DATE 14
Design, Automation & Test in Europe
24-28 March, 2014 - Dresden, Germany
The European Event for Electronic System Design & Test

Exploiting Expendable Process-Margins in DRAMs for Run-Time Performance Optimization

Karthik Chandrasekar [1], Sven Goossens [2], Christian Weis [3], Martijn Koedam [2], Benny Akesson [4], Norbert Wehn [3], and Kees Goossens [2]

[1] TU Delft, The Netherlands
[2] Eindhoven University of Technology, The Netherlands
[3] TU Kaiserslautern, Germany
[4] Czech Technical University, Czech Republic

COBRA – CA104
NEST
621429 EMC2
621353 DEWI

DRAM Access Latencies

- DRAM access times significantly influence computing systems performance
  - "Memory Wall" problem
- Recent developments in DRAM design improve memory bandwidth
  - Higher clock frequencies (DDR4), wider interfaces (WIDEIO), new architectures (hybrid memory cube)
- Latency reduction remains an open issue

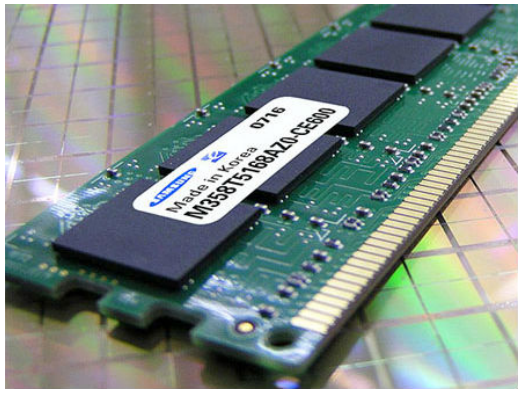

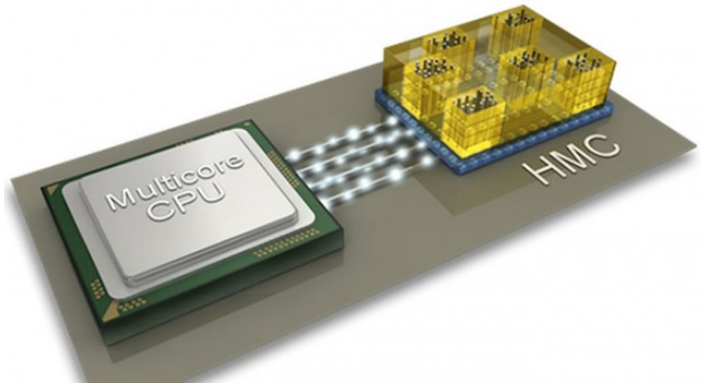

Images: extremetech.com, automatiseringgids.nl

Sources of Latency

- Access latency is a result of the DRAM timing constraints, i.e.
  - required minimum number of cycles between two memory commands
- Specified by DRAM manufacturers. Consist of:
  - Actual analog delays + pipeline stages
  - Margins accounting for **PVT variation**

- **P: process variation**
  - Not all devices are created equal. Random variations in device parameters like channel mobility / length, oxide thickness
- **V: voltage variation**
  - Lower supply voltage → larger delay
    Devices needs to work within specified voltage range (1.425 – 1.575 V)
- **T: temperature variation**
  - Higher temperatures → larger delays. Ambient + self heating
    Commercial DRAMs should work correctly at temperatures up to 85°C

Latency Margins (1)

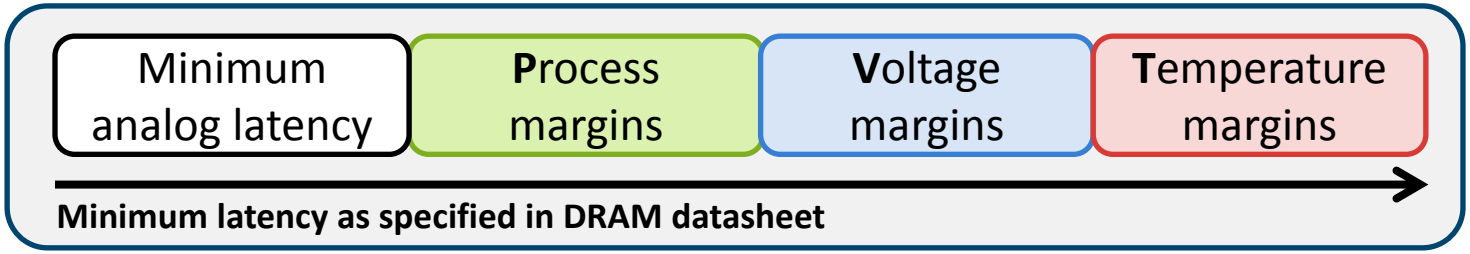

- Actual analog latencies vary significantly due to process variation:
  - Figure shows probability density function of latencies of a DDR3-800 device
  - (Using Monte-Carlo simulations on NGSPICE DRAM cross-section for 5 different timings)
  - Datasheet measures are given in brackets
- **Datasheet has to cover the worst-case variation, and is often overly pessimistic**

