BOOM: Enabling Mobile Memory Based Low-Power Server DIMMs

Doe Hyun Yoon, Jichuan Chang,
Naveen Muralimanohar, and Parthasarathy Ranganathan
Intelligent Infrastructure Lab (IIL), Hewlett-Packard Labs

ISCA 2012

Server Memory

Capacity demand

In-memory database

Memcached

Memory power

215W for 256GB

Even larger than CPU power

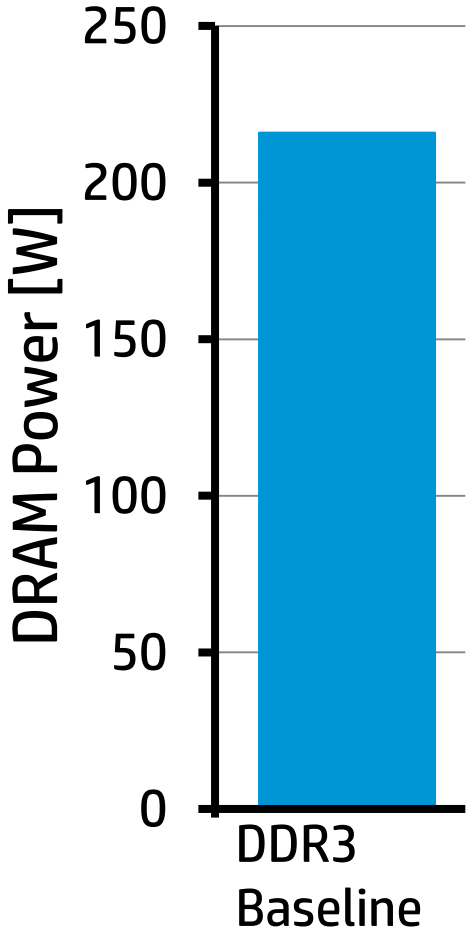

BOOM [Buffered Output On Module]

Buffered DIMM

Wide slow data bus

Mobile memory

Much lower power

Only 58W for 256GB

No compromise in performance & reliability

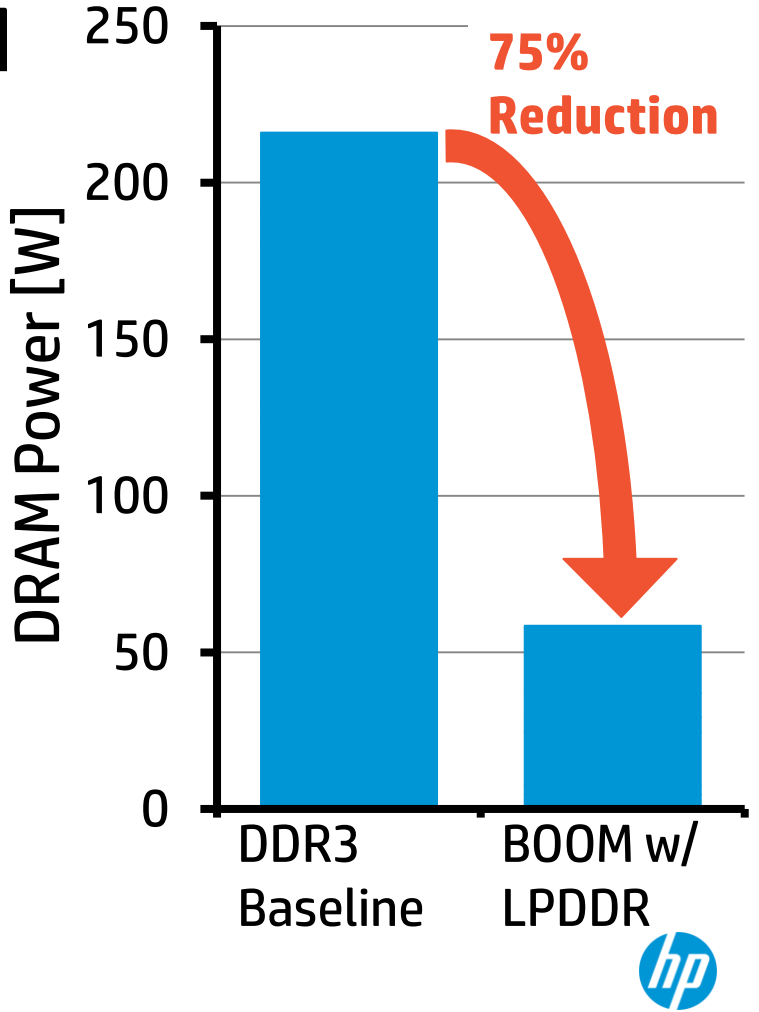

Mobile Memory for Servers

Mobile Memory Advantages

Designed for low power, comparable latency

	LPDDR2 x16	DDR3 x4/x8	DDR3 x16	
V_{DD}	1.2V	1.5V	1.5V	
I_{DD2P}	1.6mA	35mA	40mA	Background current
I/O	Unterminated	Terminated	Terminated	
tRCD	15ns	13.75ns	13.75ns	Row activation latency
tCL	15ns	13.75ns	13.75ns	Column read latency
Burst length	Burst 4	Burst 8	Burst 8	

Can we use mobile memory for servers?

(Server) Memory Requirements

High capacity

High performance – high BW

Low power

Low cost – Total Cost of Ownership (TCO)

High reliability/availability – chipkill-correct

These objectives are inter-twined among others

Mobile memory can reduce server power, but may adversely affect other objectives

Chipkill-correct

System continuously functions even at a DRAM chip failure

Single DRAM chip failure correction and double DRAM chip failure detection

At least 2 ECC DRAM chips per logical rank

Reed Solomon code combined with erasure technique

Easier to implement chipkill with narrow (x4) DRAM chips

2 x4 ECC chips per 16 x4 data chips

12.5% overhead

2 x8 ECC chips per 8 x8 data chips

25% overhead

2 x16 ECC chips per 4 x16 data chips

50% overhead

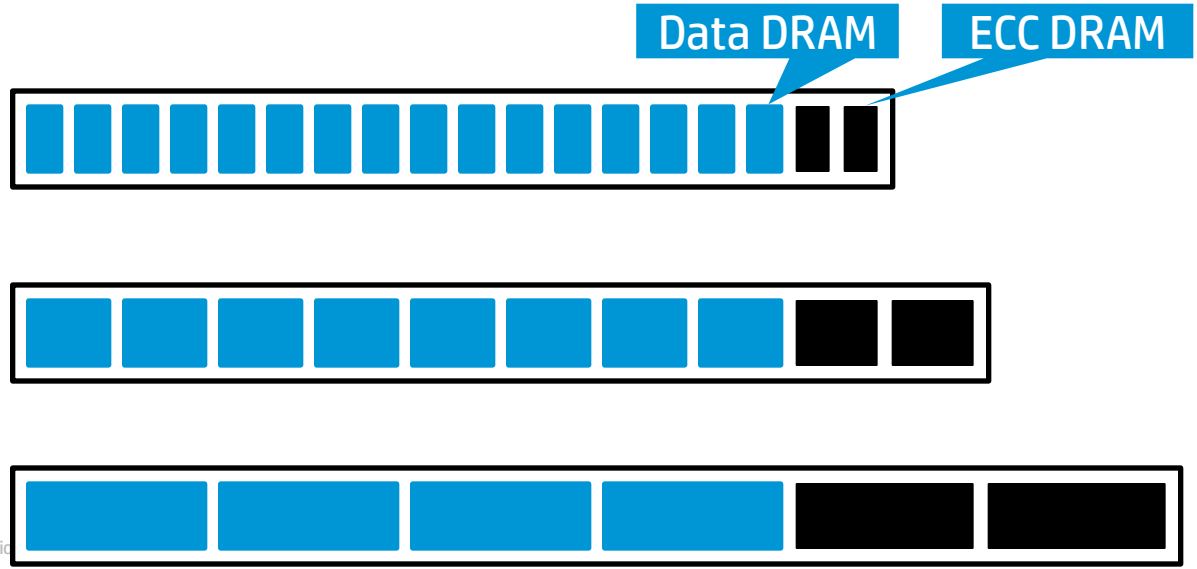

Mobile Memory for Low-Power Server DIMMs

Capacity

2-8Gb LPDDR2 is already available

Buffered DIMMs can help

Cost

Increasing market volume drives cost low

Low operating power can improve TCO

Bandwidth

0.8GHz LPDDR2 vs. 1.6GHz DDR3

1.6GHz LPDDR3 vs. 3.2GHz DDR4

Reliability

Only x16 and x32 devices

Difficult to enable chipkill-correct

How to enable high BW and chipkill-correct with mobile memory?

