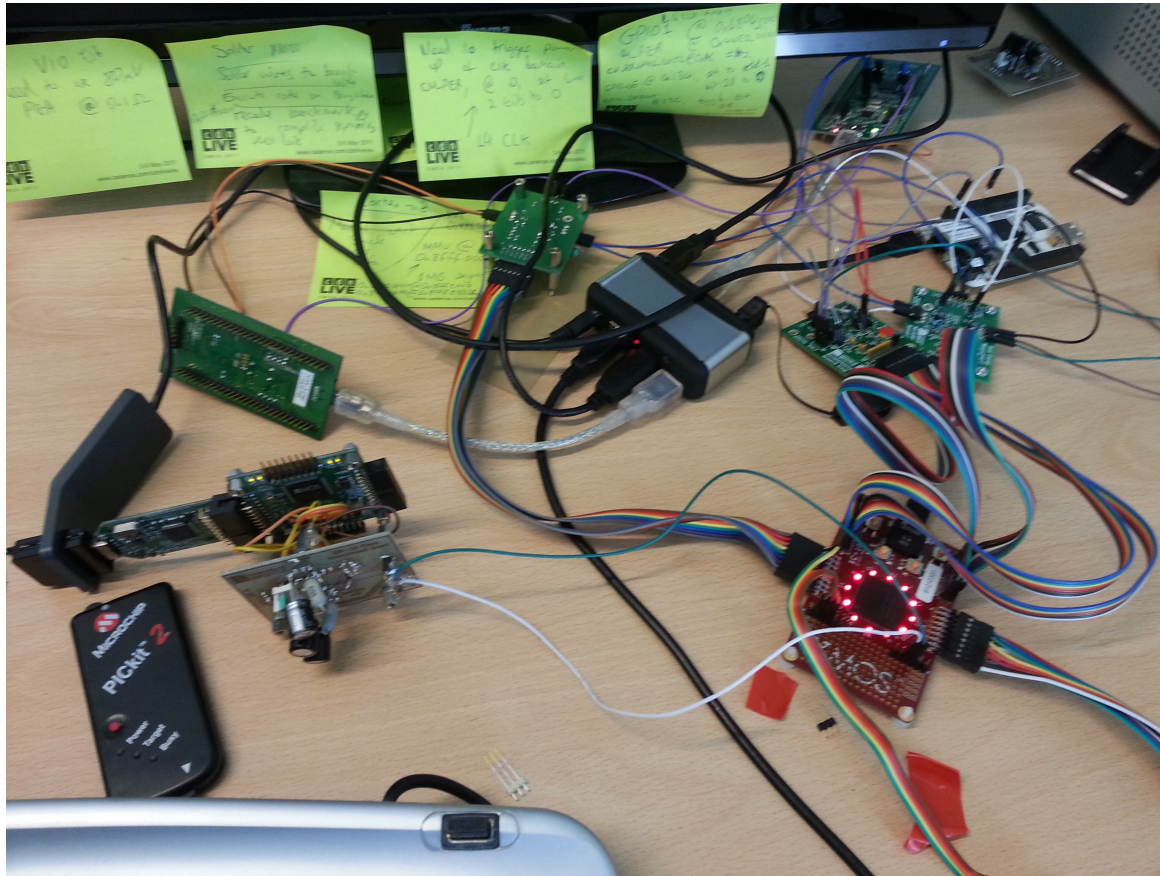


Impact of different compiler options on energy consumption



James Pallister
University of Bristol / Embecosm

Simon Hollis
University of Bristol

Jeremy Bennett
Embecosm

Motivation

- Compiler optimizations are claimed to have a large impact on software:
 - Performance
 - Energy
- No *extensive* study prior to this considering:
 - Different benchmarks
 - Many individual optimizations
 - Different platforms
- This work looks at the effect of many different optimizations across 10 benchmarks and 5 platforms.
- 238 Optimization passes covered by 150 flags
 - Huge amount of combinations

This Talk

- This talk will cover:
 - Importance of benchmarks
 - How to explore 2^{150} combinations of options
 - Demo
 - Correlation between time and energy
 - How to predict the effect of the optimizations
 - The best optimizations

Importance of Benchmarks

- One benchmark can't trigger all optimizations
- Perform differently on different platforms
- Need a range of benchmarks
- Broad categories to be considered for a benchmark:
 - Integer
 - Floating point
 - Branching
 - Memory

Our Benchmark List



Name	Source	B	M	I	FP	T	License	Category
Blowfish	MiBench	L	M	H	L	Multi	GPL	security
CRC32	MiBench	M	L	H	L	Single	GPL	network, telecomm
Cubic root solver	MiBench	L	M	H	L	Single	GPL	automotive
Dijkstra	MiBench	M	L	H	L	Multi	GPL	network
FDCT	WCET	H	H	L	H	Single	None [†]	consumer
Float Matmult	WCET	M	H	M	M	Single	GPL	automotive, consumer
Integer Matmult	WCET	M	M	H	L	Single	None [†]	automotive
Rjindael	MiBench	H	L	M	L	Multi	GPL	security
SHA	MiBench	H	M	M	L	Multi	GPL	network, security
2D FIR	WCET	H	M	L	H	Single	None [†]	automotive, consumer

Choosing the Platforms



- Range of different features in the platforms chosen
 - Pipeline Depth
 - Multi- vs Single- core
 - FPU available?
 - Caching
 - On-chip vs off-chip memory

Platforms Chosen



ARM Cortex-M0	ARM Cortex-M3	ARM Cortex-A8	XMOS L1	Adapteva Epiphany
Small memory	Small memory	Large memory	Small memory	On-chip and off-chip memory
Simple Pipeline	Simple Pipeline, with forwarding logic, etc.	Complex superscalar pipeline	Simple pipeline	Simple superscalar pipeline
		SIMD/FPU		FPU
			Multiple threads	16 cores

