Memory Performance Tutorial
Hot Chips 16

J Thomas Pawlowski
Sr. Dir. Architecture Development & Senior Fellow
NetCom Group, Micron Technology, Inc.

Agenda

- Objectives and background
- Soft Error Rate discussion
- Brief discussion of nonvolatile cells, devices, and device characteristics
- Volatile memory cell types, characteristics, general operation
- Specific device operation
  - Emphasis on factors for performance comparison
  - SRAM: DDR2, QDR II
  - New SRAM: DDR3, QDR3
  - Some background on SDRAM – refresh, address migration
  - DRAM: DDR, DDR2
  - New DRAM: GDDR3, FCRAM II, RLDRAM II, RAMBUS XDR
- Performance comparison of most memory types
  - Many operational scenarios
- Analysis and conclusions

Disclaimers

- All information presented herein is from sources not requiring an NDA. More details are available from the individual manufacturers.
- All road maps shown use estimated dates only, subject to change, consult individual manufacturers for updates.
- Many devices mentioned have trademarked names:
  - QDR™ and derivatives are trademarked by Cypress, et al
  - FCRAM™ and derivatives are trademarked by Fujitsu
  - RLDRAM™ and derivatives are trademarked by Infineon
  - RAMBUS™ and RAMBUS XDR™ trademarked by Rambus
- No statements made herein are to be taken as design advice
  - Regard them as generalizations to assist your understanding
  - If you have specific design issues, contact your friendly Micron support personnel
- Do not design any systems based on this information. Consult manufacturer datasheets.

Soft Error Rate - Definitions

- Hard error
  - An error induced by a device fault
  - Data is lost and can no longer be stored at that location
- Soft error
  - A random error induced by an event which corrupts data.
  - The device is not damaged and can correctly store data when written again.
- Single event upset - SEU
- Multiple event upset - MEU
- Soft error rate - SER
- Failure in time - FIT
  - 1 failure per 10^9 device operating hours

Causes of SER

- Radiation (α particles, B10 fission, cosmic rays)
- α particles (helium nuclei)
  - Short mean path length in silicon
    - 4 MeV - 18μm
    - 9 MeV - 70μm
  - Was virtually a non-issue on Micron devices due to material elimination
  - Typical SRAM SER due to alpha is 5 FITs/Mb at 1.8V at Micron
- B10 fission
  - Eliminate or shield against BPSG
- Cosmic rays
  - > 99 percent neutrons
  - Use ECC or else errors will be observed

Atmospheric Filtering and Secondary Effects

- Primary particles interact in atmosphere
  - Produce many generations of other particles
    - Neutrons, electrons, muons, protons (< 1 GeV)
    - Neutrons and protons have about 40x the impact in silicon
- At sea level, fairly equal mix of the 4 particles
- Neutron flux increases with altitude
  - F ~ A^N where A is ~1.25 and N is altitude in 1000s feet (e.g. N= 3 @ 3,000 feet)
    - ~2x at 3,000 feet, 3x at 5,000 feet, 9x at 10,000 feet

Atmospheric Filtering and Secondary Effects (continued)

- 90 percent of humans live at 1,700 feet elevation or lower
- 95 percent live at 4,000 feet or lower
  - Luckily, most equipment operates in lower neutron flux
- Neutron flux is also a function of longitude/ latitude
  - < 2x difference based on geographic location (see JESD89 spec)
  - Singapore is sea level and has lowest flux based on location

Nonvolatile Memory – NOR Flash

- NOR Flash memory cells
  - Common gate and drain, separate source, field effect transistors (FETs) with a floating gate.
  - Floating gate holds charge
  - When everything is biased correctly, will see the 1 or 0 caused by the floating gate, hence see 1 bit per cell.
  - Typically $8.5F^2$ ~ $15F^2$ cell size (F= process minimum feature size)

Nonvolatile Memory – NOR Flash (continued)

- NOR Flash memory multilevel cells
  - Same concept, but with tighter charge control and tighter sensing resolution, more overhead die area
  - 2 bits per cell, same cell size, equivalent result is $4.25F^2 \sim 7.5F^2$ cell size
  - Longer read and write times

V_T Distribution Example

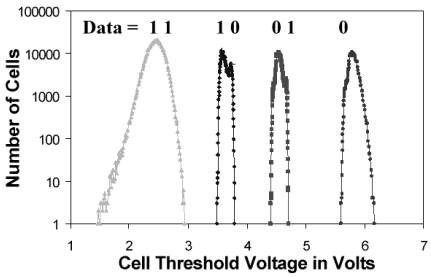

Nonvolatile Memory – NAND Flash

- NAND Flash memory cells
  - Common gate, daisy- chained source to drain
  - Floating gate holds charge as with NOR
  - Read or write current passes through chain of devices
  - Needs to operate on larger data chunks
  - Typically $4F^2$ cell size, but 1 extra device per 16, result is equivalent to $4.25F^2$ cell size

Nonvolatile Memory – NAND Flash (continued)

- NAND Flash memory multilevel cells
  - Same concept, but with tighter charge control and tighter sensing resolution, more overhead die area
  - 2 bits per cell, same cell size, equivalent result is 2.13F^2 cell size

Flash Memory Cell Comparison

	NAND	AND	NOR
Cell Array	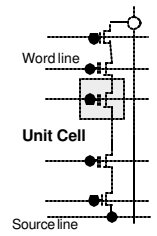	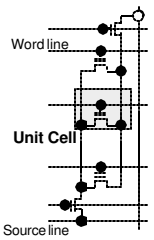	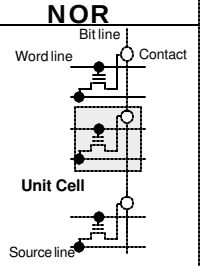
Layout	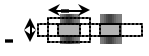	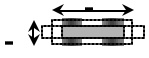	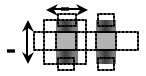
Cross-section	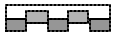	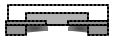	
Cell size	4F^2	8F^2	10F^2

? Cell size does not include overhead of extra "chain" device.
? All technologies capable of multiple bits per cell.

Basic NAND/ NOR Comparison

- NAND
  - Advantages:
    - Fast writes
    - Fast erases
  - Disadvantages:
    - Slow random access
    - Byte writes difficult
  - Applications:
    - File (disk) applications
    - Voice, data, video recorder
    - Any large sequential data
- NOR
  - Advantages:
    - Random access
    - Byte writes possible
  - Disadvantages:
    - Slow writes
    - Slow erase
  - Applications
    - Replacement of EPROM
    - Execute directly from nonvolatile memory

Generic Nonvolatile Memory Comparison (Today)

	NOR Flash	NAND Flash
Applications	Code, data	Mass storage
Future applications	MLC: mass storage	Code and data
Density range	Up to 512Kb	Up to 4Gb
READ latency	60ns- 120ns	25µs
Max Read bandwidth	41 MB/ s- 112 MB/ s (16b)	40 MB/ s (16b bus)
Max Write bandwidth	0.25 MB/ s	5MB/ s
Erase time	400ms (128KB blk)	2ms (128KB block)
Read device current	1.6x	1x
Write device current	3x	1x

Note that there is a wide variation among competing devices

Micron Flash Memory Comparison

Characteristic	NAND Flash MT29F2G08/16A	NOR (Q-Flash) MT28F128J3
Random Access Read	25µs (first byte) 50ns for remaining 2111	120ns
Sustained Read Speed (sector basis)	16 MB/ s (x8) or 32 MB/ s (x16)	20.5 MB/ s (x8) or 41 MB/ s (x16)
Random Write Speed	300µs/ 2112 bytes	180µs/ 32 bytes
Sustained Write Speed (sector basis)	1.5MB/ s	0.178MB/ s
Erase block size	128KB	128KB
Erase time per block (typ)	2ms	750ms

?NOR Flash memory is ideal for direct code execution (boot code)
?NAND Flash memory is ideal for file storage (e.g. data or image files. If code is stored, it must be shadowed to RAM first, as in a PC).

Volatile Memory Cell Types

- 1 Transistor, 1 Capacitor (1T-1C)
  Dynamic random access memory (DRAM) cell
  - Used on all production DRAM
  - Used on SRAM replacement devices
    - e.g. Micron CellularRAM™
- 3T DRAM cell
  - Not used much
  - Easy way to make DRAM in logic process
  - Some activity for high-performance SRAM replacement
- 6T static random access memory (SRAM) cell

SRAM Cell

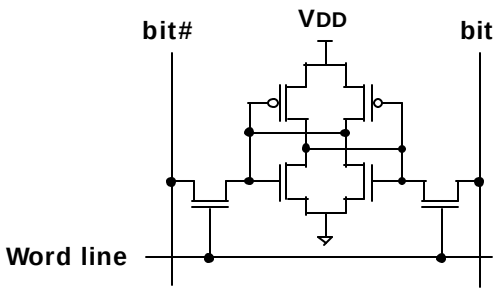

- Reads are not destructive
- Any lost charge is restored by the P-FET pull-up devices
- No refresh needed

- SRAM cell layout requires 3 horizontal lines (VDD, Vss, word line) 2 vertical lines (bit, bit#)

SRAM Cell Scaling

- Spacing between features increases on each process node
  - Required to deal with defects

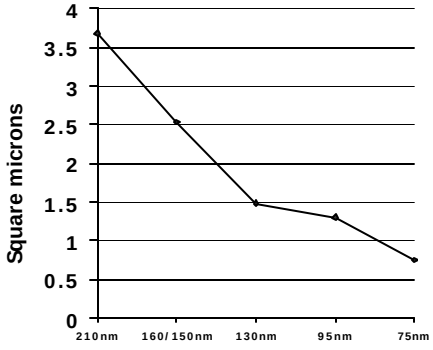

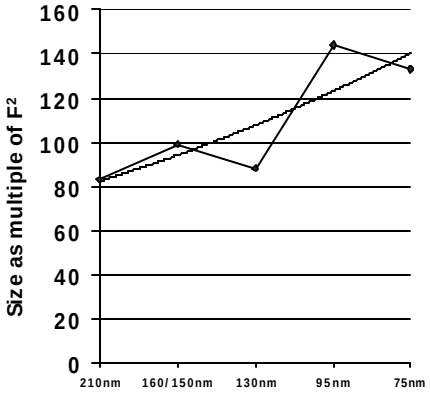

SRAM Die Size

- Depends on many factors, but…
- Can be calculated as n * A / eff
  - n = bit count, A = area of 1 bit, eff = die efficiency (e.g. 0.6 per unit die area is array bits, remainder is overhead)
- With all other factors remaining constant:
  - Die efficiency improves as density increases,
  - Worsens as I/O count increases,
  - Worsens as logic/ complexity increases
- Any memory die size can be similarly calculated

SRAM Soft Error Rate

- Typical SRAM SER = 10K FITS / Mb
  - (Sea level, Singapore, 0.13µm process, 1.8V core)
- Trend is that SER is growing with each new process node
  - Cell grows as function of minimum feature size
    - Target area to capacitance ratio increases
  - Voltage continues to reduce
    - Less cell charge
- Modeling suggests 20K- 50K FITs/ Mb on 90nm, 1.2V core
  - ~100K FITS for 36Mb SRAM on 90nm
- ECC is mandatory, or replace with DRAM

SRAM SER
18Mb 0.16µm QDR SER vs. Voltage

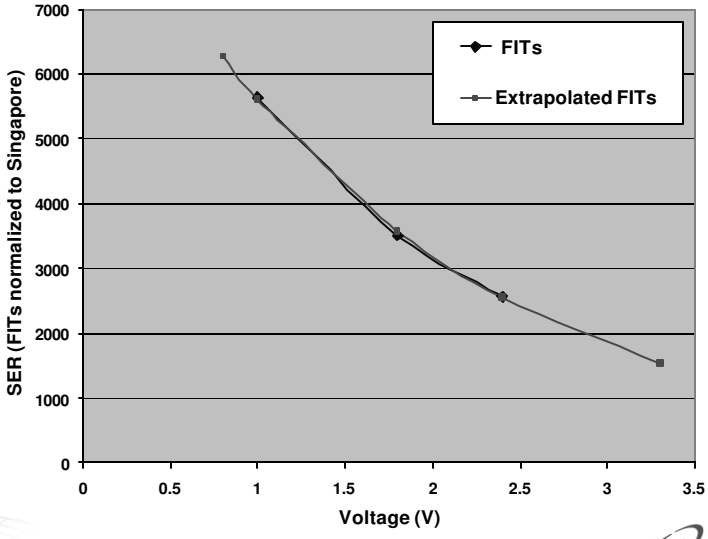

1T-1C DRAM

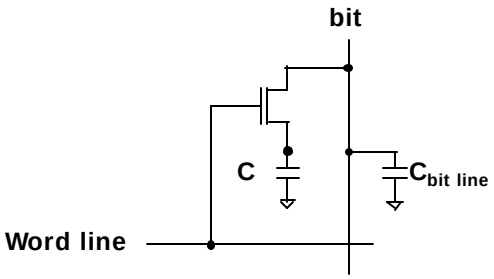

- Reads are destructive
- Bit value must be written back when read is done.
- C will discharge through leakage paths
- To restore charge perform read, elevate bit to full voltage and drive current into C

- 1T 1C DRAM cell layout requires
  - 1 horizontal line (Word line)
  - 1 vertical line (bit line)
  - Vss (or a reference node)

1T-1C DRAM Charge Redistribution

- During read, charge from C is shared with $C_{bit\ line}$
- If read begins with $V_{bit\ line} = V_{DD}/2$:
  - ? $V_{bit\ line} = (V_C - V_{DD}/2) * [C/(C + C_{bit\ line})]$
- ? $V_{bit\ line}$ is relatively small, e.g. 200mV
  - Need to amplify the voltage before sending the result off chip – sense amplifiers

Whole DRAM Array

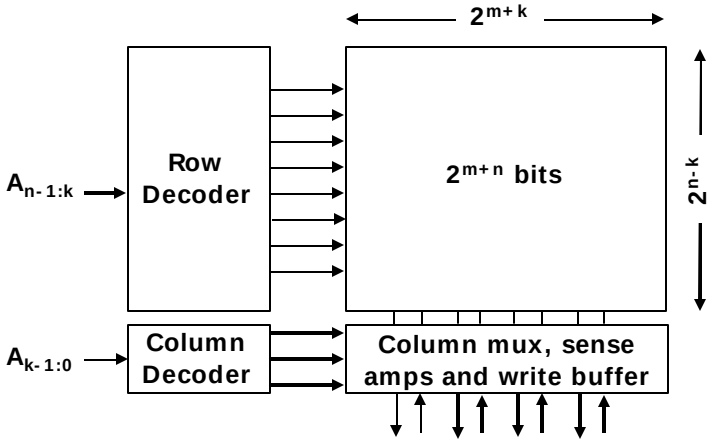

? Example: 16M words, 16 bit wide bus
$n = \log_2 16M = 24$, $k = 8$, $m = \log_2 16 = 4$

DRAM Cell Size

- Commercial DRAM production uses two different cell types
  - Both are 1T-1C but the capacitor is formed differently
- Trench capacitor is formed first
  - Like digging a trench but with a high wall surface area
  - Implants added on top to form N and P FETs, contacts, metal lines, etc.
  - Infineon DRAM, IBM eDRAM are examples: cell size is $8F^2$
- Stacked capacitor is formed with or after other chip elements
  - More vertical structure also with high wall surface area
  - Most other DRAM manufacturers do it this way, e.g. Micron, Samsung, Elpida
  - Cell size is $8F^2$ for all except new Micron DRAMs

Industry Roadmap
DRAM Cell Size Scaling

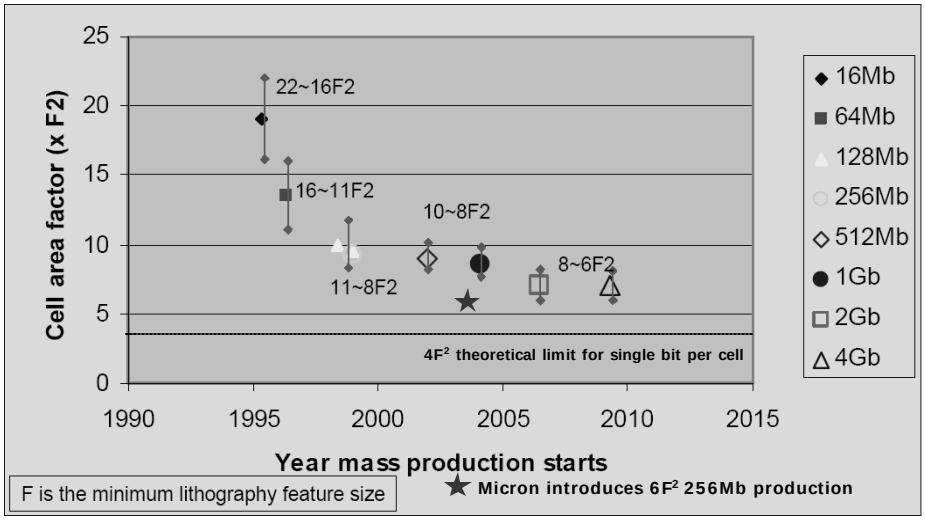

Note that there is a large difference between mass production starts versus "sweet spot" production.

