

Is VLIW the Right Architecture for Embedded Processing?

Trevor Mudge

Bredt Professor of Engineering

Advanced Computer Architecture Lab

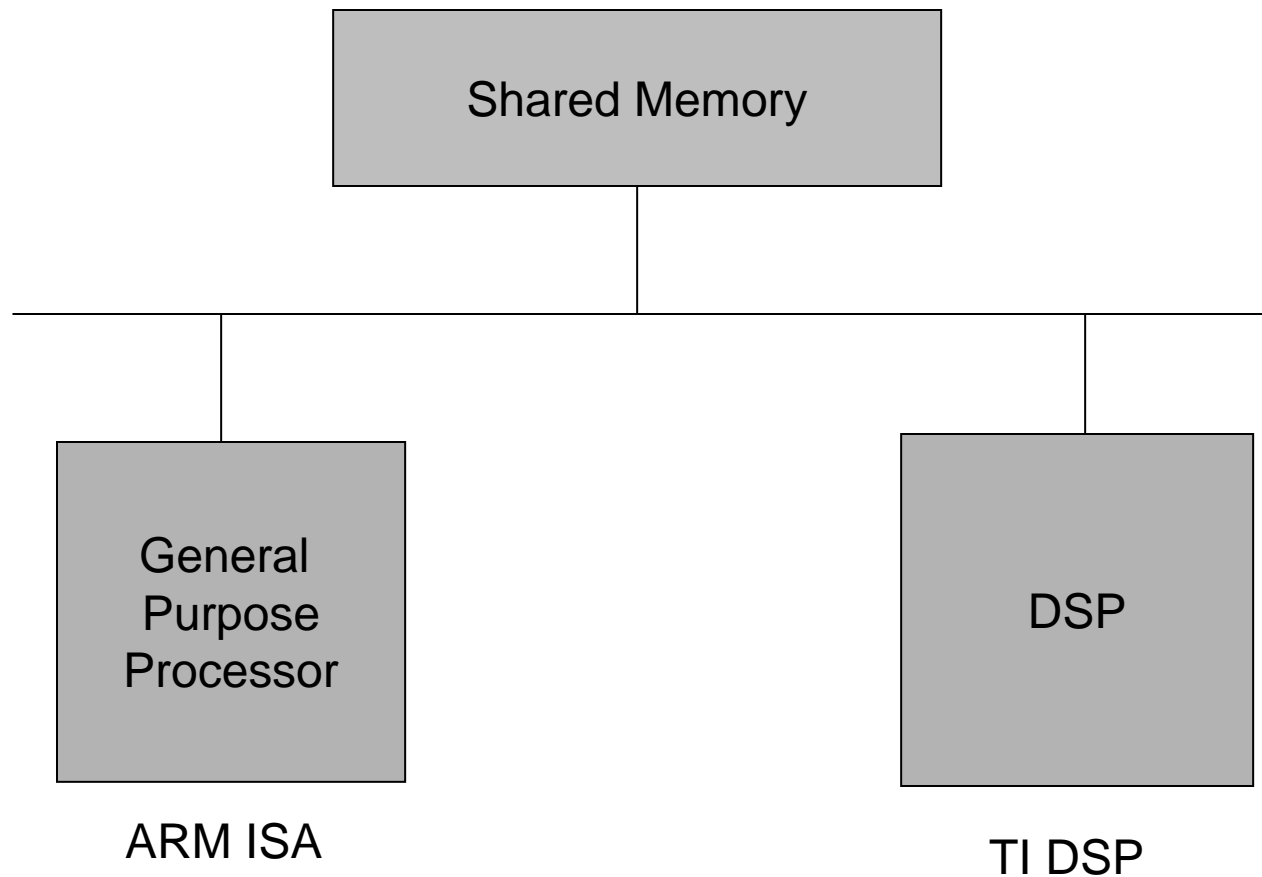
The University of Michigan of Engineering



With

- Prof. Wayne Wolf, Princeton Univ.
- Tiehan Lv, Princeton Univ.
- David Oehmke, Univ. Michigan
- Jeff Ringenber, Univ. Michigan

What Kinds of Embedded Systems?



Trends

- General purpose computers
 - Superscalar
 - Deeper pipelines
 - Out-of-order execution
 - Multi-issue
- Digital signal processors
 - VLIW
 - Simpler hardware/function unit
 - Require more compiler support
 - Program in micro-code