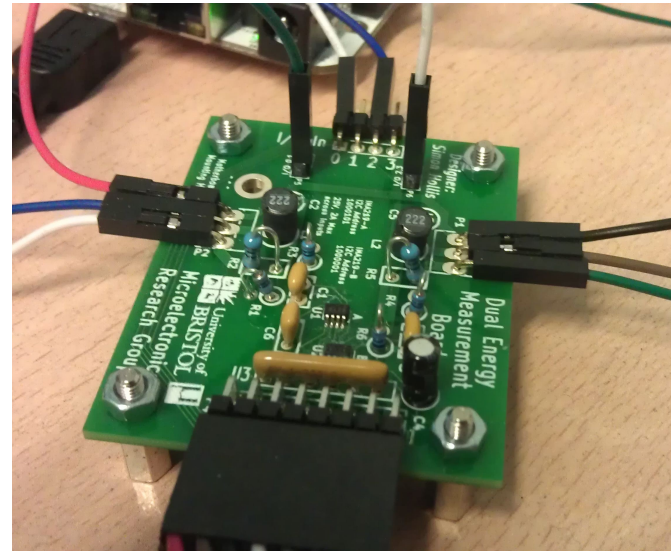
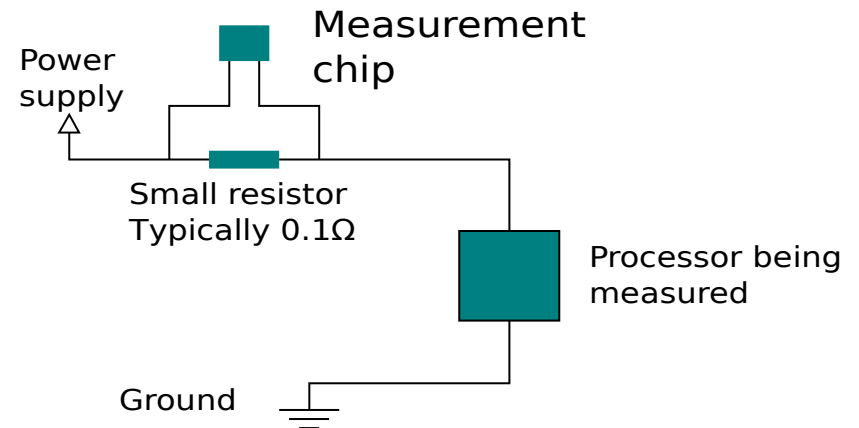
Experimental Methodology

- Compiler optimizations have many non-linear interactions
- 238 optimization passes combined into 150 different options (GCC)
- 82 compiler options enabled by O3
- How to test all of these, while accounting for the interactions between optimizations?

Fractional Factorial Designs

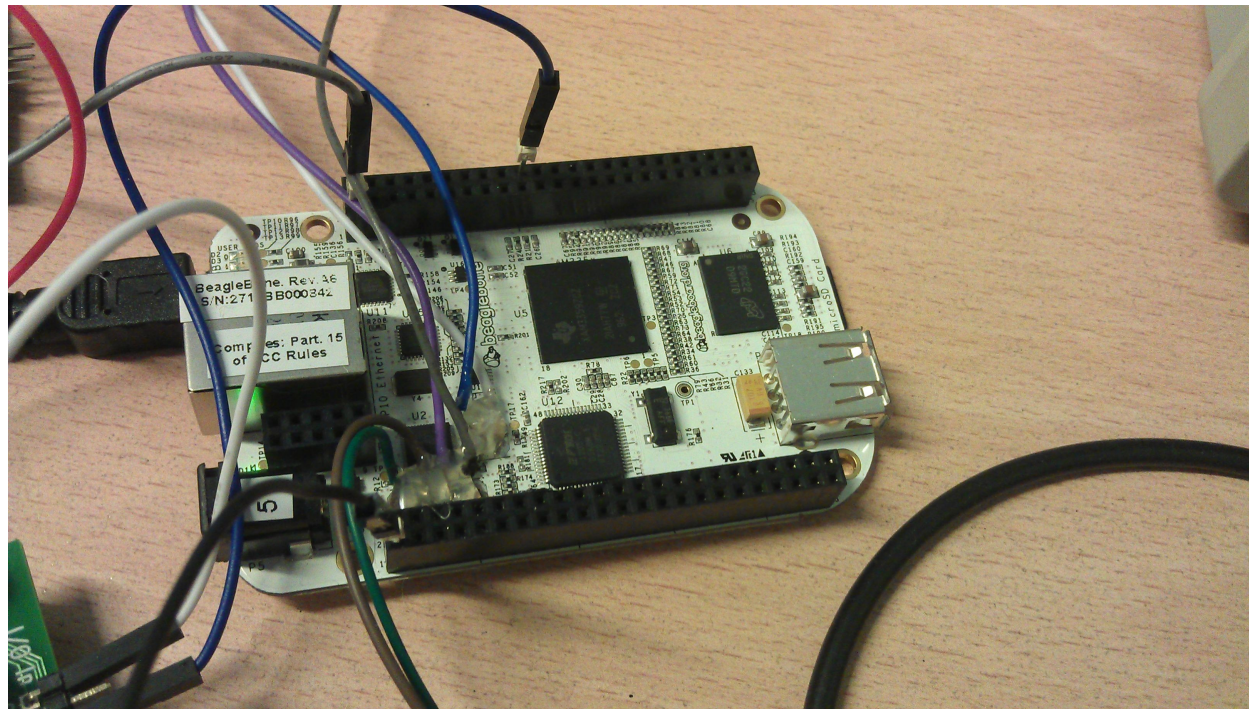
Hardware Measurements

- Current, voltage and power monitor
- 10 kSamples/s
- Low noise
- XMOS board to control and timestamp measurements
- Integrate to get energy consumption