Micron Jump to $6F^2$ Cell Ahead of the Curve

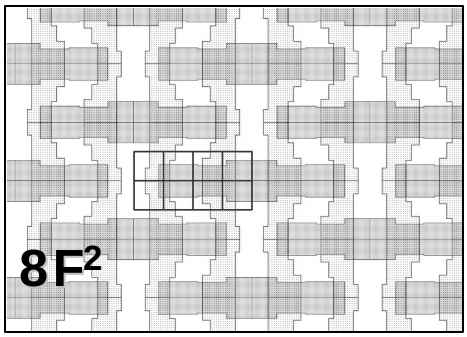

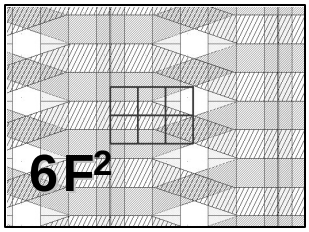

- Reduces memory array size by 25 percent
- Same cell capacitance
- Shortens bit lines and/or word lines
  - Less line capacitance means lower energy requirement, faster speed

DRAM Soft Error Rates

- 1T-1C DRAM has small cell target area and high cell capacitance
- Target area with each new process node is smaller
  - Area trend reduces SER with each new process
- Voltage is reduced every few process nodes
  - Voltage reduction trend increases SER due to reduction in stored charge
- Capacitance tends to stay about the same with each new process node
- Net result is reduction in SER with each new process node
- Typical SER is less than 1 FIT/Mb
  - e.g. 256Mb DRAM on 1.8V, 110nm process 100 FITs ~200 FITs
  - SER per component remains about the same even though density increases with time
  - Should expect lower SER on $6F^2$ than $8F^2$
  - No significant difference between trenched and stacked cell DRAMs for same cell size

DDR II and QDR II SRAM

- Motivation
  - Networking data structures
  - Improved bus efficiency
  - Improved pin bandwidth
  - Increased request rates over previous SRAMs

Networking SRAM Architectures

- Double data rate (DDR) I/O interface
- 2-word burst and 4-word burst
- Quad data rate: QDR and QDR II
  - Separate DIN and QOUT buses
  - Separate/concurrent read and write ports
  - 2 reads (1 data pair) and 2 writes (1 data pair) per clock cycle

Networking SRAM Architectures (continued)

- Double date rate: DDR and DDR II common I/O (CIO)
- Common DQ bus, common R/W bus
  - 1 request per clock cycle, stall for read-to-write transitions
- Double data date: DDR2 separate I/O (2-word burst only)
  - Separate DIN and QOUT buses
  - Separate/non-concurrent read and write ports
  - 1 request per clock cycle, no stalls
- All use HSTL I/O (1.4V - 1.9V)

QDR II R/W 2-Word Burst

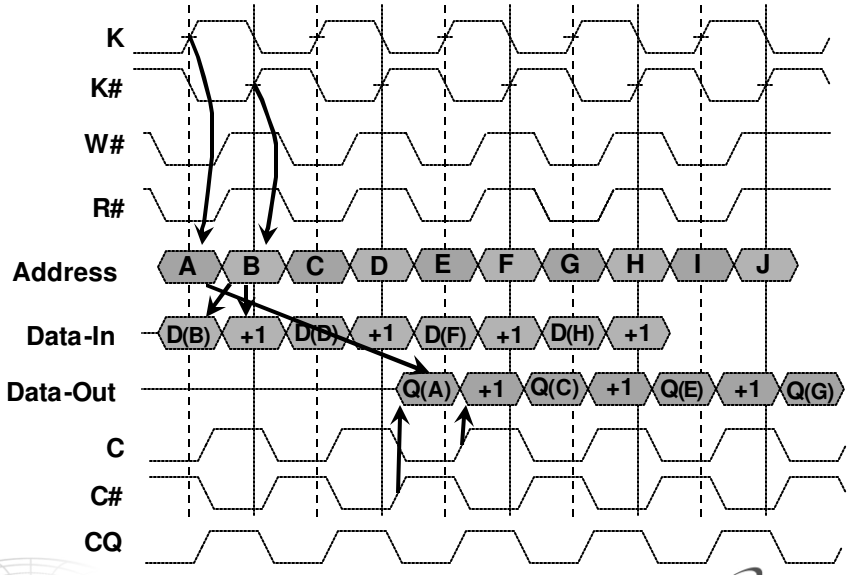

QDR II R/W 4-Word Burst

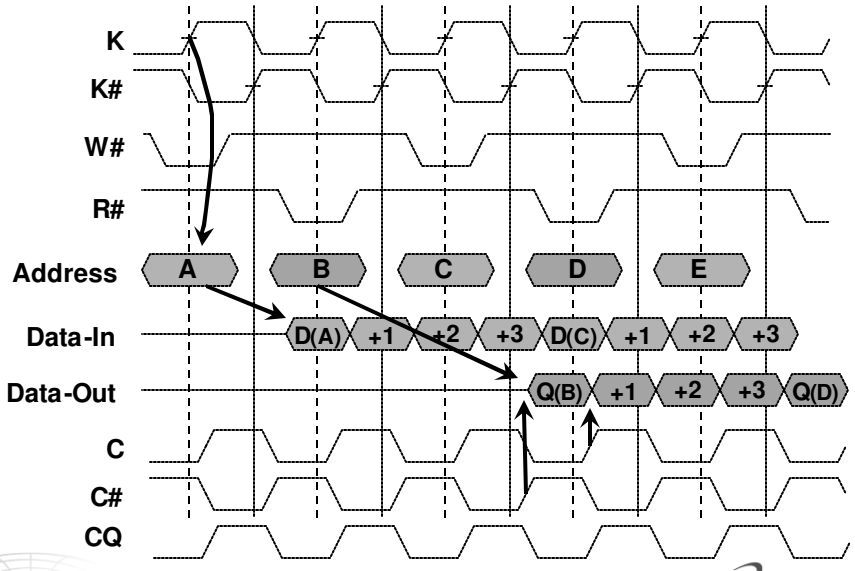

DDR II CIO R/W 4-Word Burst

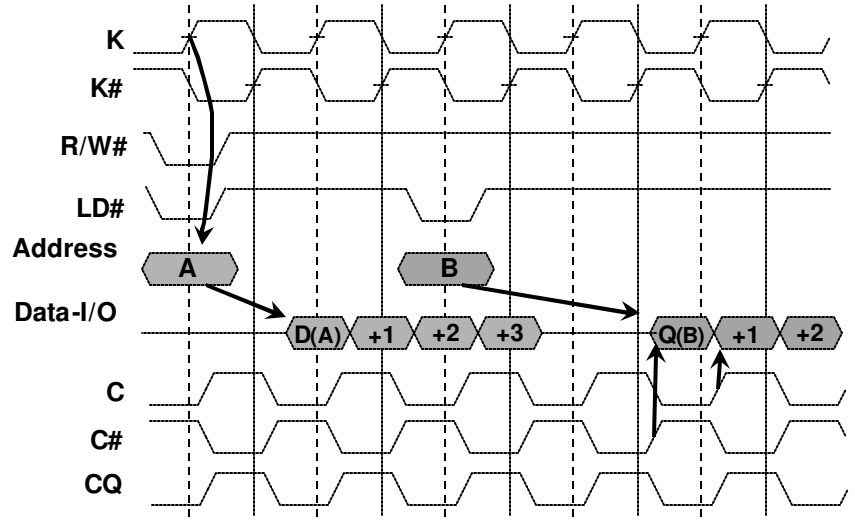

DDR II SIO R/W 2-Word Burst

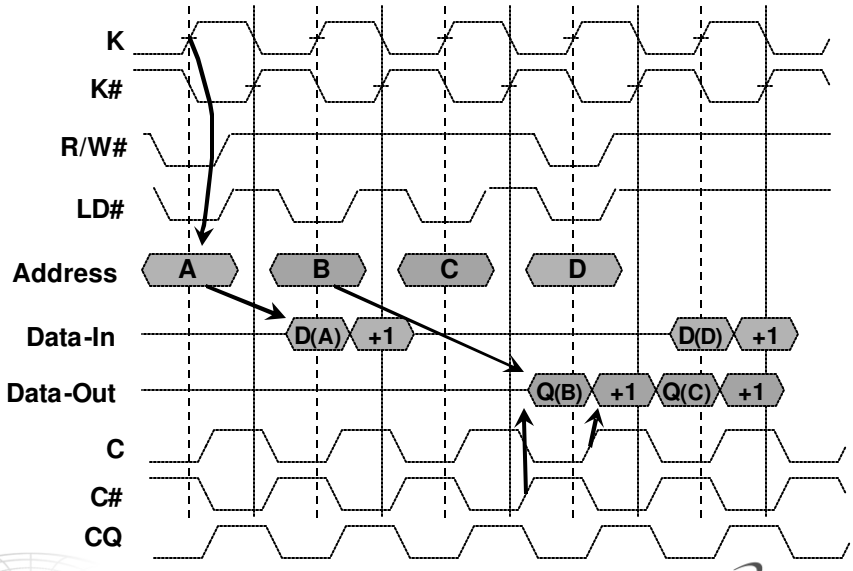

Relative Bandwidth per Pin of All Network SRAMs at 250 MHz Clock

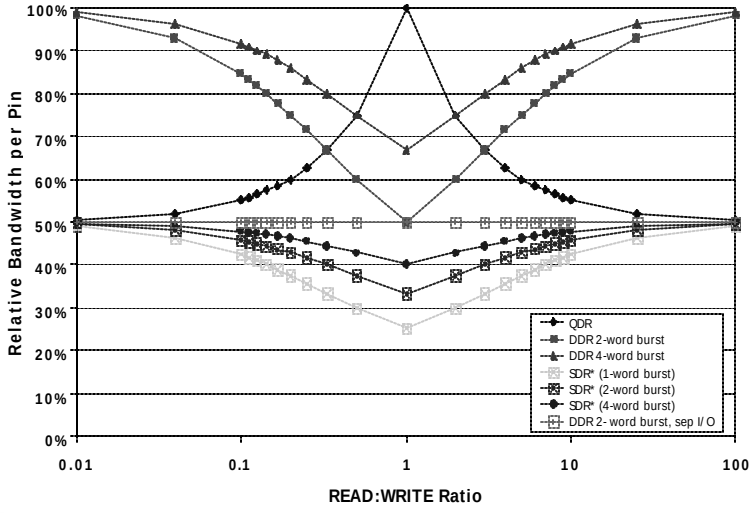

*Note: SDR represents bandwidth per pin for NoBL, ZeroSB, NtRAM, and ZBT devices

DDR II/ QDR II SRAM Performance Models

- DDR II CIO BL2 or BL4
  - R-R, W-W, W-R: BL/2 for next command
  - R-W: BL/2+1 cycle < 200 MHz, BL/2+2 >= 200 MHz
- DDR II SIO BL2
  - R-R, W-W, R-W, W-R: 1 cycle for next command
- QDR II BL2
  - Reads on K rising edges, writes on K falling edges, no restrictions
- QDR II BL4
  - R-W, W-R: 1 cycle
  - W-W, R-R: 2 cycles

Note: R = Read, W = Write, BL = Burst Length

No lower frequency limit

QDR III and DDR III

- Still being defined
- Target 250 MHz-500+ MHz clock frequency
  - QDR III 2-word burst
  - QDR III 4-word burst
  - DDR III 2-word burst
- 1.2V core

QDR III and DDR III (continued)

- 18Mb - 288Mb densities
- 9 bits, 18 bits, and 36 bits wide
- Other design details available from suppliers under NDA
  - My guesses: 4- cycle latency, VTT mid - rail termination scheme like RLDRAM II, increased ratio of echo clocks to outputs, increased ratio of input clocks to inputs
  - For today's performance analysis, latency guess is irrelevant; multiple identical addresses can be in flight simultaneously

QDR Comparison

Feature	QDR	QDRII	QDRIII
Frequency	BL2: 166 MHz BL4: 200 MHz	BL2: 250 MHz BL4: 333 MHz	BL2, BL4: 250- 500 MHz
Data Valid	1.4ns 166MHz	1.9ns @166MHz 0.98ns @ 333MHz	
Initial Latency	1.4 Cycles	1.6 Cycles	4? Cycles
Clocks	No Echo CLKs	Echo CLKs	Echo CLKs
Density	9/18/36Mb	18/36/72Mb+	36-144Mb+
Power Supply	2.5V	1.8V	1.2V

DDR III/ QDR III SRAM Performance Models

- DDR III CIO BL 2
  - R- R, W- W, W- R: BL/ 2 for next command
  - R- W: BL/ 2+ 2
- QDR III BL 2
  - Reads on K rising edges, Writes on K falling edges
  - Write address not = read address during same cycle
- QDR III BL4
  - R- W, W- R: 1 cycle
  - W- W, R- R: 2 cycles
- Note that this is speculative
  - Unknown if there are data dependencies in the final products

QDR / DDR SRAM Summary

- QDR SRAMs are optimized for systems with short-term, balanced READ and WRITE operations
- DDR CIO SRAMs are optimized for data streaming operations or READ/ WRITE unbalanced systems
- DDR SIO SRAMs are optimized for one address/ clock, 2-word burst systems
- High data availability, high cost, high SER
- Low density
- Extremely simple performance models

DRAM

- Some notes on single data rate synchronous DRAM
- DDR SDRAM (2.5V core)
- DDR2 SDRAM (1.8V core)
- RLDRAM II (1.8V core)
- FCRAM I, II and II+ (2.5V and 1.8V cores)
- GDDR3 (1.8V core)
- Rambus XDR (1.8V core)

Single Data Rate SDRAM

Density	Widths	Frequencies	Packages
64 Mb	4, 8, 16 32	133-183 MHz	54 TSOP 86 TSOP
128 Mb	4 8 16 32	133-166 MHz	54 TSOP 54 TSOP, 60 FBGA 54 TSOP 86 TSOP, 90 FBGA
256Mb	4 8 16 32	133-166 MHz	54 TSOP, 60 FBGA 54 TSOP, 60 FBGA 54 TSOP, 54 FBGA 86 TSOP, 90 FBGA
512 Mb	4 8 16	133-183 MHz	54 TSOP 54 TSOP 54 TSOP

SDRAM - Refresh

- The refresh rate is dependent on the number of rows in the device
  - Each refresh command refreshes a single row
  - Addressing is handled by an internal refresh controller; address bits are a "don't care" during an auto- refresh command
  - ^{t}REF represents the maximum time in which all rows much be refreshed at least once – typically 64ms for SDR and DDR SDRAM
  - ^{t}REFI represents the average periodic refresh time – ^{t}REF/ #of rows (8K refresh: 64ms/ 8K = 7.8125µs, 4K refresh: 15.625µs)
  - ^{t}REFC represents the absolute maximum time you can go without issuing any refresh commands – Micron DDR devices allow 9 REFI cycles (9 x ^{t}REFI), JEDEC standard is 8 REFI cycles
  - ^{t}RFC represents the auto refresh command period. (CKE must remain HIGH during this time for all MICRON devices)

2.5V DDR SDRAM

- Motivation
  - Increase bandwidth per data pin without significant cost increase
  - 266 MB/ s – 400 Mb/ s data rates
- Inspired by original DDR SDRAM

Double Data Rate SDRAM

- Micron 2.5V devices available
  - 128Mb – 1Gb
  - 4-bit, 8-bit and 16-bit wide data bus
  - 66-pin TSOP on all devices, 60-ball FBGA on some devices
  - 133 MHz - 200 MHz clock, 266 Mb/s - 400 Mb/s per data pin
    - Check specific devices for frequency availability
    - Micron leads all suppliers in product breadth

DDR SDRAM Block Diagram

4 x 16 Meg x 8

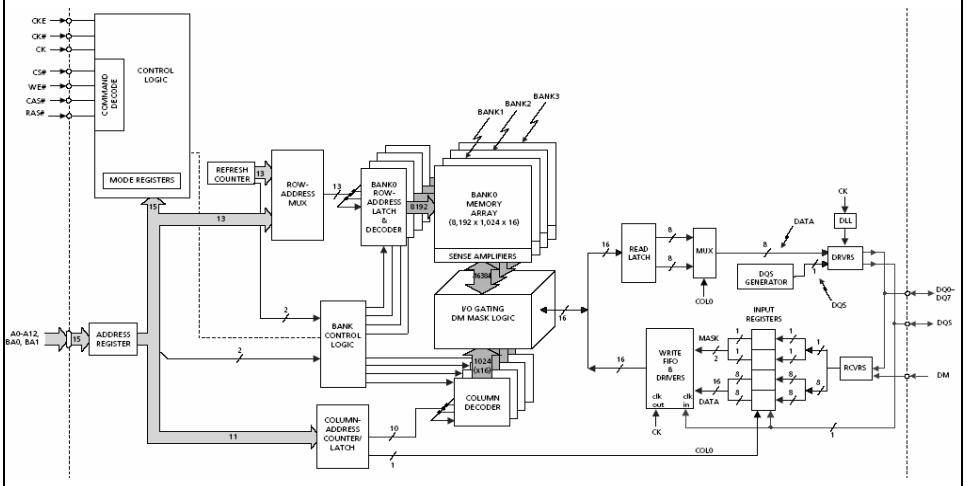

Basic DDR SDRAM Commands

NAME (FUNCTION)	CS#	RAS#	CAS#	WE#	ADDR	NOTES
DESELECT (NOP)	H	X	X	X	X	9
NO OPERATION (NOP)	L	H	H	H	X	9
ACTIVE (Select bank and activate row)	L	L	H	H	Bank/Row	3
READ (Select bank and column, and start READ burst)	L	H	L	H	Bank/Col	4
WRITE (Select bank and column, and start WRITE burst)	L	H	L	L	Bank/Col	4
BURST TERMINATE	L	H	H	L	X	8
PRECHARGE (Deactivate row in bank or banks)	L	L	H	L	Code	5
AUTO REFRESH or SELF REFRESH (Enter self refresh mode)	L	L	L	H	X	6, 7
LOAD MODE REGISTER	L	L	L	L	Op-Code	2

NOTE:

1. CKE is HIGH for all commands shown except SELF REFRESH.
2. BA0-BA1 select either the mode register or the extended mode register (BA0 = 0, BA1 = 0 select the mode register; BA0 = 1, BA1 = 0 select extended mode register; other combinations of BA0-BA1 are reserved). A0-A12 provide the op-code to be written to the selected mode register.
3. BA0-BA1 provide bank address and A0-A12 provide row address.
4. BA0-BA1 provide bank address; A0-A*i* provide column address, (where *i*=9 for x16, *i*=9,11 for x8, and *i*=9,11,12 for x4) A10 HIGH enables the auto precharge feature (non persistent), and A10 LOW disables the auto precharge feature.
5. A10 LOW: BA0-BA1 determine which bank is precharged.
   A10 HIGH: all banks are precharged and BA0-BA1 are "Don't Care."
6. This command is AUTO REFRESH if CKE is HIGH, SELF REFRESH if CKE is LOW.
7. Internal refresh counter controls row addressing; for within the Self Refresh mode all inputs and I/Os are "Don't Care" except for CKE.
8. Applies only to read bursts with auto precharge disabled; this command is undefined (and should not be used) for read bursts with auto precharge enabled and for write bursts.
9. DESELECT and NOP are functionally interchangeable.

DDR: Read with Autoprecharge

- ^{t}RCD
- ^{t}RAS
- ^{t}RC
- ^{t}RAP
- ^{t}RP

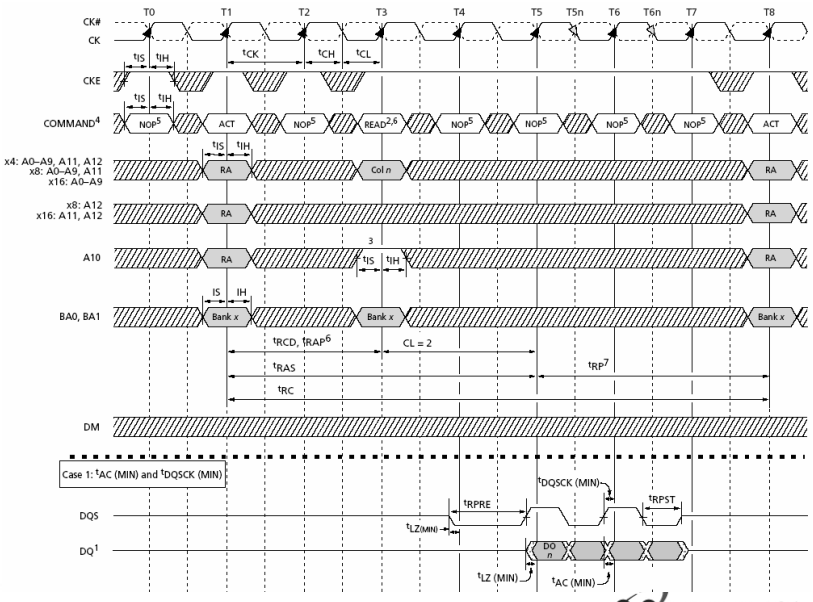

DDR Data Output Timing

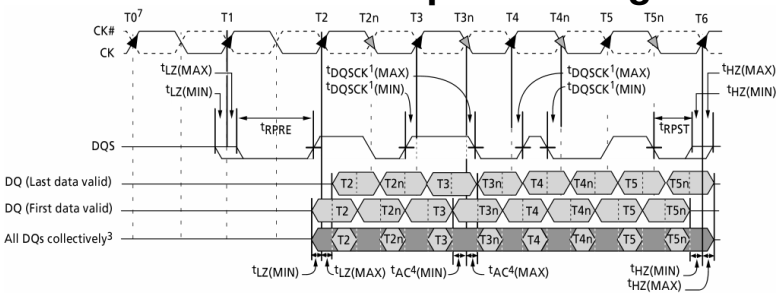

- tDQSCK is the DQS output window relative to CK.
- DQs transitioning after DQS transition define tDQSQ window.
- All DQs transition by tDQSQ after DQS edge, regardless of tAC.
- tAC is the DQ output window relative to CK.
- tLZ (MIN), tAC (MIN) and tHZ (MIN) are the first valid signal transition.
- tLZ (MAX), tAC (MAX) and tHZ (MAX) are the latest valid signal transition.
- The DLL is used to realign the data strobe to the CK/ CK# crossing.
- Data is edge aligned to the DRAM clock

DDR: DQS Preamble and Postamble

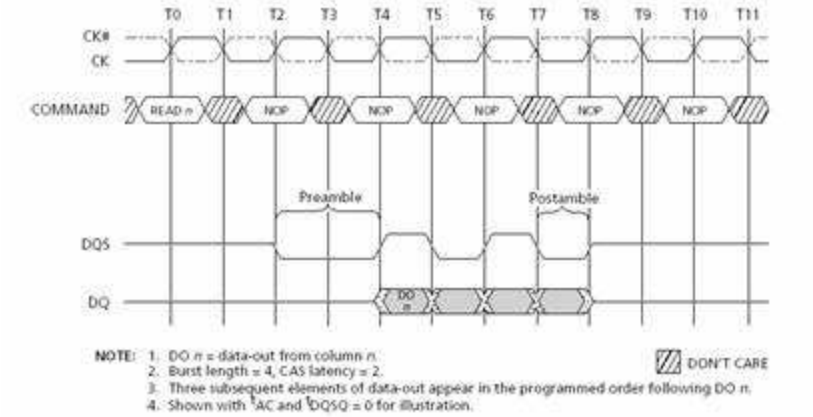

- Preamble (tRPRE) provides a timing window for the receiving device to enable its data capture circuitry while a known level is valid on the strobe signal
  - avoids false triggers of the capture circuit
- Postamble (tRPST) represents the DQS LOW time following the last transition

DDR: Read Data Valid Window

- ^{t}DQSQ represents DQ to DQ skew in relation to data strobe
- ^{t}DQSQ is derived at each DQS clock edge and is not cumulative over time
- ^{t}HP is the lesser of ^{t}CL or ^{t}CH clock transition
- The data valid window is derived for each DQS transitions and is defined as ^{t}QH minus ^{t}DQSQ.

