



Application Characterization:
A comprehensive way of analyzing and understanding
interaction between hardware and software
(applications/compiler/runtime): performance and energy

Jean Christophe Beyler, Intel ECR

BOF Submitted by:

Marie-Christine Sawley, Intel ECR & Gregori Fursin, INRIA



ECR Team

ECR team includes over 25 high level researchers

ECR Team is part of research labs in Europe and North America

ECR is a member of the Intel EMEA HPC Exascale labs together with ExaScience lab and ExaCluster lab

Part of DSG-DCSG



2 main Streams of Research

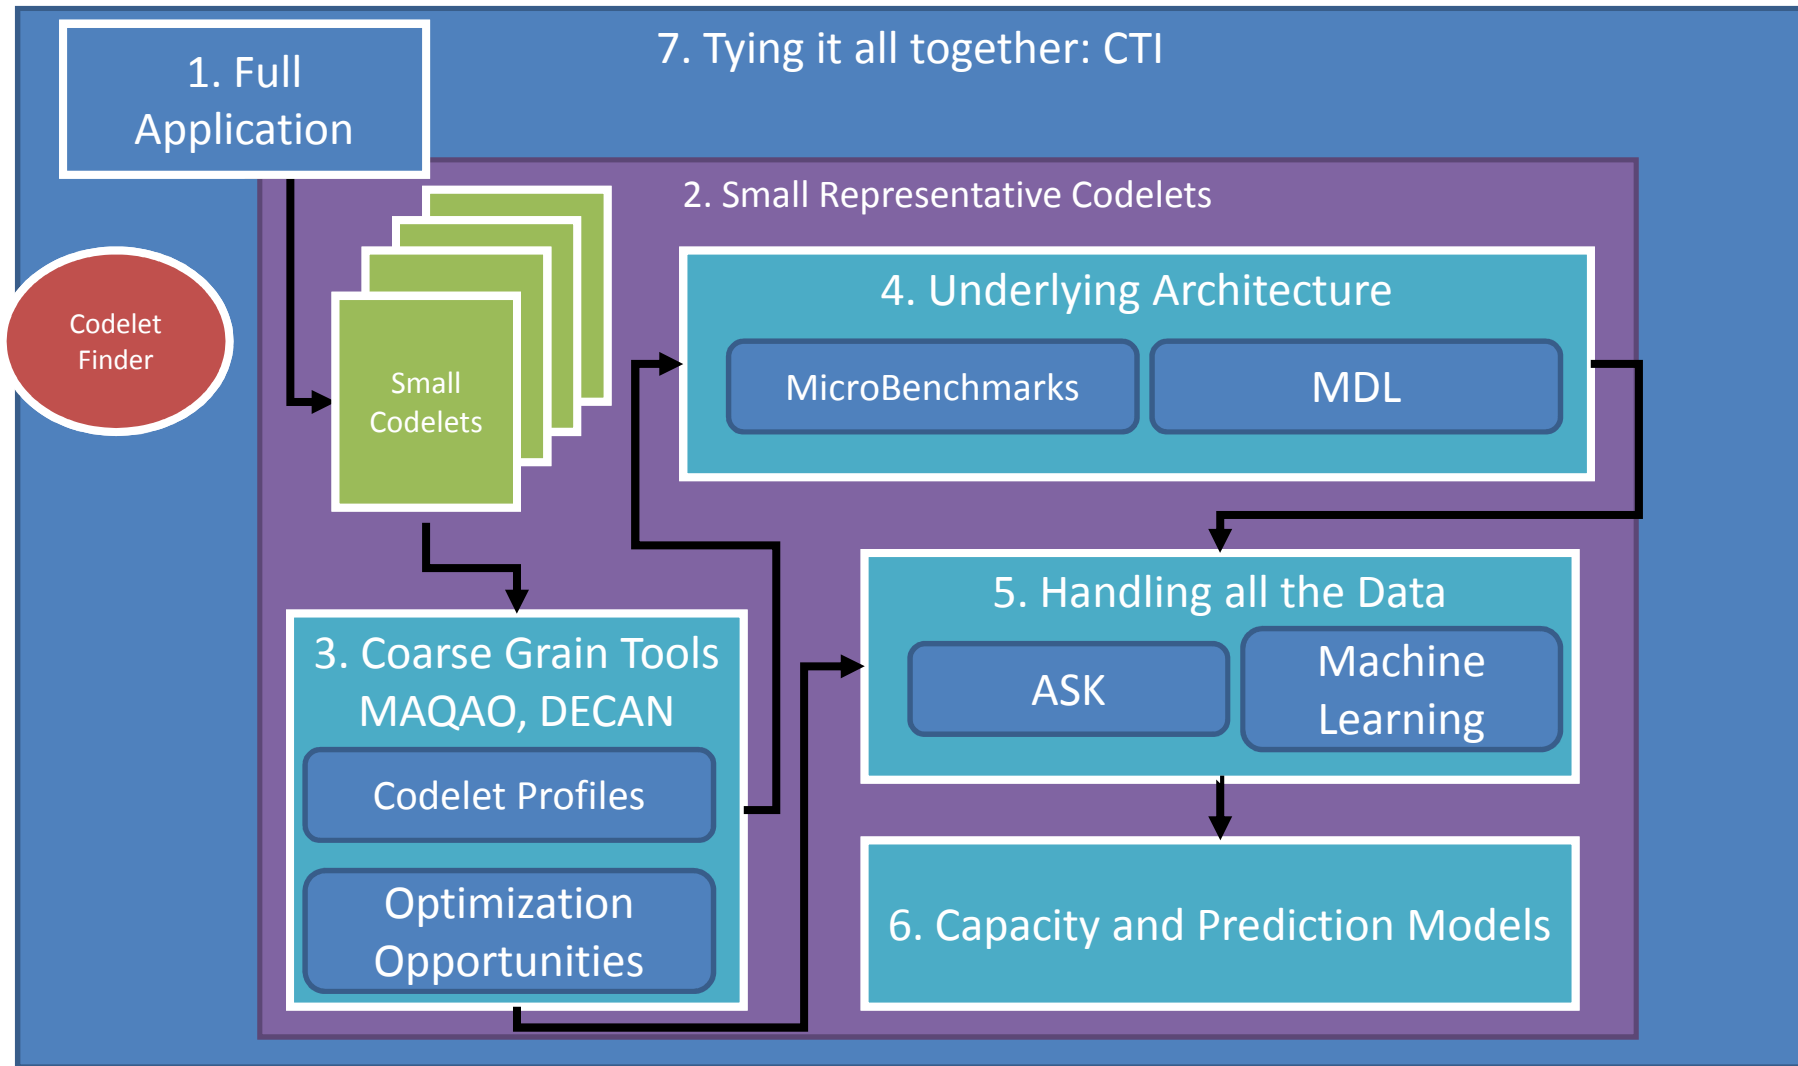
Software for application characterization and performance optimization

Extract fine grain information about the interaction of whole software with the underlying architecture

Application co-design

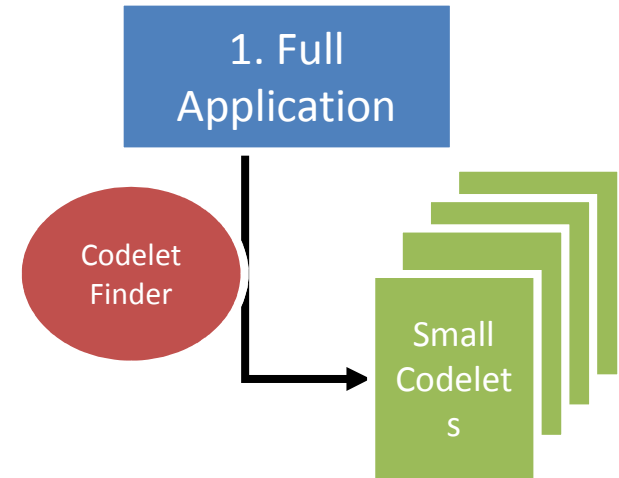
Leveraging from the low level information and the capacity of new architectures for enabling progress in science using computational power

Outline



Cutting the Application Up

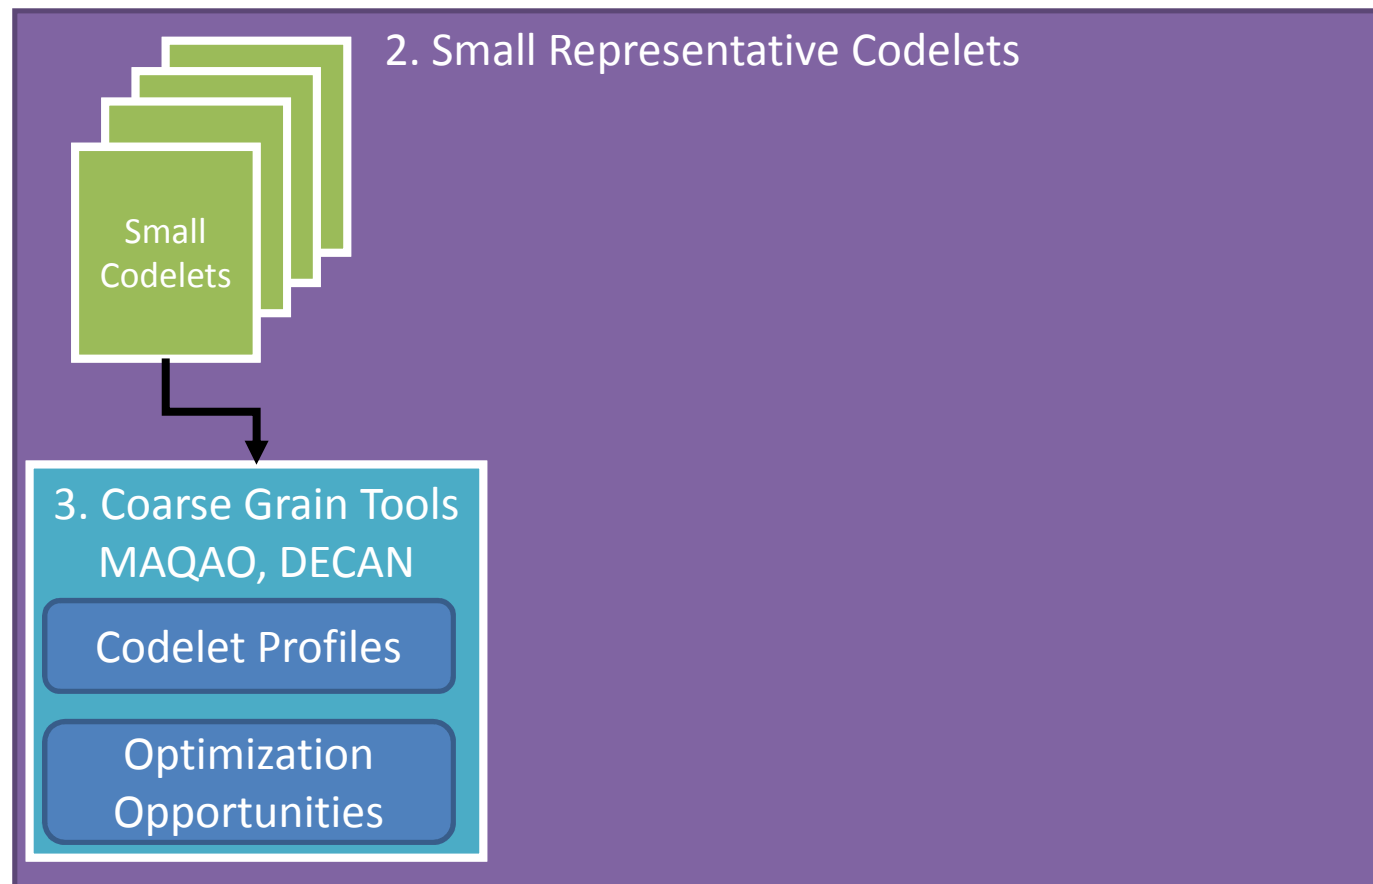
- First step: finding the hot spots
- Considering full applications is difficult
 - Study the hotspots
 - Automatic solution:
 - Using CAPS Enterprise's Codelet Finder tool
- Second step: work on the codelets separately



Codelet Finder

- Key features
 - Implemented by CAPS Enterprise
 - Handles C or Fortran codes
 - Automatically detects hotspots and extracts loops into:
 - Kernel, wrapper, data input
 - Data input is retrieved by a core dump before the kernel
- Future work
 - Allow users to modify input data easily
 - Add more supported constructs

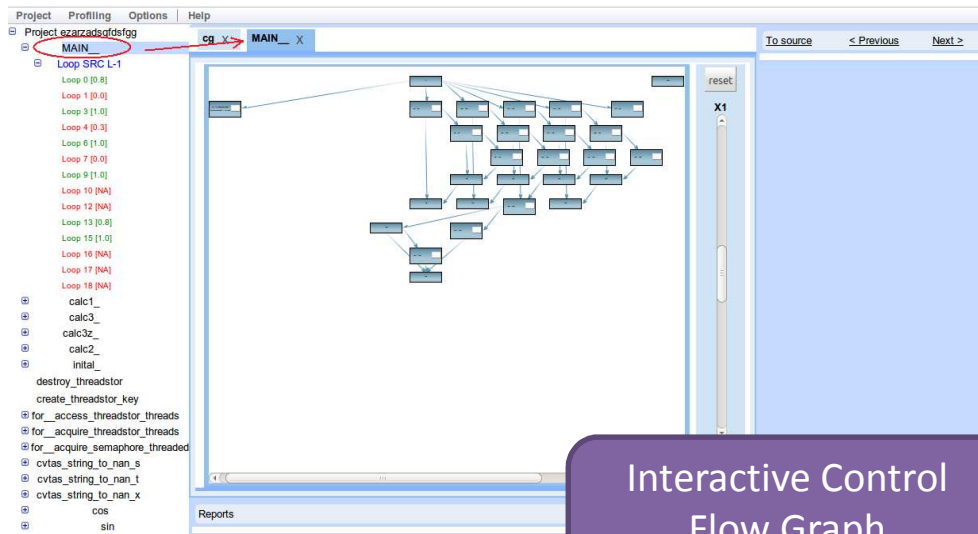
Outline



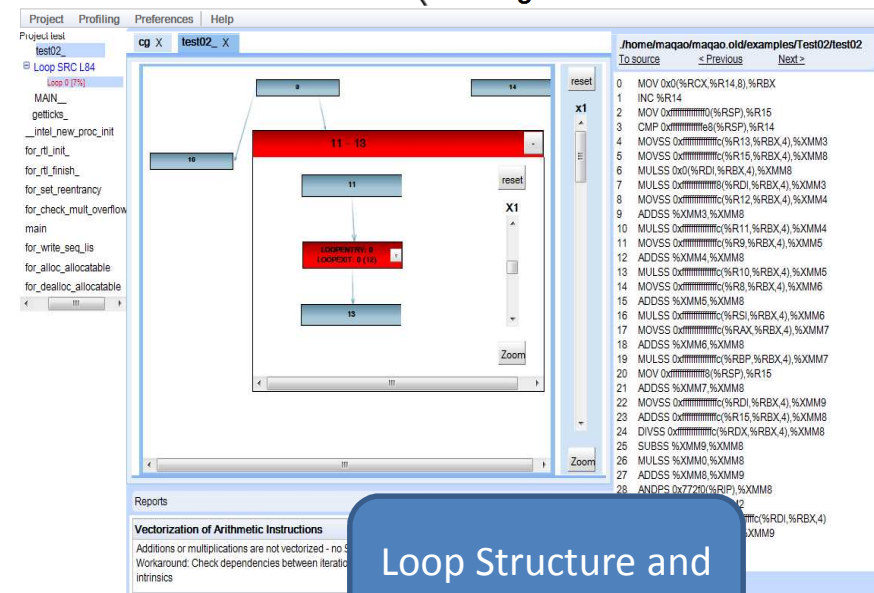
Tools to Quantify

- Logical next step
- Measure performance with profiling tools such as MAQAO
 - Provides important information
 - What can be optimized?
 - What is obtainable?

MAQAO.org Web Interface



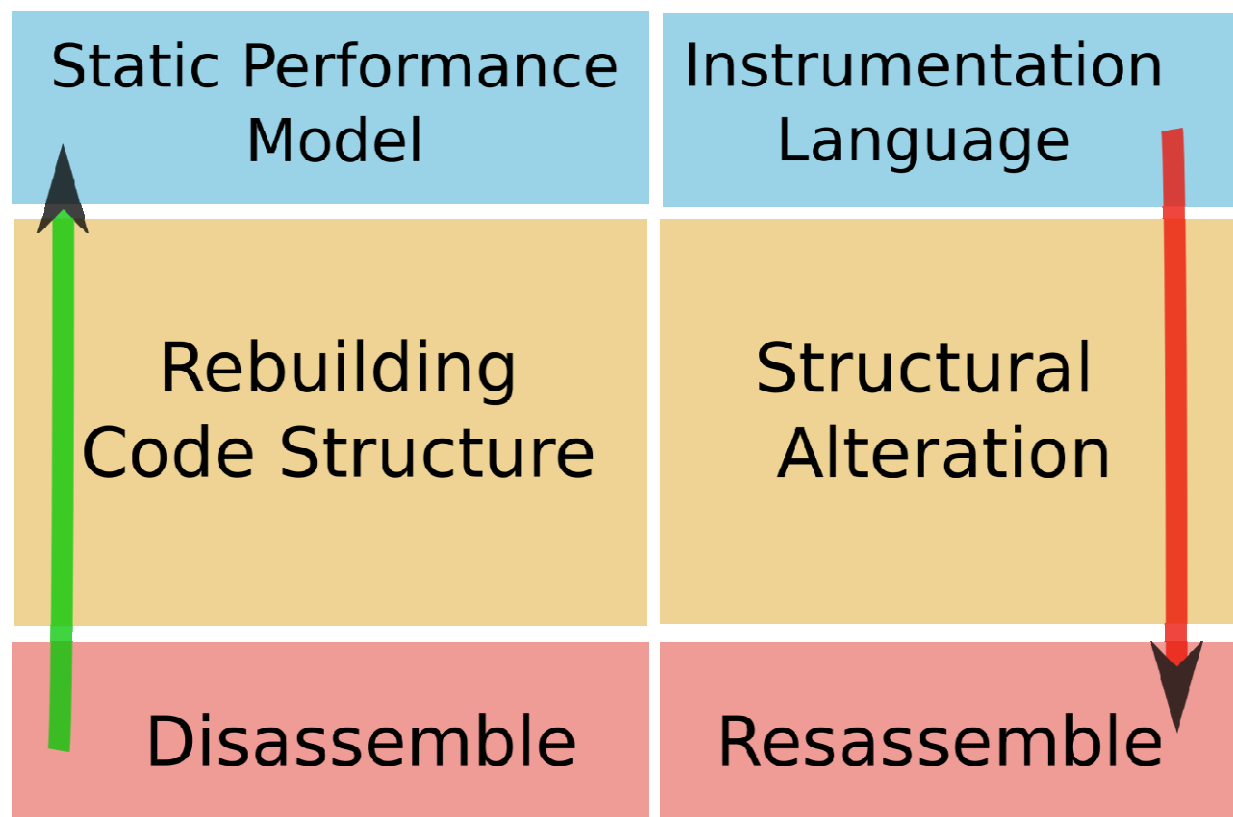
Interactive Control
Flow Graph
(Exagraph)



Loop Structure and
Assembly Code



- Disassemble or reassemble SSE and AVX binaries
- Performance model for Core2, Nehalem, and Sandy Bridge
- Low overhead profiler
 - < 100 cycles per probe
 - OMP compliant

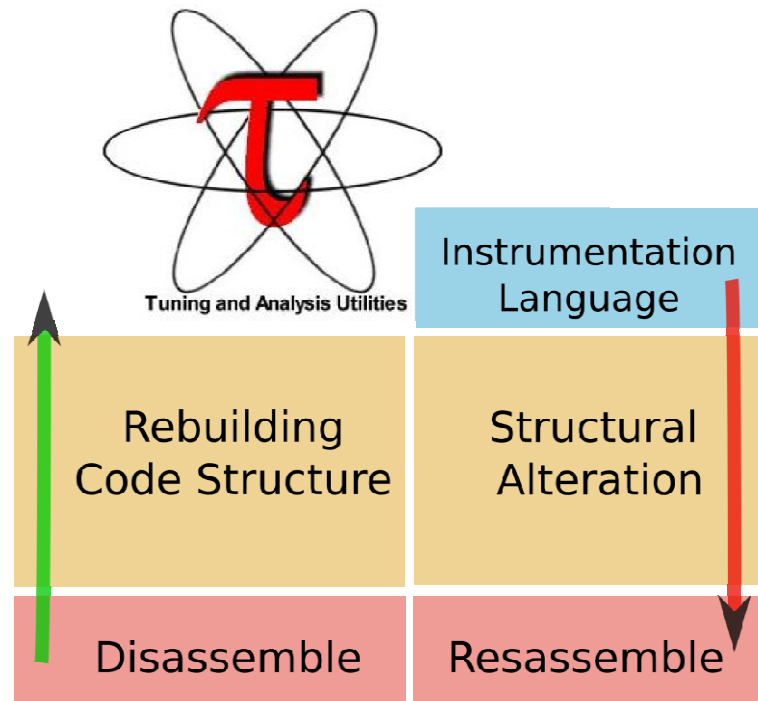




- Performance analysis
 - No pragma or source code alteration
 - Vectorization ratio
 - Detailed pipeline model:
 - Dispatch, decoder, LSD, per port pressure
 - Memory traffic
 - Aggregate memory instructions per group
 - Unrolling factor
- Static performance prediction
 - 'What if' the code is fully vectorized
 - 'What if' the data is stored in L1