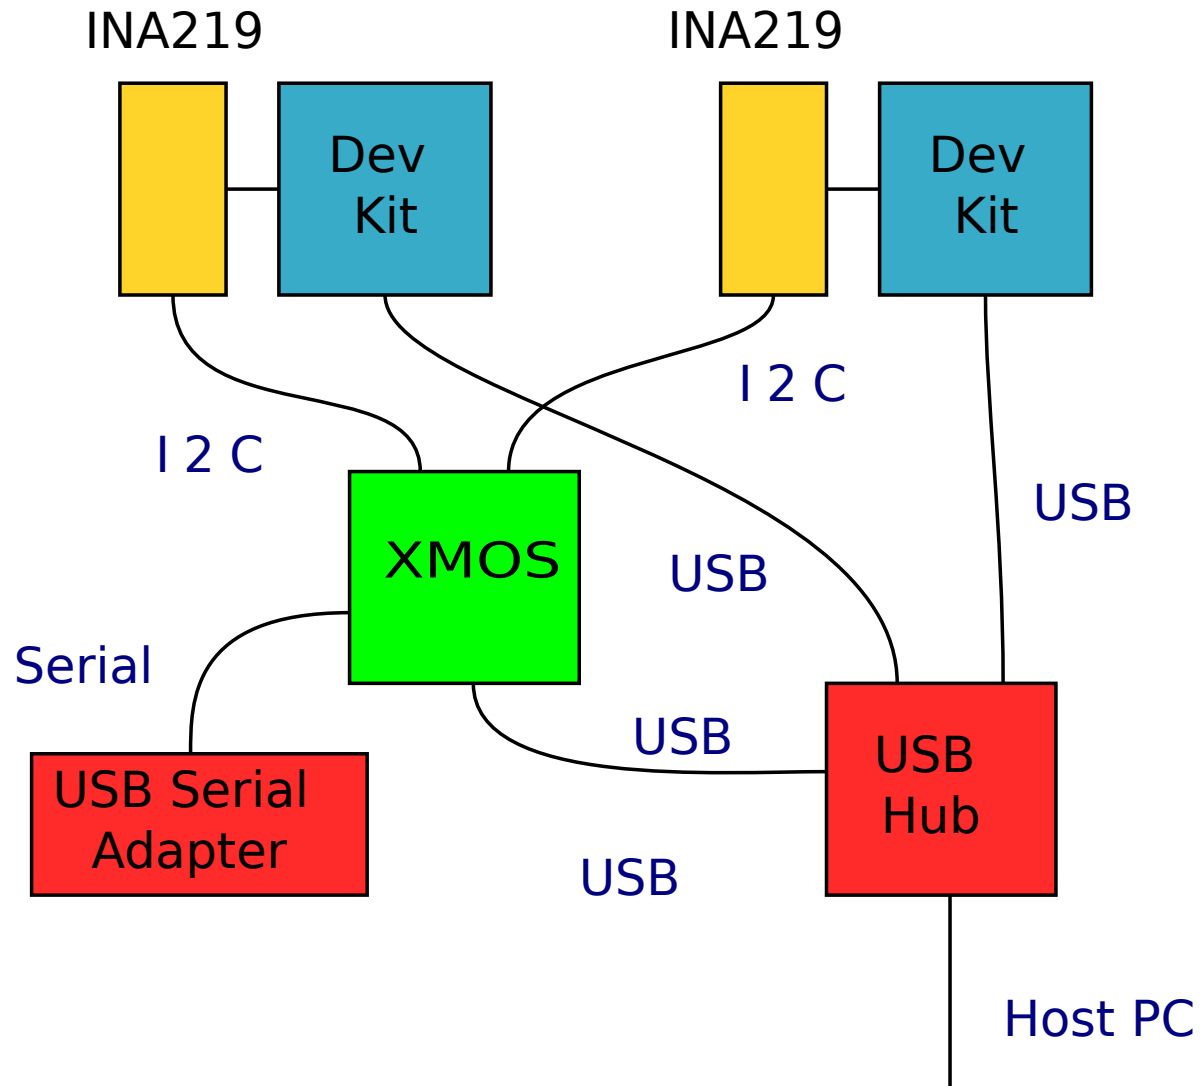


Instrumenting the Hardware

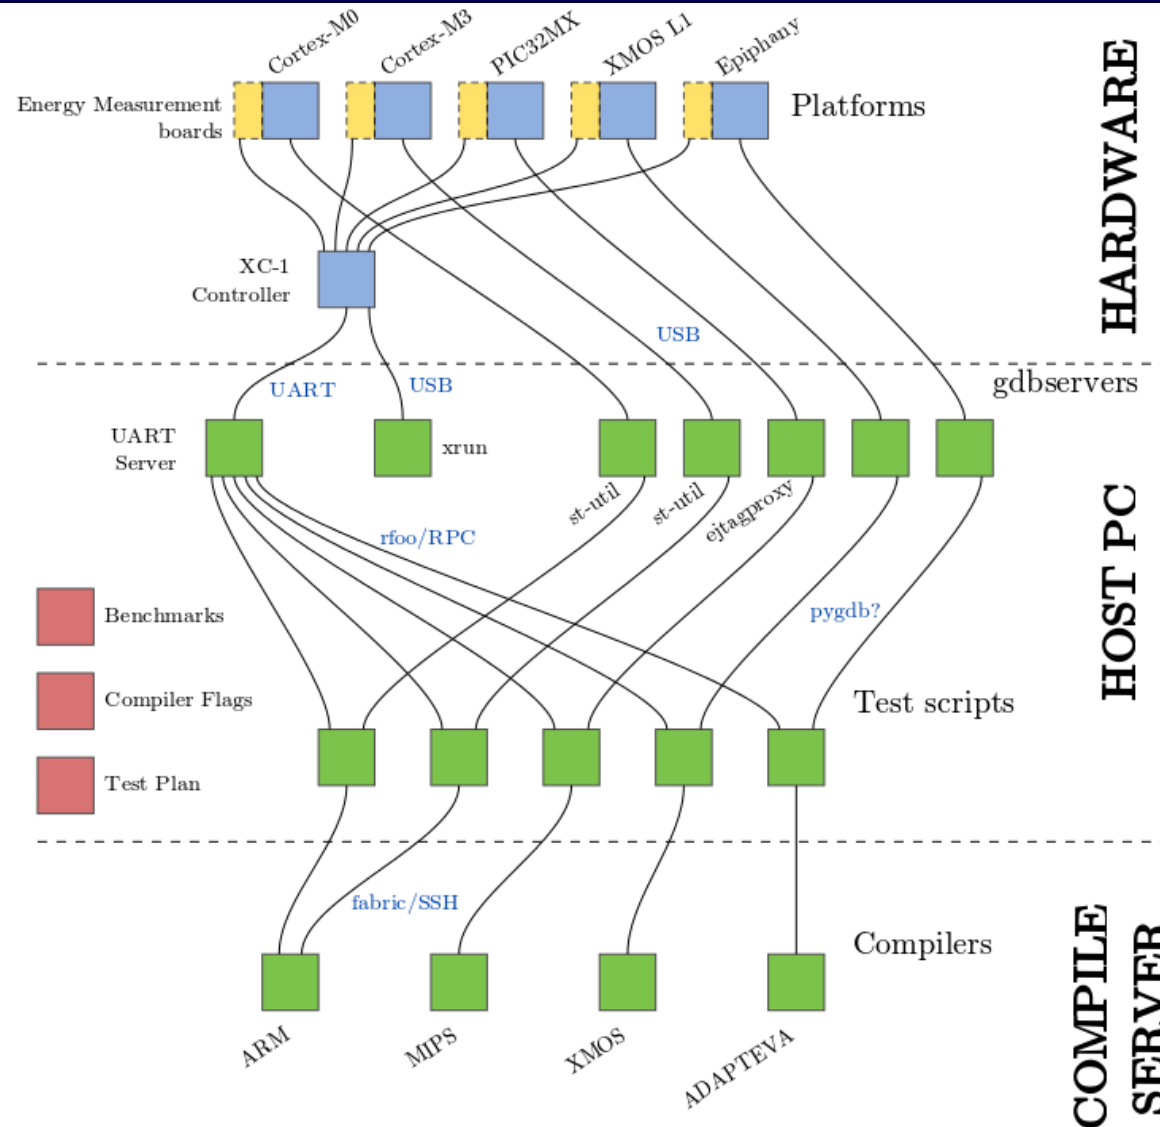
- How to attach the power measurement circuit to the hardware?
- Invasive...



