PARA DIME

Power and Energy Modelling of Multi-core Processors for System-Level Design Space Exploration

Santhosh Kumar Rethinagiri, **Oscar Palomar**, Osman Unsal, Adrian Cristal
Barcelona Supercomputing Center
Energy-Aware Computing Workshop 2014

This project and the research leading to these results has received funding from the European Community's Seventh Framework Programme [FP7/2007-2013] under grant agreement n 318693

ParaDIME Consortium

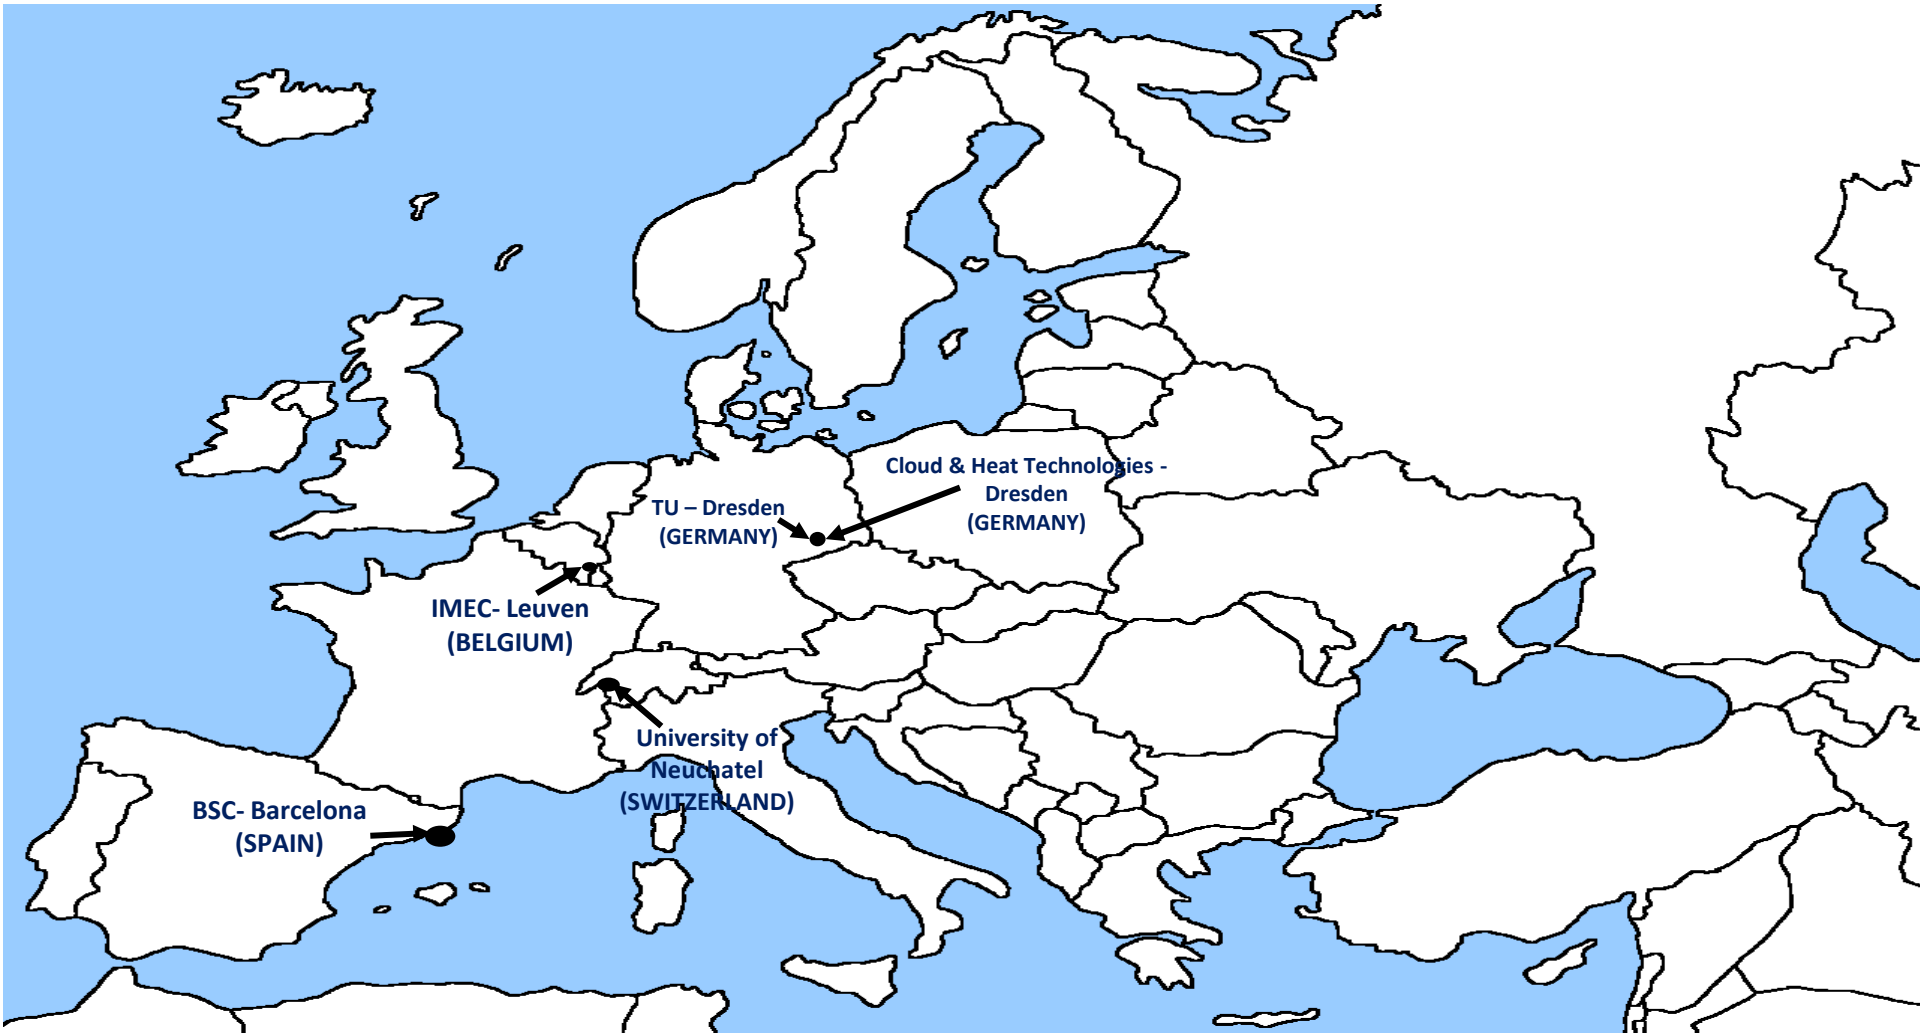

Why ParaDIME ?

- Parallel Distributed Infrastructure for Minimization of Energy
- Rising cost
  - Hardware cost
  - Programming efficiency
  - Runtime optimization
  - Energy aware data center computing

The ParaDIME Stack

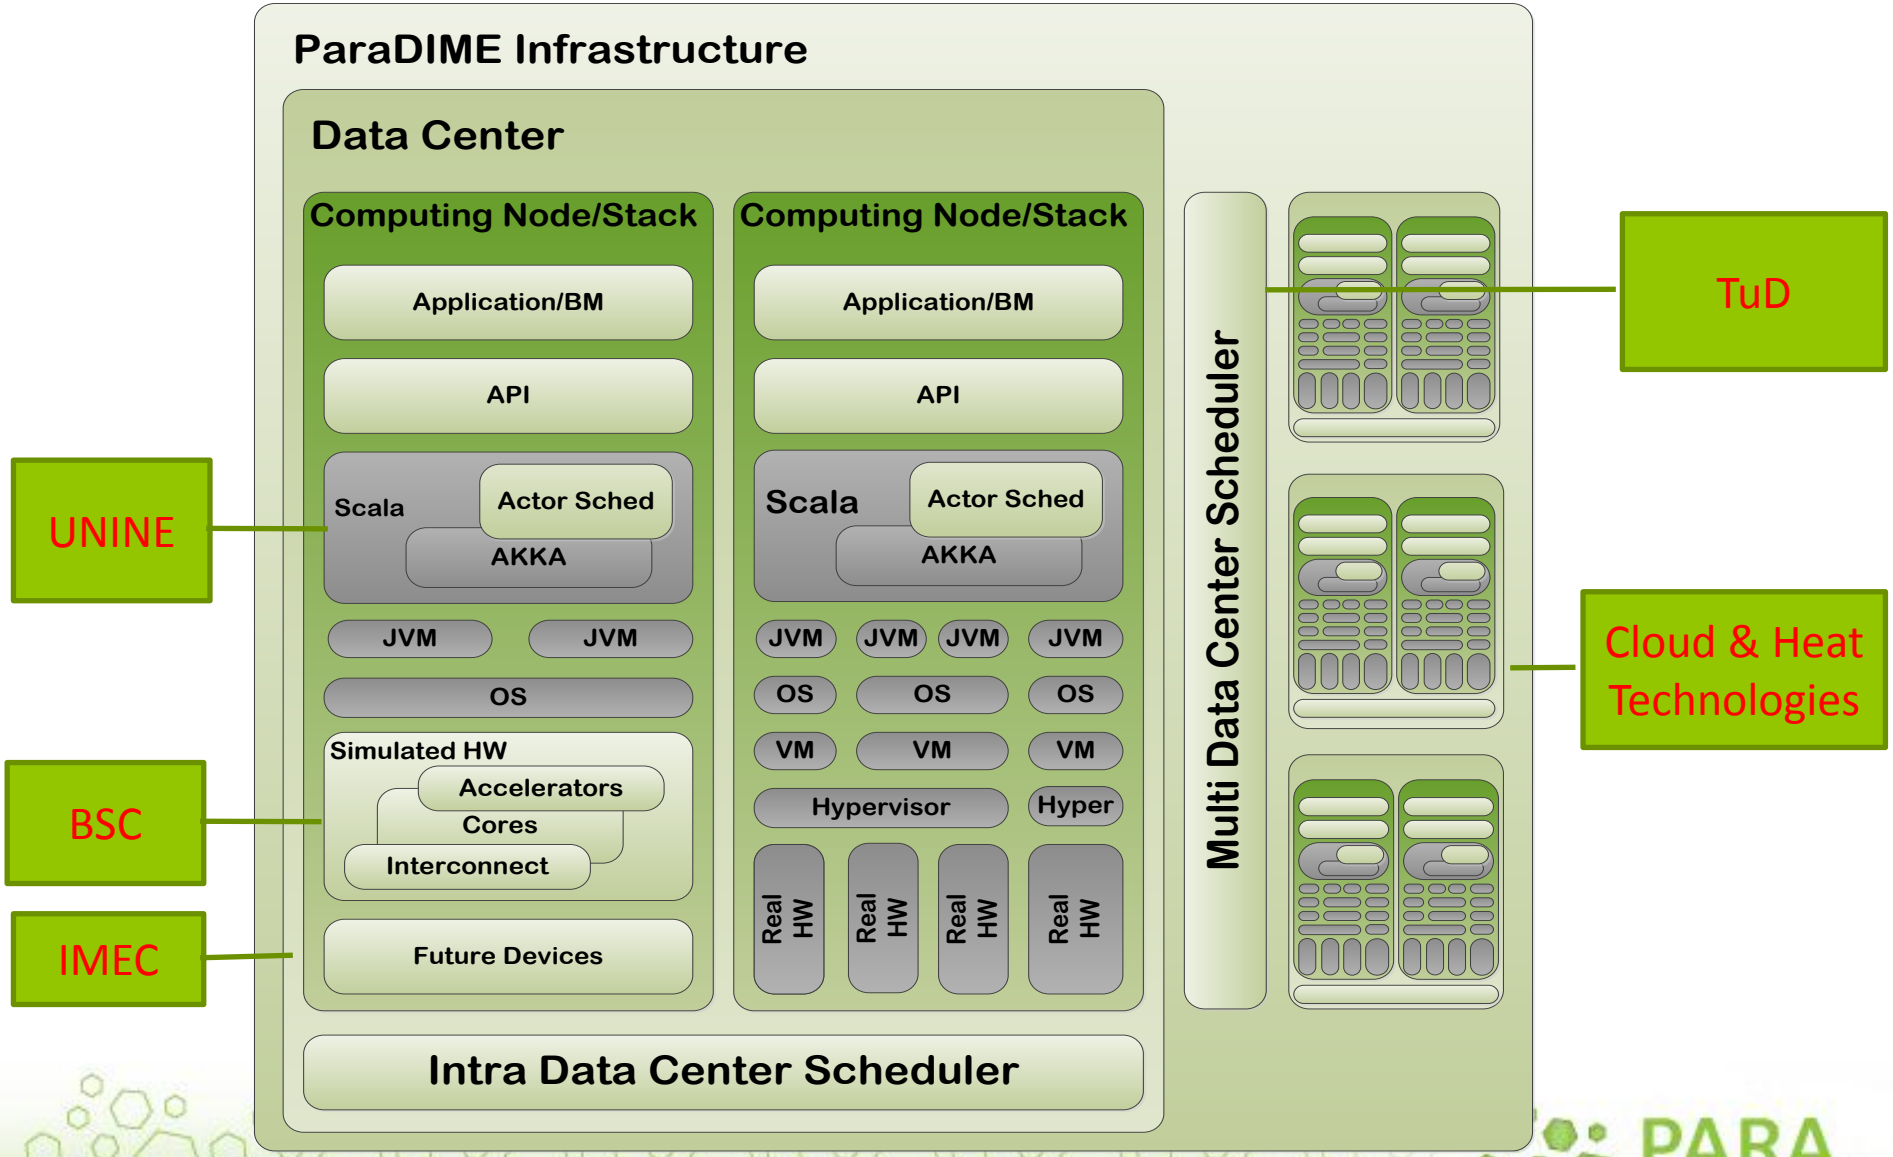

Challenges of modelling power of heterogeneous systems

- Estimating power/energy is a critical design goal for electronic devices.
- Designers today must evaluate power estimation as early as possible in the electronics design.
- Design changes are easier in the design phase and have the greatest impact on application power estimation at System-Level.
- A platform to use **different processors and components**.
- Functional level is accurate but it's a **course grain**. Restriction in terms of measuring power from the real board.
- For **fine grain**, we can achieve it from gate level simulation. Restriction applies as we don't have the tools and RTL sources. Very slow simulation speed.
- Another challenge is power law holds for a simple processor but for **complex processor system** remains debatable ?