- Modular Architecture
 - Demo of TAU using binary capacities of MAQAO at SC11



QMC == CHEM with MAQAO

```
!DIR$ VECTOR ALIGNED
do j=1,LDC
  C1(j,i)=C1(j,i)+(A(j,k_vec(1))*d11 &
    + A(j,k_vec(2))*d21 &
    + A(j,k_vec(3))*d31 &
    + A(j,k_vec(4))*d41)&
  C2(j,i)=C2(j,i)+(A(j,k_vec(1))*d12 &
    + A(j,k_vec(2))*d22 &
    + A(j,k_vec(3))*d32 &
    + A(j,k_vec(4))*d42)&
enddo
```

	B	G	H	Ai	Aj	AK	AL	AM
1	In green, SSE/AVX vectorized loops							
2								
3	Nb used reg.		Performance in L1					
4	ID	XMM	YMM	Nb cycles	FLOP/cycle	IPC	Bytes loaded / cycle	Bytes stored / cycle
5	1a	12	0	4	2	3	5	1
6	1b	0	12	10	12.8	2.1	32	6.4
7	2a	16	0	8	2	3.125	3	1
8	2b	0	10	10	12.8	2.1	32	6.4
9								
10	Nb used reg.		Performance in L1					
11	ID	XMM	YMM	Nb cycles	FLOP/cycle	IPC	Bytes loaded / cycle	Bytes stored / cycle
12	1	0	15	8	16	3.125	24	8
13	2	0	15	8	16	3.125	24	8

MAQAO static analysis before (top) and after (bottom) optimization

- Dealing with the two hottest loops in the application
 - Dense x sparse matrix multiply
- FLOP/cycle not optimal:
 - 12.8 but should be 16
 - AVX, 32 bits elements, perfect ADD/MUL balance
- Replacing LDC with its value "hard coded" allows the compiler to factor for the two matrices C1 and C2

Tools to Explain

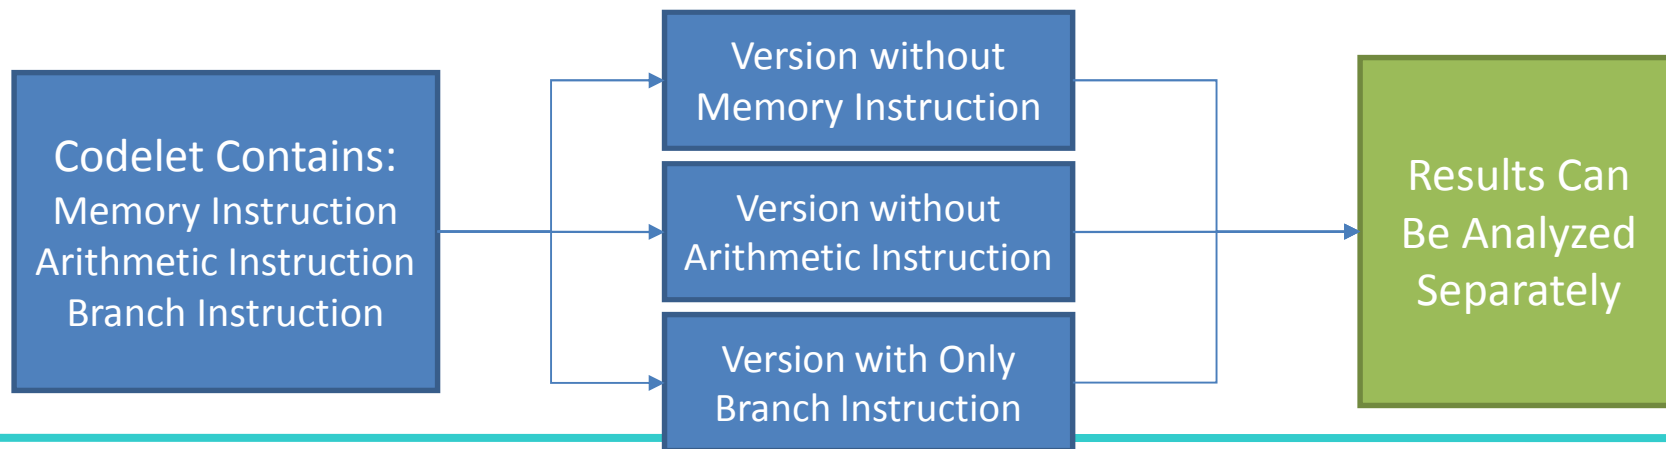
- MAQAO and similar tools provide information
 - Profiler detects hot spot
 - MAQAO goes beyond and evaluates the gap
 - Current and optimal static performance
 - It remains the discrepancy is difficult to understand
- DECAN is an exploratory tool

DECAN

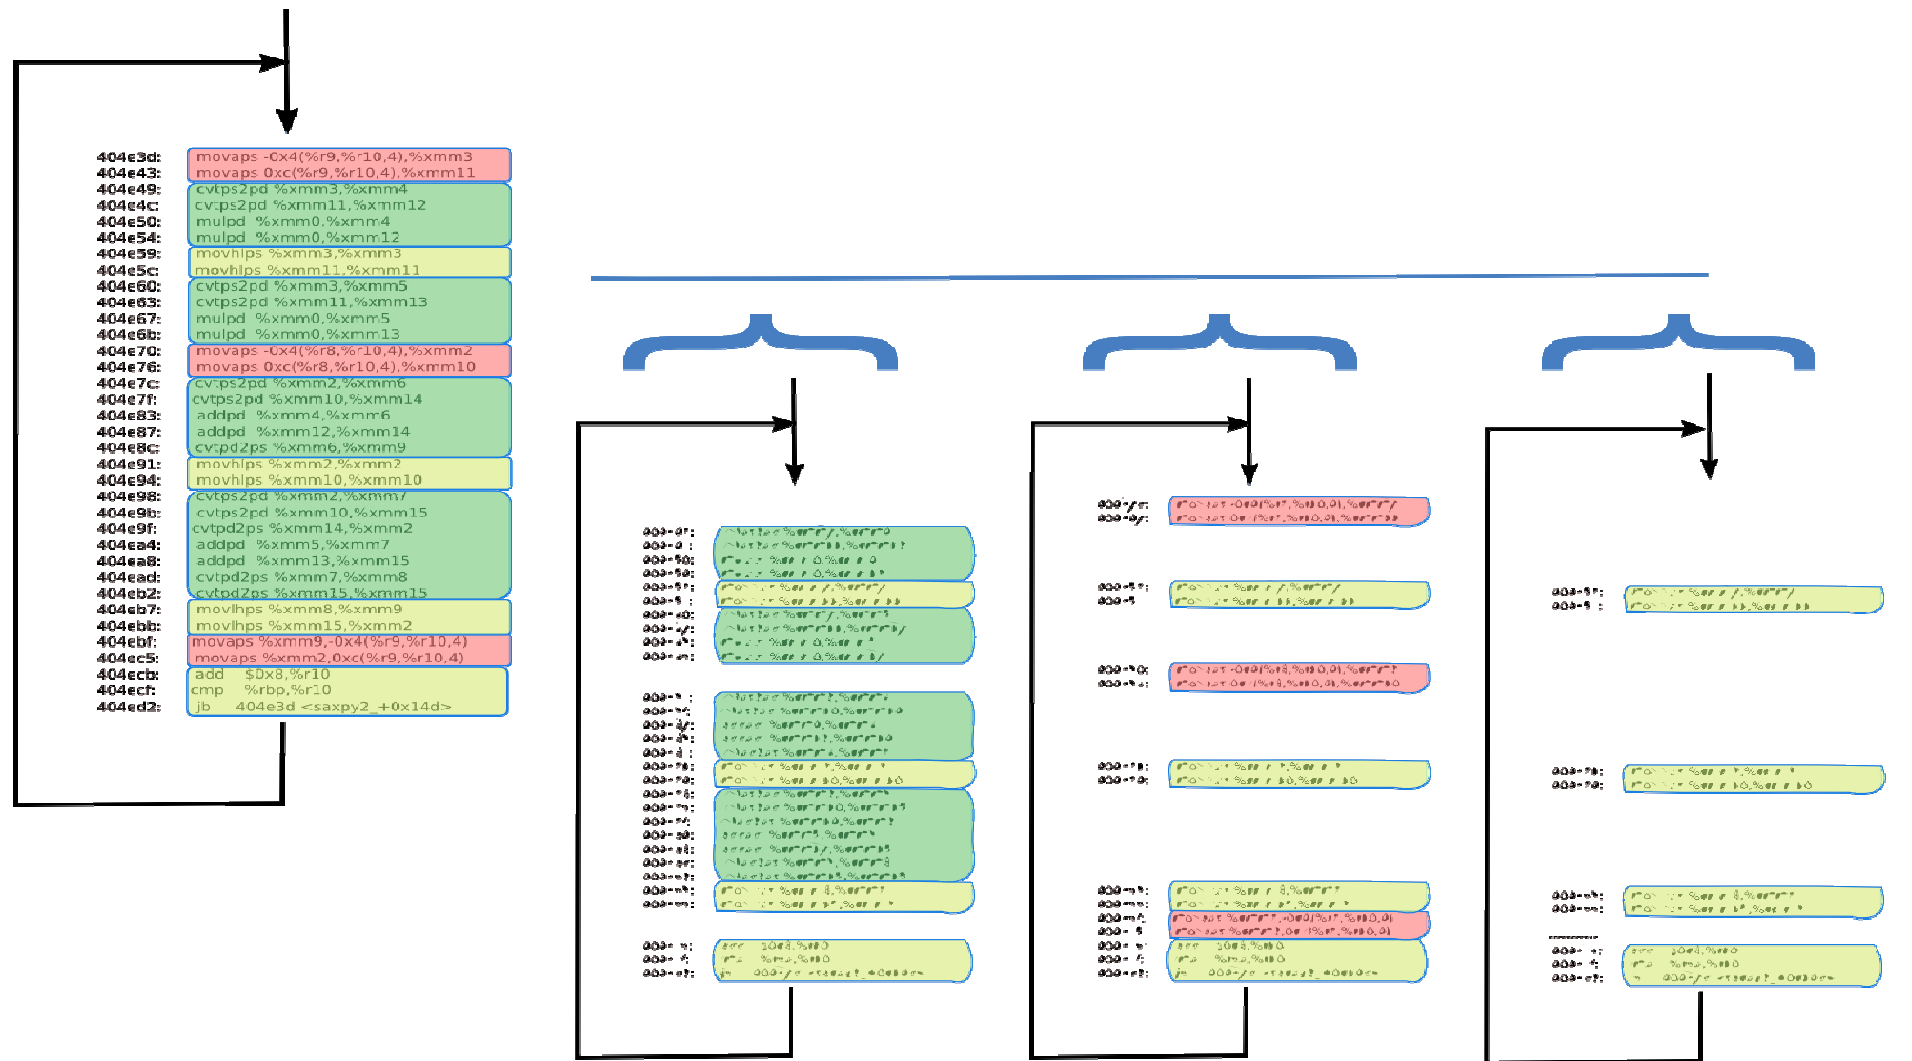
- DECAN's concept is simple
 - Measure the original binary
 - Patch and replace the selected instructions group in the original binary
 - New binary is generated for each patch
 - Measure new binaries
 - Measurements are represented in a CSV file
 - Analyze and compare

DECAN

- Codelet Decomposition
 - MISTREAM
 - All vector arithmetic instructions are deleted
 - FPSTREAM
 - All vector loads and store instructions are deleted
 - NOFPNOMISTREAM
 - All vector arithmetic, load, and store instructions are deleted



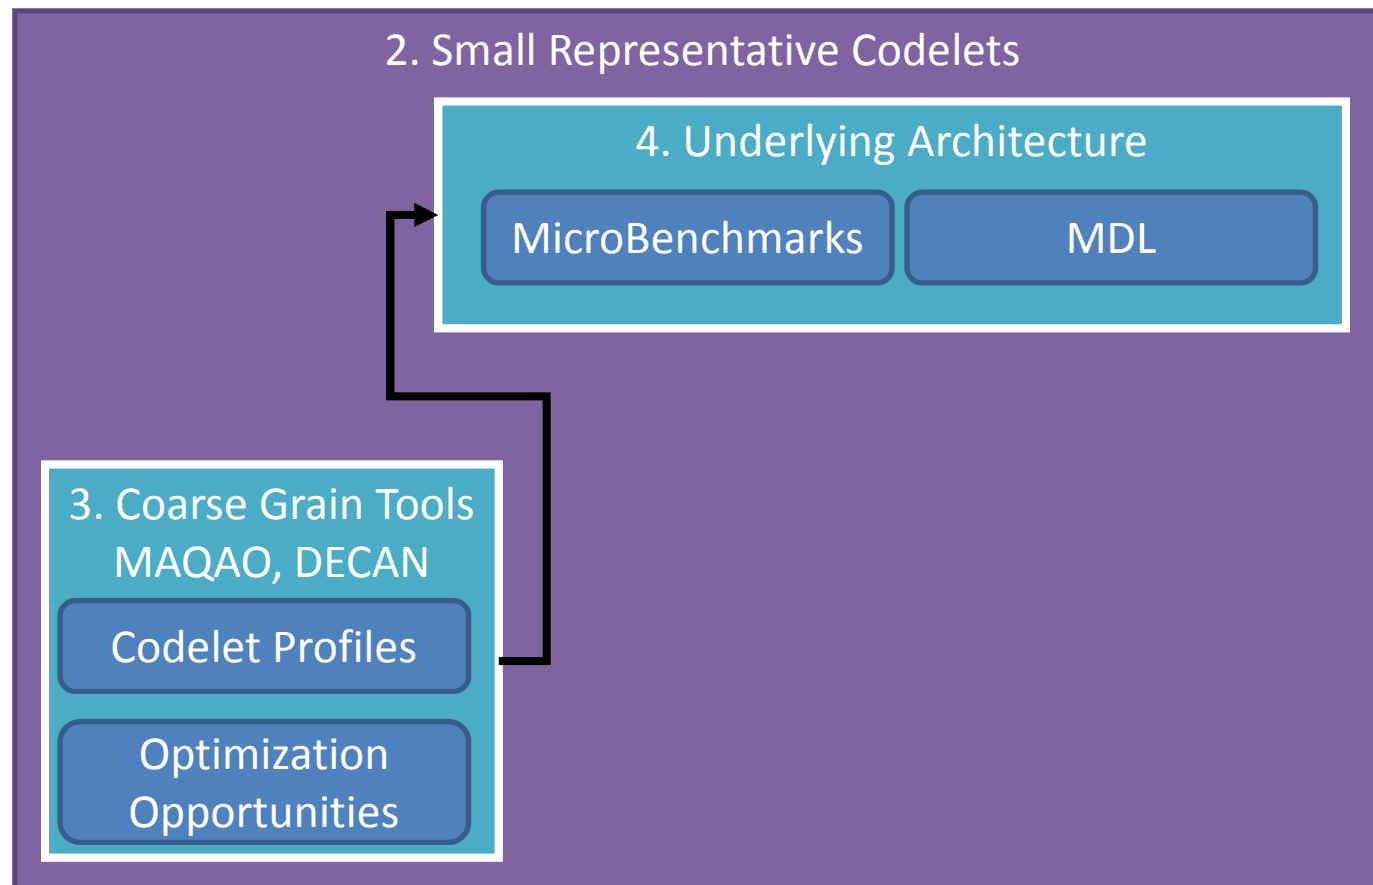
DECAN



MAQAO + DECAN Provides

- Speed-up by a factor of 4
 - Up to 37% of the peak performance on Sandy Bridge
- Vectorization ratio crucial on Sandy Bridge
- Value profiling

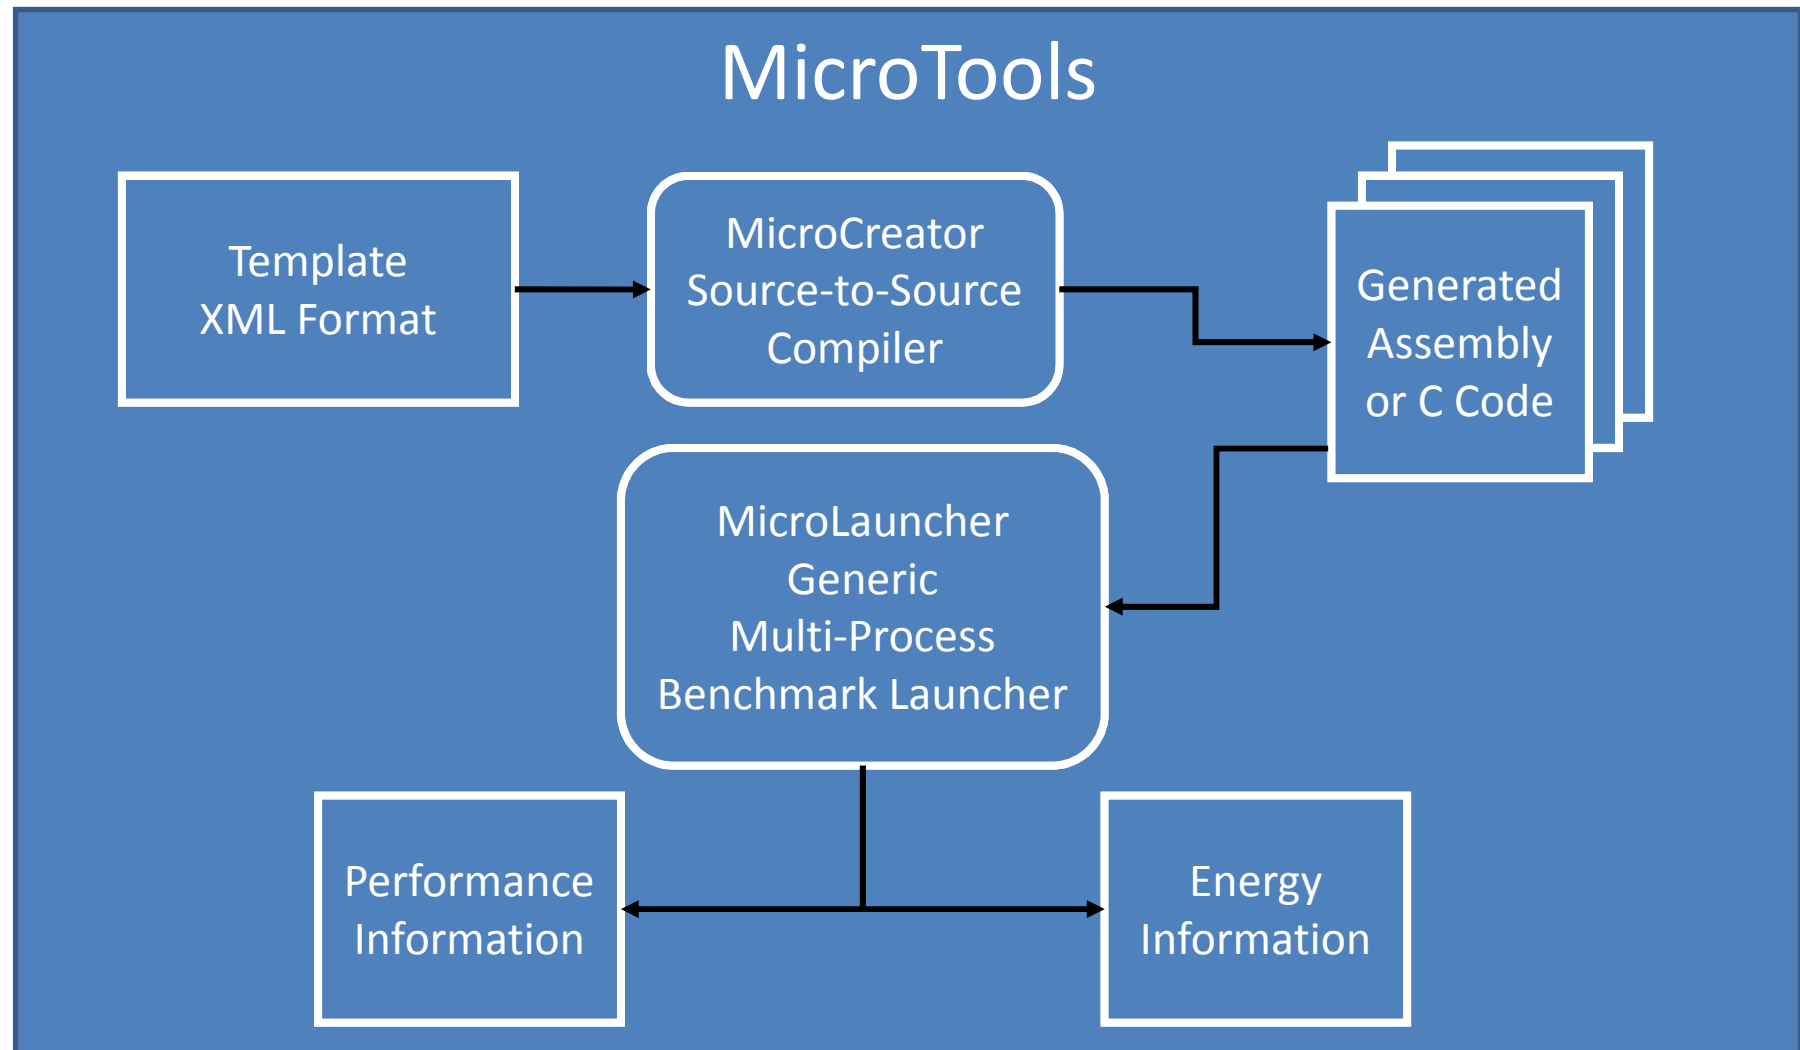
Outline



Underlying Architecture

- Understanding the target architecture
 - Gives insight on potential bottlenecks
 - Provides solutions to optimize a code
- How is it done?
 - Emulators or simulators are slow if even available
 - Microbenchmarking considers the hardware as a black box

MICROBENCHMARKS



MicroTools Usage

MicroTools enables an exhaustive exploration of architecture performance.

