

A Survey on ARM Cortex A Processors

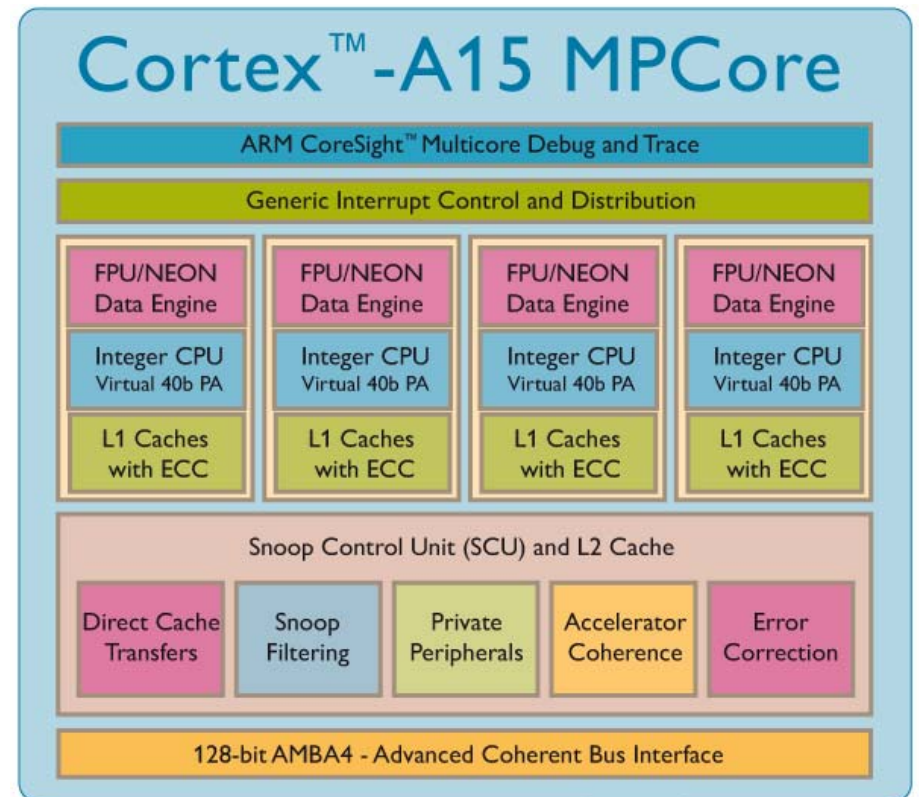
Wei Wang
Tanima Dey

Overview of ARM Processors

- Focusing on Cortex A9 & Cortex A15
- ARM ships no processors but only IP cores
 - For SoC integration
- Targeting markets:
 - Netbooks, tablets, smart phones, game console
 - Digital Home Entertainment
 - Home and Web 2.0 Servers
 - Wireless Infrastructure
- Design Goals
 - Performance, Power, Easy Synthesis