BOOM Architecture

Baseline Load-Reduced DIMM (LR-DIMM)

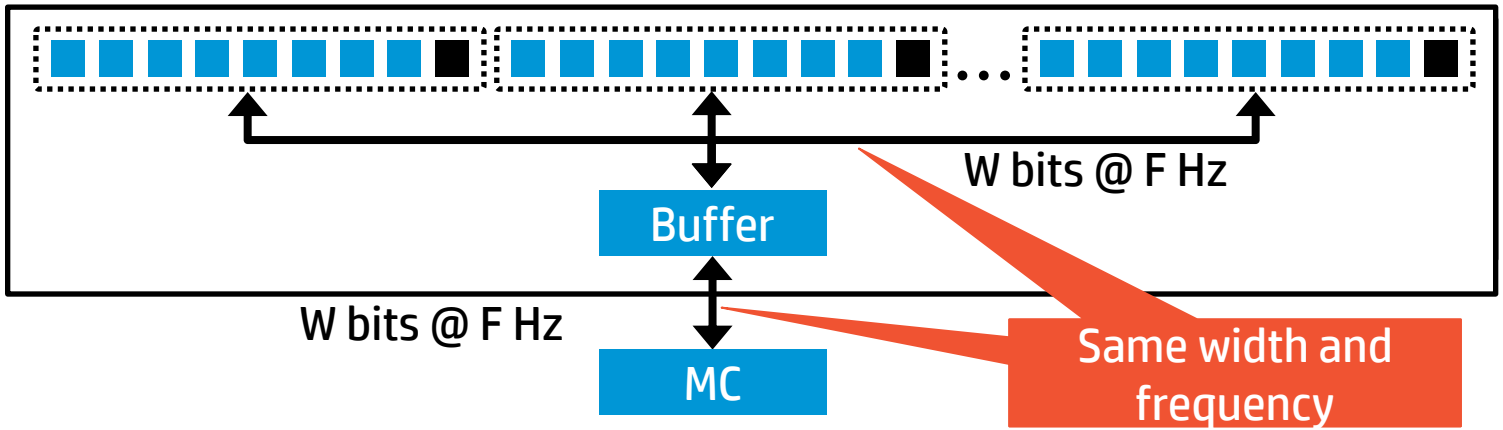

On-module buffer completely isolates external bus from internal bus

Allow a larger number of DIMMs per channel at high channel frequency

External and internal buses have the same channel width and frequency

High channel bandwidth ➔ High DRAM frequency

Baseline Load-Reduced DIMM (LR-DIMM)

