



ARM Cortex-A* Series Processors

Haoyang Lu, Zheng Lu, Yong Li, James Cortese

ARM Cortex-A* Series Processors

- Applications
- Instruction Set
- Multicore
- Memory Management
- Exclusive Features

ARM Cortex-A* series: Applications



Ford Sync



Digital TV



Networking solutions



ARM Cortex-A* series: Applications



Smartphones and
Tablets



ARM Cortex-A* series: Applications

Processors of the Cortex-A series and their applications:

	A5	A7	A8	A9	A15	A53	A57
Smart phones	*		*	*		*	*
Home Computing	*		*	*			
Smart TVs	*		*	*			
Digital Cameras	*			*			
Embedded Computing	*	*	*	*	*	*	*
Home Networking	*			*	*		
Storage	*						

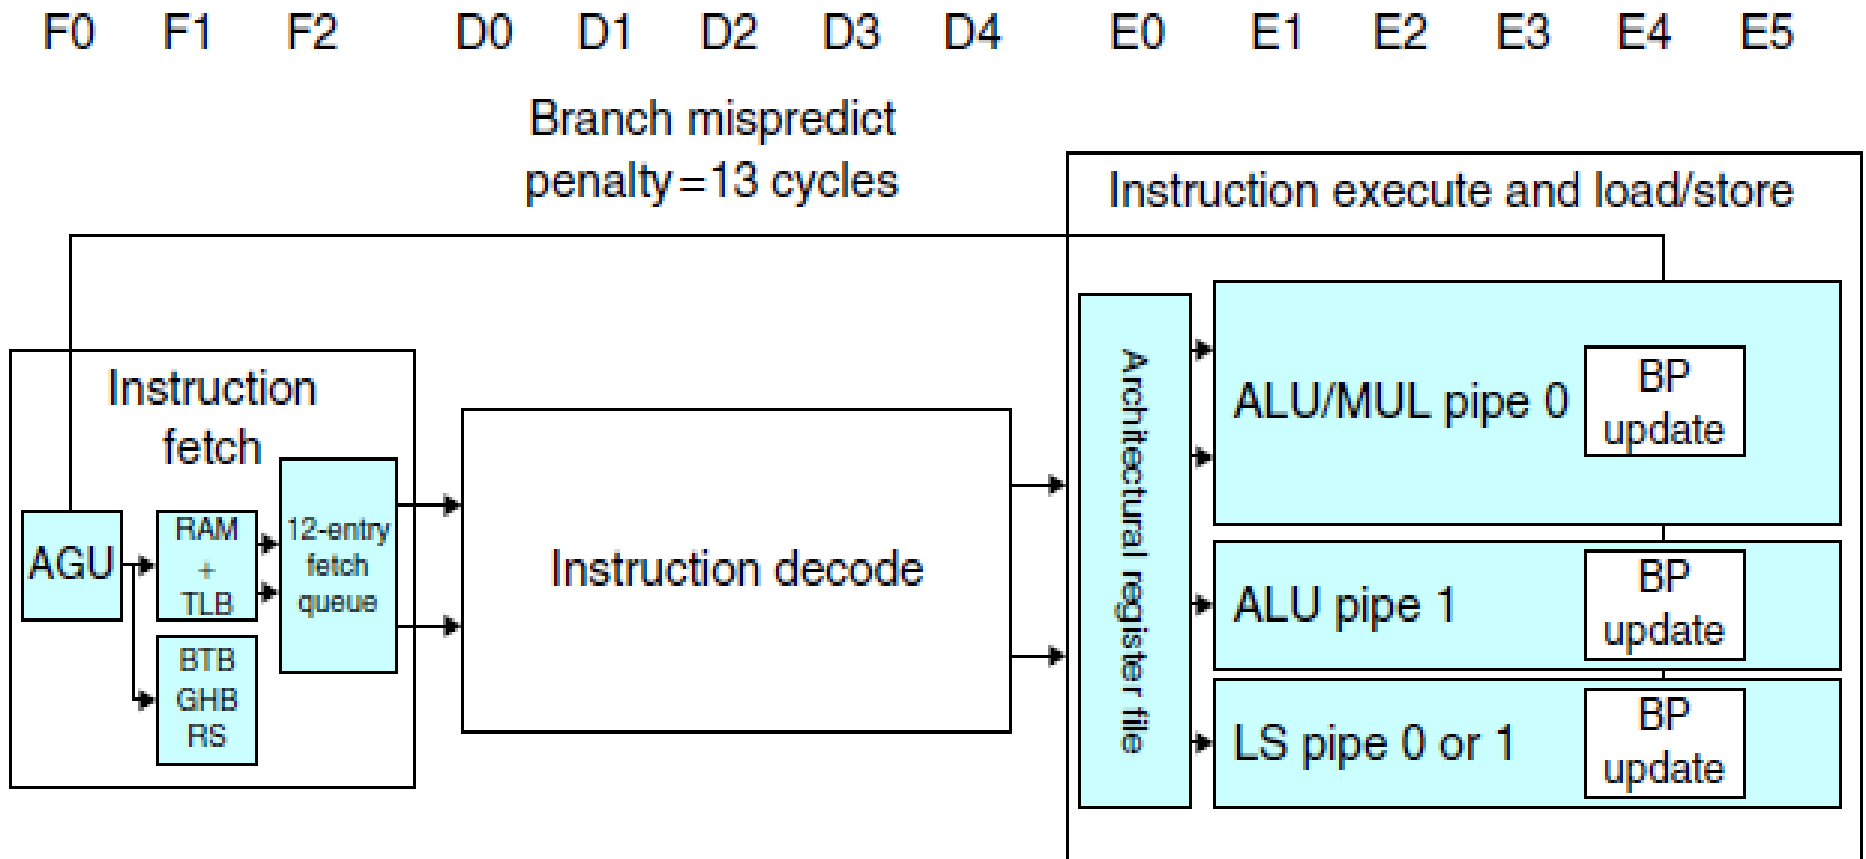
ARM: Instruction Set

- Two instruction set:
 - ARM instruction set (32-bit)
 - Thumb instruction set (mixed 16/32 bit)
- Thumb-2: Bit-field manipulation, table branches and conditional execution
- Unified Assembly Language (UAL): supports generation of either ARM and Thumb instructions from the same source code

ARM Cortex-A8 series: Pipeline

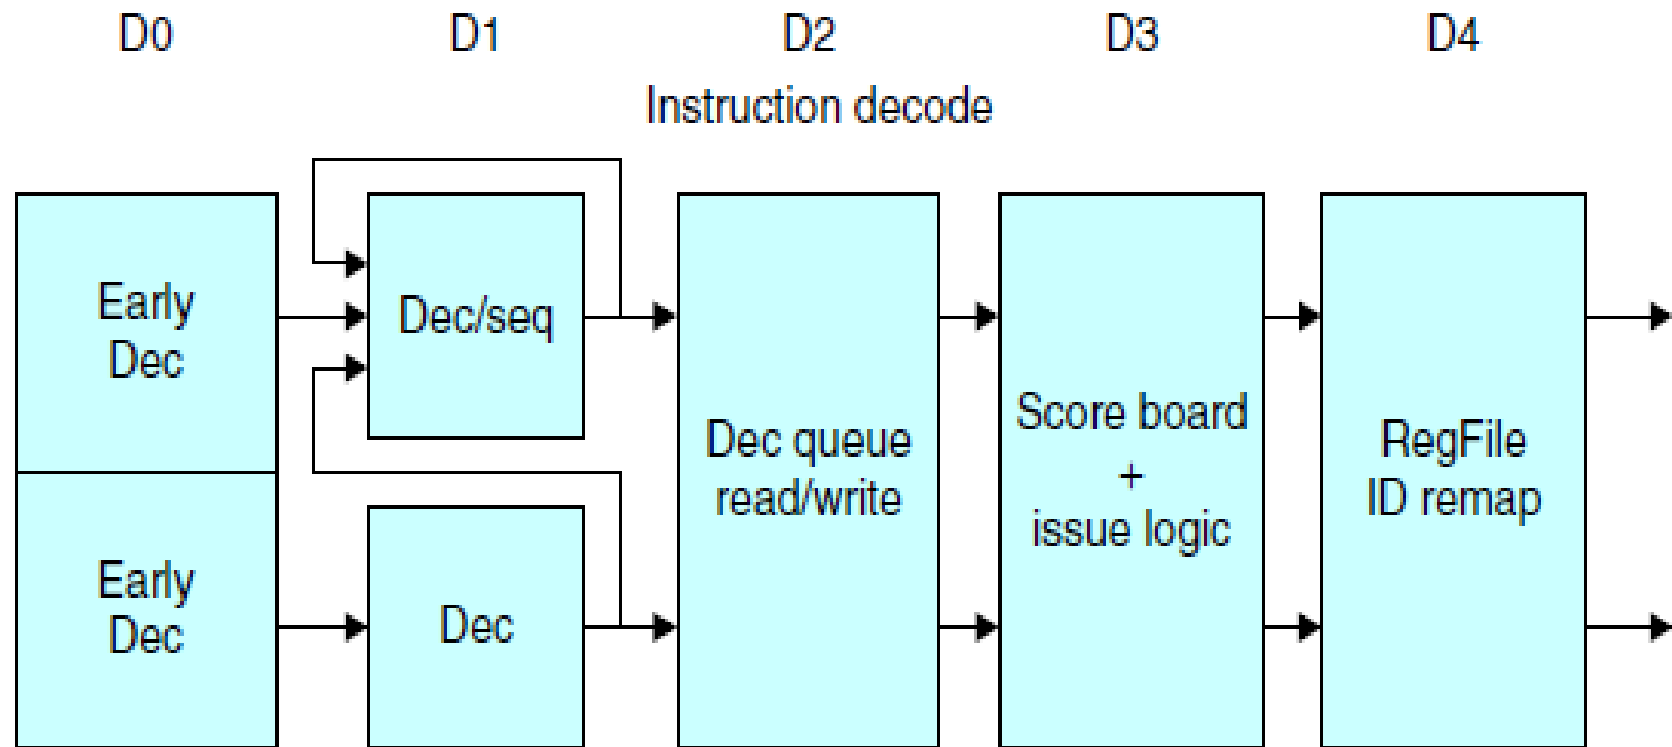
- Dual-issue
- Statically scheduled superscalar
- Dynamic issue detection – issue two instructions per clock
- Dynamic branch predictor – 512 entry branch target buffer
 - 4K-entry global history buffer
 - Mispredict penalty : 13 cycles