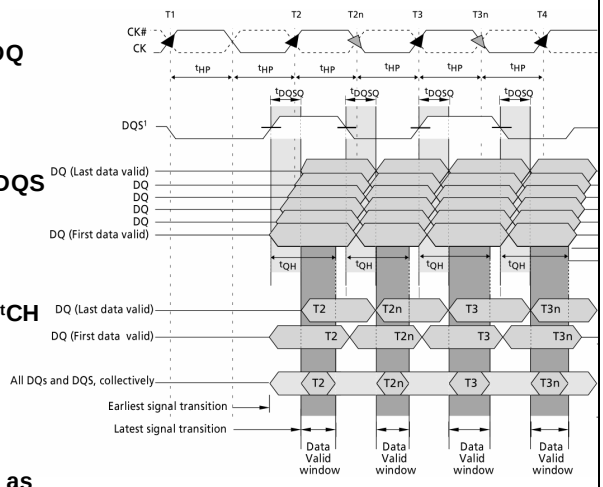

DDR: Write with Autoprecharge

- ^{t}RCD = Active to read or write delay
- ^{t}WR = Write recovery
- ^{t}RP = Precharge time
- ^{t}RAS = Active to precharge
- ^{t}RC = Active to active in same bank

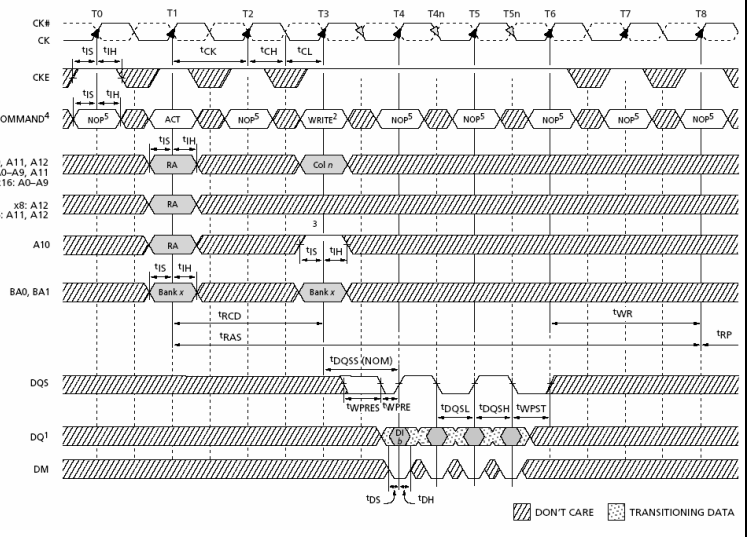

DDR: Data Input Timing

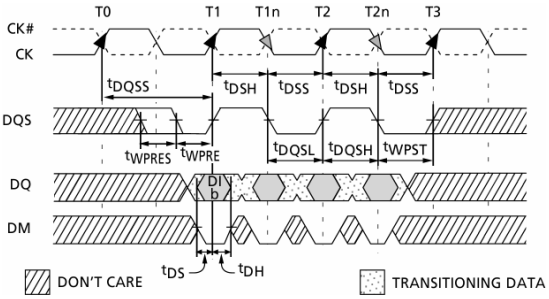

- tDQSS = Write command to first DQS latching transition
- tDSH (MIN) generally occurs during tDQSS (MIN).
- tDSS (MIN) generally occurs during tDQSS (MAX).
- WRITE command issued at T0.

DDR: Read Interrupted by Write

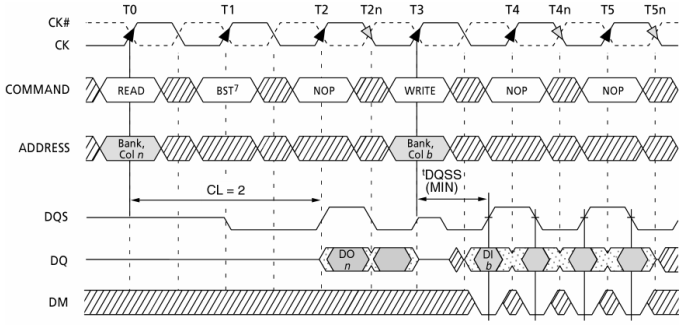

- Burst length = 4 in the cases shown
  - Applies for bursts of 8 and full page as well
  - If the burst length is 2, the BST command shown can be NOP.
- Shown with nominal tAC, tDQSCK, and tDQSQ.
- tRAS (MIN) still applies

DDR: Write Followed by Read

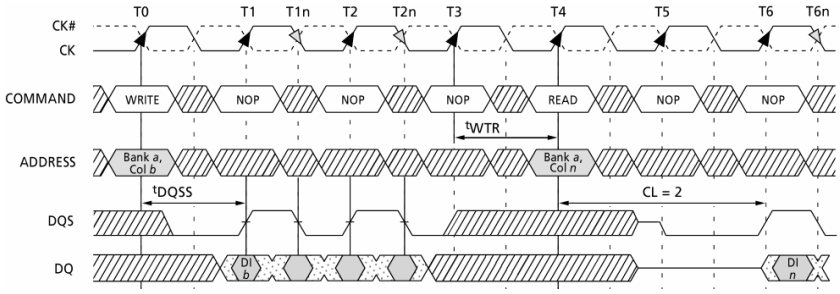

- READ command can be earlier if to a different device
- If the READ command applies to a different row in the same bank the READ command would come after ^{t}WR, ^{t}RP, and ^{t}RCD

DDR SDRAM Performance Model

Clock Frequency		200 MHz		166 MHz		133 MHz	
Clock period	T	5 ns, 7.5 max		6 ns, 13 max		7.5 ns, 13 max	
Cas Latency	CL	3		2.5		2	
		ns	T	ns	T	ns	T
ACT to same bk ACT	tRC	55	~11	60	~10	65	~9
ACT to R or W	tRCD	15	~3	15	3	20	~3
ACT to PRE	tRAS	40	~8	42	~7	40	~6
PRE Period	tRP	15	3	15	3	20	~3
W to PRE	tWR	15	3	15	3	15	2
W to R	tWTR		1		1		1
ACT to diff bk ACT	tRRD	10	2	12	2	15	2
Auto RFSH period	tRFC	70	~14	72	~12	75	~10

- 7.8125µs refresh period used (512Mb device)
- Other CAS latencies are possible but aren't used in performance model
- Some calculations require CL rounded up to next integer
- 4 banks
- T represents 1 clock cycle
- Clock period sweet spots: 5 ns, 6 ns, 7.5 ns
- Recall the disclaimer!

DDR SDRAM Performance Model

Command	To Same Bank	To Different Bank
W to R	1 + BL/2 + TWTR	BL/2
W to W	BL/2	BL/2
W to PRE	1 + BL/2 + TWR	1
W to ACT	W to PRE + TRP	1
R to R	BL/2	BL/2
R to W	CL(rounded up) + BL/2	CL(rounded up) + BL/2
R to PRE	BL/2	1
R to ACT	BL/2 + TRP	1
ACT to ACT	TRC	TRRD
ACT to R or W	TRCD	1
ACT to PRE	TRAS	1

? **Recall the disclaimer: do not base controller design on this information, consult manufacturer data sheets for latest and most accurate information**

1.8V DDR2 SDRAM

- **Motivation**
  - **Increase frequency of DDR SDRAM without significant cost increase**
  - **400 Mb/s – 800 Mb/s data rates**

1.8V DDR2 SDRAM Availability

- Micron devices available
  - 256Mb – 2Gb
  - 4-bit, 8-bit and 16-bit wide data bus
- 256Mb: 4, 8 bit in 60 ball FBGA, x16 in 80FBGA
- 512Mb – 2Gb in 92 ball FBGA
- 200 MHz- 333 MHz clock, 400 Mb/ s- 667 Mb/ s per data pin
  - Check specific devices for frequency availability
  - Micron leads all suppliers in product breadth

Basic Changes From DDR to DDR2

- Increased frequency spawns many changes
  - Lower core and I/ O voltages
  - On-die termination (ODT)
  - Off-chip driver characteristics and calibration
  - 4n prefetch (instead of 2n, sets minimum burst length)
    - To keep column performance (cost) in check
    - Result is elimination of 2- word burst
  - Additive CAS latency
    - Since T_{RR} is larger, needed mechanism to clean up request issuing
  - Bank count increase, made possible by density, needed for performance

Feature Overview: DDR vs. DDR2

Feature/ Option	DDR	DDR2	DDR2 Advantage
Package	TSOP (66-pin)	FBGA only	Better electrical performance and speed
Voltage	2.5V 2.5V I/O	1.8V 1.8V I/O	Reduces memory system power demand
Densities	64Mb –1Gb	256Mb – 4Gb	High- density components enable large memory subsystems
Internal banks	4	4 and 8	1Gb and higher will have 8 banks for better performance
Pre-fetch (MIN write burst)	2	4	Provides reduced core speed dependency for better yields
Speed (data pin)	200, 266, 333, 400 Mb/s	400, 533, 667 Mb/s	Migration to higher speed I/O
Read Latency (CAS latency)	CL 2, 2.5, 3 CLK	CL + AL CL = 3, 4, 5	Eliminating one half clock settings helps speed internal DRAM logic and improves yields
Additive Latency (Posted CAS)	N/A	AL options 0,1,2,3,4	Mainly used in server applications to improve command bus efficiency

Feature Overview (Continued)

Feature/ Option	DDR	DDR2	DDR2 Advantage
WRITE Latency	1 clock	READ Latency - 1	Improves bus efficiencies
Termination	Motherboard parallel to V_{TT}	DRAM on-die termination (ODT) optional on motherboard	ODT for both memory and controller improves signaling, and reduces system cost
Burst Lengths	2,4,8	4,8	
Data Strobes	Single ended	Differential or single ended	Improved system timing margin by reduced strobe crosstalk
Modules	184-pin unbuffered registered 200-pin SODIMM 172-pin MicroDIMM	240-pin unbuffered registered 200-pin SODIMM 244-pin MiniDIMM 214-pin MicroDIMM	Modules are the same length, with added pins

DDR- DDR2 Differences: Page Size

- Page size is the minimum number of columns accessed with a single ACTIVATE command = # columns x bus width

Density		DDR	DDR2
256Mb	Page Size	1KB	1KB
	Banks	4	4
512Mb	Page Size	2KB	1KB (x4, x8), 2KB (x16)
	Banks	4	4
1Gb	Page Size	2KB	1KB (x4, x8), 2KB (x16)
	Banks	4	8
2Gb	Page Size	-	1KB (x4, x8), 2KB (x16)
	Banks	-	8

DDR2 Differences: Activate Commands

- Activate for 4-bank DDR and DDR2

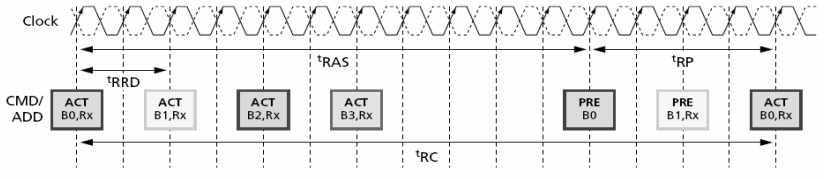

- Activate for 8-bank DDR2

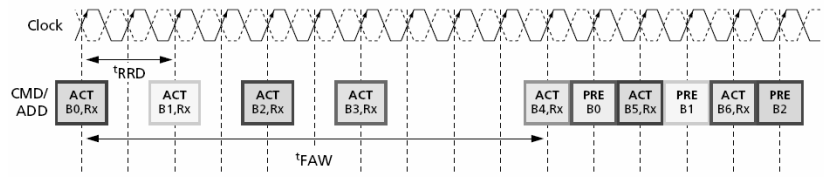

Note: tRRD is 7.5ns on 1KB page size and 10ns on 2KB page size (two clock minimum).
tFAW is 37.5ns on 1KB page size and 50ns on 2KB page size.

DDR2 8-Bank Restrictions

- Bank ACTIVE restrictions
  - No more than 4 banks may be activated in a rolling 4 * ^{t}RRD + 2 * ^{t}CK period
  - ^{t}RRD is now based on page size
    - 2KB page size = 10ns (x16 configuration on 1Gb and 2Gb)
    - 1KB page size = 7.5ns (x4, x8 configuration on 1Gb and 2Gb)
    - ^{t}RRD has 2 * ^{t}CK (MIN) at any ^{t}CK
- Bank PRECHARGE restrictions
  - Precharge(ALL) command timing equals ^{t}RP+ 1 * ^{t}CK
    - Single-bank PRECHARGE = ^{t}RP
    - 8-bank PRECHARGE(ALL) = ^{t}RP + 1 * ^{t}CK

DDR2 Read – No Additive Latency

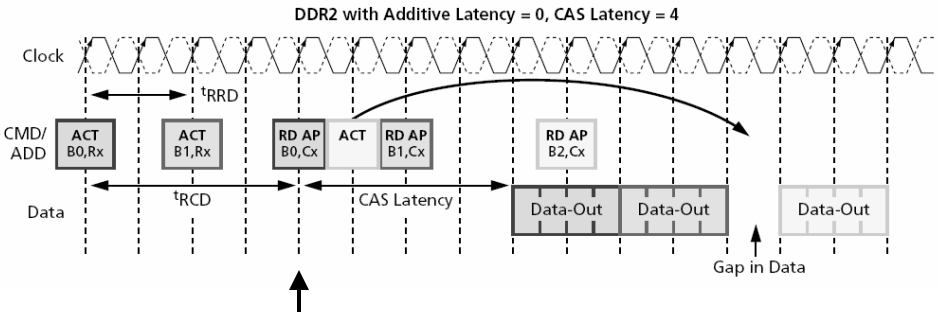

? Desire to insert ACT here to keep data bus utilized

DDR2 With Additive Latency

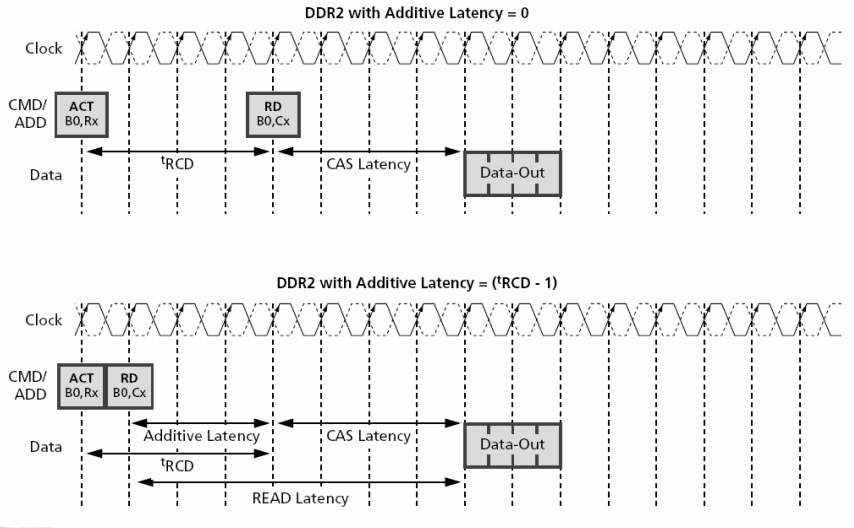

Result: Improved Bus Utilization

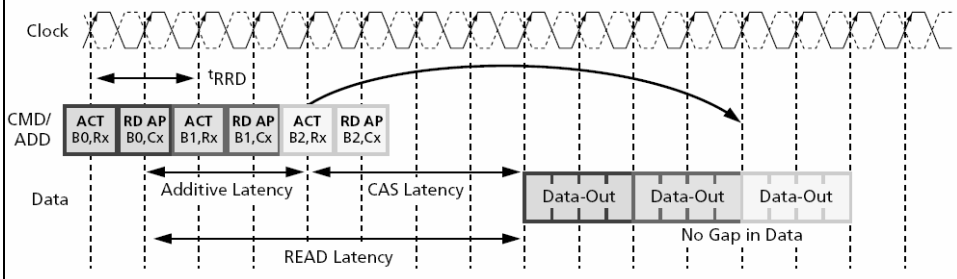

- Allows commands to be placed without conflict
- No gap in data
- Increased overall latency is a drawback in some applications

Latencies

- Only whole clock CAS latencies
- Blue indicates primary speed grades

Speed Bin	DDR2-667		DDR2-533	DDR2-400	Units
	4-4-4	5-5-5	4-4-4	3-3-3	
Parameter	MIN	MIN	MIN	MIN	
CAS Latency	4	5	4	3	^{t}CK
	12	15	15	15	ns
^{t}RCD	12	15	15	15	ns
^{t}RP	12	15	15	15	ns

? **3-3-3 terminology means:**

3-cycle CL (CAS latency)

3-cycle ^{t}RCD (ACTIVE to READ or WRITE delay, i.e. row time)

3-cycle ^{t}RP (precharge command period, i.e. row closing time)