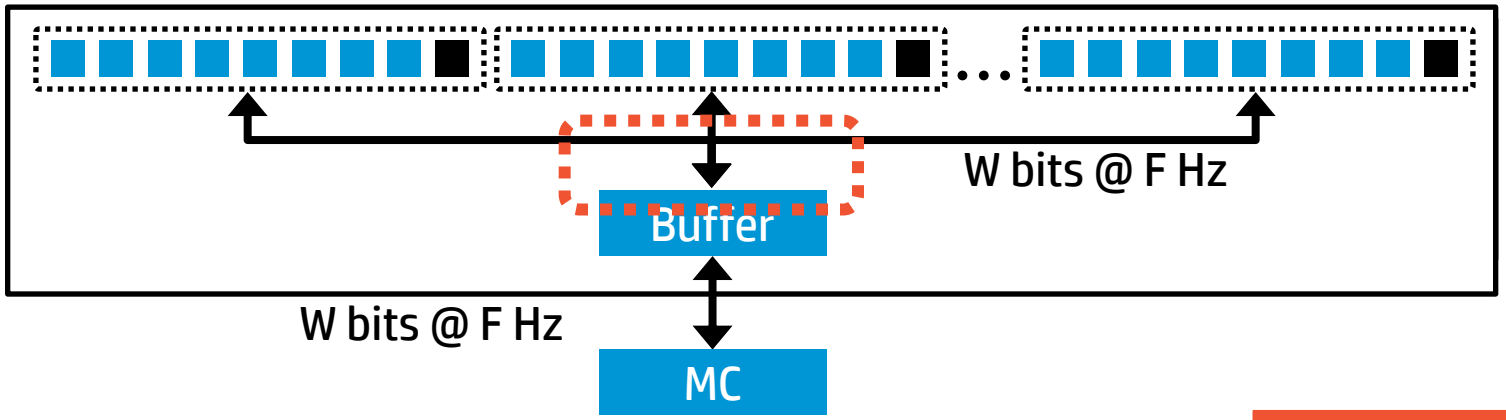

BOOM

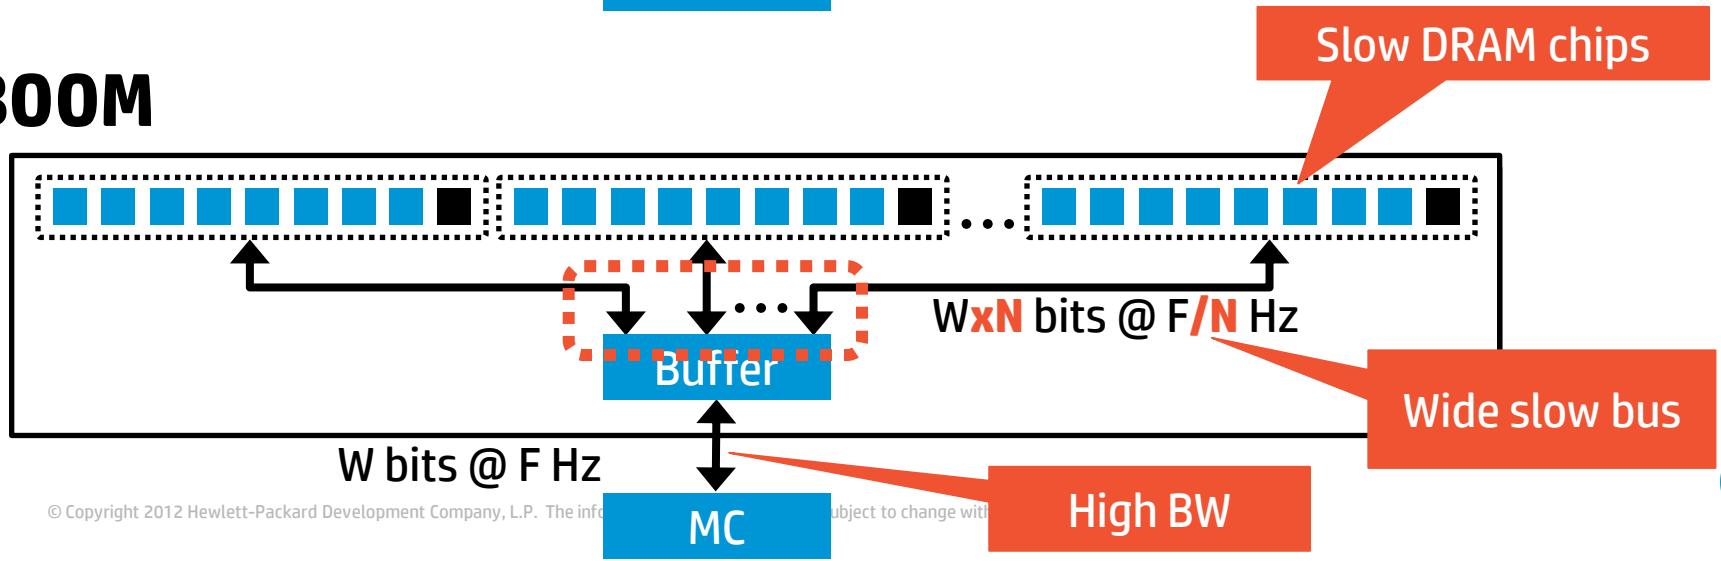

BOOM

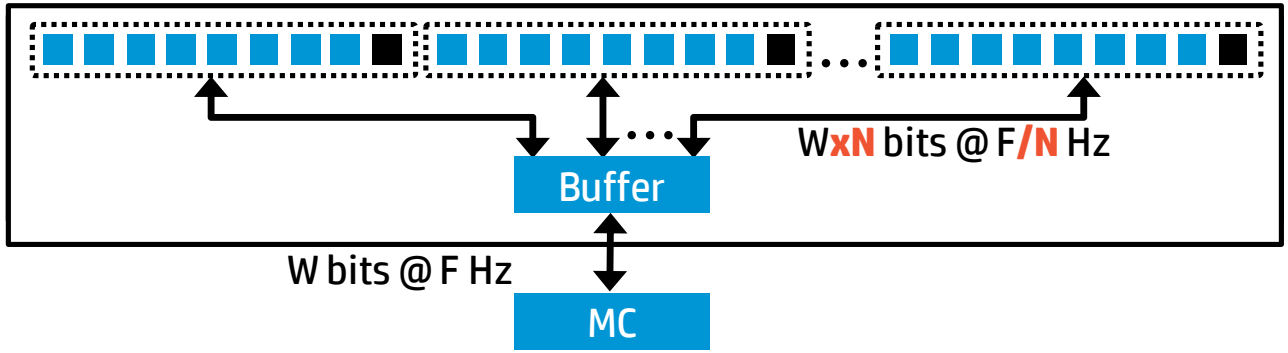

Re-organize DIMM

Commodity DRAM

Changes:

- Routing within DIMM
- Buffer chip
- ECC handling at the memory controller

N times wide internal bus

1/N DRAM frequency – Low power

Chipkill-correct for wide DRAM chips (e.g., x16)

The same external BW

N2 & N4 architectures – DRAM chips running at ½ or ¼ of the bus frequency

Baseline LR-DIMM with DDR3-1600

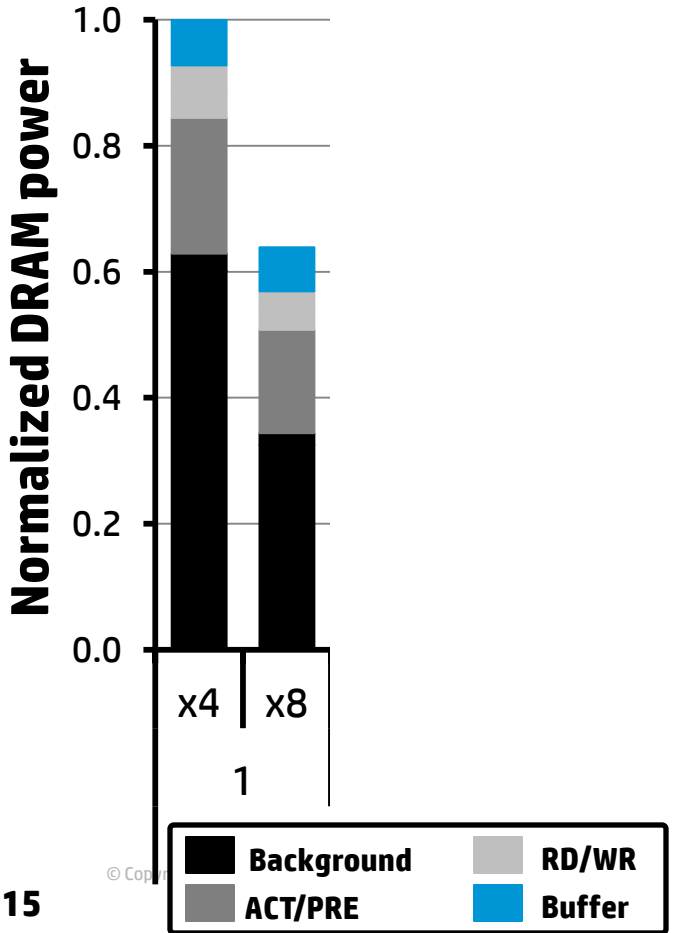

8 DDR3-1600 memory channels

Capacity: 256GB

Peak BW: 104GB/s

2 channel lock-step mode for chipkill-correct

Only 4 logical channels (128-bit wide)