1) Two real case-studies

- NOPS impact on dispatcher
- Memory + arithmetic interaction

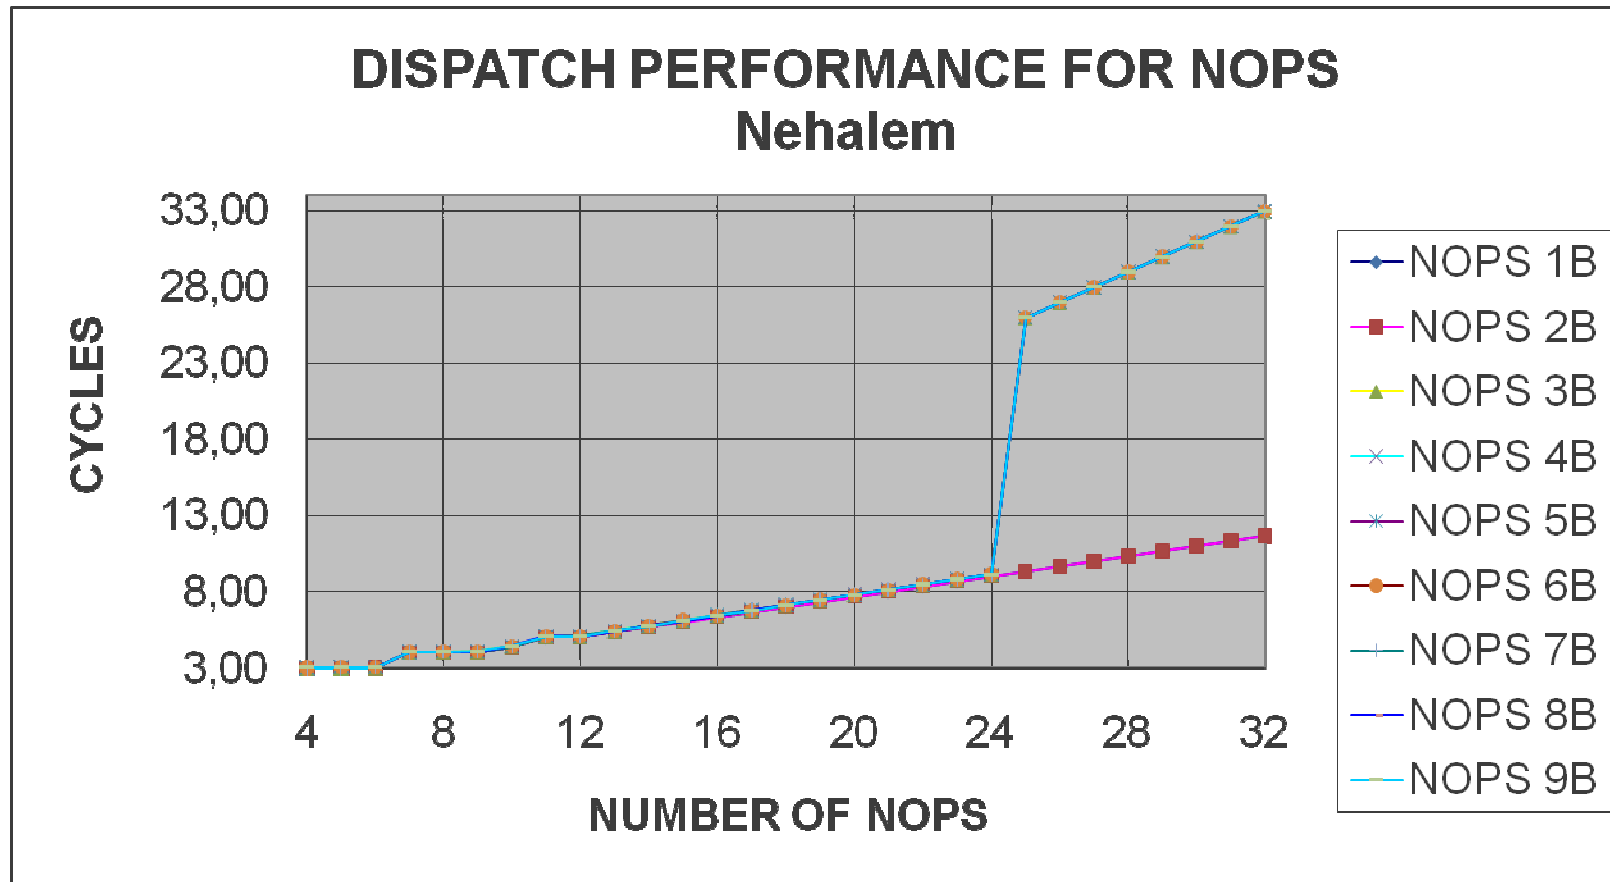
2) How to deal with all the data collected?

- Automated reports analysis

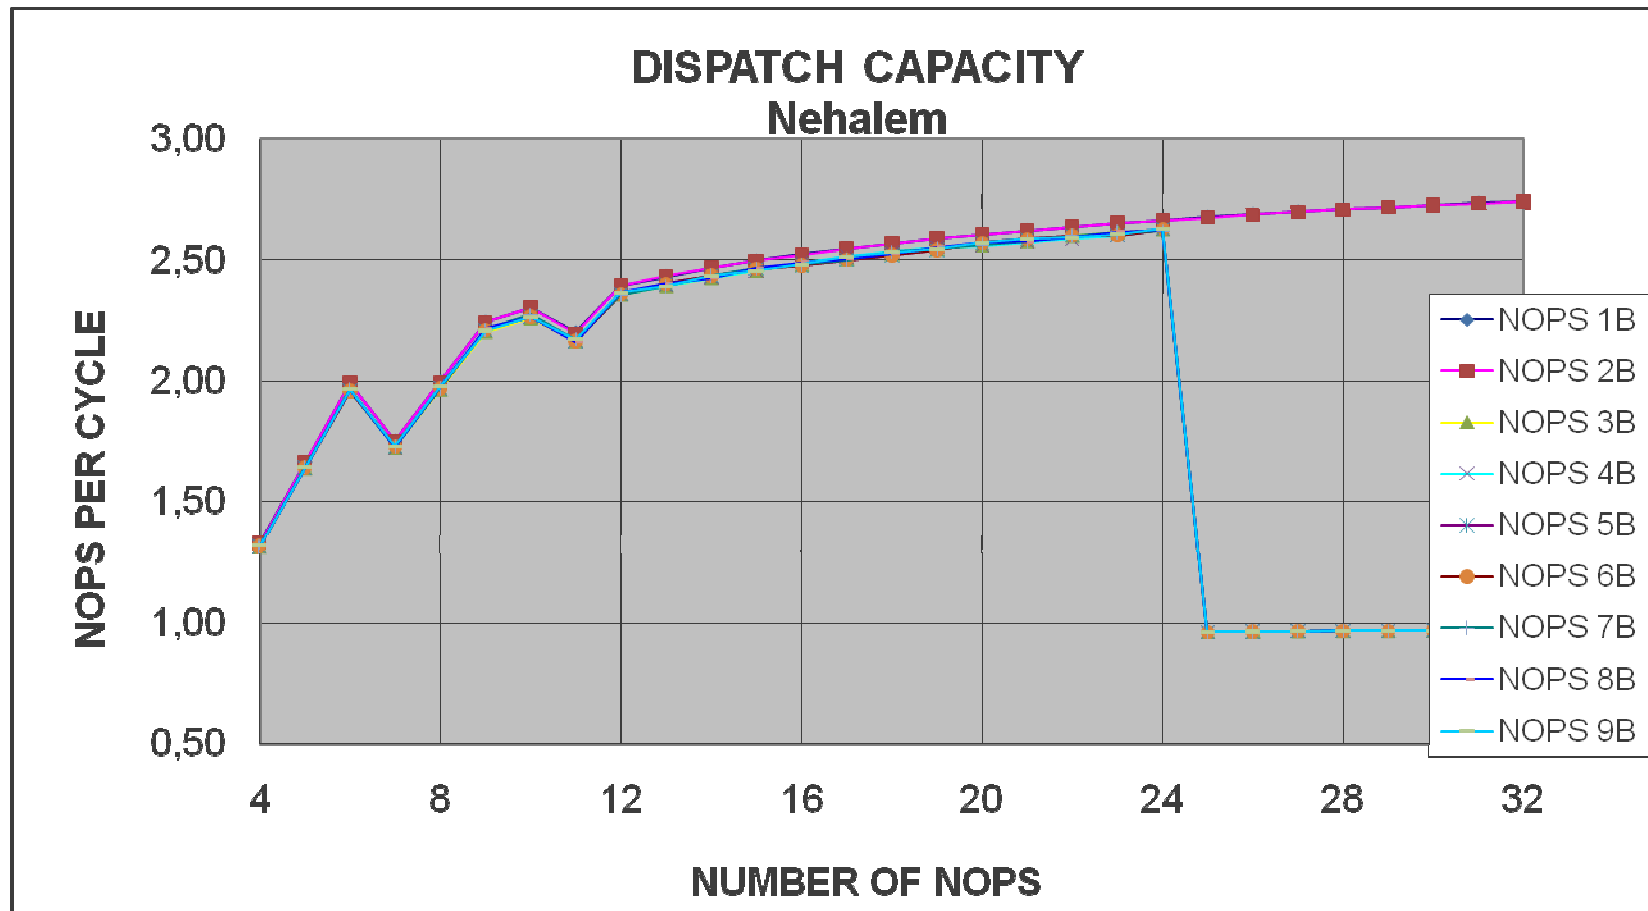
NOPS Experiment

- Goal: Evaluate the dispatch unit
- Loop body parameters
 - The NOP instruction size:
 - Varies from 1 to 9 bytes
 - The number of NOP instructions in the loop body:
 - Varies from 4 to 32
- Each loop body tested consists of the same NOP instructions repeated from 4 to 32 times

NOPS Experiment



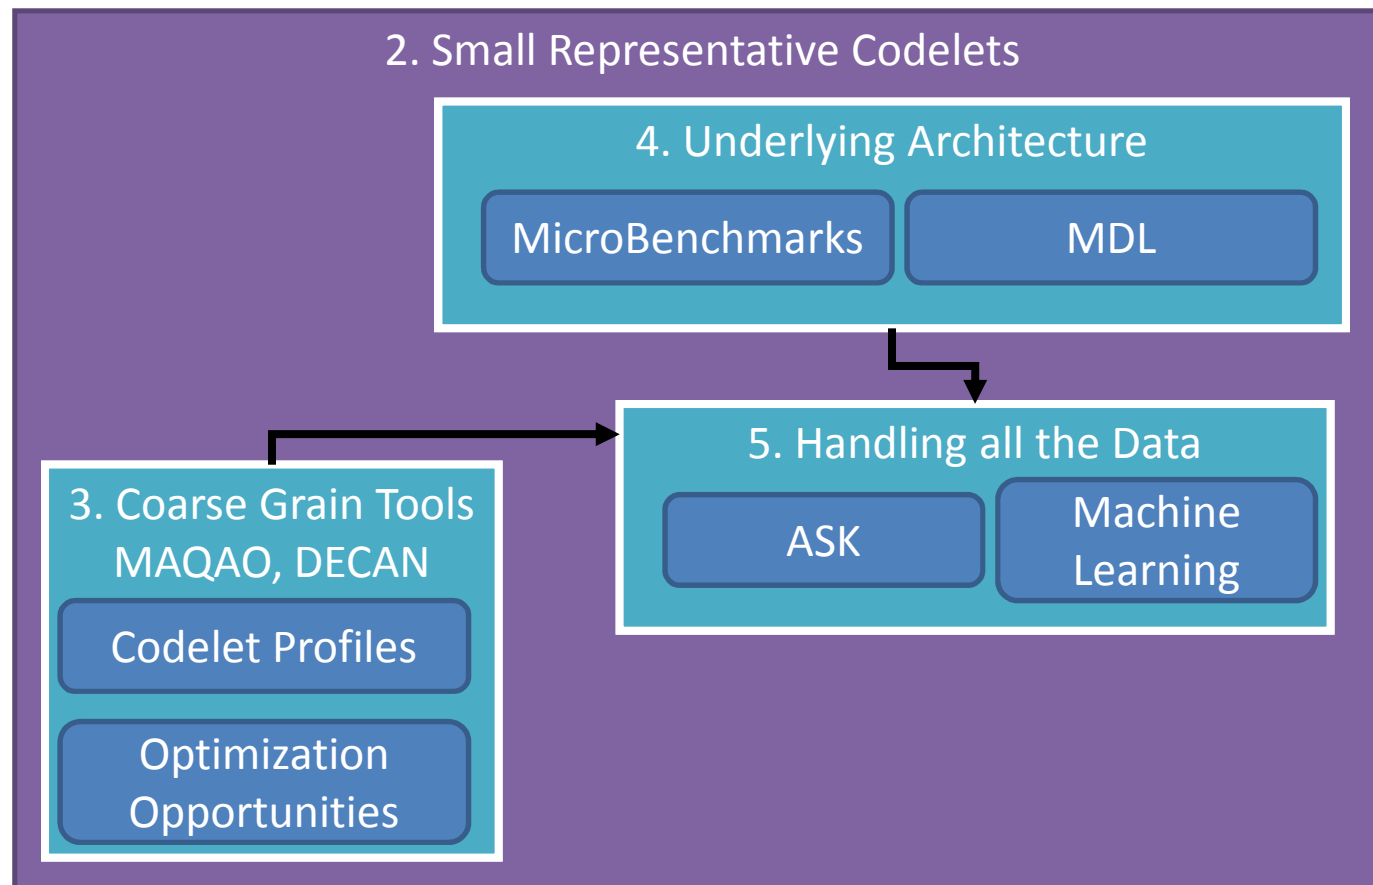
NOPS Experiment



NOPS Experiment

- Dispatcher capacity is around 3 nops/cycle
- For 1-byte and 2-byte nops
 - Dispatch behavior is linear
- For large size nops
 - Dispatch rate falls down to 0.96 instructions per cycle

Outline

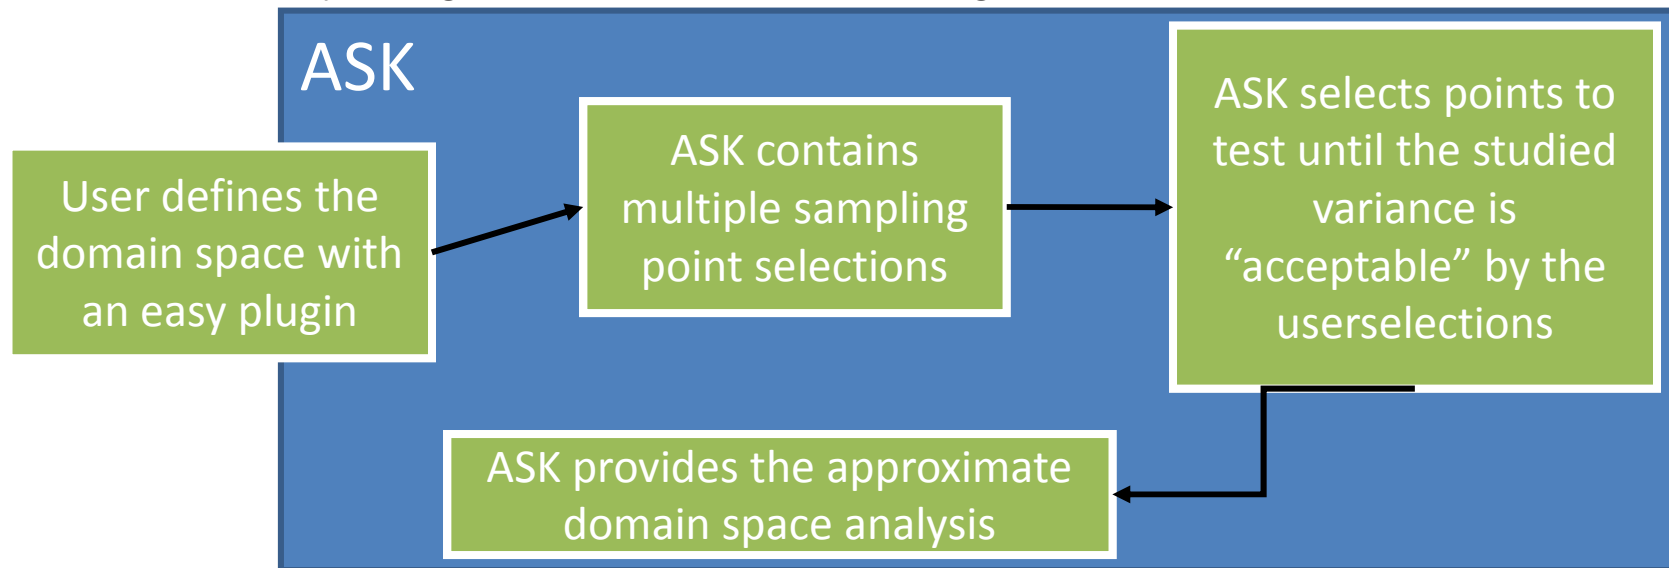


HANDLING THE DATA

- MicroCreator philosophy:
 - Exhaustively search around a given program specification
 - Sometimes, we need to reduce the space (ASK)
- A lot of data produced
 - We need automatic data handling tools
 - Validate the stability of the results
 - Identify unexpected situations to help the engineer.

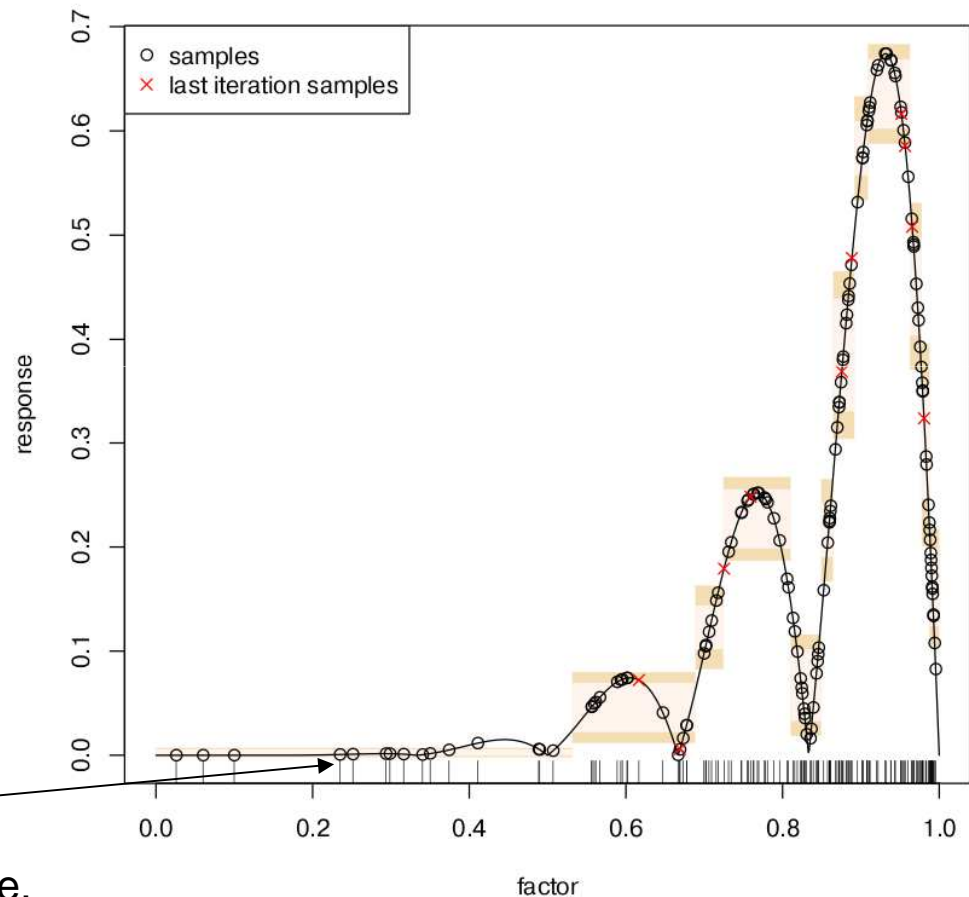
ASK (Adaptive Sampling Kit)

- ASK is a toolkit providing state of the art sampling strategies
- Modular and Extensible Pipeline:
 - Combine different static and dynamic strategies
 - Easy to add new custom sampling strategy
 - Easy integration with benchmarking tools



ASK: An example

- Search the domain space
- Find high variance regions
- Draw new points from high-variance regions



“Flat” regions are less interesting to explore.

