memory space transaction	2		t_{CYCLE}
$t_{POSTREG-WRITEDELAY}$	Delay from the end of the current register space transaction to the beginning of the next register space transaction	4		t_{CYCLE}
$t_{MODEOFFSET}$	Offset from the beginning of SMode packet to request packet for standby to active transaction	4	4	t_{CYCLE}
t_{MODESA}	Number of SMode packets to cause a transition from StandbyMode to ActiveMode	1	4	t_{CYCLE}
$t_{MODEPA[0]}$	Number of SMode packets to cause a transition from PowerDown-Mode [0] to ActiveMode	16	20	t_{CYCLE}
$t_{MODEPA[1]}$	Number of SMode packets to cause a transition from PowerDown-Mode [1] to ActiveMode	208	224	t_{CYCLE}
t_{MODEAR}	Number of SMode packets necessary to cause a transition from ActiveState to ResetState		288	t_{CYCLE}

a. Programmable - All RDRAMS will operate across the full programming range.
b. Minimum at $t_{CYCLE,MIN}$. The delay is programmable to give equivalent timings at longer t_{CYCLE}.
c. Calculated with $t_{READHIT,MIN}$
d. Calculated with $t_{WRITEHIT,MIN}$

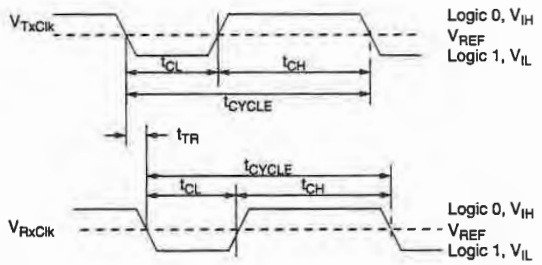
Samsung Ex. 1097, p. 830
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

Rise/Fall Timing

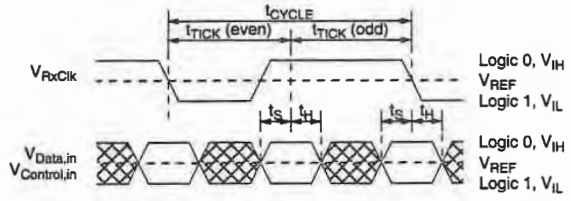


Where:
$V_{OH,MIN} = V_{TERM,MIN}$
$V_{OL,MAX} = V_{TERM,MAX} - Z_O * (I_{OL,MIN})$

Clock Timing



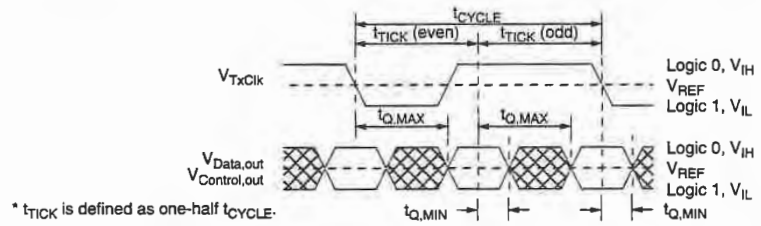
Receive Data Timing



* t_{TICK} is defined as one-half t_{CYCLE}.

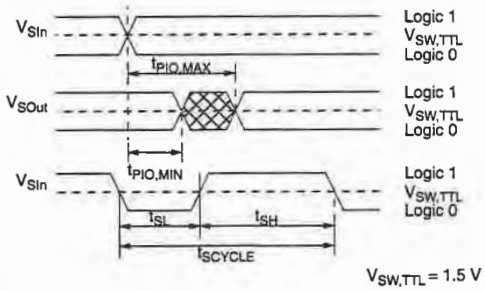
Samsung Ex. 1097, p. 831
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

Transmit Data Timing



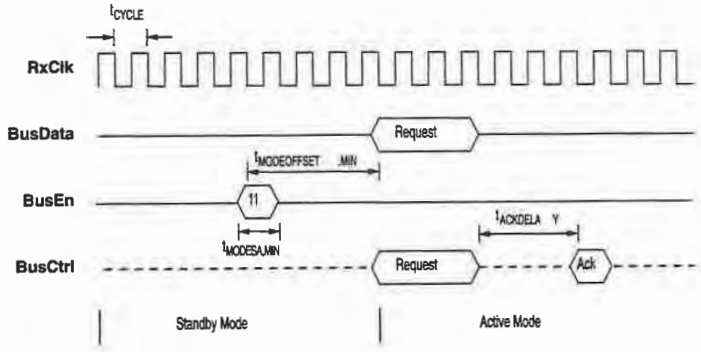
* t_{TICK} is defined as one-half t_{CYCLE}.