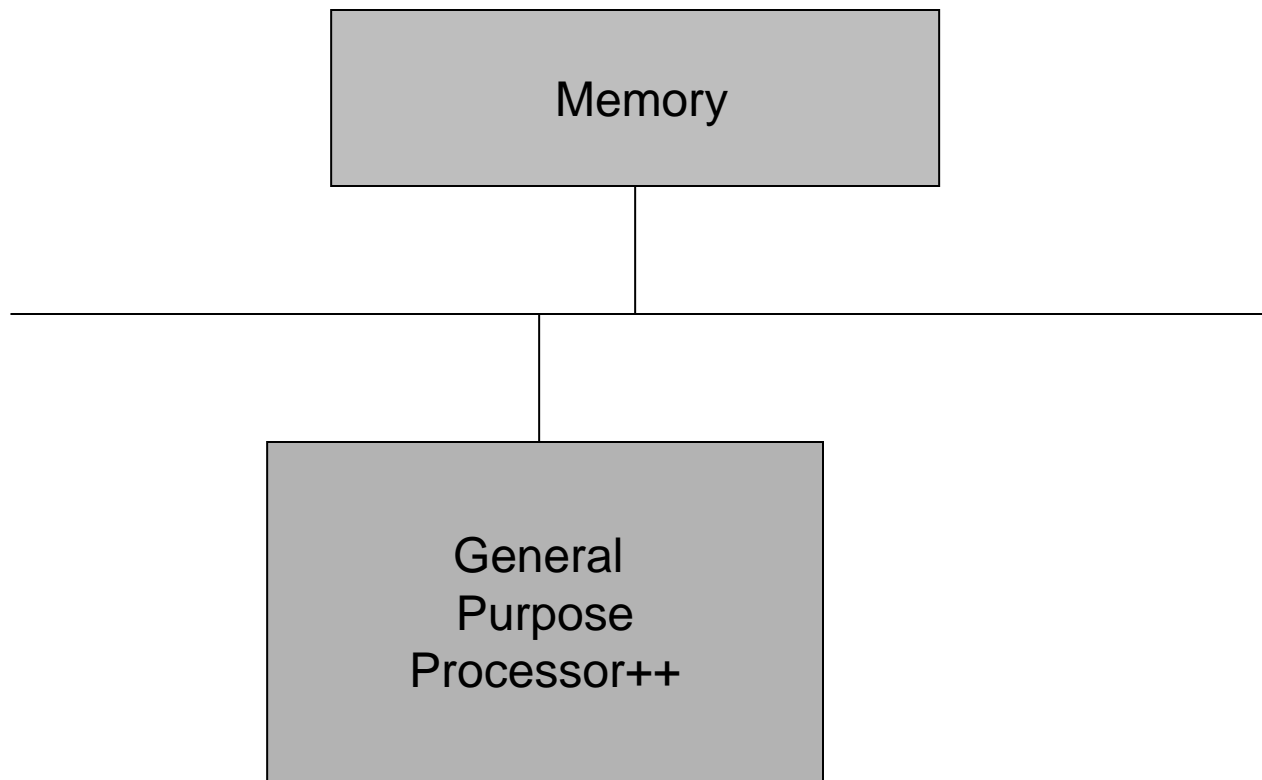
Observations

- VLIW machines unsuccessful in the GP world
- Itanium
 - Relaxed issue rules -> EPIC
 - Out-of-order memory support?
 - More like superscalar
- DSP solution to software
 - Libraries
 - Not sustainable

Thesis

- VLIW only makes sense in special purpose ultra-high performance data streaming apps
- Software is more expensive than hardware
 - Less automated
 - More faulty
 - Legacy languages with large time constants
- DSPs should look more like GP computers
 - Reduce VLIW features
 - More like GP computers
 - Easier to program

One Chip Solution?



Types of Computing

- Control intensive
 - Frequent conditional branches
 - (Super)scalar machines
- Data parallel
 - Loops
 - Vector/VLIW machines

Superscalar vs. VLIW

- Superscalar
 - complexity in the hardware
 - good performance on un-optimized software
 - known quantity
- VLIW
 - complexity in the compiler
 - requires hand coded kernels
 - portability compromised
 - life cycle costs

Software vs hardware

- Software – difficult to design and test
 - we accept buggy software
 - costly to produce
- Hardware – easier to design and test
 - we don't accept bugs
 - many steps are automated
- Conclusion
 - be conservative about moving complexity from hardware to software
 - In those applications that are almost all data parallel => VLIW
 - Small traces of control code => Superscalar

Smart Camera Algorithms

- Complete application
 - Background elimination
 - Skin area detection
 - Contour following
 - Super-ellipse fitting
 - Graph matching
- Mix of data parallel and control type algorithms

Processing a Frame




Original frame




Background elimination



Skin area detection



Contour following



Superellipse fitting

GSM Algorithms

- Coder and decoder – etsi.org
 - part of a digital cellular telecommunications system (GSM - Phase 2+) used for coding and decoding speech
 - specifically the GSM Enhanced Full Rate speech codec

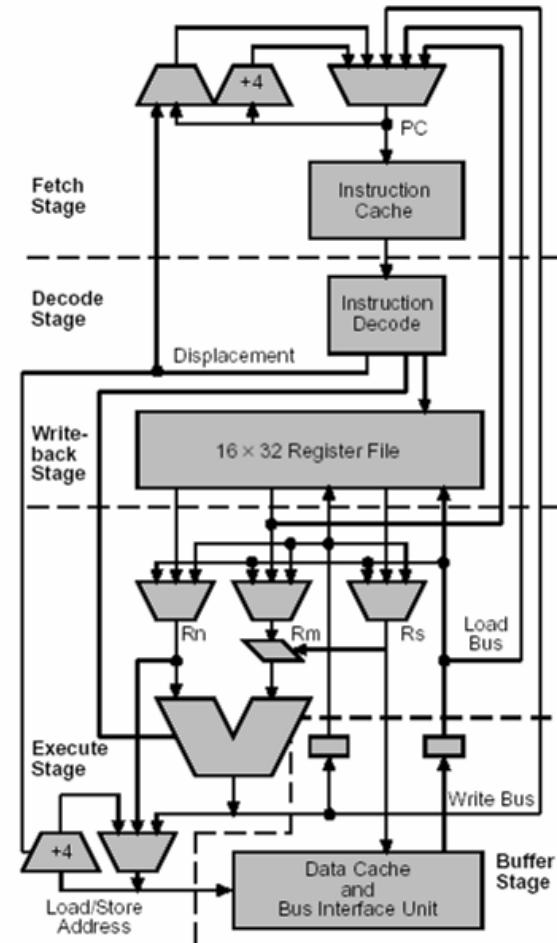
StrongARM

ARM7:

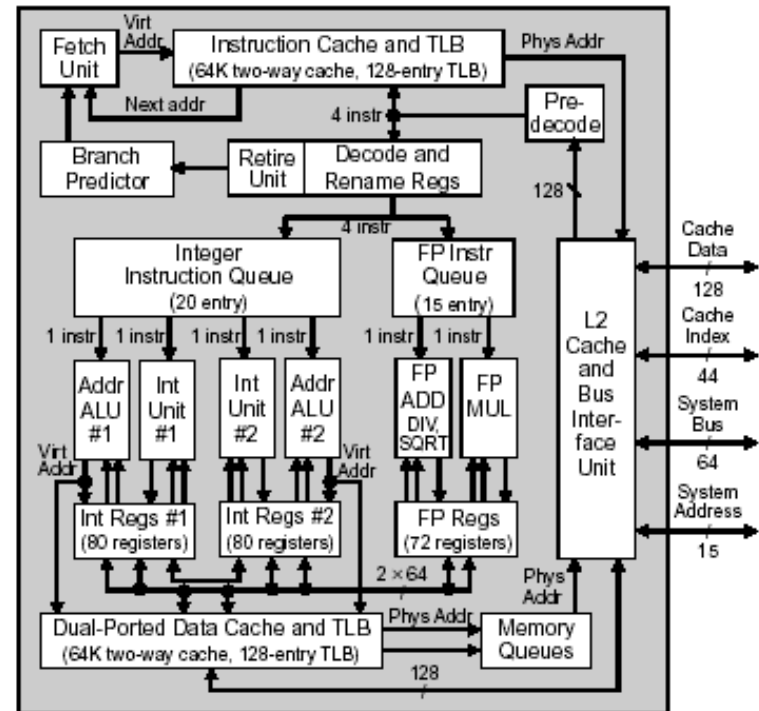
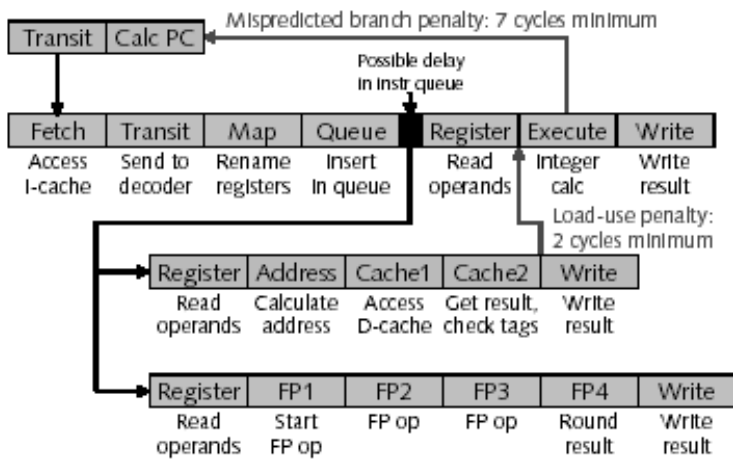
Fetch	Decode	Execute
latch instr	decode	shift/rotate ALU op D-cache access sign/zero ext commit result