Power estimation methodology and tools

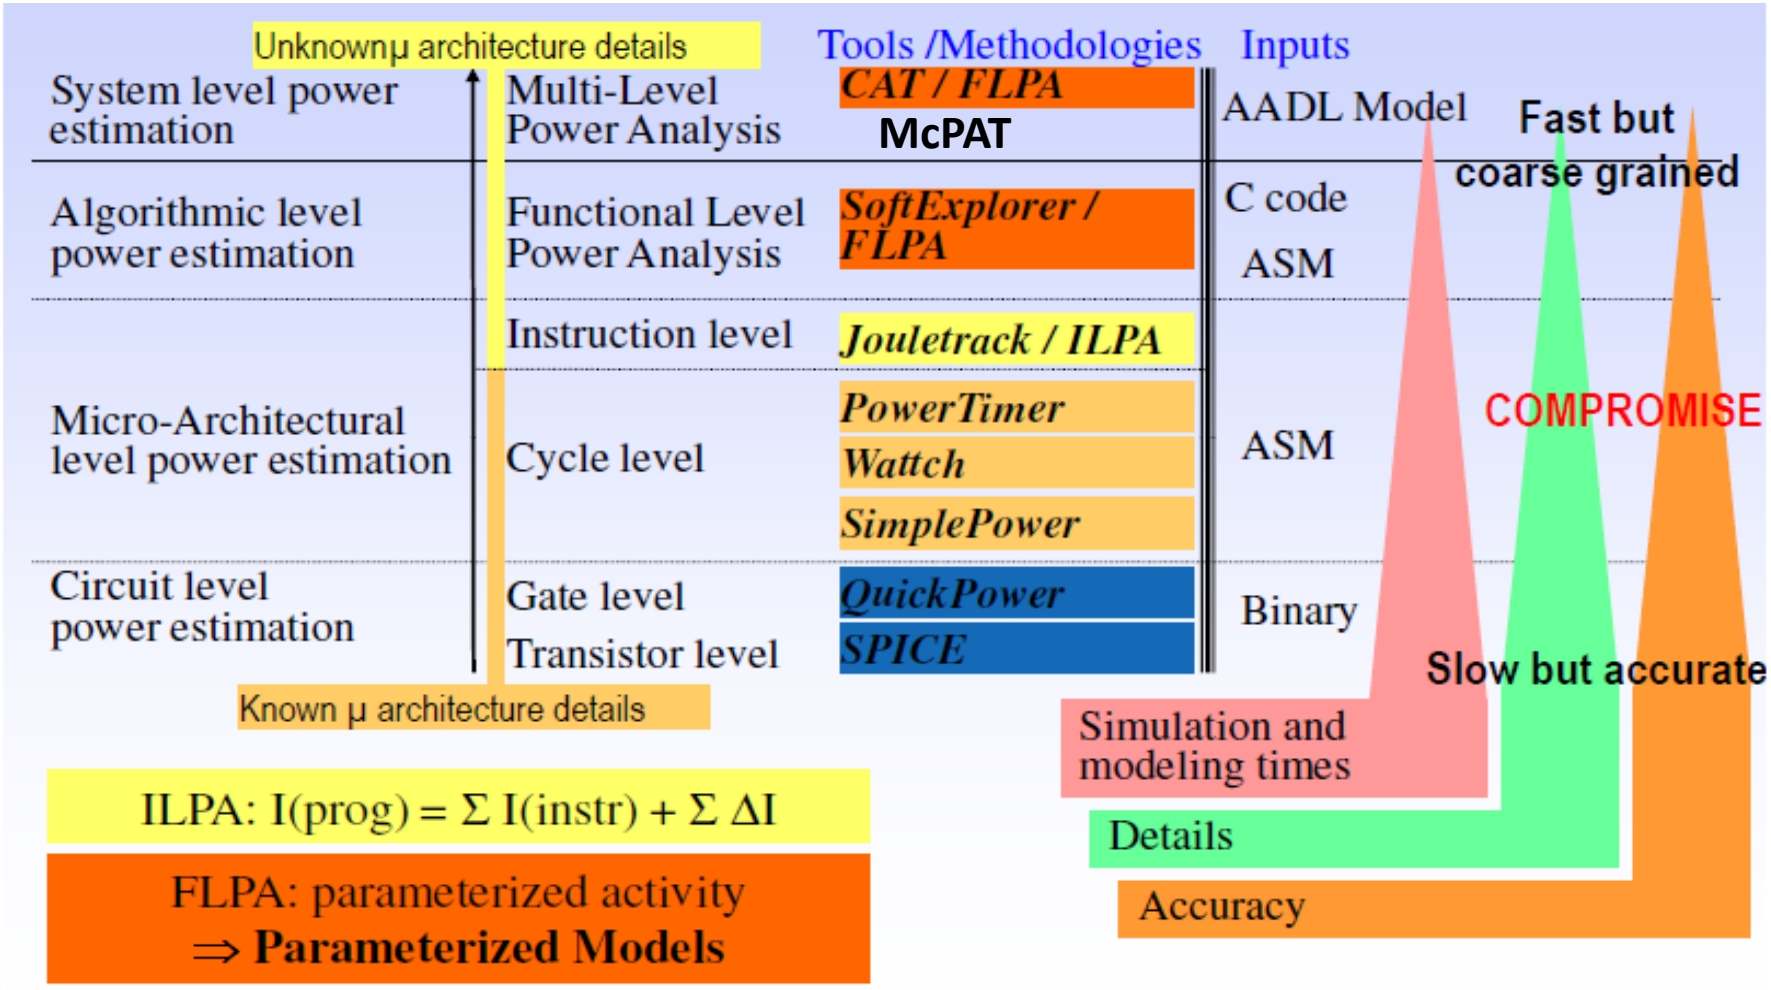

Hybrid design space exploration methodology

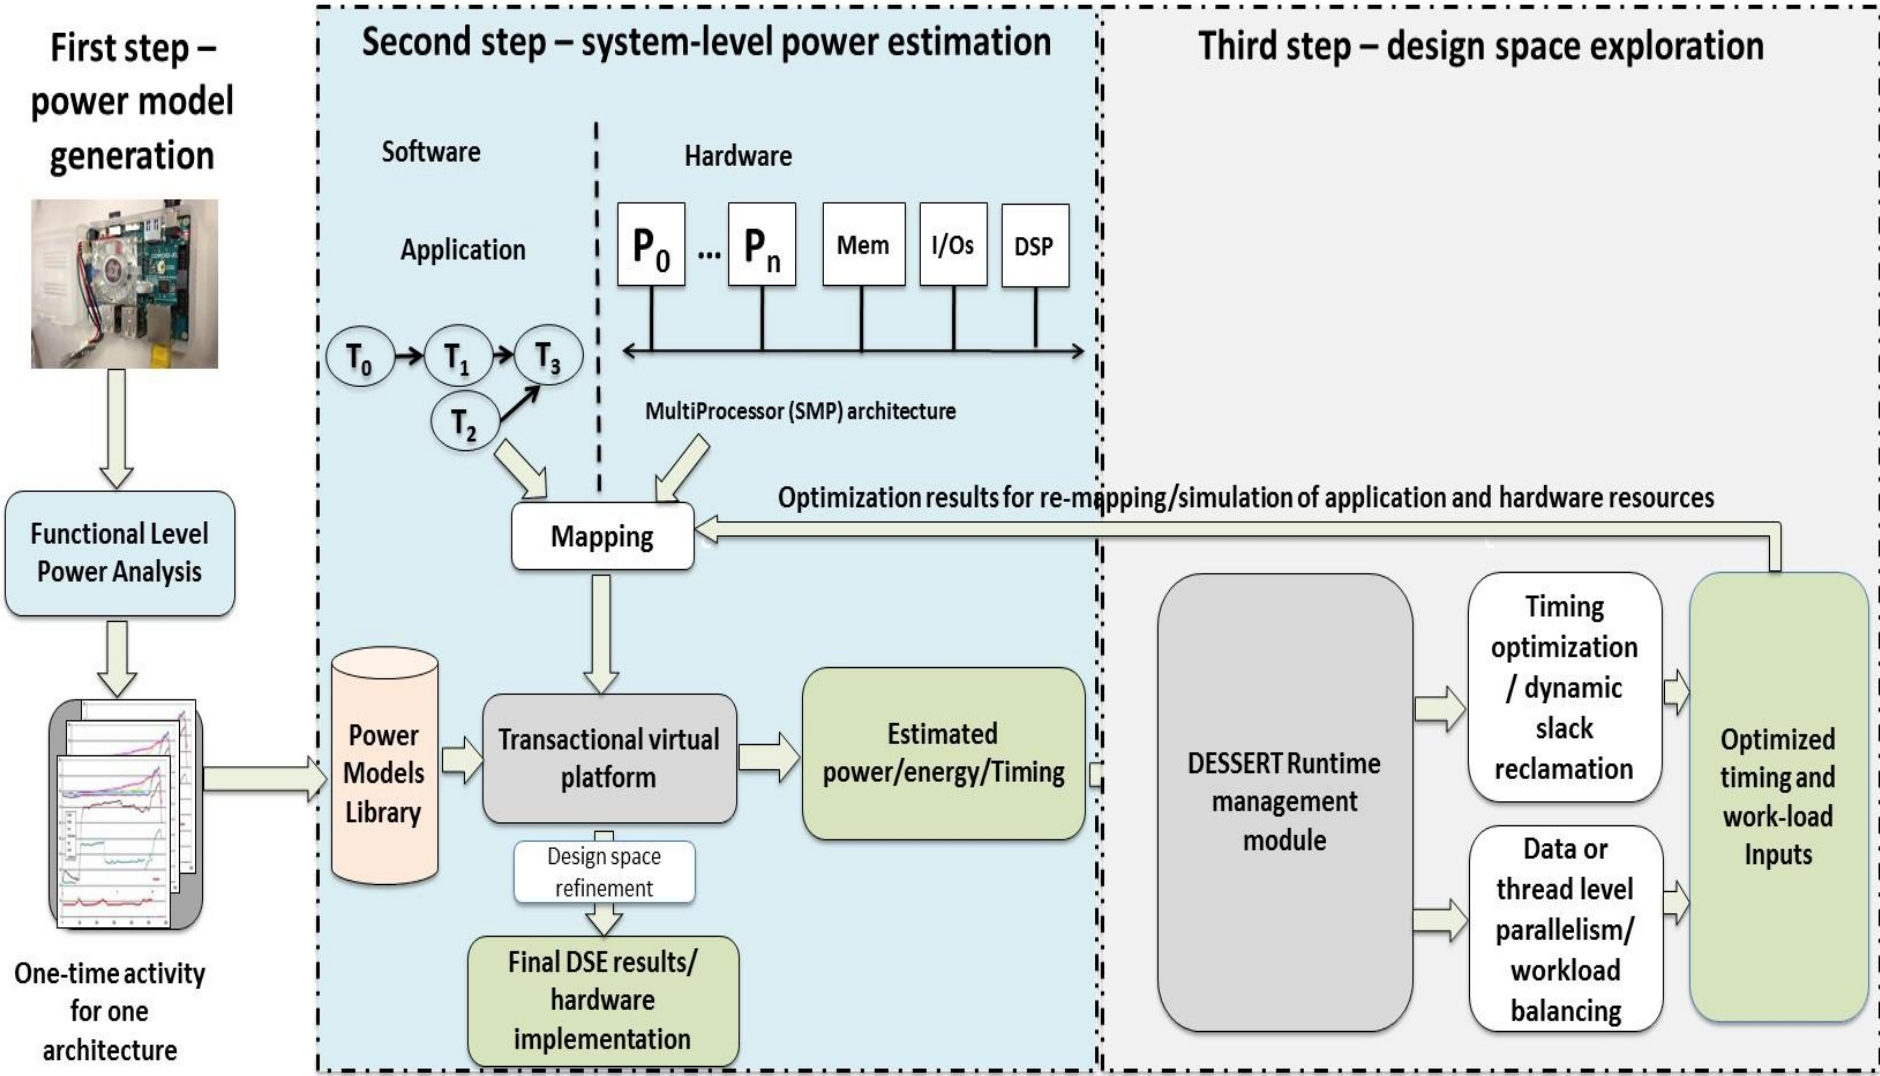

First step : FLPA (Functional Level Power Analysis)

- FLPA
  - 3 steps
  - Complex architectures
  - Simples models
  - Precise estimations
  - Processors Models
    - DSP & GPP
    - *SoftExplorer* Tool
  - FPGA Models
  - Flash Memory Models
  - System ...

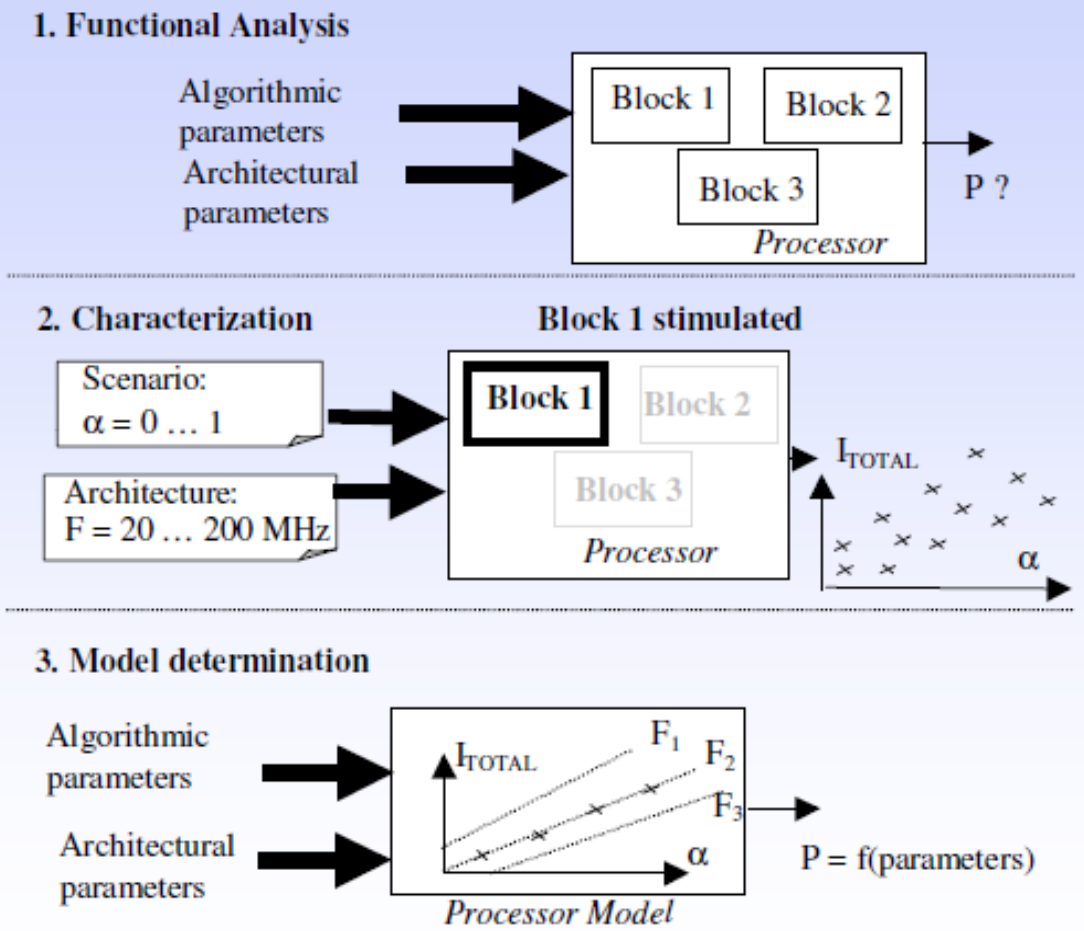

PARA DIME

Functional block (ARM Cortex-A9)