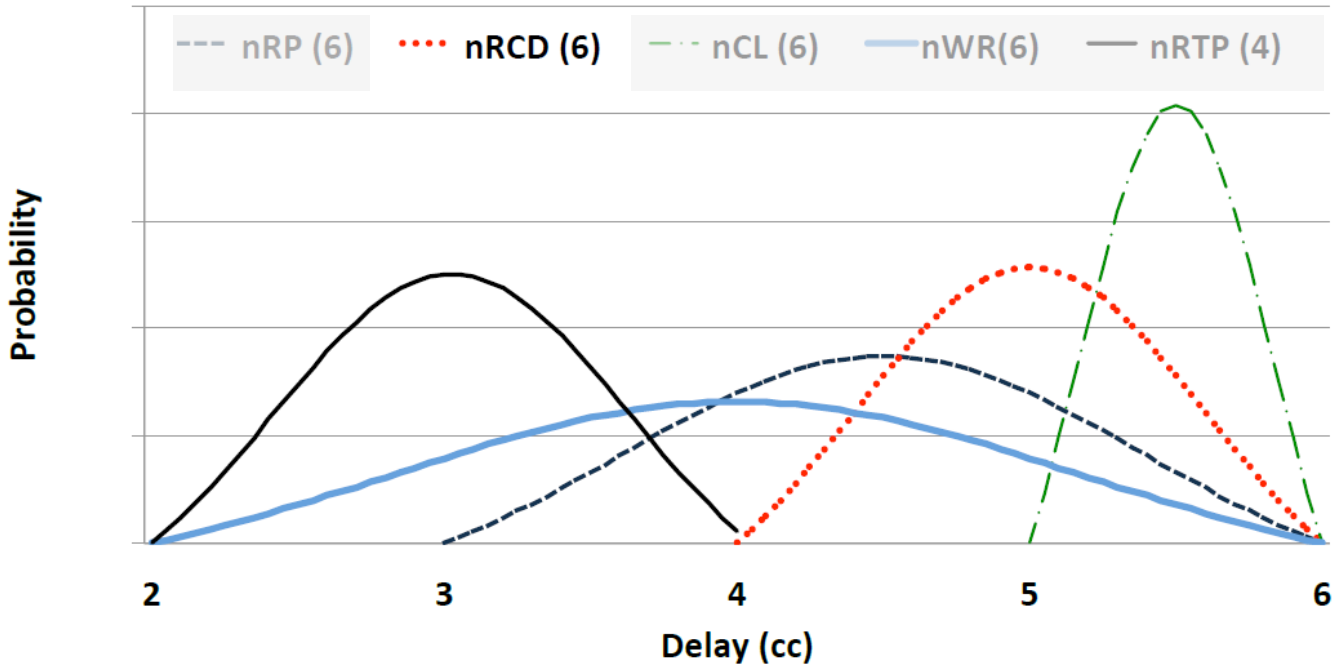

Latency Margins (2)

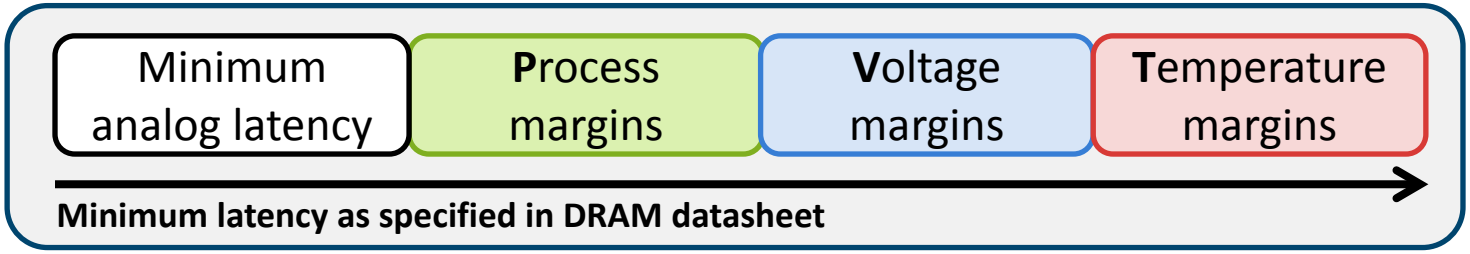

- Manufacturer only give 1 datasheet number as the timing constraint
- Most devices can perform better than specified
- **Process variation is fixed after manufacturing, and can thus be measured**

Contributions

- Generic post-manufacturing performance characterization methodology
  - Identify actual achievable performance at run-time
  - Valid under worst-case operating conditions (minimum voltage, maximum temperature)
- Removes the excessive process margins on a per device basis
- Demonstrated on 48 DDR3 devices (12 identical DIMMs, 1 vendor)
  - Derive actual device delays
  - Verify functional correctness under worst-case operating conditions