Order Influence Report

- Decomposes results per number of stores and loads
- Quickly identify configurations where performance depends on the order of instructions

MOVAPS Sandybridge L3														
#L \ #S		1		2		3		4		5		6		7
1	7.29	7.6	17.85	19.15	24.29	33.07	38.22	38.69	44.19	52.57	57.22	58.51	63.41	71.37
2	9.8	10.6	20.19	20.86	26.16	34.89	31.21	41.95	45.96	54.65	51.57	59.83		
3	12.88	13.18	21.46	22.81	27.92	37.55	33.22	45.15	36.77	56.53				
4	15.67	15.74	23.08	24.35	28.43	39.41	33.81	45.77						
5	18.56	19.01	24.89	25.7	30.55	41.53								
6	21.42	21.54	26.88	28.25										
7	24.55	25.67												

Outline

2. Small Representative Codelets

5. Handling all the Data

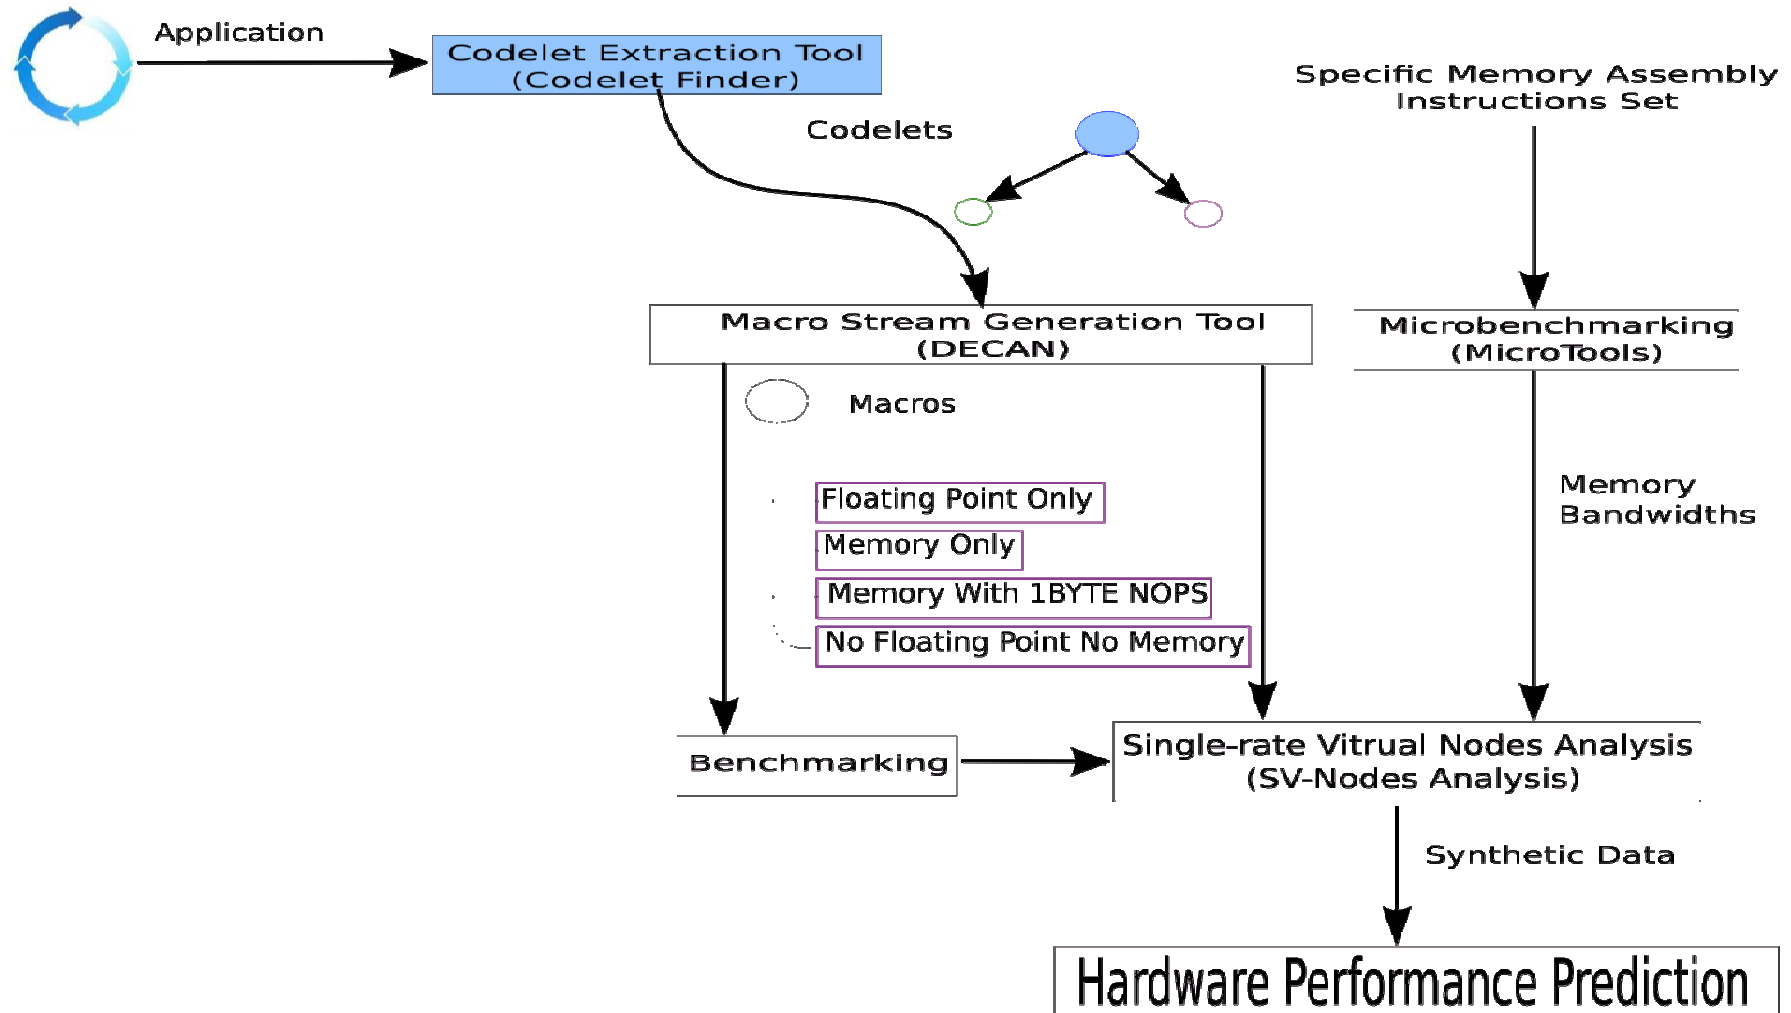
ASK

Machine
Learning



6. Capacity and Prediction Models

Capacity Model

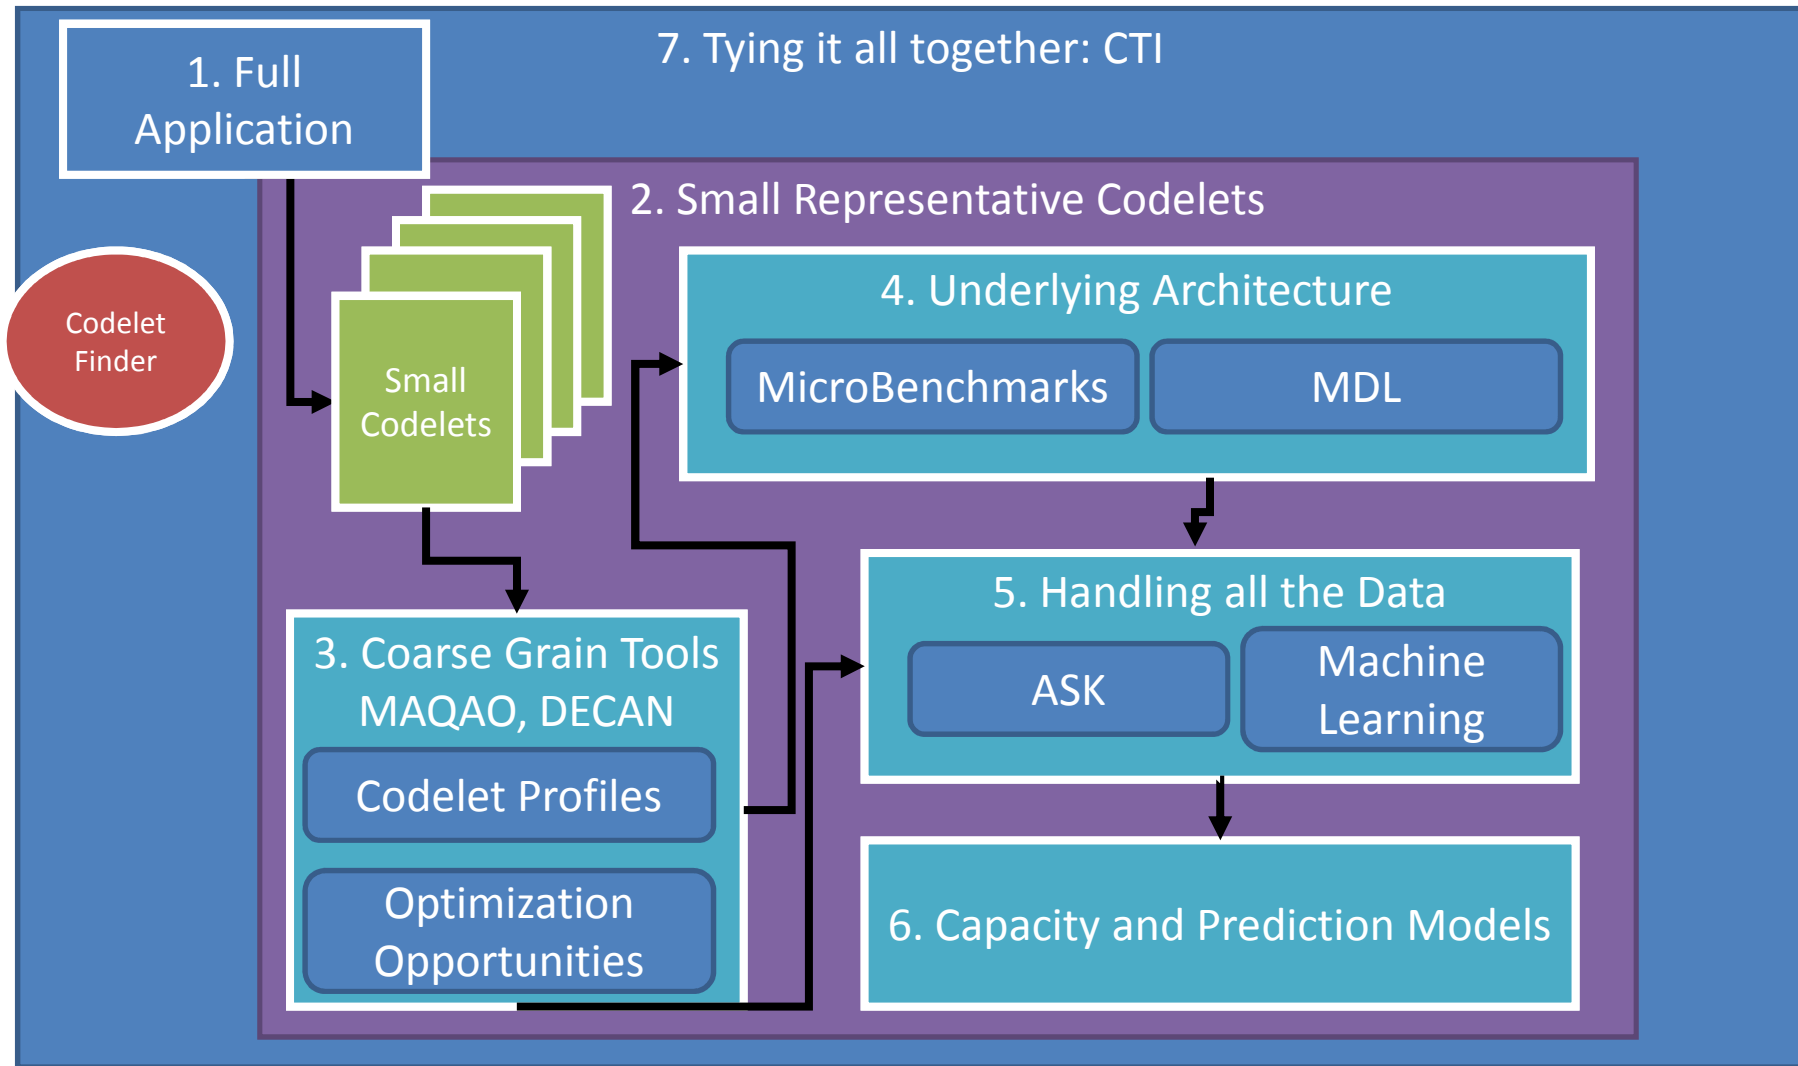


Capacity Model

- Capacity Model and MDL
 - Provide prediction and modelization of program performance
- However, alone the tools are less valuable
 - A need to centralize the data and analysis:

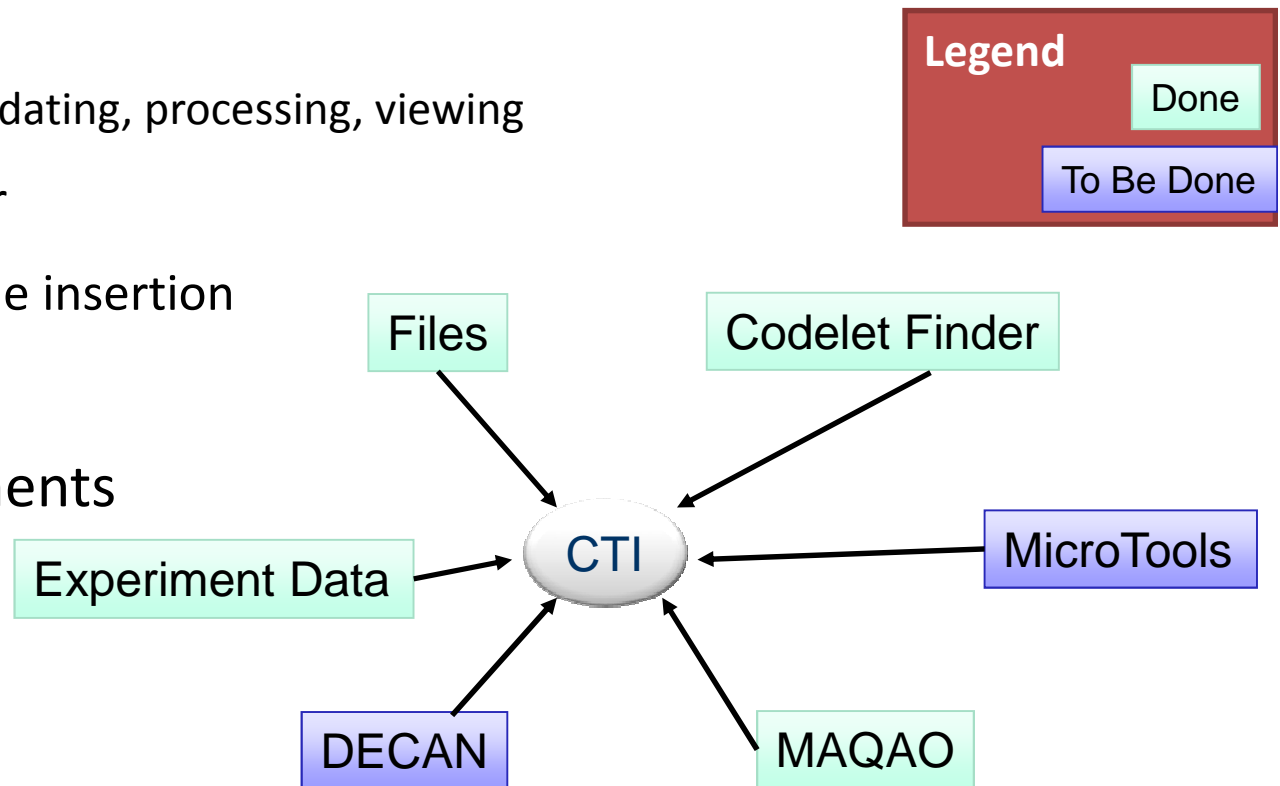
CTI : Codelet Tuning Infrastructure

Outline

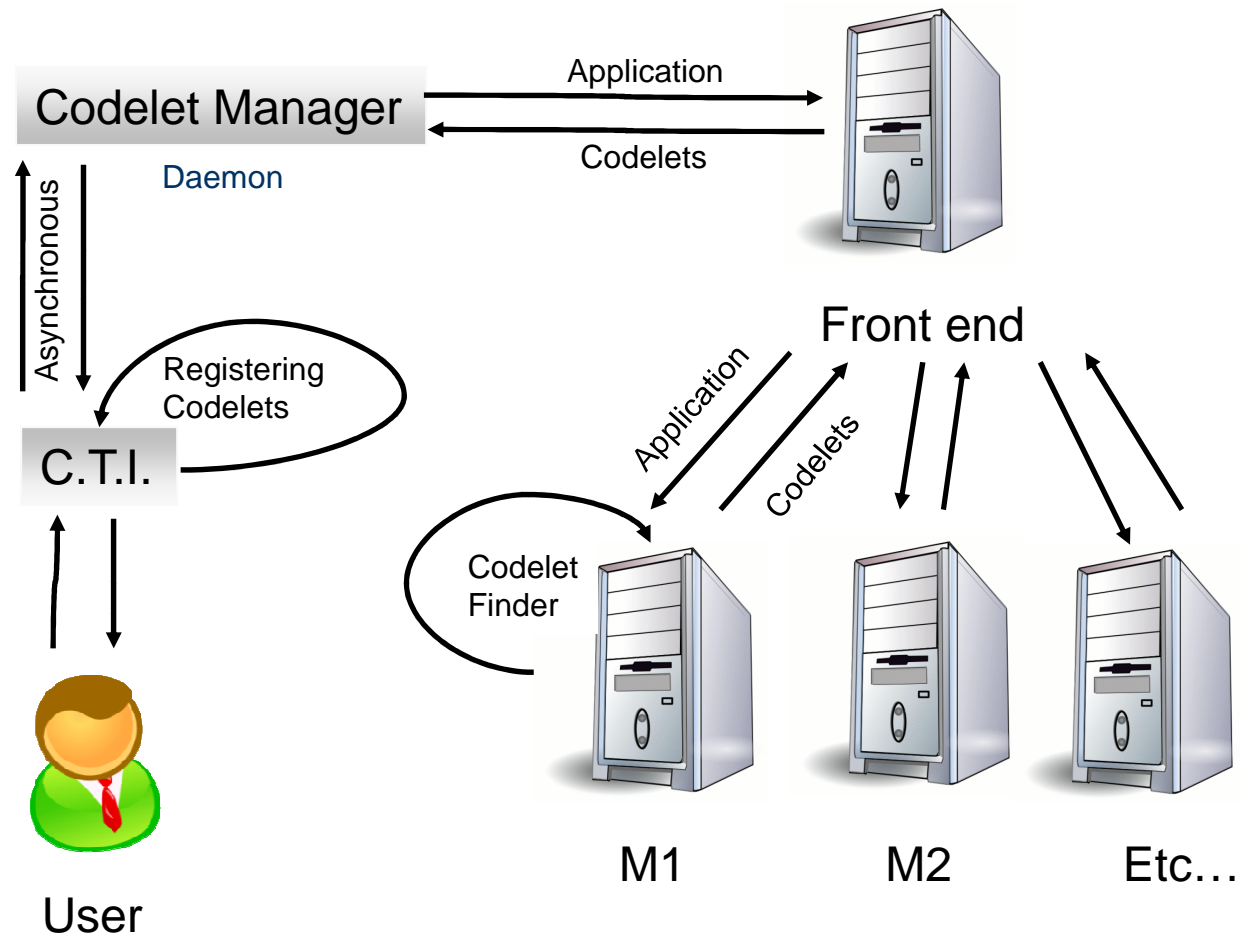


Codelet Tuning Infrastructure (CTI)

- A single place to store a huge amount of data
 - File manager
 - File sharing, updating, processing, viewing
 - Codelet manager
 - CSV automatic file insertion
 - Query the data
- Automate experiments
- Tools integrator



CTI Codelet Manager



CTI Codelet Manager

- Launching the « create » process

Codelet Tuning
Services

Exascale 
computing research

Welcome ftalbart | Log out

[Hot links](#) | [Browse Repositories](#) |

[Full view](#) | [Create new codelets](#)

test

Repository = /home/users/ftalbart/test/.ctr/data/de1e6e3b-88ca-4d8e-8a0d-6864546b8ac2

UID = de1e6e3b-88ca-4d8e-8a0d-6864546b8ac2

Plugin command = *init*

name	directory	codelets
test	/home/users/ftalbart/test	

Create new codelets

Platform:



```
johndoe@toto:~$cti application create_codelets test borodine
The process is running !
DATA_CTI_UID=53e7cbd5-f429-4d40-aff6-c767ca3abc7a
```