StrongArm:

Fetch	Decode	ALU	Cache	Writeback
calc PC l-cache access	decode cc mux clk ctrl CAM bypass reg file ctrl	shift/rotate ALU op	D-cache access sign/zero ext post result	commit write



Alpha 21264



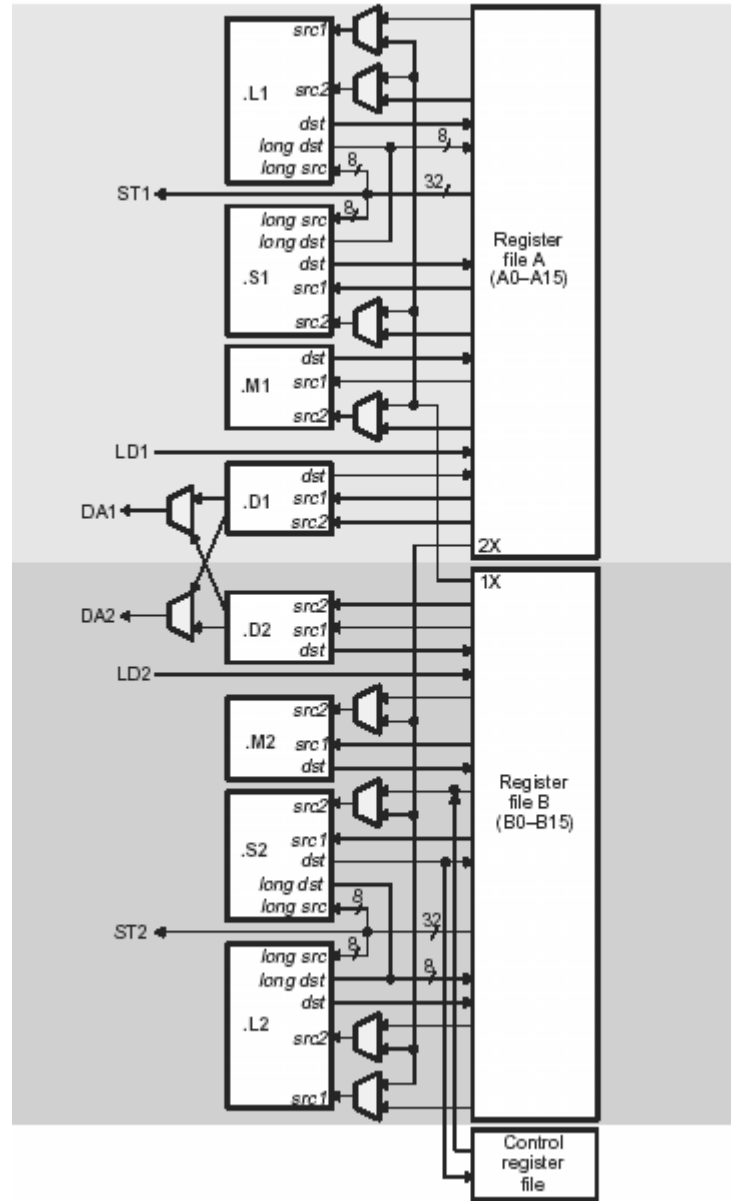
SA1 and Alpha Core Configs

	SA1 Core	Alpha Core
Instruction Fetch Queue Size	8	32
Extra Branch Misprediction Latency	1	3
Branch Predictor	Not Taken	Gshare (4096 entry, 8 bit hist)
BTB	NA	1024 sets, 4 way associative
Pipeline	in order	out of order
Decode Width	1	4
Issue Width	1	4
Commit Width	1	4
Register Update Unit Size	4	128
Load/Store Queue Size	4	64
Memory Disambiguation	perfect	perfect
Cache	perfect	perfect
Memory Latency	1	1
Memory Width	4 bytes	8 bytes
Pipelined Memory	no	yes
# Integer ALUs	1	4
# Integer Mult	1	1
# Memory Ports	1	2
# Floating Point ALUs	1	4
# Floating Point Mult	1	1

TI C62x

- VLIW
 - Up to 8, 32-bit instructions per cycle
 - RISC-like ISA
- 2 - Cluster Architecture
- Per cluster:
 - 16 General Purpose Registers
 - 4 Fully-Pipelined Functional Units
 - One crosspath to other cluster
- Predicated execution
- Multi-cycle latency instructions