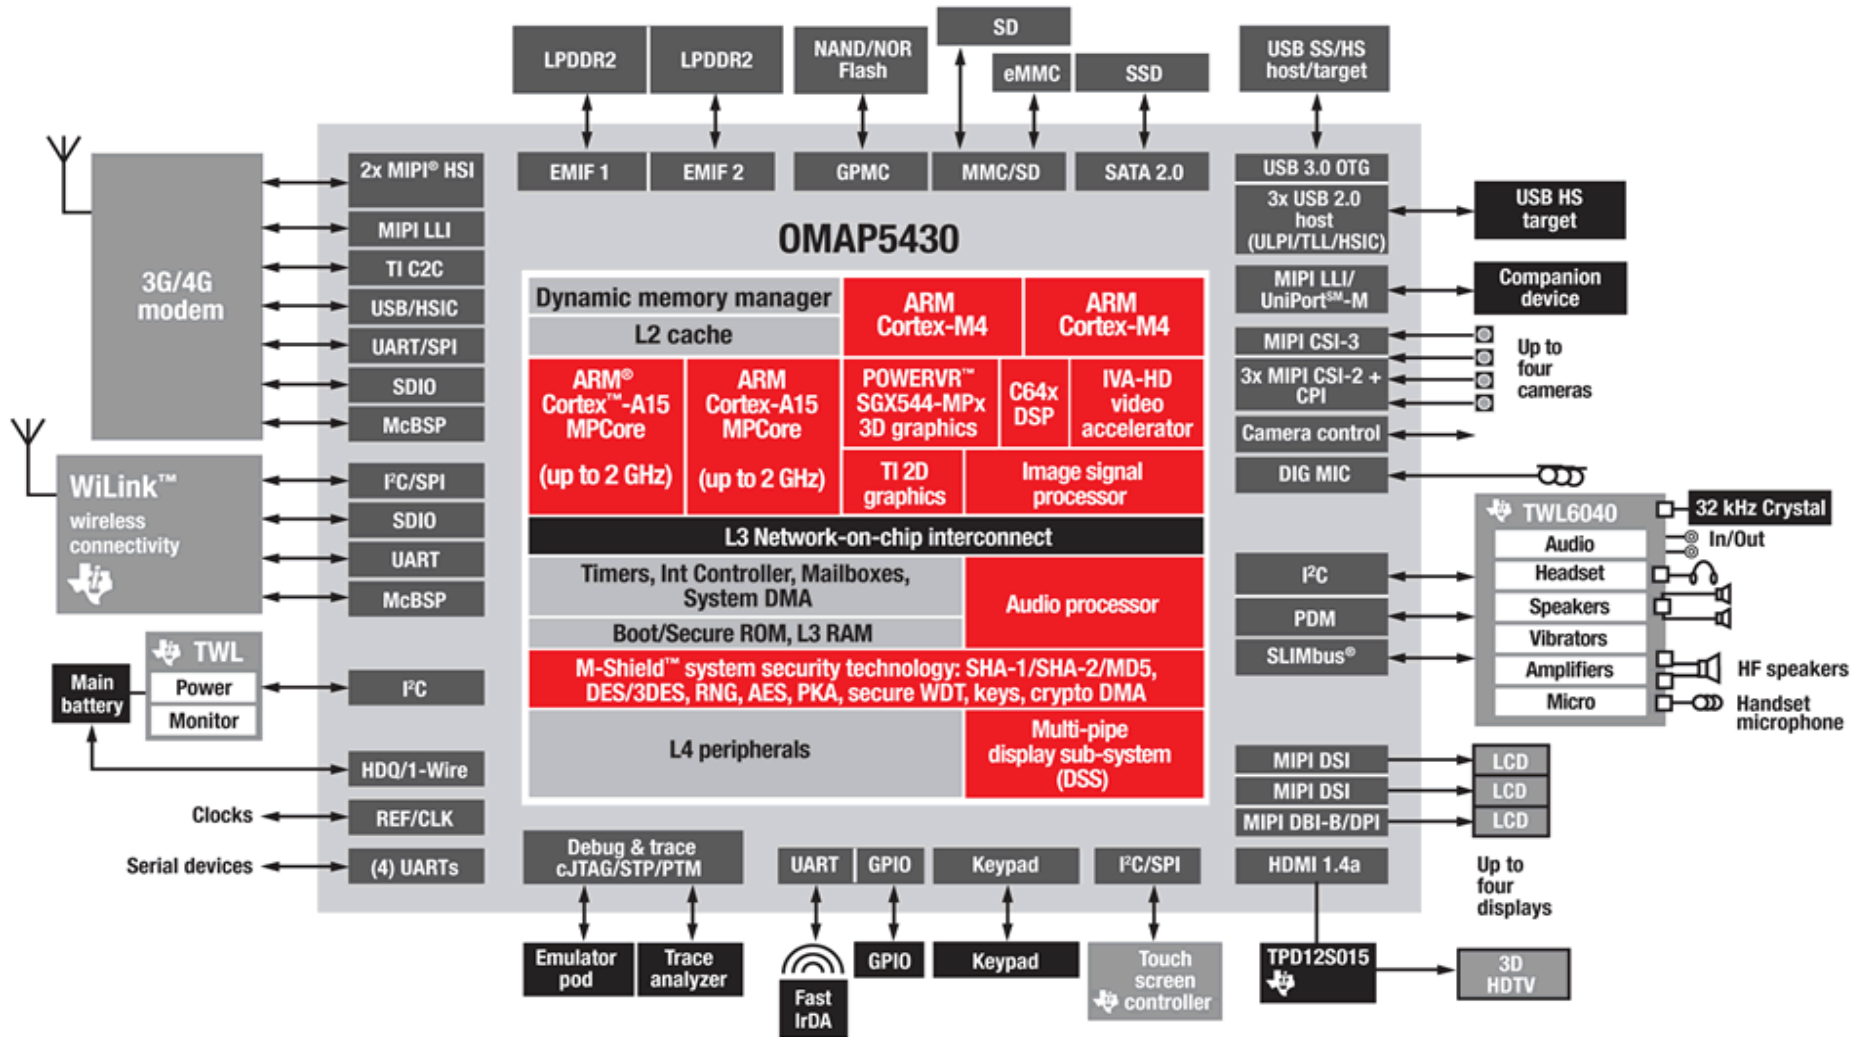
ARM Cortex A9/A15

- 1-4 Cores
- Out-of-Order Superscalar
- Branch predictor
- 32KB L1 I/D caches
- ~4MB L2 caches with Coherency
- NEON(SIMD) & FPU
- 32/28nm (A15)
45nm (A9)



