Predictive Models for Thermal Modelling

Andrea Bartolini
DEI, Università di Bologna
<a.bartolini@unibo.it>

Layout

- Motivation
- Static Thermal Modelling
- Dynamic Thermal Modelling
- Thermal Management
- Datacenter automation

Motivation

Intel Pentium - *1 core*

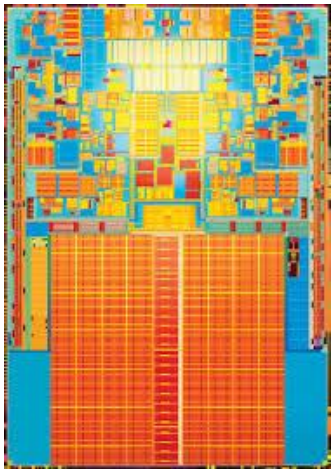

Same Power Budget !

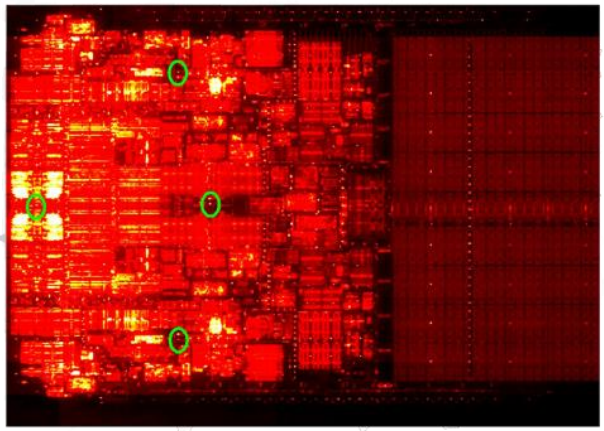

1. ↓ frequency
2. ↓ # active core at the same time
3. ↑ cooling cost

End of Dennard's Scaling
=> Power density increases
=> Thermal issues

Moore's Law:
"The number of transistors and resistors on a chip doubles every 24 months"

↑ theoretical max frequency
↑ #cores

Intel KNL – 72 cores

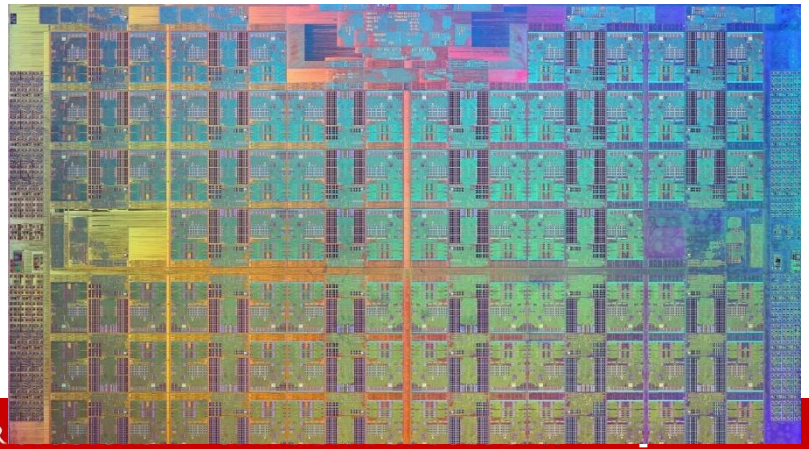

Datacentre: The problem scales up

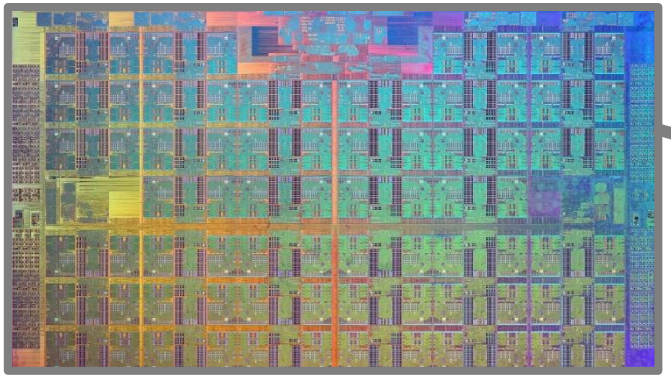

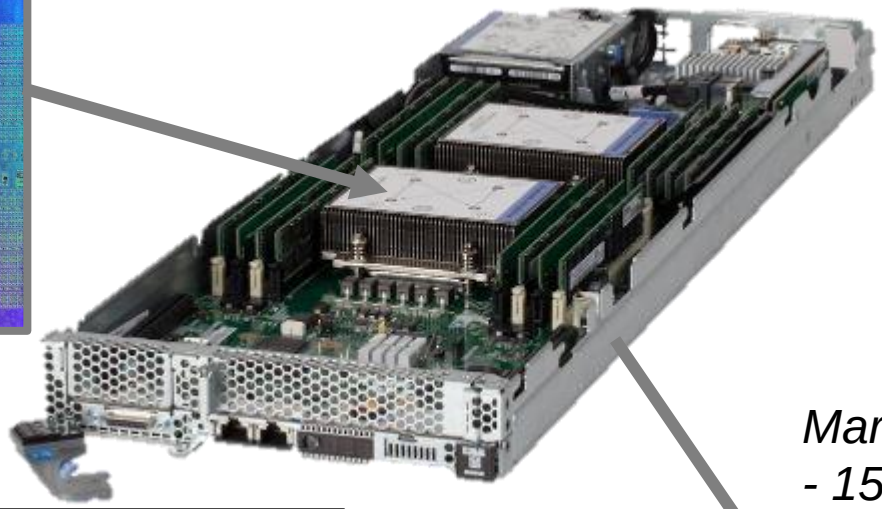

Marconi @Cineca
- 1512 nodes
- 54432 cores

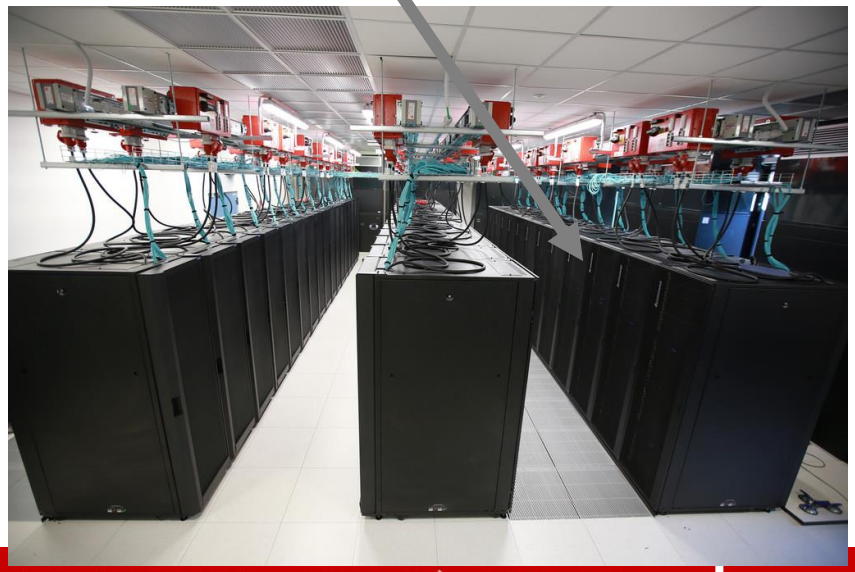

Tihane-2
(most powerful supercomputer 2013-2015)
- 3120000 cores,
- 17.8MW IT only => 24MW w. cooling

Datacentre: The problem scales up

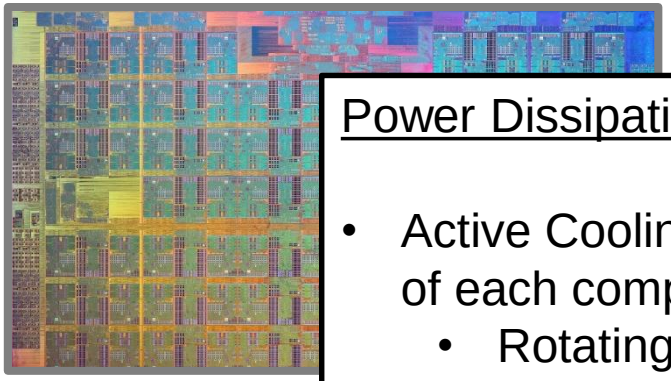

Power Dissipation Problem:

- Active Cooling => forced air/liquid flow on the surface of each computing device
  - Rotating Fans, Pumps are needed
- Heat transferred to the fluid needs to be removed from the room
  - Chillers and Air conditioners
  - Direct exchange with the ambient (Free Cooling)
- Variable workload on each core and on the entire machine

=> Worse case static design is inefficient!