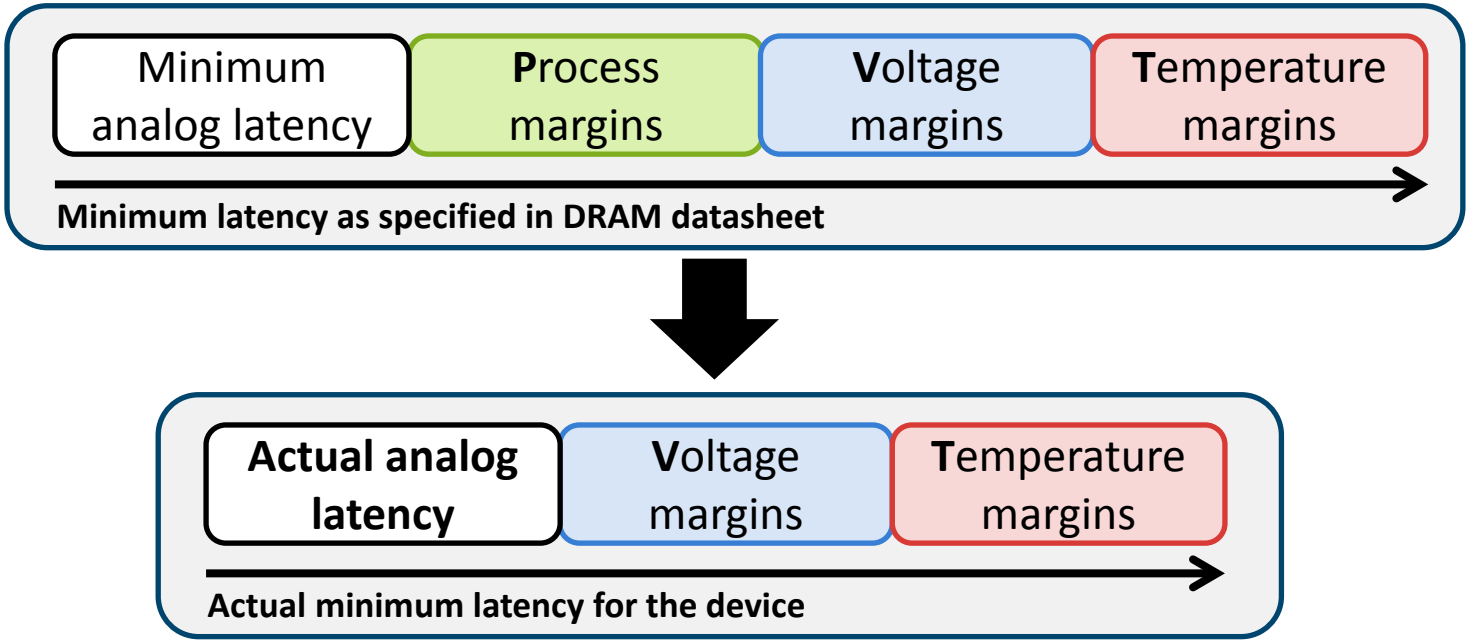

Outline

Introduction

DRAM Background

Performance Characterization Methodology

Experiments / Results

Conclusions

DRAM Background

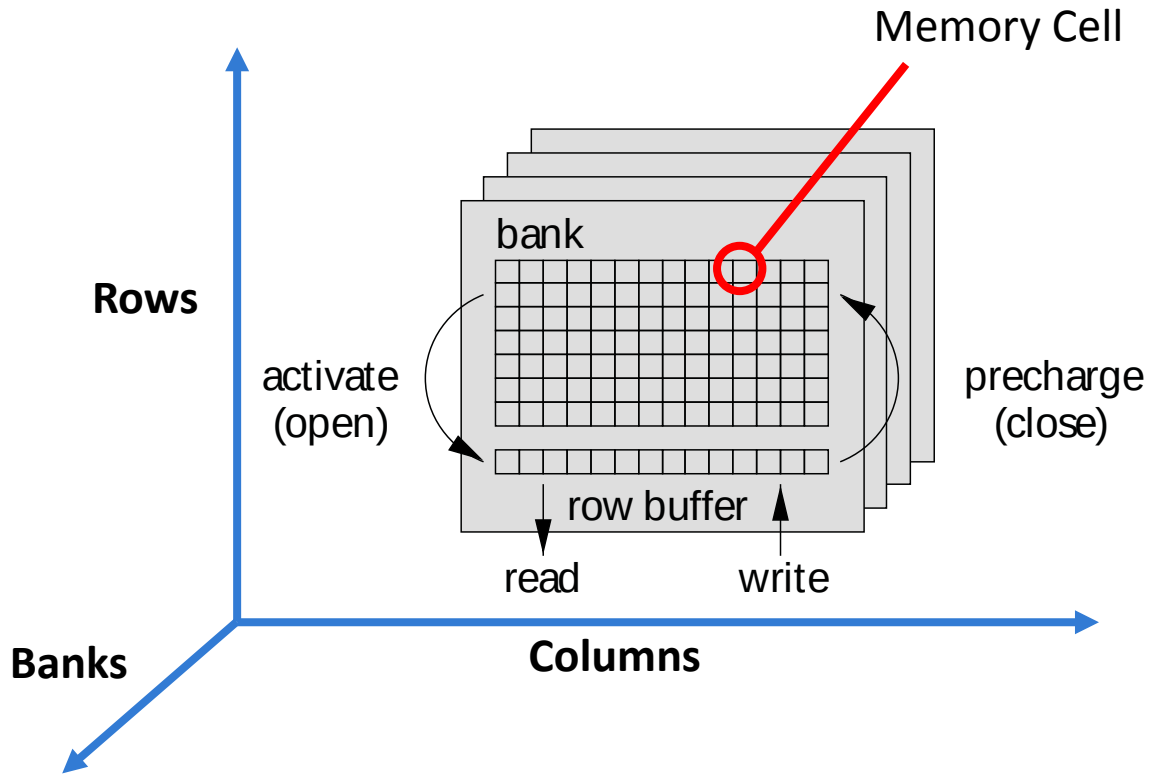

6 commands:

- activate (ACT)
- precharge (PRE)
- read (RD)
- write (WR)
- refresh (REF)
- NOP

- Each bank has **row buffers** (intermediary for reading or writing into the memory cell)
- A row has to be **activated** before it is accessible. Moves data from cells to row buffer
- To open a different row, the old one has to be closed by **precharging**
- **Row buffers are shared between sets of successive rows**
- Timing constraints enforce a minimum distance between the commands

Timing Constraints

- Commands issued to the DRAM must satisfy certain timing constraints, e.g.:

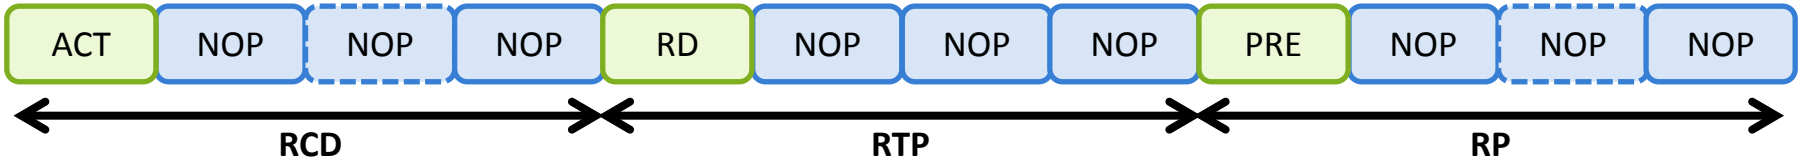

- For example, for a Micron DDR3-800 device (fast-core), these constraints are:

Constraint	Description (minimum time between)	time (cc)
RCD	Activating to read/write command	6
RP	Precharge to activate	6
RTP	Read to precharge	4
CL	Read to first data	6

- Devices are speed binned based on the CL, RCD and RP parameters:

Speed bin	Frequency (MHz)	CL- RCD - RP [ns]
800 [Slow-Core]	400	15 - 15 - 15
800 [Fast-Core]	400	12.5 - 12.5 - 12.5
1066 [Slow-Core]	533	15 - 15 - 15
1066 [Fast-Core]	533	13.125 - 13.125 - 13.125

- Devices failing the operating frequency or **one** of the timing requirement are down-binned → variability in the actual performance is expected.