Serial Configuration Pin Timing



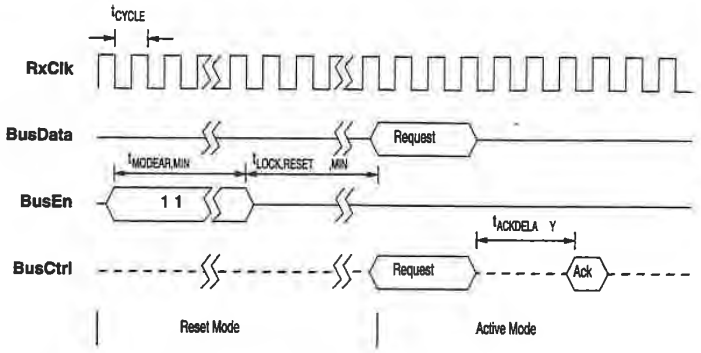
$V_{SW,TTL}$ = 1.5 V

Standby Mode to Active Mode Timing

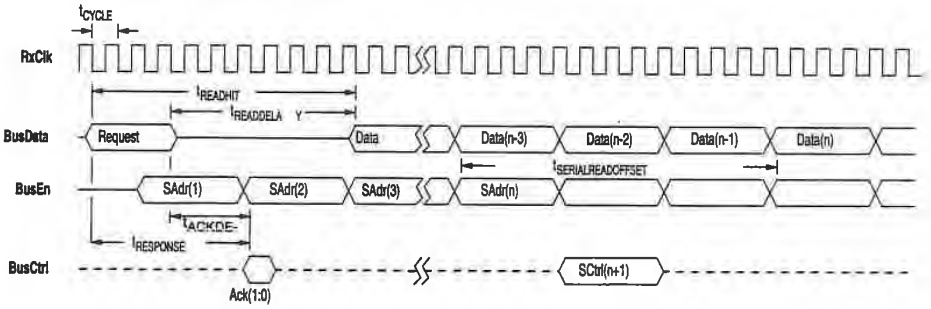


Samsung Ex. 1097, p. 832
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

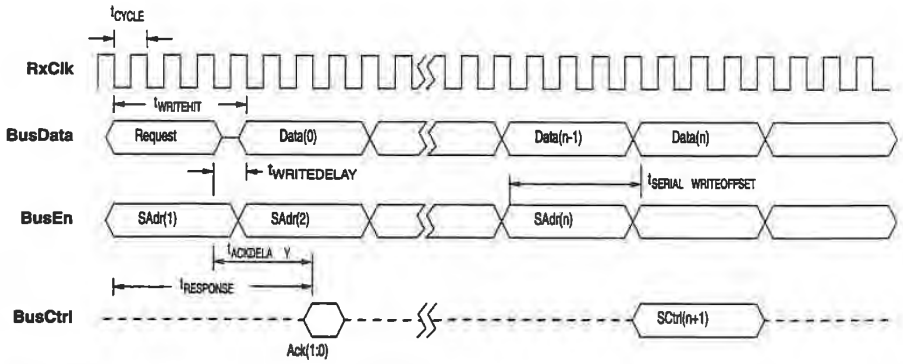
Reset Timing



Read Hit Timing Diagram

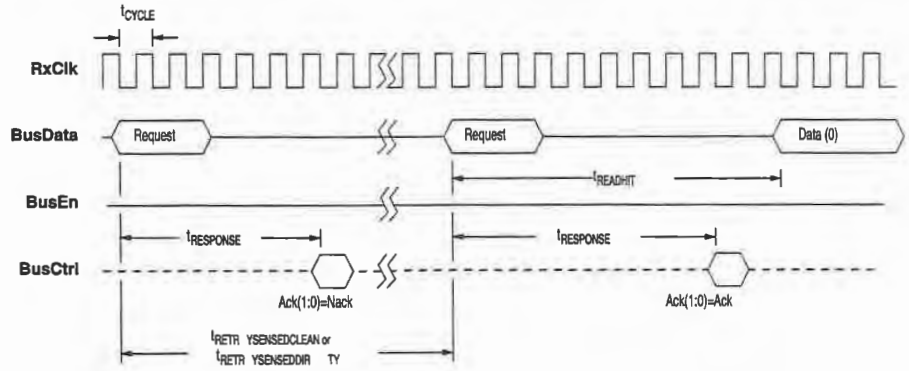


Write Hit Timing Diagram

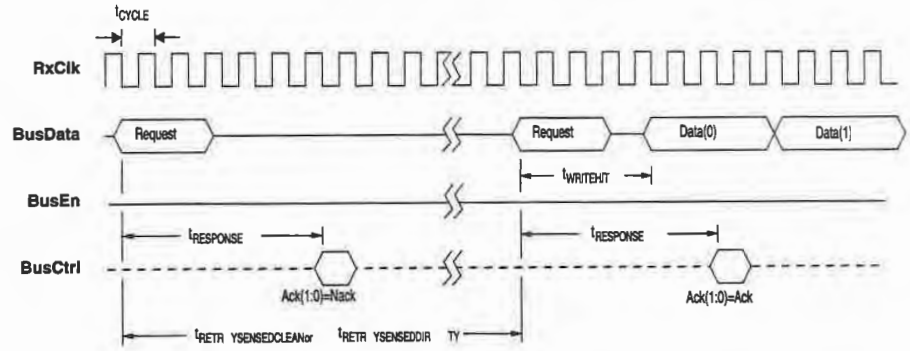


Samsung Ex. 1097, p. 833
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

Read Miss Timing Diagram

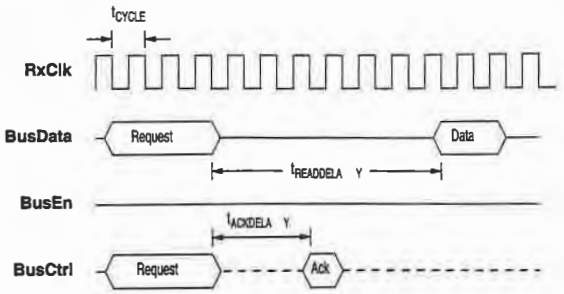


Write Miss Timing Diagram