Tihane-2
(most powerful
- 3120000 core
- 17.8MW IT or

@Cineca
nodes
cores

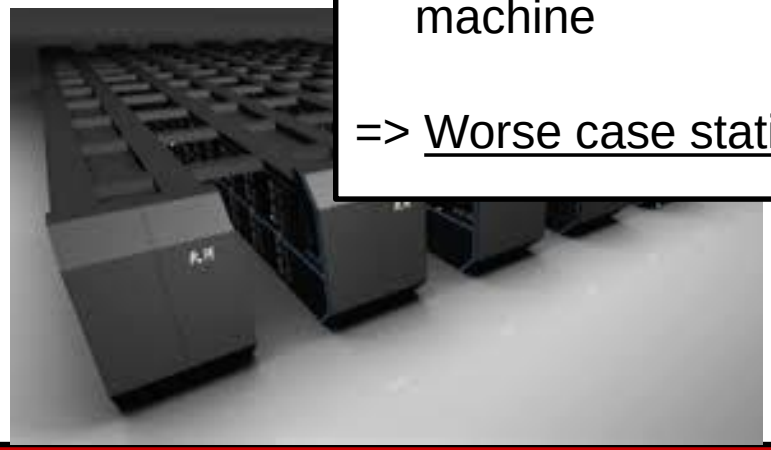

Datacentre: The problem scales up

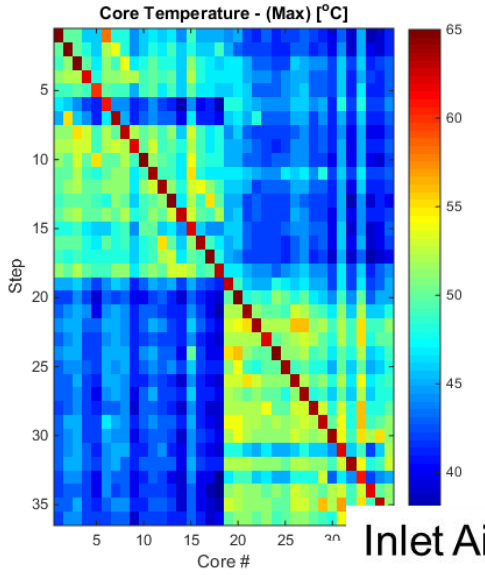

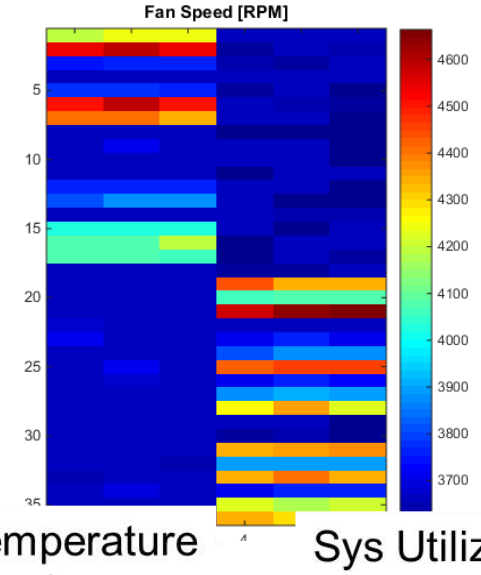

on the surface

e removed from i @Cineca

Inlet Air Temperature

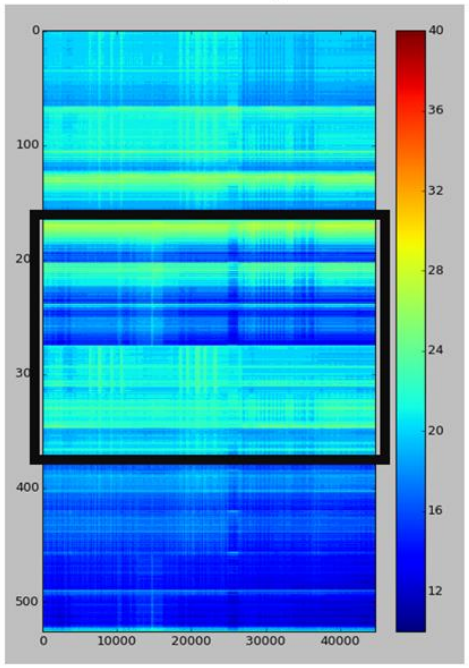

Sys Utilizzation

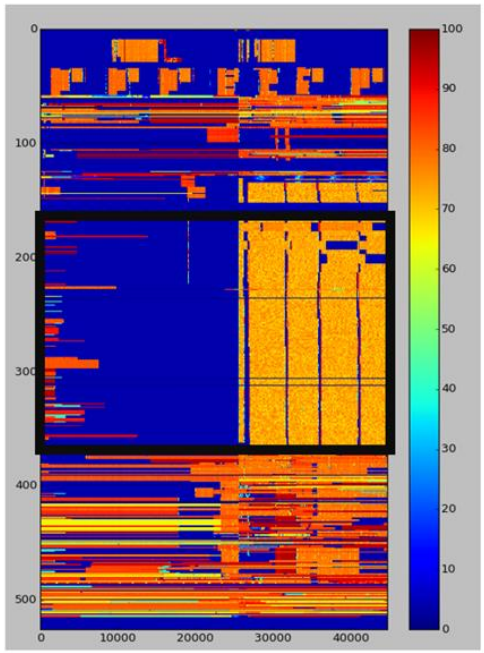

Outet Air Temperature

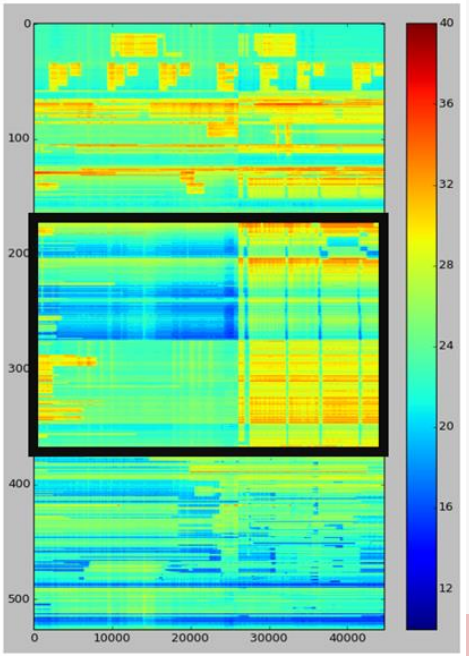

Finane-2
(most powerful
- 3120000 core
- 17.8MW IT or

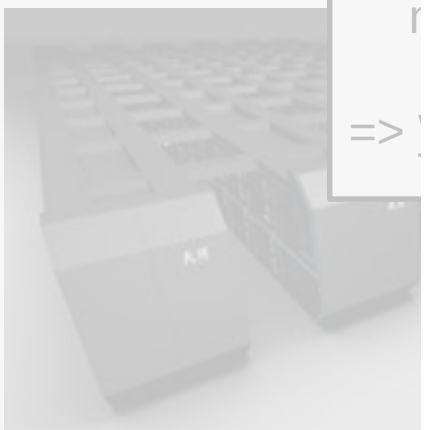

=>

Dynamic Power

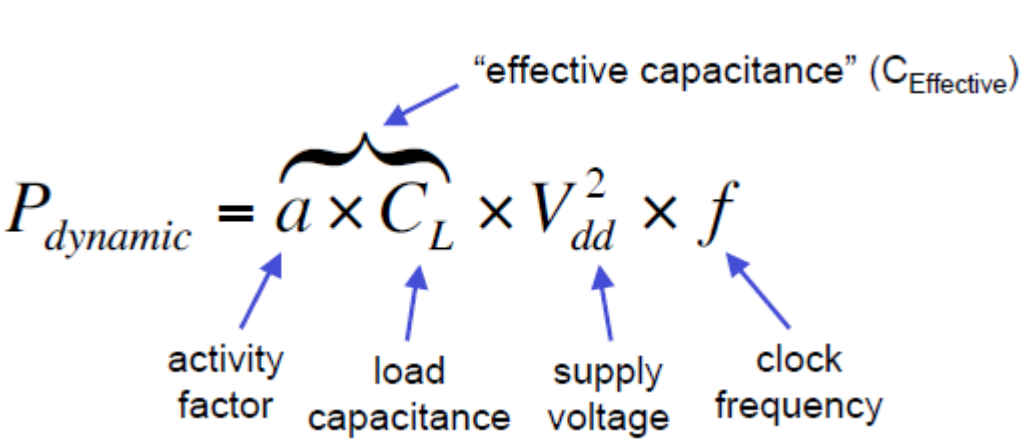

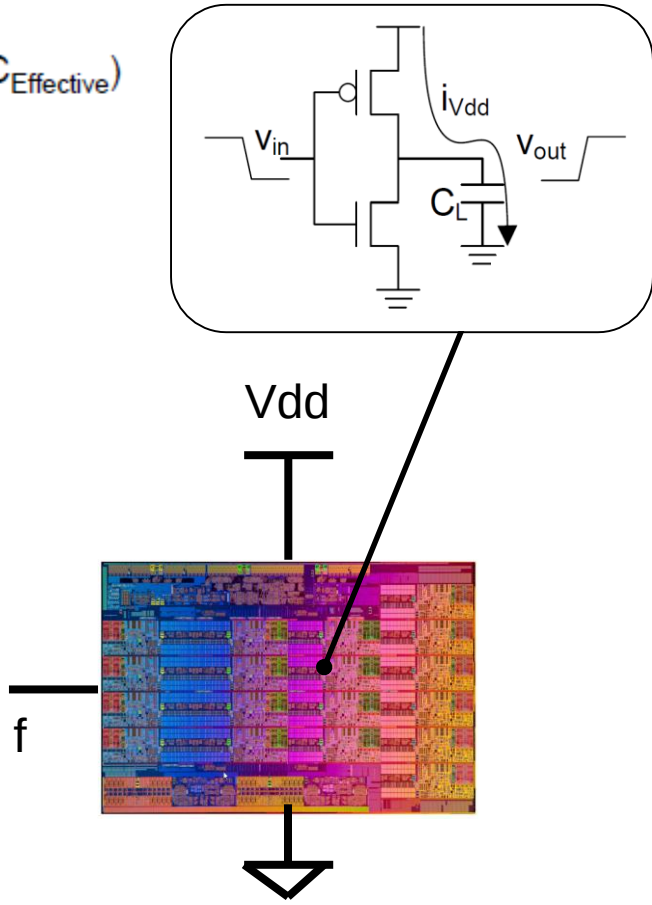

- **Linear ↓ with ↓ $C_{Effective}$**
- **Linear ↓ with ↓ f**
- **Quadratic ↓ with ↓ V_{dd}**
- **Cubic ↓ with ↓ both V_{dd} and f**

Sub-threshold Leakage Current

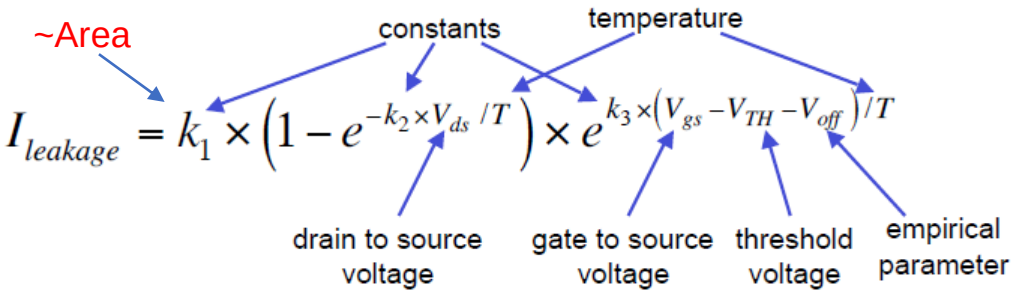

- **Exponential ↓ with ↓ V_{gs} (~V_{dd})**
- **Exponential ↓ with ↑ V_{TH}**
- **Exponential ↓ with ↓ T**

David H. Albonesi ACACES10

Alpha-Power Thermal Model

Delay:

$$D_p = \frac{C_{out}\, V_{dd}}{I_{ON}} = \frac{C_{out}\, V_{dd}}{\mu(T)[V_{dd} - V_{th}(T)]^{\alpha}}$$

Carrier Mobility:

$$\mu(T) = \mu(T_0)\left(\frac{T_o}{T}\right)^m$$

Threshold Voltage:

$$V_{th} = V_{th}(T_0) - k(T - T_0)$$

$T \uparrow$	$\mu \downarrow$	$V_{th} \downarrow$

Delay Trend

- For wires, the resistivity is linearly dependent from T
  - Delay increases as T increases
- For Low VT (LVT) design ($V_{dd} >> V_{th}$)
  - μ dominates w.r.t. V_{th}
  - Delay Increases as T increases
- For High VT (HVT) design ($V_{dd} \approx V_{th}$)
  - V_{th} dominates w.r.t μ
  - Delay decreases as T increases, Indirect Temperature Dependence (ITD)

Thermal Behavior of CMOS gates

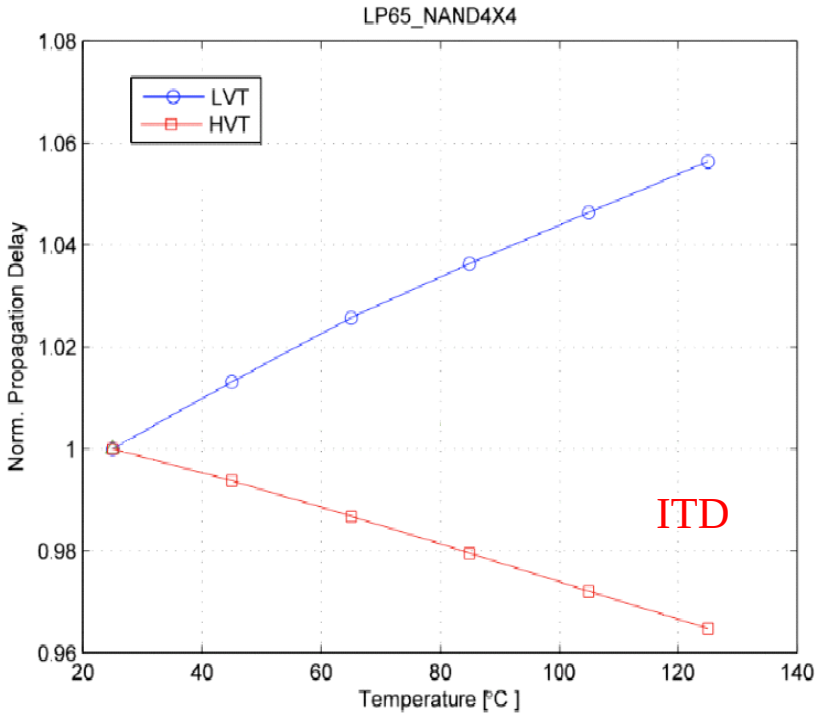

TEST Chip

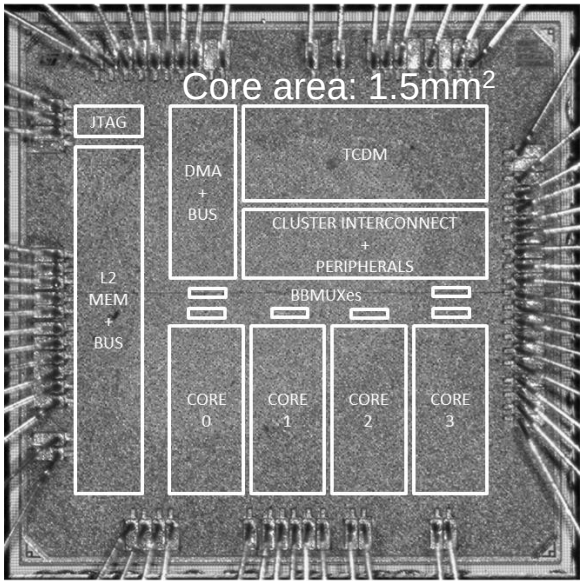

A 60 GOPS/W, −1.8 V to 0.9 V body bias ULP cluster in 28 nm UTBB FD-SOI technology, D. Rossi,et al., Solid-State Electronics, 2016

Impact of body bias on maximum frequency and leakage power over the supported range of voltage supply.vb

