



Design, Automation & Test in Europe
12-16 March, 2012 - Dresden, Germany

The European Event for Electronic
System Design & Test

DRAM Selection and Configuration for Real-Time Mobile Systems

Manil Dev Gomony, Christian Weis, Benny Akesson,
Norbert Wehn and Kees Goossens

Eindhoven University of Technology, The Netherlands
University of Kaiserslautern, Germany




Outline

- **Introduction**
- **Our approach**
- **Analysis results**
- **Proposed methodology**
- **Conclusions**

Mobile platforms


- **Multi-processor platforms for mobile systems**
 - Real-time and non-real-time applications
 - **Strict power budget**
 - **Dynamic Random Access Memory (DRAM) is shared**



DRAM subsystem requirements

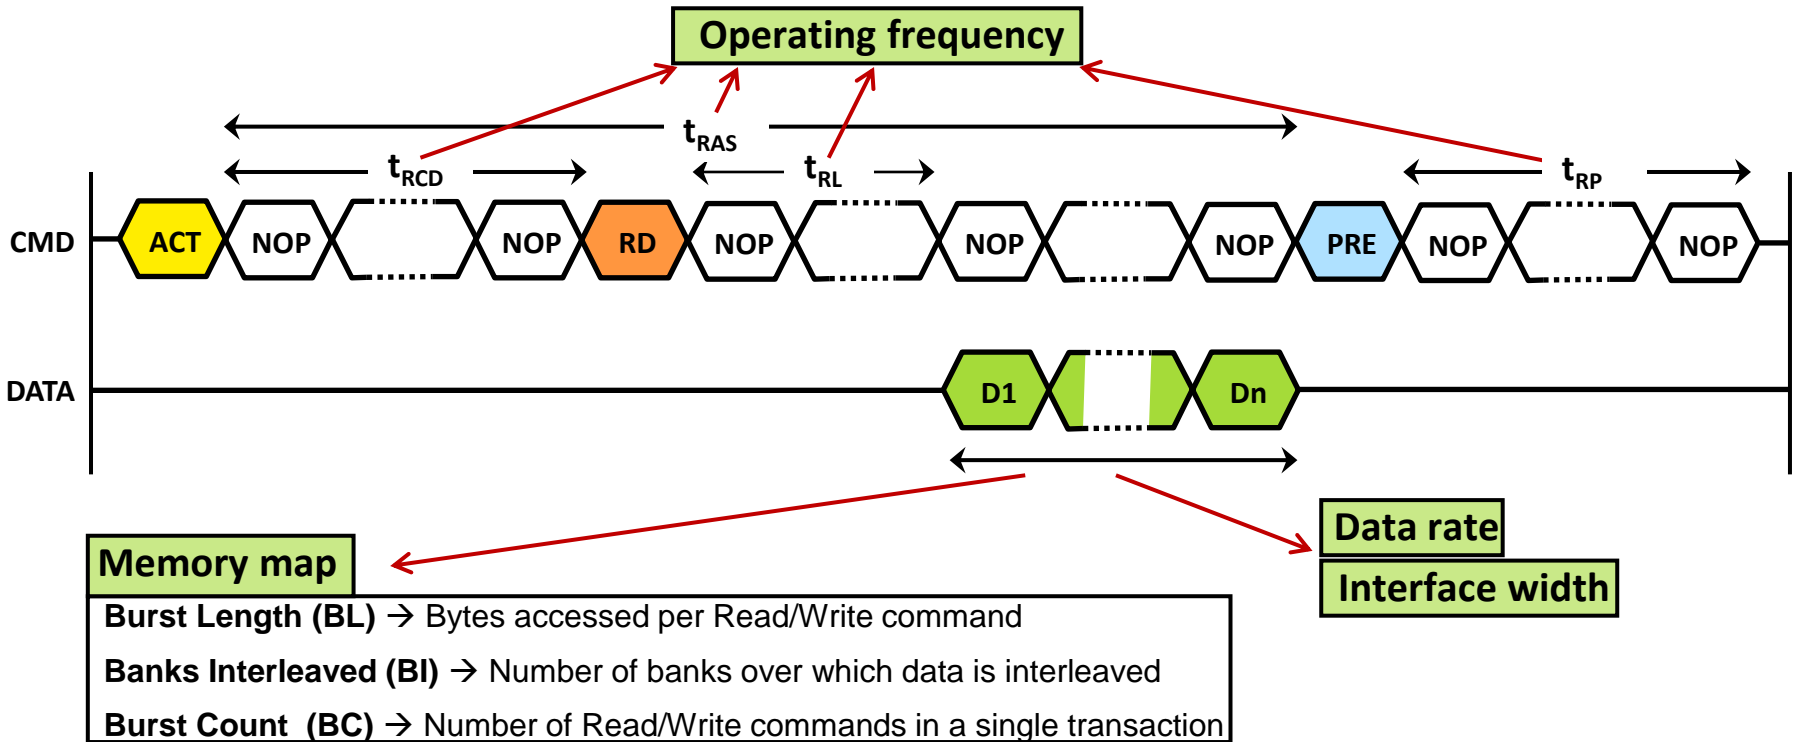
- **DRAM subsystem in mobile platforms must:**
 - **Guarantee bounds on bandwidth to real-time applications**
 - **Real-time memory controllers**
 - **Provide best average-case performance to non-real-time applications**
 - **Meet the power budget**