Texas Instrument OMAP5

TI OMAP5430 SoC



Comparison of ARM, Atom, i7

	Cortex A15 (no L2, 32nm)	Cortex A9 (no L2, 40nm)	Atom N270 (45nm)	i7 960 (45nm)
Number of Cores	2 (4 maximum)	2 (4 maximum)	1 Core, 2 HT threads	4 Cores, 8 HT threads
Frequency	1Ghz – 2.5 Ghz	800Mhz (Po) 2Ghz (Per)	1.6 Ghz	3.2 Ghz
Out-of-Order?	Yes	Yes	No	Yes
L1 cache size	32KB I/D	32KB I/D	32KB I/D	32KB I/D
L2 cache size	N/A	N/A	512KB	1MB + 8MB L3
Issue Width	4	4	2	4?
Pipeline Stages	?	8	16	14 ~ 24 (?)
Supply Voltage	?	1.05V (Per)	0.9 – 1.1625 V	0.8-1.375 V
Transistor Count	?	26,00,000?	47,000,000	731,000,000
Die size	?	4.6 mm ² (Po) 6.7 mm ² (Per)	26 mm ²	263 mm ²
Power Consumption	?	0.5 W (Po) 1.9 W (Per)	2.5W (TDP)	130W (TDP)

Comparison of ARM SoC, Atom, i7

	TI OMAP5 (28nm)	Nvidia Tegra 2 (40nm)	Atom N450 (45nm)	I7 2600S (32nm)
CPU Cores	2 x A15 2 x M4	2 x A9	1 Core, 2 HT threads	4 Cores, 8 HT threads
CPU Freq.	2Ghz (A15)	1Ghz	1.66Ghz	2.6Ghz
GPUs ASICs	Video, Audio, Encryption, Display, 2D/3D	8x GPUs, Audio, Video, ISP	1 GPU	1 GPU
L2	?	1MB	512KB	1MB+8MB
Die Size	?	49mm ²	66mm ²	?
Transistors	?	260,000,000	123,000,000	?
Package Size	17 x 17 mm ²	23 x 23 mm ²	22 x 22 mm ²	37.5 x 37.5 mm ²
Power Consumption	?	150~500mW ?	5.5W (TDP)	65W (TDP)

Power/Performance Optimization as a SoC

- Application-specific SoC design
 - Integrate different ASICs
 - Customize Cortex Processors
 - Reduced memory bandwidth & frequency
- Mixing High V_t / Low V_t transistors
- Twisting floorplan, routing, clock tree design
- Power gating/Clock gating/DVFS
 - Four modes: Run, Standby, Dormant, Shutdown
 - Fine-grained pipeline shutdown
 - Faster register save and restore (state save/restore)
 - Power domains & voltage domains