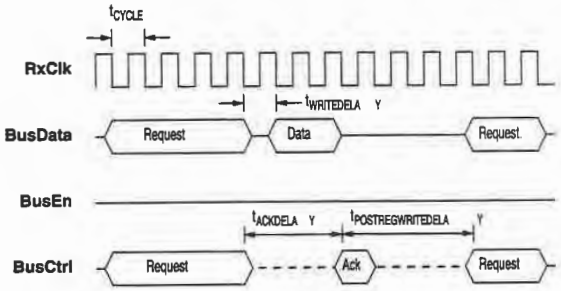


Samsung Ex. 1097, p. 834
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

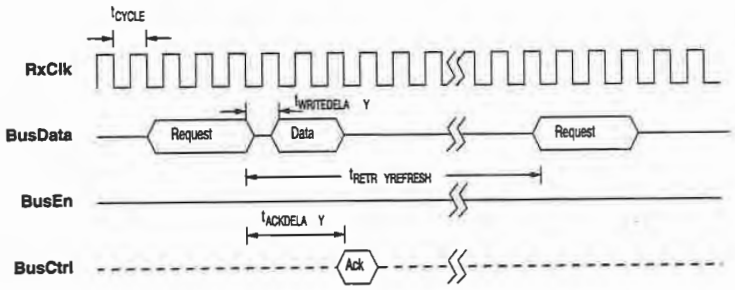
Register Read Timing Diagram



Register Write Timing Diagram

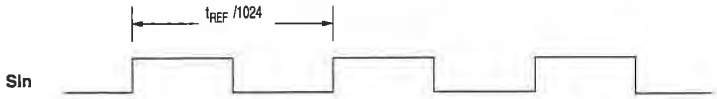


Manual Refresh Using Setrr

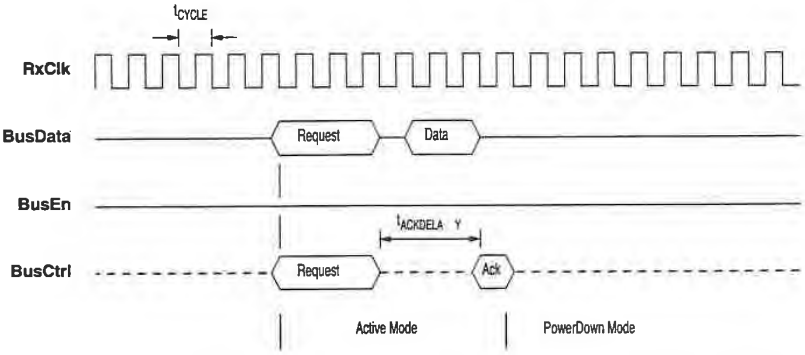


Samsung Ex. 1097, p. 835
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

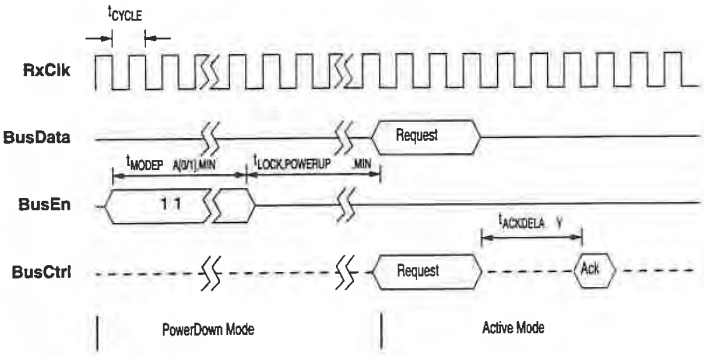
Refresh Timing for Powerdown Mode Using Sin, SOut



Power Down Timing (Register Write)



Power Up Timing



Samsung Ex. 1097, p. 836
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