Execution Core



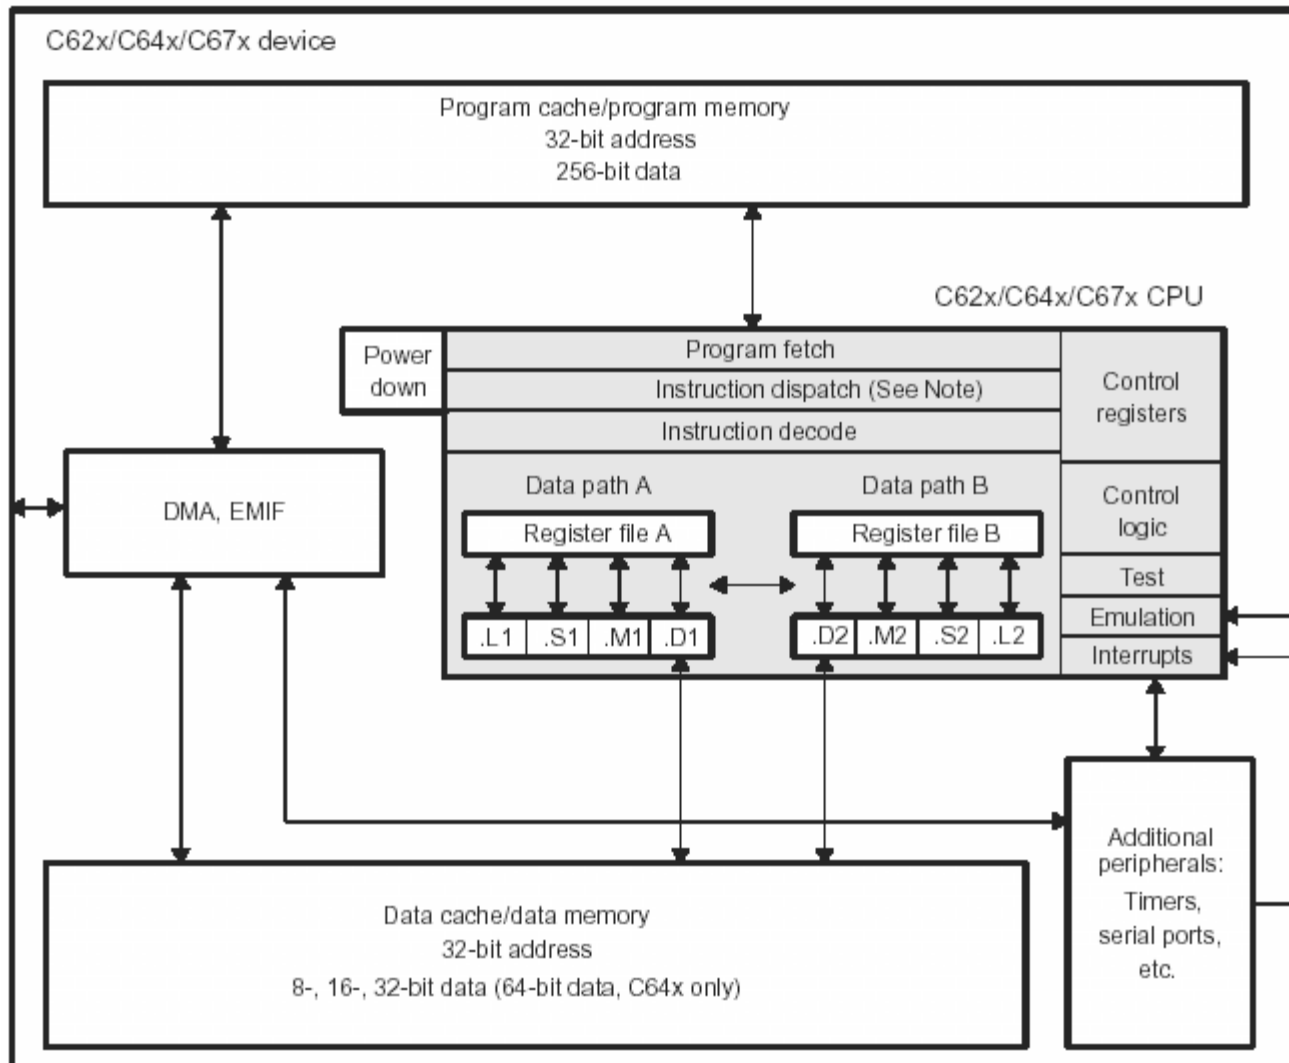
Functional Units

- L-Unit
 - 32/40-bit Arithmetic
 - 32-bit Logical Operations
 - 32/40-bit Compare Operations
 - Leftmost 1 or 0 counting for 32-bit
 - Normalization count for 32 and 40-bit
- D-Unit
 - 32-bit Add and Subtract (linear and circular addressing)
 - Loads and Stores with 5-bit constant offset
 - Loads and Stores with 15-bit constant offset (.D2 unit only)