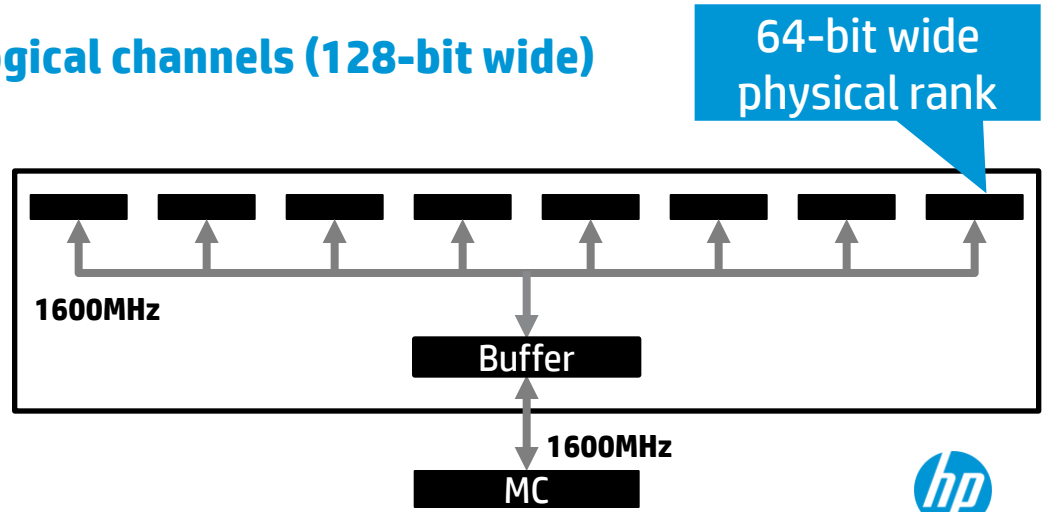

BOOM N2 with DDR3-800

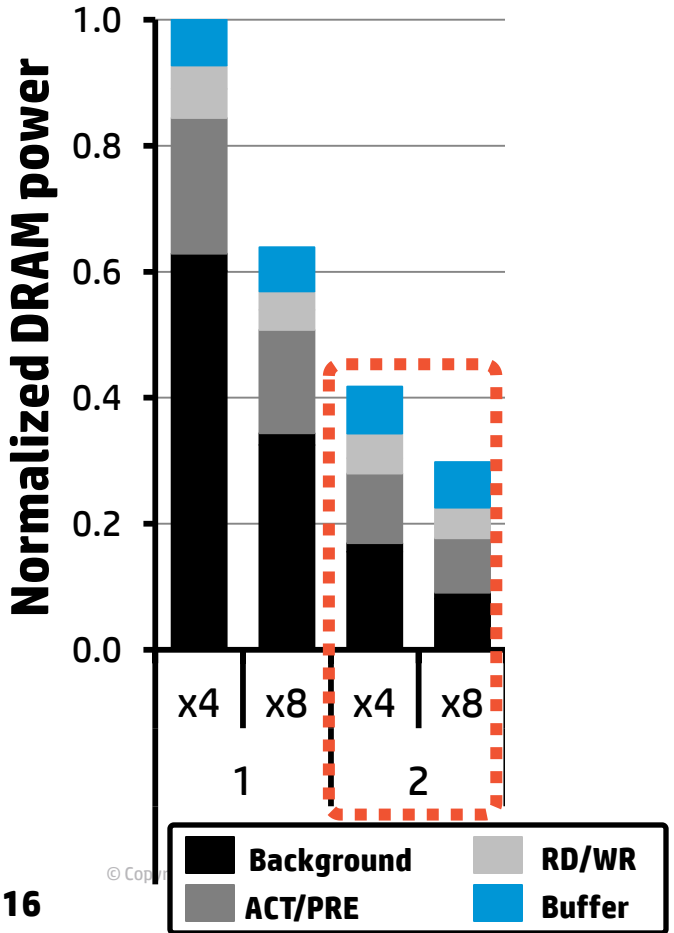

1/2 DRAM frequency – reduce background power

Same BW, Same capacity

128-bit wide Internal bus

No lock-step mode, but support chipkill-correct

8 logical channels (64-bit wide)