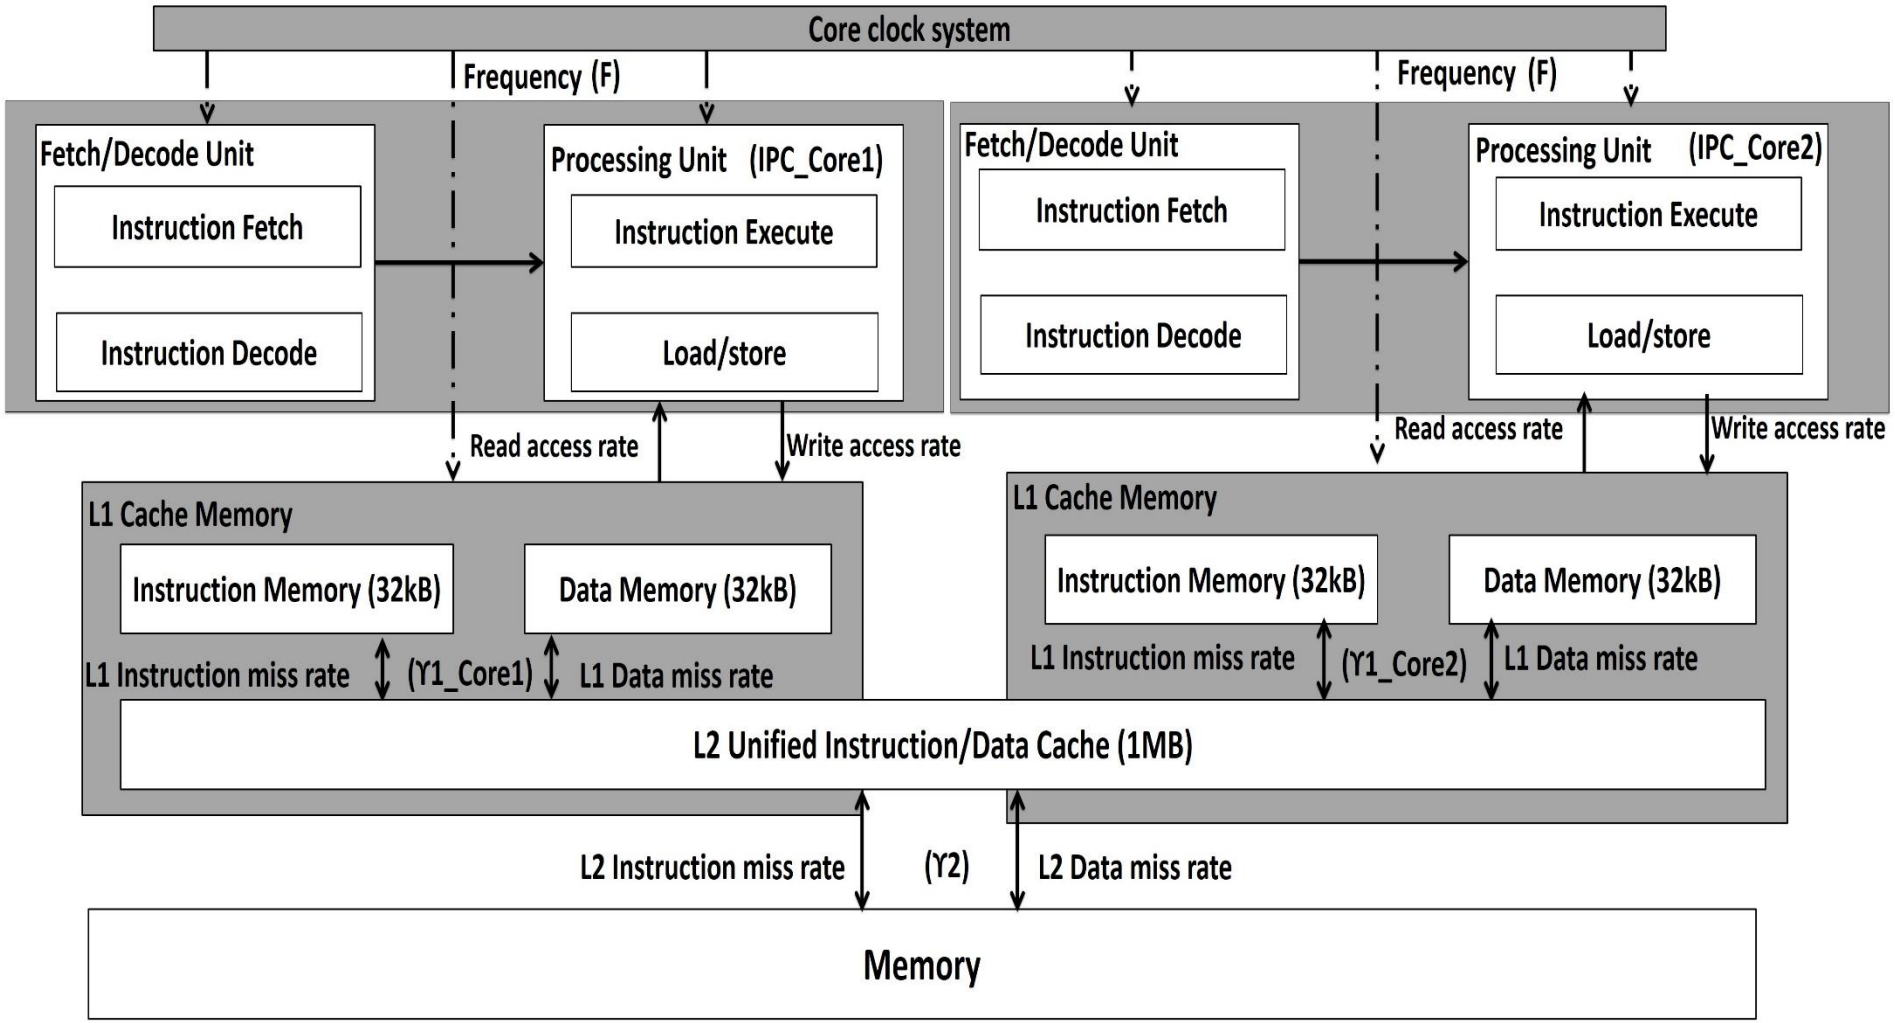

PARA DIME

Generic Power Model Parameters

The Parameters which influence the power in a system.

	Name	Description
Algorithmic	τ	External memory access rate
	γ_1	L1 cache miss rate for a processor
	γ_2	L2 cache miss rate for a processor
	SCU	Snoop control unit counter for ARM Cortex-A9 multi-core
	γ_1	L1 cache miss rate for a processor
	ρ	parallelism rate for DSP processor fetch stage access
	β	DSP processor processing rate between instruction memory unit and processing unit
	PSR	Pipeline stall rate between instruction memory unit and processing unit
	IPC	Instruction Per Cycle
	α	area utilization for a CLB
Architectural	$F_{processor}$	Frequency of the processor
	F_{bus}	Frequency of the bus
	N	Number of cores