DRAM overview



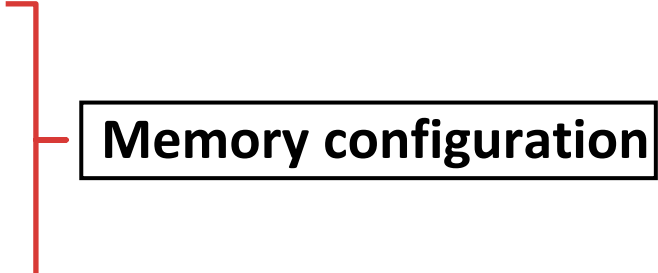
- Data is stored in storage cells consisting of a capacitor-transistor pair
- Storage cells are arranged to form a memory array
- Memory array and row buffer constitute a bank
- Data is accessed by issuing memory commands

DRAM overview



$$\text{Memory efficiency} = \frac{\text{Clock cycles containing useful data}}{\text{Total clock cycles}}$$


Mobile DRAMs

- **Low-Power Double Data Rate (LPDDR)**
 - **Low-Power Double Data Rate 2 (LPDDR2)**
 - **3D-Stacked Wide-IO DRAM (3D-DRAM)**
 - **Performance and power consumption depends on:**
 - **Operating frequency**
 - **Interface width**
 - **Memory map (BI, BC and BL)**
- 
- Memory configuration**

Our contributions

- It is difficult to determine the memory configuration for a given set of mixed real-time applications
- Our contributions:
 - We show the trends in real-time performance of mobile DRAMs across and within generations
 - We propose a methodology to select the DRAM configuration for a real-time mobile system

Our approach




- **Analyze the trends in worst-case bandwidth, average-case execution time and power consumption**
- **From the analysis, derive a methodology for selection of memory configuration**

Memory devices

- Fastest and slowest device in each of the following memory generations:
 - LPDDR
 - LPDDR2
 - 3D-DRAM
- 3D-DRAM configurations are generated using the *3D-DRAM generator model* from *University of Kaiserslautern, Germany*

Worst-case bandwidth results



- **LPDDR, LPDDR2 and 3D-DRAM guarantee up to 0.75 GB/s, 1.6 GB/s and 3.1 GB/s**
- **3D-DRAM has higher efficiency with increasing request size, because of its wider interface**

Memory map selection

- Selection criteria of memory map (BI, BC and BL):

- 1. Access granularity \leq request size**

- Access granularity = $BI \times BC \times BL \times IO \text{ Width}$
- Data fetched from memory is not discarded

- 2. Interleave data to the maximum number of banks (BI) to exploit bank-level parallelism**