First Multi-core ever in RVT FDSOI

PULP
Parallel Ultra Low Power

http://iis-projects.ee.ethz.ch/index.php/PULP

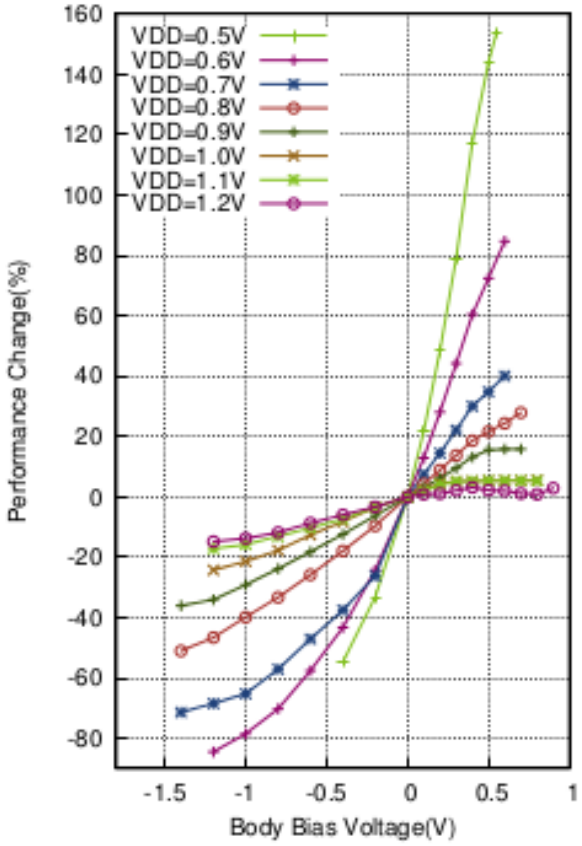

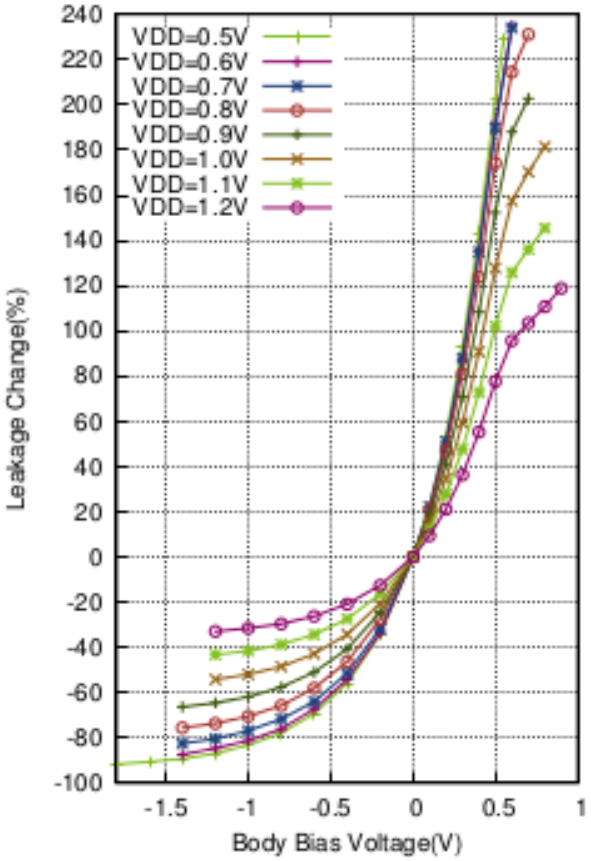

Temperature Impact

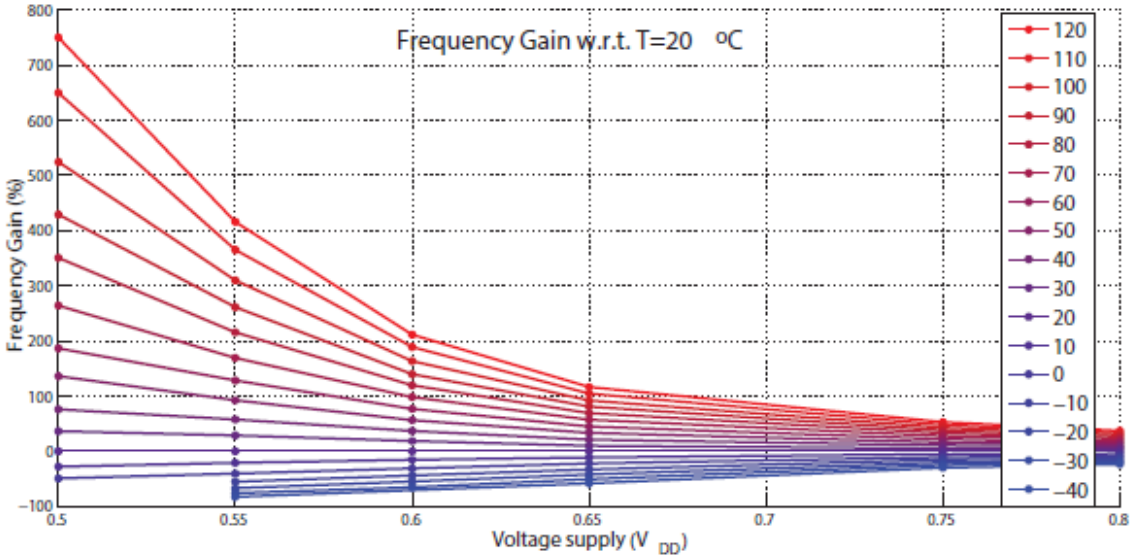

Silicon - up to 800x performance variation at high temperatures

Fitting of Compact models

Vbb,Vdd,Temperature

⇩

Power, Max Freq

$$F_{MAX} = b_1 \cdot log\left[e^{b_2 \cdot V_{DD} - b_3 + b_4 \cdot V_{BB}} + 1\right] \cdot V_{DD}$$

$$P_{DYN} = a_1 + a_2 \cdot F_{MAX} \cdot V_{DD}^2$$

$$P_{LKG} = c_1 \cdot e^{c_2 \cdot V_{DD} - c_3 + c_4 \cdot V_{BB}} \cdot e^{c_5 \cdot V_{DD}} + c_6$$

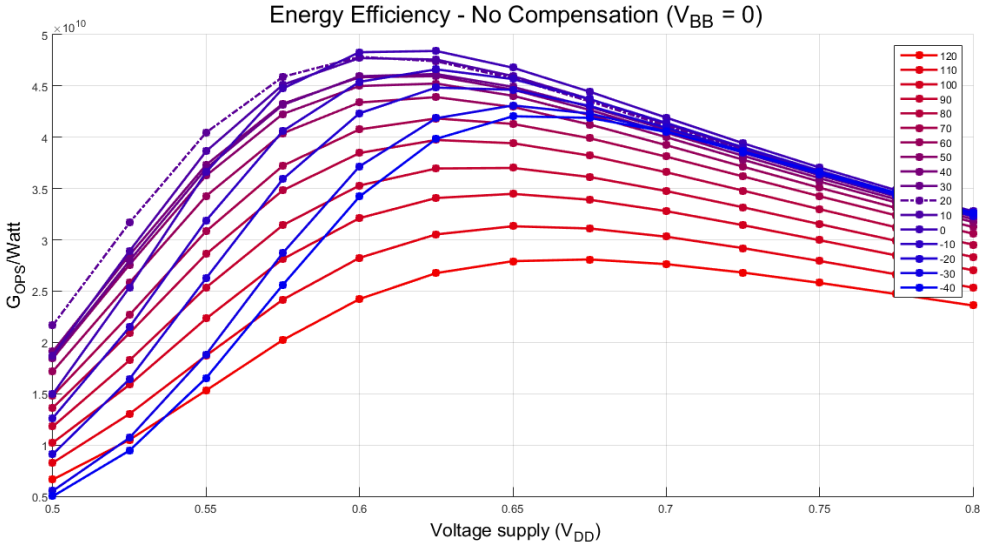

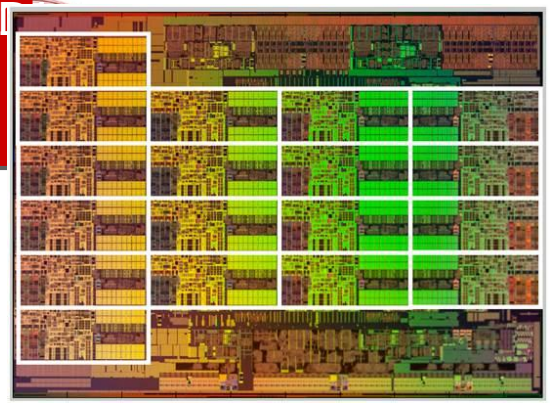

Haswell - E5-2699 - Air Cooled

18 cores
PowerVirus

~10°C
Spatial variation

+100 Watts Fan
+ 20% Power

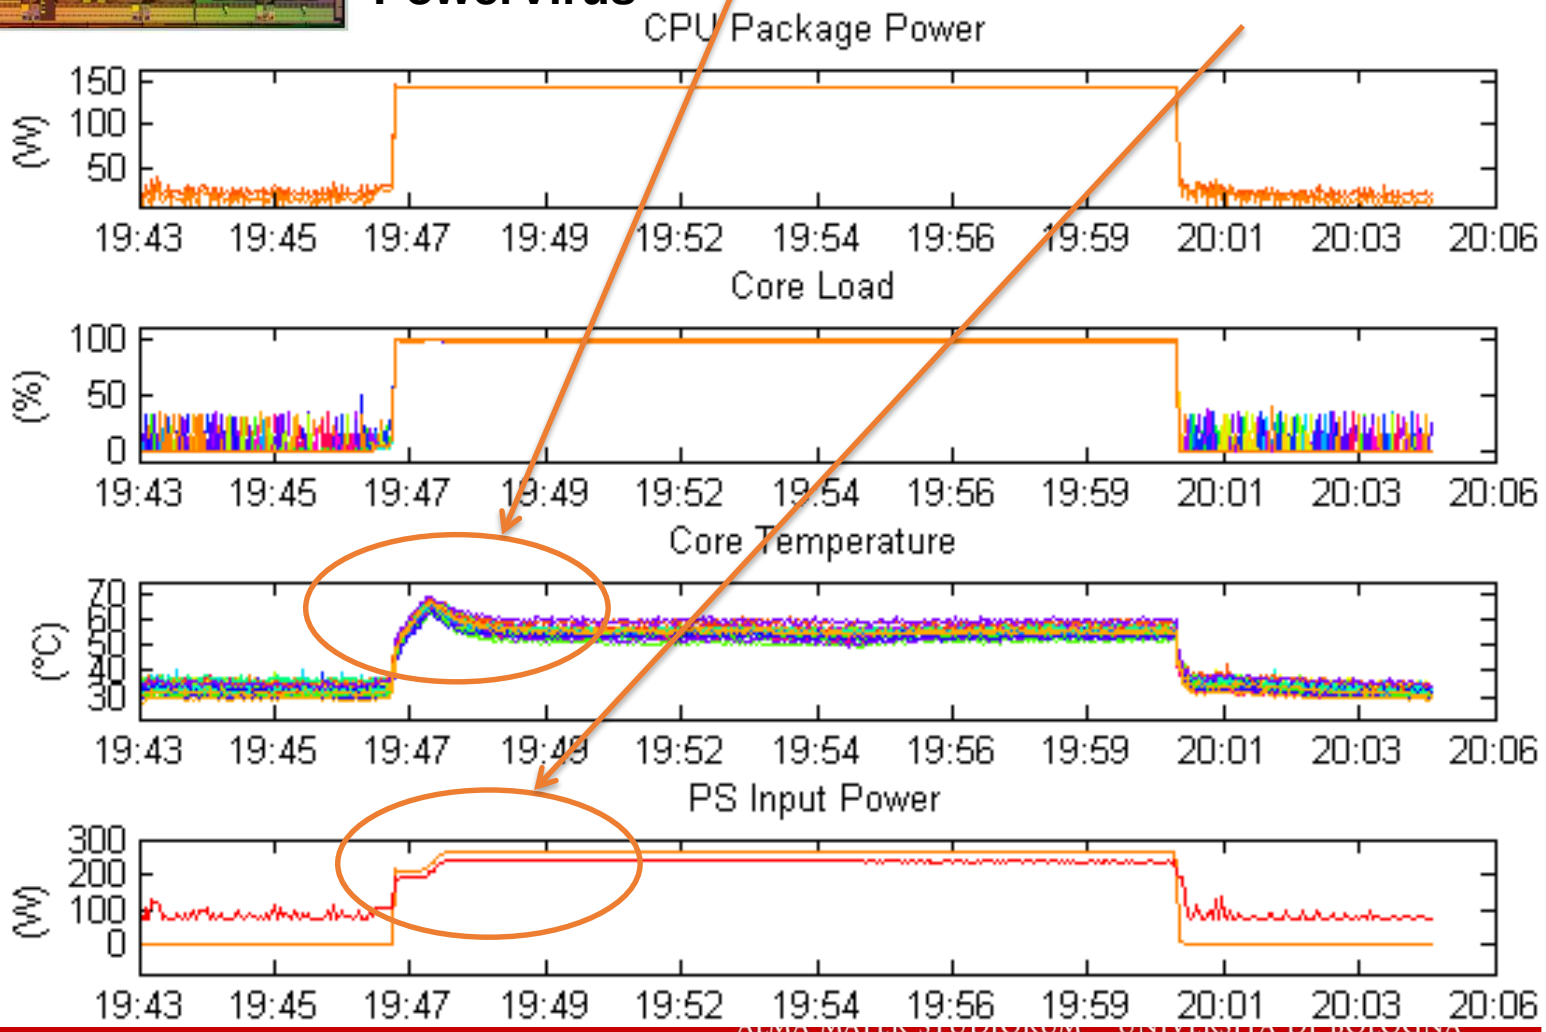

Thermal Impact

#1 Core Rotating Power Virus

Up to 20 C Temperature difference on DIE
~ 30 C Temperature difference in between sockets
- Thermal neighbours exists!