Power measurement environment

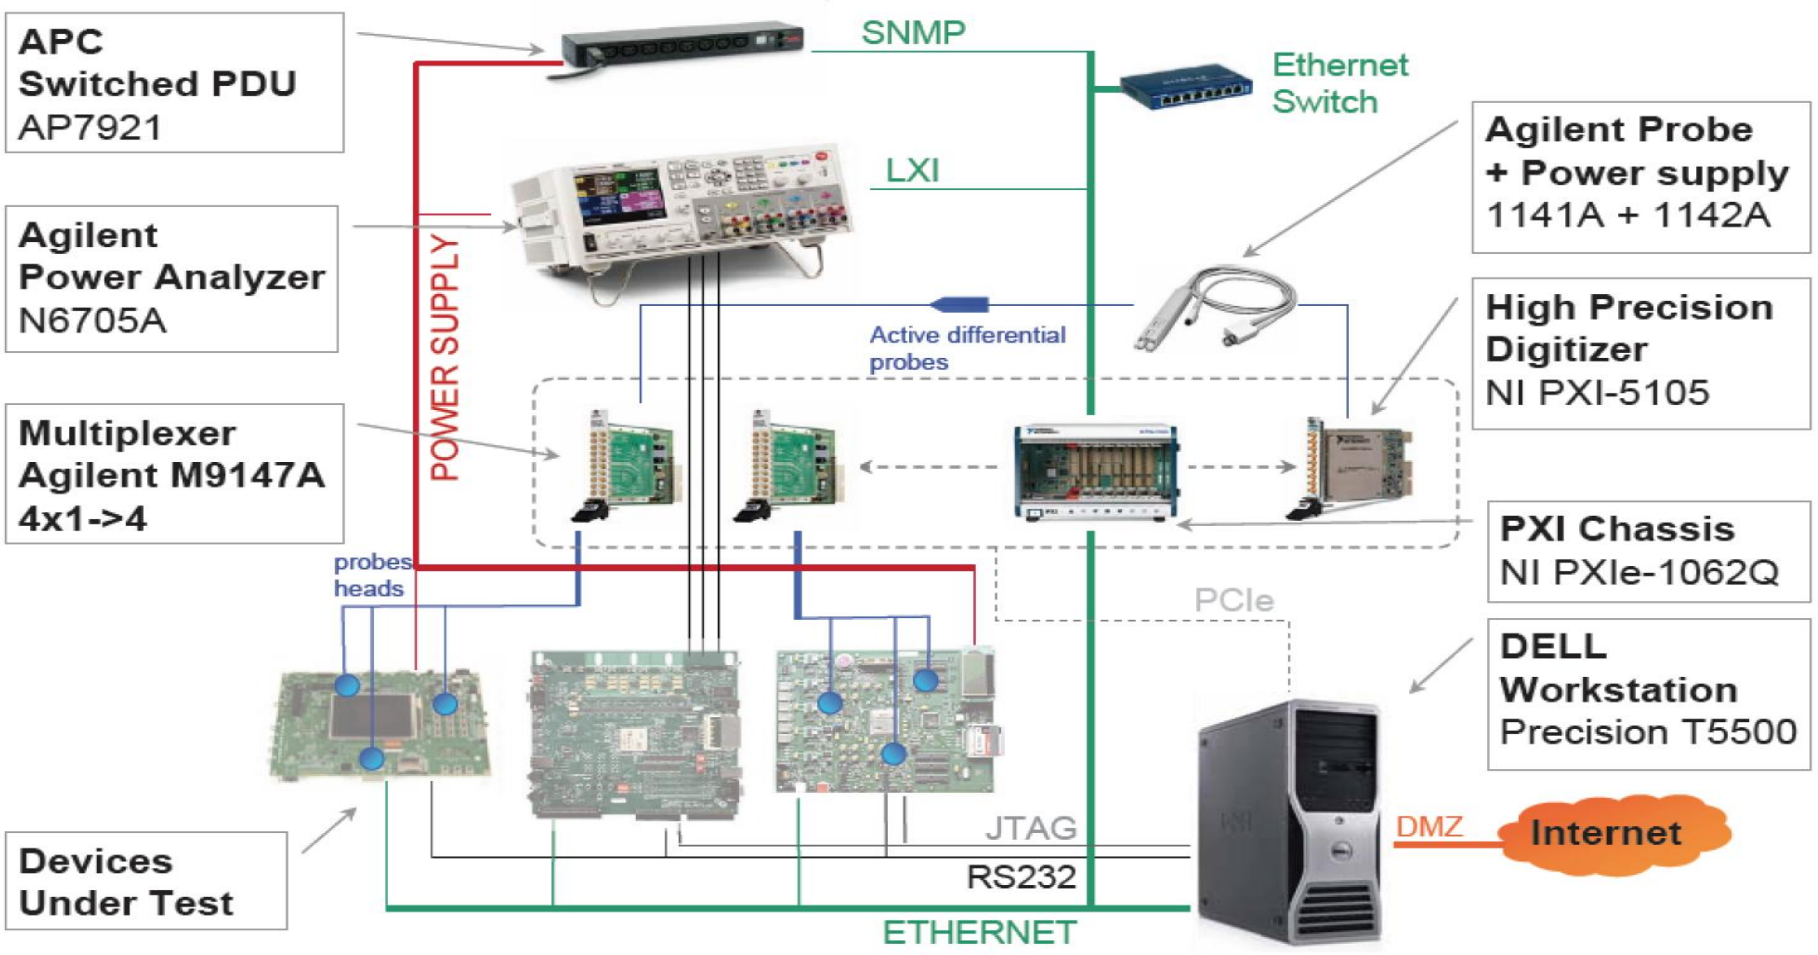

Courtesy: Open-People project

Variation of Instruction Per Cycle (IPC) in Power for ARM Cortex-A8

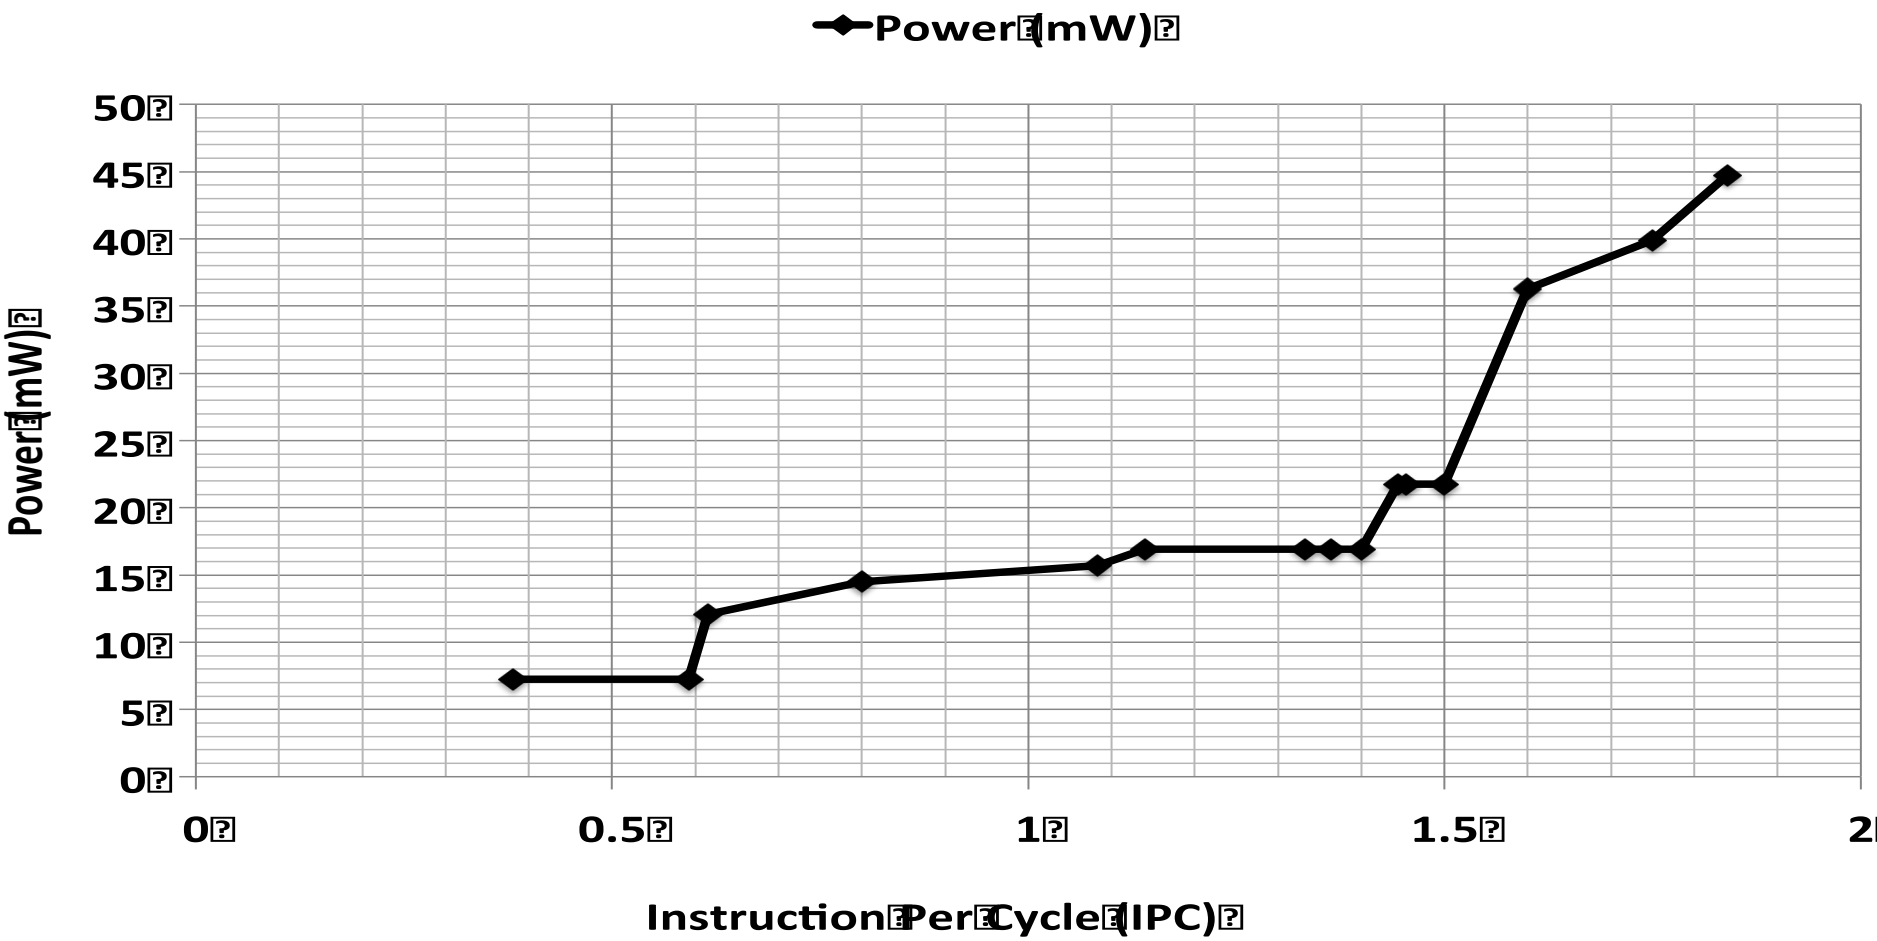

PARA DIME

Power consumption models generated with FLPA

Processors	Power models
ARM Cortex-A7	P(mW) = 0.39 $F_{Processor}$ + 3.65 IPC + 0.66 ($\gamma 1$) + 2.12 ($\gamma 2$) + 8.27
ARM Cortex-A7 (dual/quad-core)	P(mW) = a $F_{processor}$ + $b \sum_{i=1}^{4}(\gamma 1_{c_i})$ + c $\sum_{i=1}^{4}(IPC_{c_i})$ + d ($\gamma 2$) + 5.24
ARM Cortex-A15 (single core)	P(mW) = 0.56 $F_{processor}$ + 8.95 IPC + 1.2 ($\gamma 1$) + 3.5 ($\gamma 2$) + 15.93
ARM Cortex-A15 (dual/quad-core)	P(mW) = a $F_{processor}$ + $b \sum_{i=1}^{4}(\gamma 1_{c_i})$ + c $\sum_{i=1}^{4}(IPC_{c_i})$ + d ($\gamma 2$) + 17.45
Heterogeneous architecture	E(mJ) = $\sum_{i=1}^{n} E_{p_i} + E_{mem} + E_{sync}$ + $E_{I/O}$

Second Step: System Level Power Analysis

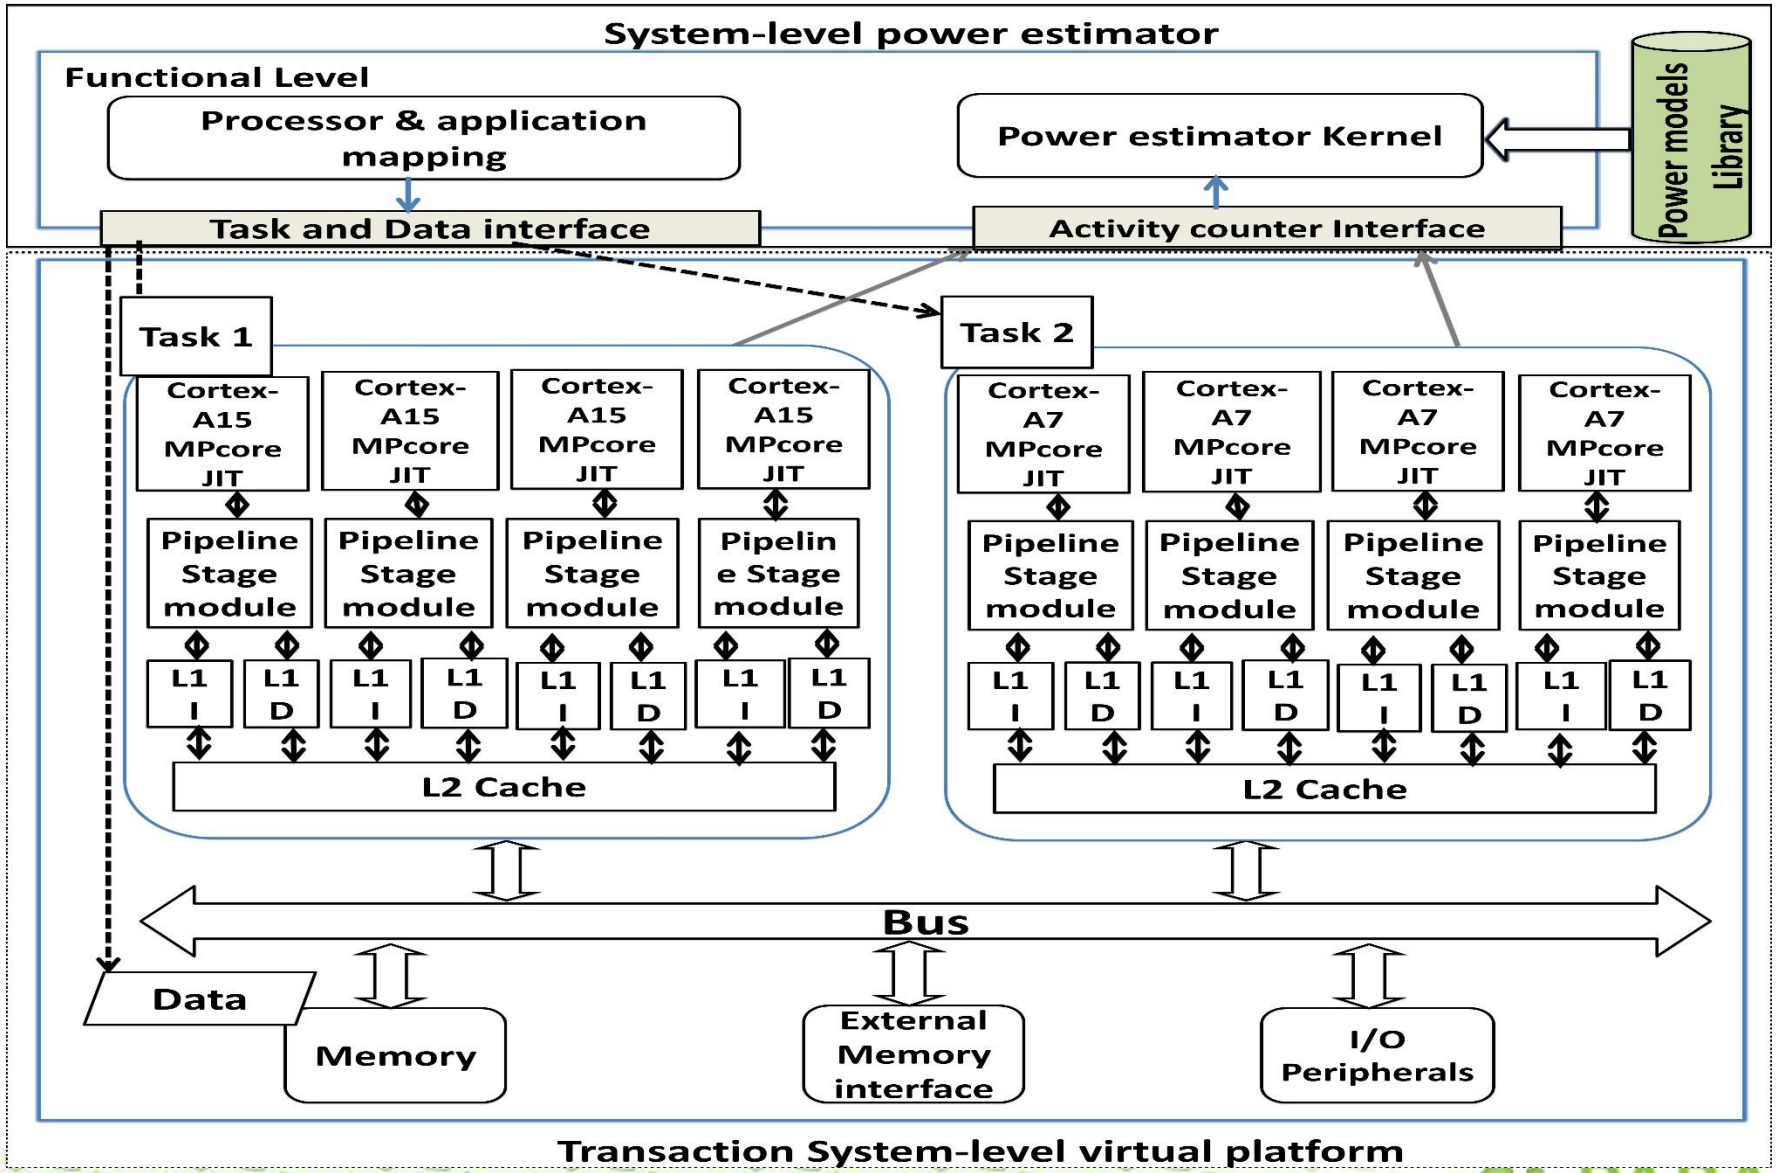

PARA DIME

Result Interface

armCortexA9 Global Stats : armCortexA91

[+] General Info:
Simulation Length: 0,034798196 (s)
Global Utilization Rate: 100,00 (%)

Bootstrap: 0,16% (0,000056044 s)
Idle Time: 0,00% (0,0 s)
All Tasks Time: 74,89% (0,026020107 s)

Number of software tasks: 5
DEMUX1: 23,52% (0,008172901 s)
IDCT1: 10,51% (0,003650125 s)
IQZZ1: 5,72% (0,001988166 s)
LIBU1: 4,11% (0,001426854 s)
VLD1: 31,03% (0,010782061 s)
Estimated Context Switch Rate: 11,38%

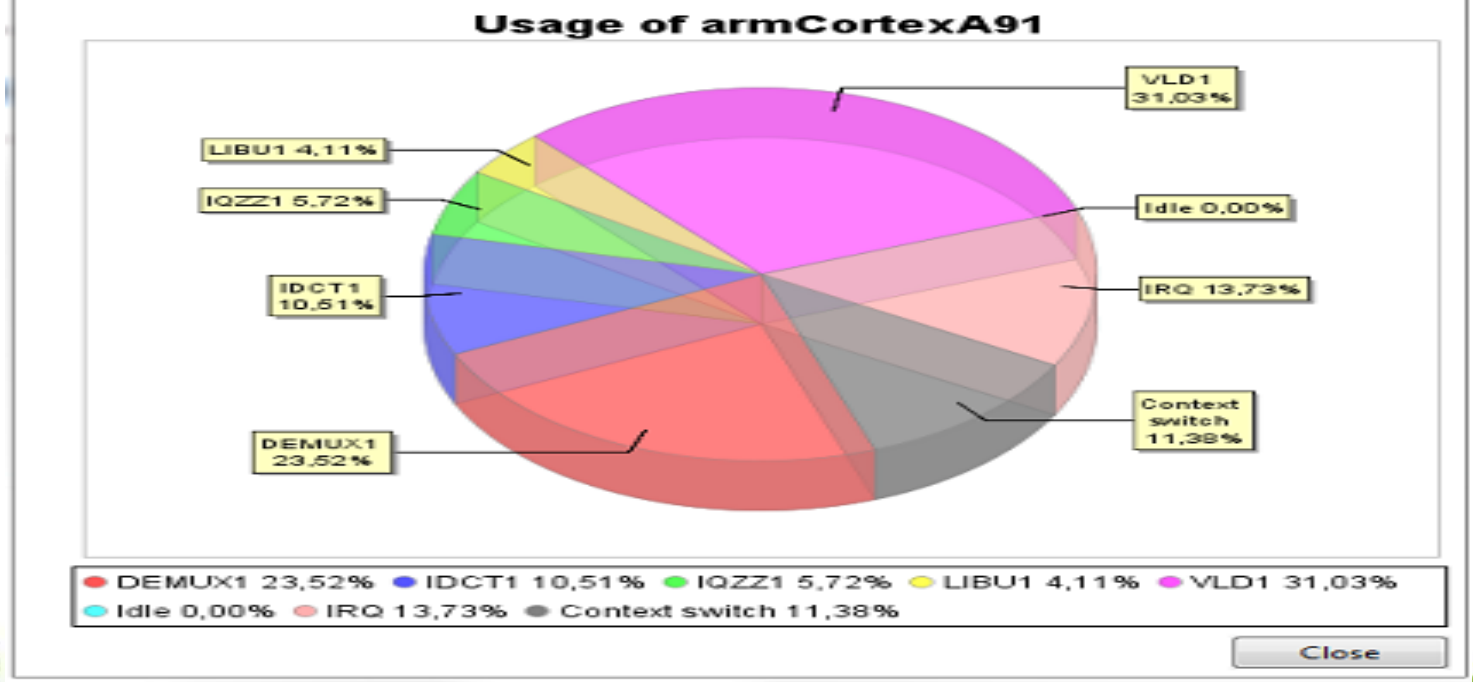

Close

Result Interface

//	Cache miss rate (L1)	Cache miss rate (L2)	IPC	Power (mW)	Energy (mJ)	Power optimization
Estimation results	6.1	0.15	1.28	474,4	860,7	DVFS
Core details	ARM Cortex–A8	1 core	No multicore	JPEG	Power optimized 20 %	Work Load - NO