CTI Codelet Manager

- Checking the codelets list

Codelet Tuning
Services

Exascale 
computing research

Welcome ftalbart | Log out

[Hot links](#) | [Browse Repositories](#) |

Full view

c01a8629-c459-4fed-8ce0-b0cd41b1b44b

Repository = /home/users/ftalbart/test/.ctr/data/c01a8629-c459-4fed-8ce0-b0cd41b1b44b

UID = c01a8629-c459-4fed-8ce0-b0cd41b1b44b

Plugin command = *init*

codelets	application	platform
(test_Codelet1_6_10_2011_14:39:42)	test	chopin
(test_Codelet2_6_10_2011_14:39:42)		

Franck Talbart, Yuriy Kashnikov, Pablo Oliveira, William J. ... gori Fursin
Contributors: David Wong, David Kuck

(C)opyright, Exascale Computing Research Center(Intel/CEA/GENC...
All rights reserved
Feedback Documentation
CTI version 0.9.71af1a1



johndoe@toto:~\$ cti view data <UID|Alias>

CTI Codelet Manager

- Checking the codelet files

Services

Welcome ftalbart | Log out

Hot links | Browse Repositories | Search

Full view | Download

864bf952-c711-406b-8f7e-41fe87e1ec5d

Repository = /home/users/ftalbart/test/.ctr/data/864bf952-c711-406b-8f7e-41fe87e1ec5d

UID = 864bf952-c711-406b-8f7e-41fe87e1ec5d

Plugin command = *init*

data	files
View complex data	files/codelet_mlo4qvzd_1.mb.gz
	files/codelet_mlo4qvzd_1.mb.gz.d/TIMESTAMP
	files/codelet_mlo4qvzd_1.mb.gz.d/input.mb.gz
	files/codelet_mlo4qvzd_1.mb.gz.d/Makefile
	files/codelet_mlo4qvzd_1.mb.gz.d/text.c
	files/codelet_mlo4qvzd_1.mb.gz.d/cxx_wrapper.cc
	files/codelet_mlo4qvzd_1.mb.gz.d/cxx_wrapper.h
	files/codelet_mlo4qvzd_1.mb.gz.d/main.c
	files/codelet_mlo4qvzd_1.mb.gz.d/asm_wrapper.s
	files/codelet_mlo4qvzd_1.mb.gz.d/hmppcf/MicroBencherFileParser.h
	files/codelet_mlo4qvzd_1.mb.gz.d/hmppcf/MicroBencherFileCommon.def
	files/codelet_mlo4qvzd_1.mb.gz.d/hmppcf/MicroBencherFileUnparser.h
	files/codelet_mlo4qvzd_1.mb.gz.d/hmppcf/MicroBencherFileCommon.h
	files/codelet_mlo4qvzd_1.mb.gz.d/hmppcf/src/MicroBencherFileCommon.cc
	files/codelet_mlo4qvzd_1.mb.gz.d/hmppcf/src/MicroBencherFileUnparser.cc
	files/codelet_mlo4qvzd_1.mb.gz.d/hmppcf/src/MicroBencherFileParser.cc

Content of the file

Go back to the entry

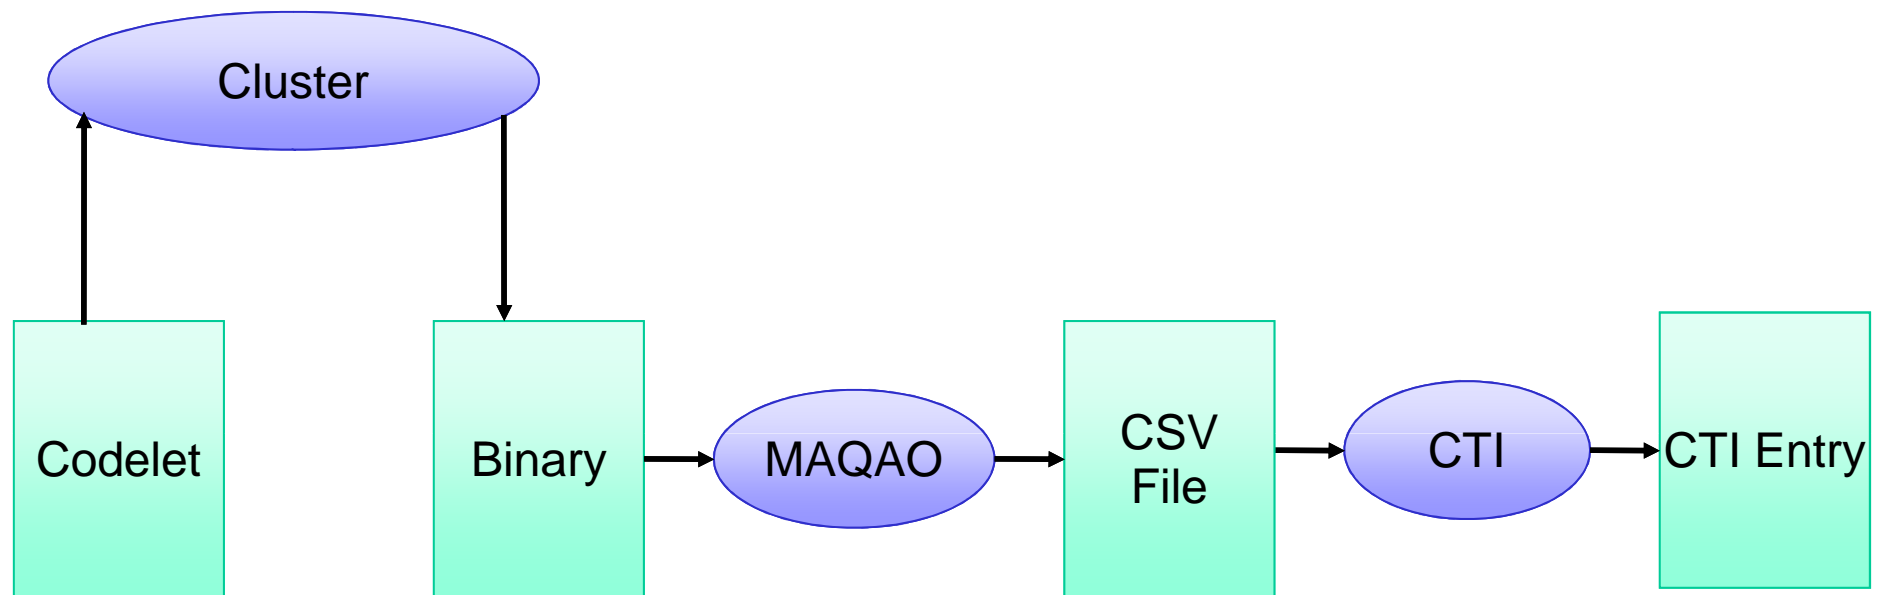
```
void codelet_mlo4qvzd(int *tab, int * __hmppcf_addr__i)
{
    int i = * __hmppcf_addr__i;
    for (i = 0 ; i < 10 ; i++)
        tab[i] = i;
    * __hmppcf_addr__i = i;
}
```

johnndoe@toto:~\$ cti import_files get <UID|Alias>

Franck Talbart, Yuriy Kashnikov, Pablo Oliveira, William Jalby, Grigori Fursin
Contributors: David Wong, David Kuck

(C)opyright, Exascale Computing Research Center(Intel/CEA/GENCI/UVSQ), 2010-2011
All rights reserved
Feedback Documentation

CTI MAQAO Integration



1) Codelet to Binary

The codelet is sent to the cluster for compilation

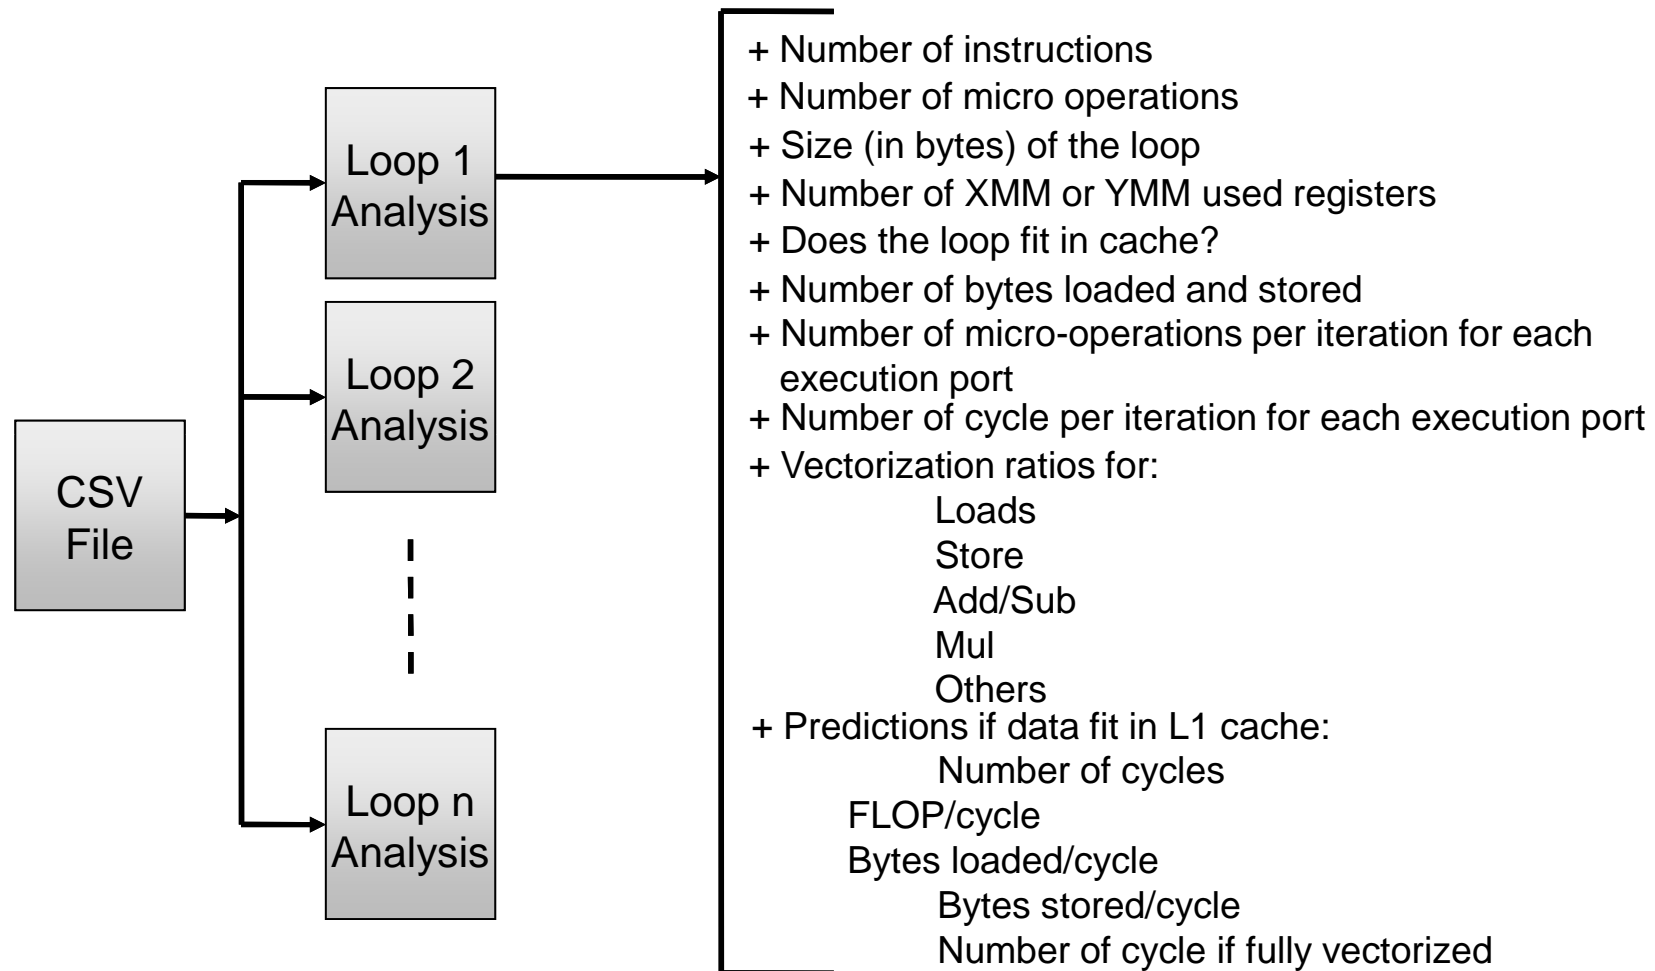
2) Binary to CSV

Binary is analyzed using MAQAO and a CSV file with results is produced

3) CSV to CTI entry

CSV file is loaded into CTI repository and saved in an entry

CTI MAQAO Integration



Overall Conclusions

- Studying a large application is difficult
 - Dividing the application into codelets
 - Using tools such as MAQAO and DECAN help understand the codelet's performance and behavior
 - Understanding the underlying architecture with MicroTools and the MDL help detect hardware bottlenecks
 - Analyzing all the data is only possible with automatic tools and infrastructures such as CTI

Acknowledgements

- Pablo Oliviera, UVSQ
- José Noudohouenou, UVSQ
- Franck Talbart, UVSQ
- William Jalby, CTO, UVSQ
- Jean-Thomas Acquaviva, Lead of the Performance Tools
- Bettina Krammer, Head of Software Tools, UVSQ
- Marie-Christine Sawley, Lab Director, Intel

ECR Contacts

- Address

UVSQ, 45 Av. des Etats-Unis, Buffon building, 5th floor
78 000 Versailles, France

- Web site: www.exascale-computing.eu

- Team

William Jalby, CT , william.jalby@uvsq.fr

Jean Christophe Beyler, Lead of Application Characterization,
jean.christophe.beyler@intel.com

