

Retention-Aware Placement in DRAM (RAPID): Software Methods for Quasi-Non-Volatile DRAM

Ravi K. Venkatesan

Stephen Herr, Eric Rotenberg



Center for Embedded Systems Research (CESR)
Department of Electrical & Computer Engineering
North Carolina State University


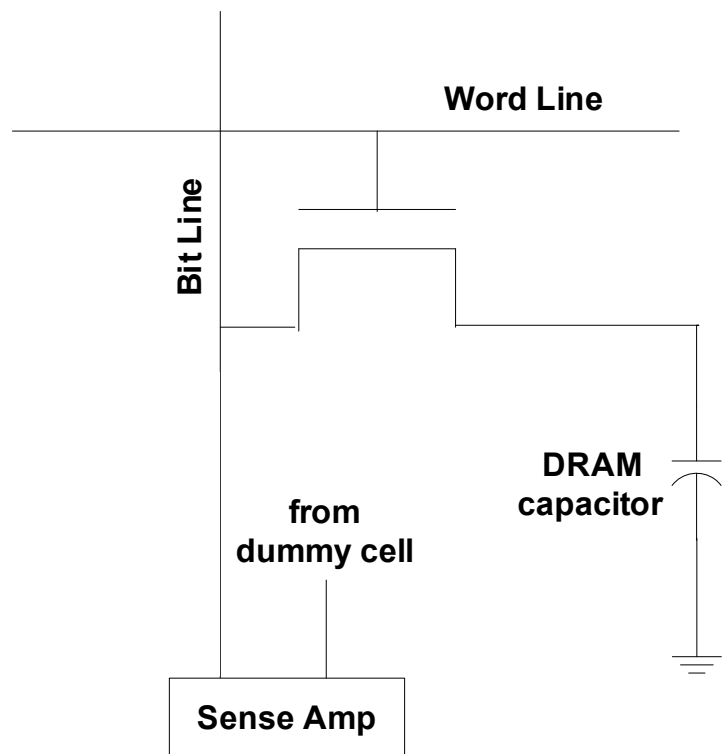
DRAM in Next-Generation Mobile Devices

- Feature-rich next-generation mobile devices
- Increase in memory requirements
- DRAM displacing SRAM
 - Samsung SGH i700 Triband Windows Smartphone : 32 MB DRAM
 - Motorola i930 cell phone has 32 MB DRAM


Dram Makers Prep Multichip Cell Phone Memories. ElectroSpec, 1/3/2003.
Innovative Mobile Products Require Early Access to Optimized DRAM Solutions. A Position Paper by Samsung Semiconductor, Inc., 2003.

- DRAM continuously drains battery even in standby
 - Periodic refresh needed to preserve stored information
- Need to reduce DRAM refresh power