ARM Cortex-A8 series: Pipeline



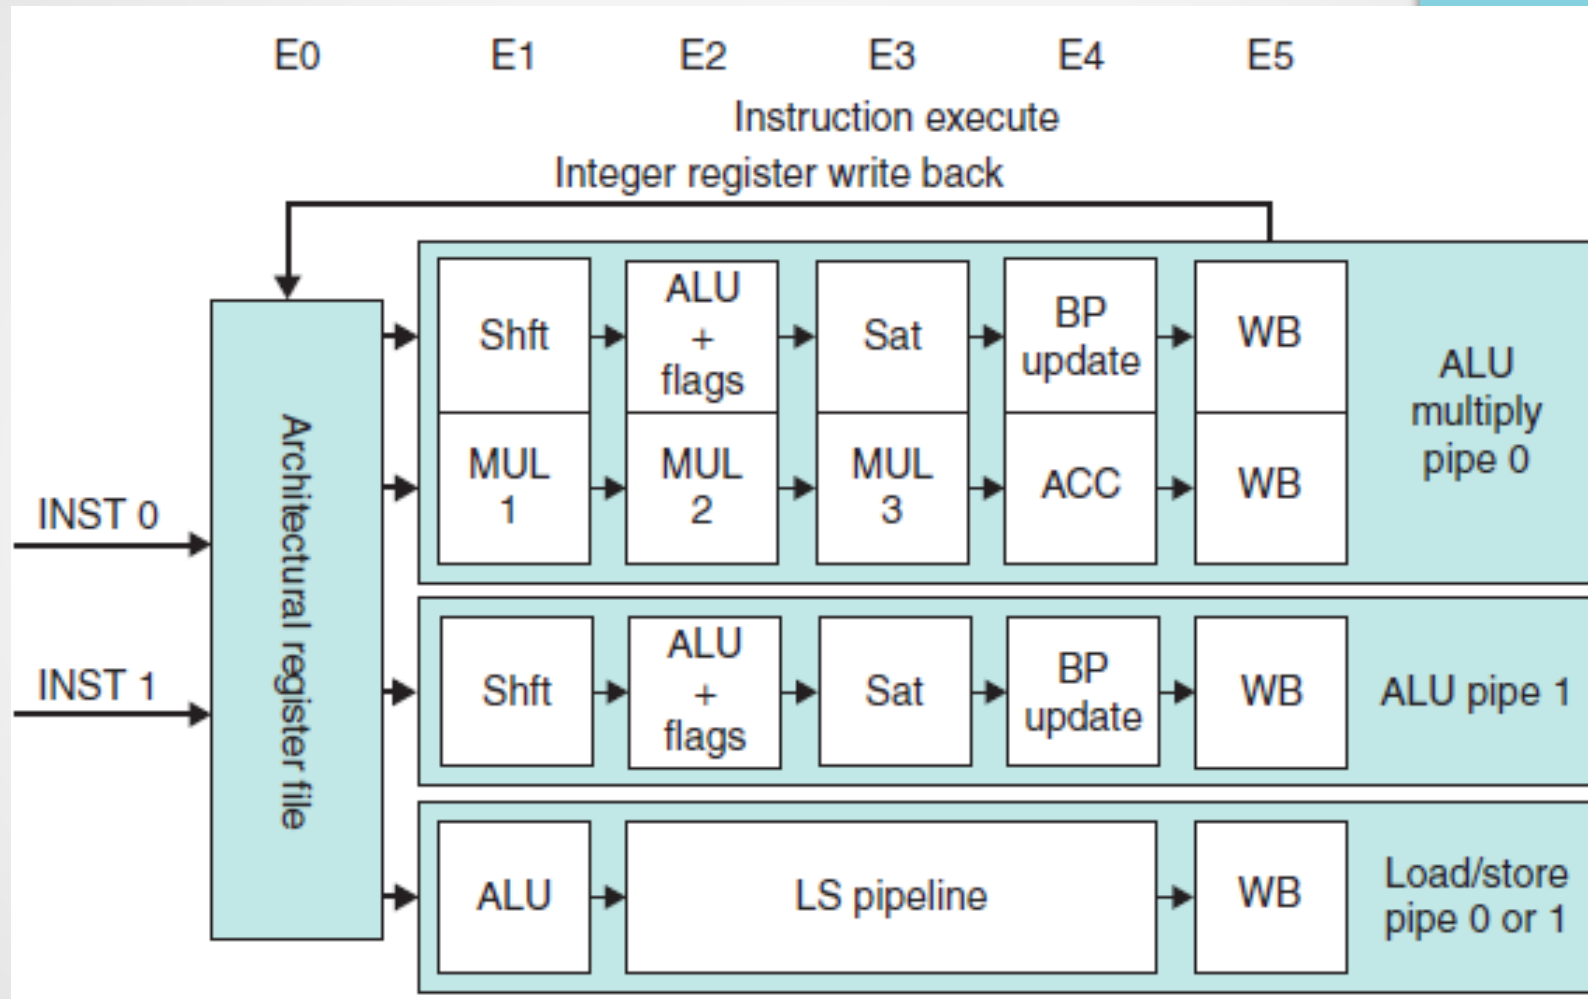
13-stage pipeline

ARM Cortex-A8 series: Pipeline



5-stage Instruction Decode

ARM Cortex-A8 series: Pipeline



Instruction decode execution

ARM Cortex-A8 series: Pipeline

- Ideal CPI is 0.5 according to its dual-issue
- Stalls:
 - Functional hazards, which occur when two instructions selected for issue simultaneously use the same functional pipeline.
 - Data hazards, which are detected early in the pipeline and may stall either both instructions
 - Control hazards, which arise only when branches are mispredicted, the penalty is 13 cycle.