Hardware



Software



Results

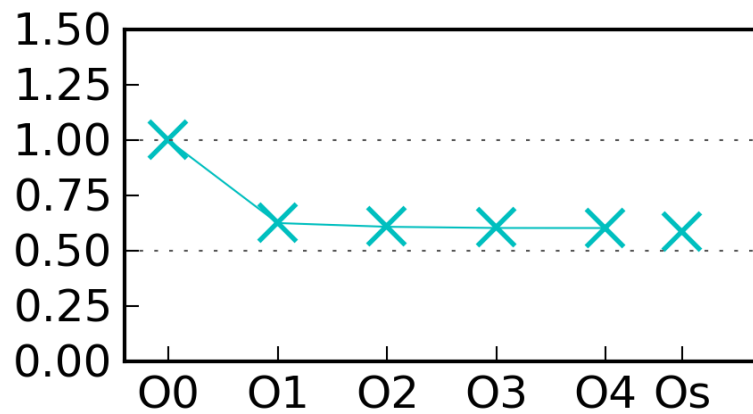
- Energy consumption \approx Execution time
 - Generalization, not true in every case
- Optimization unpredictability

- No optimization is universally good across benchmarks and platforms

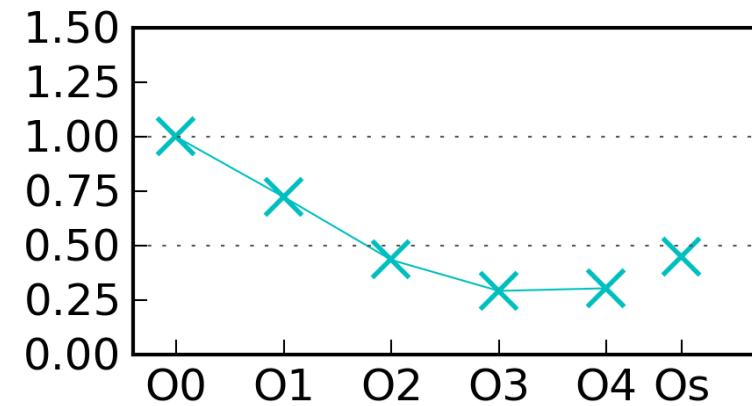
Overview



FDCT, Cortex-M0



FDCT, Cortex-A8

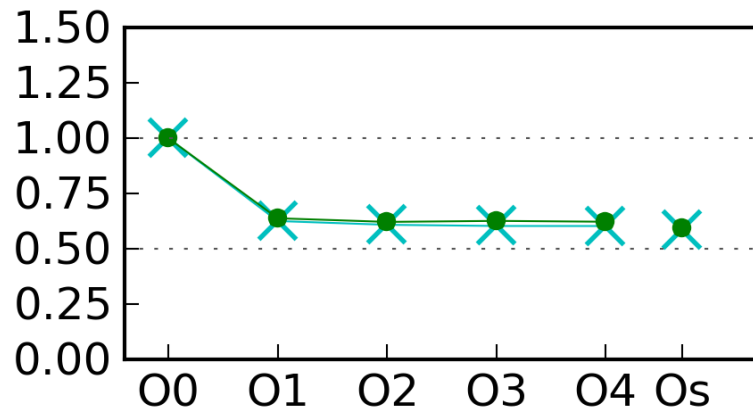


× Execution time

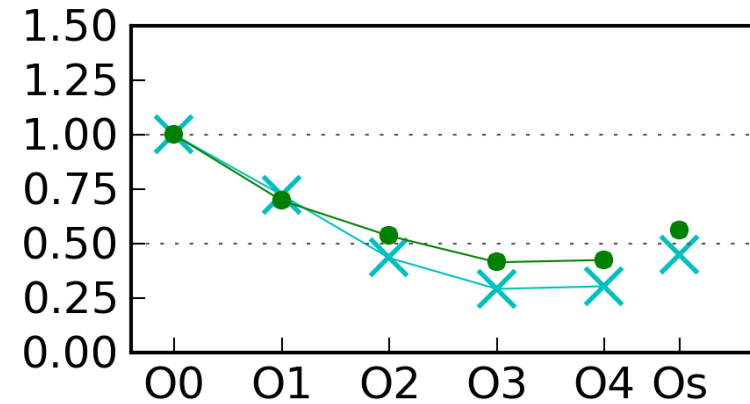
Overview



FDCT, Cortex-M0



FDCT, Cortex-A8

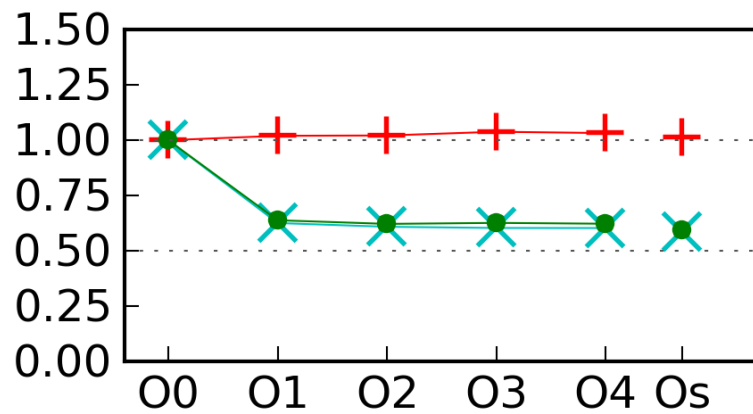