DRAM – Cells

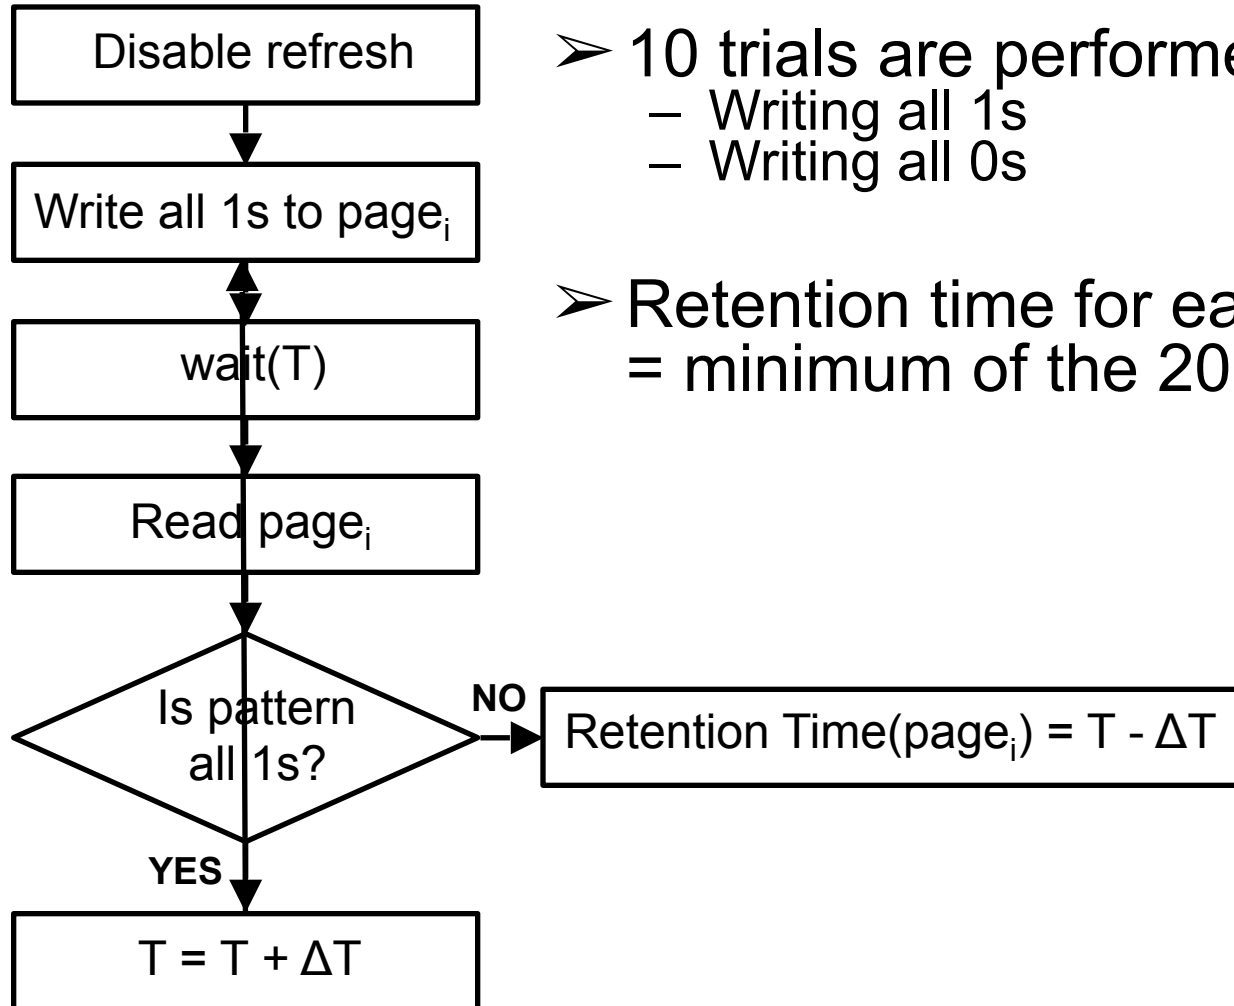


Retention Time Measurement



ISSI DRAM Chip
(IS42S16800A)
128 Mb

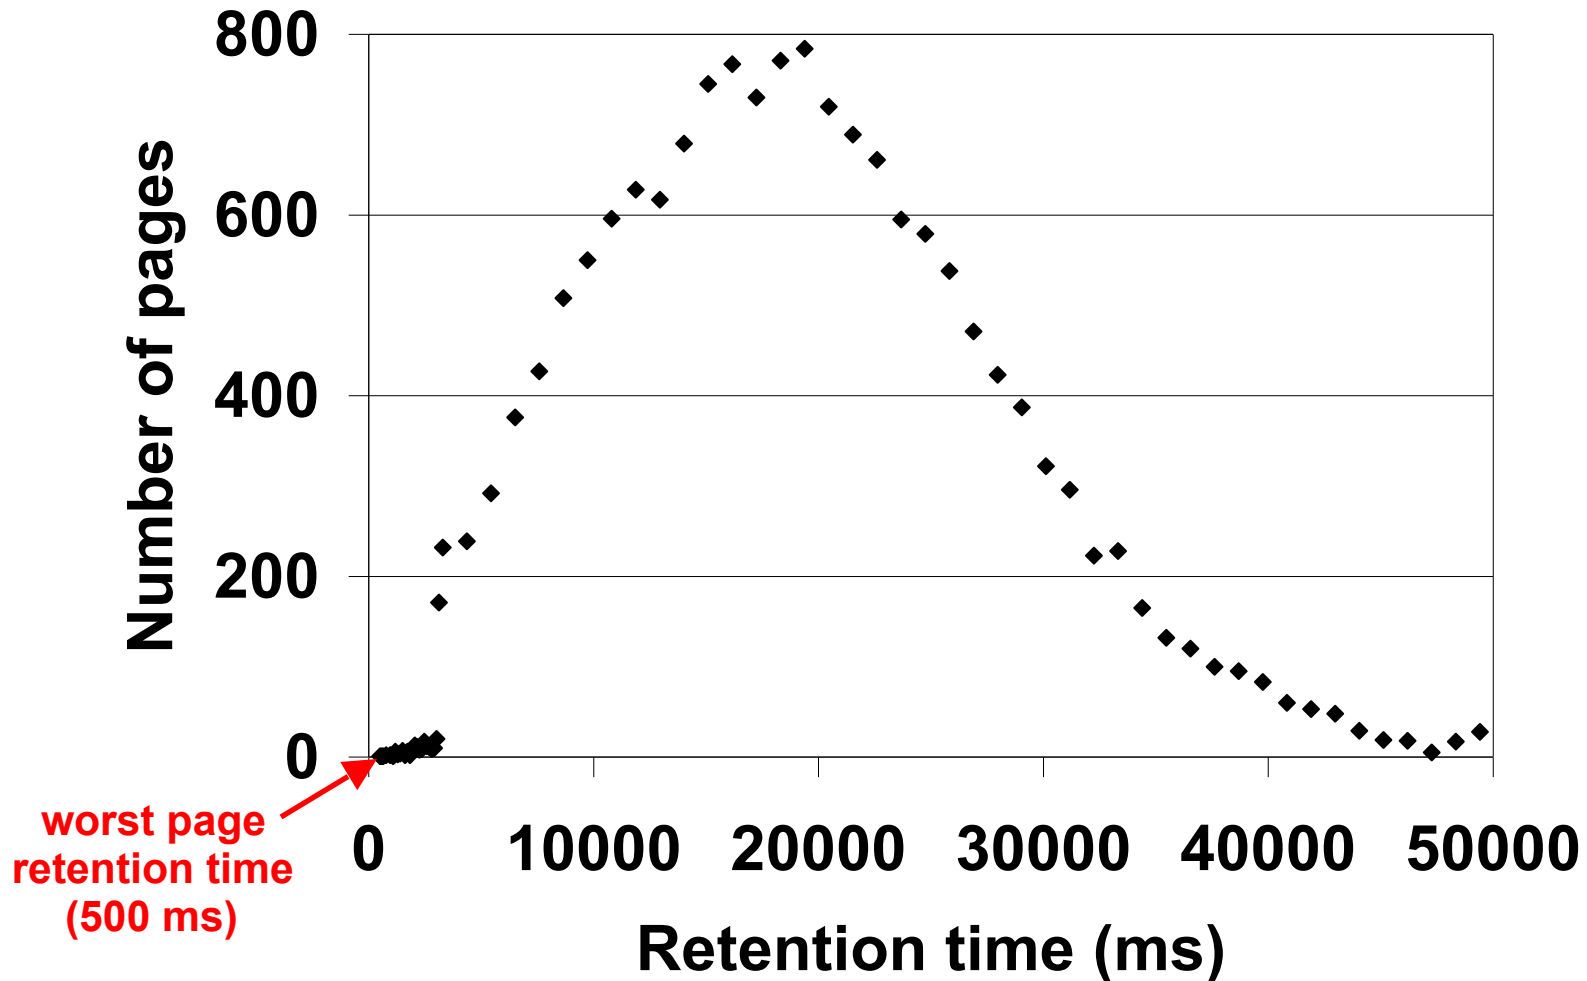