Outline

Introduction

DRAM Background

Performance Characterization Methodology

Experiments / Results

Conclusions

Performance Characterization Methodology

- Goal: determine actual achievable performance under worst-case operating conditions
  - **Basic idea: Iteratively reducing the timings and check function correctness**
    - write a known data pattern into the memory
    - read it back to verify it was correctly written and read

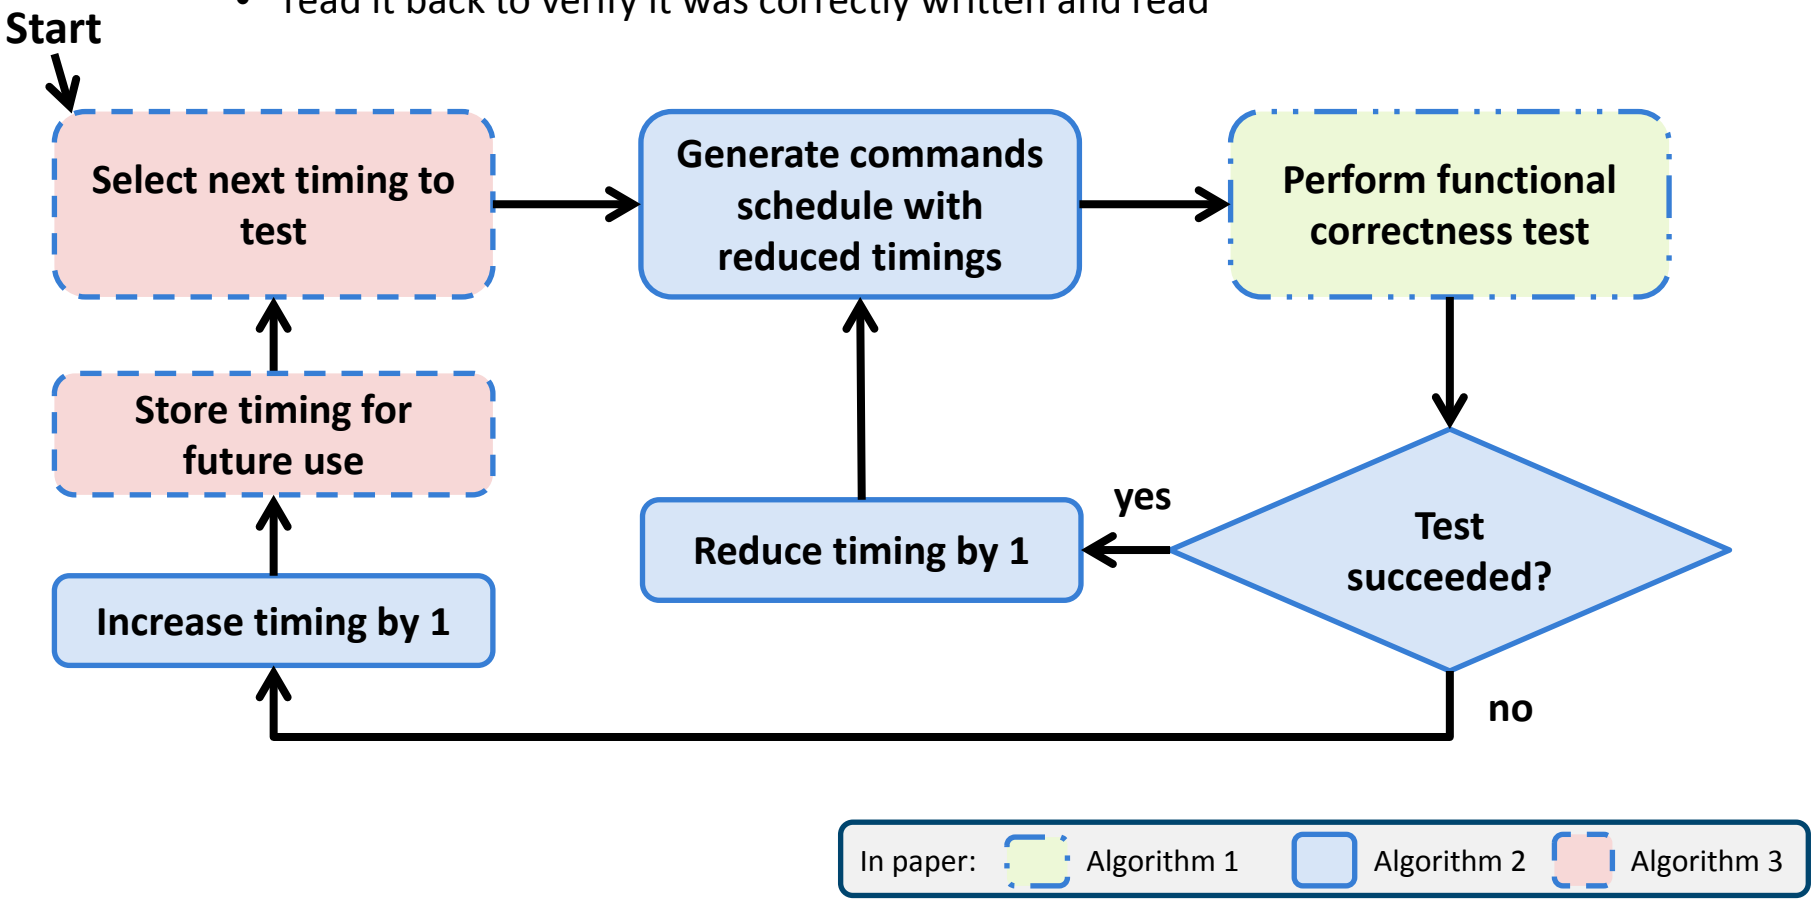

Functional Testing Approach

1. Verify DRAM operations are complete, i.e. that data is stored in memory array (and not for example retained only in the row buffer)
   → **Use the fact that row-buffers are shared between adjacent rows**
   Writing inverted data into them

2. Test for cross talk and stuck-at faults, interference between memory cells
   → **Test data has neighboring bits flipped, or all bits set / unset**