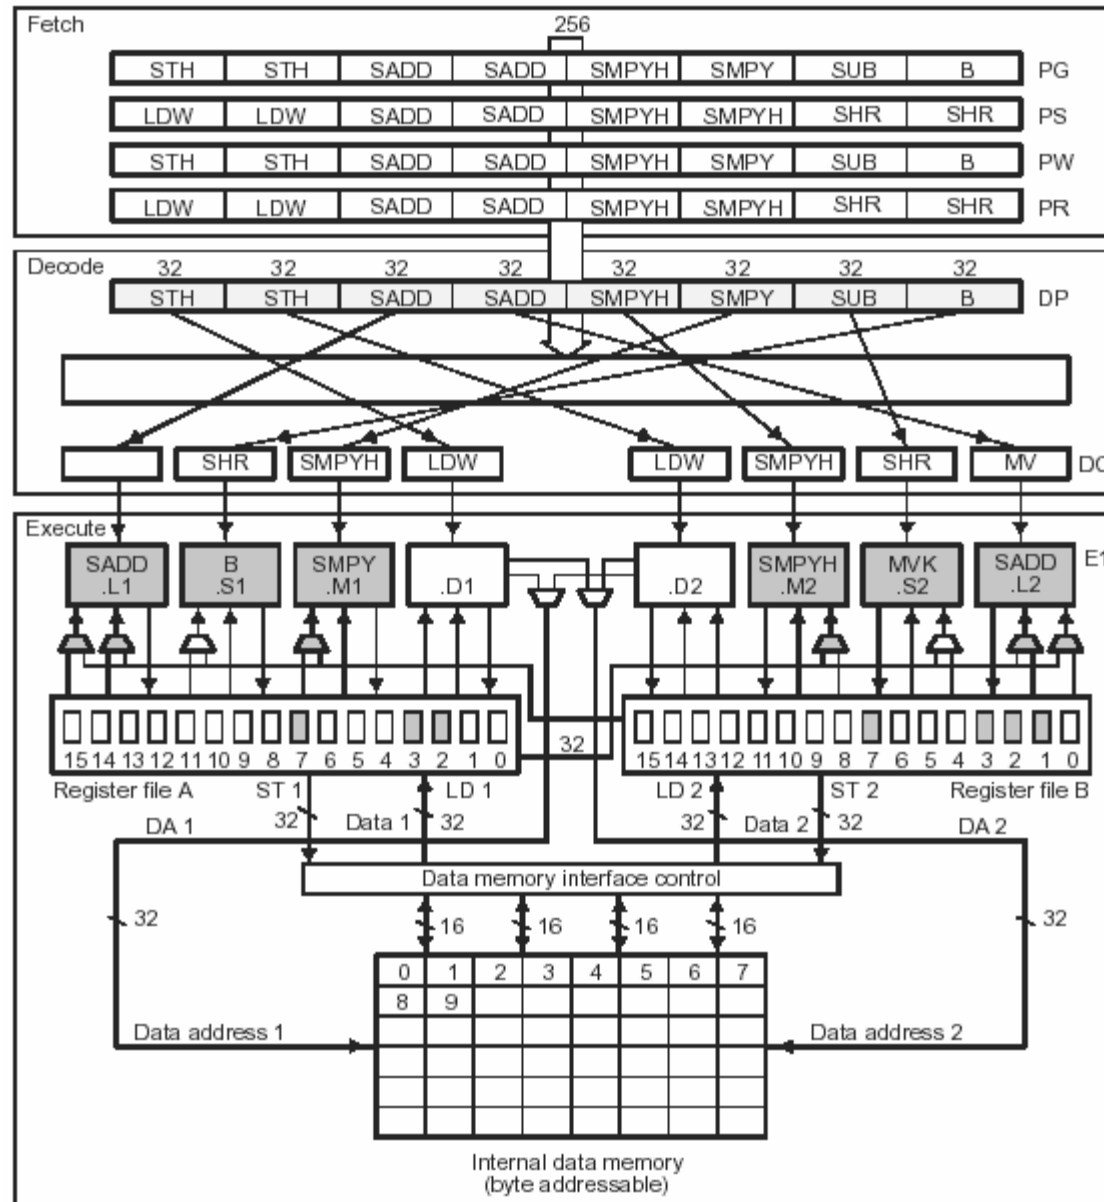
Functional Units

- S-Unit
 - 32-bit Arithmetic
 - 32-bit Logical Operations
 - 32/40-bit Shifts
 - 32-bit Bit-field Operations
 - Branches
 - Constant Generation
 - Control Register Access (.S2 unit only)
- M-Unit
 - 16x16 Multiply

The Core

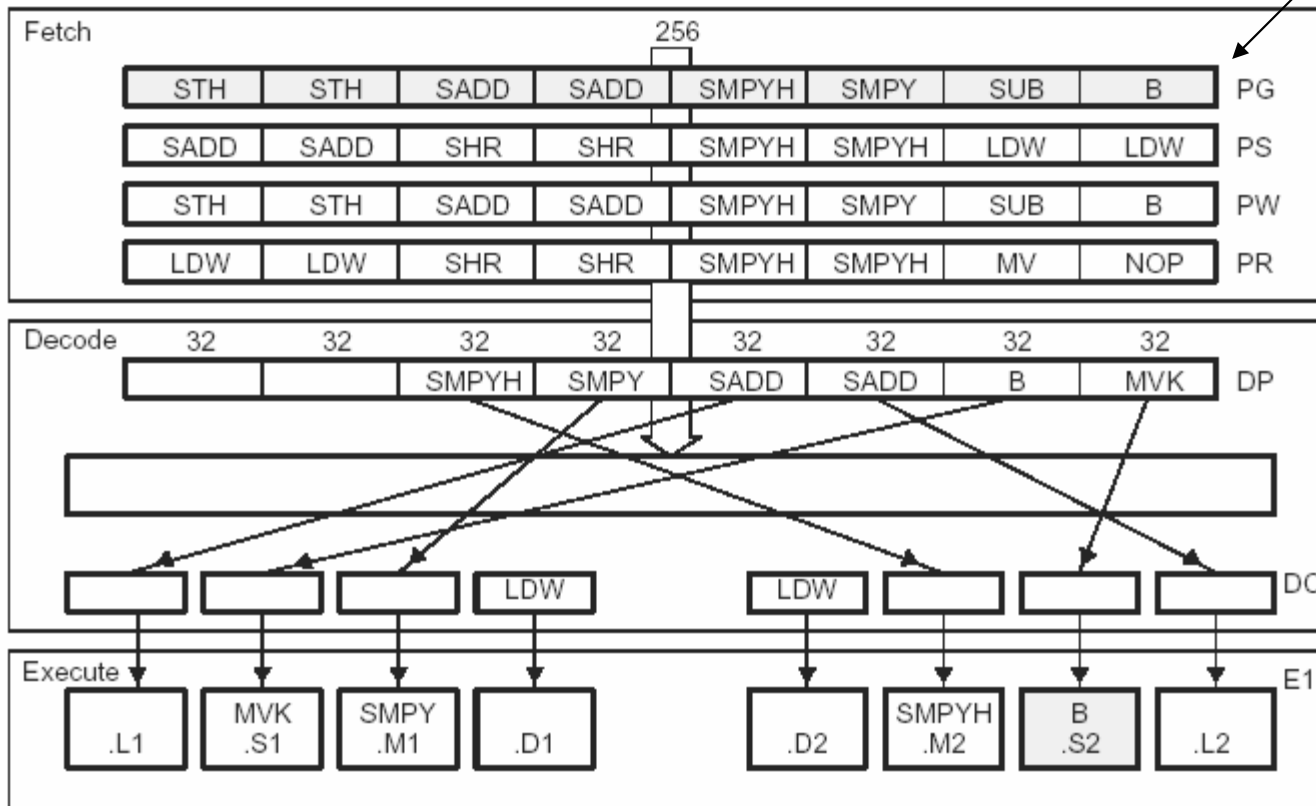


The Pipeline



Branch Execution

Execution Target



Clustering/Architectural Features

- Only one FU can use a crosspath at a time
- Limits placed on which FUs can use the crosspath for their srcs
 - L Unit: src1, src2
 - S Unit: src2
 - M Unit: src2
 - D Unit: None
- D Unit from other cluster can generate memory address for current cluster
- Load/Store with 15-bit offset only available to .D2 Unit and can only use registers B14/B15
- MVC (move to control register) only available to .S2

Code Compression Support

- Multi-cycle NOPs
 - Expands into up to 9 single cycle NOPs
 - Branches can interrupt the execution
 - Allows for compacting of strings of NOPs in the code

ISA

- RISC like
- Small set of fixed latency instruction types
 - Whole pipeline stalls on a cache miss
- Extras
 - Support for 40-bit arithmetic
 - Support for bit field manipulation (extract, set, clear, bit-count)
 - Support for saturation
 - Support for 5-bit constants for source 2 instead of a register