×× Execution time
●● Energy consumed

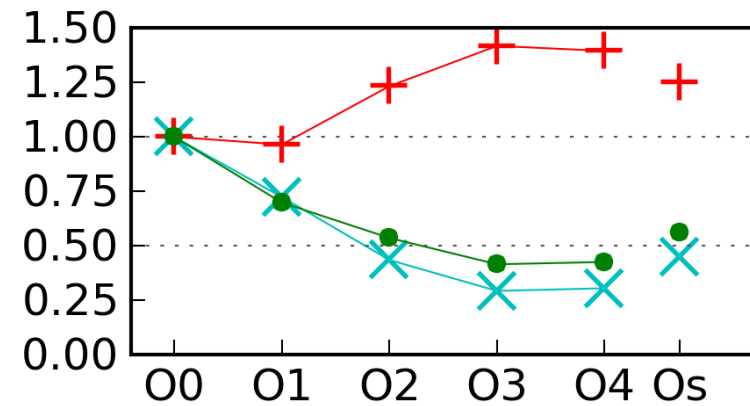
Overview



FDCT, Cortex-M0

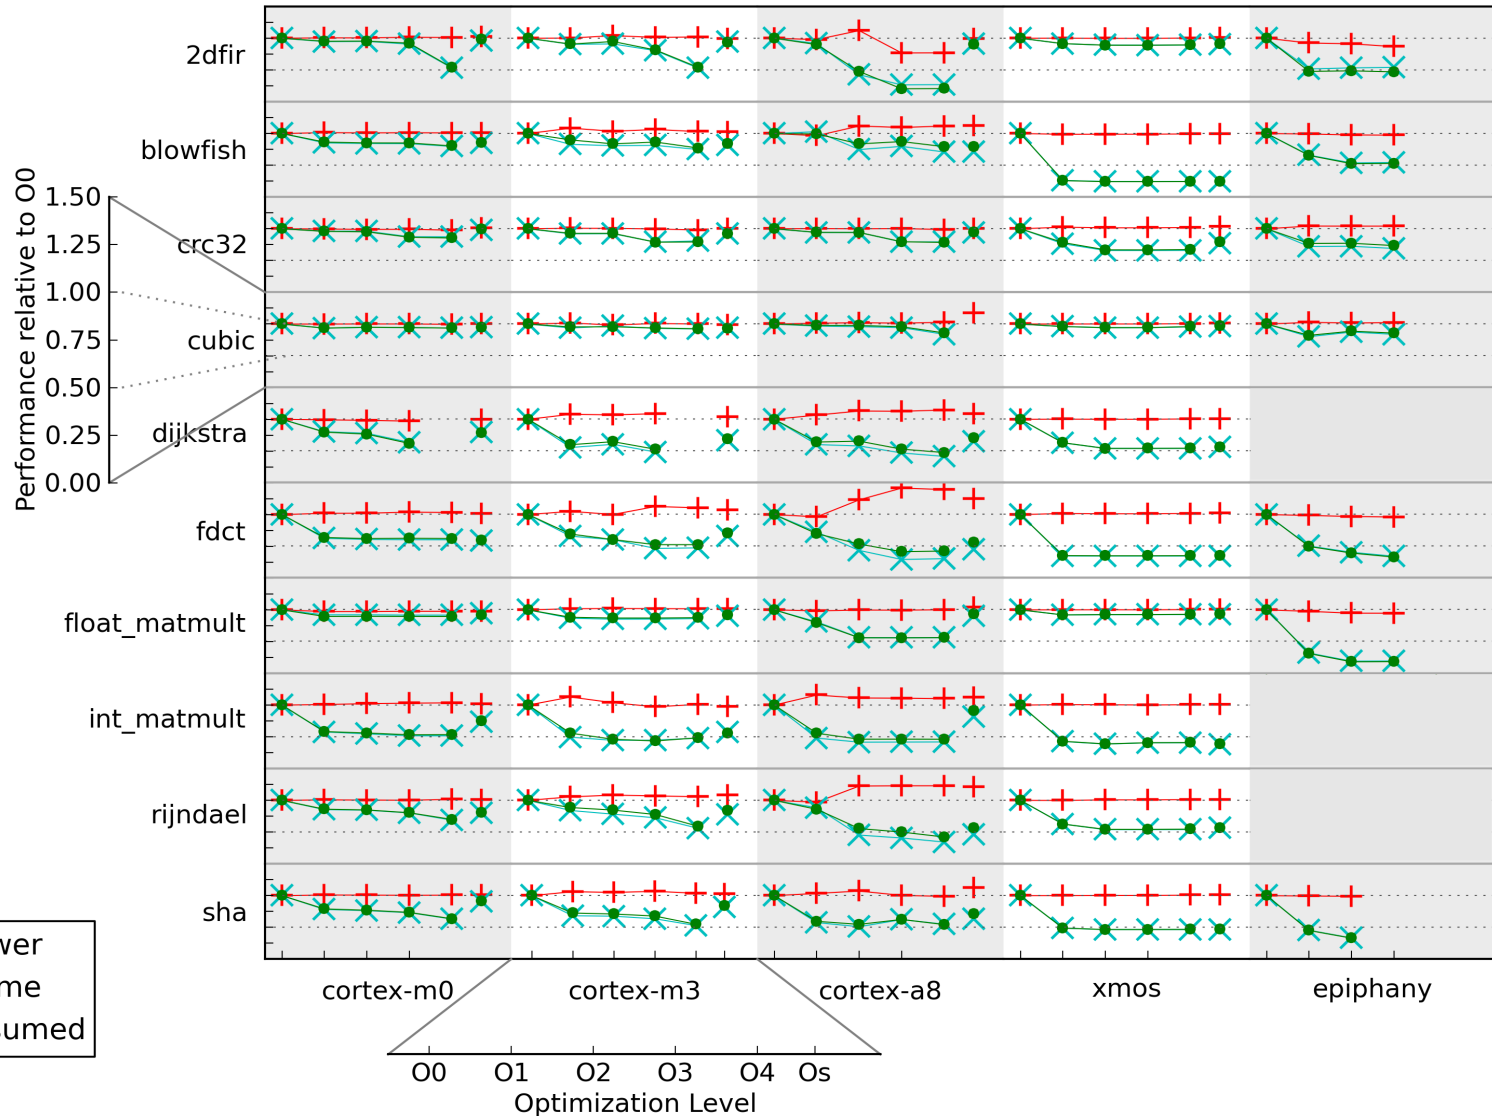


FDCT, Cortex-A8



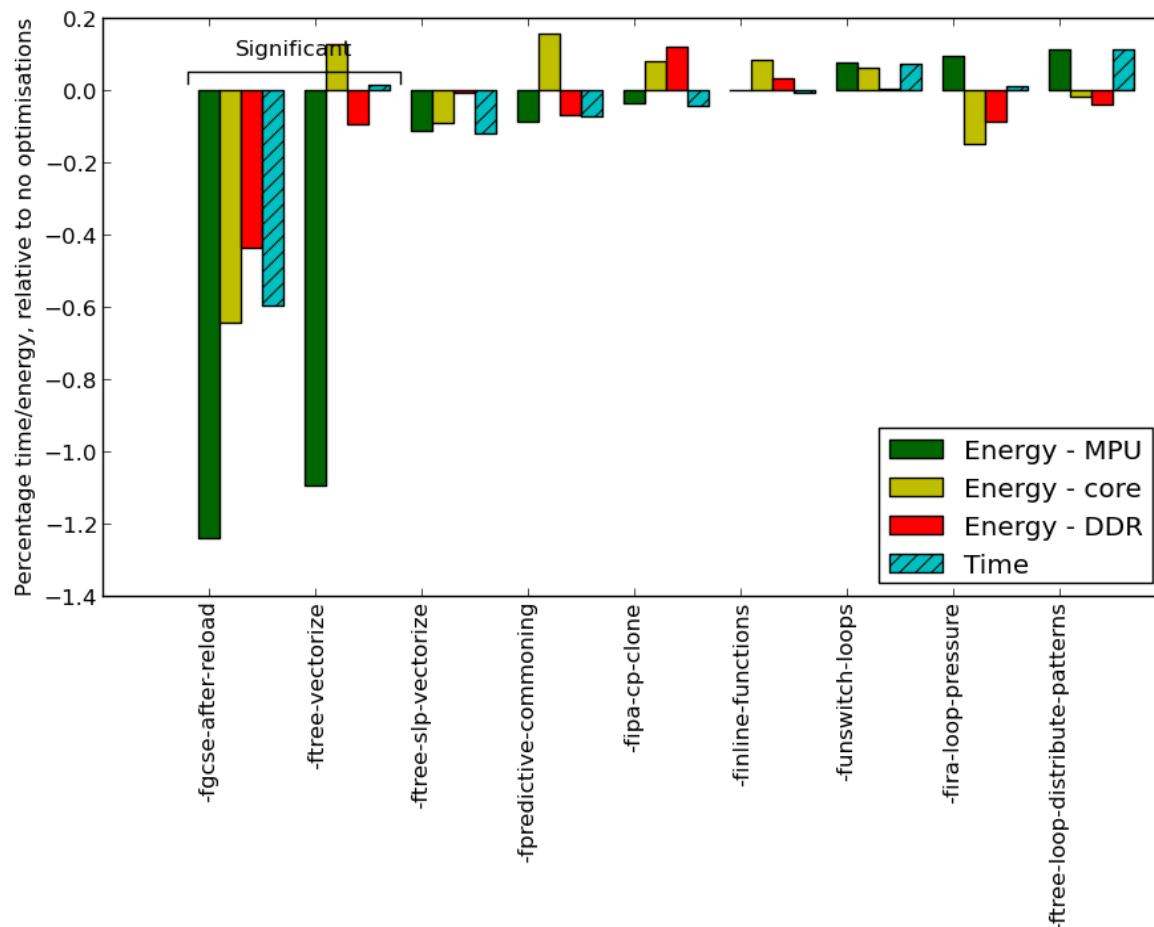
- ++ Average power
- ×× Execution time
- Energy consumed