Arm Cortex-A series: Multicore

Cortex™ -A9

ARM CoreSight™ Multicore Debug and Trace

ARMv7 32b CPU

NEON™
Data Engine

Floating Point
Unit

16-64k
I-Cache

16-64k
D-Cache

Core
1

2

3

4

ACP

SCU

Dual 64-bit AMBA3 AXI

Arm Cortex-A series: Multicore

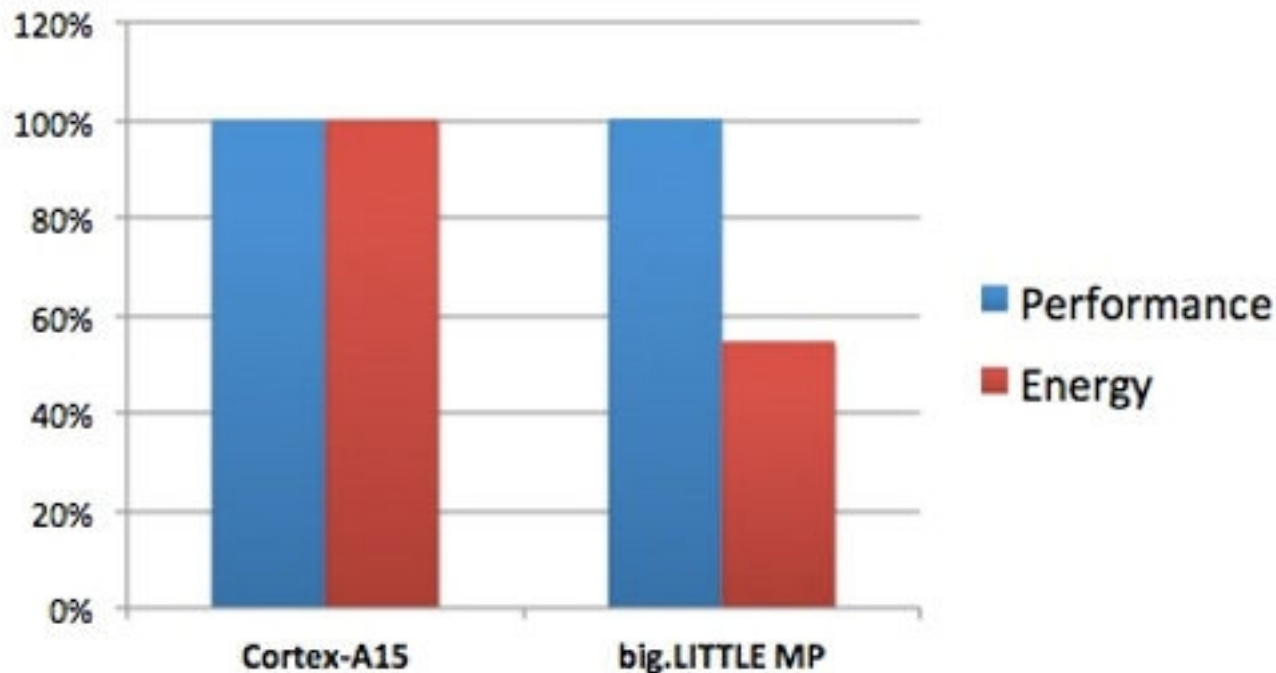
Multicore configurations are controlled and managed by the Snooping Control Unit (SCU). The SCU makes sure that level 1 cache coherence is achieved. Additional levels of coherence are achieved with an Accelerator Coherence Port (ACP).

Arm Cortex A series: Multicore

- big.LITTLE technology : a powerful processor is paired with a less powerful processor;
eg. A15 and A7, or the A53 and A57

Arm Cortex A series: big.LITTLE

big.LITTLE Performance and Energy - Web Browsing + mp3 audio



The processing is divided between the two processors to achieve increased efficiency but no decrease in performance.

Memory: A8 VS Intel i7

- Cortex-A8

	Size	Associativity	Latency
L1	16 / 32KB	4-way	Two words per cycle
L2	0 / 128 / 256 / 512 / 1024KB	8-way	
TLB	32	Fully associative	

- Intel i7

	Size	Associativity	Latency
L1	32KB	4-way I 8-way D	4 cycles, pipelined
L2	256KB	8-way	10 cycles
L3	2MB per core	16-way	35 cycles
ITLB	128	4-way	1 cycle
DTLB	64	4-way	1 cycle