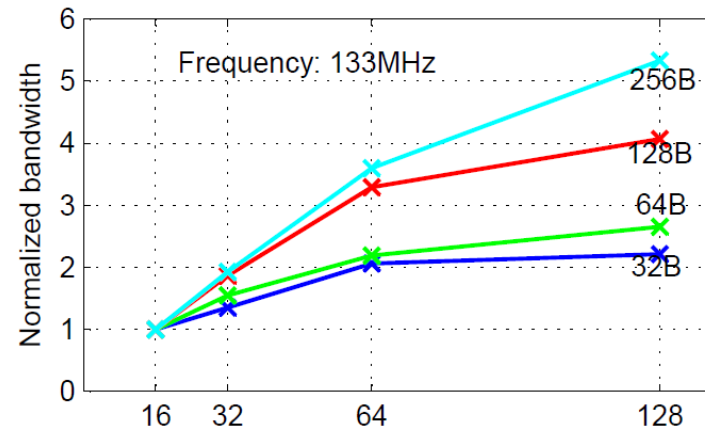
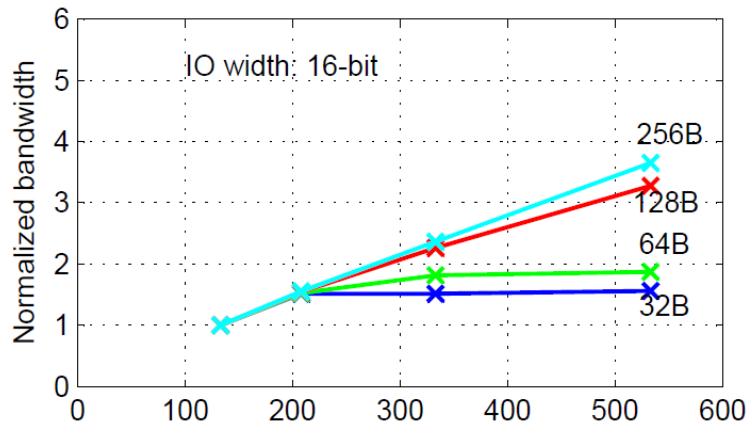
- Bank-level parallelism amortizes overhead

- 3. After satisfying 1 and 2, increase BC**

- Maximum efficiency in a single transaction


IP3-4: Memory-Map selection for Firm Real-Time SDRAM Controllers
Wednesday 16:00-16:30, Room: Ground Floor

Frequency and IO width selection



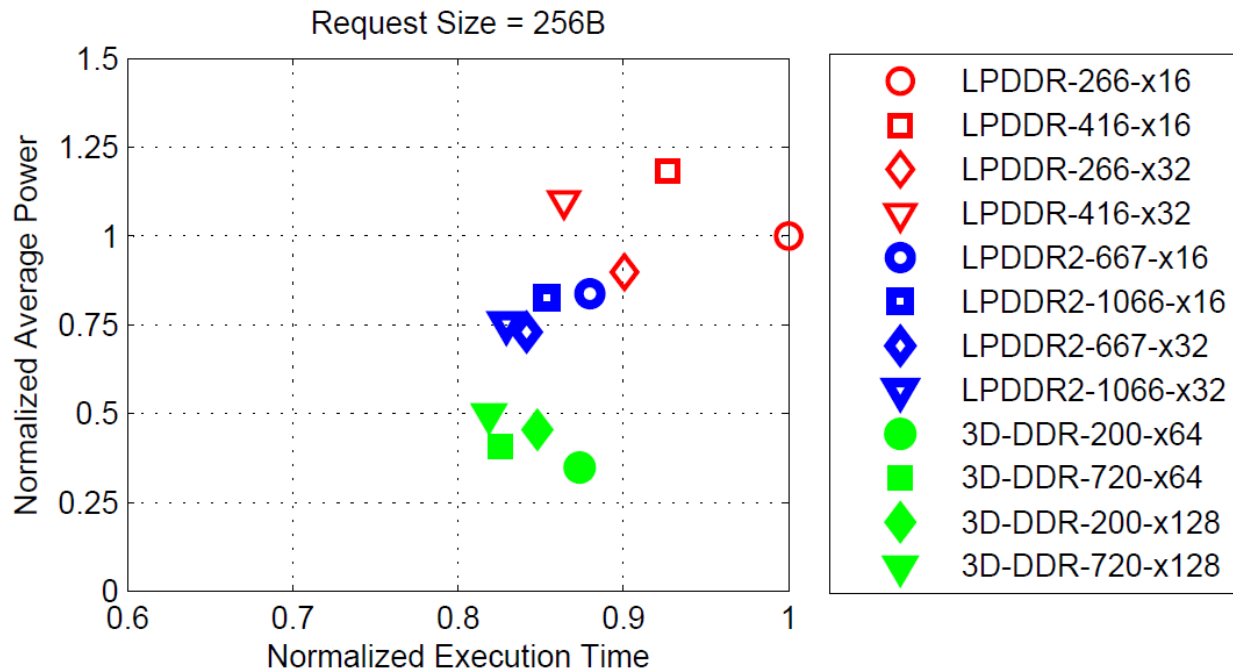
- Operating frequency increases → overhead increases
- Interface width increases → overhead remains constant
- IO width and operating frequency selection:
 1. Select the widest interface as long as the access granularity is less than or equal to request size
 2. Select a higher operating frequency