Overview



When Time \neq Energy

- Complex pipeline
- -free-vectorize
 - NEON SIMD unit
 - Much lower power

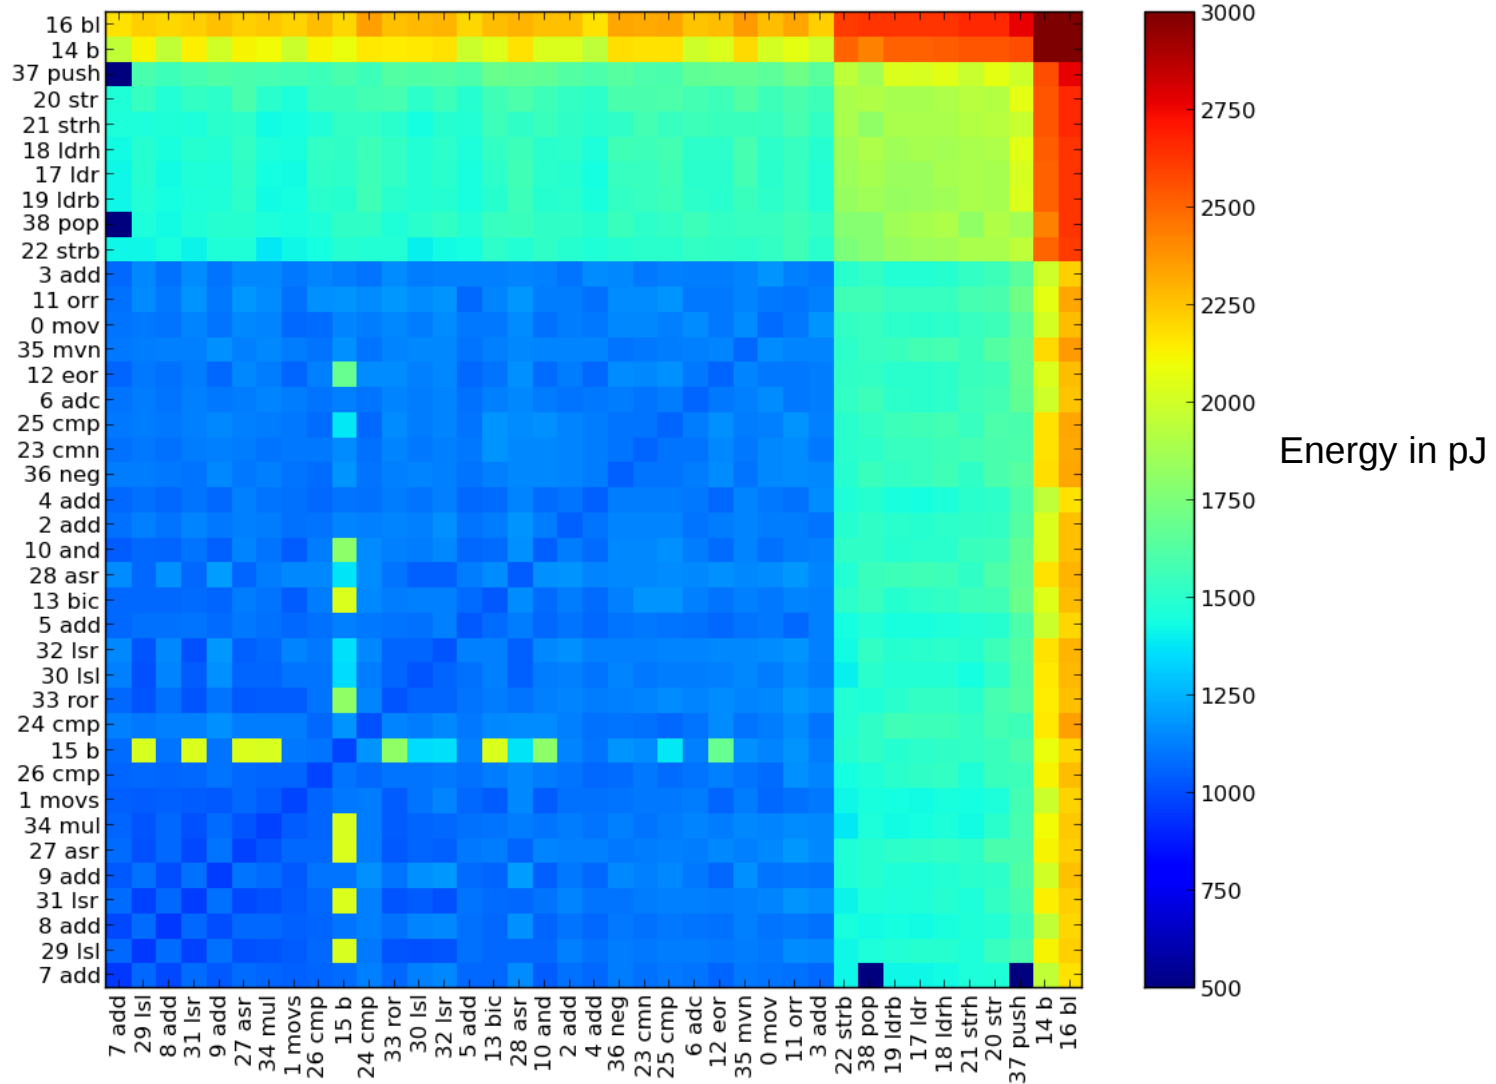


O3 Flags, 2DFIR, Cortex-A8

Conclusion: Mostly, Time \approx Energy

- Highly correlated
- Especially so for 'simple' pipelines
- Little scope for stalling or superscalar execution
- Complex pipelines:
 - Still a correlation
 - But more variability
 - SIMD, superscalar execution
- To get the most optimal energy consumption we need better than “go fast”

Case Study: Cortex-M0



What does this mean?

For the Compiler Writer

- Current optimization levels (O1, O2, etc.) are a good balance between compile time and performance/energy.
- Never completely optimal
- Machine learning
 - MILEPOST
 - Genetic algorithms
- Current optimizations targeted for performances
- Few (if any) optimizations in current compilers designed to reduce energy consumption