DESSERT

Timing details

Optimize

Estimate

Export .CSV File

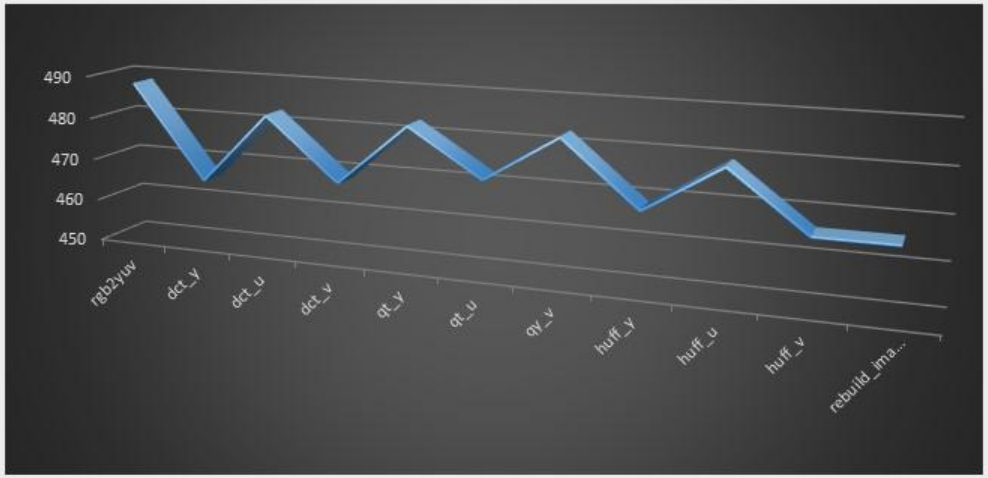

Results and Comparison (Power estimation)

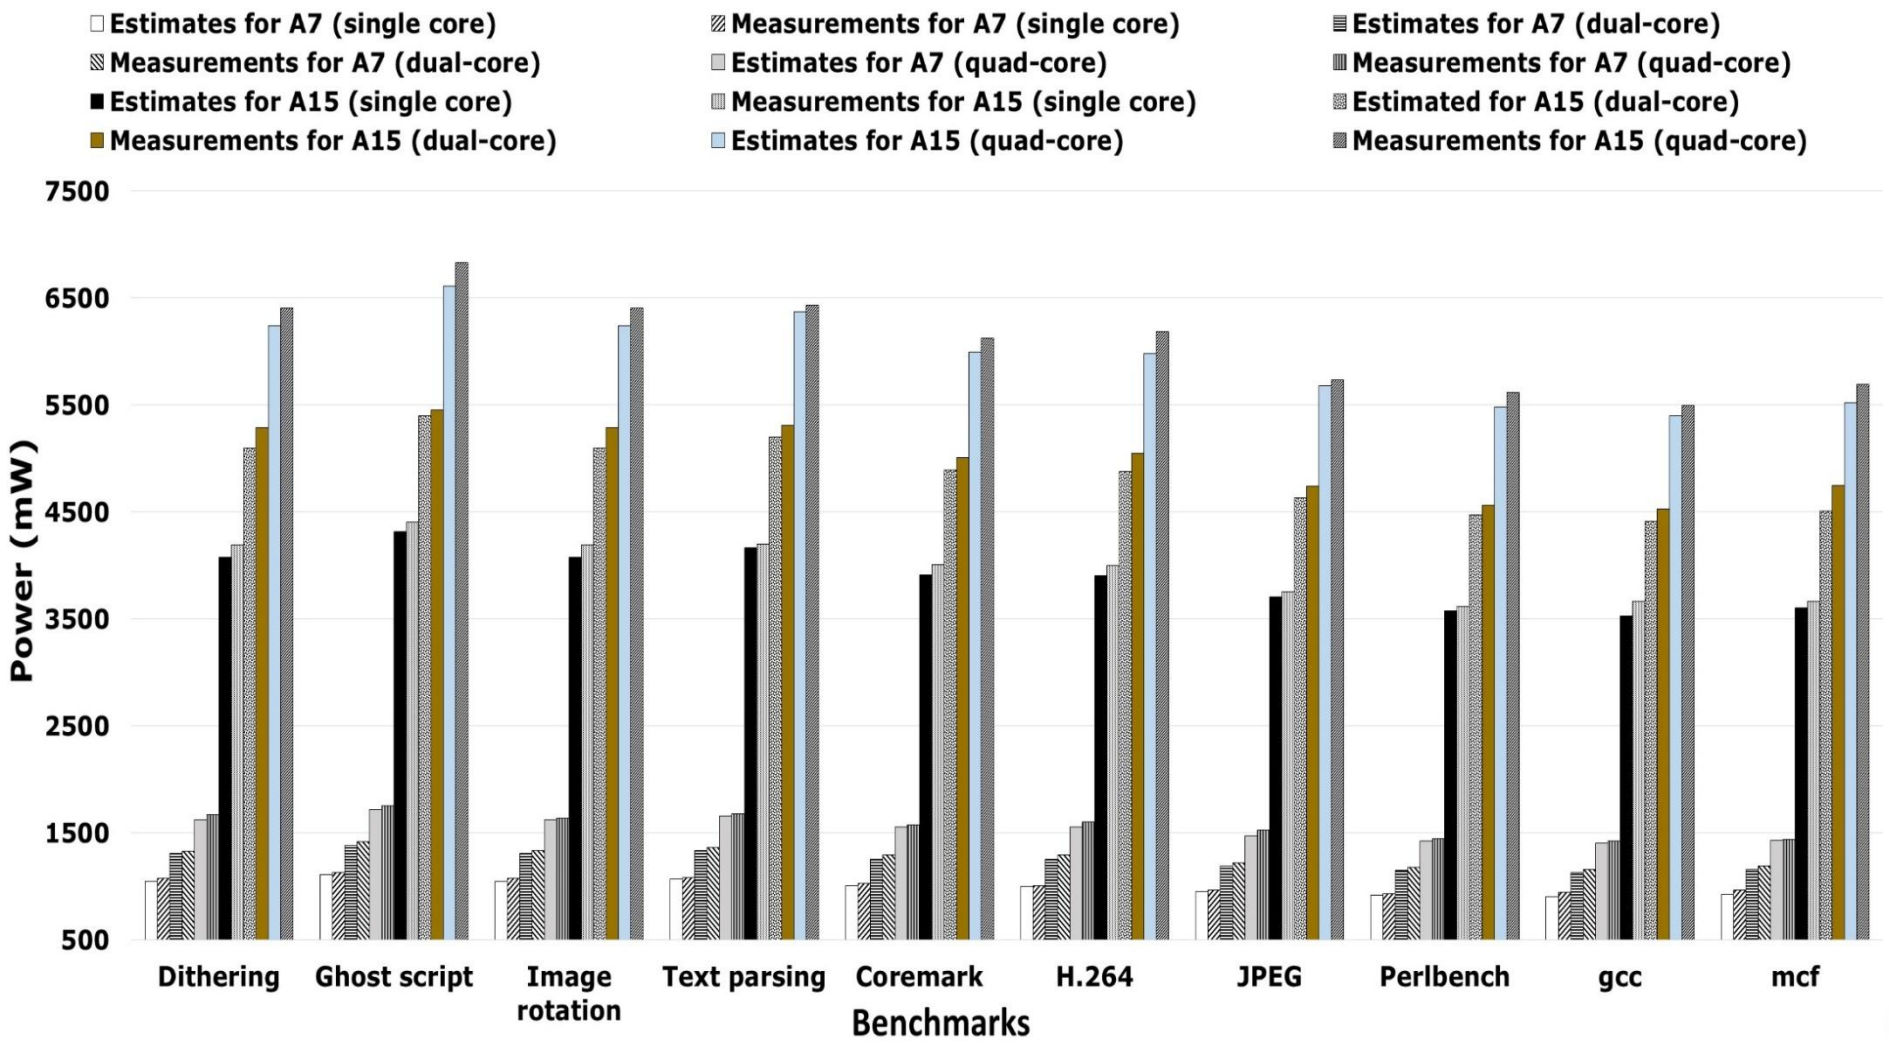

PARA DIME

Results and comparison (Energy)

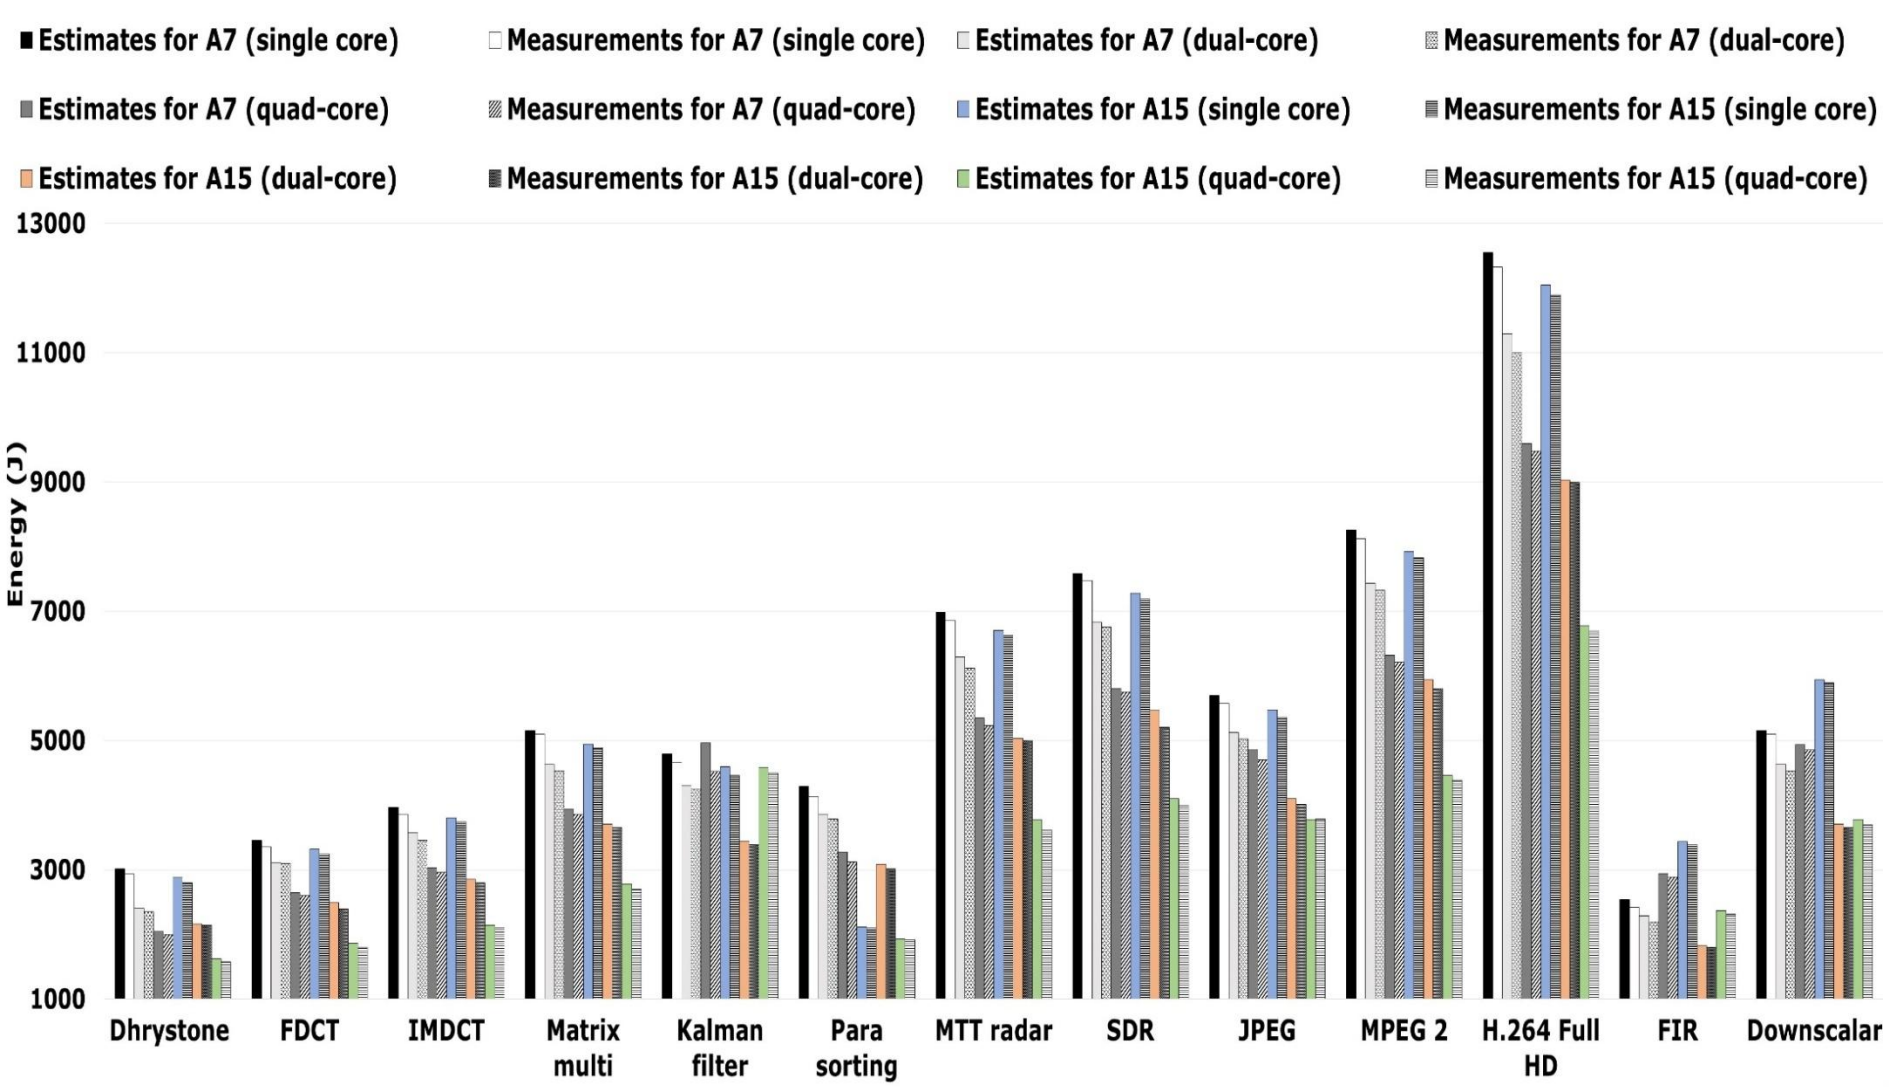

PARA DIME

Third step: Auto optimization

- DVFS
  - Runtime – Inter task DVFS
  - Programmer annotation based DVFS
- Work-load balancing based on task
  - Runtime
  - Programmer based request