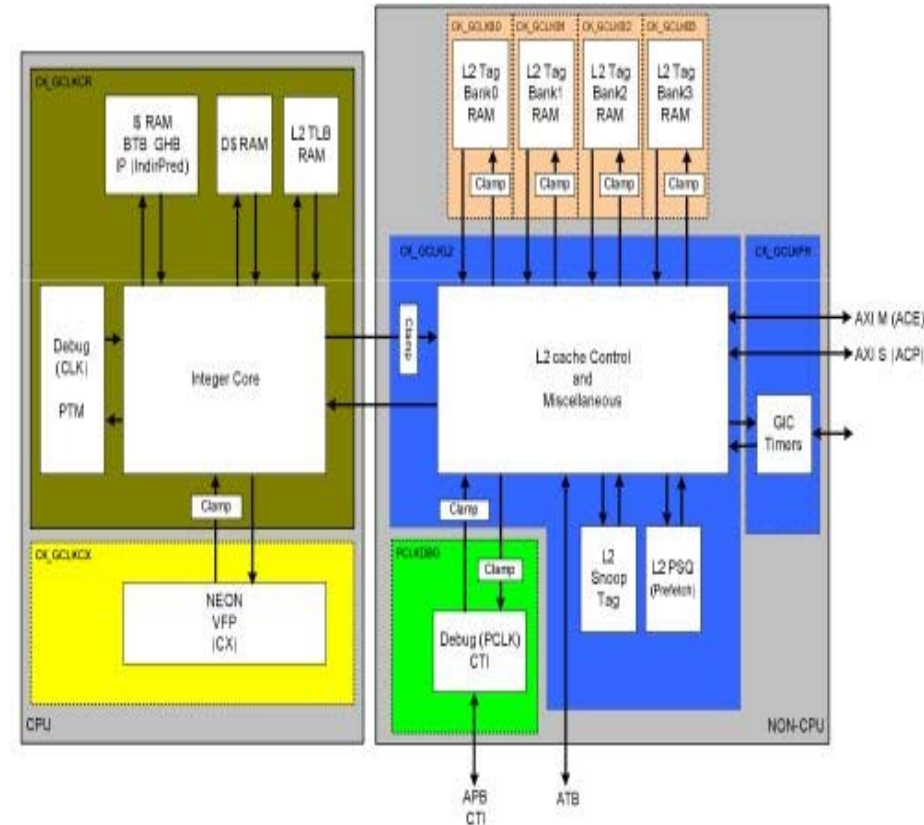
Power Saving as SoC: Power Gating

- Different power domains

- Cores
- NEON/VFP
- Debug Interface
- L2 cache tags (per bank)
- L2 cache control
- Interrupt Controllers

- Impact of power gating

- 3% reduction in performance
- 2% increase in area
- 4% increase in dynamic power
- 95% decrease in power when turned off



Power/Performance as a CPU

- Performance Enhancement (power hungry techniques)
 - Dynamic issue design
 - 4-way superscalar
 - Complex Branch predictor
 - Large L1/L2 caches
- Power savings
 - Accurate branch prediction
 - Micro TLB
 - RISC
 - SIMD, Jazzelle RCT etc.

ARM Instruction Set Architecture

- ARM processor architecture supports 32-bit ARM and 16-bit Thumb ISAs
- ARM architecture -- RISC architecture
 - Large uniform register file
 - Load/store architecture
 - Simple addressing modes
 - Auto-increment and auto-decrement addressing modes
 - Load and Store multiple instructions
- Instructions can also be "conditionalised" based on condition code in Application Program Status Register

ARM Instruction Set Architecture

- Thumb
 - Extension to the 32-bit ARM architecture
 - Features a subset of the most commonly used 32-bit ARM instructions compressed into 16-bit opcodes
 - Excellent code-density for minimal system memory size, reduced cost and power efficiency
 - Designers have the flexibility to emphasize performance or code size
 - "Thumb-aware" core is a standard ARM processor fitted with a Thumb decompressor in the instruction pipeline
- ARM uses the **Universal Assembly Language**

DSP

- ISA extension
- Features: new instructions to load and store pairs of registers, 2-3 x DSP performance improvement over ARM7
- Eliminates the need for additional hardware accelerators
- Provides high performance solution with low power consumption
- Reuses existing OS and application code
- Supports including servo motor control, Voice over IP (VOIP) and video & audio codecs