3. Test the entire DRAM (All columns, rows and banks bits, all decoder logic, etc.)
   → **Iterate over all addresses**

4. Assure functional correctness under worst-case operating conditions
   → **Identify and retain voltage and temperature margins**

Functional Correctness Test

For all bank, row pairs, columns, do:

For data patterns in {0xA5A5, 0x5A5A, 0xFFFF, 0x0000}

Step 1: write inverted patterns in adjacent rows

Step 2: read data patterns and compare to reference

Report failure if read data does not match reference

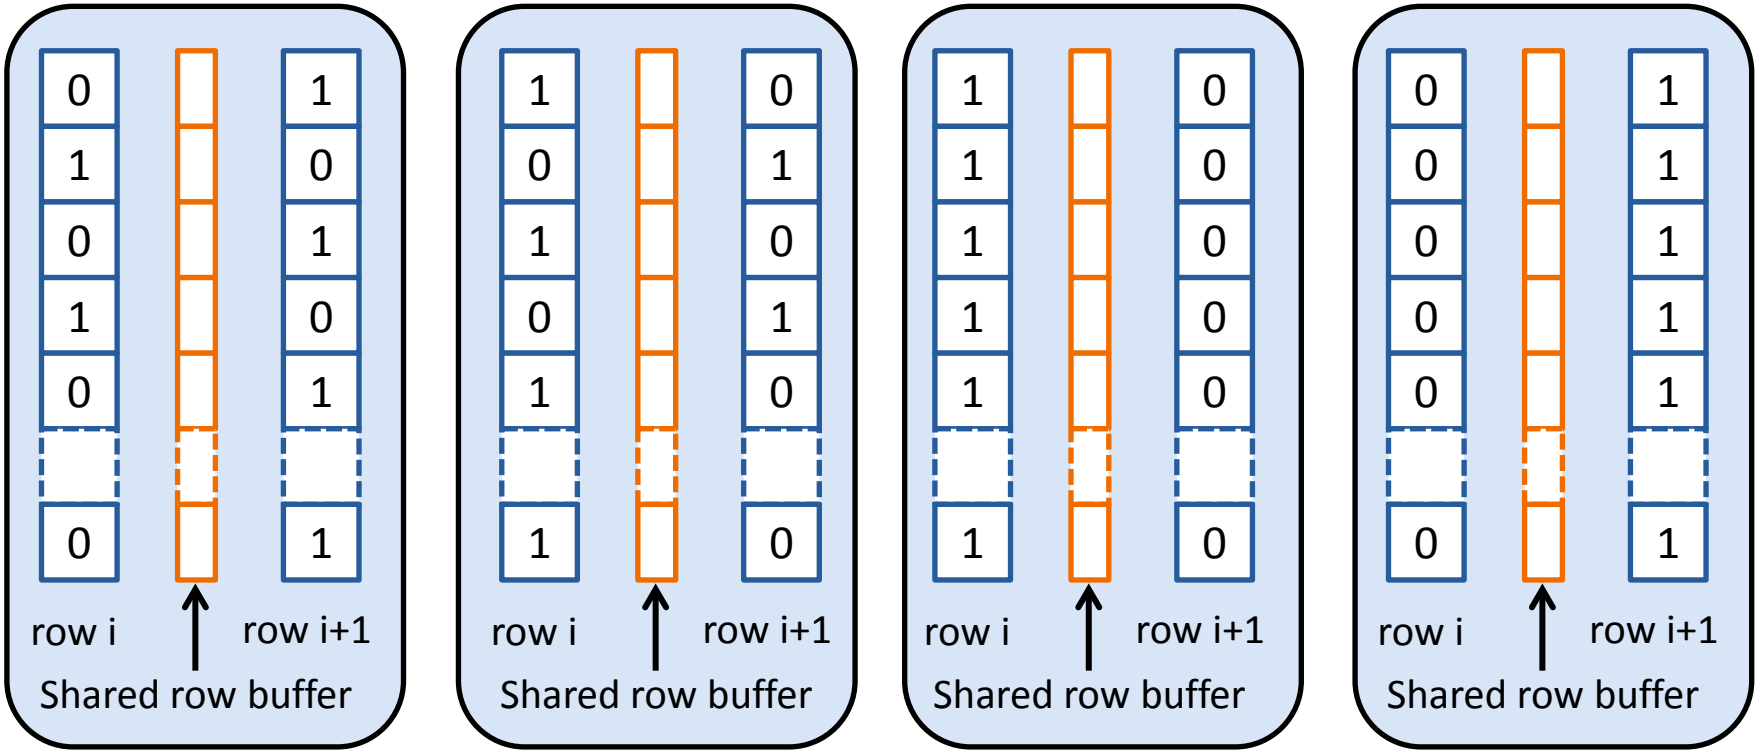

Functional Testing Approach

1. Verify DRAM operations are complete, i.e. that data is stored in memory array (and not for example retained only in the row buffer)
   → **Use the fact that row-buffers are shared between adjacent rows Writing inverted data into them**