Task scheduling

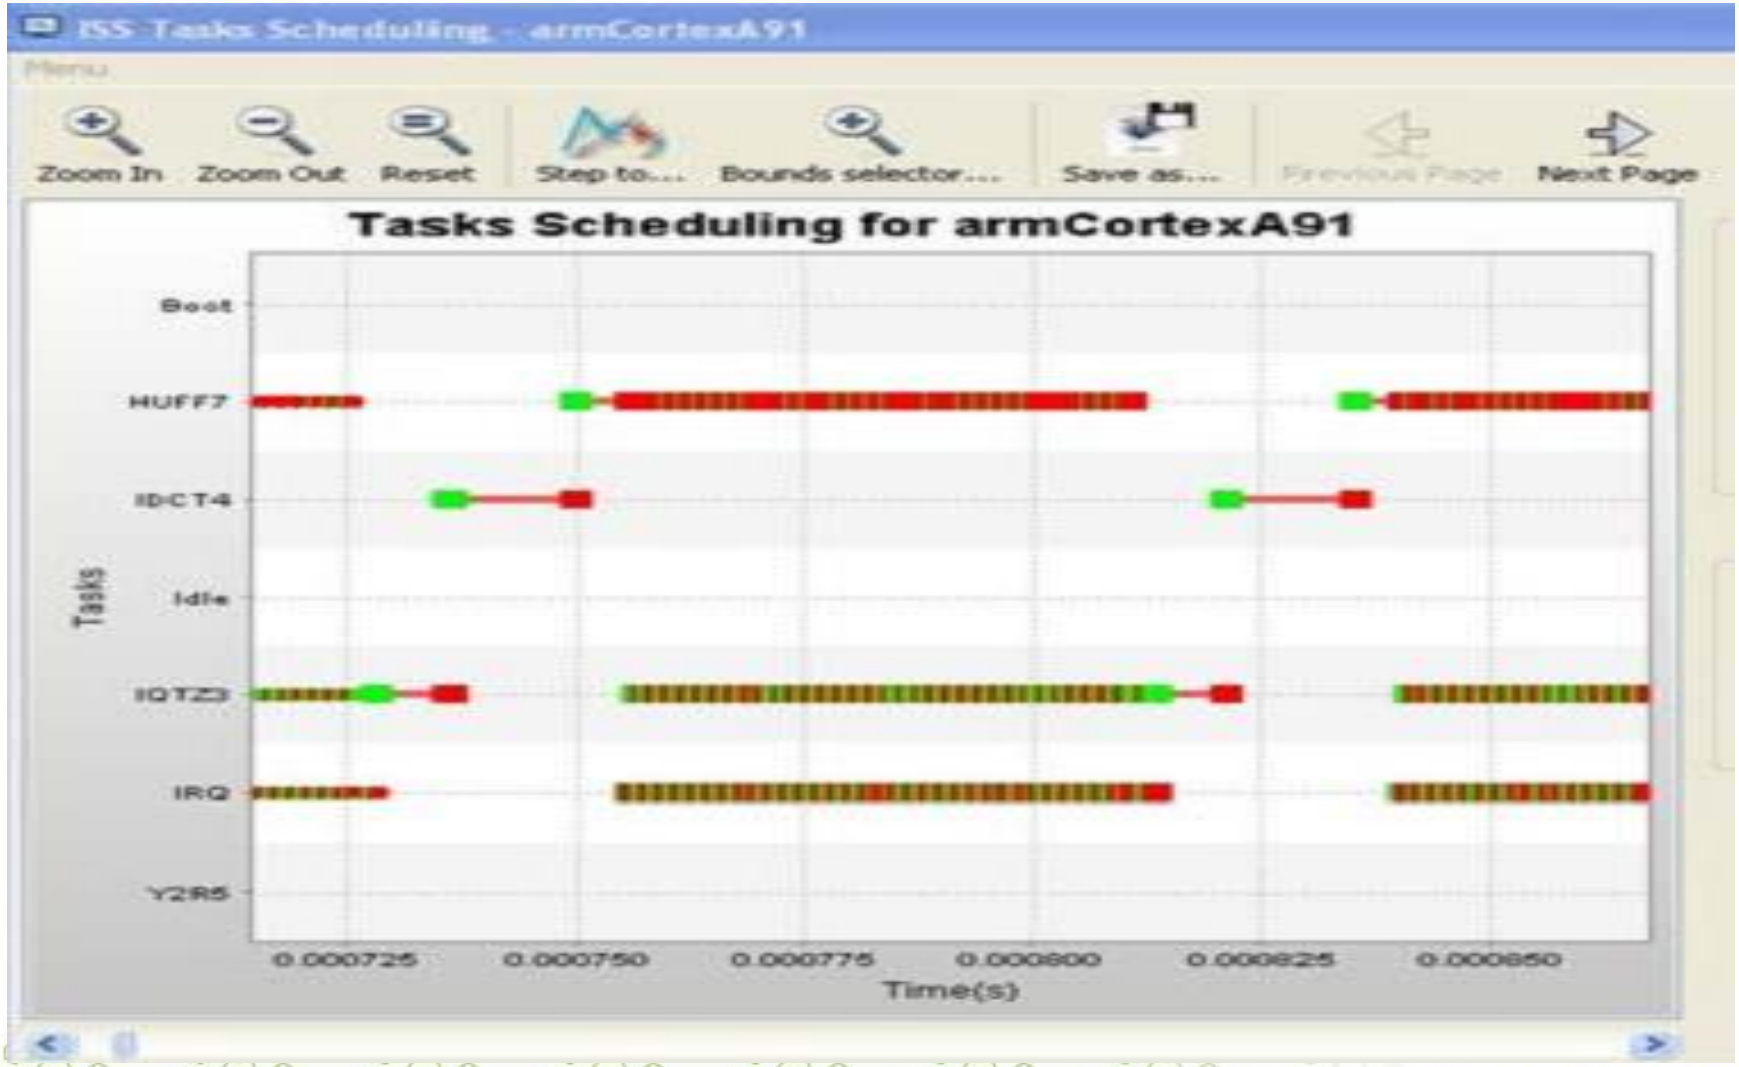

Optimization based on work load balancing

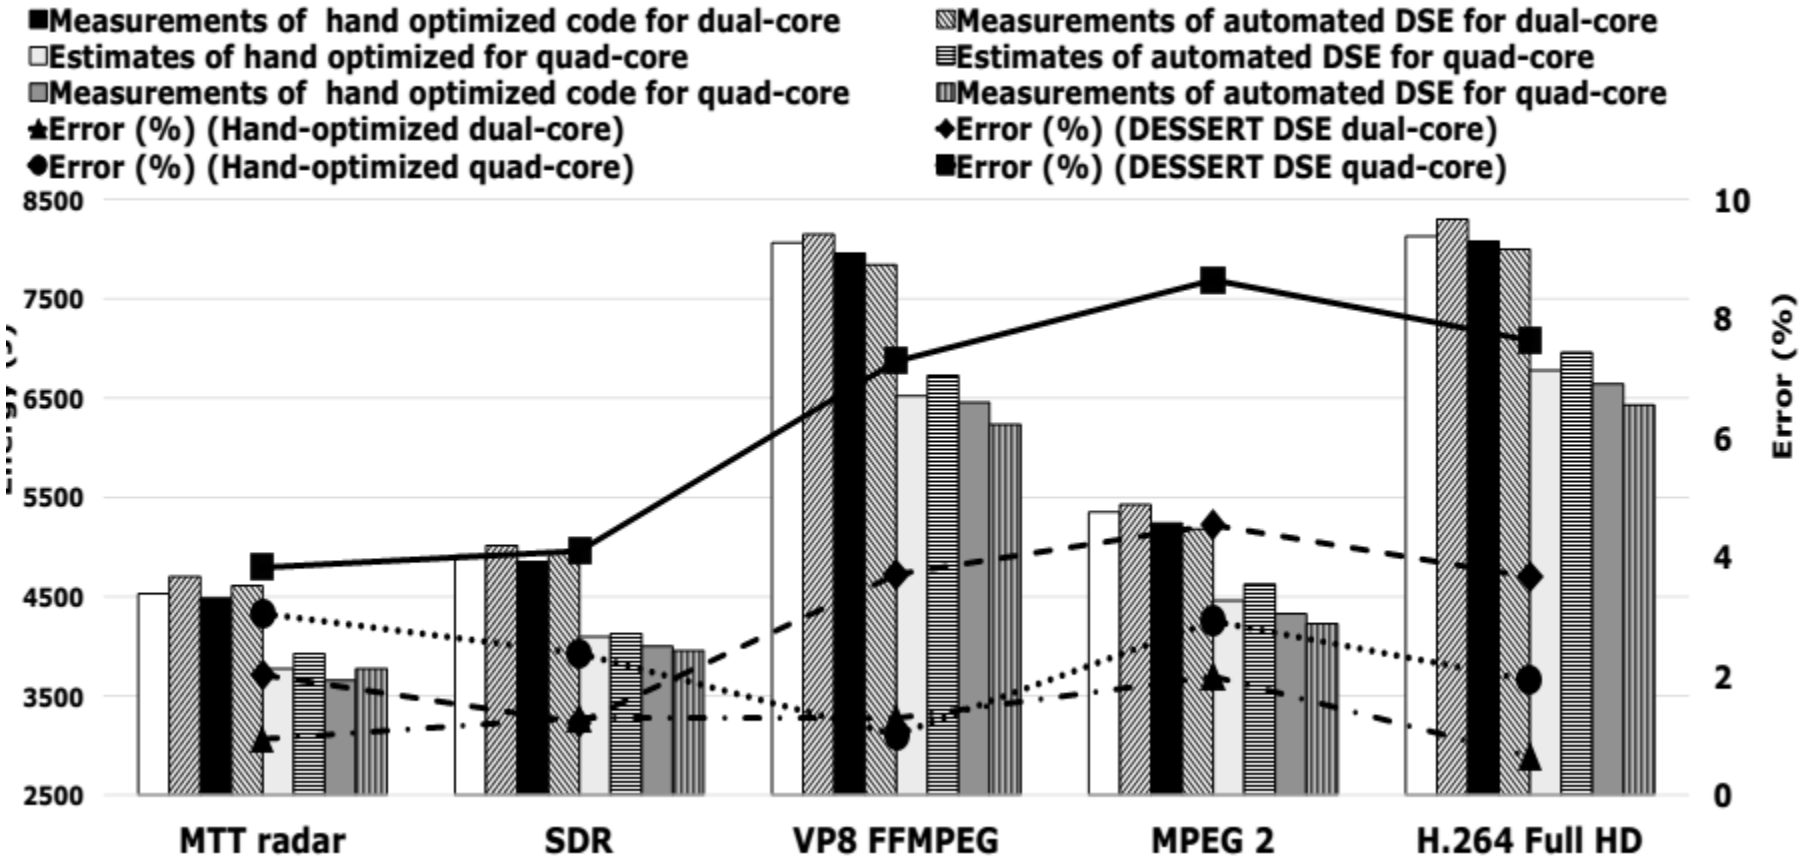

Optimization (Inter task DVFS)

Mixed processor core	Power estimated (mW)	Power measured (mW)	Error (%)
2 processor core	3854	3914	1.5
4 processor core	4587	4626	0.84
8 processor core	5327	5398	1.3

Conclusion

- In our tool, we have proved that our estimates are accurate.
- Adaptable for any kind of complex processor system.
- Added advantage, rapid prototyping of the components and porting of the applications made easy.
- Estimating power and designing applications made easy and time efficient.