Mechanical Drawings

The RDRAM is available in both horizontal and vertical surface mount plastic packages. Dimensions for the Horizontal surface mount plastic package are shown below.

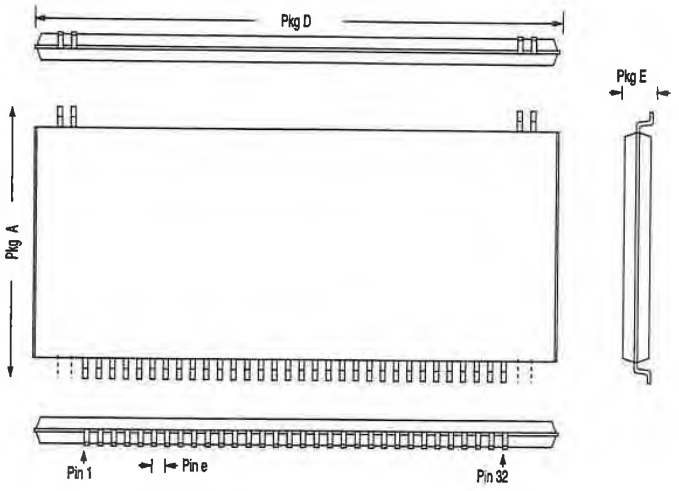


Figure 15: SHP-32 Package

The next figure shows the footprint of the SHP-32 package. Plane R-R is the electrical reference plane of the device on the center line of the SMT pads.

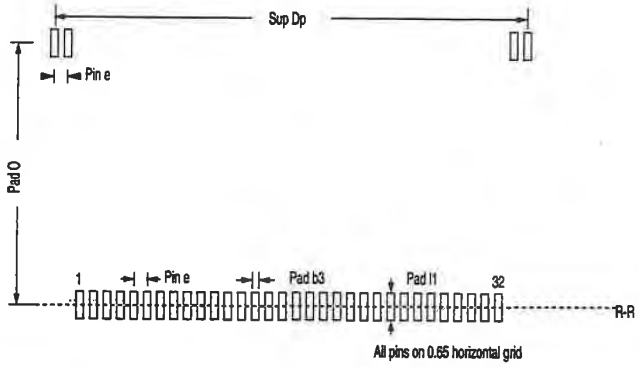


Figure 16: SHP-32 Footprint

Samsung Ex. 1097, p. 837
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

This table summarizes the values of the package and footprint dimensions[3].