More ISA

- Limitations
 - C62x has no floating point
 - Missing Integer Divide instruction (emulation provided by function call)
 - Integer Multiplies are only 16-bit (have to emulate 32-bit multiplies)
 - No Branch-and-Link (have to explicitly store the return address)
- Predication
 - 5 GPR registers used (A1,A2,B0,B1,B2)
 - Check either for Zero or Not Zero



Simulation Results

- Added a validated TMS320C6200 core to SimpleScalar (www.simplescalar.org)
- SS includes StrongARM core
- SS includes Alpha 21264 core


Functional Unit Usage



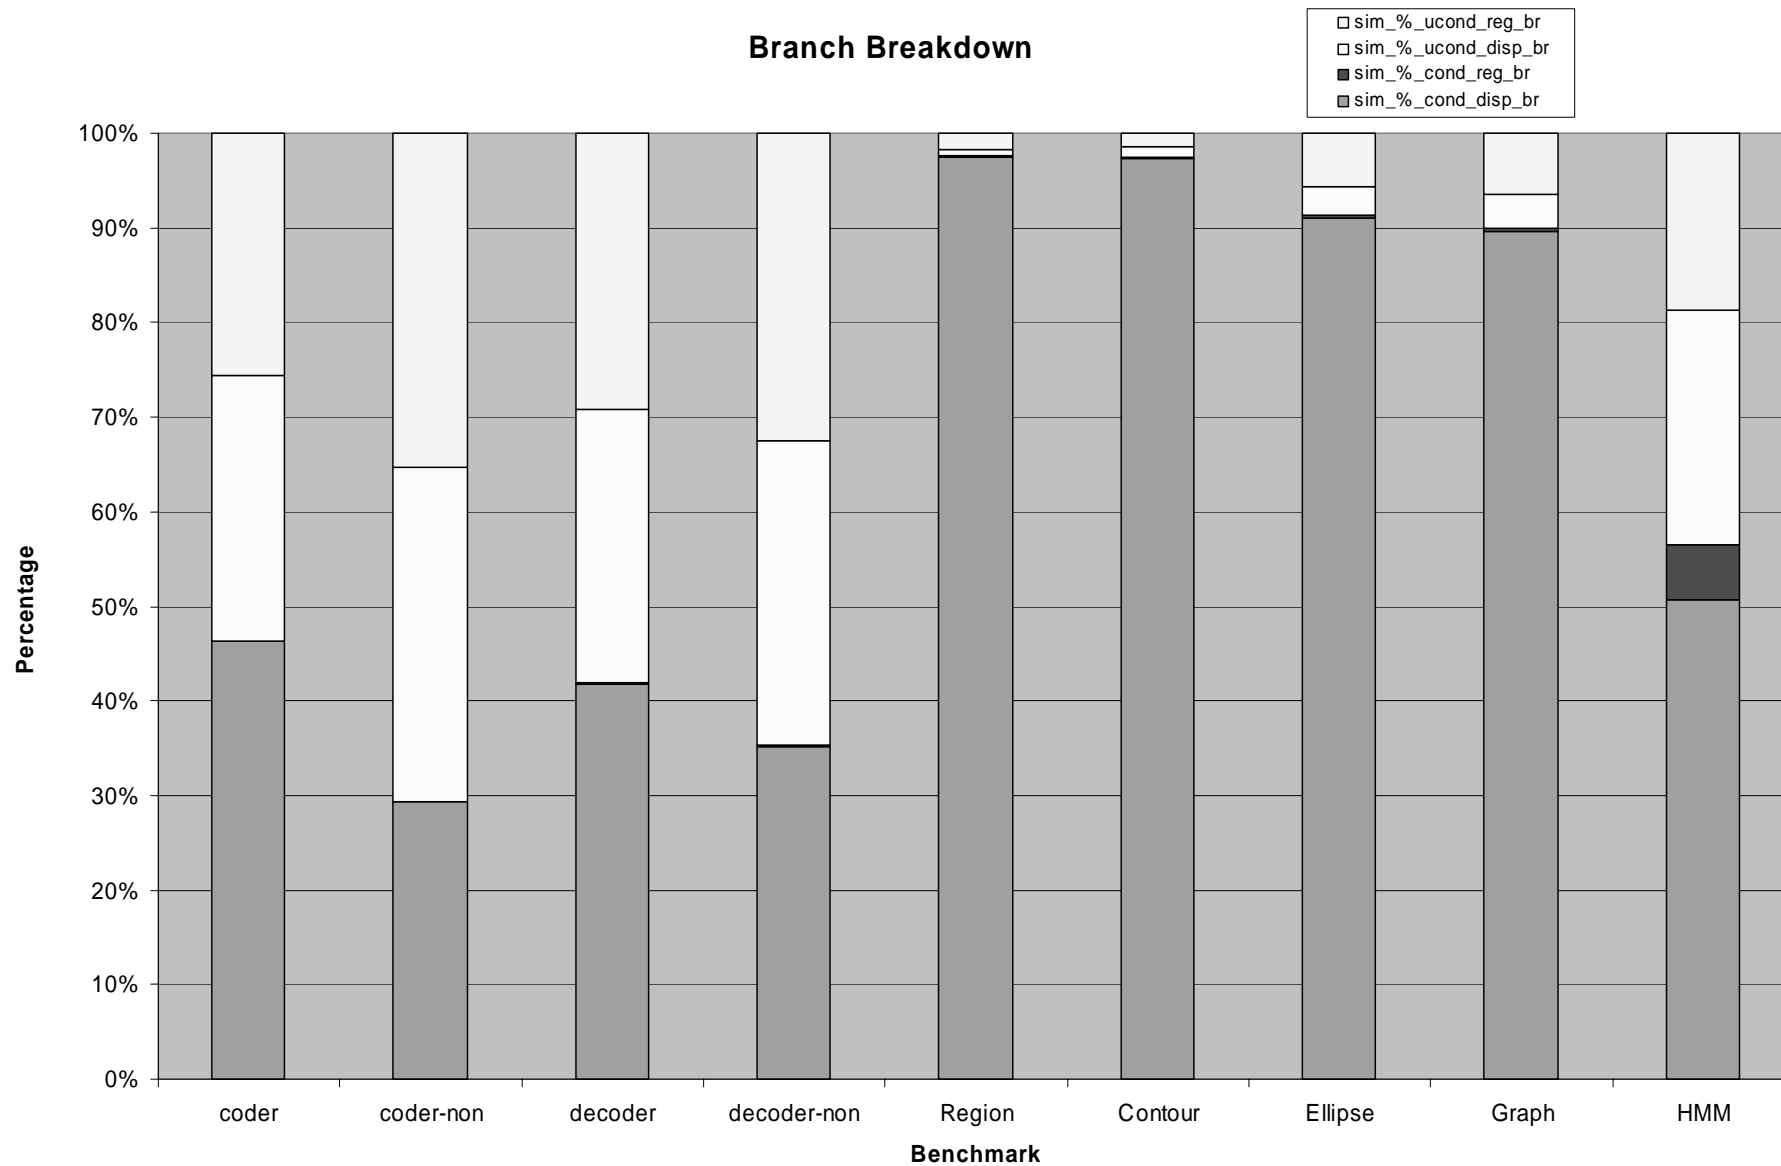
Cluster Utilization





Inst Class Breakdown





Branch Breakdown



NOP cycle info




IPC info

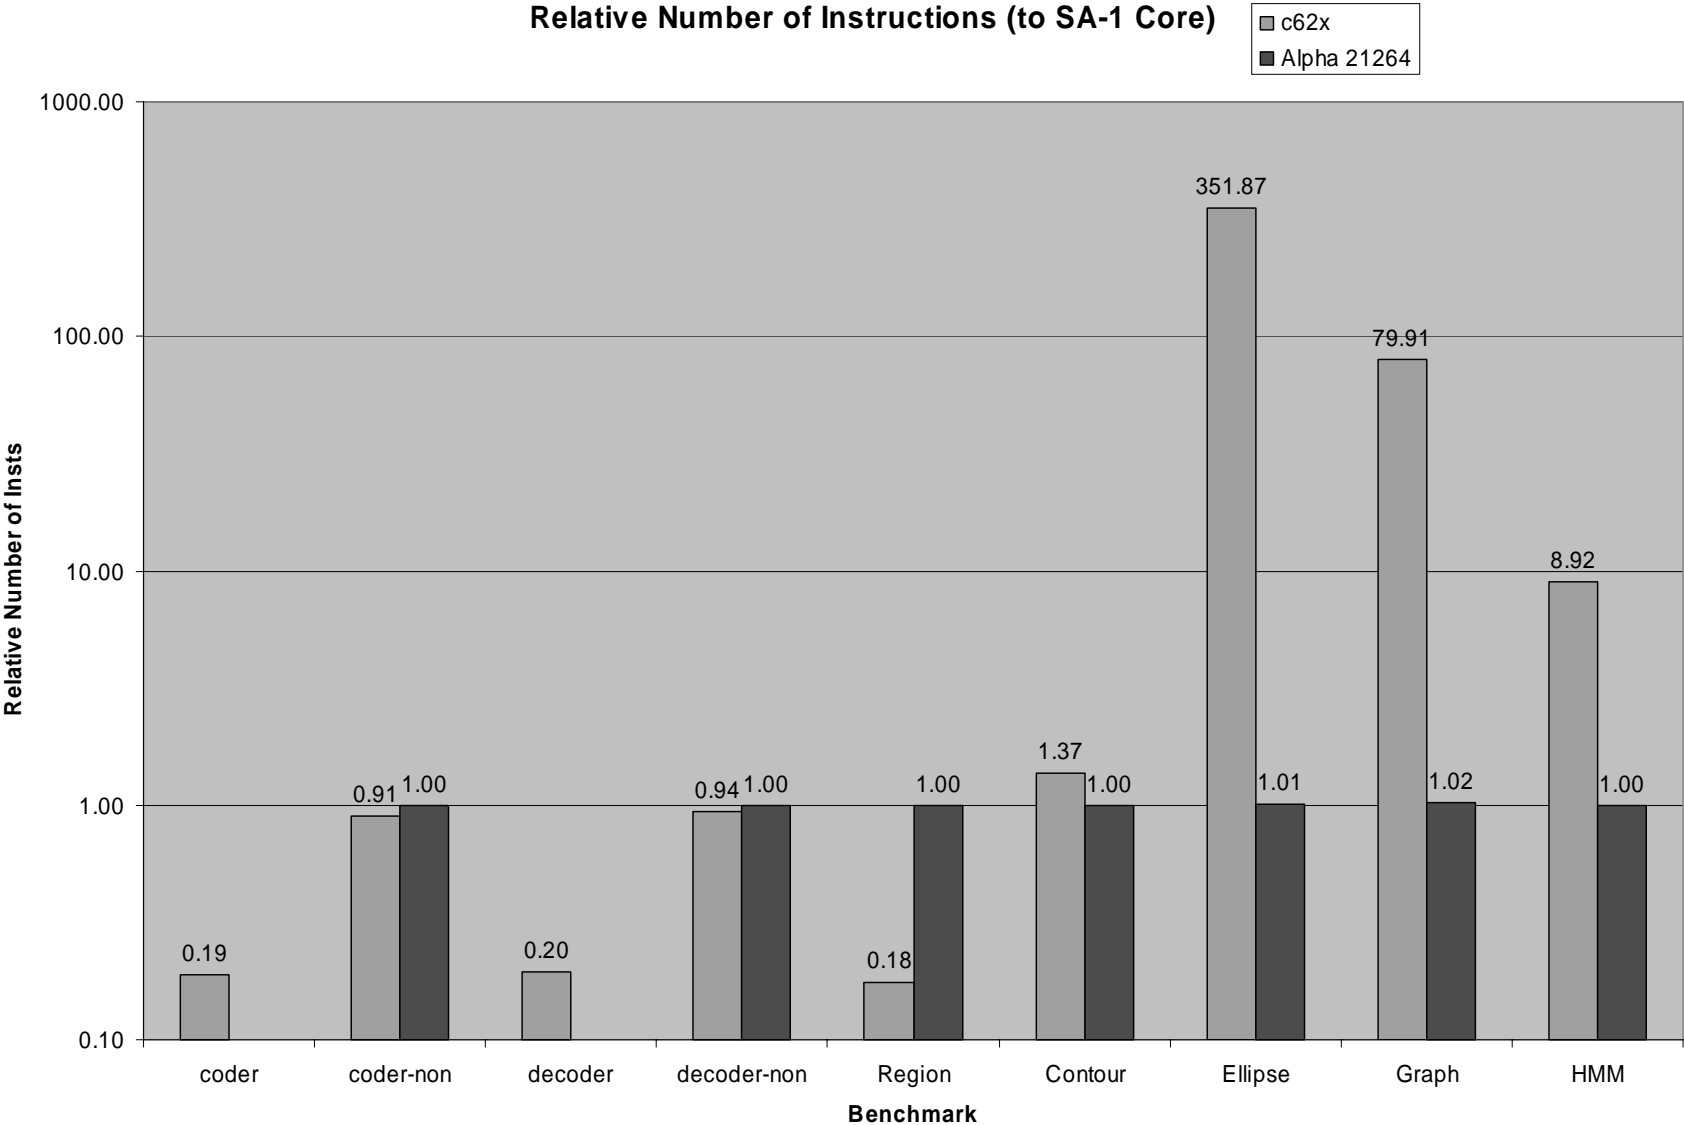


StrongARM vs C62 Stats


ARM Inst Class Breakdown




Relative Number of Instructions (to SA-1 Core)