2. Test for cross talk and stuck-at faults, interference between memory cells
   → **Test data has neighboring bits flipped, or all bits set / unset**

3. Test the entire DRAM (All columns, rows and banks bits, all decoder logic, etc.)
   → **Iterate over all addresses**

4. Assure functional correctness under worst-case operating conditions
   → **Identify and retain voltage and temperature margins**

Command Generation

- For each combination of timing parameters, a functional test is executed
- Memory controller uses modified read and write schedules. Example of read schedule:

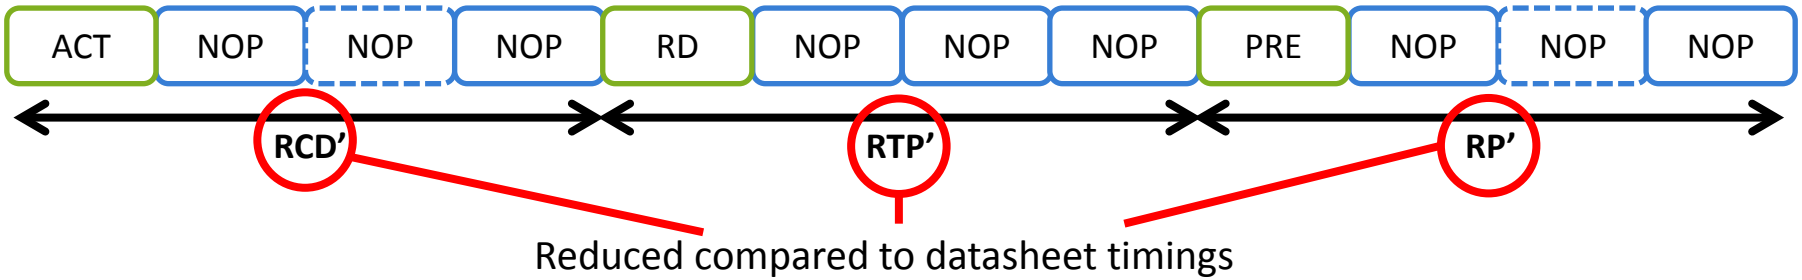

Constraint	Description (minimum time between)	Source
RCD'	Activating to read/write command	test driven
RP'	Precharge to activate	test driven
RTP'	Read to precharge	test driven
WR'	Write recovery to precharge	test driven

Characterization Algorithm

- **We reduce RCD, RP, RTP and WR**

Constraint	Description (minimum time between)
RCD	Activating to read/write command
RP	Precharge to activate
RTP	Read to precharge
WR	Write recovery to precharge

- **Don't change:**
  - Refresh related constraints (REFI, RFC), and
  - constraints that limit maximum supply current (RRD, FAW), and
  - constraints connected to pipeline stages in the DRAM (CL, CWL)

- **(paper describes the order in which constraints are minimized)**

Voltage and Temperature Compensation

- Voltage and temperature margins must be re-applied if algorithm is not executed under worst-case operating conditions
- Derive required voltage/ temperature margins with NGSPICE DRAM model
  - Find required timings at nominal and worst-case conditions, resp.
  - % difference is required timing margin for timings found at nominal operating conditions
  - Conservatively round timing up to an integer number of clock cycles once applied

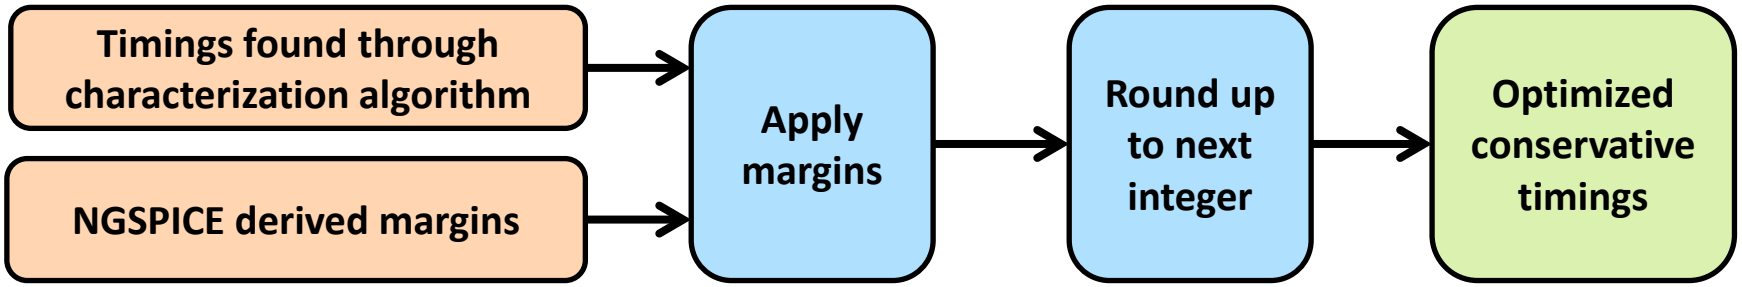

Conservative Margin Compensations for a *DDR3-800, 1 gigabit* device:

Constraint	Nominal (1.5V, +27°C) [ns]	Worst case (1.425V, +85°C) [ns]	Difference (%)
WR	4.69	5.38	14.7
RP	5.88	6.68	13.5
RCD	8.84	10.02	13.3
RTP	8.84	10.66	20.6