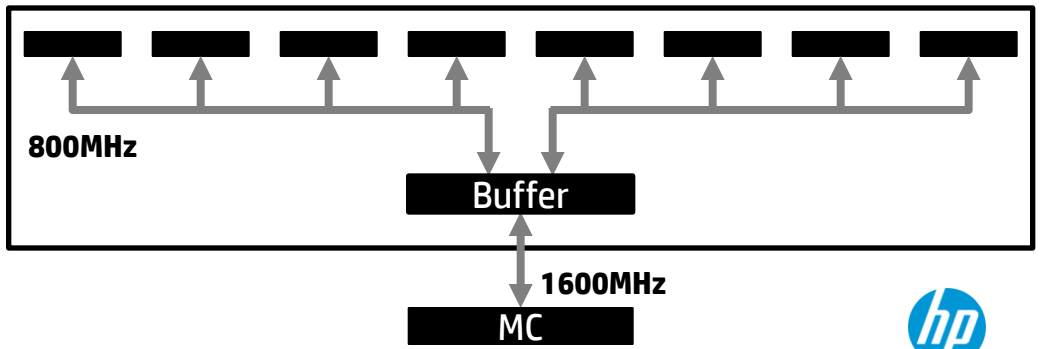

BOOM N4 with LPDDR2-400

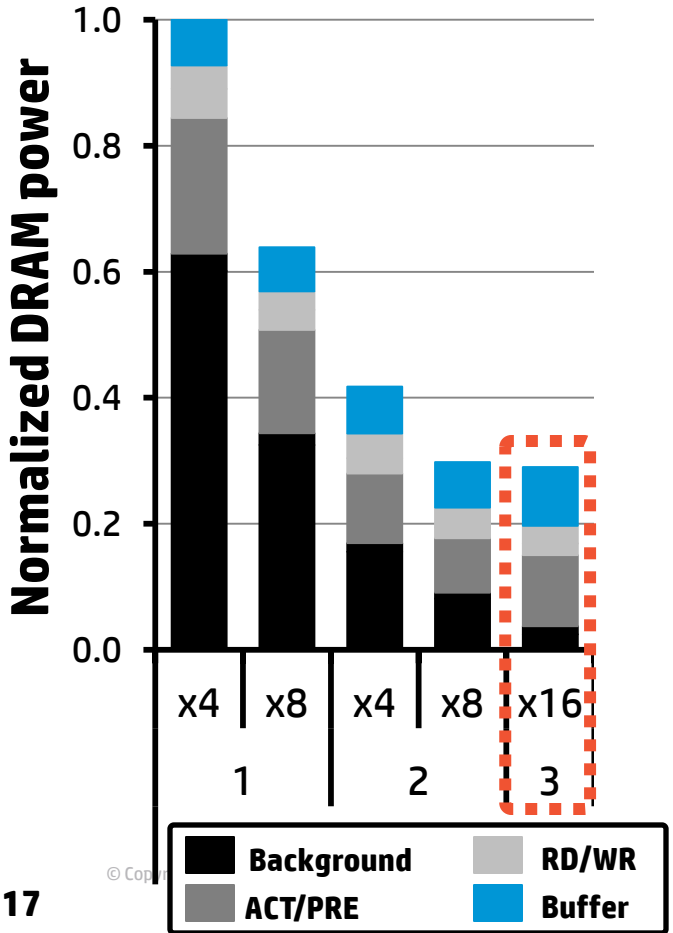

1/4 DRAM frequency

LPDDR2 further reduces background power

Still, Same BW, Same capacity

Internal 256-bit wide channel

No lock-step mode – 8 logical channels

Support chipkill-correct even for x16 LPDDR2 DRAM

Wide internal channel increases ACT/PRE power

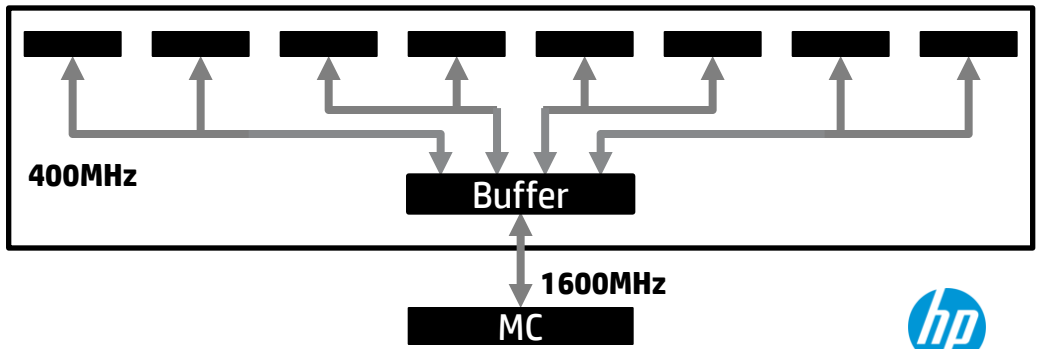

BOOM N4 + Sub-Ranking

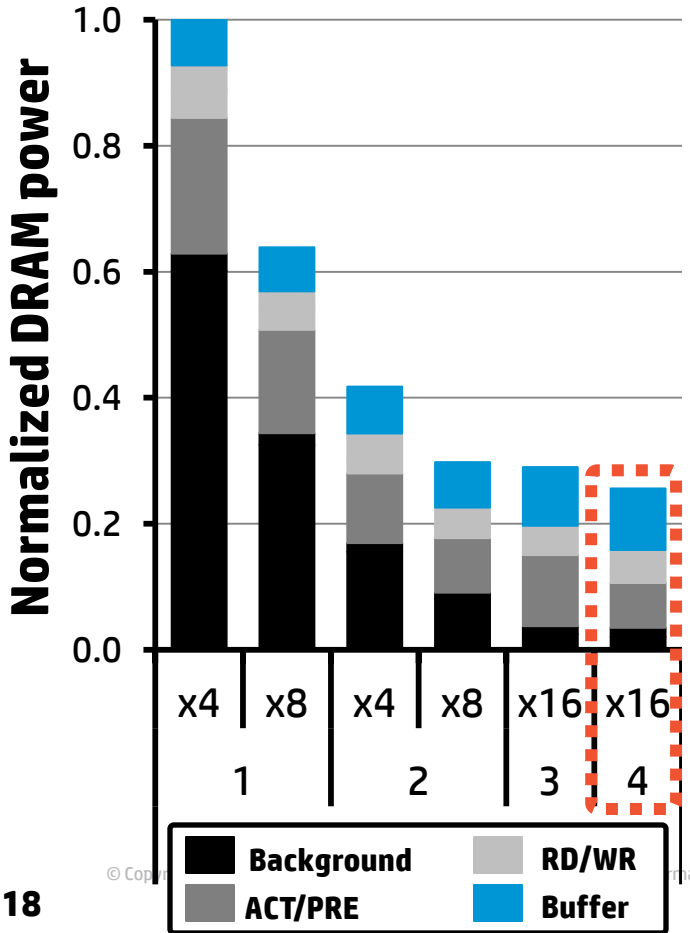

Sub-ranking (or rank-subsetting)