Retention Time Characterization Algorithm



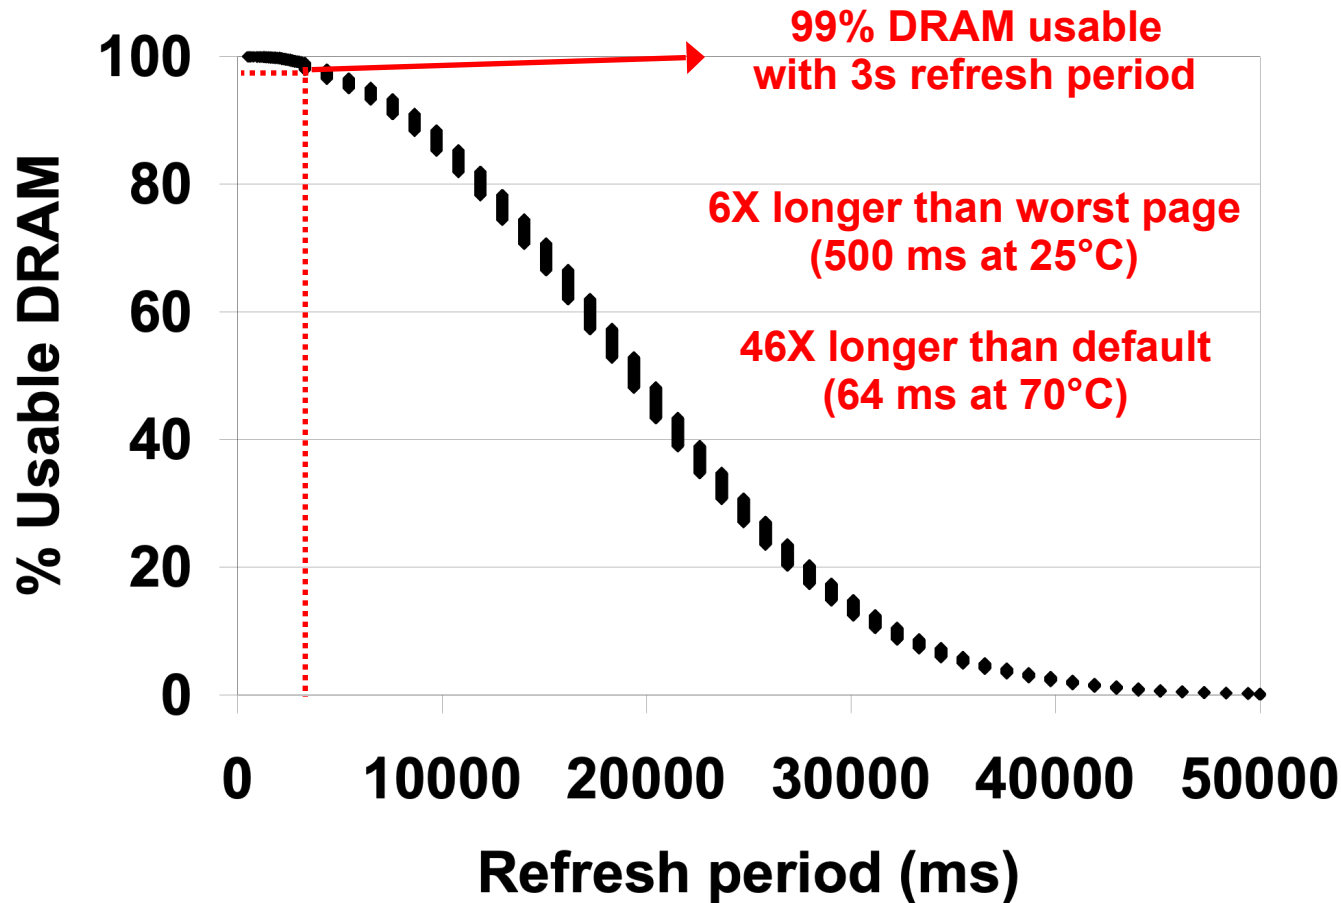
- 10 trials are performed for
 - Writing all 1s
 - Writing all 0s