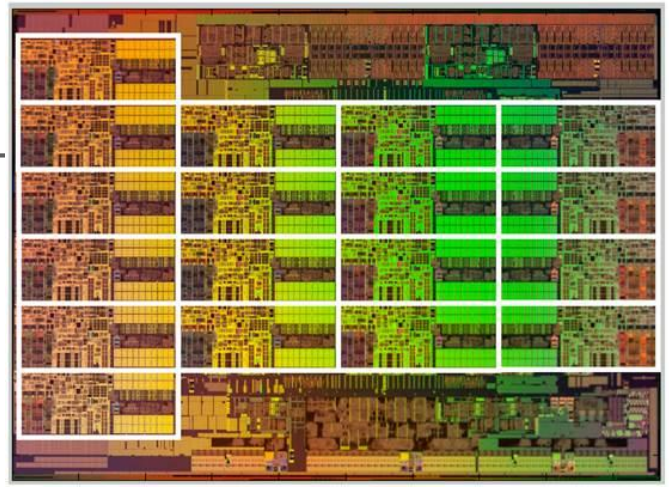

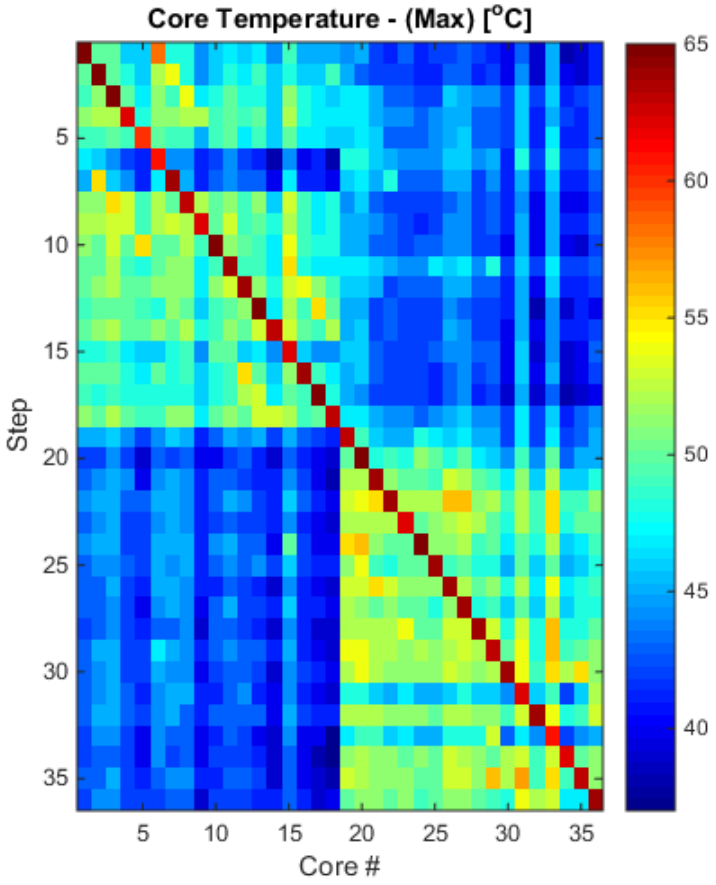

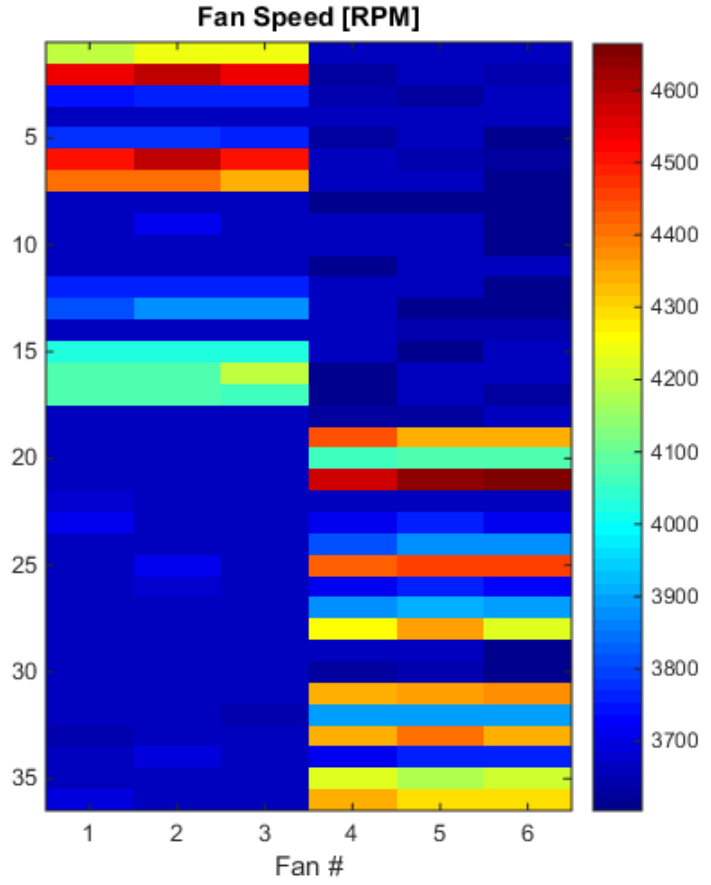

Haswell - PowerVirus #1

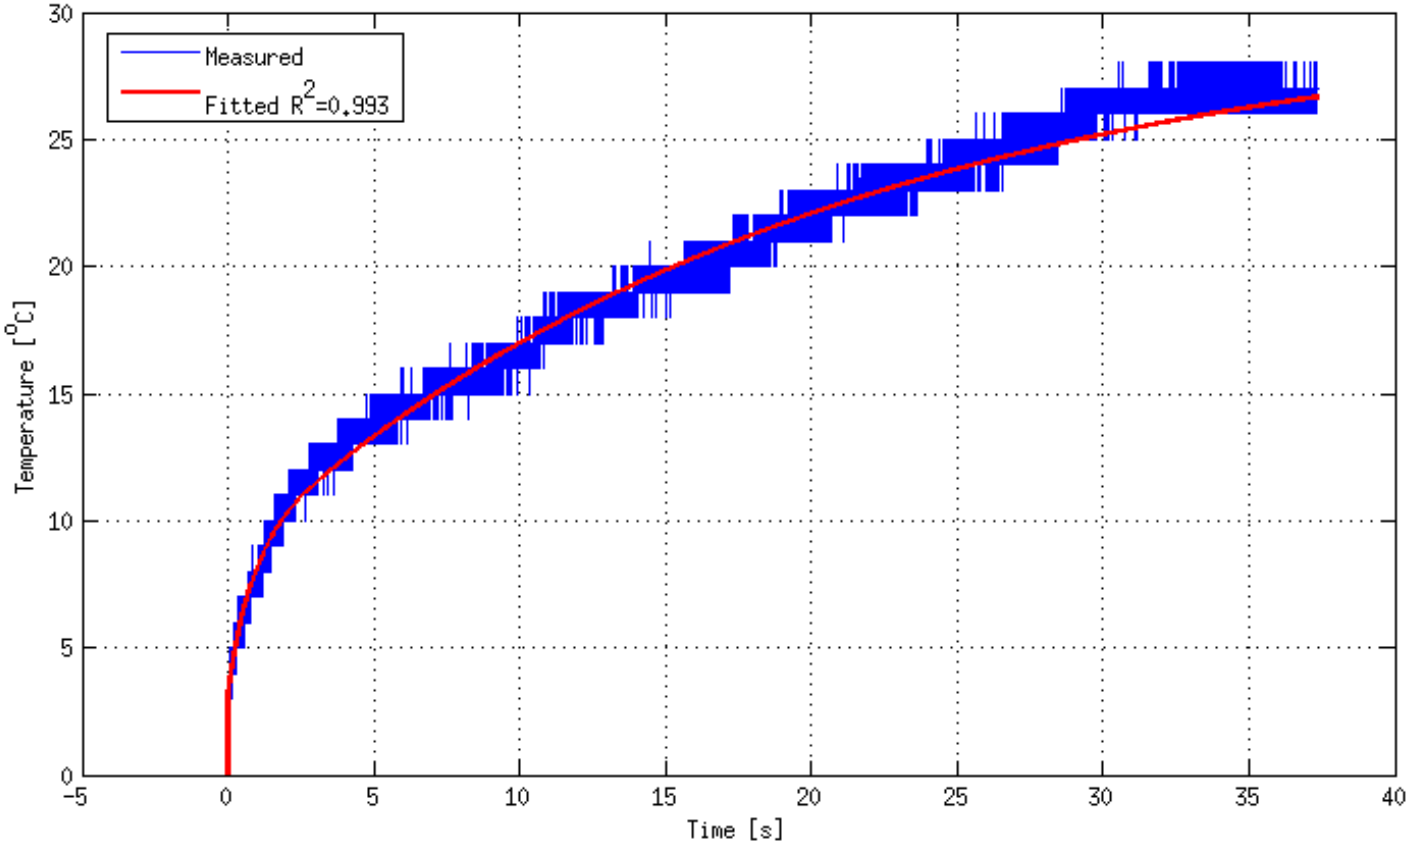

Haswell - PowerVirus #1

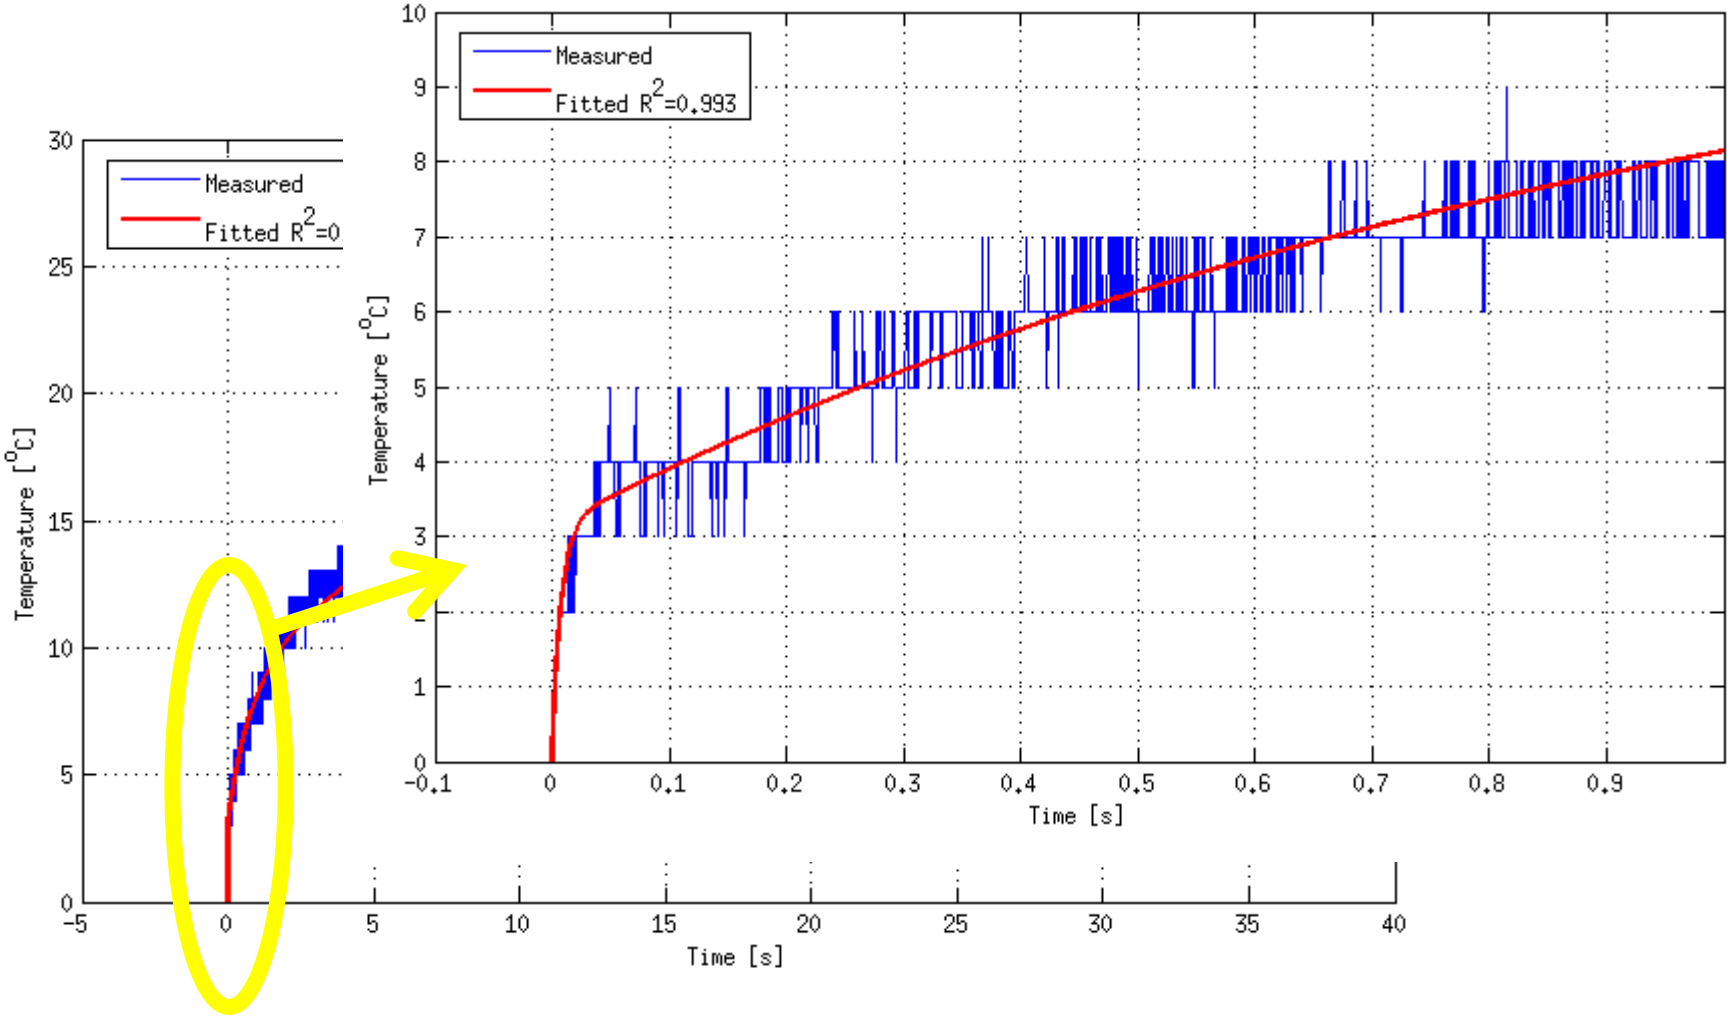

Haswell - PowerVirus #1