TI vs ARM IClass Breakdown




IPC Comparison




Number of Cycles Relative To sa1core - nottaken


sa1core - perfect
 c62x
 Alpha 21264 - gshare
 Alpha 21264 - perfect

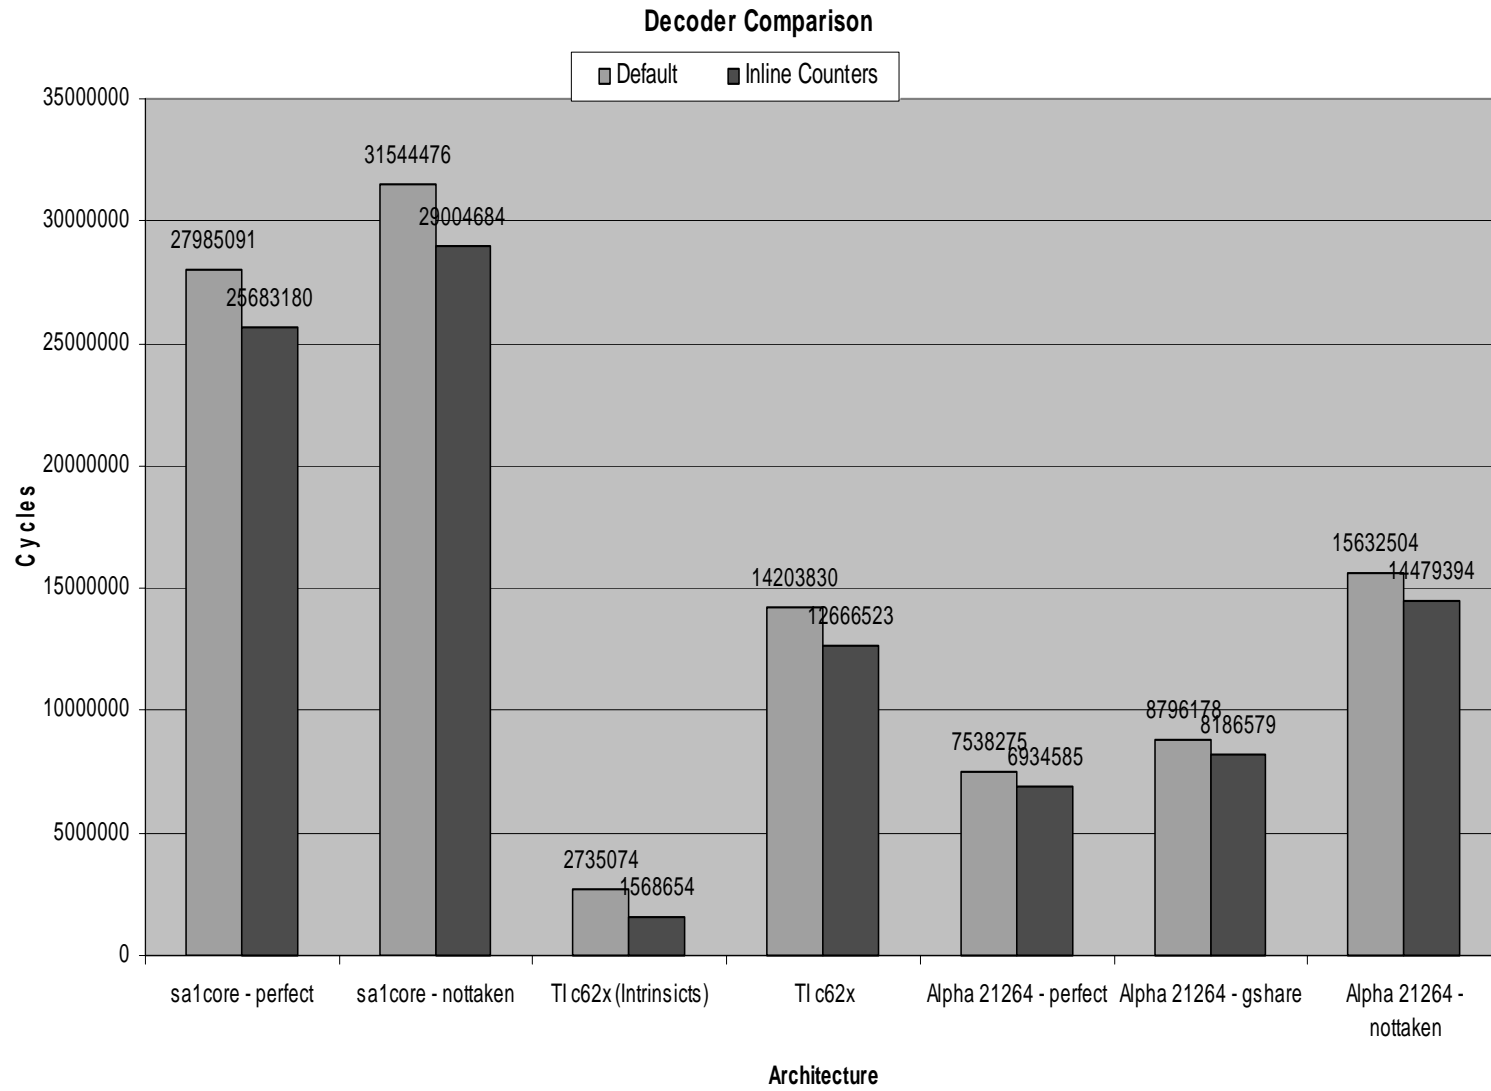


Diagnosis

- C62 compiler does not inline aggressively enough
- Codec is a better comparison because it does not include floating point







Conclusion: Intrinsic more important

- In this case intrinsic include
- Saturated add/sub/mul
- Absolute value
- Various 16 field extracts

Question?

- Is there a “convergent” architecture that combines both and covers a large class of embedded applications

The End