- Retention time for each page = minimum of the 20 trials


DRAM – Retention Time Distribution



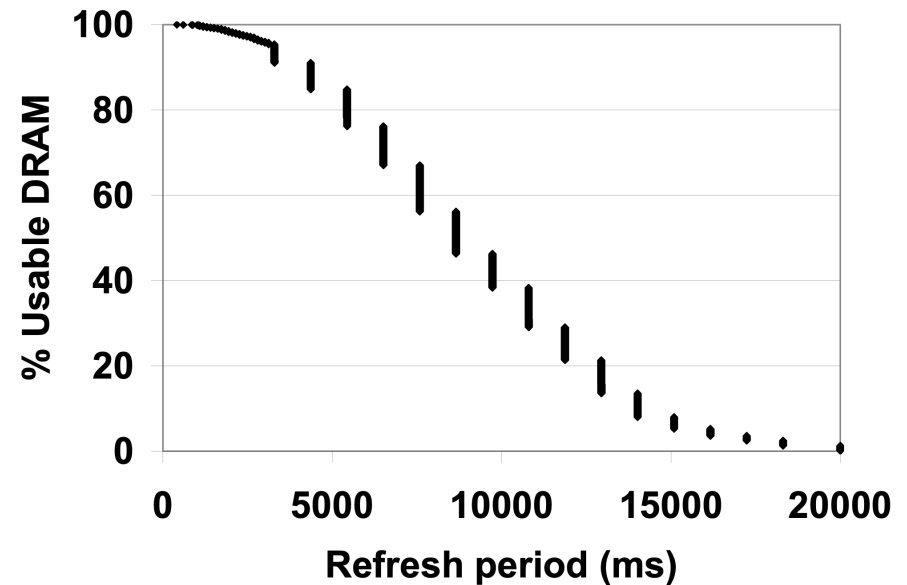
Cumulative Retention Time Distribution




DRAM – Retention Time Distribution 45°C




worst page
retention time
(320 ms)



DRAM – Retention Time Distribution 70°C



**worst page
retention time
(150 ms)**



DRAM Refresh Approaches

Three classes of hardware techniques:

- Conventional
 - Refresh period based on worst cell
 - Worst-case temperature
- Temperature-Compensated Refresh (TCR)
 - Refresh period based on worst cell
 - Dynamically adjusts for temperature
- Better than worst-case design
 - Multiple refresh periods (refresh period per block of cells)
 - Can be coupled with TCR

RAPID (Retention-Aware Placement in DRAM)