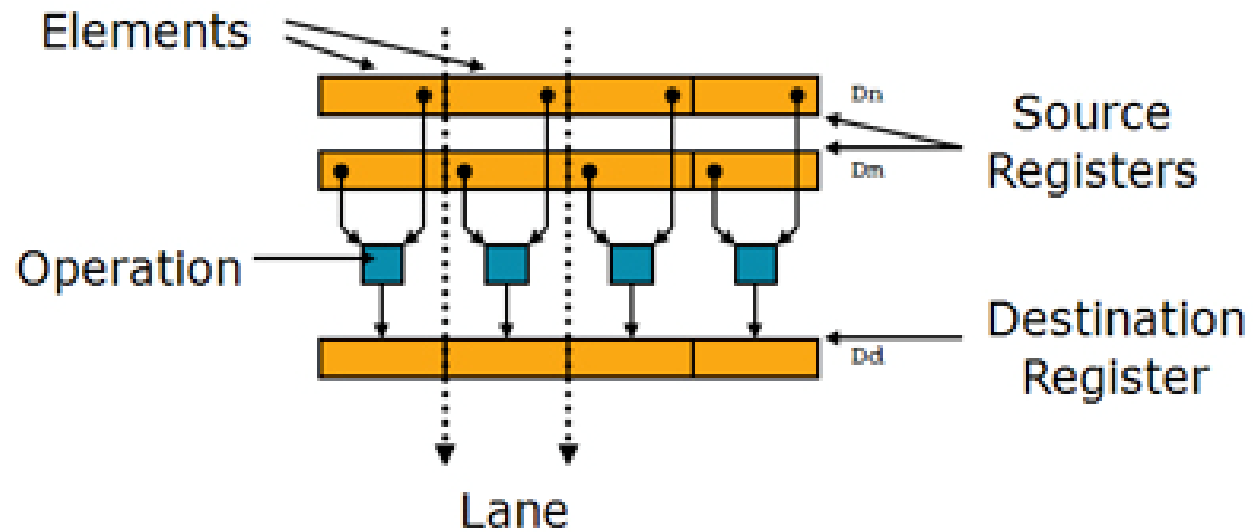
SIMD

- 75% higher performance for multimedia processing in embedded devices
- “Near zero” increase in power consumption
- Simultaneous computation of 2x16-bit or 4x8-bit operands
- Offers single tool-chain and processing device, transparent of OS

NEON

- Cleanly architected and works seamlessly with its own independent pipeline and register file
- Large NEON register file with its dual 128-bit/64-bit views enables efficient handling of data
 - Minimizes access to memory, enhancing data throughput
- Designed for autovectorizing compilers and hand coding
- Provides flexible and powerful acceleration for consumer multimedia applications
 - Supports the widest range of multimedia codecs used for internet applications

NEON



Vector Floating Point Architecture

- Coprocessor extension to the ARM architecture
- Supports floating point operations in half-, single- and double-precision floating point arithmetic
- Fully IEEE 754 compliant with full software library support
- Supports execution of short vector instructions but these operate on each vector element sequentially
- Three-dimensional graphics and digital audio, printers, set-top boxes, and automotive applications

Jazzelle

- Combined hardware and software solution for accelerating execution
- Software -- fully featured multi-tasking JVM
- Hardware -- coprocessor CP14 provides support for the hardware acceleration
- Jazzelle DBX technology for direct bytecode execution – Direct interpretation bytecode to machine code
- Jazzelle RCT technology supports efficient AOT and JIT compilation with and beyond Java

Jazelle

- Jazelle DBX and RCT are cache and memory efficient, maintaining low power
- Jazelle DBX is a robust and proven solution and easy to integrate
- Jazelle RCT provides an excellent target for any run-time compilation technology
- Developers' Flexibility
 - Resource constraint device: Jazelle DBX only
 - On high-end platforms, Jazelle RCT alone with JIT and AOT

Conclusion

- Aggressive power hungry design targeting at high single thread performance
 - Out-of-Order Execution
 - Wide superscalar
 - Large caches with coherency protocols
- Power saving techniques for ARM CPUs
 - RISC
 - ISA Optimization: Thumb, Thumb2, ThumbEE
 - Application-Specific Components: SIMD, DSP, VFPU, Jazzelle
- Power saving techniques for SoC chips
 - Fine-grained power gating & clock gating & DVFS
 - Fine-grained pipeline shutdown
 - fast registers saving/restoring
 - Customizable CPU components
 - Mixing high Vt and low Vt transistors

Reading materials

- ARM Cortex-A9 Technical Reference Manual
- ARM Cortex-A9 MPCore Technical Reference Manual
- Keys to Silicon Realization of Gigahertz Performance and Low Power ARM Cortex-A15, Lamber A. et. al., ARM Technology Conference 2010
- 2GHz Capable Cortex-A9 Dual Core Processor Implementation, http://www.arm.com/files/downloads/Osprey_Analyst_Presentation_v2a.pdf
- Circuit Design: High performance AND low power, the ARM way, http://www.arm.com/files/downloads/Enabling_High_Performance_CPU_Implementation.pdf
- ARM MPCore Architecture Performance Enhancement, http://www.arm.com/files/downloads/MPF_2008_Japan_-_ARM_Cortex-A9_Final.pdf
- Cortex-A9 Processor Microarchitecture, http://www.arm.com/files/downloads/Cortex-A9_Devcon_2007_Microarchitecture.pdf
- Details of a New Cortex Processor, Revealed, http://www.arm.com/files/downloads/Cortex-A9_Devcon-talk_Introduction_FINAL-02.pdf
- ARM Cortex-A9 Performance, <http://www.arm.com/products/processors/cortex-a/cortex-a9.php>