The ParaDIME Stack

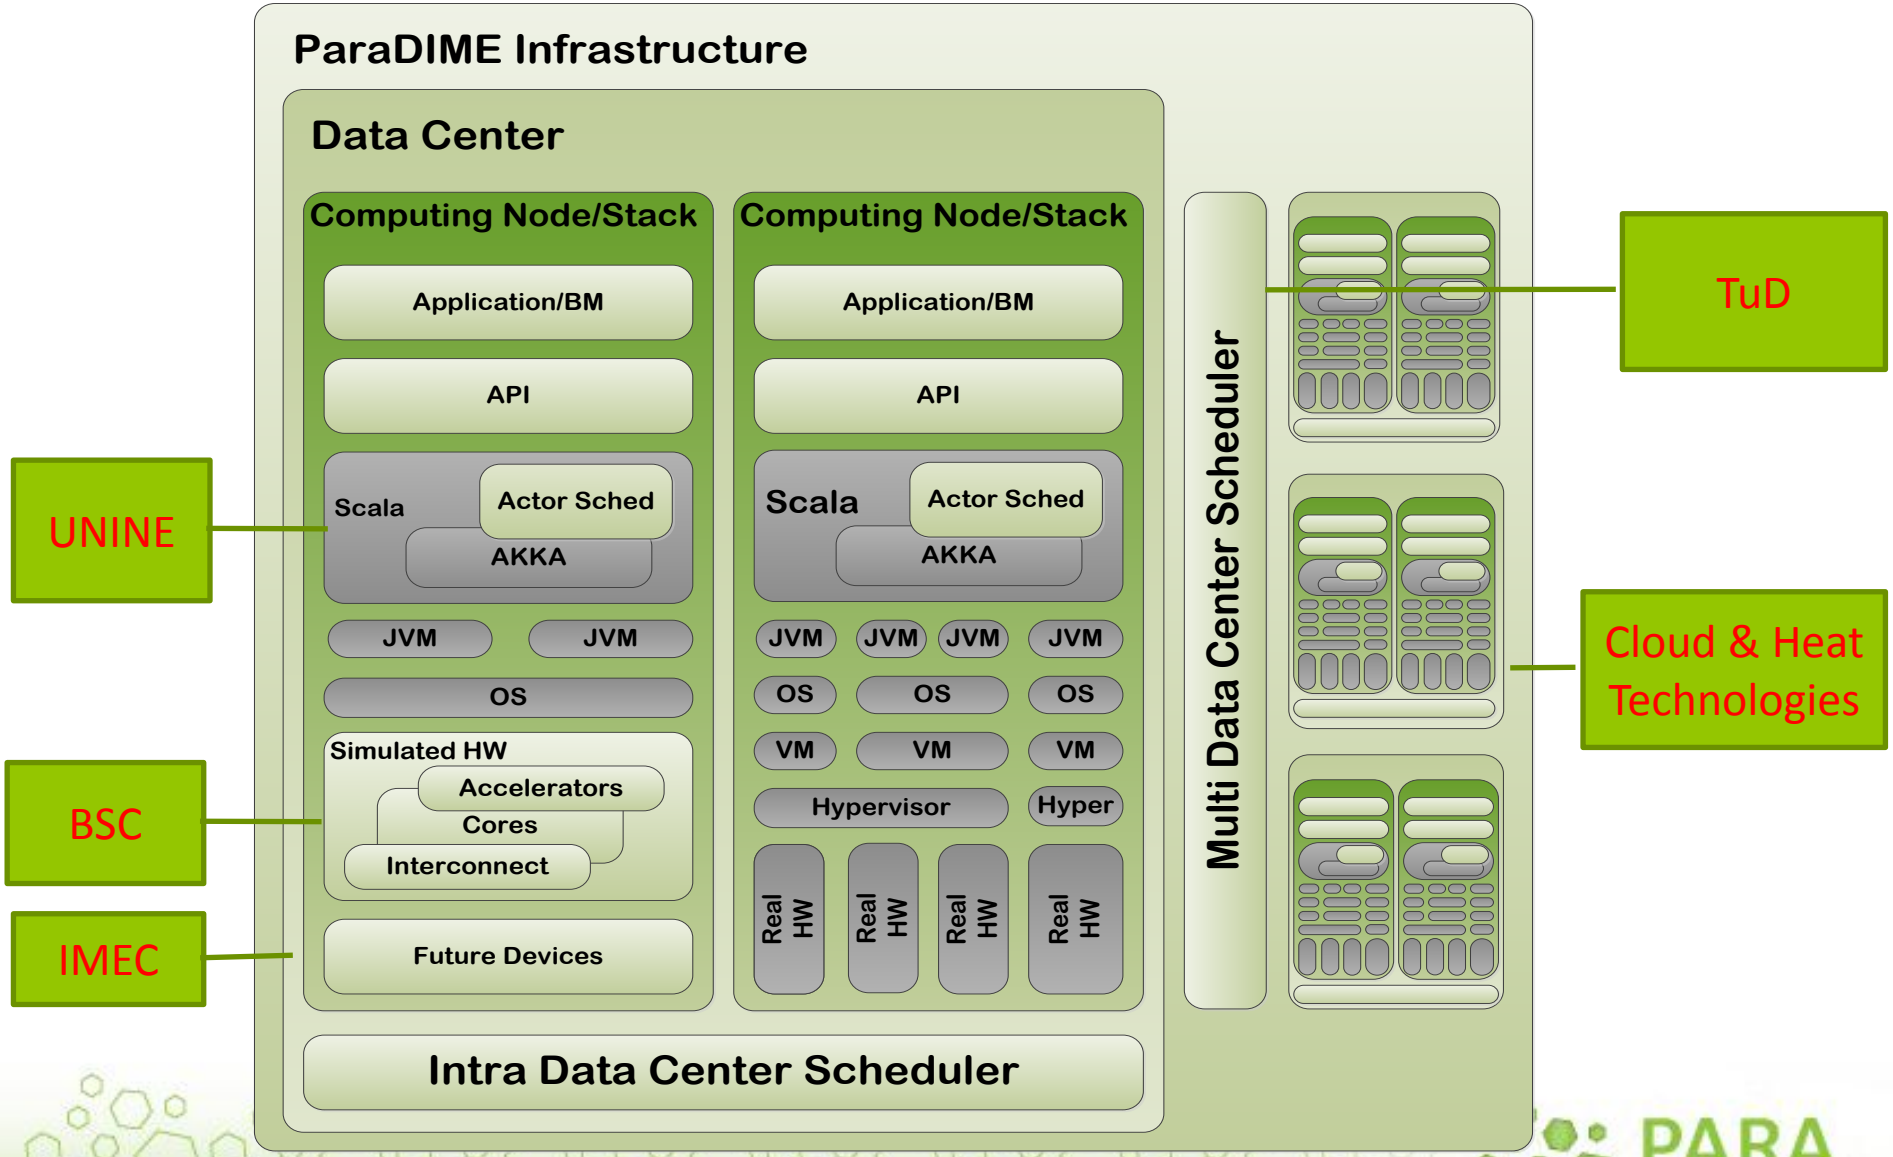

Hardware Architecture

- Energy-Efficient Message Passing
  - Message passing microarchitecture
  - Message passing accelerator
    - Task passing
- Operation Below Safe V_{dd}
  - Automatic HW lowering of V_{dd}
  - SW-guided (low-power annotation)
  - Errors?
- Heterogeneous Computing
  - Architectural level
  - Device level