- Novel software-only approach

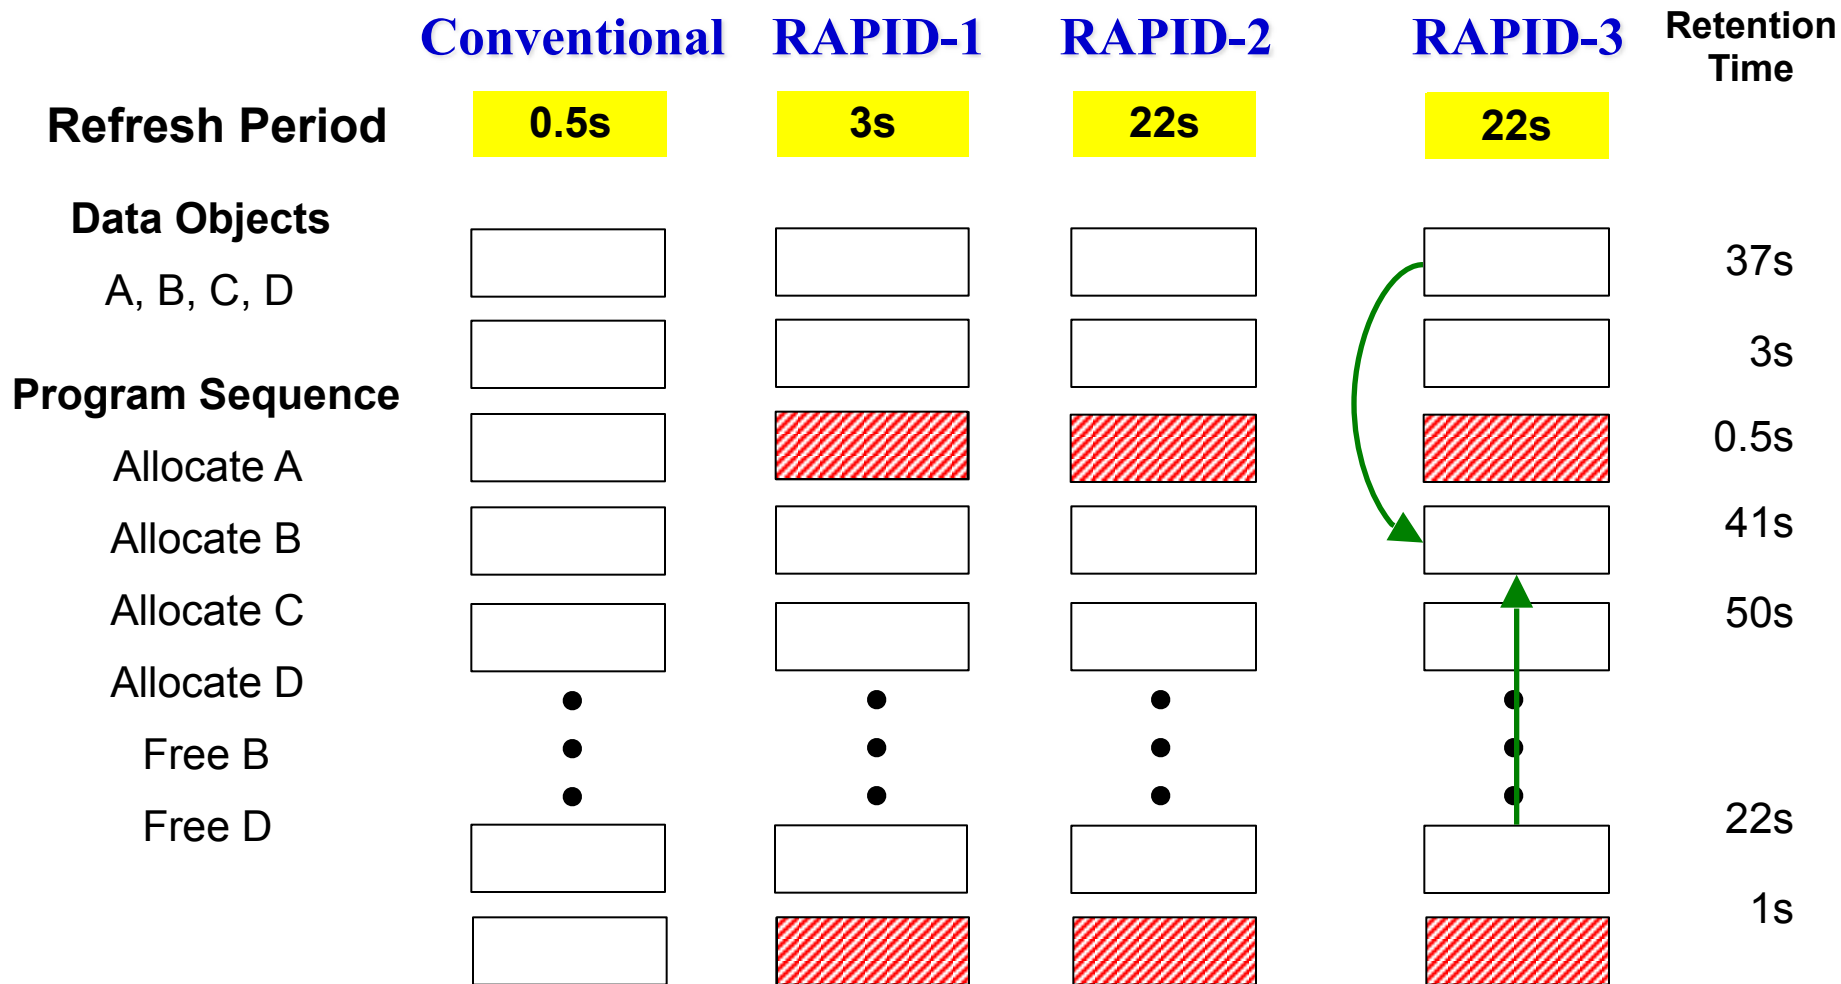
KEY IDEA:

- Allocate longer-retention pages before allocating shorter-retention pages
- Single refresh period selected based on shortest-retention page among populated pages, rather than shortest-retention page overall

NOTE:

- All pages are still refreshed
- Extended refresh period safe only for populated pages
- OK for overall correctness

RAPID



RAPID - Recap

- Novel software-only approach
 - Can exploit off-the-shelf DRAMs to reduce refresh power
- RAPID - 1
 - Eliminate outlier pages
- RAPID - 2
 - RAPID - 1 and
 - Allocate longer-retention pages before allocating shorter-retention pages
- RAPID - 3
 - RAPID - 2 and
 - Continuously reconsolidate data to longest-retention pages possible

Coupling RAPID with Off-the-shelf DRAMs

- RAPID requires a single refresh period that can be adjusted
- Refresh options in off-the-shelf DRAMs:

SELF-REFRESH

- DRAM issues refresh commands internally at a refresh period that is fixed or based on a particular temperature range
- Typically not programmable
- Low power
- Used in standby operation

AUTO-REFRESH

- Regularly timed refresh commands issued by memory controller
- Typically programmable
- High power
- Used in active mode

Coupling RAPID with Off-the-shelf DRAMs (cont.)

Neither refresh option is fully satisfactory

- Self-refresh not programmable
- Auto-refresh programmable, but power-inefficient


Key observation

- DRAMs allow enabling/disabling self-refresh via configuration register

Coupling RAPID with Off-the-shelf DRAMs (cont.)

Solution

- Set up a periodic timer interrupt (INT1)
 - Period = RAPID refresh period
- Set up another interrupt (INT2)
 - Triggers 64 ms (self-refresh period) after INT1
- INT1 triggered:
 - Enable self-refresh
- INT2 triggered:
 - Disable self-refresh