Marie-Christine Sawley, Co-design, marie-christine.sawley@intel.com

Bettina Krammer, Tools, bettina.krammer@uvsq.fr

- Collaboration partners

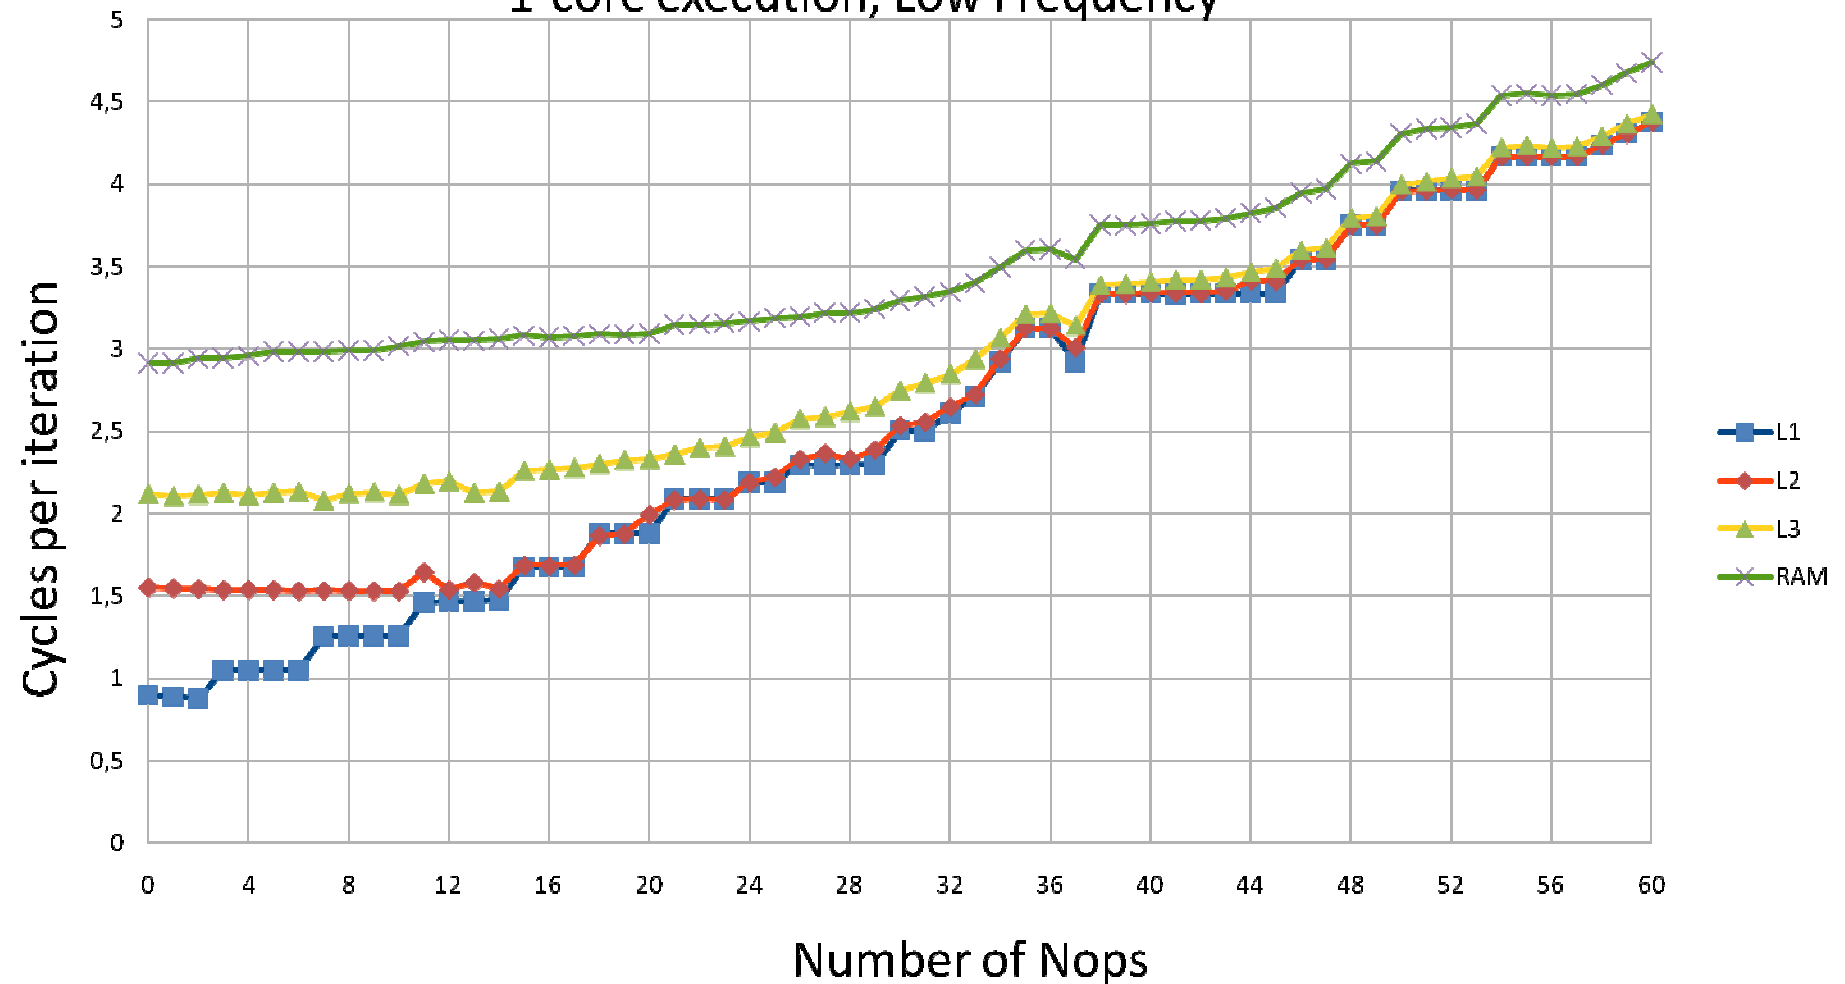


Thank you

DECAN + MICROTOOLS

saxpy2MIS

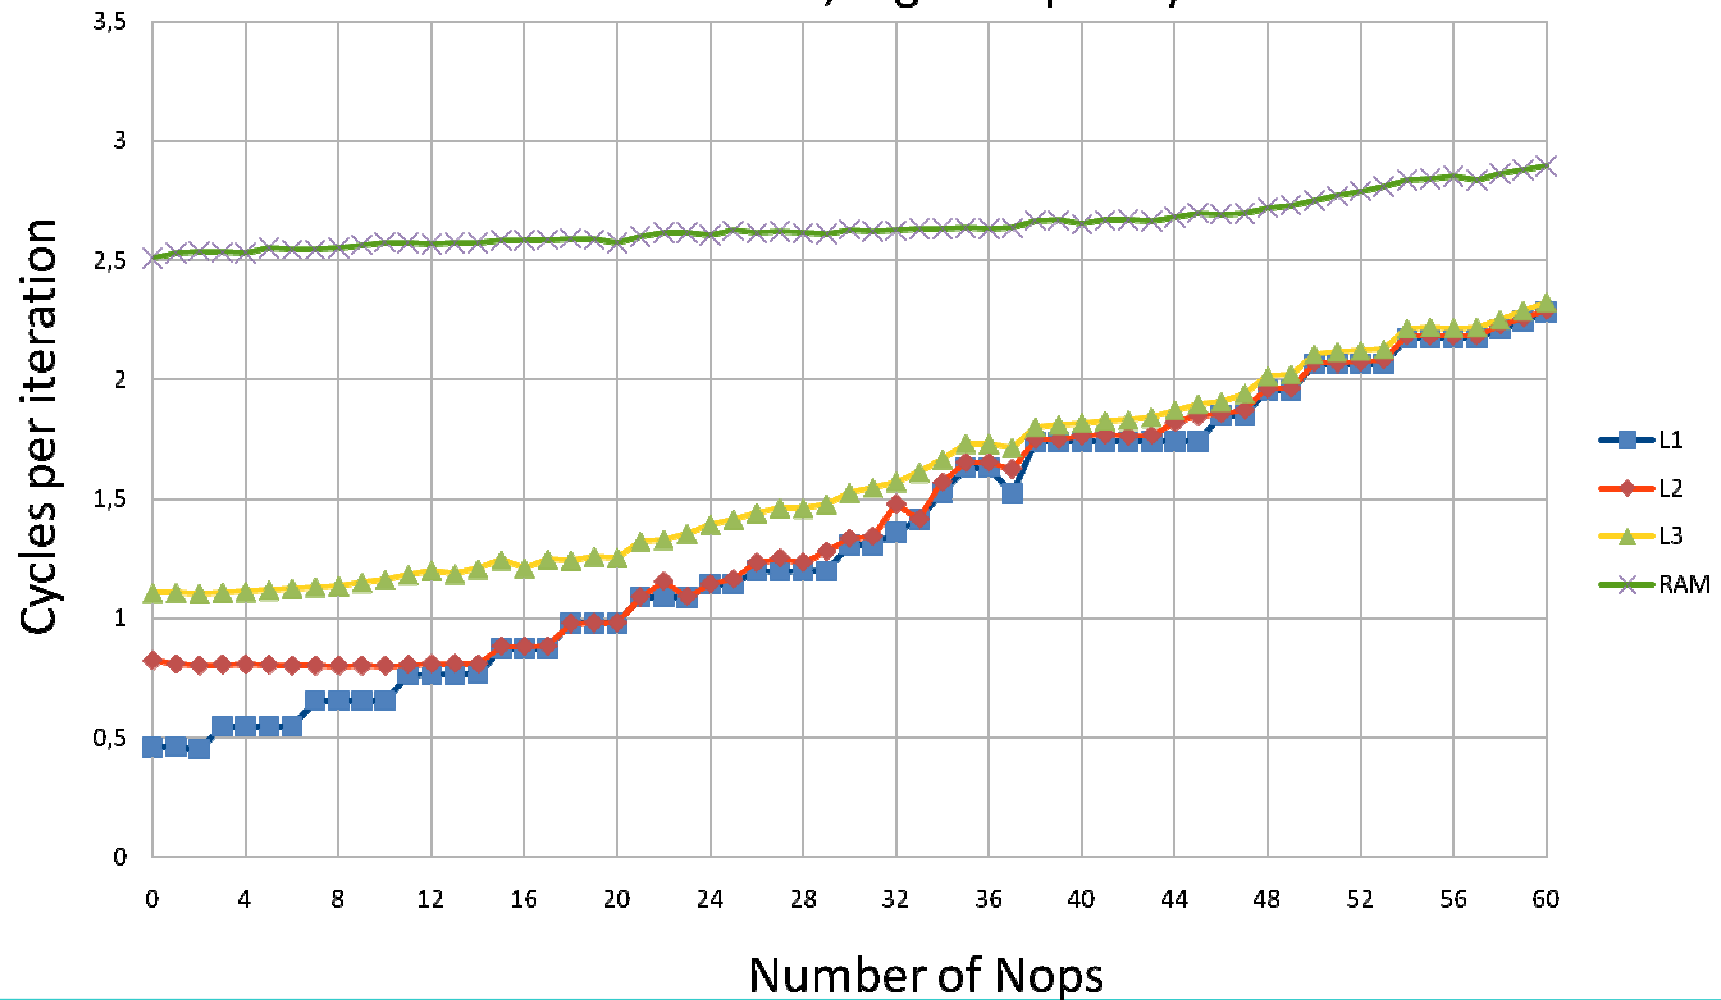
1-core execution, Low Frequency



DECAN + MICROTOOLS

saxpy2MIS

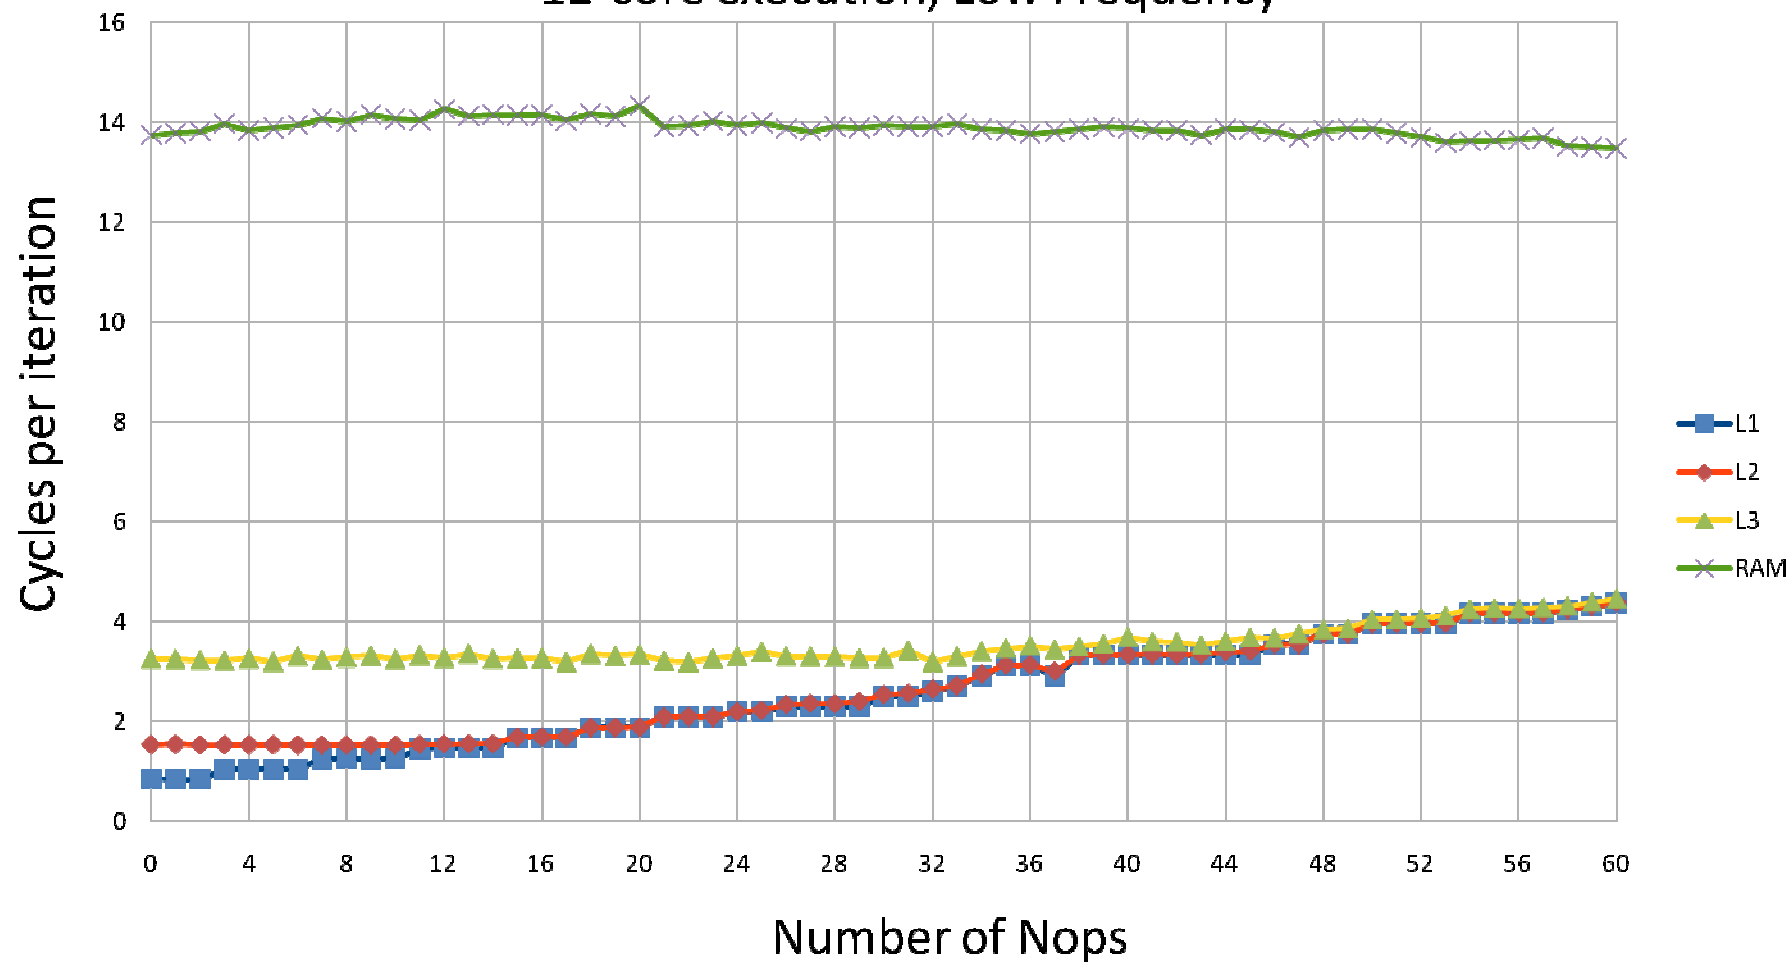
1-core execution, High Frequency



DECAN + MICROTOOLS

saxpy2MIS

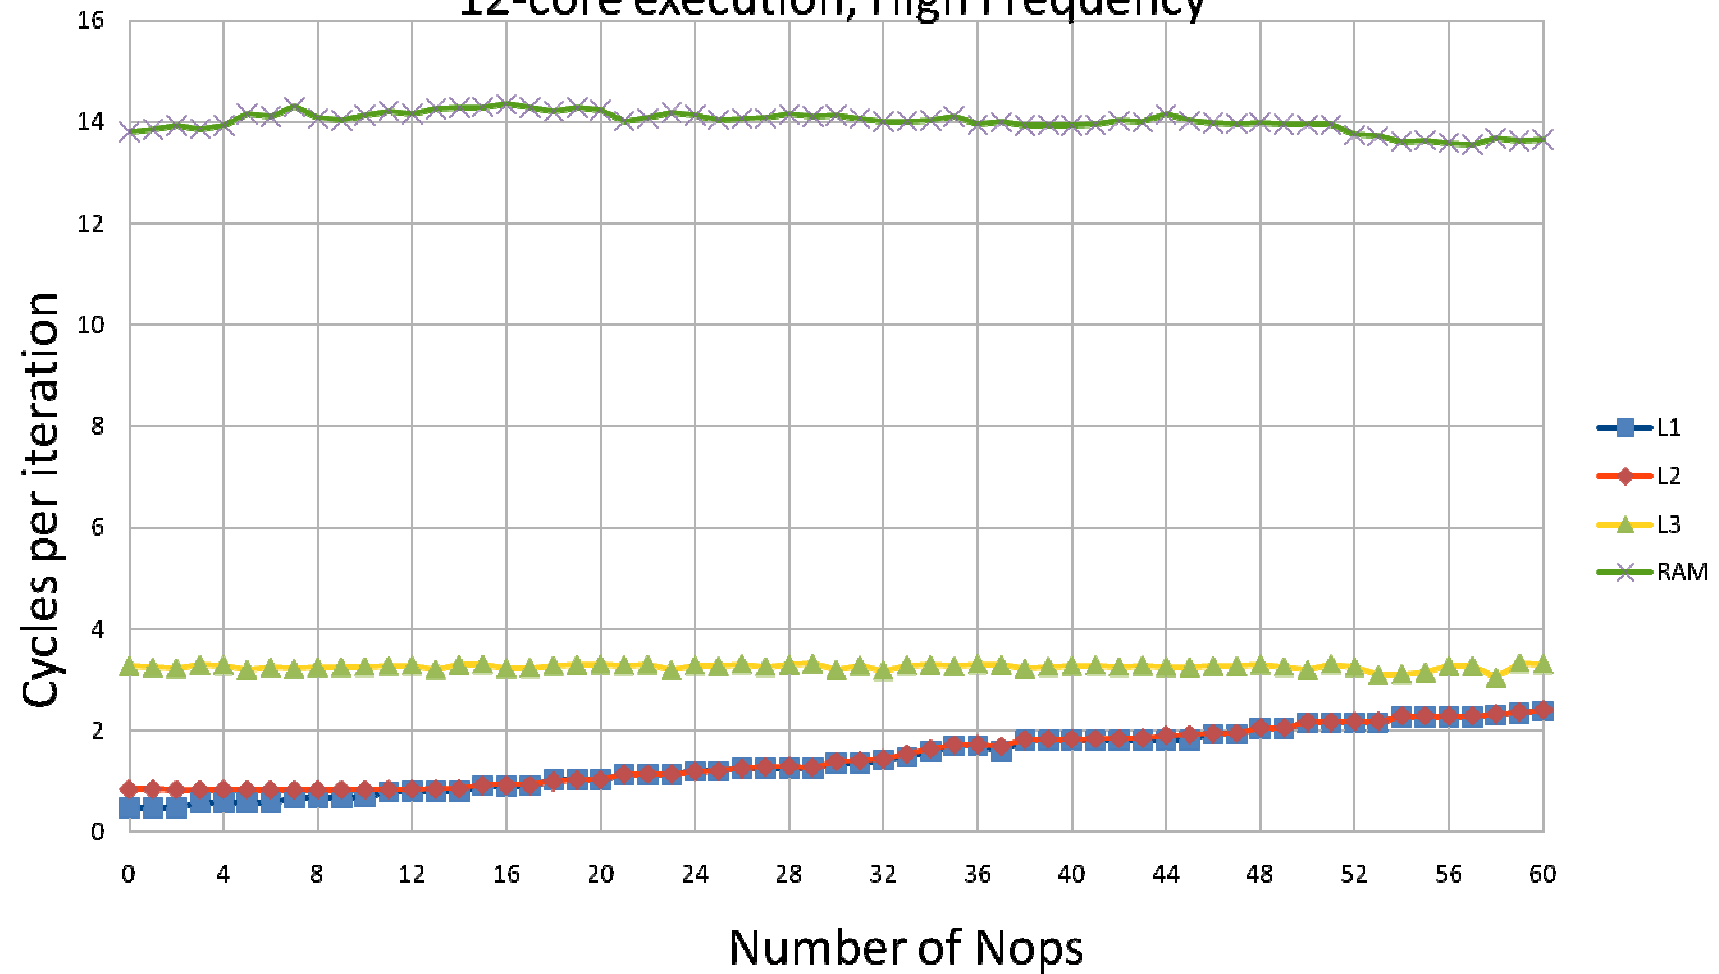
12-core execution, Low Frequency



DECAN + MICROTOOLS

saxpy2MIS

12-core execution, High Frequency



Handling Stores

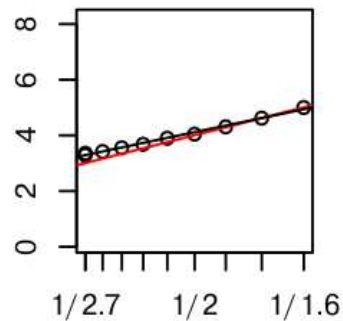
- L3 Nehalem
 - Pure load patterns scale better with frequency

Handling Stores

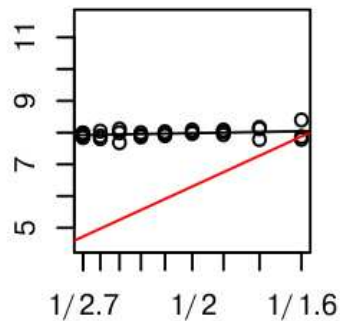
- L3 Nehalem
 - Pure load patterns scale better with frequency

MOVSD

2 LOADS



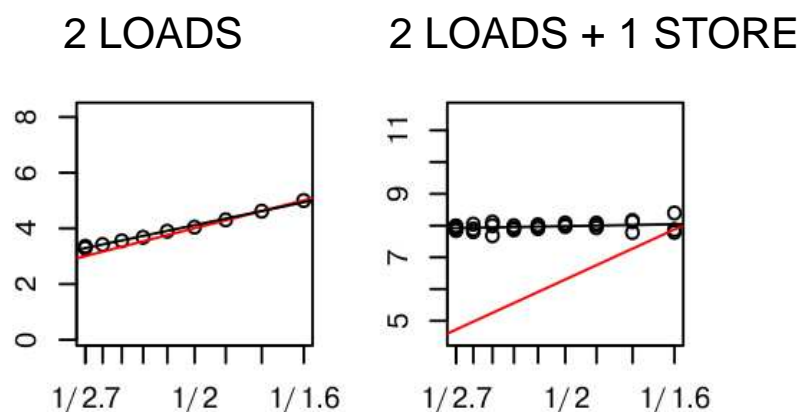
2 LOADS + 1 STORE



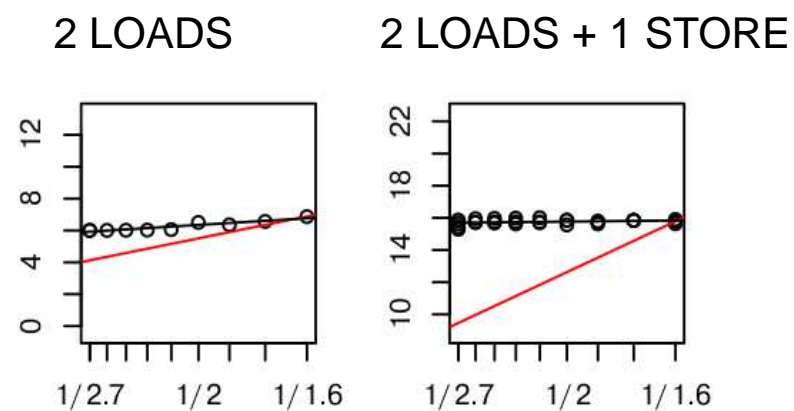
Handling Stores

- L3 Nehalem
 - Pure load patterns scale better with frequency

MOVSD



MOVAPS



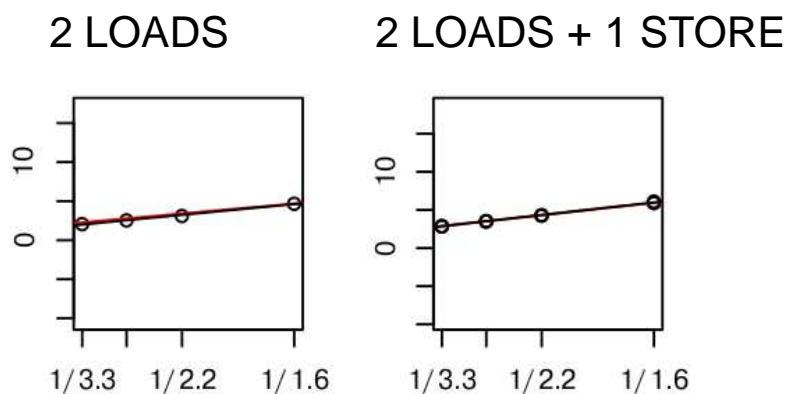
Handling Stores

- L3 Sandy Bridge cache scales much better
- Perfect scaling for all the store and load configurations

Handling Stores

- L3 Sandy Bridge cache scales much better
- Perfect scaling for all the store and load configurations

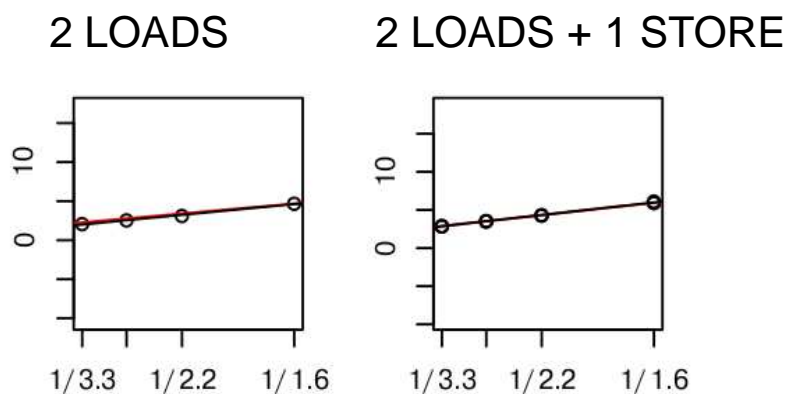
MOVSD



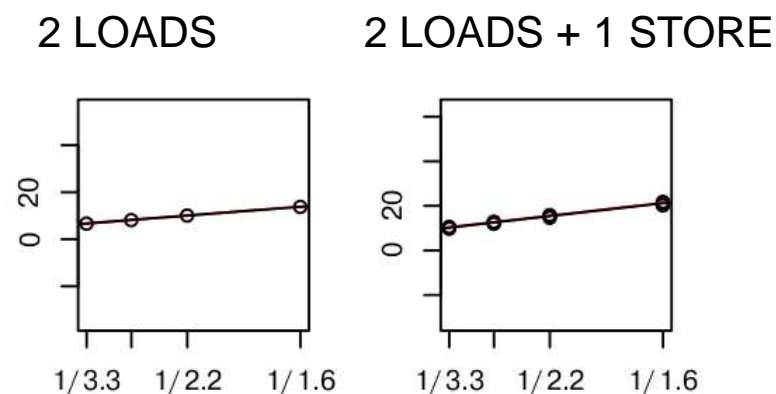
Handling Stores

- L3 Sandy Bridge cache scales much better
- Perfect scaling for all the store and load configurations