Access two 128-bit wide sub-ranks independently

Reduce ACT/PRE power

Increased complexity in scheduling

Higher ECC overhead

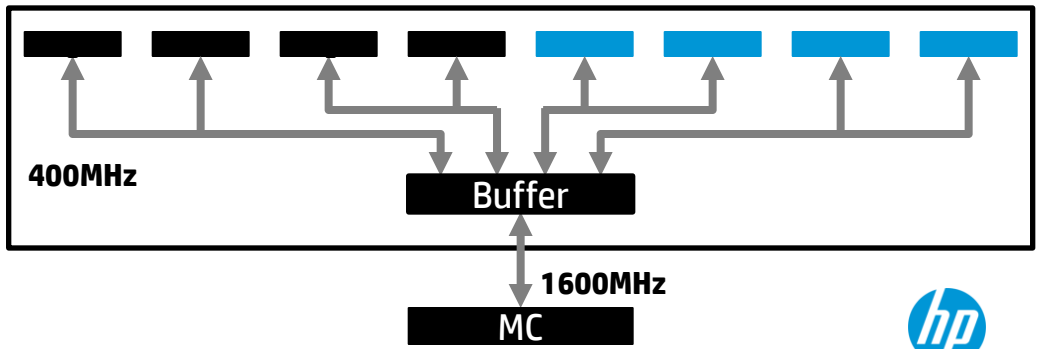

DRAM Power reduction with BOOM

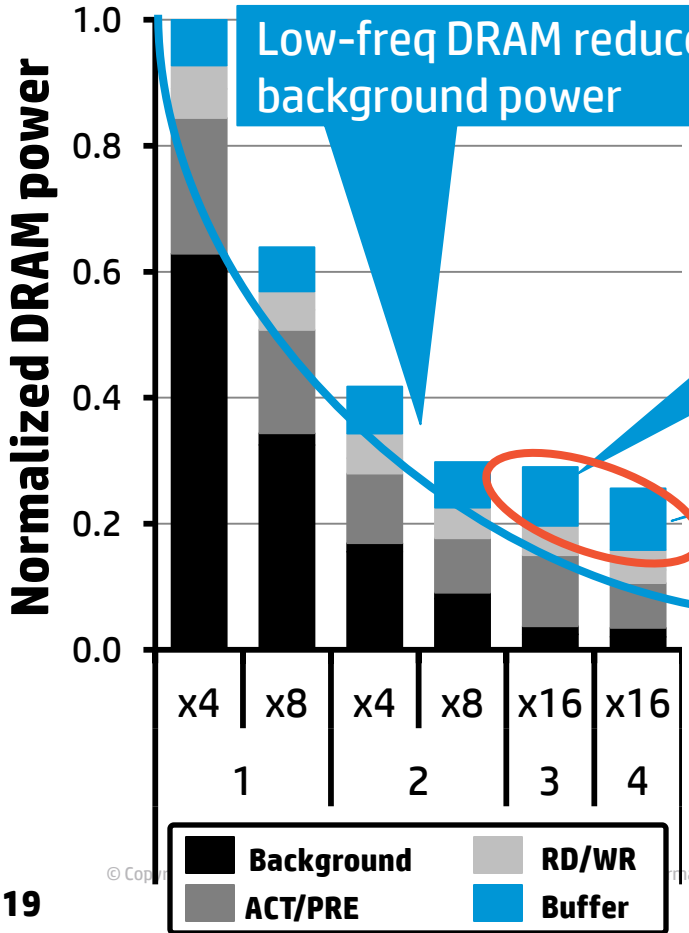

Low-freq DRAM reduces background power

LPDDR2 further reduces background power, but N4 arch increases ACT/PRE power

Sub-ranking saves ACT/PRE power

Buffer and I/O power is the next challenge

75% DRAM power saving w/ BOOM