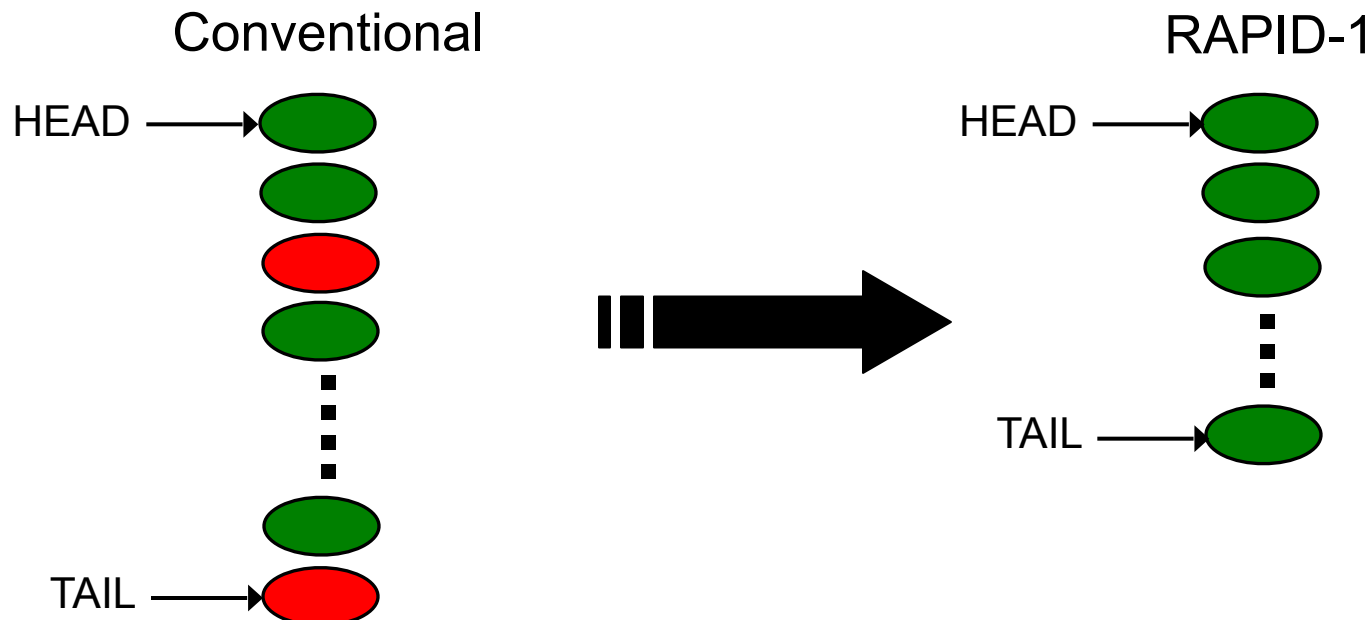


RAPID – Implementation

- Modifications to routines which allocate and deallocate physical pages in memory

INACTIVE LIST

(free-list of inactive physical pages)

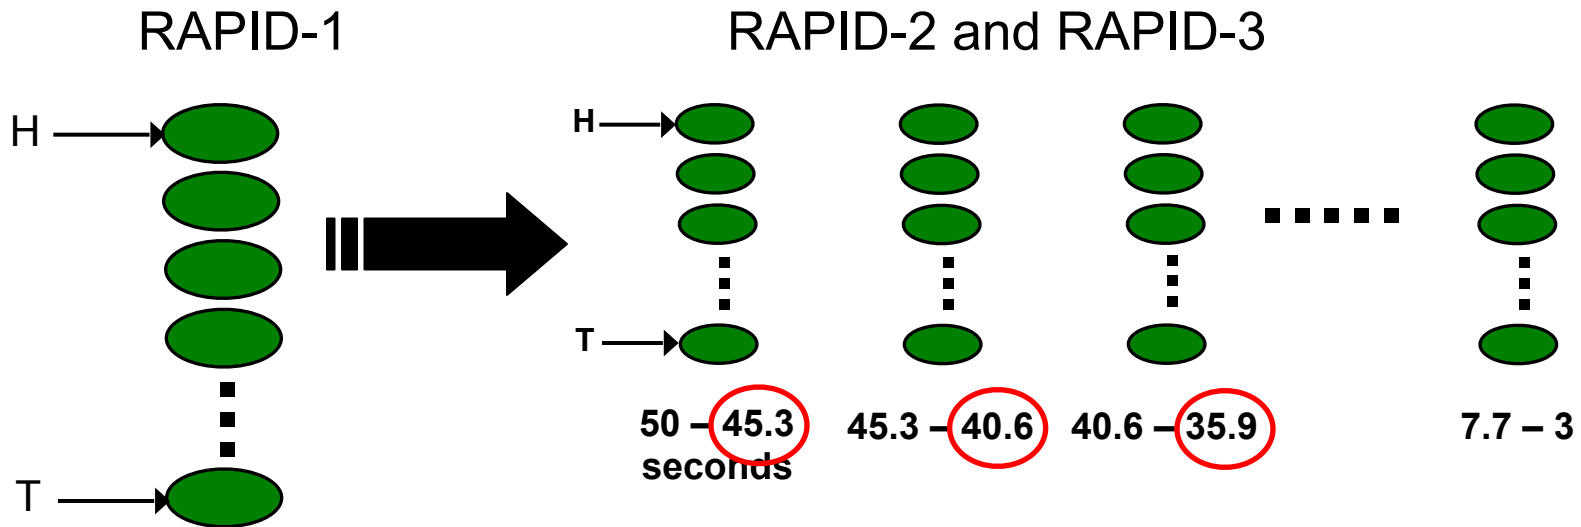


RAPID-1: 1% of outlier pages are excluded from Inactive List

RAPID – Implementation (cont.)

INACTIVE LIST

(free-list of inactive physical pages)



Refresh Period = 35.9s

RAPID-2 and RAPID-3

Single Inactive List transformed into Multiple Inactive Lists

Related Work

➤ Characterization

[Hamamoto et al. 1998, IEEE Tran Electron Devices] On Retention Time Distribution of DRAM

➤ Dual Period Refresh

[Yanagisawa 1988, US Patent] Semiconductor Memory

➤ Multiperiod Refresh

[Kim and Papaefthymiou 2001, 2003, IEEE Tran VLSI Systems] Block-Based Multiperiod Refresh

➤ Multiperiod Refresh + Selective

[Ohsawa et al. 1998, ISPLED] Optimizing DRAM refresh Count for Merged DRAM/Logic LSIs

➤ Placement + Single Refresh Period + Selective

[Kai et al. 2002, US Patent] Semiconductor Circuit and Method of Controlling the Same