MOVSD



MOVAPS



Energy

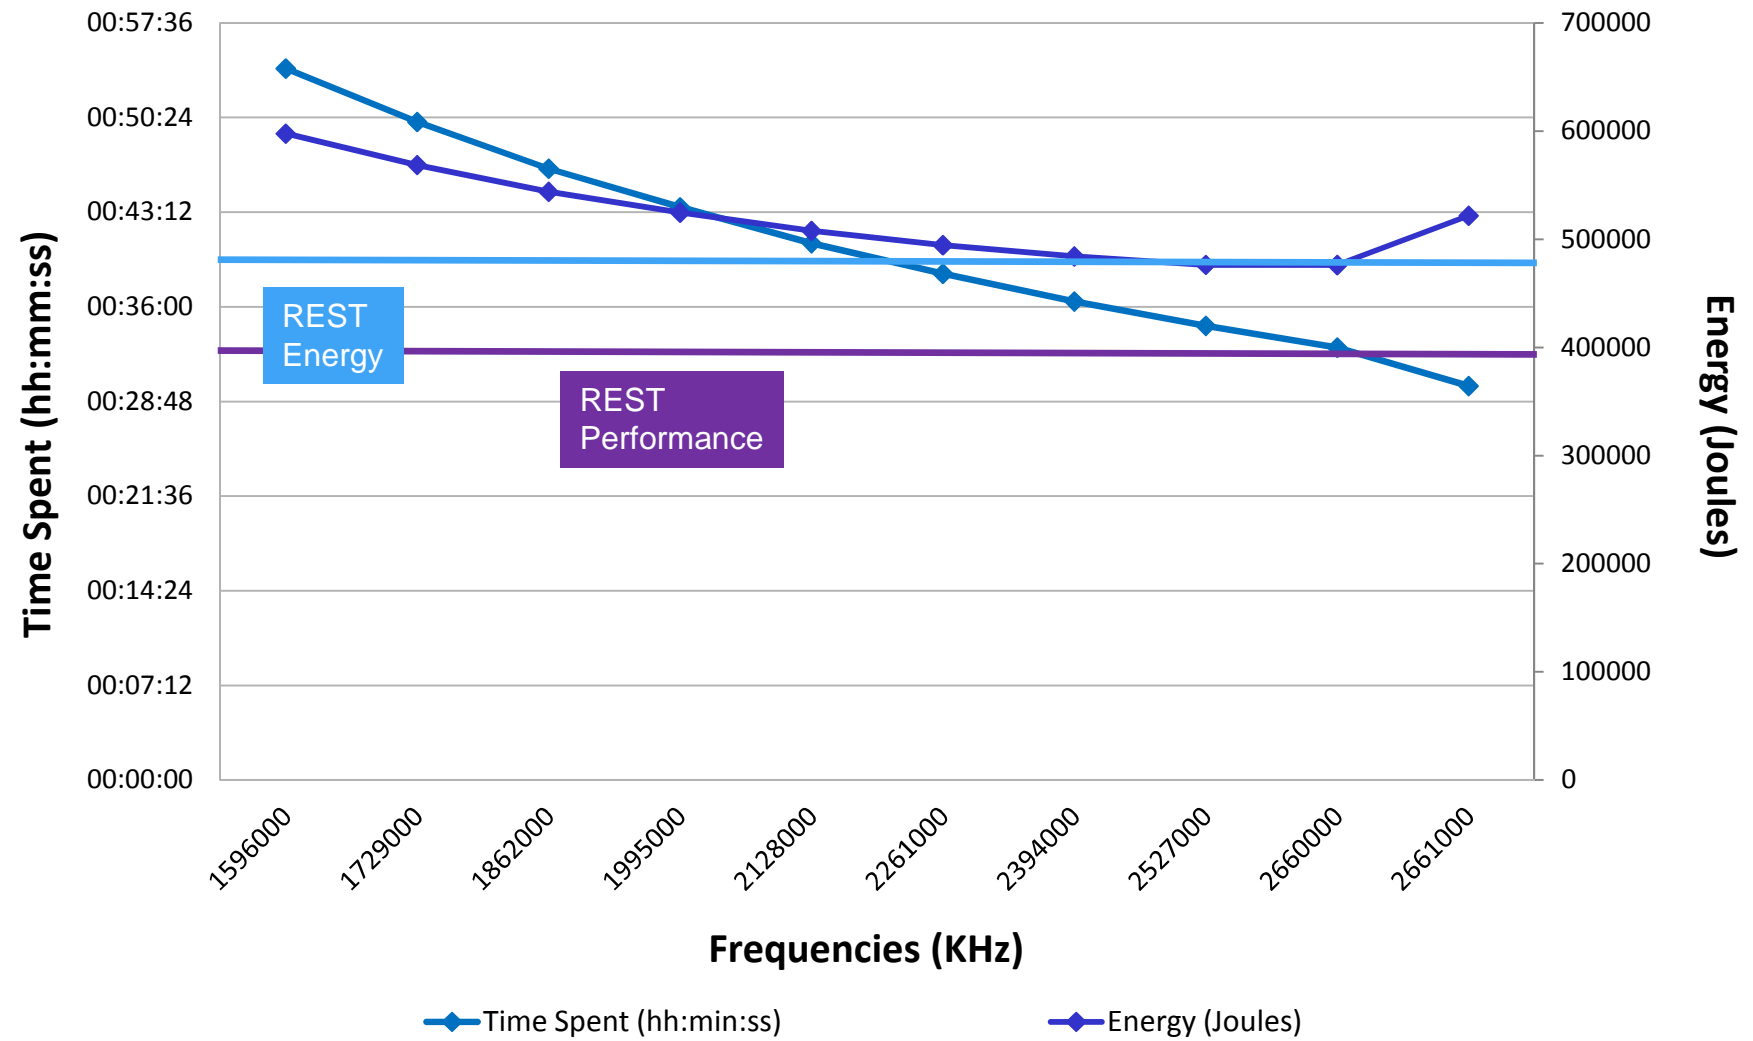
Results

Used a dual-Nehalem, 6-cores each, nine possible frequencies

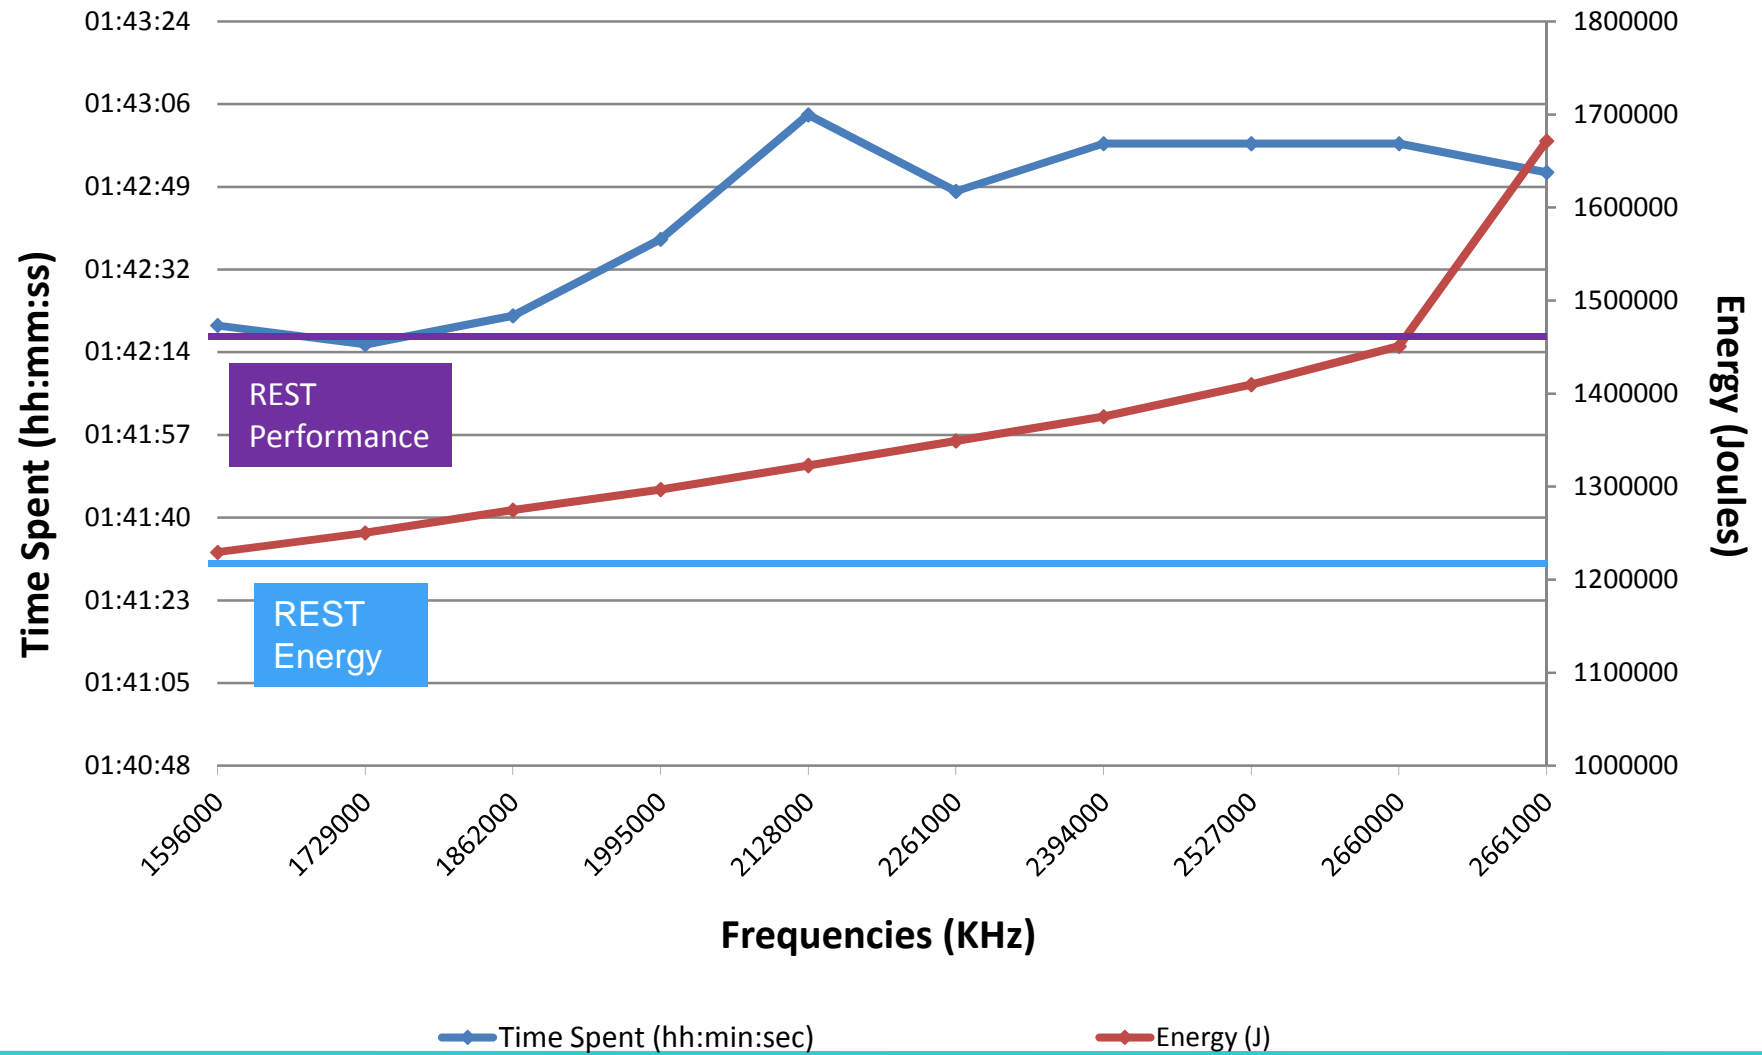
Used a Sandy Bridge Quad-core, sixteen frequencies

Compared with O3 execution, the Linux governors and static frequencies

Gromacs Nehalem (Energy)