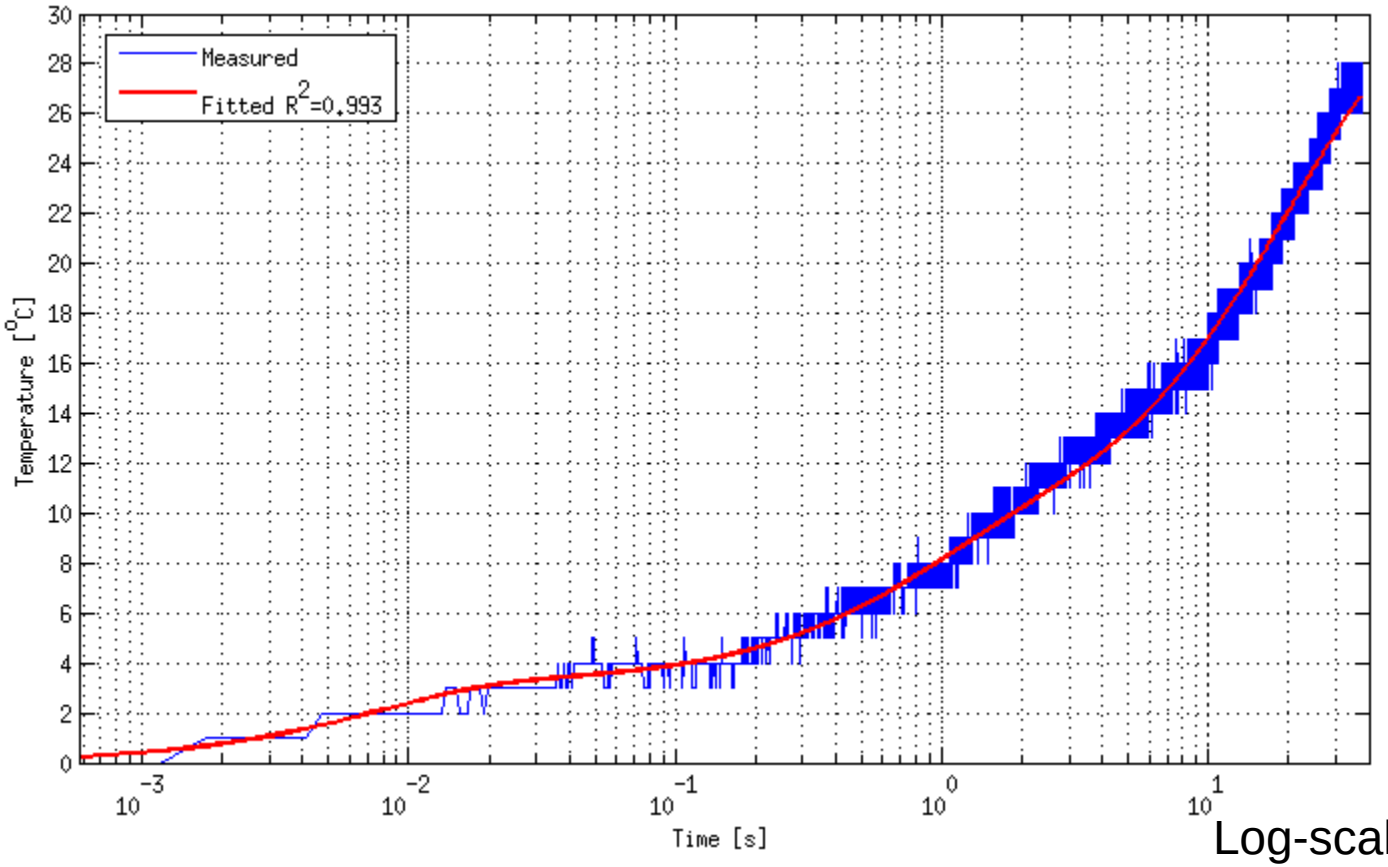

Log-scale

Thermal transient: time constants Model fitting

Tc@0.334ms	Time (s)
t1	0.0076
t2	0.7716
t3	20.0745

Layout

- Motivation
- Static Thermal Modelling
- **Dynamic Thermal Modelling**
- Thermal Management
- Datacenter automation