Outline

Introduction

DRAM Background

Performance Characterization Methodology

Experiments / Results

Conclusions

Experimental Setup

- Tested methodology on 12 512MB Micron DDR3-1066 DIMMs
  - Each DIMM has 4 1Gb x16 DDR3 devices (chips) → 48 devices tested
- Using Xilinx ML605 FPGA board, with a MicroBlaze processor to drive the test
  - Hardware accelerator for functional correctness test
  - Reconfigurable memory controller
- 4 seconds / test (4 GB / test on accelerator to memory controller interface)
- Total test duration depends on how quickly devices fail (typically completes in < 2 min)

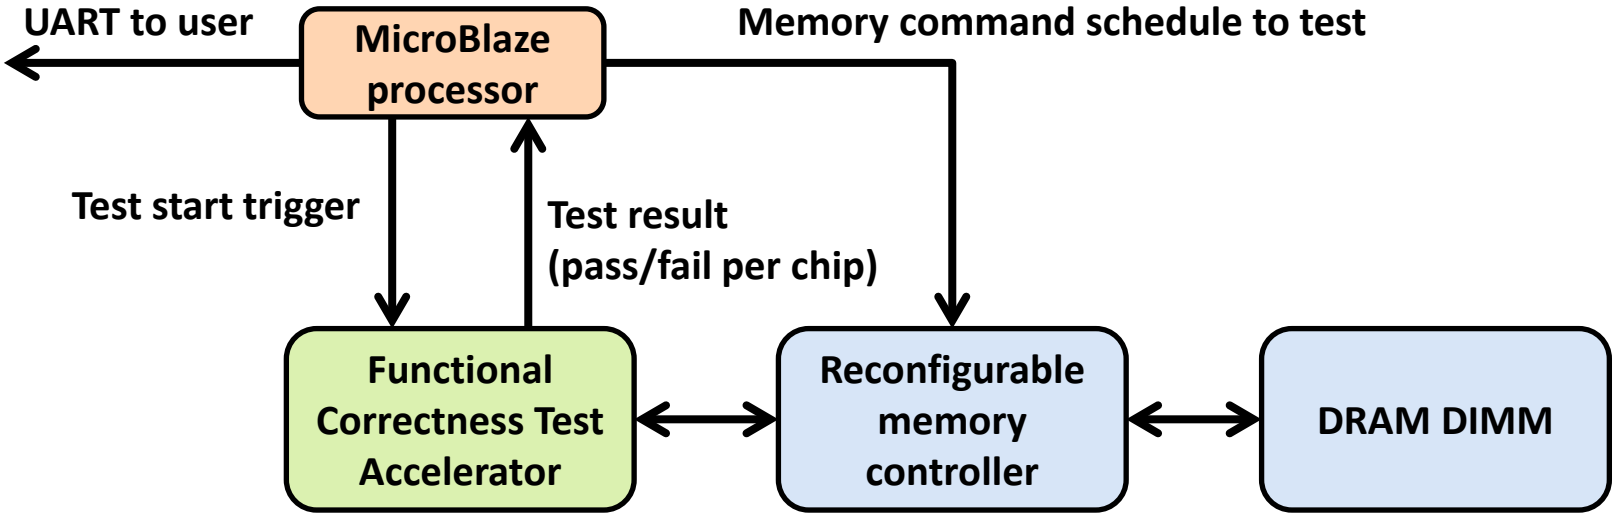

Results at Nominal Conditions

- Results grouped per DIMM
- 1 line for each timing constraint under test

Rightmost column: datasheet value

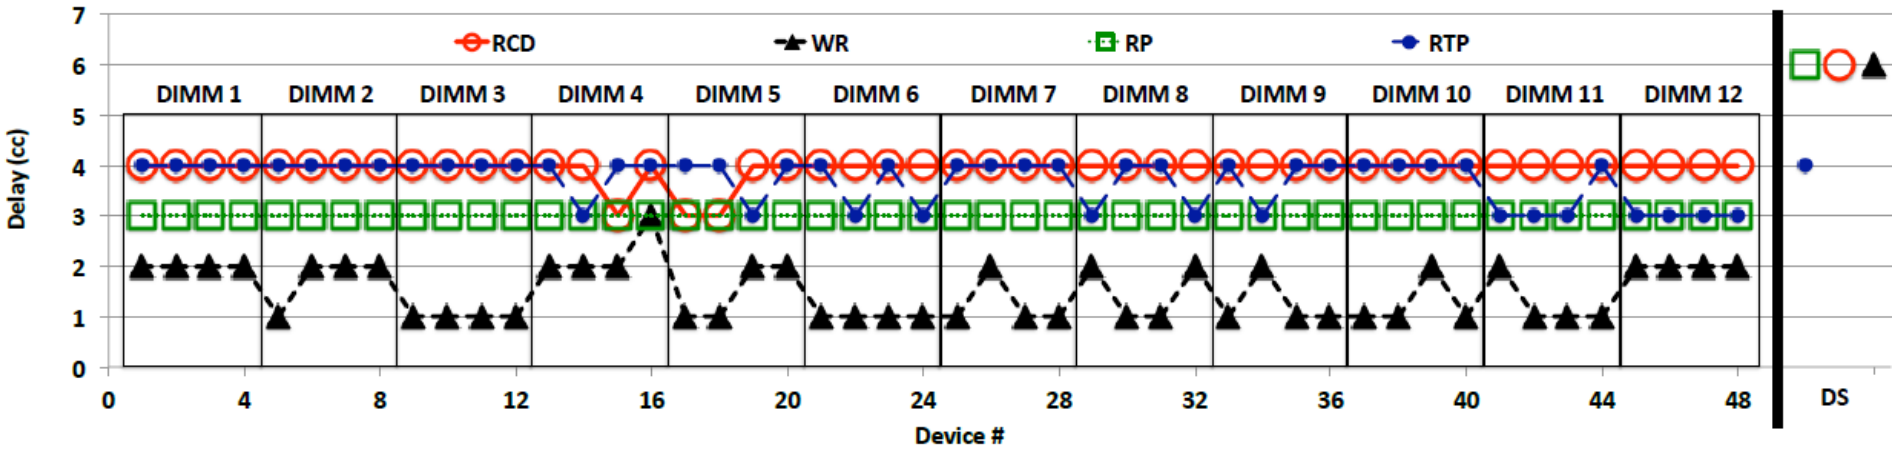

- Variation observed across DIMMs and across chips on a DIMM
- Datasheet values are very pessimistic
- (Still need to apply voltage and temperature margins)