- **DDR2-800 will be 5-5-5, 12.5ns**

DDR2 READ and WRITE Latencies

Figure 11: READ Latency

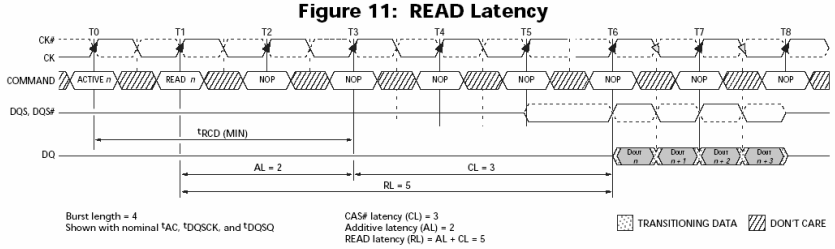

Figure 12: Write Latency

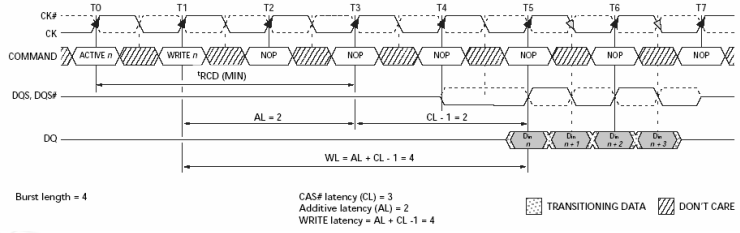

DDR2 Read - Precharge

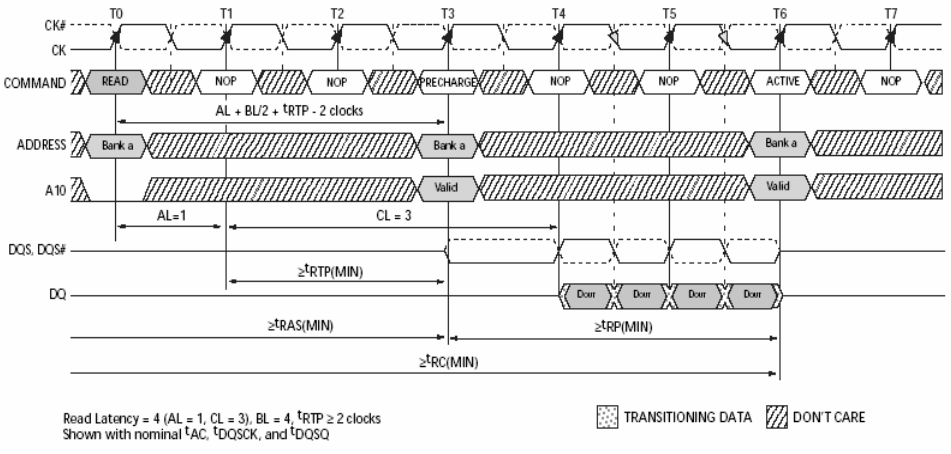

DDR2 Write to Precharge

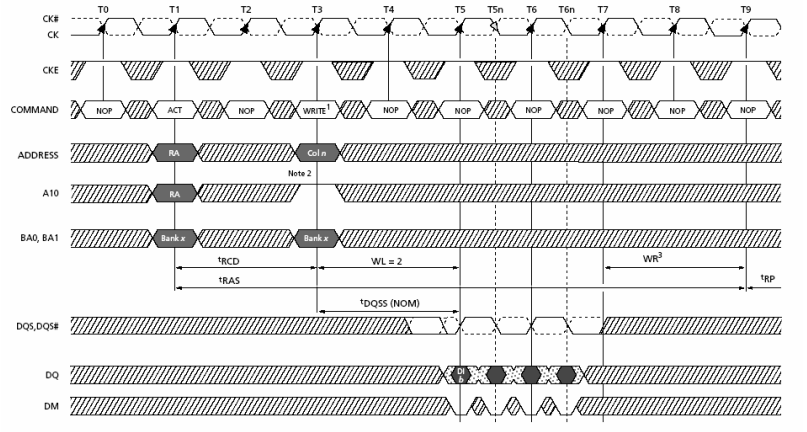

NOTE: 1. Burst length = 4, additive latency = 0, and WRITE latency = 2.
2. Enable auto precharge.
3. WR is programmed via EMR [11,10,9] and is calculated by dividing tWR (ns) by tCK and rounding up to the next integer value.

DDR2 Read to Write

Figure 24: READ to WRITE

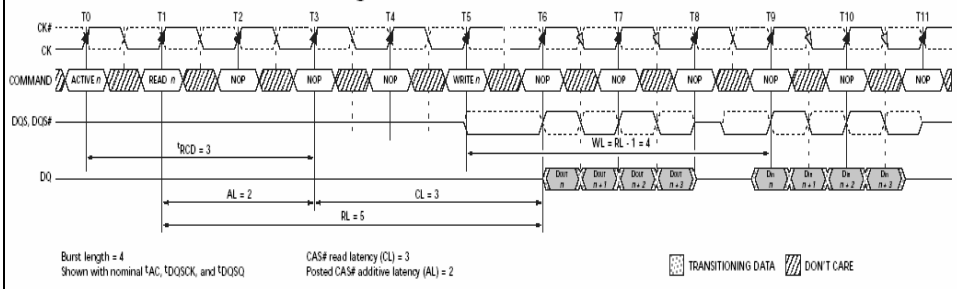

DDR2 Refresh

Density	256Mb	512Mb	1Gb	2Gb	4Gb	Units
Refresh count (x4x8, x16)	8K,8K	16K,8K	16K,8K	32K,16K	TBD	
Refresh cycle	64	64	64	64	64	ms
Refresh interval	7.8	7.8	7.8	7.8	7.8	µs
^{t}RFC	75	105	127.5	197.5	TBD	ns

- External refresh interval is maintained at 8K refresh per 64ms cycle
- DRAM is managing more internal row/ bank refresh than external commands given
  - ^{t}RFC time is increased for each density

On-Die Termination (ODT)

- ODT is a new termination scheme for DDR2 in which the controller and the DRAM have internal termination for the DQ, DQS/ DQS# and DM signals
- The DRAM termination is turned on and off by the controller depending on system loading and READ or WRITE operations
- ODT allows the system bus to achieve improved signal integrity

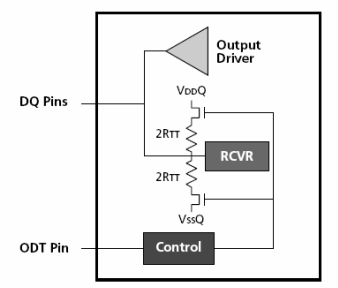

Off-Chip Driver (OCD)

- OCD calibration
  - Goal is 18 ohms ± 1.5 ohms
  - A new mode that allows the controller to measure and adjust the output driver strength of the DRAM
  - System guarantees memory channel timing and margin
  - Must recalibrate often for temperature variations
- Using OCD default, DRAM parametrics are guaranteed over full process, voltage, and temperature

DDR2 SDRAM Performance Model

Clock Frequency		333 MHz		266 MHz		200 MHz	
Clock period	T	3 ns, 8 max		3.75 ns, 8 max		5 ns, 8 max	
Cas Latency	CL	5		4		3	
		ns	T	ns	T	ns	T
ACT to same bk ACT	tRC	54	18	55	15	55	11
ACT to R or W	tRCD	15	5	15	4	15	3
ACT to PRE	tRAS	39	13	40	11	40	8
PRE Period	tRP	15	5	15	4	15	3
W to PRE	tWR	15	5	15	4	15	3
W to R	tWTR	7.5	3	7.5	2	10	2
ACT to diff bk ACT	tRRD 1KB pg	7.5	3	7.5	2	7.5	2
ACT to diff bk ACT	tRRD 2KB pg	10	4	10	3	10	2
Auto RFSH period	tRFC	105	35	105	28	105	21
Internal R to PRE	tRTP	7.5	3	7.5	2	7.5	2

- Refresh every 7.8125µs.
- Other CAS latencies are possible but not used in the performance model
- T represents 1 clock cycle, all values are shown for fast corner, must calculate
- Clock period sweet spots: 3 ns, 3.75 ns, 5 ns
- Recall the disclaimer!

DDR2 SDRAM Performance Model

Command	To Same Bank	To Different Bank
W to R	(CL- 1) + BL/ 2 + TWTR	BL/ 2
W to W	BL/ 2	BL/ 2
W to PRE	(CL- 1) + BL/ 2 + TWR	1
W to ACT	W to PRE + TRP	1
R to R	BL/ 2	BL/ 2
R to W	BL/ 2 + 2	BL/ 2 + 2
R to PRE	AL + BL/ 2 + TRTP - 2	1
R to ACT	R to PRE + TRP	1
ACT to ACT	TRC	TRRD
ACT to R or W	TRCD - AL	1
ACT to PRE	TRAS	1

- For optimal operation AL = TRCD – 1. Read latency = AL + CL.
- Write latency = read latency – 1. BL = 4 or 8.
- Auto precharge not used in current model.
- Recall the disclaimer.

BREAK

You deserve it!

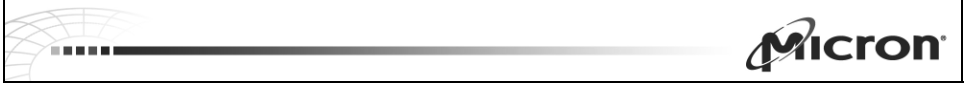

Motivation for Faster ^{t}RC Memories

- SRAM/ DRAM size ratio growth makes future SRAM too expensive per bit
- Personal Computer industry is not demanding improvements
- Most companies not addressing the need for faster ^{t}RC, YET!

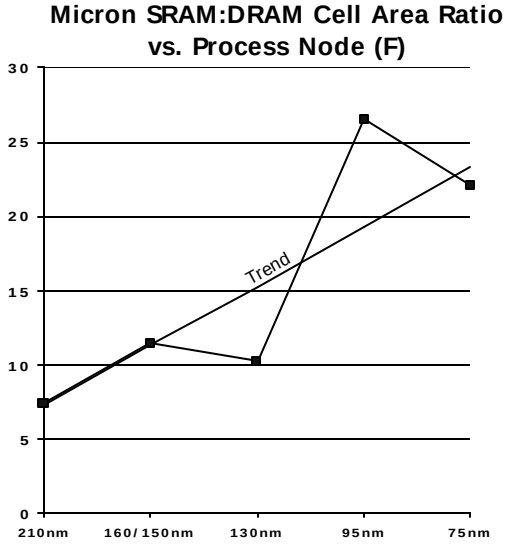

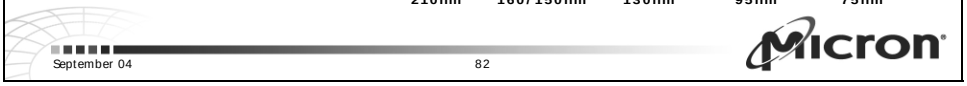

RLDRAM II

- RLDRAM = reduced latency DRAM
- Motivation
  - Address increasingly large percentage of systems that need:
    - Greater request rates
    - Short burst lengths
    - SRAM-like performance but with DRAM density and cost
    - Have low probability of using open rows
  - Maximize number of applications supported
- 2- 533MHz clock frequencies

RLDRAM II Command/ Response Examples

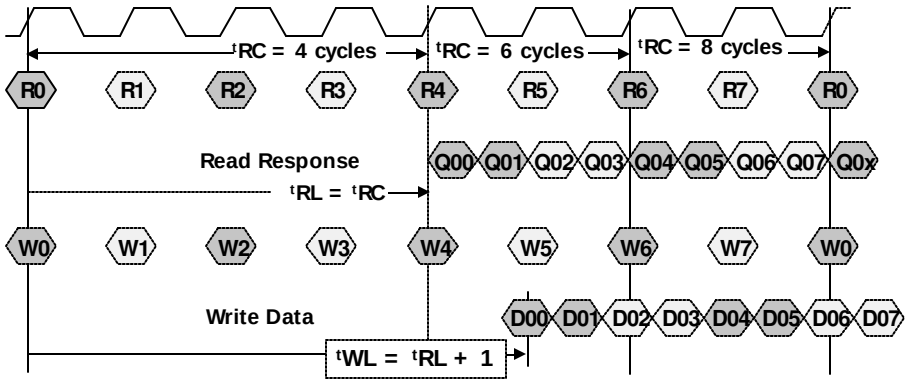

?Note: R0 means read from bank 0. Q00 means first data from bank 0. W0 means write to bank 0. D01 means second data to bank 0.

RLDRAM II Common I/O Cyclic Bank Switching

- x9, x18, x36, 2-, 4- and 8-word bursts

Burst Length	2	4	8
Data Bus Utilization	100% less refresh	100%	100%
MIN Banks Utilized 400 MHz/300 MHz/200 MHz	8/6/4	4/3/2	2/1~2/1

? Design tips:

- Organize data to minimize bank conflicts
- Consider the bank signals as the lowest address bits

RLDRAM II Minimized Bus Turnaround 200 MHz, 2-Word Burst, CIO Data Bus

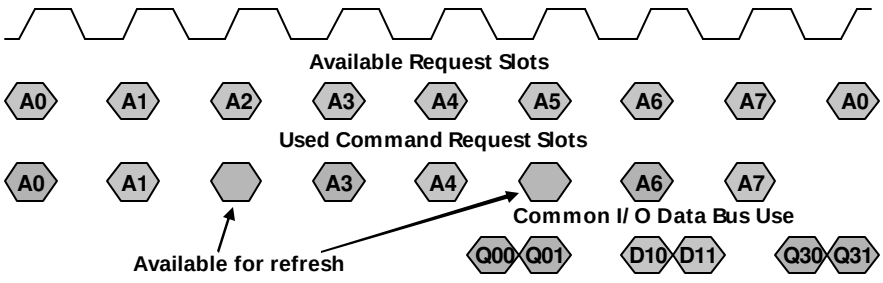

- Worst-case utilization of 66.7 percent at 2-word burst with no unidirectional data streaming
- True for ALL frequencies

Fast Bus Turnaround for High Utilization

- High bandwidth is essential in *all* systems
- Excessive latency can prevent a system from achieving rated bandwidth
- Bus turnaround and latency are reduced in RLDRAM for maximized usable bandwidth

Burst Length	2	4	8
Minimum Bus Utilization	2/3	4/5	8/9

? RLDRAM II minimum bus utilization = BL/ (BL + 1)

Eliminated Bus Turnaround 4-Word Burst, SIO Data Bus

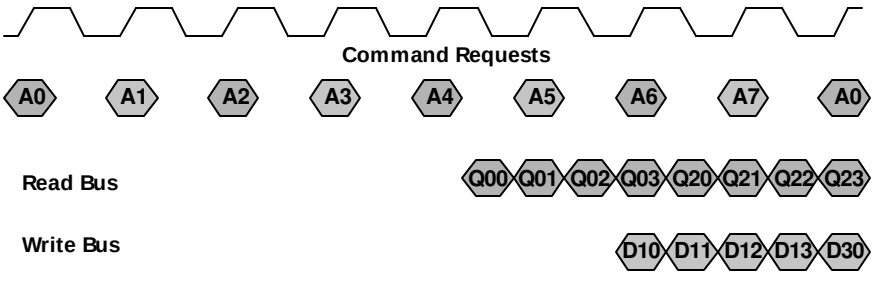

- 100 percent data bus utilization with 4-word burst
- True for ALL frequencies
- Utilization reduced only by refresh requests
- 8-word burst also has 100 percent utilization, unaffected by refresh

RLDRAM II Separate I/O

- Optimal packet buffer
- x9, x18 on each input and output data bus
- No turnaround cycles

Burst Length	2	4	8
Data Bus Utilization	50% less refresh*	100% less refresh	100%
MIN Banks Utilized 400 MHz/300 MHz/200 MHz	8/6/4	8/6/4	4/2~4/2

*Limited by address bandwidth.

Address Multiplexed Operation

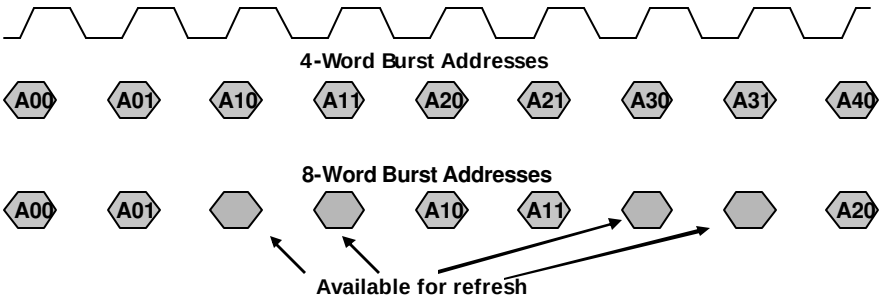

- No performance penalty for 4- or 8-word burst
- Data bus utilization drops to half for 2-word burst

RLDRAM Advantages

- Fastest DRAM ^{t}RC (15ns - 20ns)
- Highly tuned and flexible feature set
  - Excellent command/ data bus utilization
  - Mode selectable 2, 4 and 8 - word burst operation
  - Common I/ O or Separate I/ O data bus versions
  - 1.5V or 1.8V I/ O for operation like SRAM or DDR2 DRAM
  - ODT for clean, high-frequency operation
  - Programmable, self- calibrating output impedance for bus tuning
  - Mode selectable ^{t}RC to match target frequency
- Low system power – fewer resources turn on
- Scalable
- SER is four orders of magnitude less than SRAM

RLDRAM Public Road Map

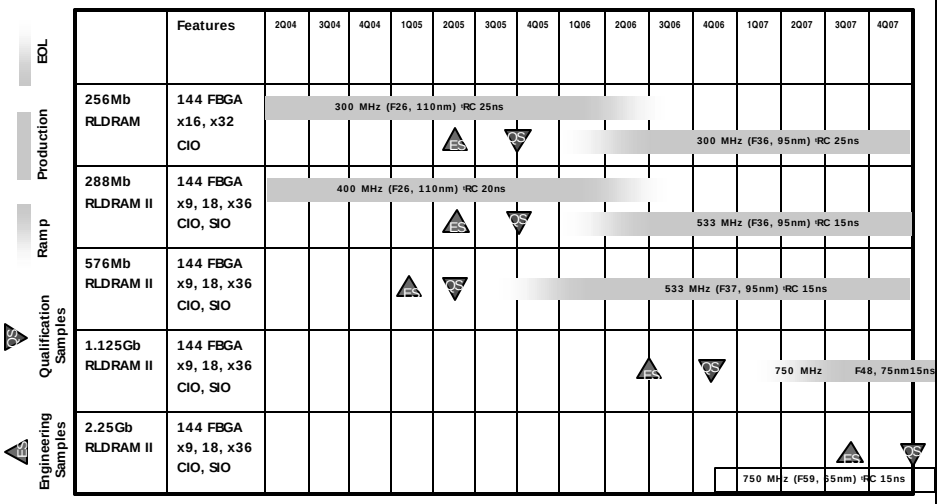

Speeds are shown in clock rates. All information is subject to change without notice. Dates are estimates only. Rev. 02/ 04
There are more devices on roadmap than shown. Contact Micron for further information.

RLDRAM II Performance Model

Clock Frequency		533 MHz		400 MHz	
Clock period	T	1.875 ns		2.5 ns	
ACT to same bk ACT	tRC	15	~8	20	~8
Auto RFSH period	= tRC	15	~8	20	~8

For lower frequencies, T RC = tRC / tCK
Examples:

f	T RC	f	T RC
200	3	150	3
267	4	200	4
333	5	250	5
400	6	300	6
467	7	350	7
533	8	400	8

- 488.28125ns periodic refresh = 3.90625µs period for each bank
- Minimum frequency = refresh rate = 2.048 MHz
- 8 banks
- Recall the disclaimer!