Heterogeneous system-level environment

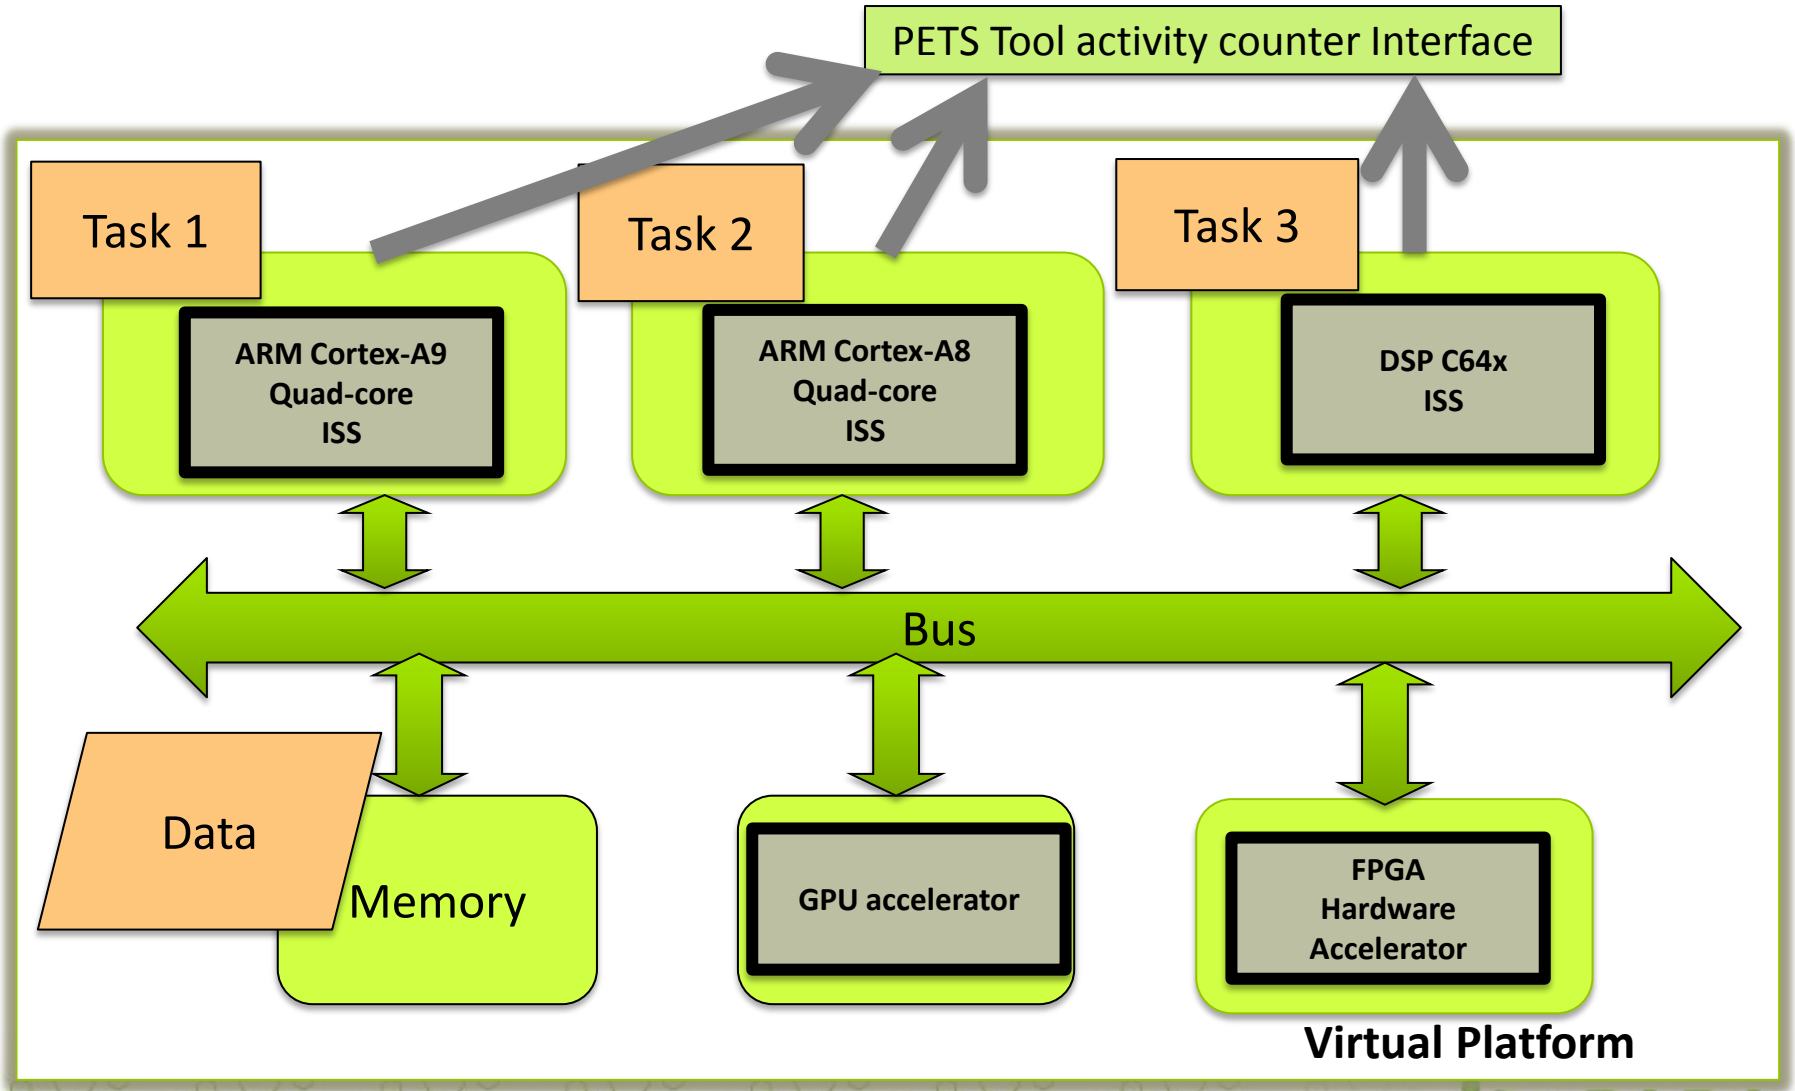

Heterogeneous computing results and comparison

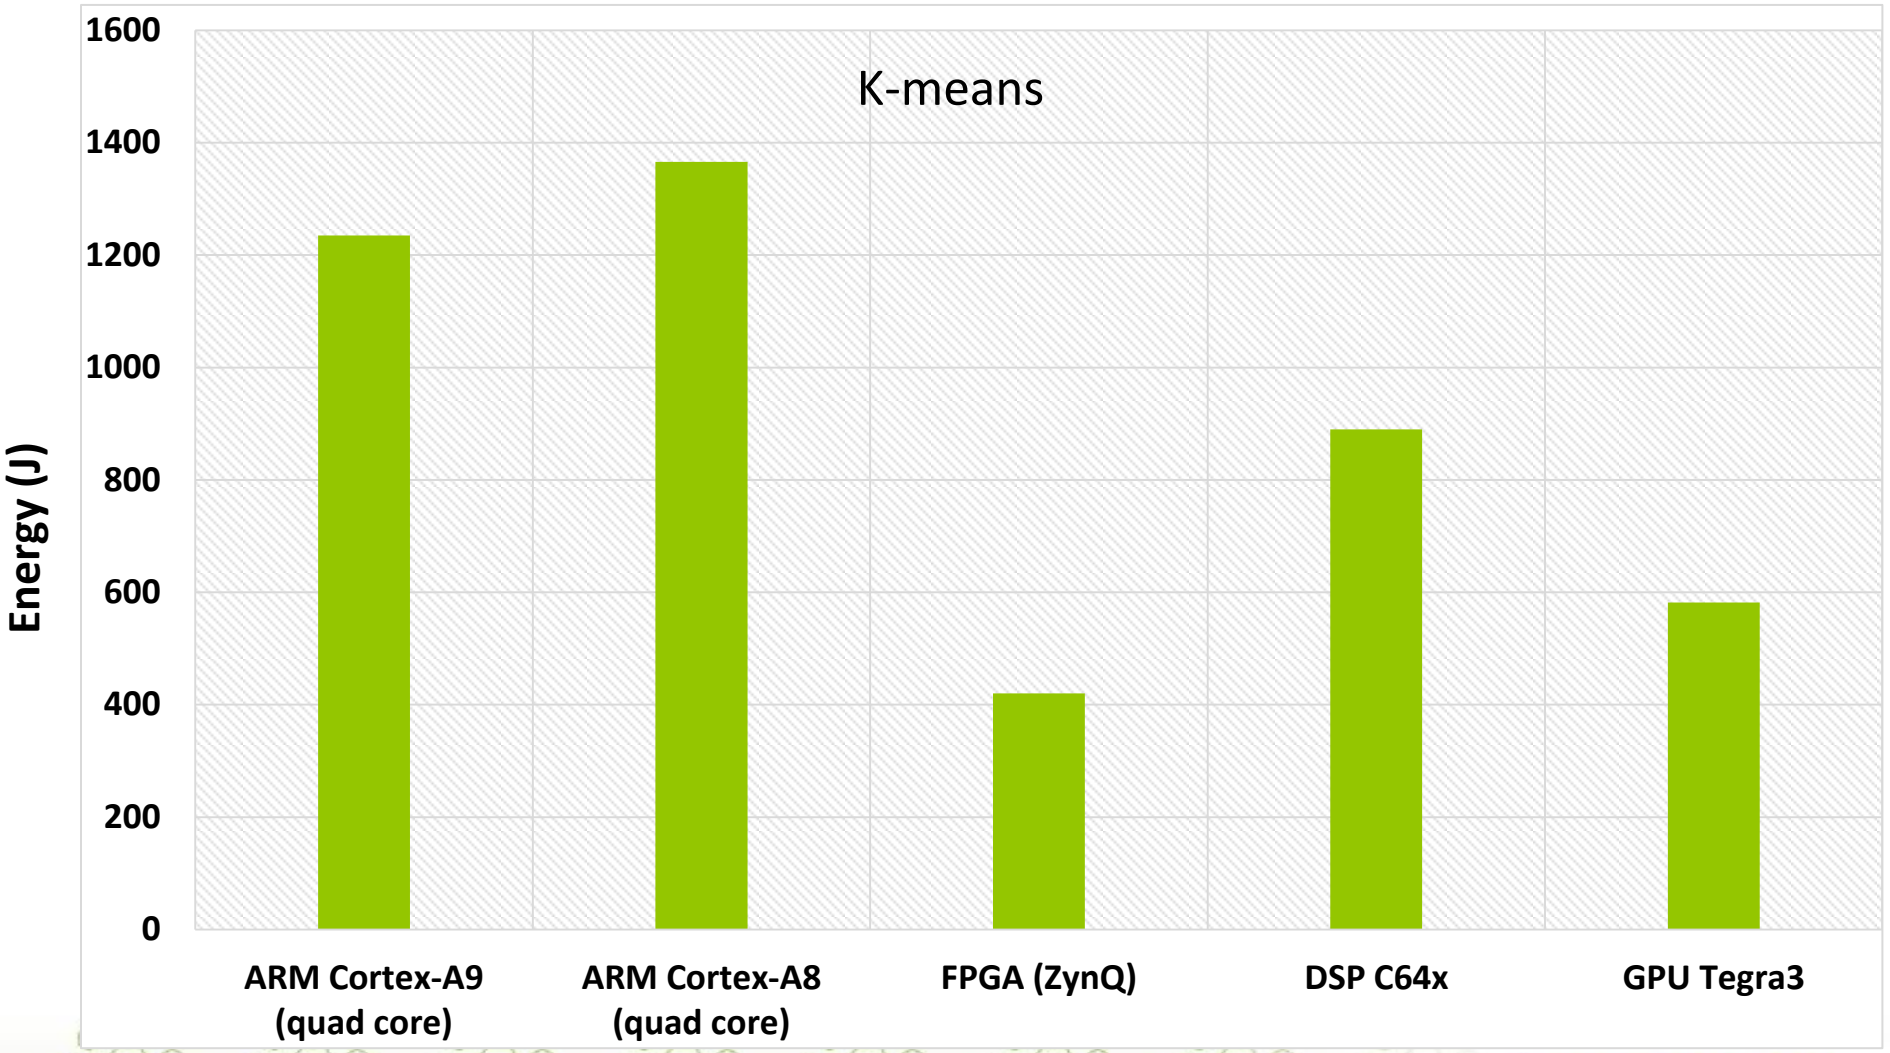

Programming model

- Message passing programming model
  - Actor model (Akka+Scala)
- Annotations to provide information to the hardware
  - Operation below Safe Vdd
  - Approximate Computing

```
@Storage(Array("precise=false", "VF_relax=true"))
var x = 5
@Calculation(Array("VF_relax=true", "VF_det=DMR",
VF_corr=TM"))
def calc(first:Array[double])
```

  - Rewrite/Expand annotated code with Scala Macro Annotations

PARA DIME

Power and Energy Modeling of Multi-core Processors for System-Level Design Space Exploration

Oscar Palomar *(BSC)*

Energy-Aware Computing Workshop 2014

This project and the research leading to these results has received funding from the European Community's Seventh Framework Programme [FP7/2007-2013] under grant agreement n 318693

Computation of power and measurement of voltage for OMAP

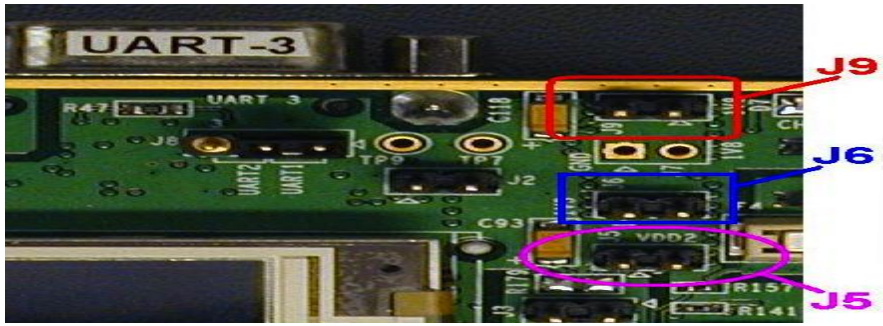

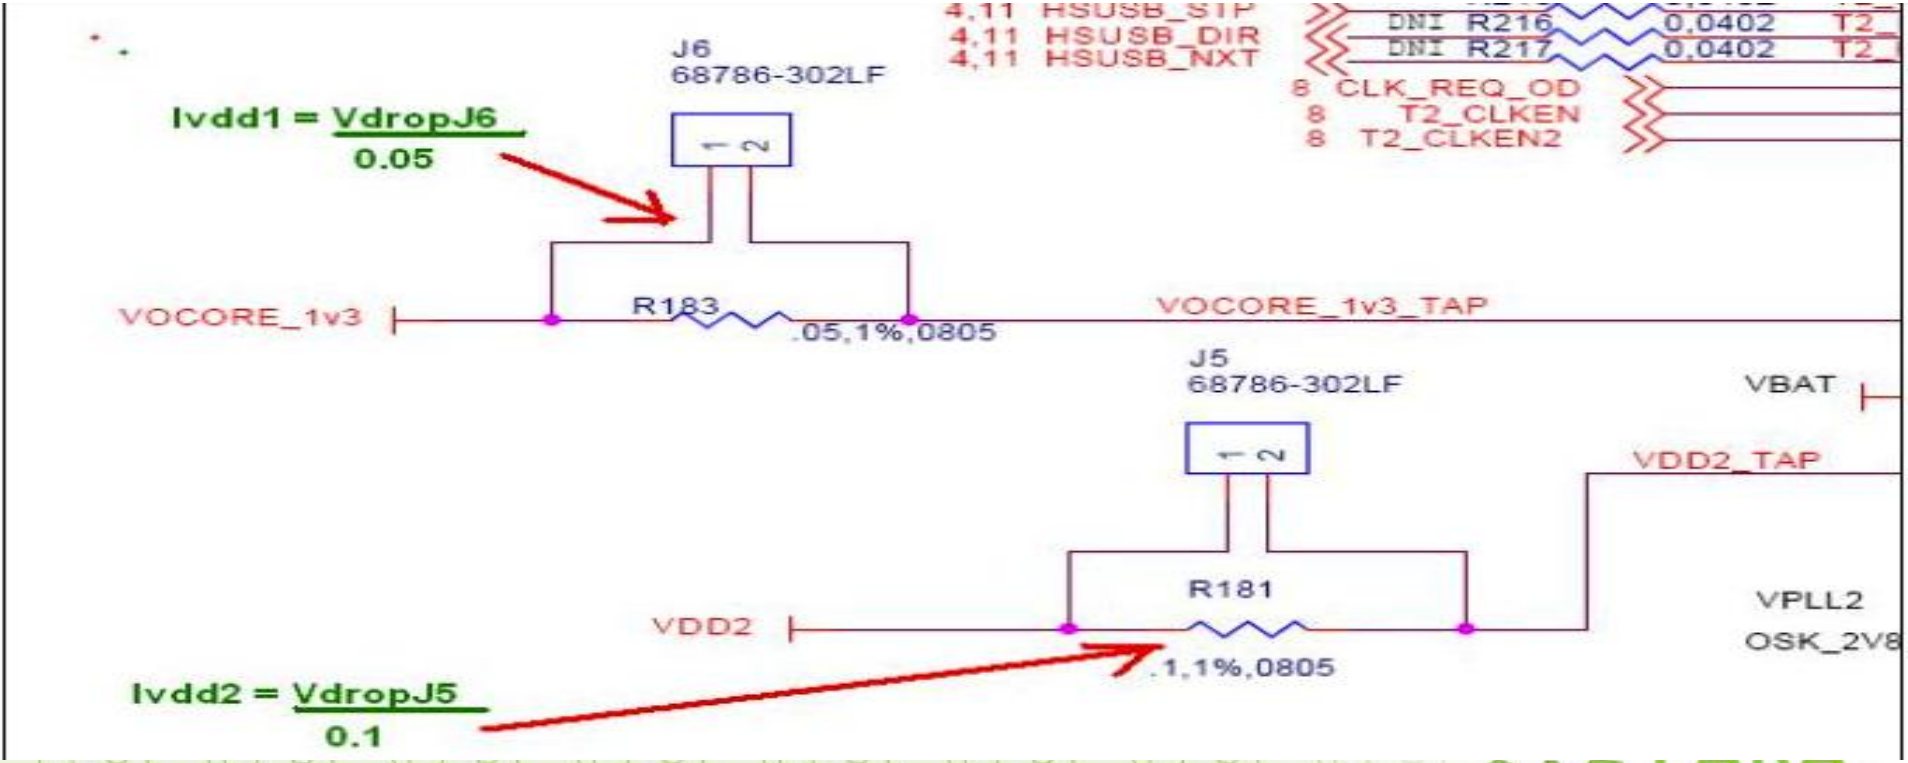

Power measurement environment

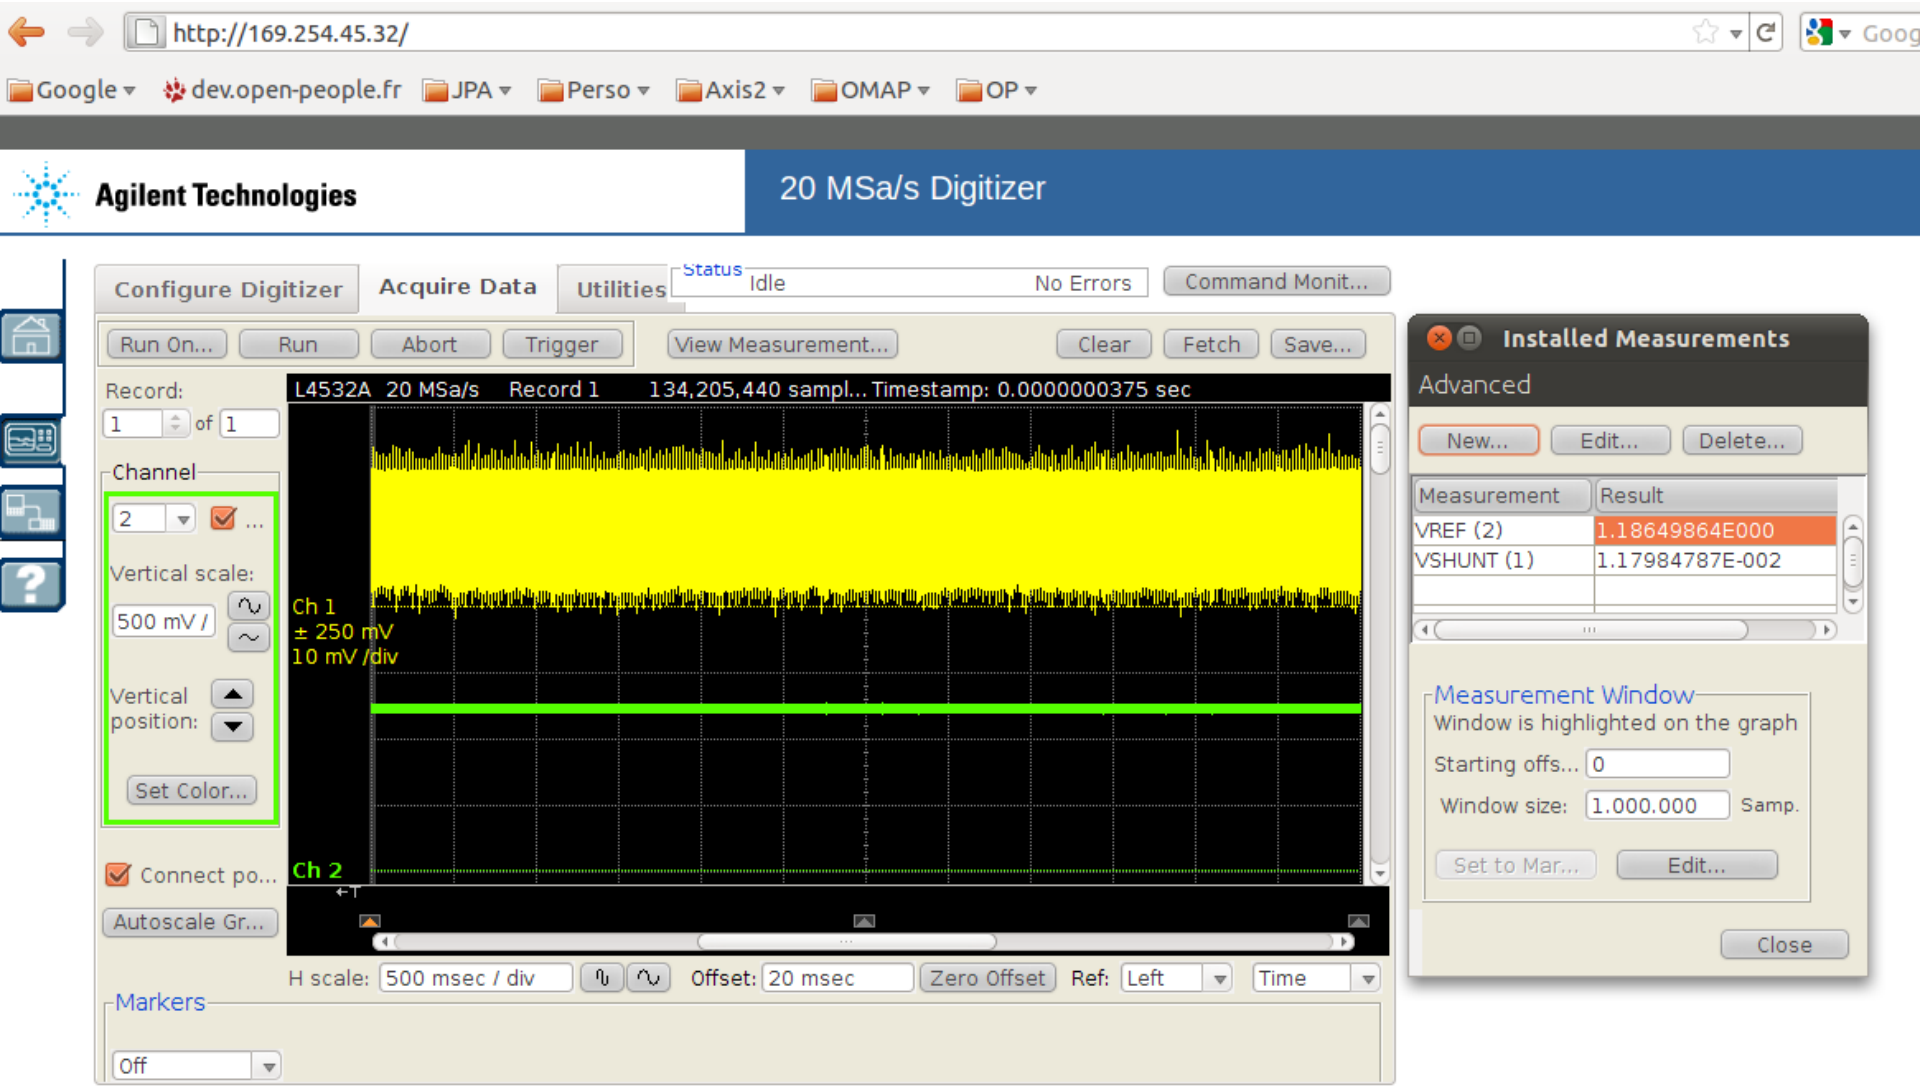

PARA DIME