Cortex-A8 Features

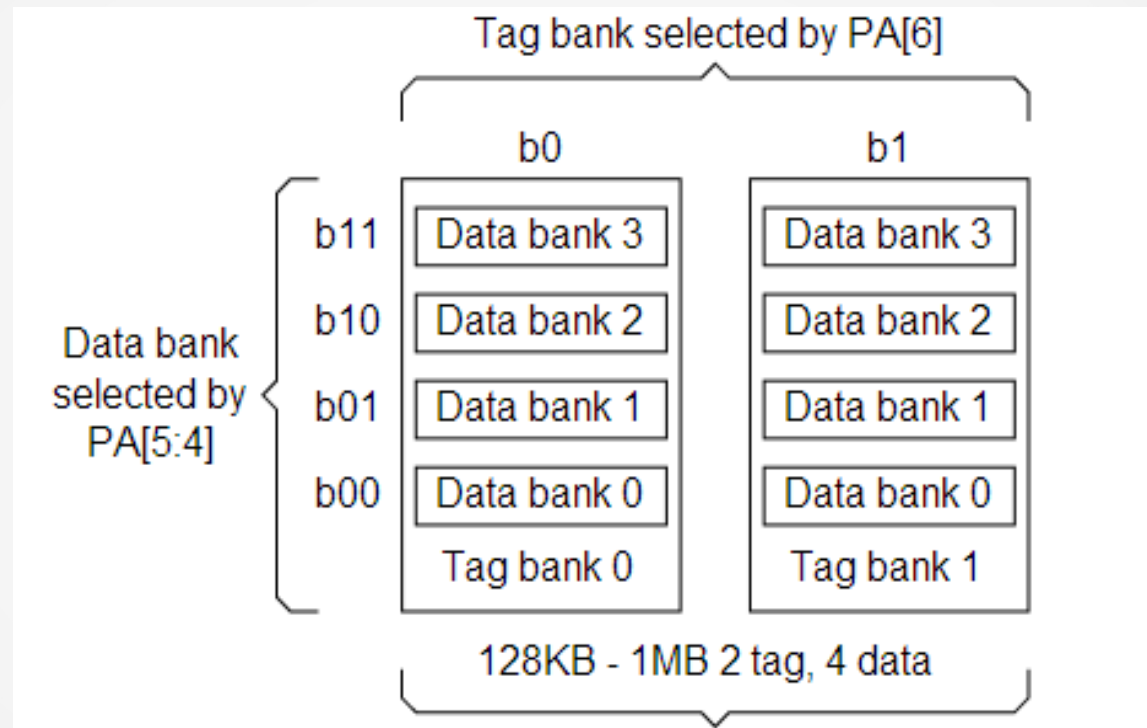
- L1 Caches

- physically tagged, and virtually indexed for instruction and physically indexed for data
- fixed line length of 64 bytes
- two words per cycle
- parity error detection

- L2 Cache

- physically indexed and tagged
- fixed line length of 64 bytes
- programmable preloading engine
- parity detection on the tag arrays
- Error Correction Code on data arrays
- partitioned into multiple banks to enable parallel operations