1. Baseline @ 1600
2. BOOM N2 @ 800
3. BOOM N4 @ 400
4. BOOM N4 @ 400 + sub-ranking

Implications for Future Memory Systems

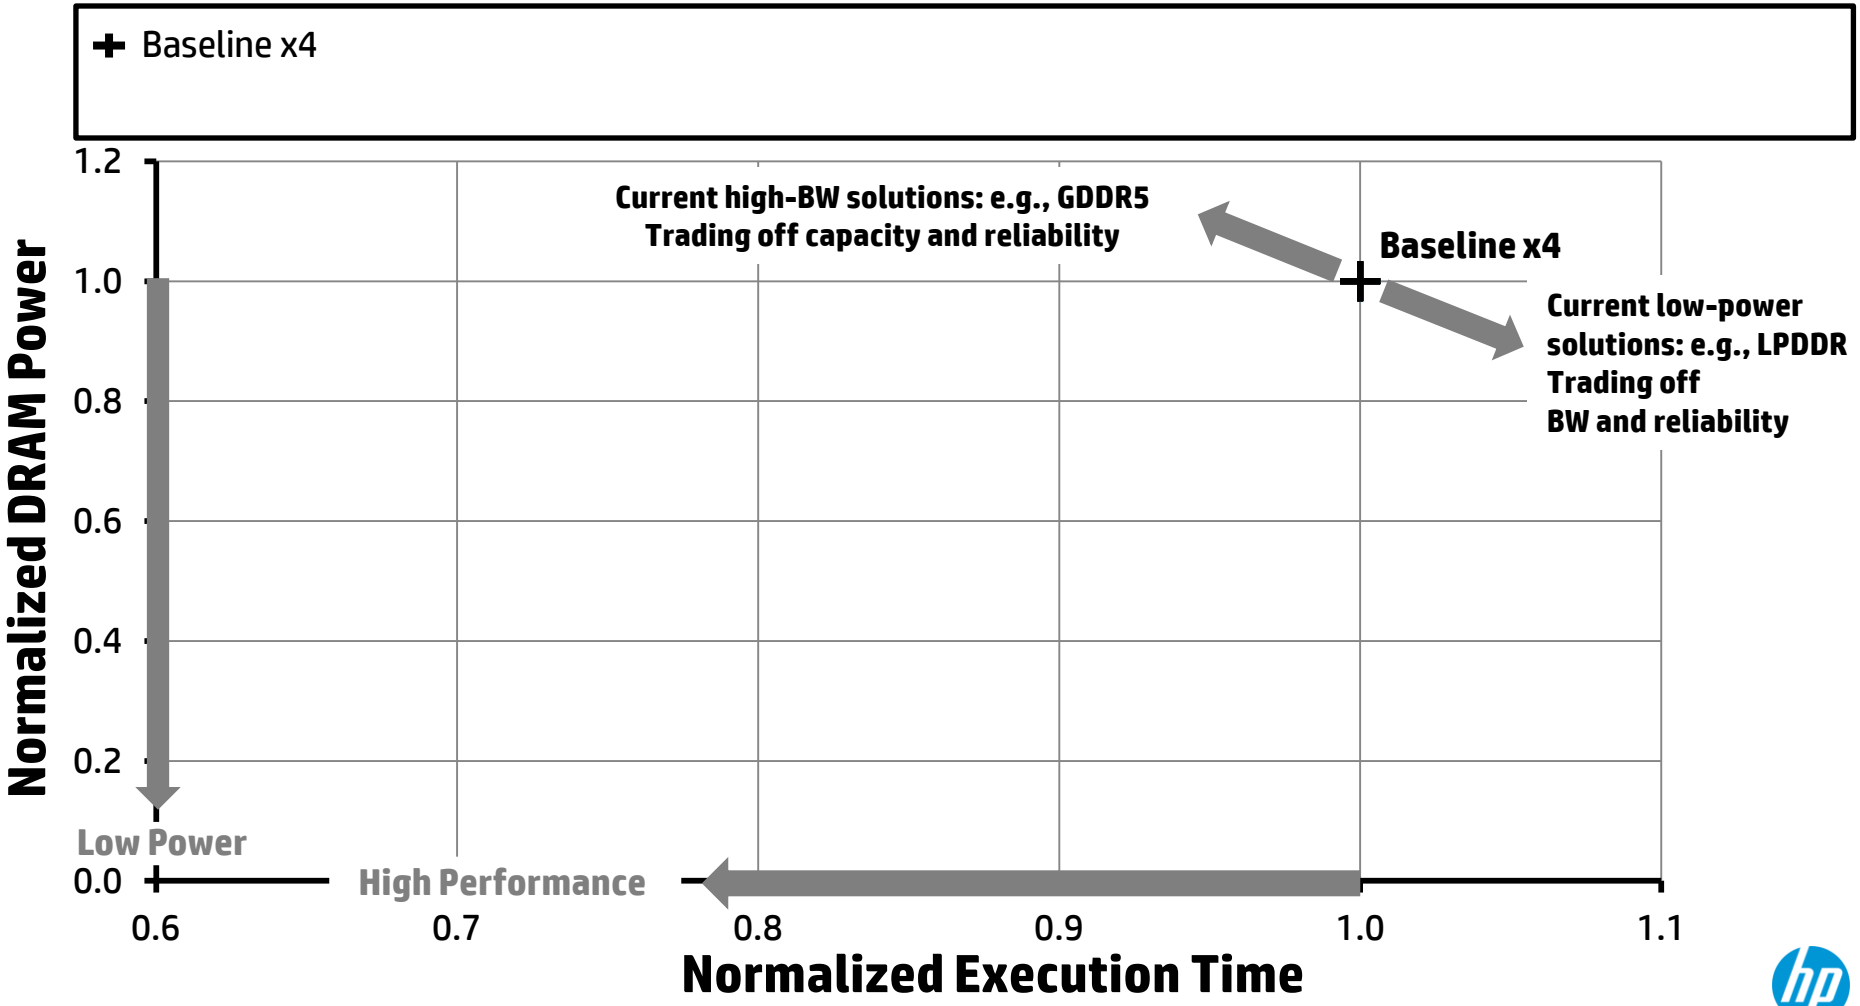

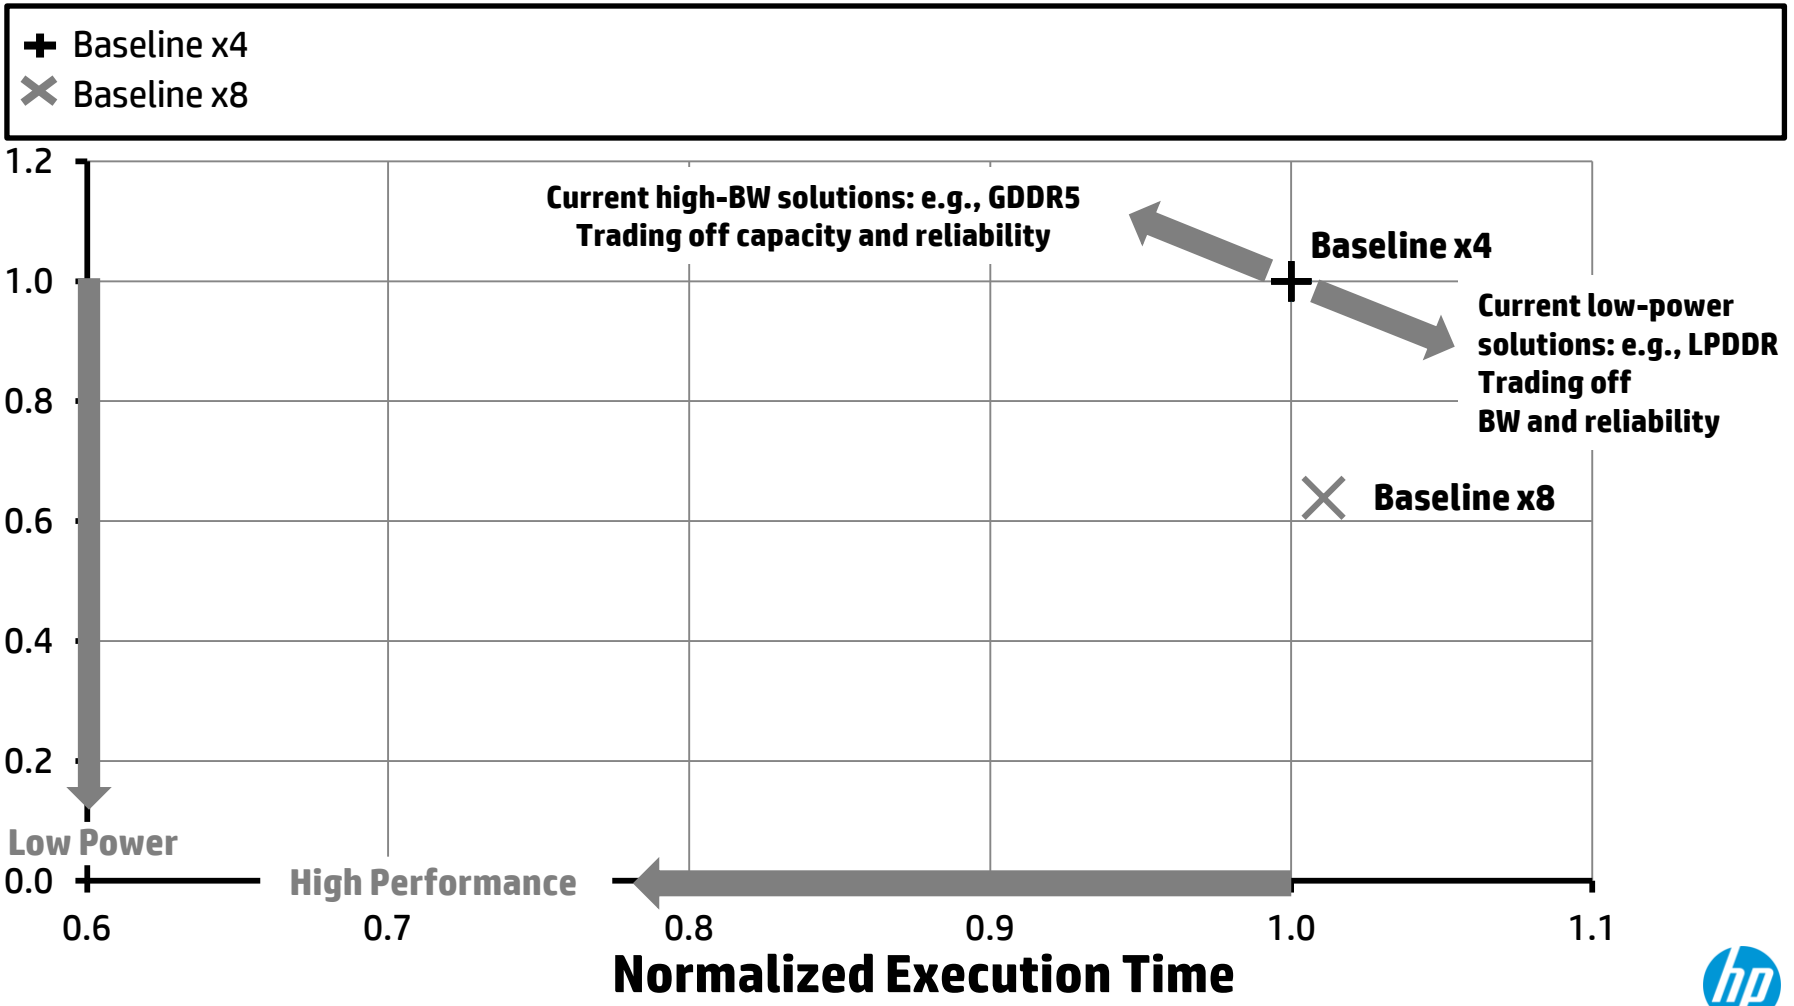

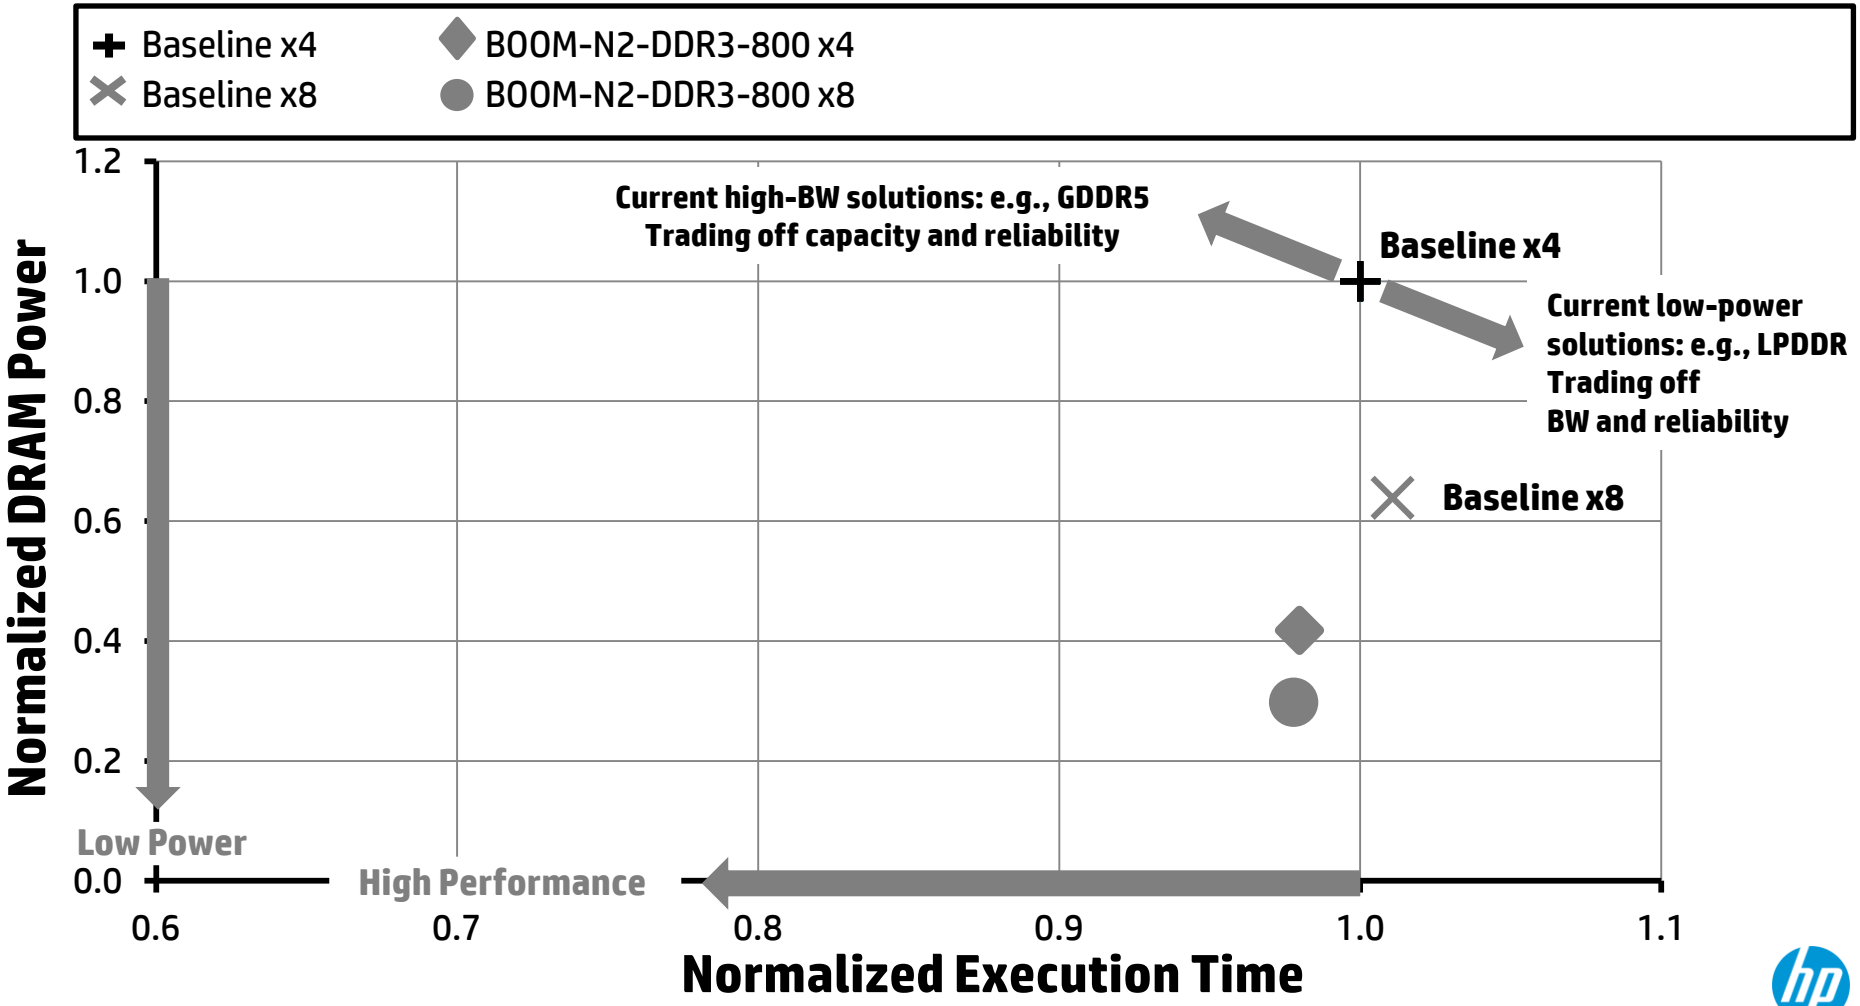

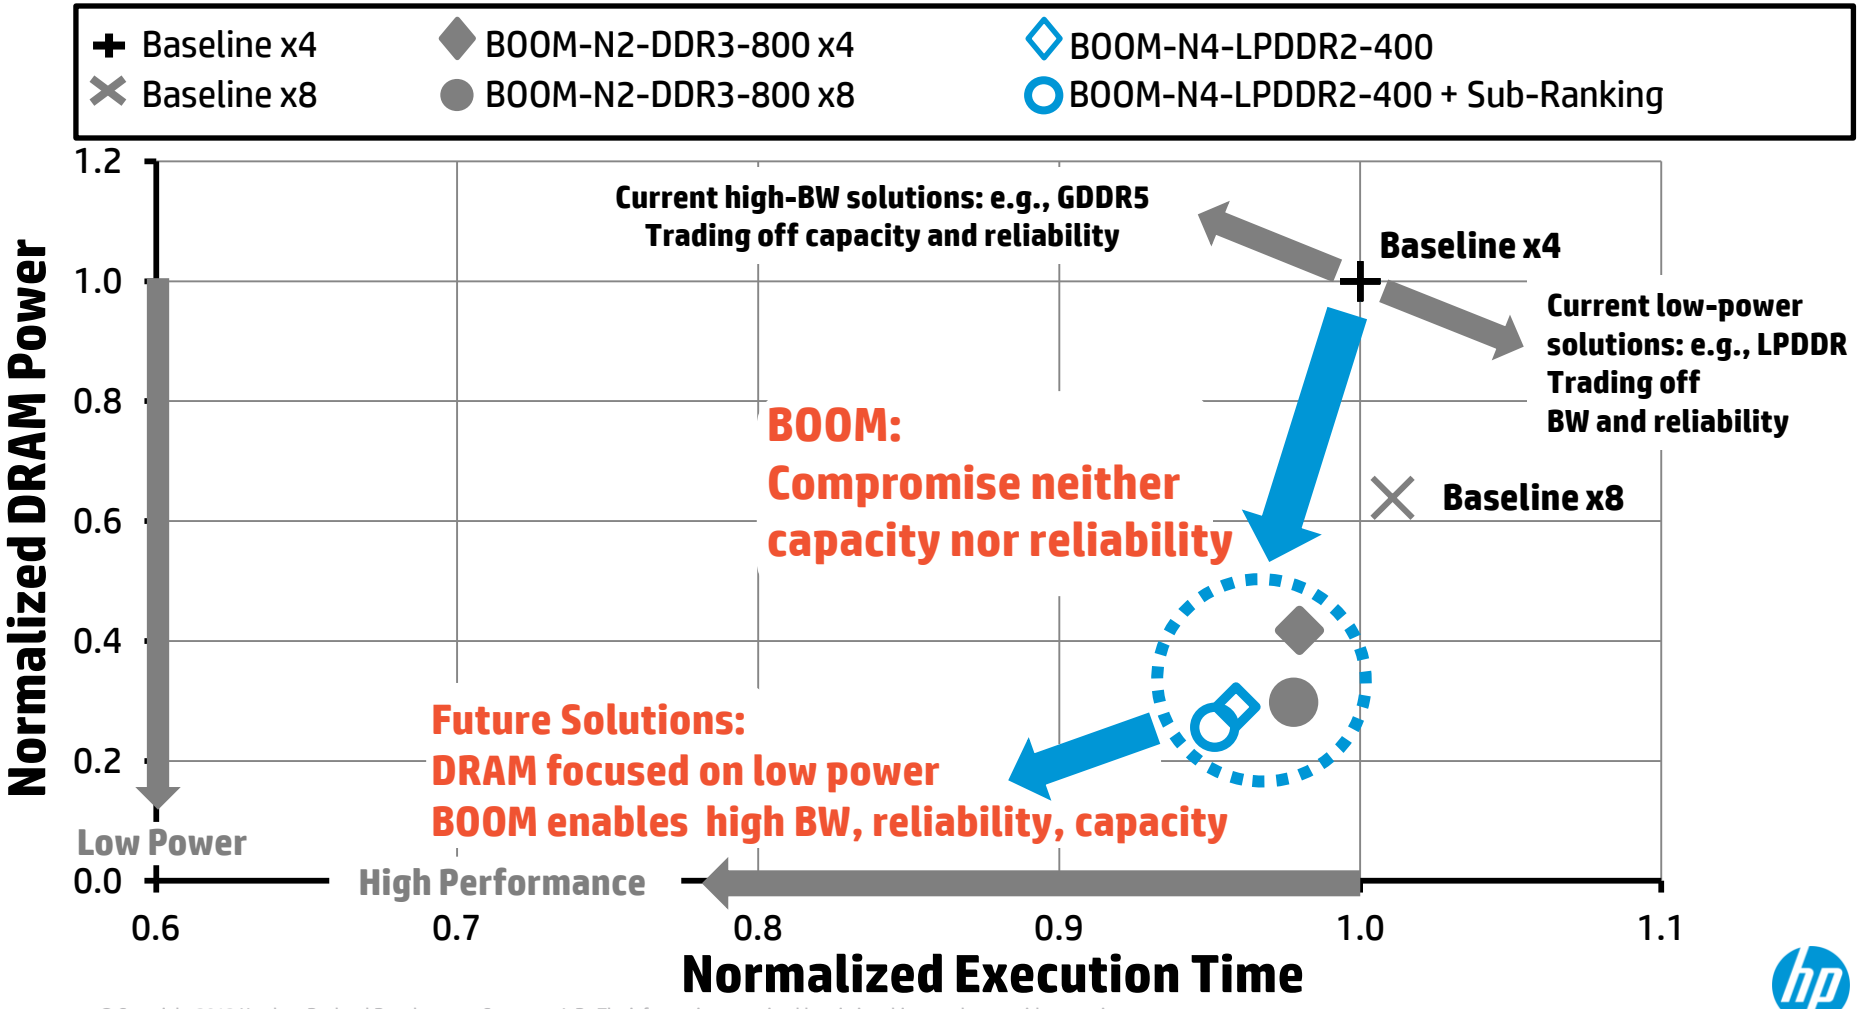

Conclusions

BOOM enables MOBILE MEMORY for servers

Wide slow internal bus

Commodity DRAM chips

DRAM power reduction

75% saving on average, across memory intensive applications

Much lower background power

Energy-proportional memory system

No compromise in BW and reliability

Same external bandwidth, Same capacity

Support chipkill-correct for wide DRAM chips

More details [ECC, access granularity, row-buffer size, etc] in the paper

BOOM: Enabling Mobile Memory Based Low-Power Server DIMMs

Doe Hyun Yoon, Jichuan Chang,
Naveen Muralimanohar, and Parthasarathy Ranganathan
Intelligent Infrastructure Lab (IIL), Hewlett-Packard Labs

ISCA 2012