Thermal Model

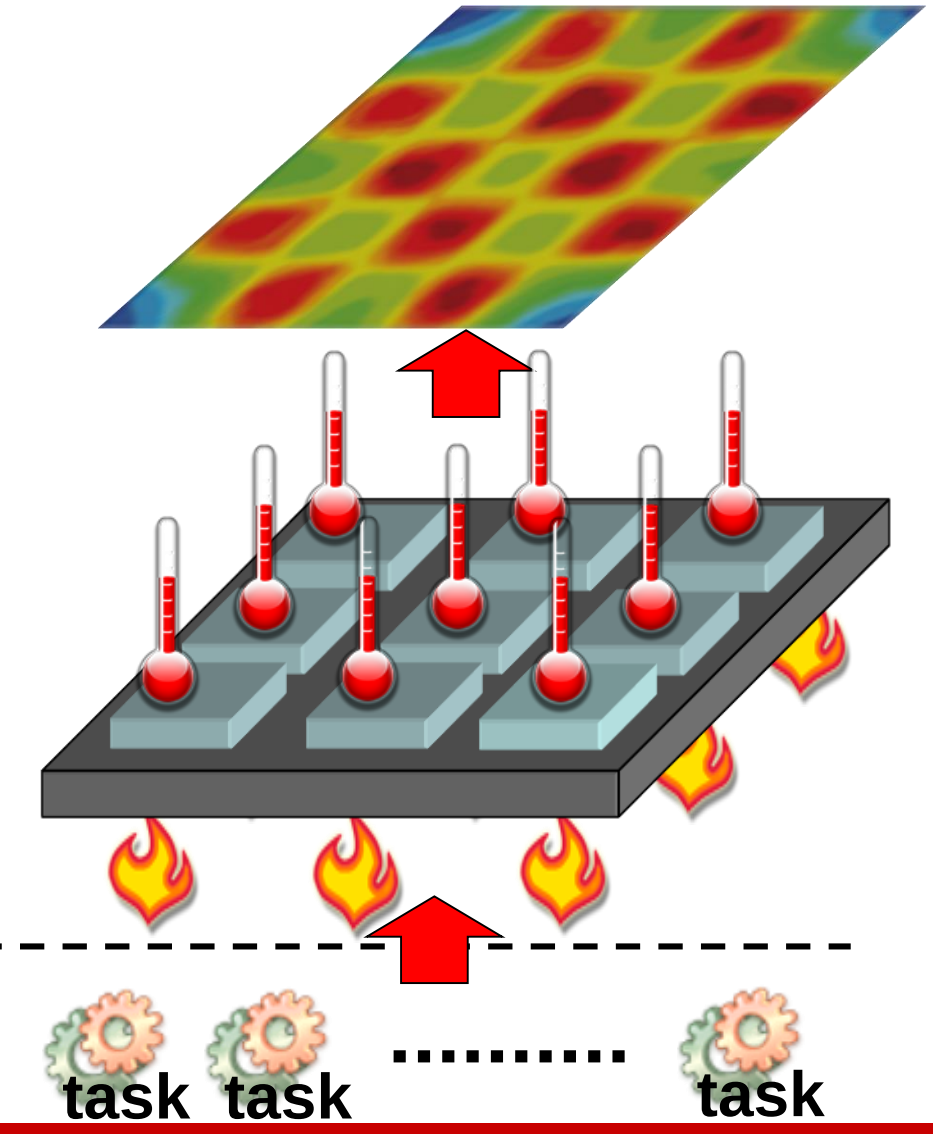

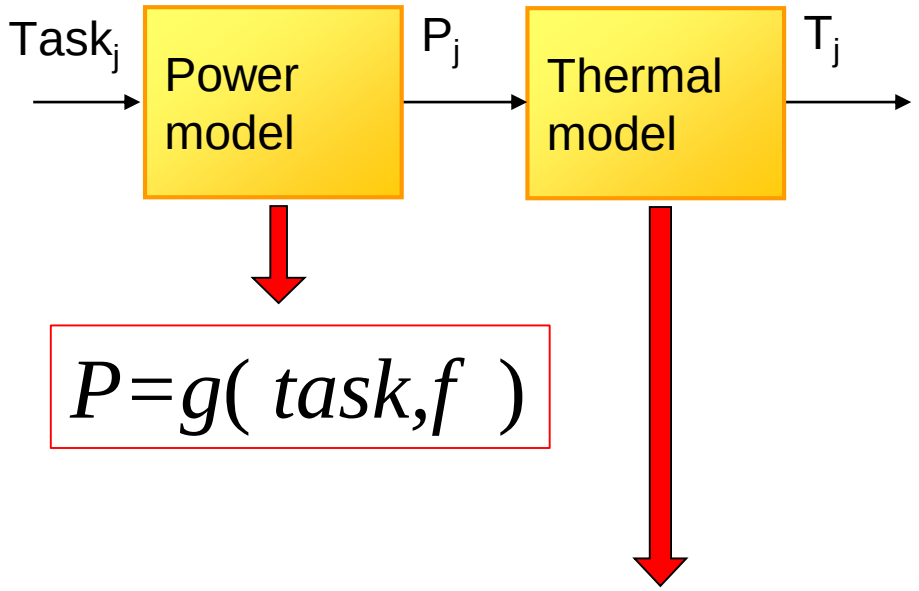

$$P = g(\ task, f\ )$$

Dynamic Model

$$ss:\ T[n+1] = AT[n] + BP[n]$$

$$tf:\ T[n+1] = \frac{B(z^{-1})}{A(z^{-1})}P[n]$$

Static Thermal Model

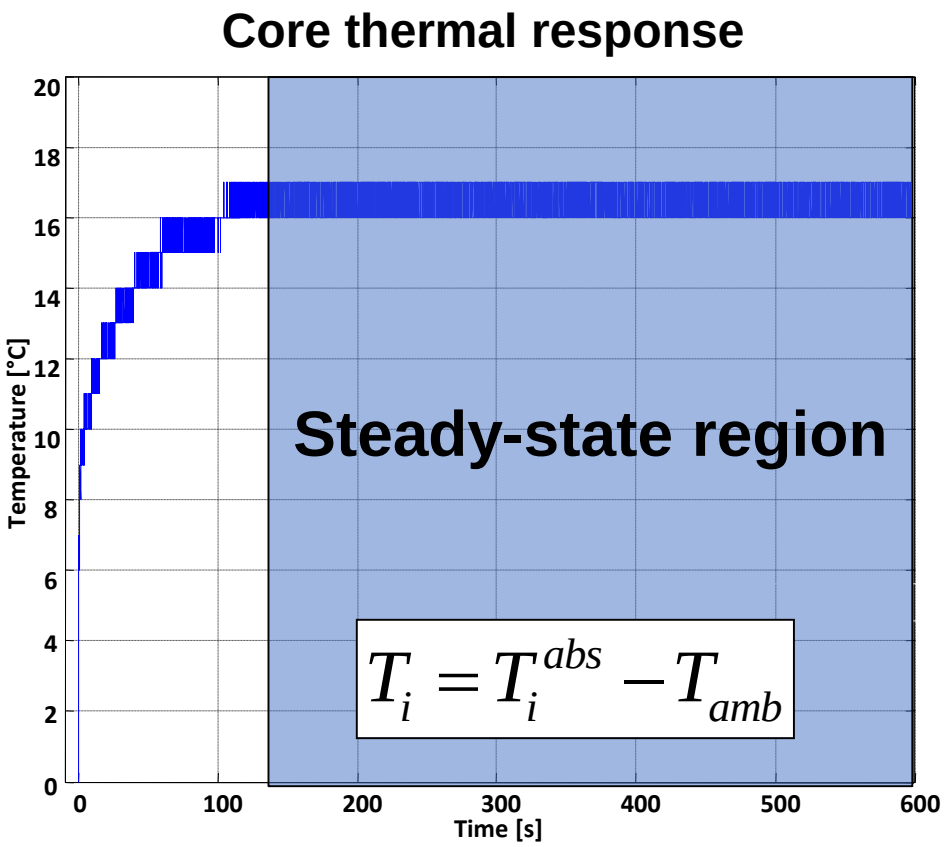

Dynamic Model

$$T[n+1] = A \cdot T[n] + B \cdot P[n]$$

Steady-state condition:

$$T[n+1] = T[n] = T$$

Static thermal model:

$$T = SG \cdot P$$

Steady-state Gain matrix:

$$SG = (I - A)^{-1} \cdot B$$

Platform

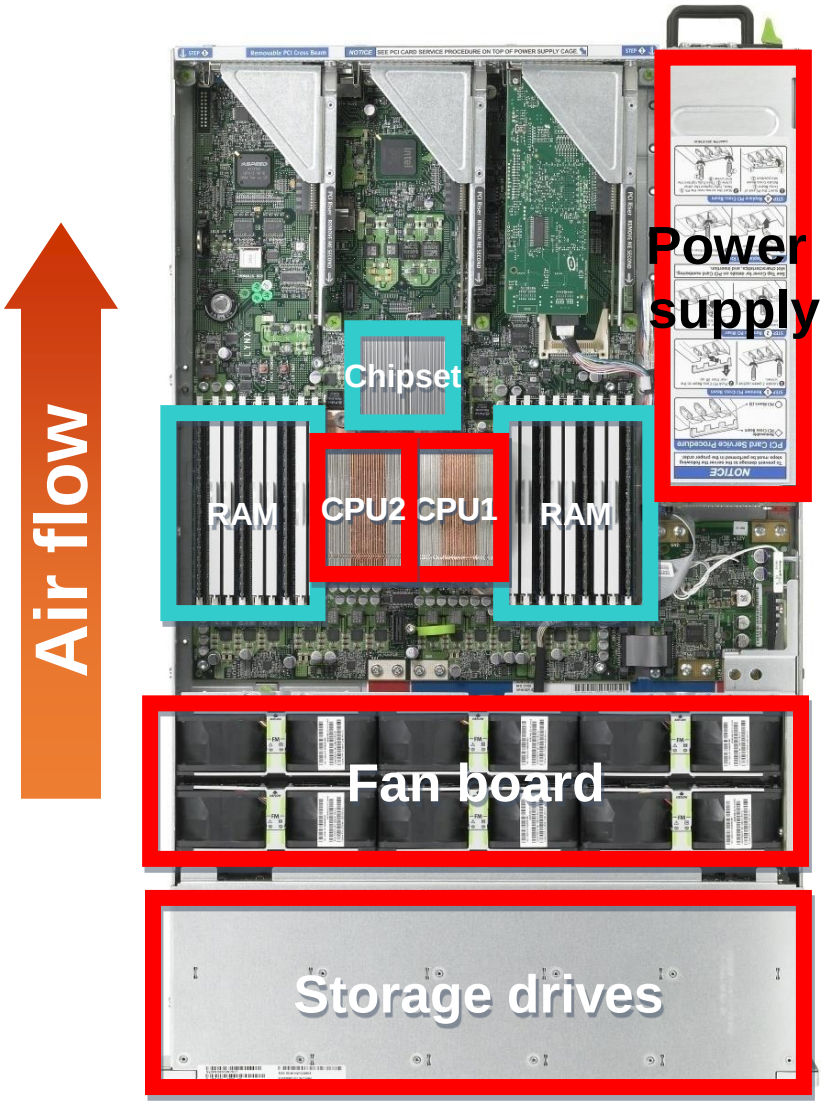

SUN FIRE X4270

- Intel Nehalem 5500
- 8core/16thread
- 1.6÷2.9GHz
- 95W TDP
- IPMI

Test#1 - Power Mode

Power Data acquisition methodology

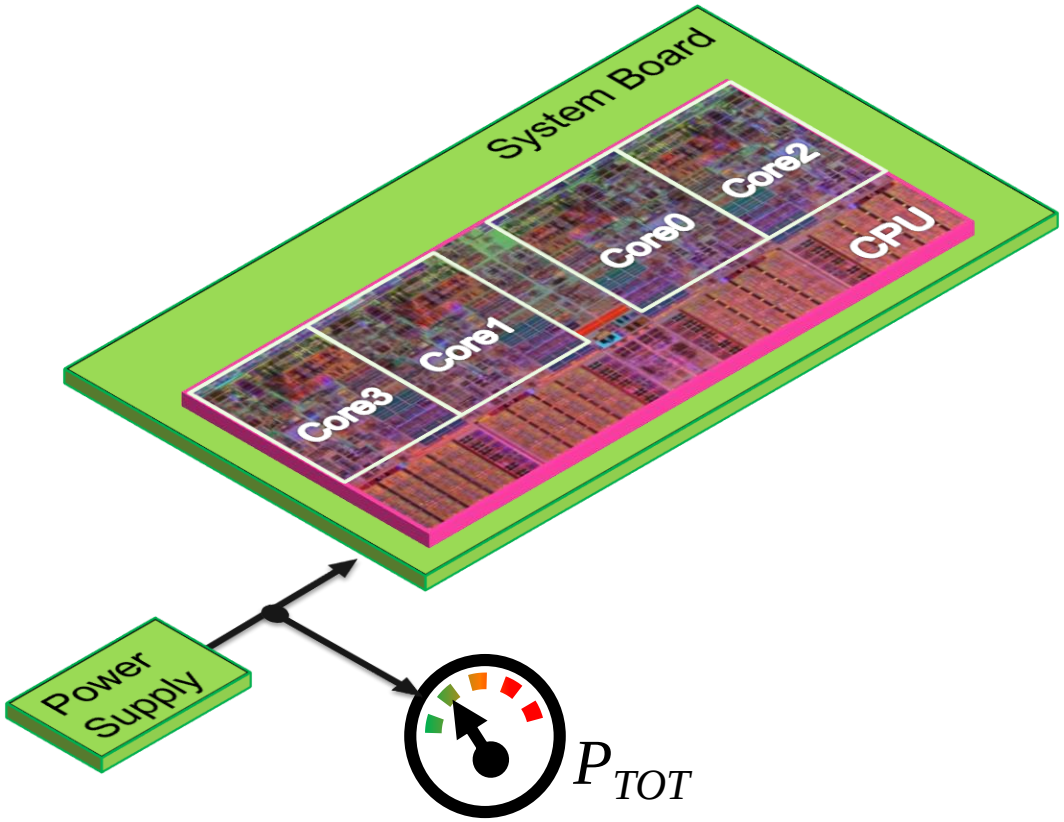

- Real multicore platform (4 core per CPU)
- System level power measurements only
- Benchmark based data profiling

Test#1 - Power Mode

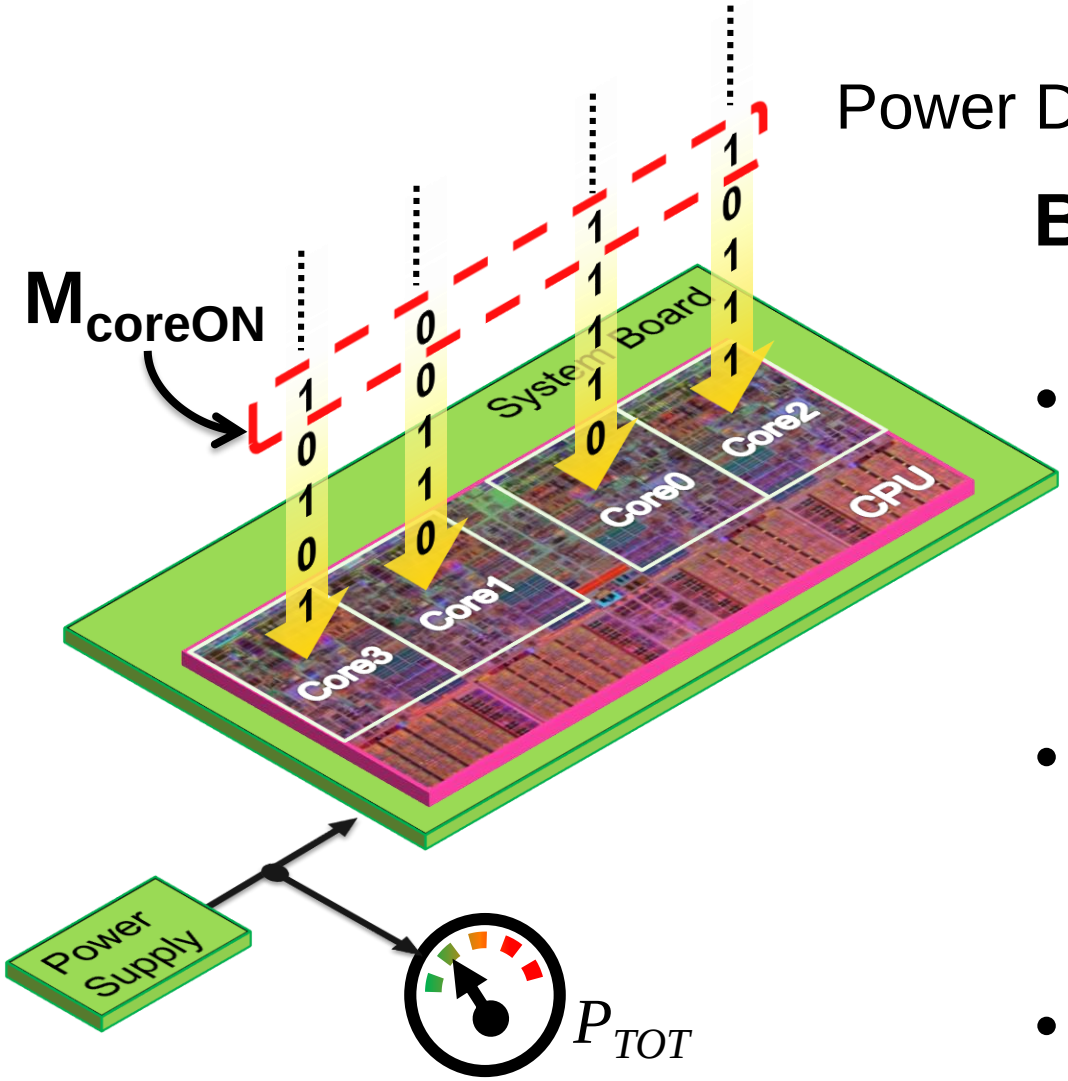

Power Data acquisition methodology

Benchmark: Per core Workload allocation

- "1" = Core in fully busy state = Maximum core power consumption (PowerVirus process)
- "0" = Idle state = minimum core power consumption
- M_{coreON} = #Active cores