RLDRAM Performance Model

Command	To Same Bank	To Different Bank
Common I/O Device		
R to R	TRC	BL/2
R to W	Max(BL/2, TRC)	BL/2
W to W	TRC	BL/2
W to R	Max(BL/2+1, TRC)	BL/2 + 1

- Write latency = read latency + 1
- BL = 2, 4 or 8- word burst length.
- 9, 18, 36b bus
- Note that 4 cycles are required to transfer BL8 data, hence some bus conflict limitations apply when ^{t}RC is very short.
- Notice for total bus turnaround 1 cycle is lost regardless of BL.
- Recall the disclaimer!

RLDRAM Performance Model

Command	To Same Bank	To Different Bank
Separate I/O Device		
R to R	Max(BL/2, TRC)	BL/2
R to W	Max(BL/2, TRC)	1
W to W	Max(BL/2, TRC)	BL/2
W to R	Max(BL/2, TRC)	1

- Write latency = read latency + 1.
- BL = 2, 4 or 8 - word burst length.
- 9, 18b buses (e.g. 18b D, 18b Q)
- Recall the disclaimer!

FCRAM I, II, AND II+

- FCRAM = fast cycle RAM
- Motivation
  - Address increasingly large percentage of systems that:
    - Need greater request rates
    - Require SRAM- like performance but with DRAM density
    - Have low probability of using open rows
  - Cost leveraging on DDR DRAM concepts

Note: most slide information is from Toshiba presentations at Denali MemCon, October 2003 and May 2004

FCRAM I and II Overview

Compatible with DDR-I SDRAM | Target on Networking Customer

	Network FCRAM-I		Network FCRAM-II
Memory Density (# of bank)	256M / 512Mb (4bank / 8bank)		288 / 576Mb (4bank / 8bank)
I/O Organization	x8, x16		x9*1, x18, x36
Clock Frequency	200MHz	266MHz+	333MHz+
Random Cycle time	25ns	22.5ns	20ns
Data Strobe	Bi-directional DQS		Uni-directional DS&QS
Vdd	2.5V		2.5V
VddQ	2.5V	1.8~1.5V	1.8V ~ 1.5V
I/O Interface	SSTL-2	SSTL-1.8, HSTL	SSTL-1.8, HSTL
Package	66pin TSOP-II*2 60ball mBGA	60ball mBGA	60ball mBGA(x18) 144ball mBGA(x36)

*1: support by 110nm

*2: only for 256Mb

Source: Ohshima- san, Toshiba

Micron

288Mb Network FCRAM-II "D"-die Features

- Process: 130nm
- Organization:
  - x18: 4M words x 4 Banks x 18 bits
  - x36: 2M words x 4 Banks x 36 bits
- Maximum Clock Frequency: 333MHz(666Mbps)
- Pipeline Architecture: Fast Core cycle, tRC=20ns
- Function :
  - Read /CAS Latency(CL) : 4,5 and 6
  - Write /CAS Latency : CL-1
  - Burst Length : 2 and 4
  - Double Data Rate
  - **Uni-directional Data Strobe (DS, QS)**
  - Interface(VddQ):SSTL-1.8(Half-strength)&HSTL(1.8~1.5V)
- Power Supply Voltage: 2.5V
- Package:
  - x18 : 60ball(=4x15) mBGA, 1.0 x 1.0mm Ball pitch
  - x36 :144ball(=8x18) mBGA, 0.8 x 1.0mm Ball pitch, JTAG
- Schedule:
  - x18…**Now Mass production**
  - x36…**Now CS Sampling**

288Mb Chip & Package Photos

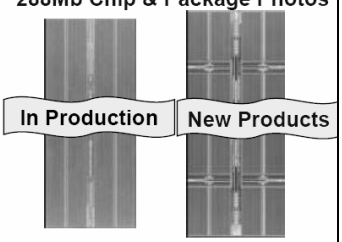

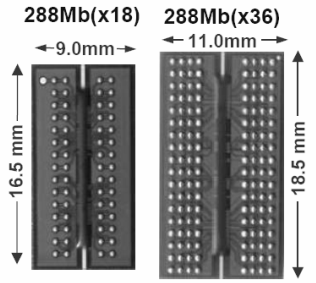

TOSHIBA

Innovation-driven, Customer-focused Growth

Denali MEMCON SAN JOSE OCTOBER 7-8

512Mb Network FCRAM1 "A"-die Features

- Process : 130nm
- Organization :
  - x8…..8M words x **8 Banks** x 8 bits
  - x16…4M words x **8 Banks** x 16 bits
- Pipeline Architecture: Fast Core cycle, tRC=22.5ns
- Maximum Clock Frequency: 266MHz(533Mbps)
- Function :
  - Read /CAS Latency(CL) : 3,4 and 5
  - Write /CAS Latency : CL-1
  - Burst Length : 2 and 4
  - Double Data Rate
  - Data Strobe:
    x8…..Bi-directional DQS, /DQS
    x16…Bi-directional LDQS, UDQS
- Power Supply Voltage: 2.5V
- Interface(VddQ): two version products
  1) SSTL1,8/HSTL(1.8V~1.5V), 2) SSTL2(2.5V)
- Package: 60ball mBGA, 1.0x1.0mm ball pitch
- Schedule:
  1) SSTL1.8/HSTL I/F: **CS Sampling**, **VP in 3Q'04**
  2) SSTL2 I/F: **CS in June'04, VP in 3Q'04**

CS Available Now

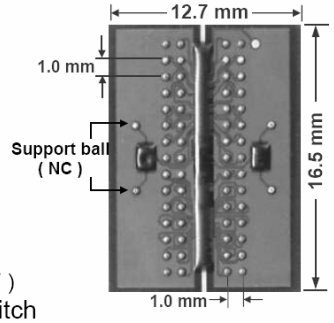

TOSHIBA

Denali MEMCON BOSTON MAY 13

576Mb Future FC2+ spec ~ now finalized ~

- Number of banks : 8 bank
- I/O Organization : x9, x18, x36 with one chip solution
- Clock Freq & tRC : 400MHz+(800+Mbps) & tRC=<20ns
- Function :
  - Read/CAS Latency(CL) : 5,6 and 7(8)
  - Write/CAS Latency(WL) : CL-1
  - Burst Length : 2,4 and 8
  - Data Strobe : Uni-directional Differential DS,QS
  - New Features
    1) ZQ Type OCD, 2) ODT, 3) QVLD
    4) Multi Bank Write, 5) External Bank Refresh
- Power Supply Voltage: 1.8V(1.5V capability considered)
- Interface(VddQ): SSTL-1.8 & HSTL (1.8~1.5V)
- Package : 144 ball mBGA, 0.8 x 1.0mm ball pitch with JTAG
- Schedule: ES in 2Q'05, CS in 3Q'05, VP in 4Q'05

FCRAM II+ Operation

FCR-489, Apr.2004, TOSHIBA

Ext. Bank Refresh with single refresh counter set

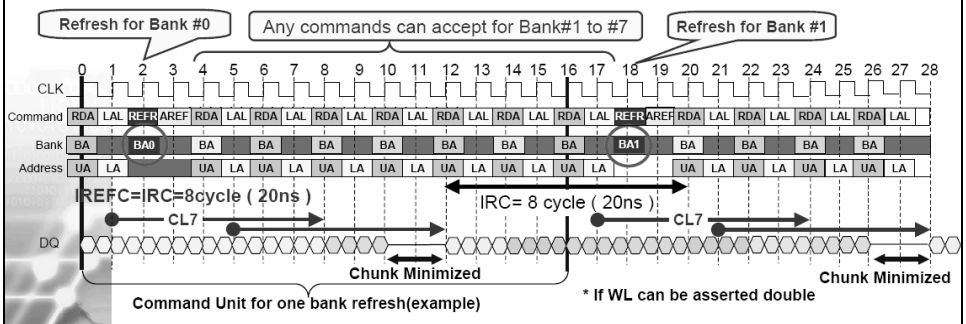

	Refresh Counter Set	Disturb Rate@BL4	Disturb Rate @BL8
FC2+(New)	1	1.024%(Best) or 0.512%*	0(Best)

FCR-489, Apr.2004, TOSHIBA

Multi Bank Write(MBW) Performance and Features

■ Features

-Selectable by EMRS2 set command(A7, A6)

-Write cycle : Same data write to two banks(ignored BA2 address)

-Read cycle : Bank Interleave for BA2=0/1

→**Multi Bank Write feature is best fit for higher read duty (~100% read) applications such as look up table memory.**

■ Performance Comparison

Parameter		BL4	BL8
tRC @CL4		5 CLK(CL+1)	7 CLK(CL+3)
Bus Efficiency @Read	w/o MBW	40%	57%
	with MBW	80%	100%

Multi Bank Write and Bank Interleave Read operation

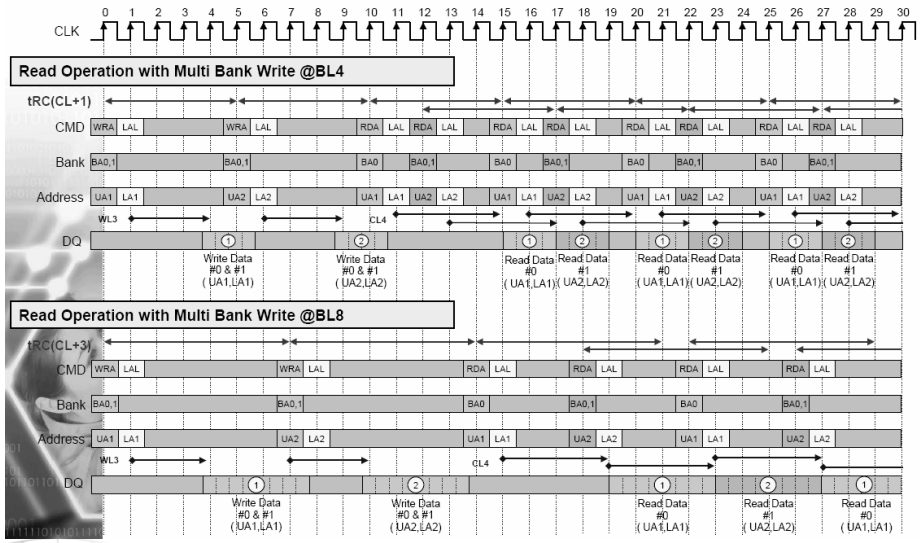

Network FCRAM Roadmap

Note) This information is subject to change without notice. mBGA is lidless(bare chip) BGA.

△=ES ◬=CS ▲=MP Under study "Die Rev.", (Design Rule), Clock Freq. / Random cycle time

Category	Features (Package (pin/ball pitch))	Density (I/O)	CY2003 1Q–4Q	CY2004 1Q–4Q	CY2005 1H–2H	CY2006 1H–2H	CY2007 1H–2H
Network FCRAM1	66pin TSOP II-400(0.65mm) ■ 4bank ■ x8 : Bi-directional DQS ■ x16: Bi-directional LDQS,UDQS ■ SSTL-2 I/F 60ball mBGA(1.0x1.0 mm) ■ 4bank(256Mb), 8bank(512Mb) ■ x8 : Bi-directional DQS, /DQS ■ x16: Bi-directional LDQS,UDQS ■ SSTL-2 I/F(only for 512Mb) ■ SSTL-1.8 I/F and HSTL	256Mb (x8/x16)	"C"(0.175D um) 200MHz / 25ns 66pin TSOP	"E" (110nm) 333MHz / 20ns			
		512Mb (x8/x16)	60ball mBGA	"A" (130nm) 266MHz / tRC=22.5ns	"B" (110nm) 333MHz / 20ns "A" (90nm)		
		1Gb (x8/x16)		60ball mold BGA	"A" (90nm)		
Network FCRAM2 & Network FCRAM2+	60ball mBGA(1.0x1.0 mm) ■ 4bank(288Mb), 8bank(576Mb) ■ Uni-directional DS/QS (per x18) ■ SSTL-1.8 I/F and HSTL I/F 144ball mBGA(1.0x0.8 mm) ■ Same features as 288Mb(x18) ■ JTAG test mode New features for FCRAM2+ ■ ODT ■ BL8 ■ Differential DS and QS ■ QVLD ■ ZQ type OCD	288Mb (x9/x18/x36)	"D" (130nm) 333MHz / 20ns, x18 I/O "D"(130nm) 333MHz / 20ns, x36 I/O	x18: 60ball mBGA x36:144ball mBGA → 144ball mold BGA(CS in May'04) "E" (110nm) 400MHz / 20ns, All I/O			
		576Mb (x9/x18/x36)	x9/x18: 60ball mold BGA x36: 144ball mold BGA Network FCRAM2+	"B" (110nm) 400MHz / 20ns, All I/O	"B" (110nm) 400MHz / 20ns, All I/O "A" (90nm)		
		1.152Gb (x9/x18/x36)		144ball mold BGA	"A" (90nm)		

FCRAM II Performance Model

Clock Frequency	-30	333 MHz		285.7 MHz		250 MHz	
Clock period	T	3 ns		3.5 ns		4 ns	
Cas Latency	CL	6		5		4	
	tRC	~21	7	~21	6	~20	5
		ns	T	ns	T	ns	T
Clock Frequency	-33	300 MHz		266.7 MHz		222.2 MHz	
Clock period	T	3.33 ns		3.75 ns		4.5 ns	
Cas Latency	CL	6		5		4	
	tRC	~23.3	7	~22.5	6	~22.5	5
		ns	T	ns	T	ns	T
Clock Frequency	-40	250 MHz		222.2 MHz		200 MHz	
Clock period	T	4 ns		4.5 ns		5 ns	
Cas Latency	CL	6		5		4	
	tRC	~28	7	~27	6	~25	5
4-bank refresh, all devices tRFC			25		23		19
		ns	T	ns	T	ns	T

- 3.9µs period for 4-bank refresh
- Maximum clock period is 7.5ns on all devices
- Recall the disclaimer!

FCRAM II Performance Model

Command	To Same Bank	To Different Bank
Common I/O Device		
R to R	TRC	2
R to W	TRC	BL/2+2
W to W	TRC	2
W to R	TRC	2

- TRC = CL + 1
- Write latency = CL – 1
- BL = 2 or 4-word burst length.
- Notice for total bus turnaround 2 cycles are lost regardless of BL
- Recall the disclaimer!

FCRAM II+ Performance Model

Clock Frequency		? MHz					
Clock period	T	?					
Cas Latency	CL	TRC - 1					
	tRC	20	8				
		ns	T				
		20	6	300MHz			
		20	7	350MHz			
		20	8	400MHz			
		20	9	450MHz			

- ? ^{t}RC range is 6- 9 cycles, calculate as ^{t}RC/ ^{t}CK.
- ?This is speculative, based only on presentations, not Toshiba data sheets. ^{t}RC might be less than 20ns
- ?Refresh takes ^{t}RC cycles, required every 488ns, 8 bank device These assumptions may be incorrect.
- ?Recall the disclaimer!

FCRAM II+ Performance Model

Command	To Same Bank	To Different Bank
Common I/O Device		
R to R	TRC	max(BL/2 , 2)
R to W	TRC	BL/2+2
W to W	TRC or TRC+2 if BL8 MBW	max(BL/2 , 2)
W to R	TRC or TRC+2 if BL8 MBW	max(BL/2 , 2)

- ? This is speculative, based only on presentations, not Toshiba data sheets.
- ? Write latency = CL – 1
- ? BL = 2, 4 or 8- word burst length.
- ? MBW is Multi Bank Write, i.e. dual- bank write not used in today's performance scenarios
- ? Notice for total bus turnaround 2 cycles are lost regardless of BL
- ? Recall the disclaimer!

GDDR3

- GDDR3 = graphics double data rate DRAM version 3
- Motivation
  - Provide vastly higher pin bandwidth for wide-bus graphic applications
  - Simplify DRAM commands
  - Optimize die size

GDDR3 Features

- Clock frequency of 500 MHz- 700 MHz
- Single- ended read and write strobes
- RDQS and WDQS per byte
- On-die termination
- Dynamic programmable impedance output driver
- Duty cycle correction on clock input
- 1.8V core
- 4x 2M x 32 (256Mb)

GDDR3 Features (continued)

- 1.8V pseudo-open drain I/O
- ^{t}RAS lockout
- Concurrent AUTO PRECHARGE
- 4 banks
- 4K refresh
- Burst length 4
  - 8 is defined but not supported by all manufacturers
- Sequential burst type only

High-Speed Advantages

- Dynamically controlled impedance output driver
- The DRAM controls the on-die termination for the reads
- On-die termination on all address and control pins
- Single-ended read and write strobes
- Reduced WRITE latency
- Designed specifically for high-speed, point-to-point applications

Key Timing Parameters

- ^{t}CK = 500 MHz to 700 MHz
- ^{t}RAS = 30ns
- ^{t}RC = 60ns
- ^{t}RCD = 16ns
- ^{t}RP = 13ns
- ^{t}RRD = 6ns
- ^{t}WR = 4 ^{t}CK
- WRITE latency = 1 to 4 clocks
  - Set in mode register
    - 3 to 4 clocks for low-power operation
    - 1 to 2 clocks with increased operating power

Read Data

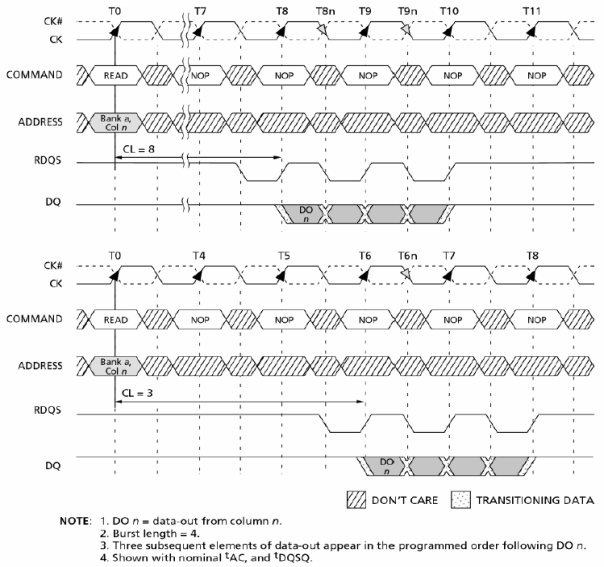

NOTE: 1. DO *n* = data-out from column *n*.
2. Burst length = 4.
3. Three subsequent elements of data-out appear in the programmed order following DO *n*.
4. Shown with nominal ^{t}AC, and ^{t}DQSQ.