Average-case experimental setup



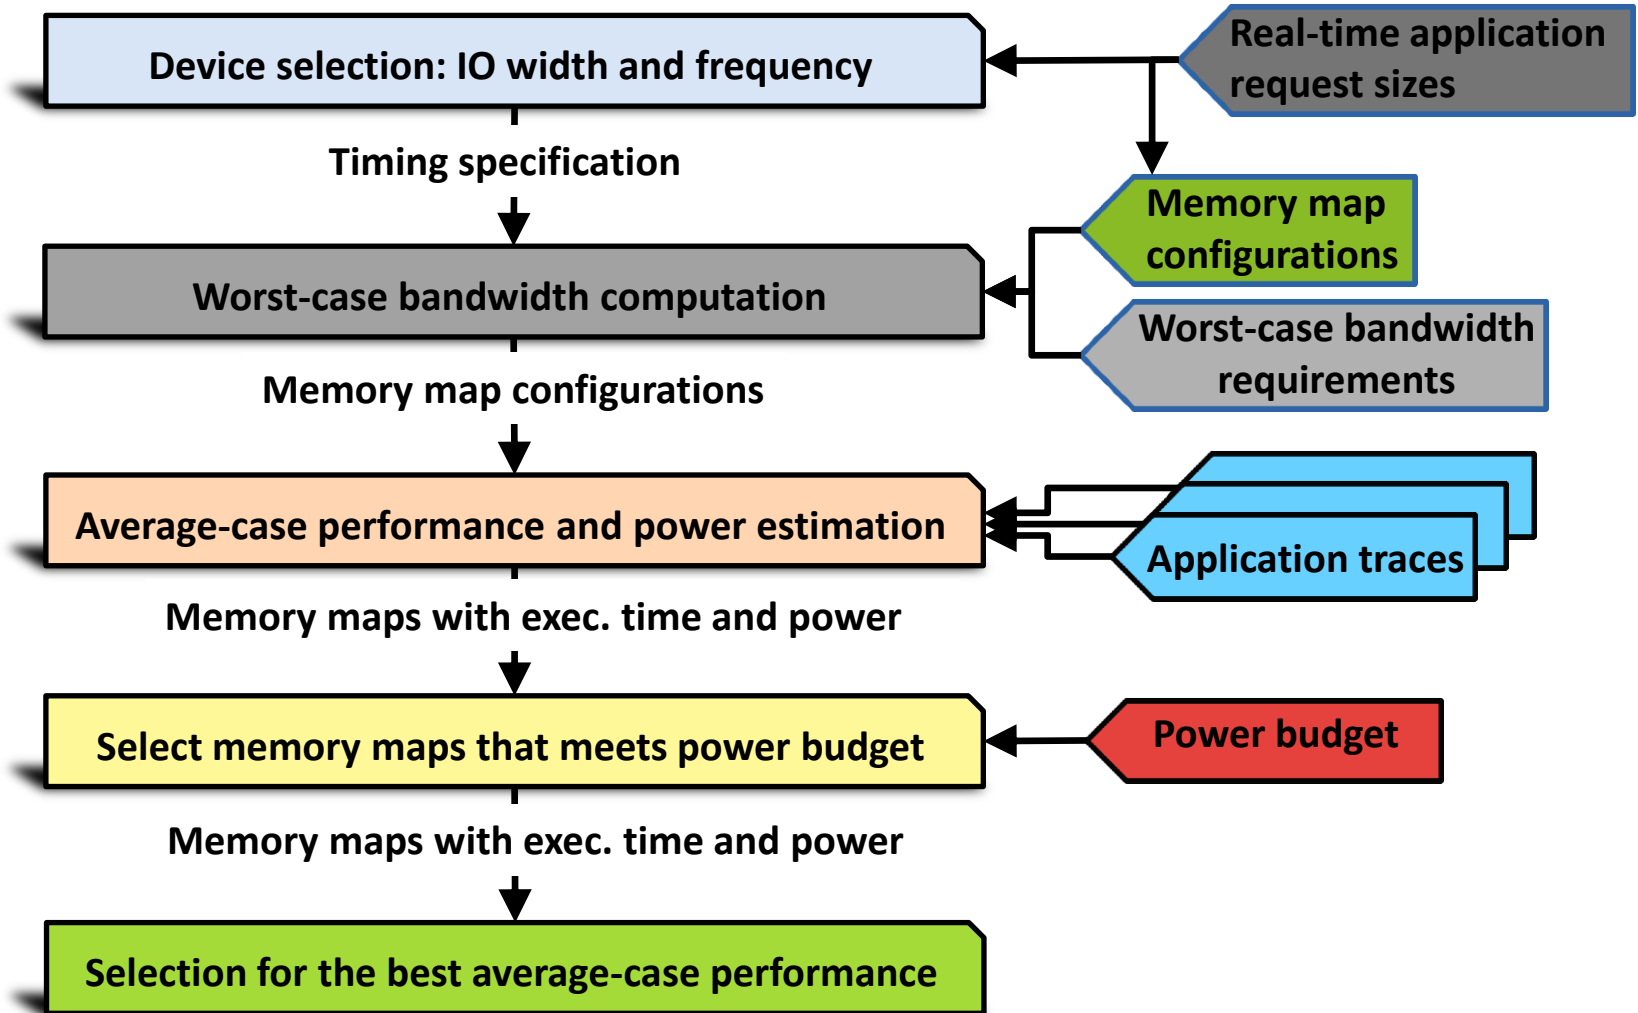
- **Application trace: memory requests by running H.263 video decoder application in SimpleScalar**
- **Real-time memory controller: Predator**
- **Request sizes: 32B, 64B, 128B, 256B**

Average-case analysis results



- Compared to LPDDR, LPDDR2 and 3D-DRAM have up to
 - 25% and 67% lower power consumption
 - 18% and 25% improvement in execution time
- Wider interface and lower operating speed → better performance at a lower power consumption

Proposed methodology



Conclusions

- We analyzed the real-time performance of mobile DRAM across and within generations

Memory	Worst-case bandwidth	Power savings w.r.t LPDDR-266-x16	Performance gain w.r.t LPDDR-266-x16
LPDDR-416-x32	0.75 GB/s	-15%	14%
LPDDR2-1066-x32	1.6 GB/s	25%	18%
3D-DRAM-720-x128	3.1 GB/s	67%	25%

- We proposed a methodology for selecting DRAM configuration for a real-time mobile system
 - Satisfies worst-case bandwidth requirements
 - Provides best average-case performance
 - Meets power budget

Questions?

m.d.gomony@tue.nl