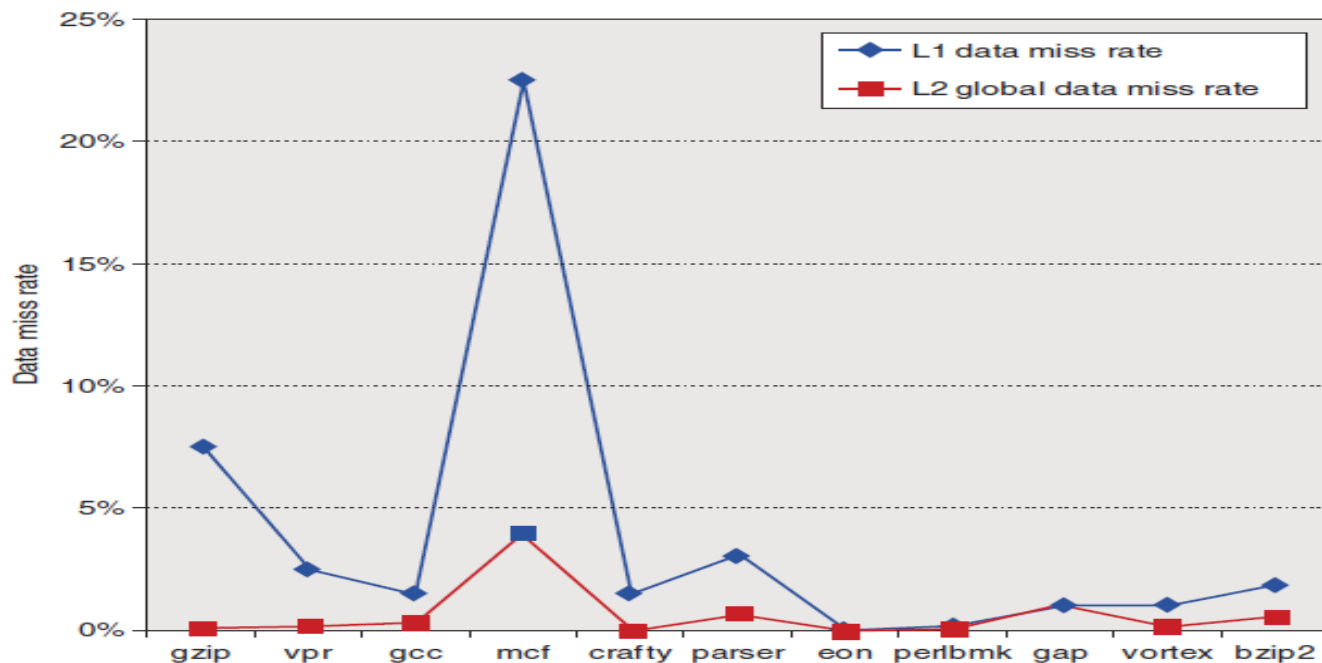
Cortex-A8 Features



Structure of L2
Cache

Cortex-A8 Performance

- simulated with 32 KB primary caches and a 1 MB eight-way set associative L2 cache using the integer Minnespec benchmarks
- instruction cache miss rates are close to zero for most and under 1% for all of them
- For the data cache test, there are significant L1 and L2 miss rates



Intel i7 Features

- L1 instruction cache, L1 data cache, and a L2 cache in each core
- support up to three memory channels of bandwidth over 25 GB/sec
- 48-bit virtual addresses and 36-bit physical addresses, a maximum physical memory of 36 GB