Write Data

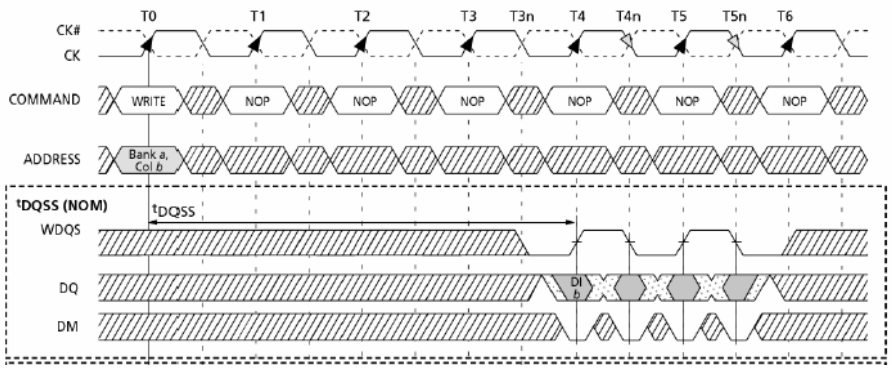

Read to Write

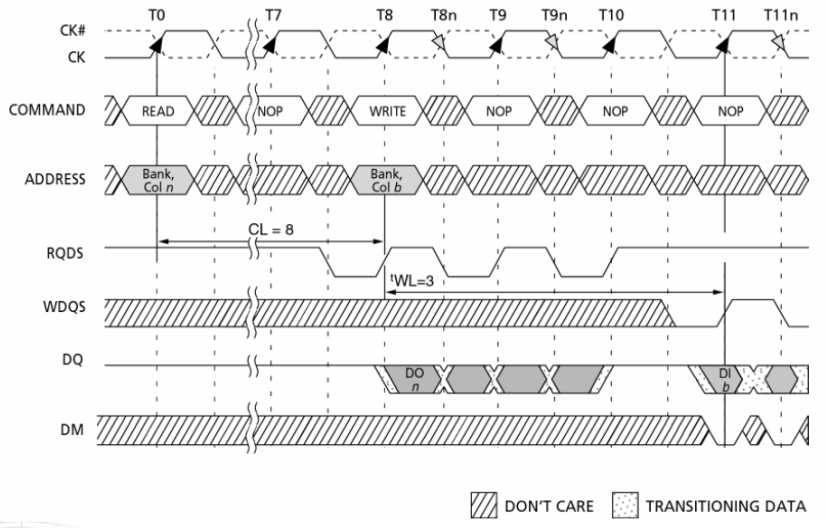

GDDR3 SDRAM Performance Model

Clock Frequency		700 MHz		600 MHz		500 MHz	
Clock period	T	1.429 ns		1.667 ns		2 ns	
Cas Latency	CL	9		8		7	
		ns	T	ns	T	ns	T
ACT to same bk ACT	tRC		31		27		21
ACT to R	tRCDR		10		9		7
ACT to W	tRCDW		6		5		4
ACT to PRE	tRAS		22		19		15
PRE Period	tRP		9		8		6
W to PRE	tWR		9		8		7
Internal W to R	tWTR		3		3		3
ACT to diff bk ACT	tRRD		8		7		5
Auto RFSH period	tRFC		39		33		27
Last data in to R	tCDLR		5		4		3
Write Latency	WL		5		5		4

- WL is programmable; value shown is used in model
- T represents 1 clock cycle
- Maximum clock period is 3.33 ns
- Recall the disclaimer!

GDDR3 SDRAM Performance Model

Command	To Same Bank	To Different Bank
W to R	WL + BL/2 + T WTR	1
W to W	BL/2	BL/2
W to PRE	WL + BL/2 + TWR	1
W to ACT		1
R to R	BL/2	BL/2
R to W	CL + BL/2	CL + BL/2 + 1 - WL
R to PRE	2	1
R to ACT		1
ACT to ACT	TRC	TRRD
ACT to R	TRCDR	1
ACT to W	TRCDW	1
ACT to PRE	TRAS	1

- Refresh is 32ms, 4K = 7.8125µs periodic.
- BL = 4. 4 banks.
- Auto precharge not enabled.
- Recall the disclaimer.

RAMBUS XDR™

Rambus XDR

FCR-489, Apr.2004, TOSHIBA

XDR DRAM Features

Preliminary

- Process : 130 nm / 110 nm
- Organization :
  - 512Mb… 4M words x 8Banks x 16bits
  - 256Mb… 2M words x 8Banks x 16bits
- Function :
  - 1clock Request Packet
  - 2clock Data Packet (16 bit burst)
  - Page Size : 2K Byte
  - tREFI : 0.49us (= 16ms/32K)
  - Bank Interleaved Refresh
  - Frequency (CLK): 300MHz / 400MHz / 500MHz
  - Frequency (DQ) : 2.4GHz / 3.2GHz / 4.0GHz
  - Octal Data Rate (ODR)
  - Write Mask
  - Early Read After Write
  - Powerdown Self-refresh support
  - Dynamic Width Control: x16/x8/x4

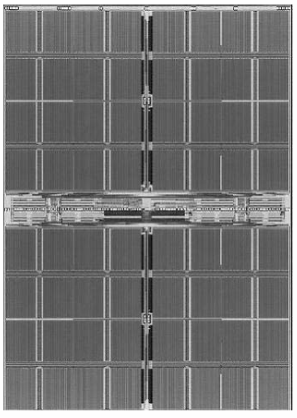

Source: Ohshima-san, Toshiba

XDR DRAM Package Information

108ball molded BGA (Lead free)

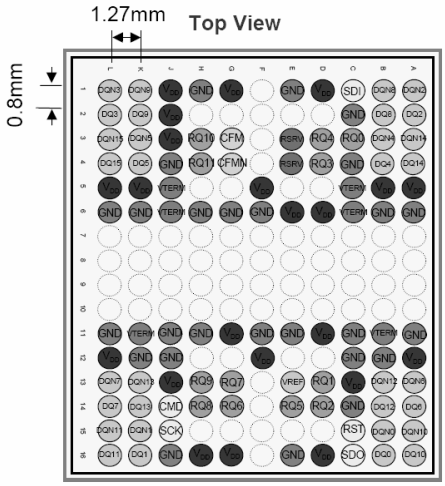

XDR DRAM Application-1

Dynamic Width Control allows system MB scaling without impacting performance

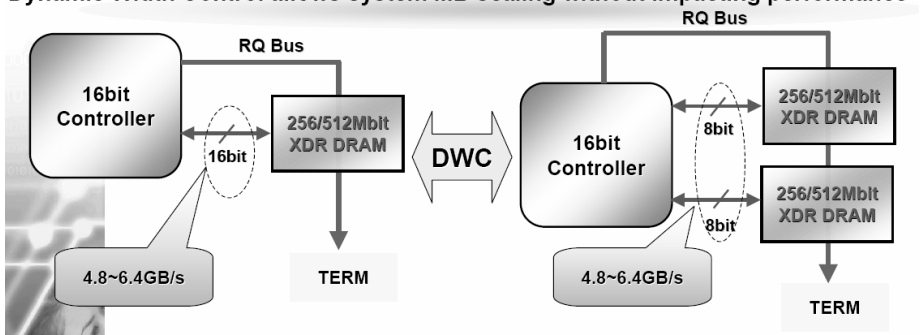

- **XDR Solution realizes 4.8~6.4GB/s with 90 controller pins (incl. power/ground)**

✓ **Solves the procurement nightmare, minimizes inventory loss, saves pin count, and saves development cost of having multiple DRAMs and controllers.**

✓ **XDR DRAM will realize 2 to 3 chip system solution with significant BW for Digital Consumer Systems in 2005 timeframe.**

Toshiba XDR™ DRAM Roadmap

Preliminary

△ WS △ ES ◬ CS ▲ MP Under Study

Features / Package	Density (I/O Width)	CY2003 1Q 2Q 3Q 4Q	CY2004 1Q 2Q 3Q 4Q	CY2005 1Q 2Q 3Q 4Q	CY2006 1Q 2Q 3Q 4Q
molded BGA108 (1.27x0.8 mm pitch) ■ 8 banks ■ 2.4G/3.2G/4.0Gbps bandwidth ■ Dynamic Width Control ■ Early Read After Write ■ DRSL I/F	512Mb (x4/x8/x16)		WS Only 1st (130 nm)		
molded BGA (1.27x0.8 mm pitch) ■ 8 banks ■ 2.4G/3.2G/4.0Gbps bandwidth ■ Dynamic Width Control ■ Early Read After Write ■ DRSL I/F	256Mb [2] (x4/x8/x16) 512Mb [2] (x4/x8/x16)			2.4G/3.2Gbps/(4.0Gbps) 2nd (110 nm) △ ◬ ▲ 2.4G/3.2Gbps/(4.0Gbps) 2nd (110 nm) △ ◬ ▲	

Note: 1. This information is subject to change without notice.

2. 110nm XDR device density's development priority is to be decided.

XDR Performance Model

Clock Frequency	400MHz (2.5ns) Device A or B			300MHz (3.3ns) Device A	
Clock period	T	A device (T)		B device (T)	
		same b	different	same b	different
ACT to ACT	tRC / tRR	16	4	20	4
ACT to R	tRCD-R	5	1	7	1
ACT to W	tRCD-W	1	1	3	1
ACT to PRE	tRAS	10	1	13	1
PRE Period (P to A)	tRP	6	-	7	-
R to Q	tCAC	6	-	7	-
W to D	tCWD	3	-	3	-
R to R or W to W	tCC	2	2	2	2
R to W	tdRW	8	8	9	9
W to R	tdWR	9	2	10	2
W to PRE	tWRP	10	1	12	1
R to PRE	tRDP	3	1	4	1
Auto RFSH period	tRFC	16	-	20	-
Refresh interval		488ns		488ns	

- T represents 1 clock cycle
- Maximum clock period is 3.83ns (261 MHz)
- Note: 500 MHz spec (XDR4000Mb/s) is excluded from performance analysis – too far out! Recall the disclaimer!

XDR Examples

- 300 MHz Device A: 53.3 ns ^{t}RC
- 400 MHz Device B: 50 ns ^{t}RC
- 400 MHz Device A: 40 ns ^{t}RC
- Future
  - 500 MHz Device B: 40 ns ^{t}RC
  - 500 MHz Device A: 32 ns ^{t}RC

Performance Analysis Methodology

- Proprietary software written in Visual Basic includes
  - User- defined advanced address generation
    - Unlimited multiple- thread capability
    - Unique request definition for each thread
    - Address generation per probability inputs
  - Memory behavioral model
    - The "Performance Model" for each memory type
  - Controller behavioral model optimized for each memory device

Performance Analysis Methodology (continued)

- Objective is to compare all memory devices fairly
- Expose each one to the identical request steam
- Controllers independently allowed to optimize request stream limited by common constraints
  - e.g. allowed to switch threads if stalled
    - None of today's examples use multi-threading
- Frequency sweep of all devices for each scenario
  - Include "sweet spot" for each memory

Component Choices

- 512~576Mb device wherever possible, otherwise nearest lower density
- DDR SDRAM:
  76.9-133.3 MHz CL2, to 166.7 MHz CL2.5, to 200 MHz CL3
- DDR2 SDRAM:
  125-200 MHz CL3, to 266.7 MHz CL4, to 333.3 MHz CL5
- RLDRAM II:
  15 ns ^{t}RC for CIO and SIO: 2-200 MHz 3T, to 266.7 MHz 4T,
  to 333.3 MHz 5T, to 400 MHz 6T, to 466.7 MHz 7T, to 533.3 MHz 8T
- FCRAM2: -30 (22.5 ns ^{t}RC): 133.3-250 MHz CL4, to 285.7 MHz CL5,
  to 333.3 MHz CL6 (note RC = CL+1)
- FCRAM2+: -25 (20 ns ^{t}RC): 200-300 MHz CL5, to 350 MHz CL6,
  to 400 MHz CL7, to 450 MHz CL8
- QDR2 SRAM: 0-250 MHz
- DDR2 SRAM: 0-333.3 MHz
- QDR3 & DDR3 SRAM: 250-500 MHz

Frequencies Analyzed

Clock Frequency (MHz)	Clock Period (ns)	Frequencies of Interest Especially Pertain To:
76.9	13	DDR SDRAM
100	10	
125	8	DDR2 SDRAM
133.3	7.5	DDR SDRAM, FCRAM2
150	6.667	RLDRAM2
166.7	6	DDR SDRAM
200	5	DDR SDRAM, DDR2 SDRAM, RLDRAM2, FCRAM2
222.2	4.5	FCRAM2
250	4	RLDRAM2, FCRAM2, ALL SRAM
261.1	3.83	XDR
266.7	3.75	DDR2 SDRAM, RLDRAM2, FCRAM2
285.7	3.5	FCRAM2
300	3.33	RLDRAM2, FCRAM2, FCRAM2+, GDDR3, XDR
333.3	3	DDR2 SDRAM, RLDRAM2, FCRAM2, QDR2/DDR2 SRAM
350	2.857	RLDRAM2, FCRAM2+
400	2.5	RLDRAM2, FCRAM2+, XDR
450	2.222	FCRAM2+
466.7	2.143	RLDRAM2
500	2	GDDR3, QDR3/DDR3 SRAM
533.3	1.875	RLDRAM2
600	1.667	GDDR3
700	1.429	GDDR3

Shaded areas signify precise frequency or period

Performing a Fair Comparison

- Must have
  - Same number of data bus signals
  - Same data size response per request
- Examples
  - Differential data devices have double the pin count
    - Compare x16 diff I/O device with x32 single-ended
  - Separate I/O devices have two data buses
    - Compare x8 SIO (2 buses, x8 each) with x16 CIO

Fair Comparison Chart
32b Base Data Size (16b x 2 Word Burst)

Device	BL2	BL4	BL8	BL16	BL32	BL64
16b DDR SDRAM, 16(18)b RLDRAM2CIO & FCRAM2+	2	4	8	2x 8	4x 8	8x 8
16b DDR2 SDRAM	-	4	8	2x 8	4x 8	8x 8
16b FCRAM2, 16(18)b DDR2 SRAM	2	4	2x 4	4x 4	8x 4	16x 4
8(9)b RLDRAM2SIO	4	8	2x 8	4x 8	16x 8	32x 8
8b XDR	-	-	16	2x 16	4x 16	8x 16
8(9)b QDR2 & QDR3 SRAMs	4	2x 4	4x 4	8x 4	16x 4	32x 4
16(18)b DDR3 SRAMs	2	2x 2	4x 2	8x 2	16x 2	32x 2
8(9)b DDR2SIO SRAM	2x 2	4x 2	8x 2	16x 2	32x 2	64x 2

Fair Comparison Chart
64b Base Data Size (32b x 2 Word Burst)

Device	BL2	BL4	BL8	BL16	BL32	BL64	BL128
32b GDDR3 SDRAM	-	4	8	2x 8	4x 8	8x 8	16x 8
16b XDR	-	-	16	2x 16	4x 16	8x 16	16x 16
32(36)b RLDRAM2CIO & FCRAM2+	2	4	8	2x 8	4x 8	8x 8	16x 8
16(18)b RLDRAM2SIO	4	8	2x 8	4x 8	8x 8	16x 8	32x 8
32b FCRAM2	2	4	2x 4	4x 4	8x 4	16x 4	32x 4
16(18)b QDR2, DDR2SIO, QDR3 SRAMs	4	2x 4	4x 4	8x 4	16x 4	32x 4	64x 4
32(36)b DDR2 & DDR3 SRAMs	2	4	2x 4	4x 4	8x 4	16x 4	32x 4

Latest Results

- All latest scenarios use "Even Distribution" of requests
  - E.g. 1R 1W means 1 read from a random address, then 1 write from a random address
  - Random address also means random bank #
- If Burst Length exceeds device burst length
  - Additional data is from open row (e.g. DDR, DDR2, GDDR3, XDR DRAMs)
  - Additional data is from next adjacent bank for devices with no open rows (e.g. FCRAM, RLDRAM)
  - Additional data is from next address for SRAMs

Corrected Scenario A – Random BL2

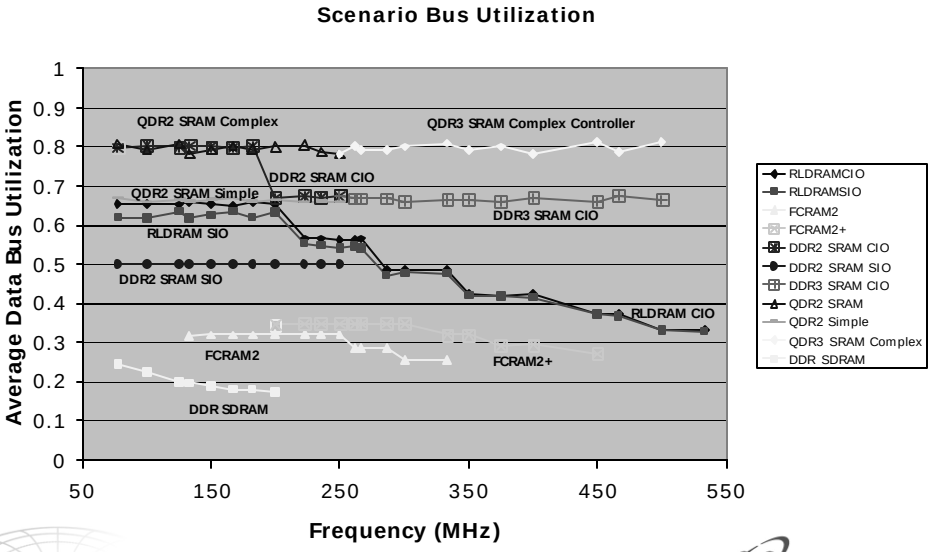

Corrected Scenario B – Random BL4

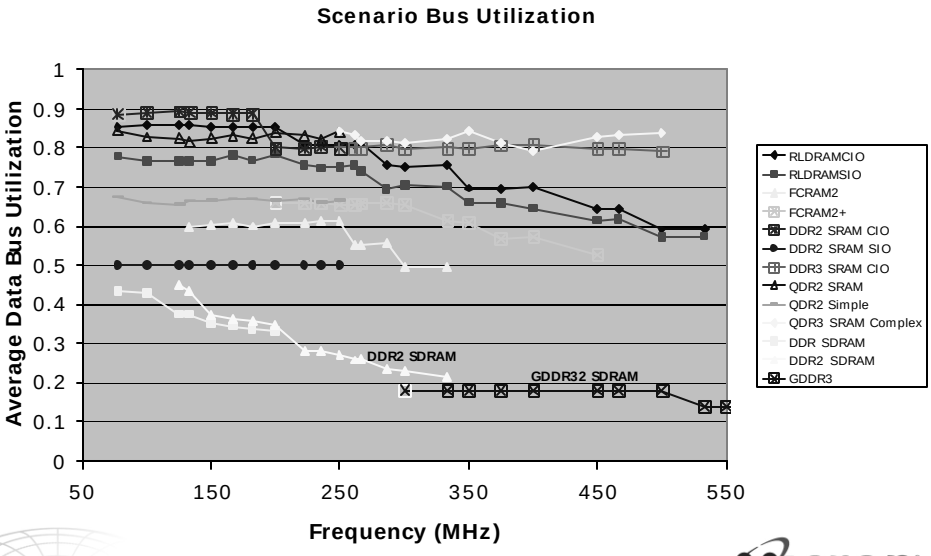

Corrected G – Random 9R 1W BL2, Locality*

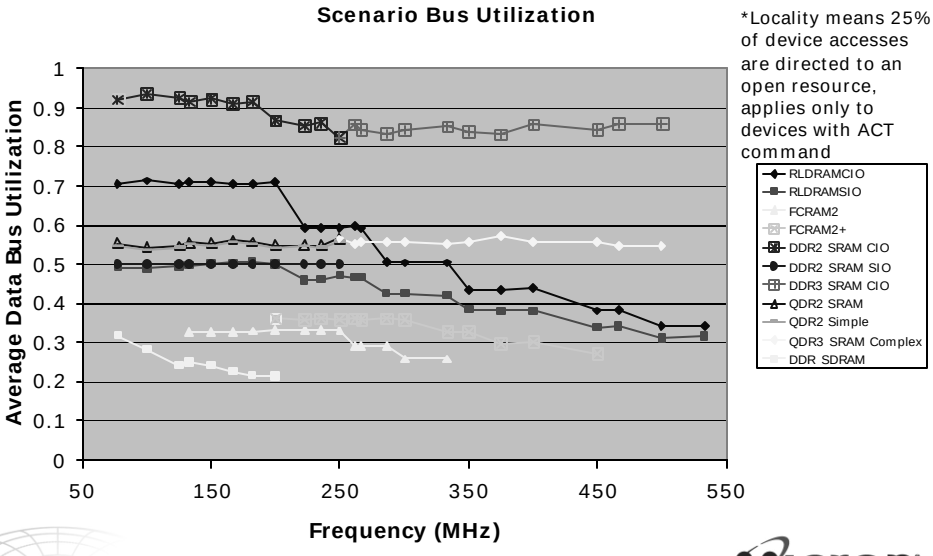

*Locality means 25% of device accesses are directed to an open resource, applies only to devices with ACT command