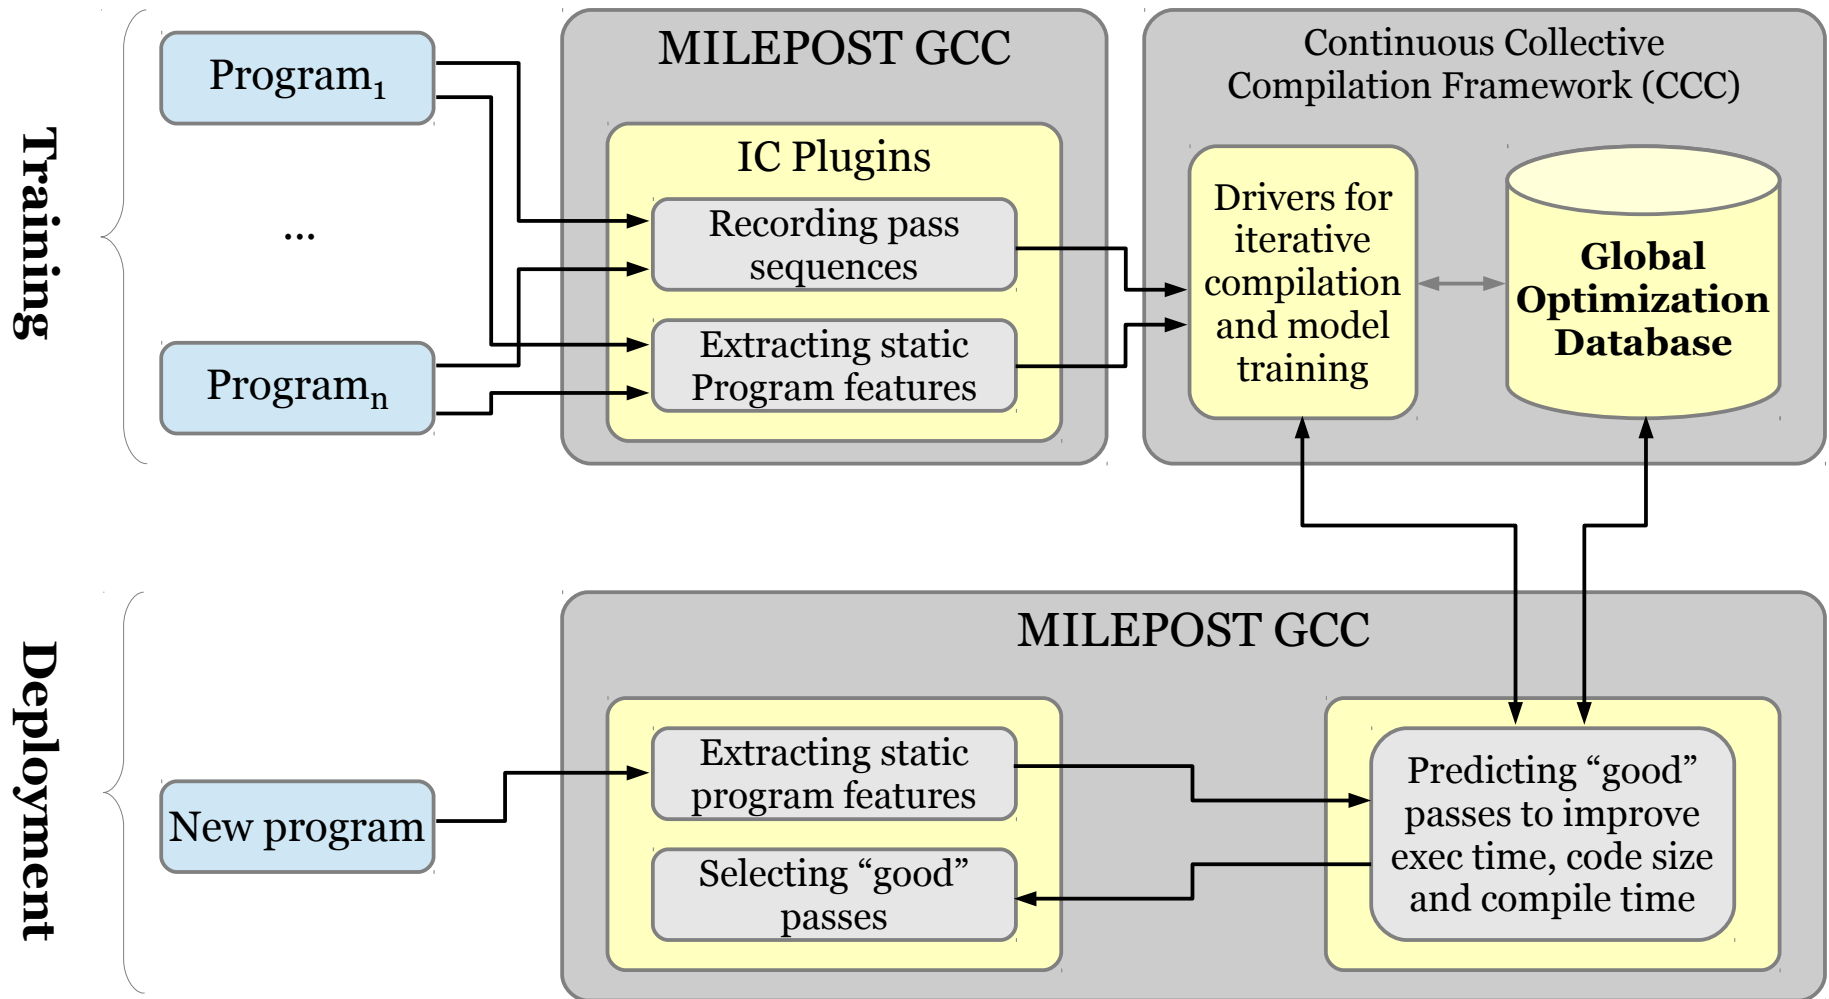
What does this mean?



For the embedded developer

- Try the optimization levels – O3 is a good bet
- Use hardware peripherals
- SIMD
- Power Modes
 - Sleep
- Memory
 - Closer to the processor the better
 - Exploit RAM

MILEPOST GCC



From Fursin et al, 2008

Conclusion

- Time \approx Energy
 - True for simple pipelines
 - Mostly true for complex pipelines
 - Good approximation
- Optimization unpredictability
 - Difficult to model the interactions between optimizations
- Commonality across platforms
 - Instruction set plays a role
 - Common options for the ARM platforms, but not Epiphany

Questions and Demonstration

james.pallister@bristol.ac.uk

simon@cs.bris.ac.uk

jeremy.bennett@embecosm.com

All data at: www.jpallister.com/wiki