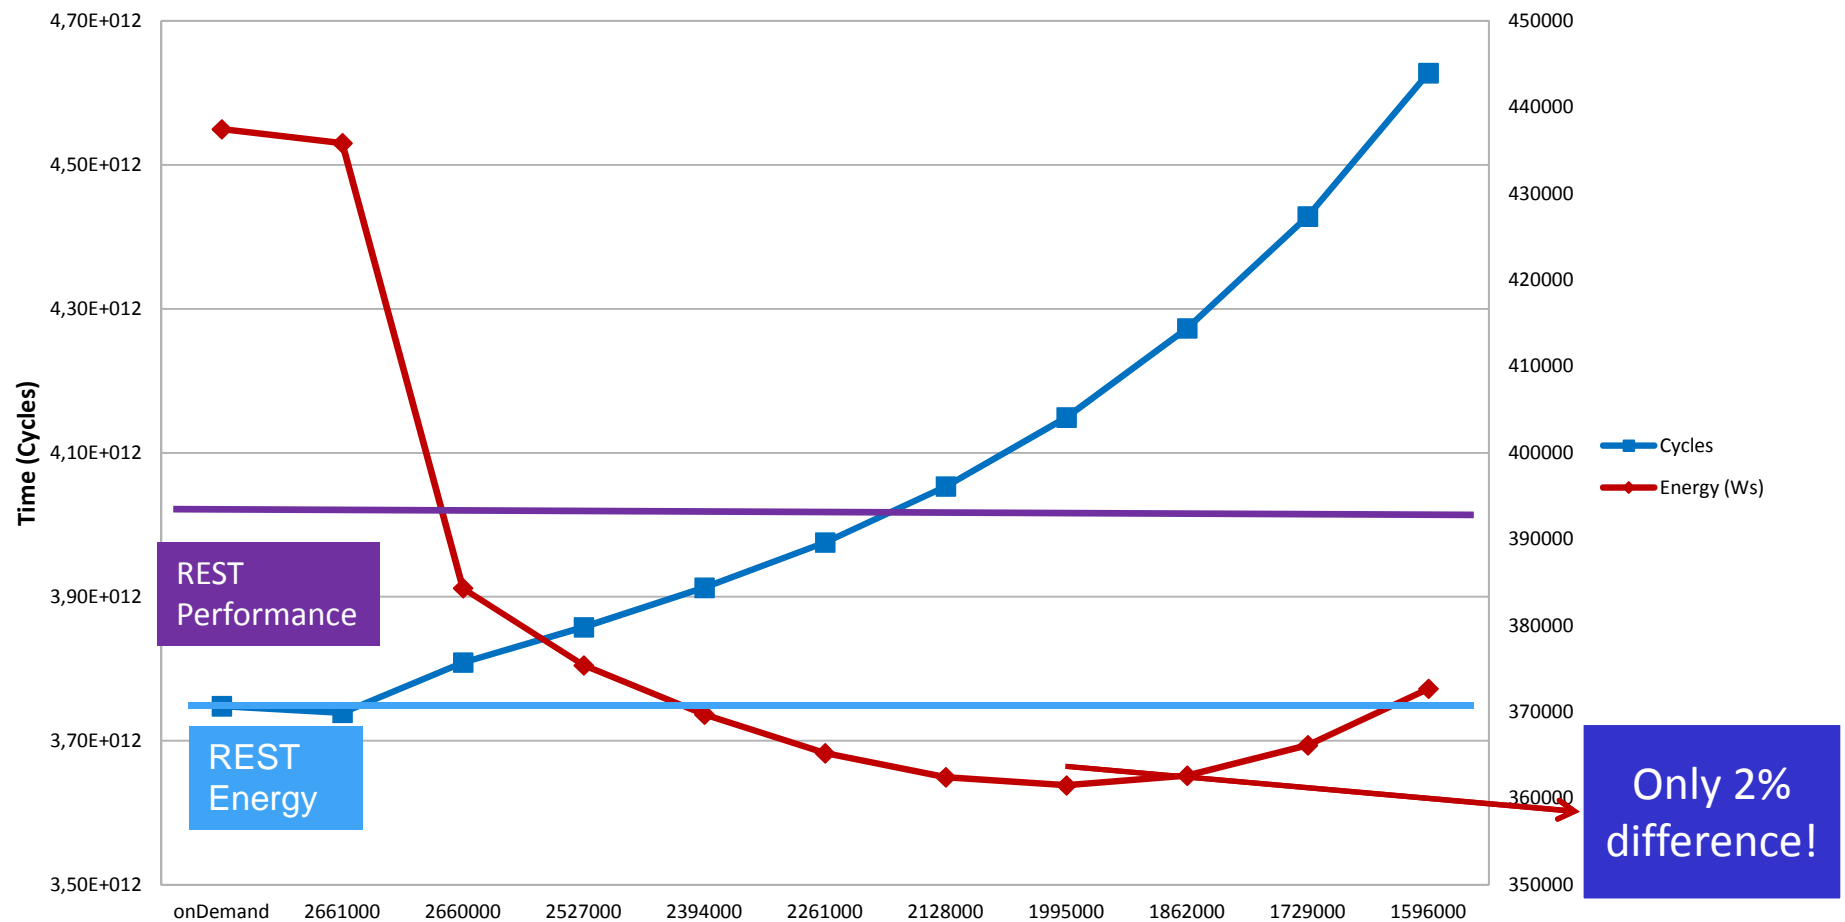


Libquantum Nehalem Energy

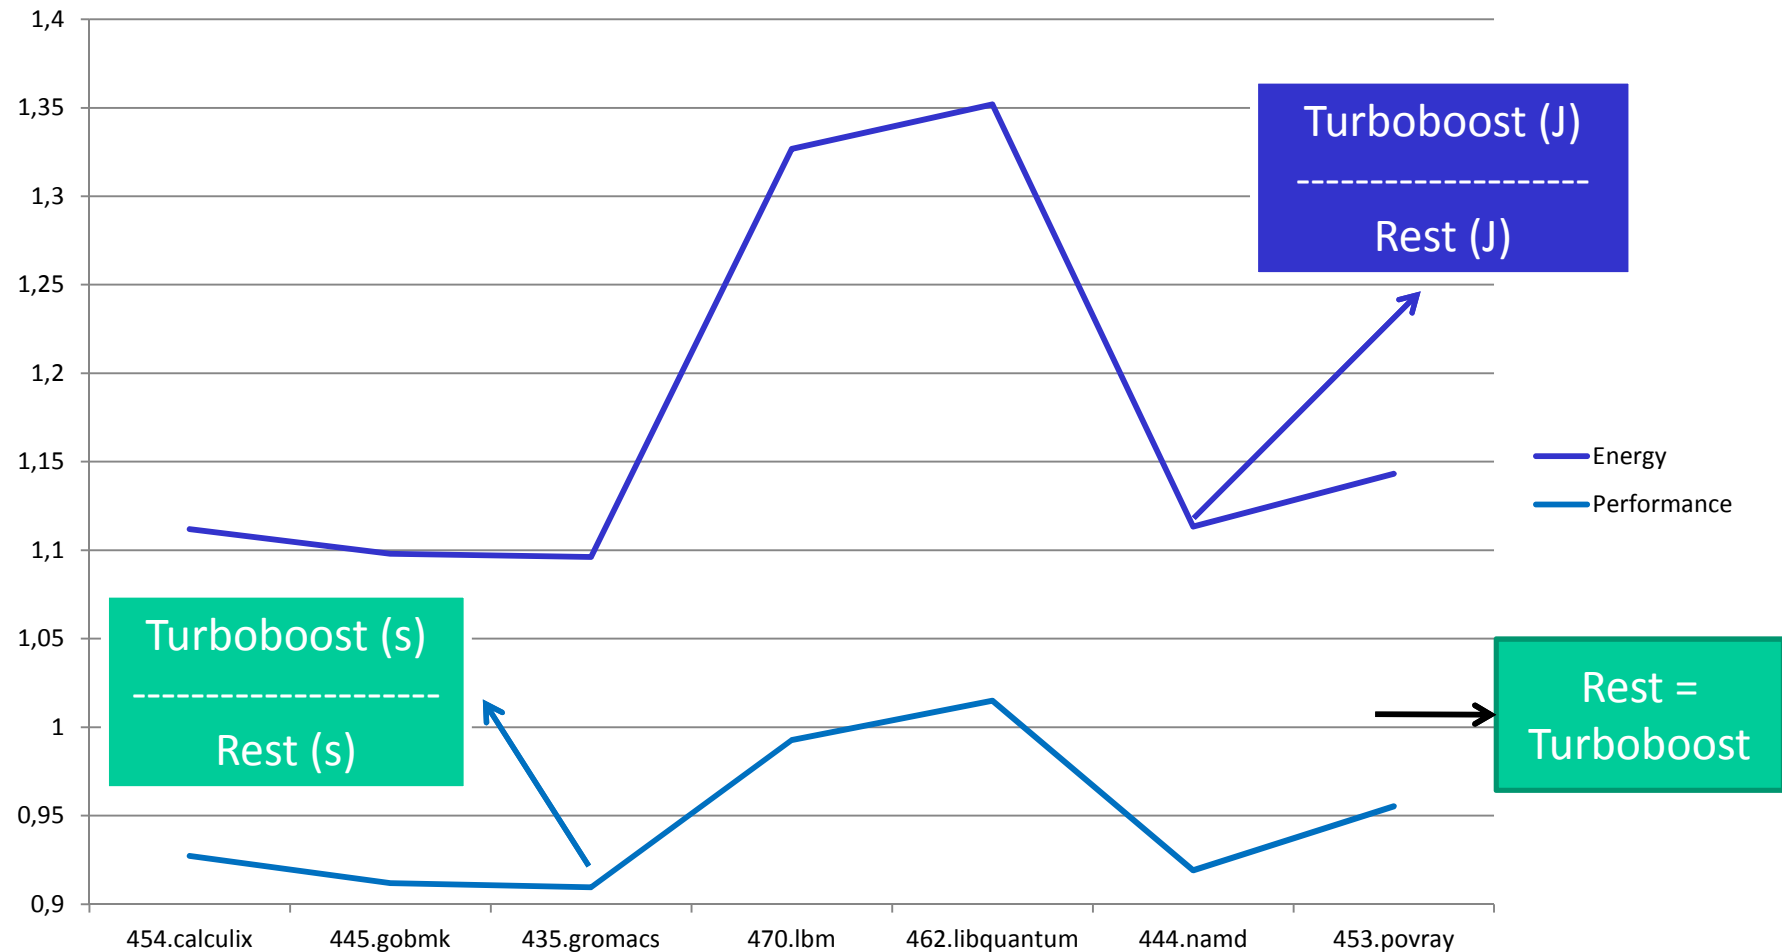


RTM on Nehalem

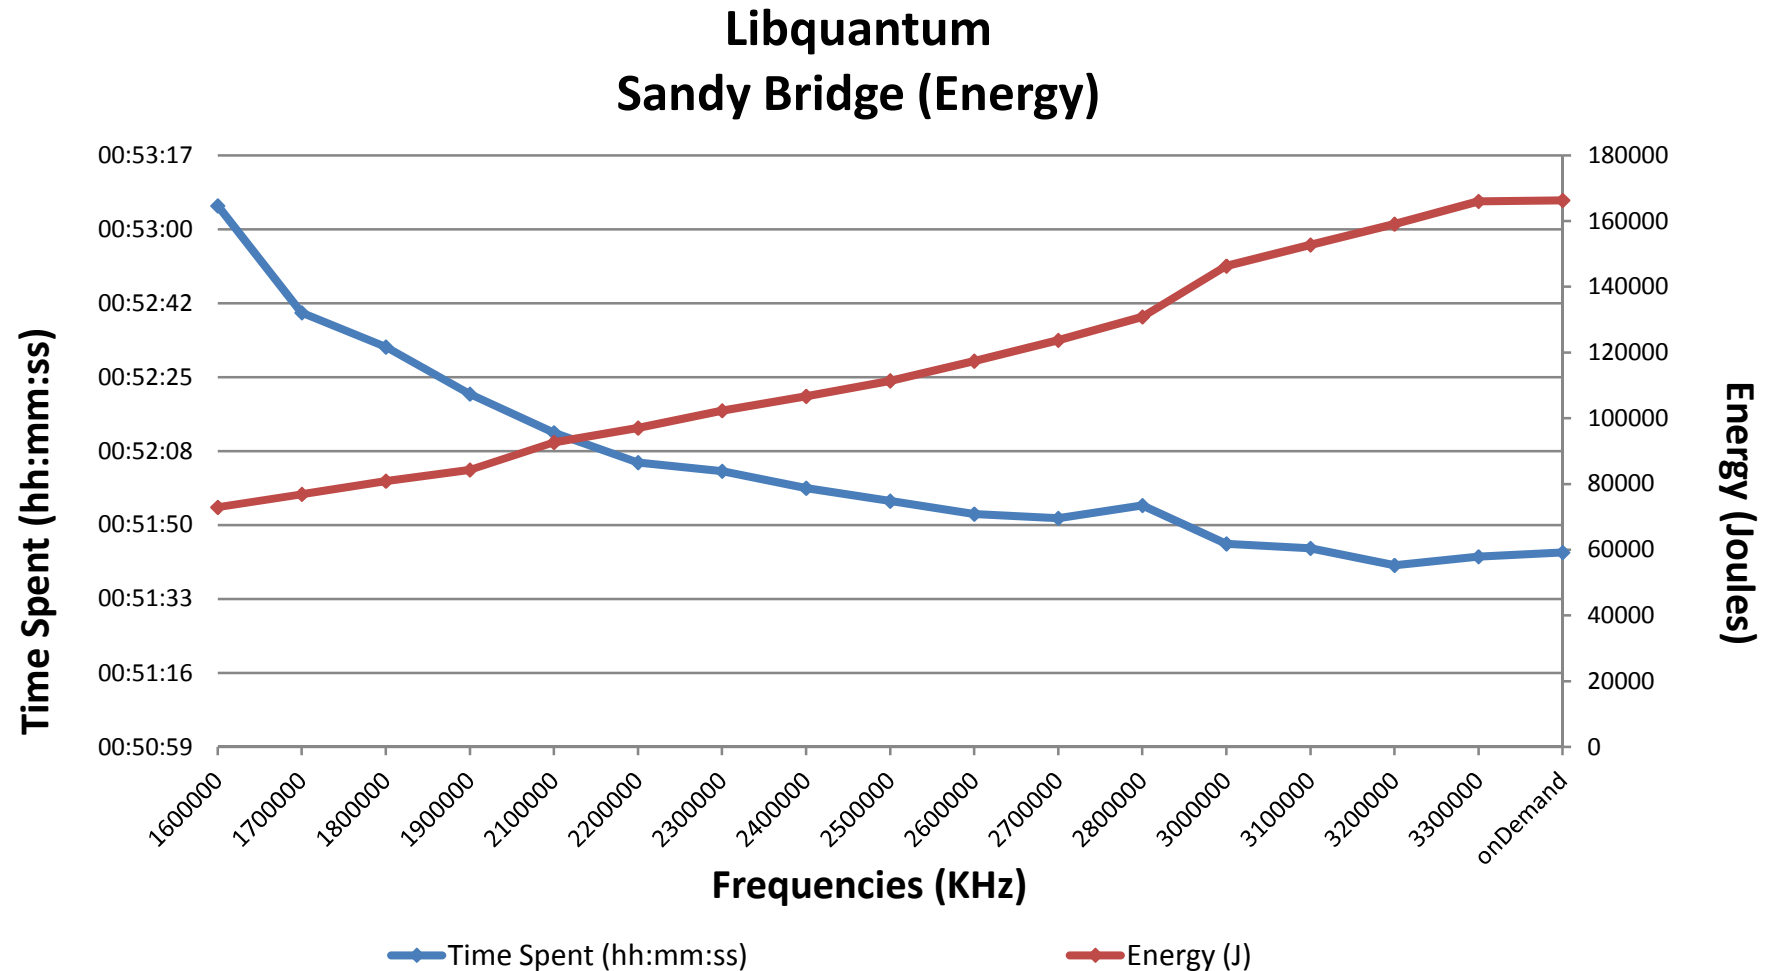
RTM Power and Performance



Ratio Turboboost on REST

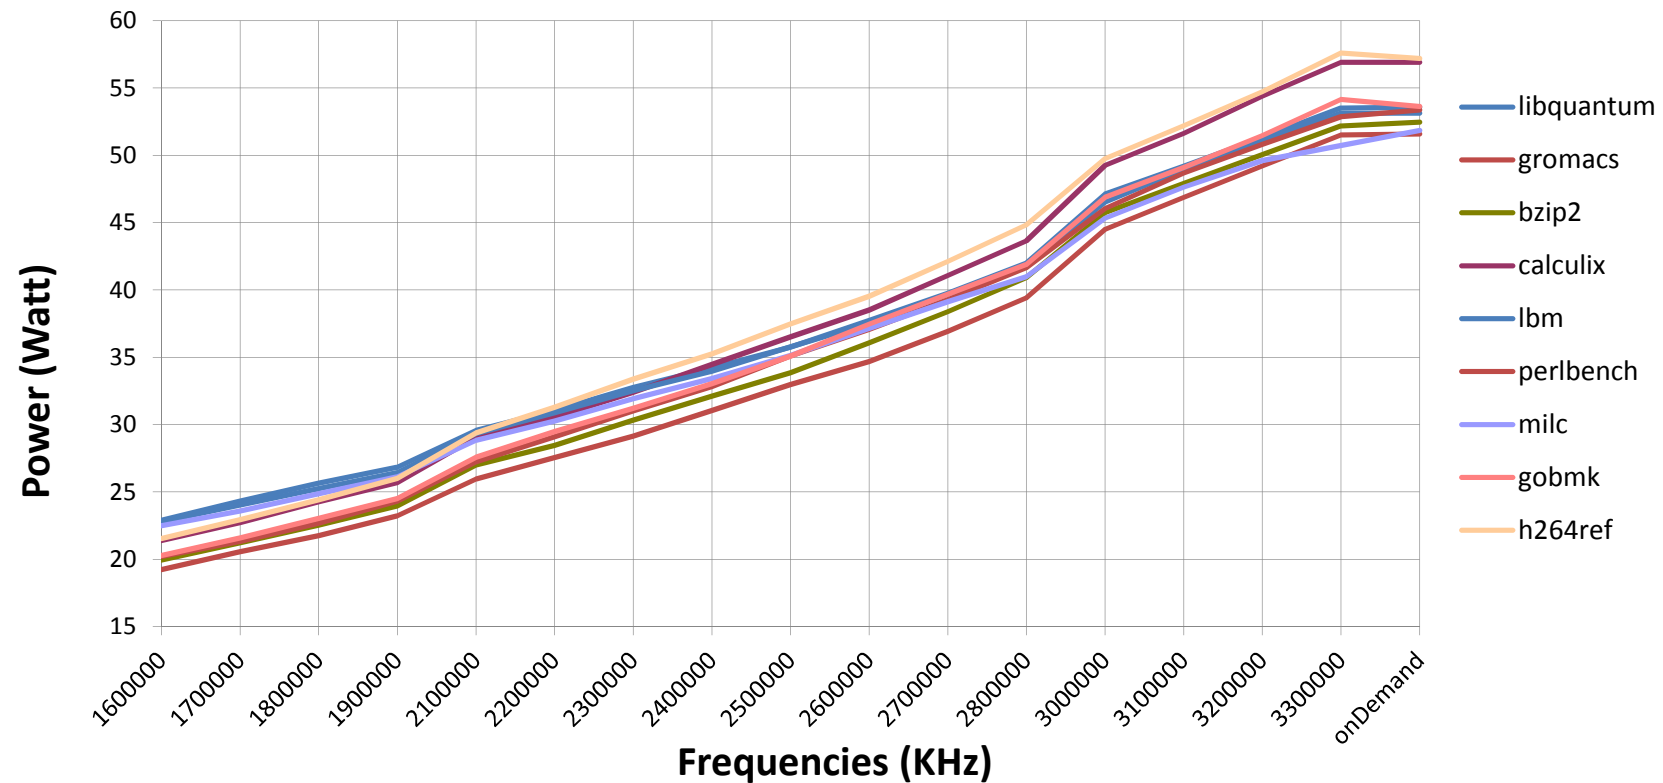


What about Sandy bridge?



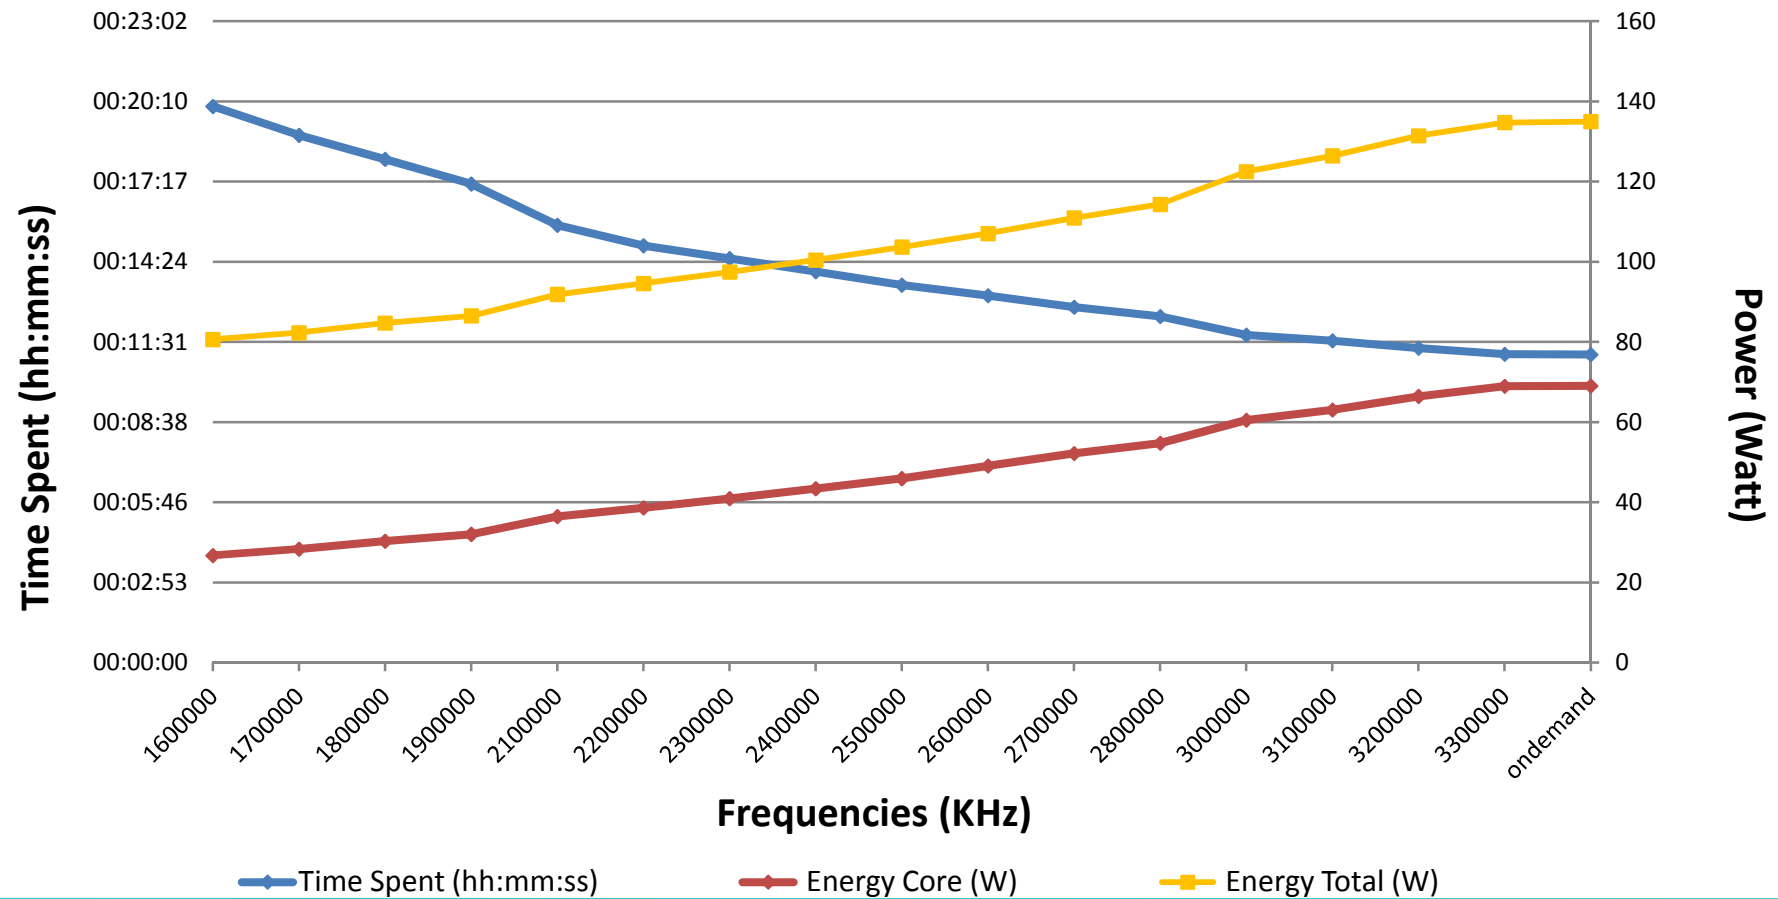
What about Watts?

Watt Usage on Sandy Bridge SPEC2006



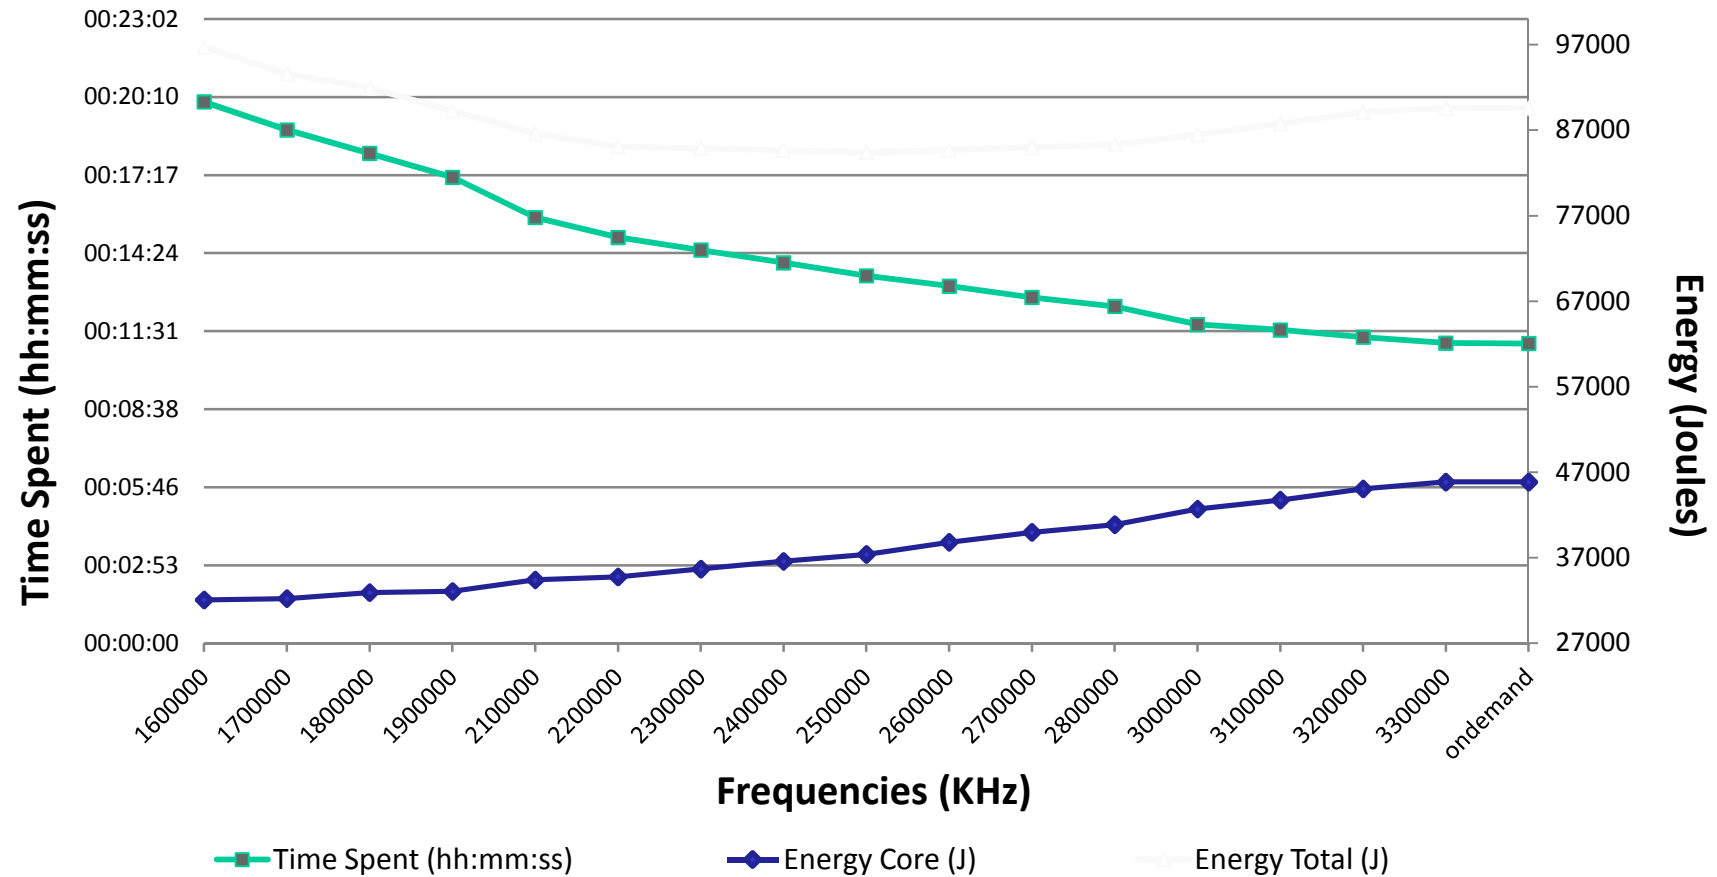
Sandy Bridge's Hardware counters 1/2

QMC=Chem
Sandy Bridge (Power)



Sandy Bridge's Hardware counters 2/2

QMC=Chem
Sandy Bridge (Energy)



UpDownbench profiled with Eprof (Hardware counter from PAPI)