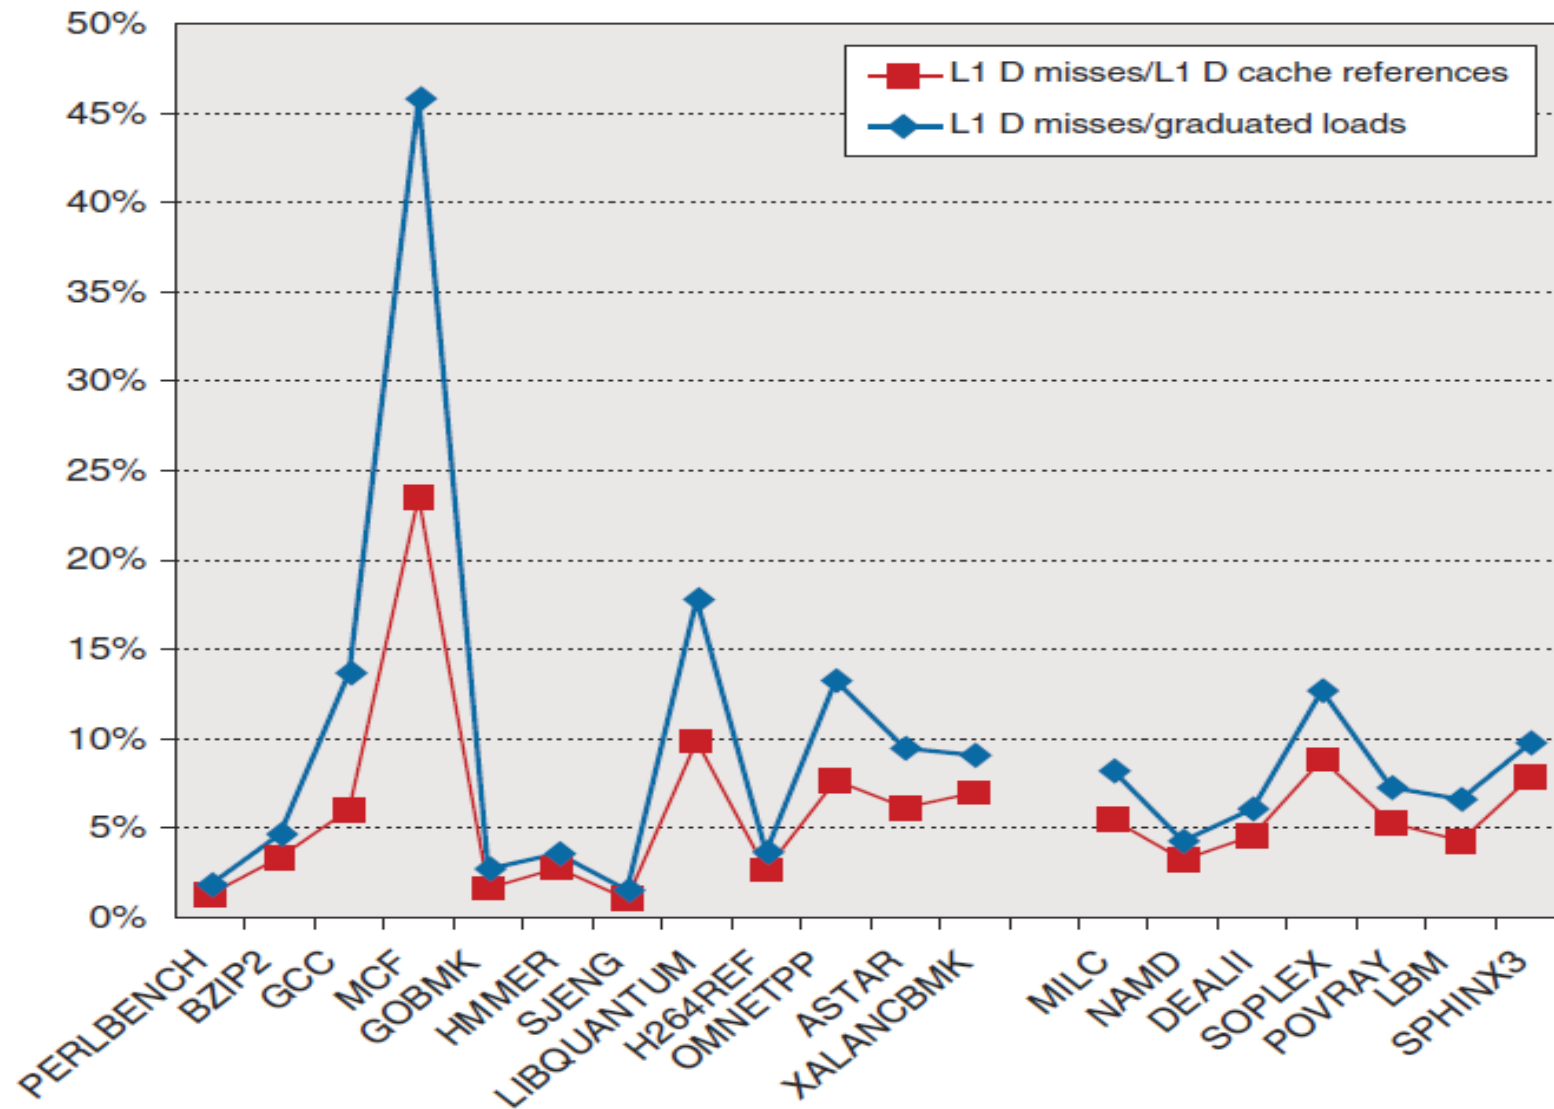
I7 Level 1 Data Cache Features

- a write-back write-allocate cache
- Store Forwarding - forward data directly from the store operation to load
- Memory Disambiguation - predict that a load does not depend on a preceding store
- Data Prefetching

Intel i7 Performance

- evaluated by 19 of the SPEC CPU2006 benchmarks
- L1 instruction cache miss rate varies from 0.1% to 1.8%, averaging just over 0.4% - Since the i7 does not generate individual requests for single instruction units, but instead prefetches 16 bytes of instruction data (between four and five instructions typically).
- L1 data cache misses are shown in two ways:
 - relative to the number of loads that actually complete - graduation
 - relative to all the L1 data cache accesses from any source.
- the miss rate when measured against only completed loads is 1.6 times higher (an average of 9.5% versus 5.9%)

Intel i7 Performance




ARM Exclusive Features - NEON

- NEON technology is used in ARM Cortex™-A series processors to enhance user's multimedia experiences.
- It can highly enhance the multimedia and signal processing algorithms which are frequently required by multimedia applications

ARM Exclusive Features - NEON

- The Advanced SIMD instructions perform packed SIMD operations:
 - Registers are considered as vectors of elements of the **same data type**.
 - Instructions perform the **same operation** in all lanes.

ARM Exclusive Features - NEON



ARM Exclusive Features - VFP

- ARM Floating Point architecture (VFP) provides hardware support for floating point operations ARM Cortex™-A series processors.
- VFP architecture v3 is an enhancement to v2:
 - Double the double-precision registers
 - Instructions of fixed-point and floating-point conversion

ARM Exclusive Features - VFP

- ARM Cortex-A8 has a cut down *VFPLite* module instead of a full VFP module, and require roughly ten times more clock cycles per float operation

ARM Cortex-A* series Processor

Thanks