[Le et al. 2005, US Patent] Bank Address Mapping According to Bank Retention Time in DRAM

[Kawasaki et al. 2005, US Patent] DRAM and Refresh method thereof



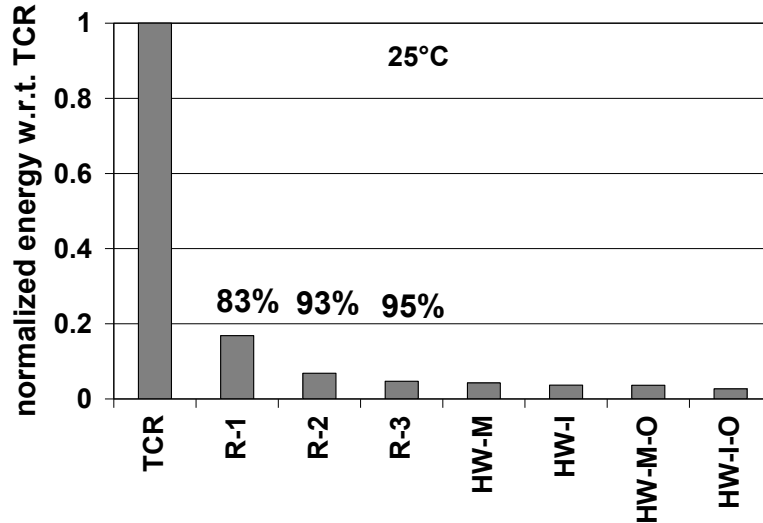
Evaluation Methodology

- Active Mode (5%), Idle Mode (95%)
- Simulated timeline = 24 hours
- Divide 24-hour timeline into time slices
- Probability [time slice = active period] = 5%
- Probability [time slice = idle period] = 95%
- Inject random number of allocations/frees in active period
 - Vary DRAM utilization across timeline

Custom Hardware Techniques


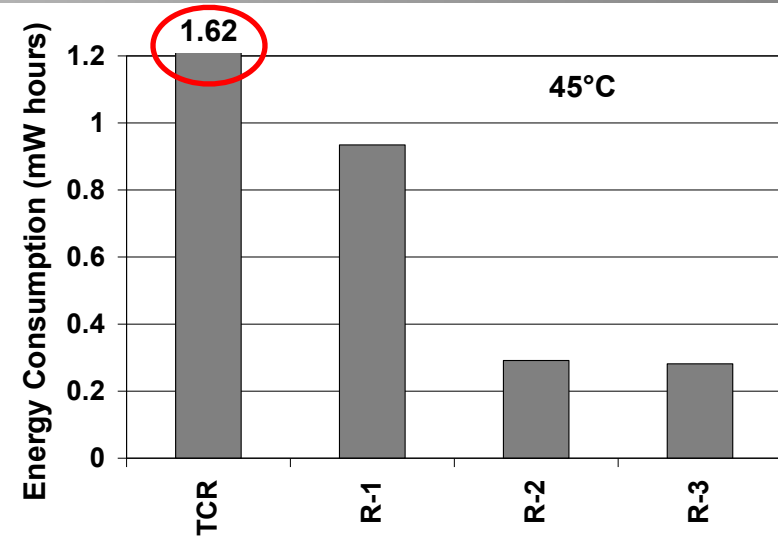

- HW-Multiperiod (HW-M)
 - Each DRAM page is refreshed at a tailored refresh period, that is a multiple of shortest refresh period
- HW-Multiperiod-Occupied (HW-M-O)
 - Same as HW-Multiperiod
 - Refresh only currently occupied pages
- HW-Ideal (HW-I)
 - Each DRAM page is refreshed at its own tailored refresh period
- HW-Ideal-Occupied (HW-I-O)
 - Same as HW-Ideal
 - Refresh only currently occupied pages

RAPID – Refresh Energy Savings






- 83-95% energy savings at 25°C
- Similar energy savings across all temperatures
- Nearly as effective as custom hardware techniques

RAPID vs. TCR



- Worst-case non-temperature adjusted R-1 yields same or better energy than TCR
- R-1 simpler, cost-effective alternative to TCR
- R-2 and R-3 yield much higher energy savings than TCR

Refresh Energy – Different DRAM Utilizations



- R-1, HW-M and HW-I constant
- R-2 and R-3 yield more energy savings than HW-M and HW-I as utilization decreases
- R-2 and R-3 approach energy of HW-I-O
- HW-I-O not retention-aware

Summary – RAPID

- RAPID reduces refresh power to vanishingly small levels
 - Quasi-non-volatile DRAM
- Software-only technique
 - No custom DRAM support required
 - Can exploit off-the-shelf DRAM
- Approaches energy levels of idealized techniques, which require hardware support