Verifying Compensated Timings (1)

- Emulate worst-case operating conditions:
  - **Temperature:** Increase to 85°C heating the DRAM module with a 4 W **peltier** element
    - Peltier output power is controlled by an Arduino board
  - **Voltage:** Supply voltage lowered with JET-5466 SODDR3 extender board, to 1.42 V

- We use the fastest DIMM from the first experiment (= most removed excessive margins)
- Stress tested with functional correctness test (running continuously)

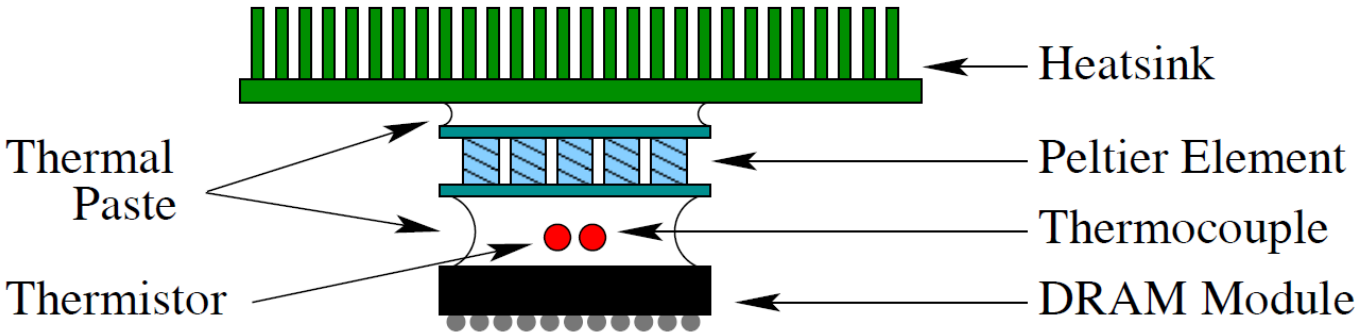

Verifying Compensated Timings (2)

- **First experiment: Uncompensated timings:**
  - Test fails immediately (even though it worked at nominal conditions)
- **Second experiment: Timings compensated with temperature / voltage margins:**
  - Un-interrupted test run for over 4 hours
  - Transferred more than 16 TB of data, no failure observed
- **Current drawn compared to datasheet:**

Current type:	Measured [mA]	Datasheet [mA]
IDD0	300	360
IDD1	524	584

  - Even with optimized timings, current is still within specs
  - Current number in datasheets account for process variation, hence the generous margins

- **Bandwidth/latency and energy benefits:**

	Read request	Write request
Schedule length reduction	**33 %**	**25.9 %**
Energy reduction (using power model)	**17.7 %**	**15.4 %**

Outline

Introduction

DRAM Background

Performance Characterization Methodology

Experiments / Results

Conclusions

Conclusions

- **Proposed a generic, post-manufacturing DRAM characterization methodology**
- **Identify and prune excessive process-margins on a per device basis**
- **Derive required voltage and temperature margins to compensate the pruned timings for worst-case operating conditions**
- **Evaluated methodology for 48 DDR3 devices**
- **Verified functional correctness under worst-case conditions**
- **Bandwidth (+33%) and energy benefits (-17%) have been demonstrated**

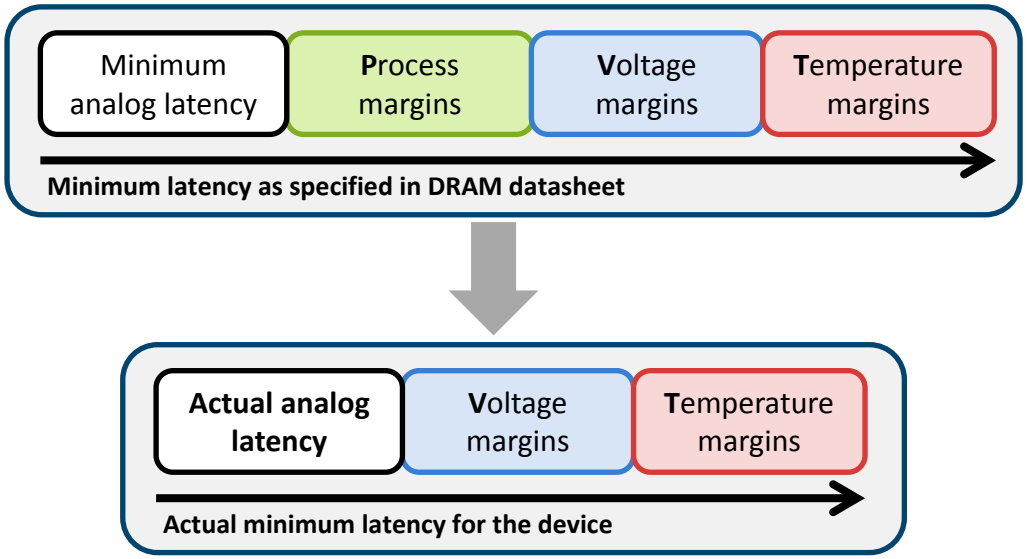

We appreciate the help of our fellow researchers at TU Delft, TU Eindhoven and University of York who sacrificed some of their ML605 time to run our algorithms.

The test bitstream used in our experiments will be made available at www.drampower.info