Table 9: SHP-32 Package Dimensions

Symbol	Parameter	Min	Max	Unit
Pin e	Pin pitch	0.65	0.65	mm
Pkg D	Package width	24.9	25.3	mm
Pkg A	Package total height	12.9	13.1	mm
Pkg E	Package thickness	-	1.7	mm
Pad b3	SMT pad width	0.30	0.40	mm
Pad l1	SMT pad length	1.2	1.4	mm
Sup Dp	Support pad outer pitch	22.75	22.75	mm
Pad O	SMT pad offset	12.5	12.5	mm

The next figure summarizes the dimensions of the EIAJ standard SVP-32 package as used in the RDRAM. Refer to the EIAJ specifications for more details of the package dimensions and recommended footprint.

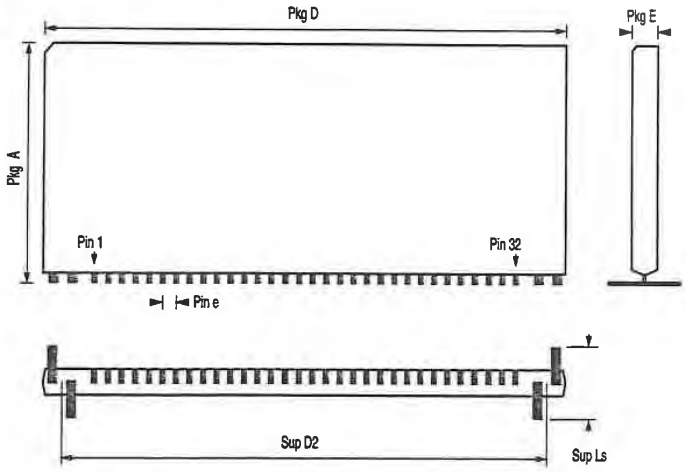


Figure 17: SVP-32 Package

3. All the support pad dimensions are provisional.

Samsung Ex. 1097, p. 838
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454

The figure below shows the footprint of the SVP-32 package. Plane R-R is on the center line of both the package and support leads, and offset from the center line of the SMT pads. Plane R-R is also the electrical reference plane of the device.

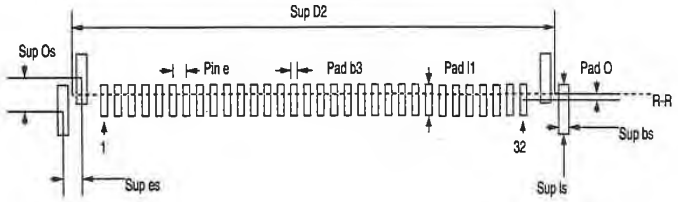


Figure 18: SVP-32 Footprint

This table summarizes the values of the package and footprint dimensions[4].

Table 10: SVP-32 Package Dimensions

Symbol	Parameter	Min	Max	Unit
Pin e	Pin pitch	0.65	0.65	mm
Pkg D	Package width	24.9	25.3	mm
Pkg A	Package total height	11.3	11.8	mm
Pkg E	Package thickness	1.2	1.4	mm
Sup Ls	Support lead span	3.4	3.6	mm
Sup D2	Support lead spacing	23.15	23.25	mm
Pad b3	SMT pad width	0.27	0.35	mm
Pad l1	SMT pad length	1.4	1.55	mm
Sup bs	Support pad width	0.45	0.55	mm
Sup ls	Support pad length	2.1	2.3	mm
Pad O	SMT pad offset	0.25	0.35	mm
Sup Os	Support pad offset	1.55	1.65	mm
Sup es	Support pad pitch	0.90	0.90	mm

4. All the support pad dimensions are provisional.

Samsung Ex. 1097, p. 839
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454



© Copyright June 1995 Rambus Inc. All rights reserved.

Rambus, RDRAM, RSocket, RTransceiver, RModule, and the Rambus Logo are trademarks of Rambus Inc.
Pentium is a registered trademark of Intel Corporation.

All other trademarks referenced are property of their respective owners.

DL0023-01

Rambus Inc.
2465 Latham Street
Mountain View, California USA
94040

Telephone: 415-903-3800
Fax: 415-965-1528

Samsung Ex. 1097, p. 840
Samsung Electronics Co., Ltd. v. Netlist, Inc., IPR2023-00454