Corrected H – Random 9R 1W BL4 Locality

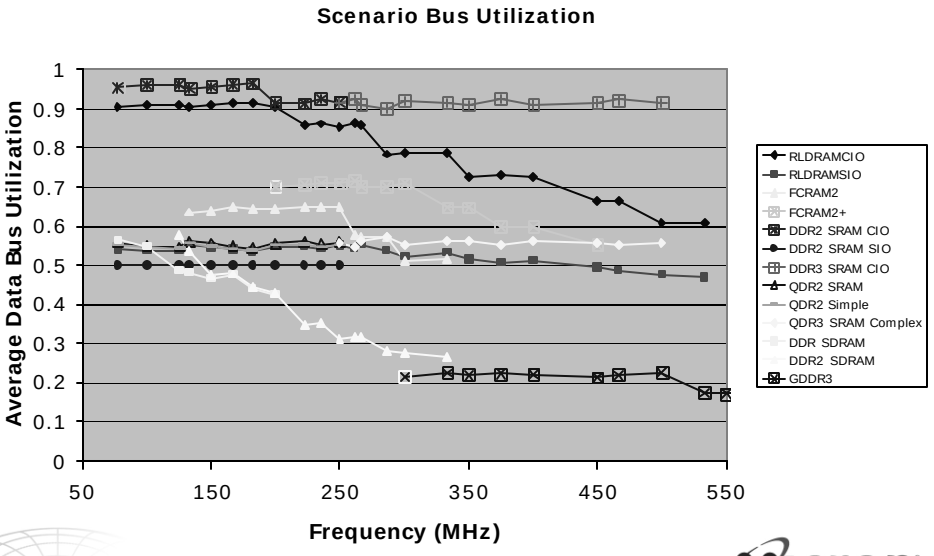

Corrected I – Random 99R 1W BL2, Locality

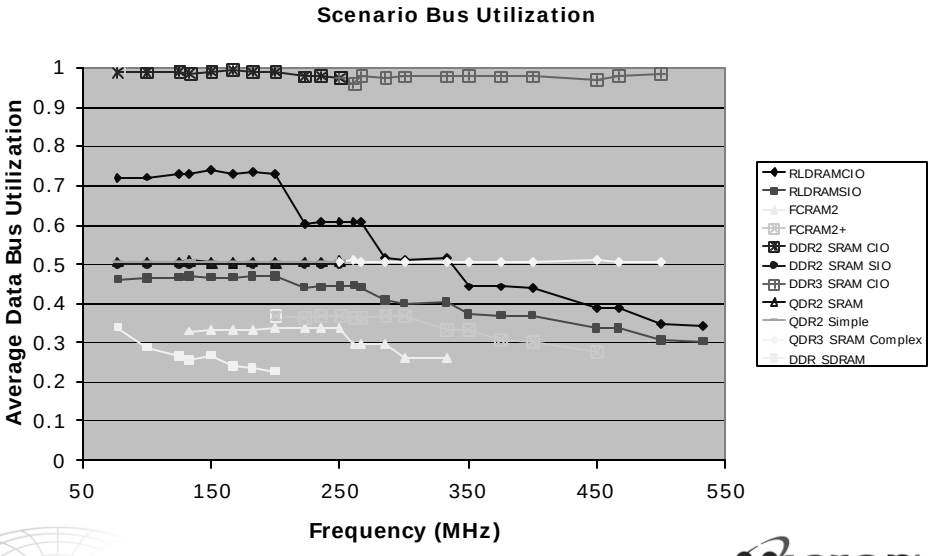

Corrected J– Random 99R 1W BL4 Locality

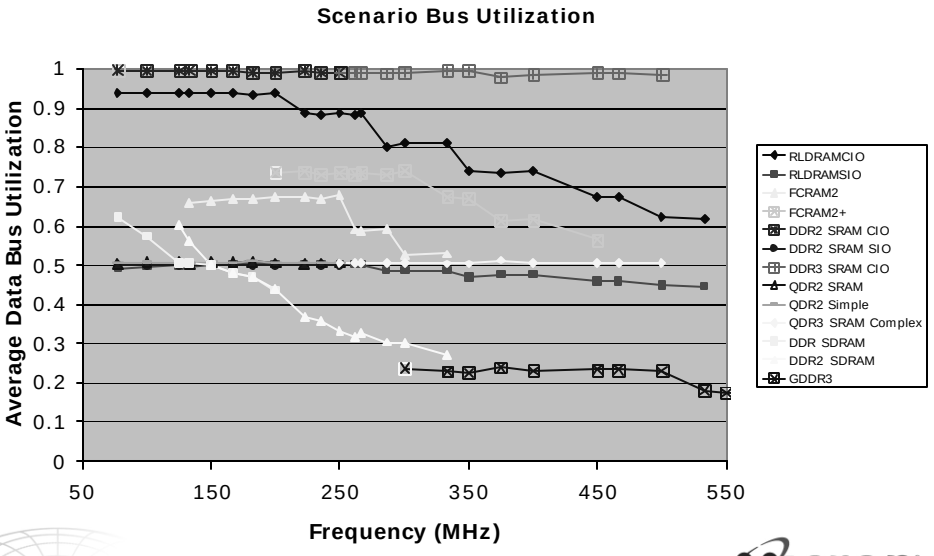

Evenly Distributed 1R 1W BL2

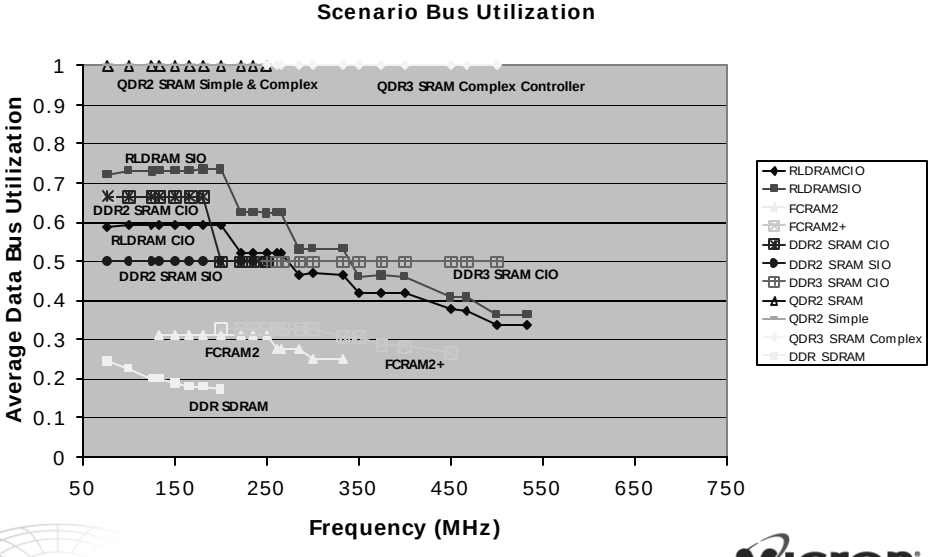

Evenly Distributed 1R 1W BL2

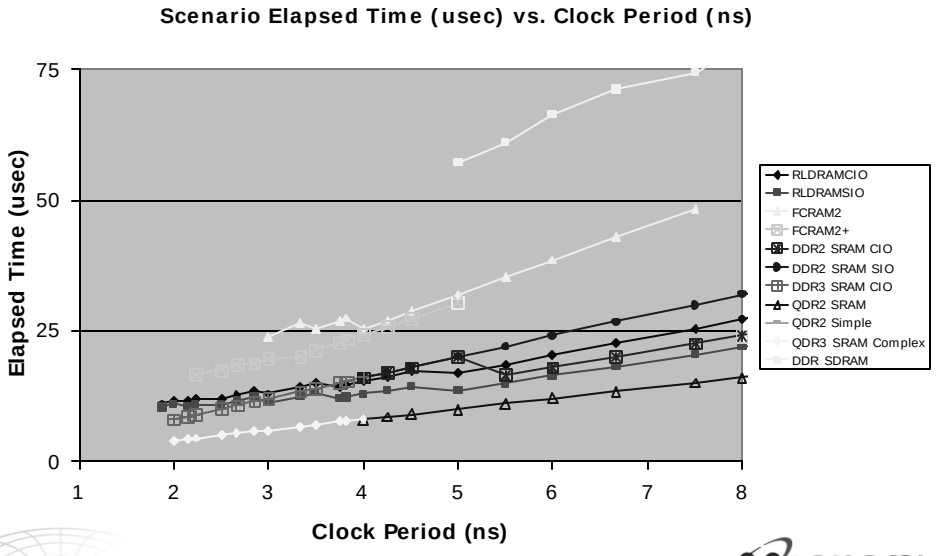

Evenly Distributed 1R 1W BL4

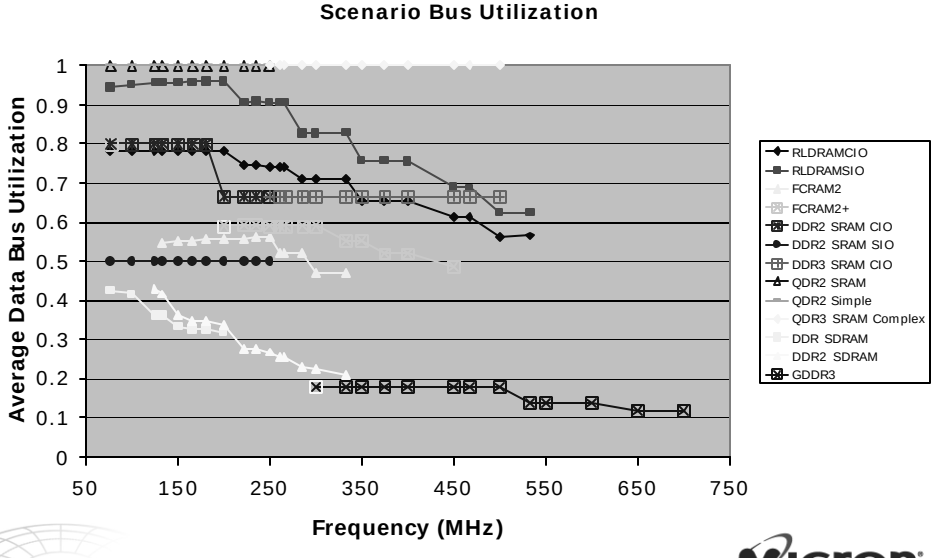

Evenly Distributed 1R 1W BL4

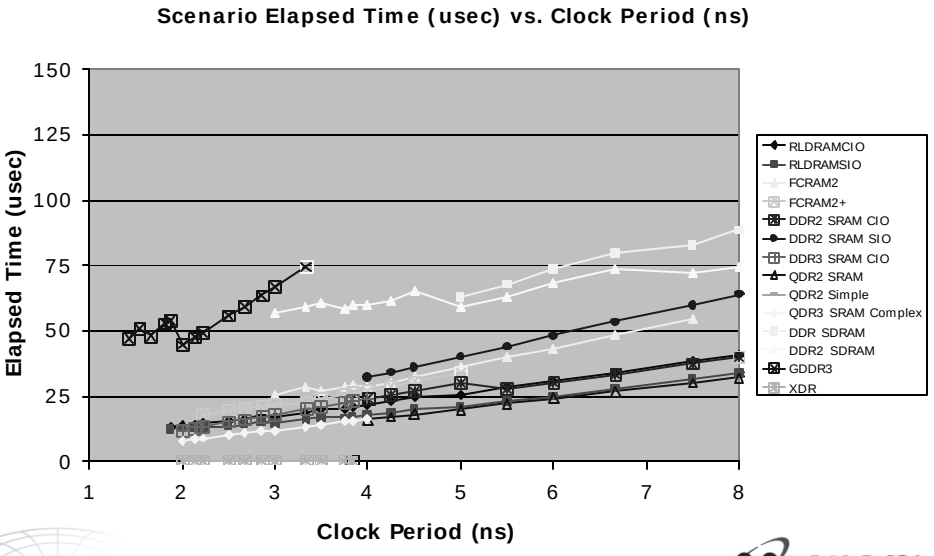

Evenly Distributed 1R 1W BL8

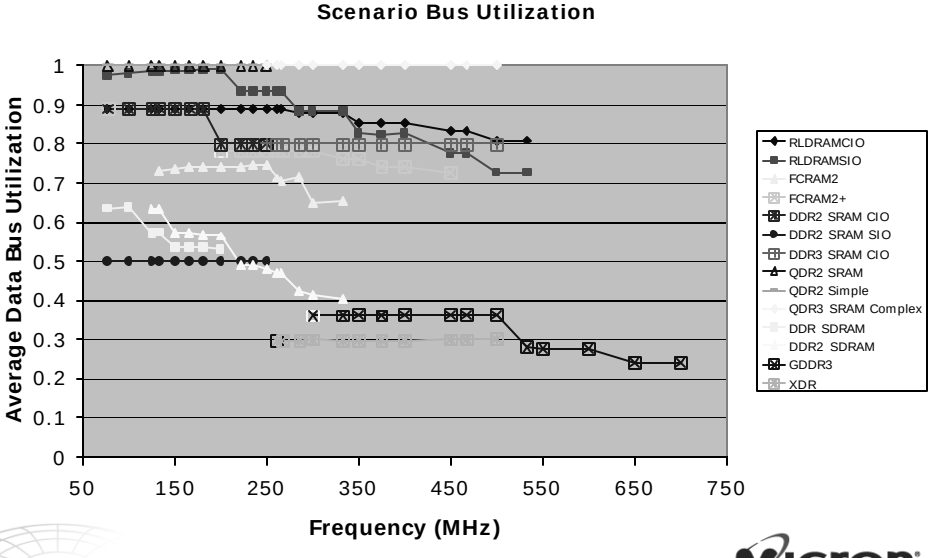

Evenly Distributed 1R 1W BL8

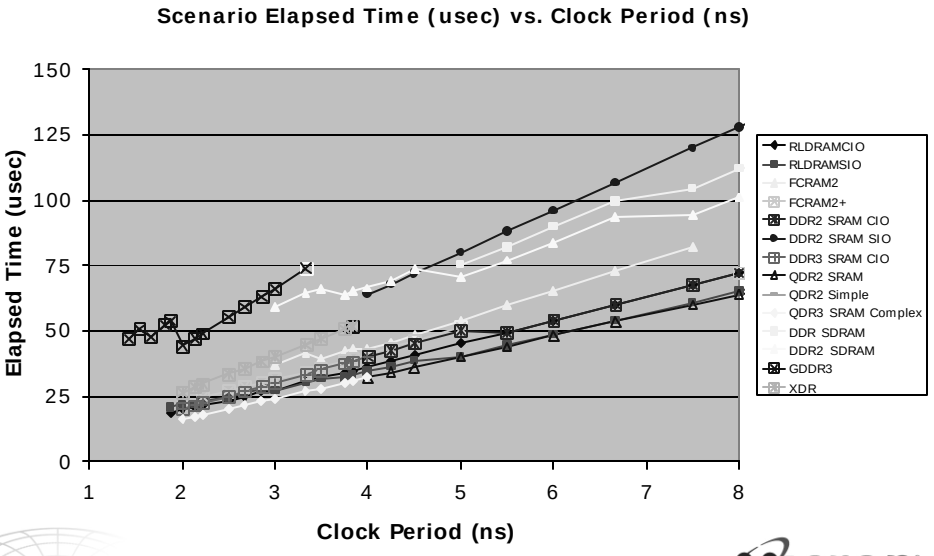

Evenly Distributed 1R 1W BL16

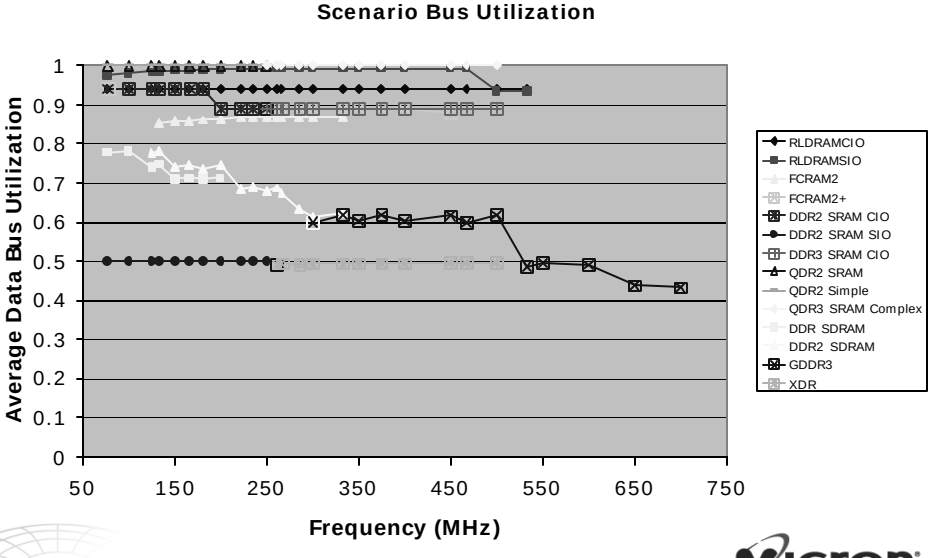

Evenly Distributed 1R 1W BL16

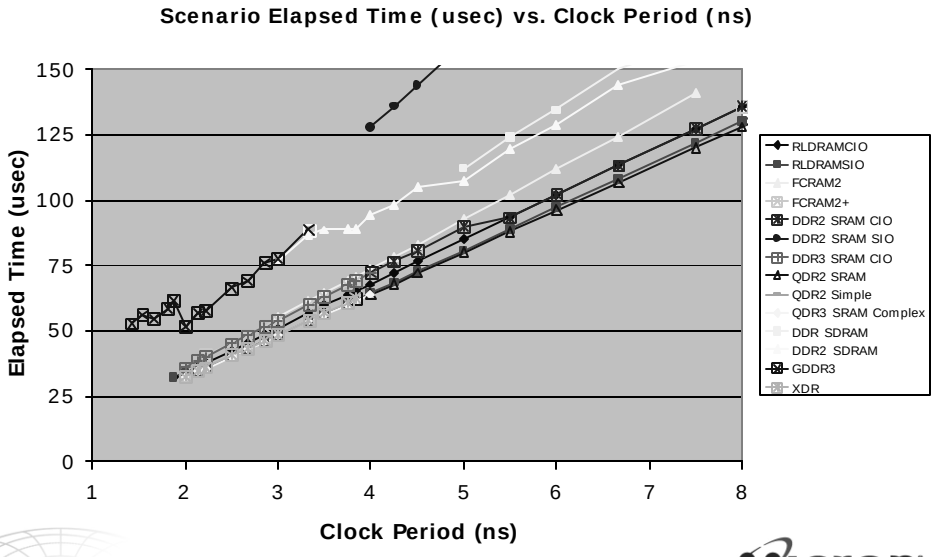

Evenly Distributed 1R 1W BL32

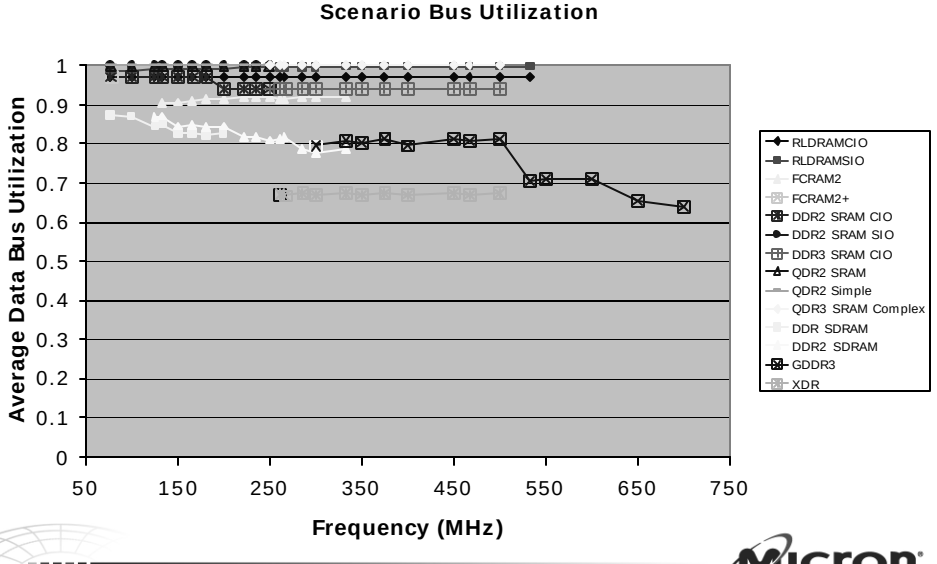

Evenly Distributed 1R 1W BL32

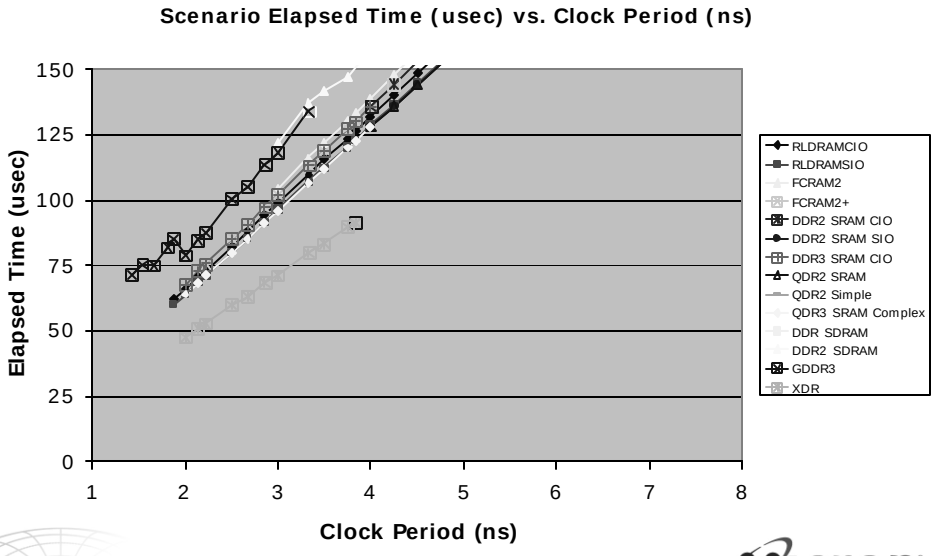

Evenly Distributed 1R 1W BL64

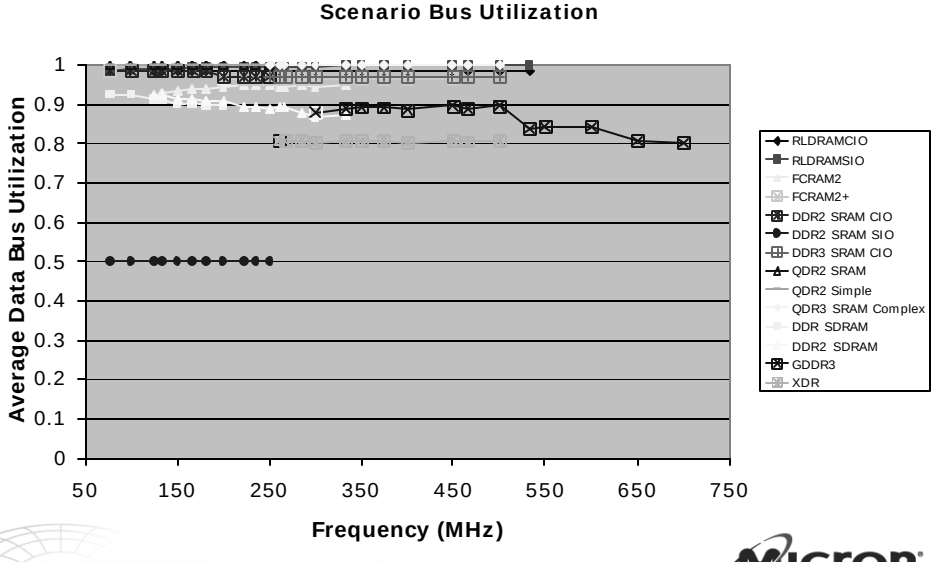

Evenly Distributed 1R 1W BL64

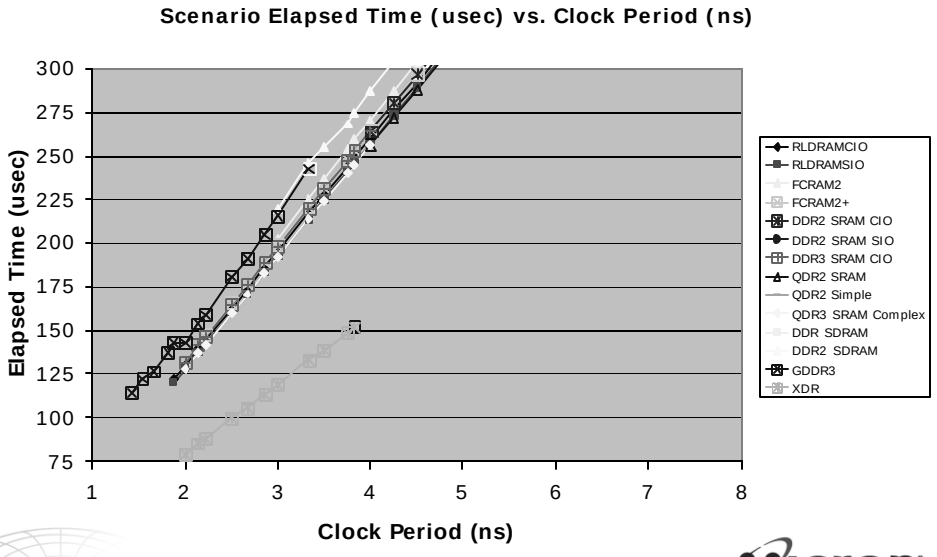

Evenly Distributed 2R 1W BL2

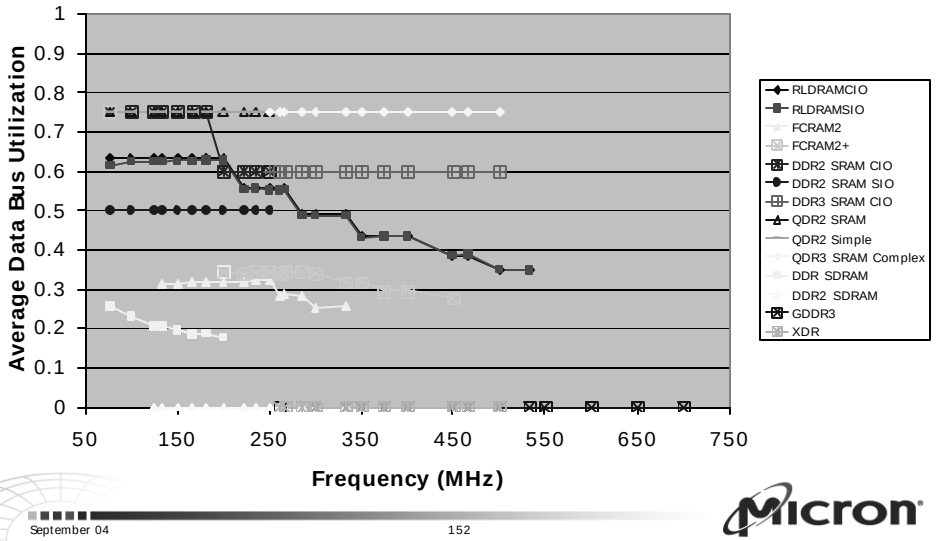

Evenly Distributed 2R 1W BL4

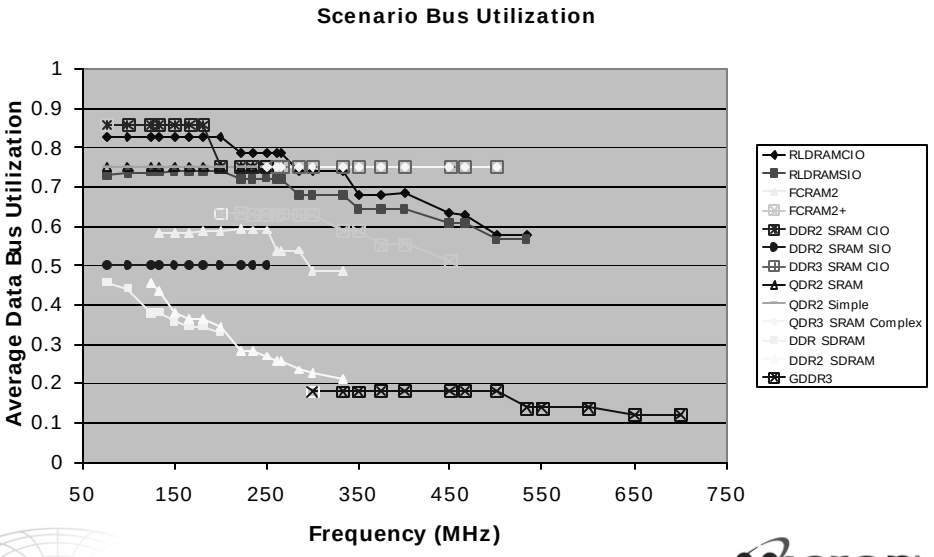

Evenly Distributed 2R 1W BL8

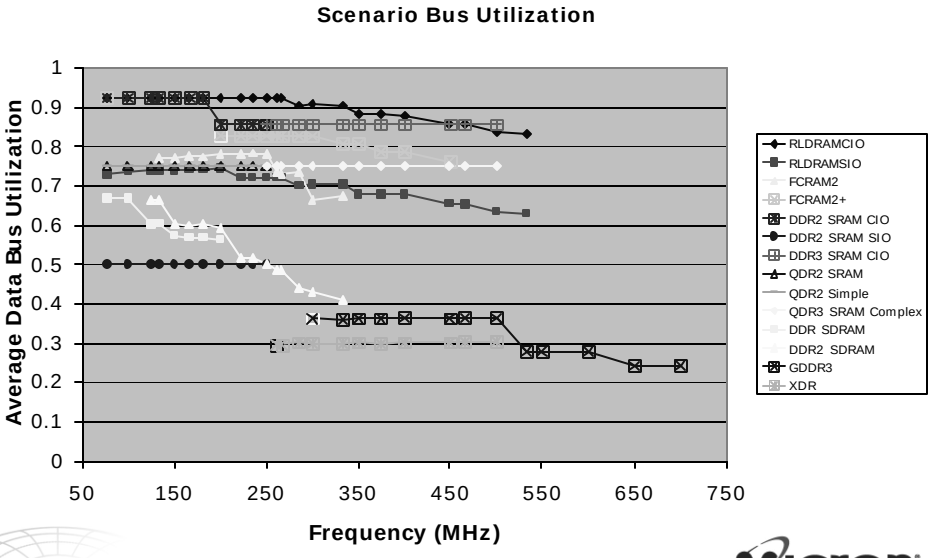

Evenly Distributed 3R 1W BL2

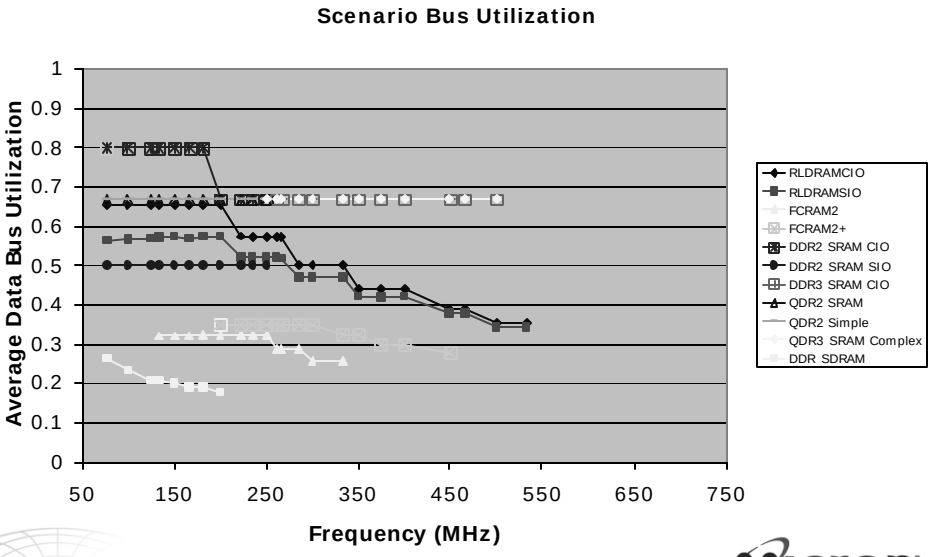

Evenly Distributed 3R 1W BL4

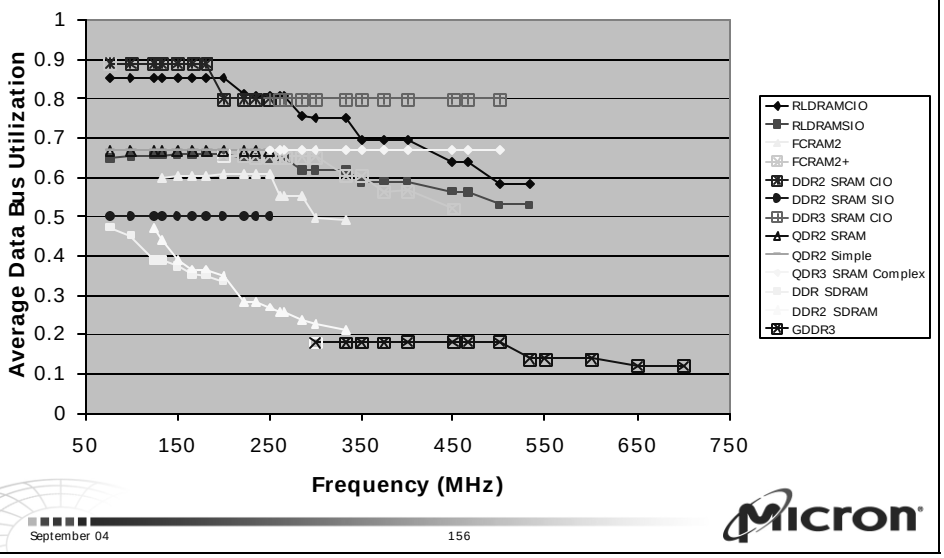

Evenly Distributed 3R 1W BL8

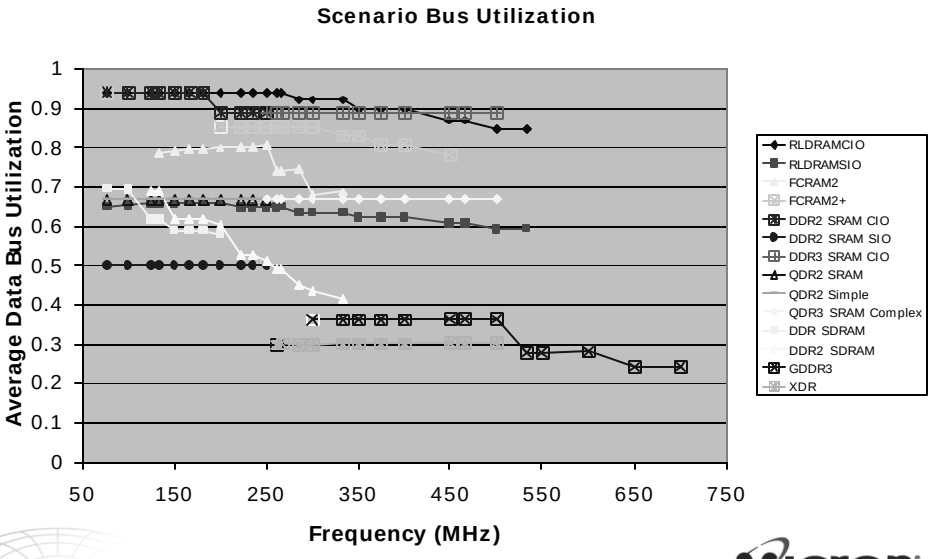

Performance Results Notes

- Long latency QDR3 and DDR3 performs like their predecessors
  - Differences will be seen when pipeline stalls are modeled
  - Thread stalls model system waiting for dependent data
  - Can't issue new command until data is received and manipulated
  - Longer latency devices affected most by this
- Observed influence of controller design on results
  - Any knowledge of what comes next helps
    - Controller can add latency to gain this knowledge
    - Done only to a very limited degree in this analysis
      - E.g. QDR complex controllers stall R issue until W available but will not reorder operations otherwise
      - E.g. Low latency DRAMs don't issue refresh to a bank with a pending request

Conclusions

- "Commodity" DRAMs improve performance with
  - Greater locality
    - Fewer bank misses
  - Fewer bus turnarounds
    - Longer burst lengths and greater R:W ratios
- XDR outperforms all other memories once data chunking gets extreme
  - E.g. Some specialized graphics applications
- RLDRAM II outperforms other DRAMs in every other scenario
  - Rivals SRAM performance
- SRAM still outperforms other memories for 2-word burst
  - Has the edge at 4-word burst above 300~400MHz but not less than that
  - No longer has much advantage otherwise except for determinism

Conclusions (continued)

- It is easy to predict SRAM behavior
- Not so easy to predict DRAM behavior
  - Probabilistic resource availability
  - Require appropriate controller and device model to assess
    - Similar controller intelligence is essential when comparing memory devices
- Future: continued DRAM improvements
  - Eventual disappearance of discrete SRAM
    - Perhaps this decade!

Acknowledgements

- Micron
  - Jeff Janzen, Chris Johnson, Computing & Consumer Group
  - Todd Dinkelman, Jim Cooke, Negin Roohi, NetCom Group
  - Jim Lane (Technical Writer)
- Toshiba
  - Shigeo Ohshima
- QDR Co-development Group (Cypress, IDT, NEC, Renasas, Samsung)
  - Mike Pearson
- RAMBUS
  - XDR data sheets
- Stanford University
  - Christos Kozyrakis for his valuable review and suggestions

Micron Technology