Test#1 - Power Mod

Power Data acquisition steps

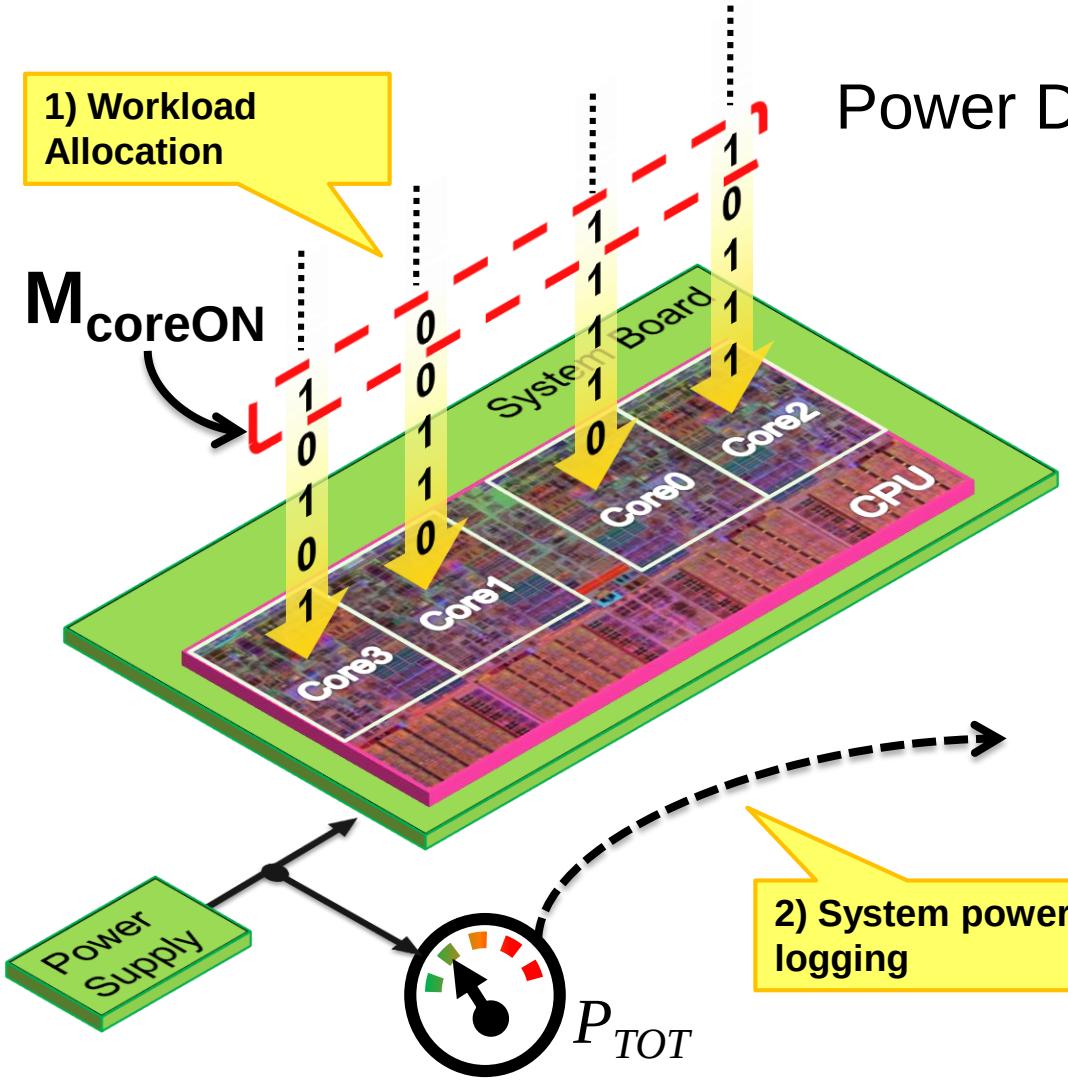

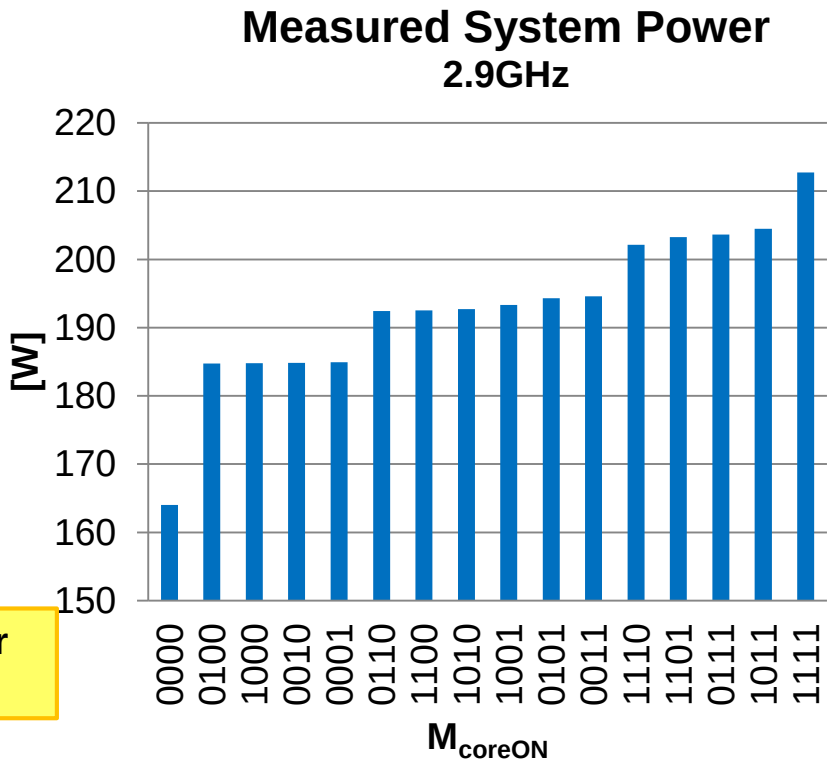

Test#1 - Power Mod

Power Data acquisition steps

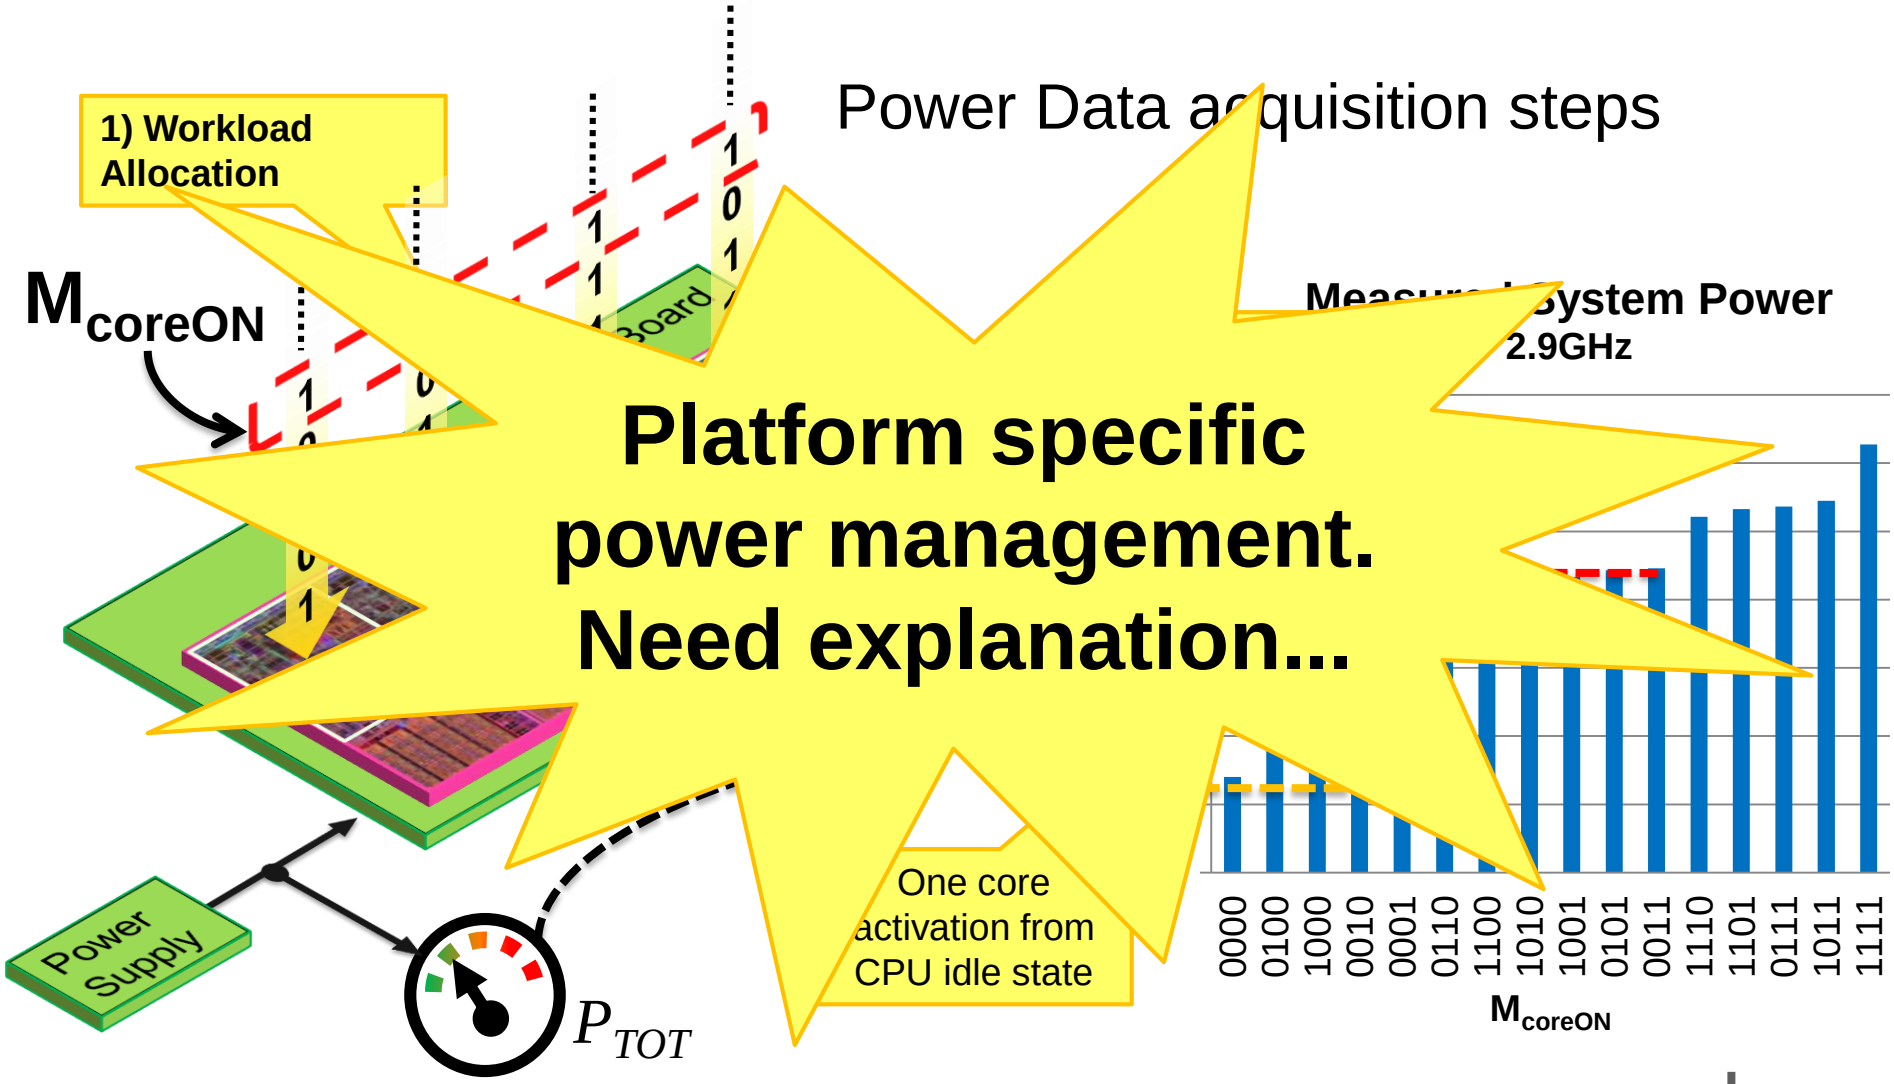

PM: breakthrough

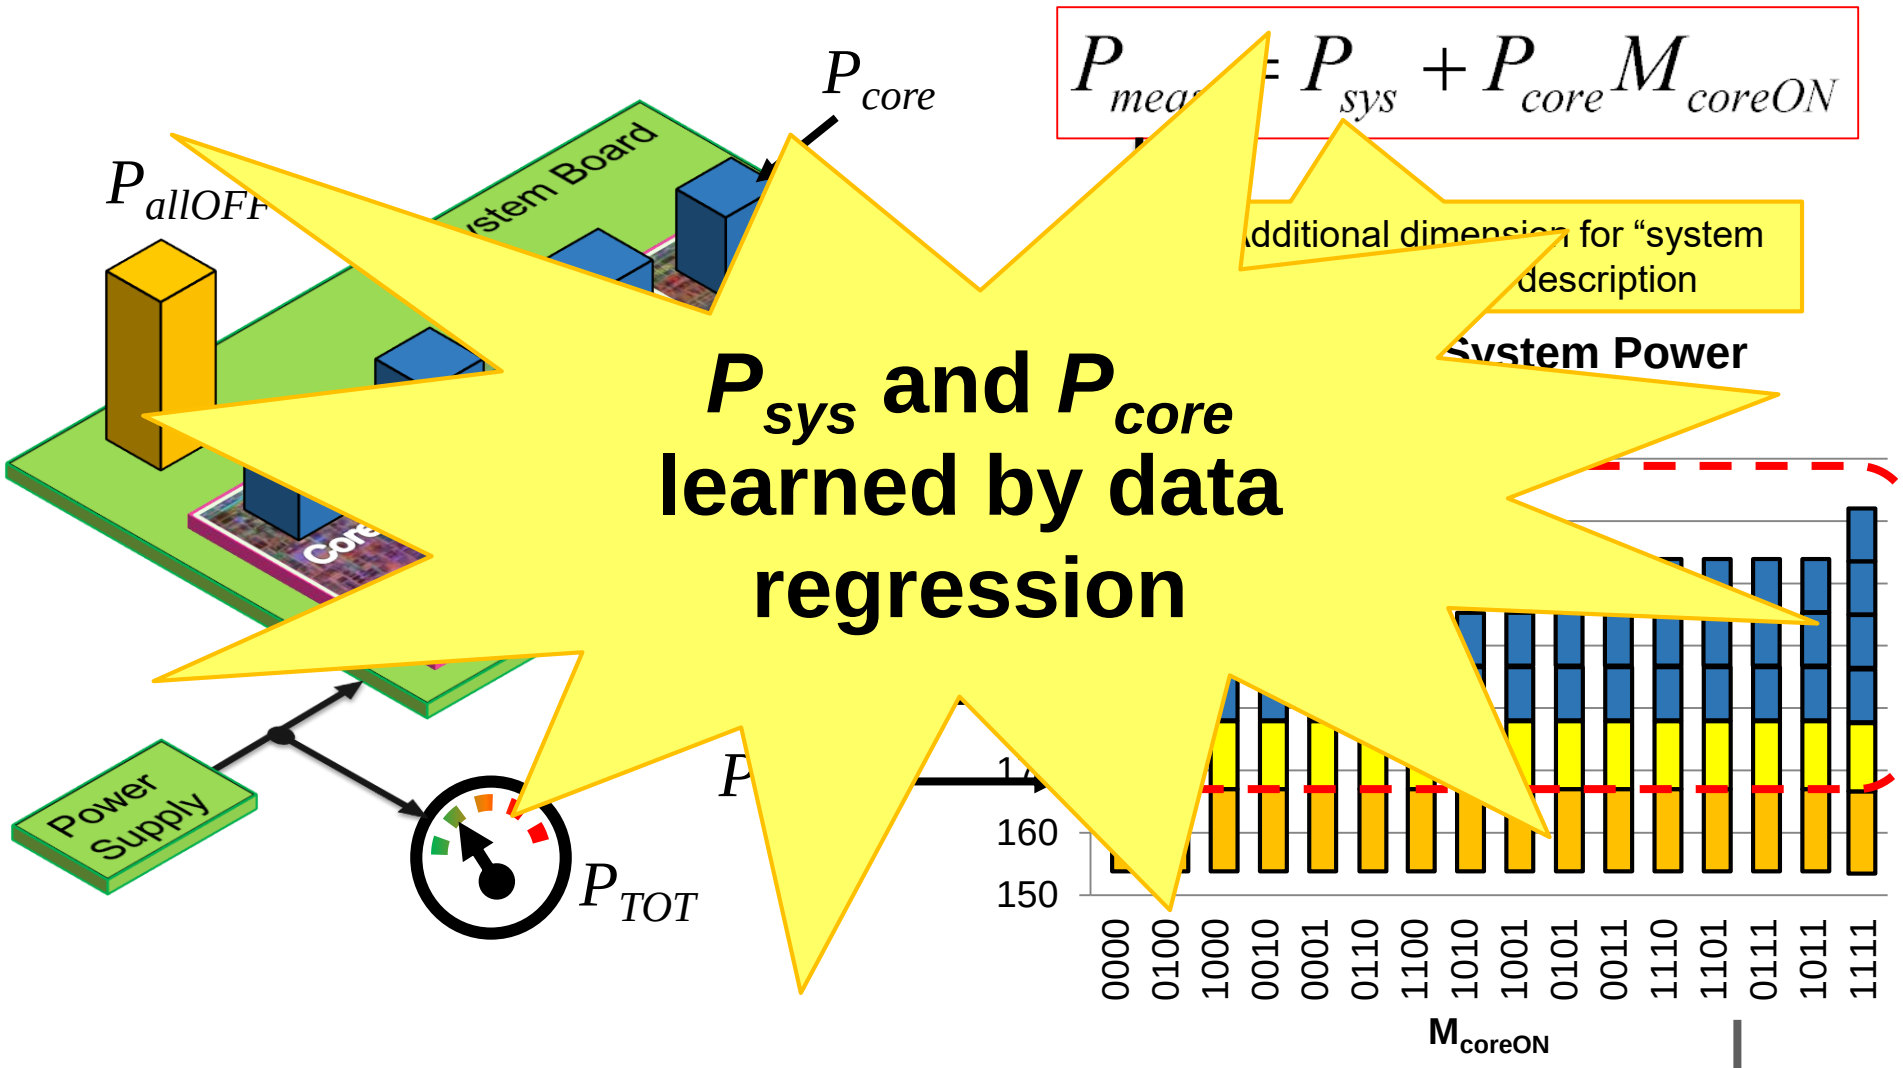

Test#2: Thermal Model

Thermal data acquisition methodology

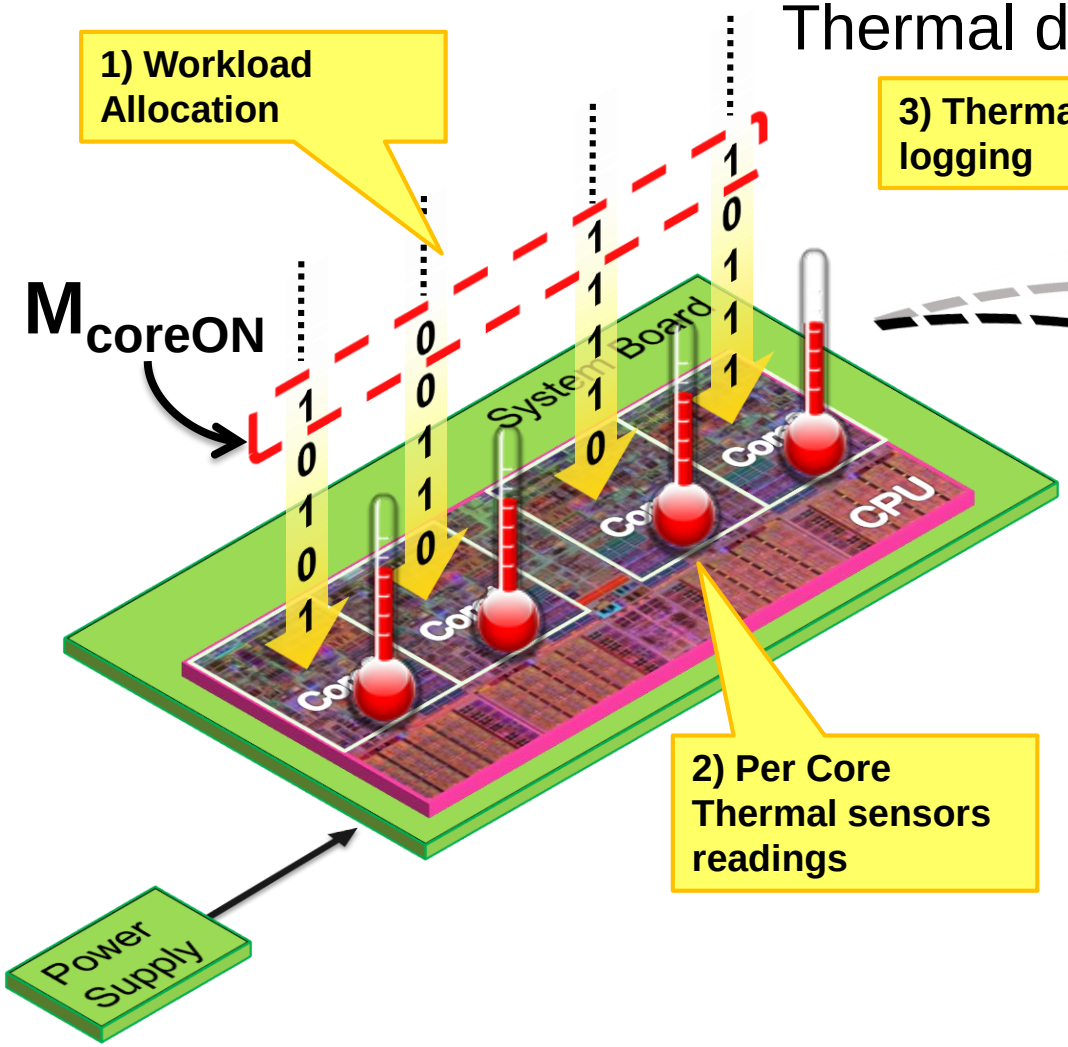

3) Thermal data logging

Measured cores temperatures (°C)

M_{coreON}	core2	core0	core1	core3
0 0 0 0	28,96	28,12	27,97	26,94
0 1 0 0	33,03	39,71	32,02	31,17
0 0 1 0	33,00	32,81	39,32	32,00
1 0 0 0	40,01	34,27	32,08	31,82
0 0 0 1	32,95	33,18	33,49	39,00
1 0 1 0	41,10	36,09	41,20	33,98
0 0 1 1	34,35	35,00	42,93	41,11
1 1 0 0	42,54	43,01	34,75	33,66
0 1 1 0	35,89	42,01	42,01	34,26
1 0 0 1	41,77	36,13	36,45	40,94
0 1 0 1	35,82	42,00	36,70	40,89
0 1 1 1	37,20	44,06	45,36	43,45
1 0 1 1	44,05	38,54	45,75	43,99
1 1 0 1	44,87	45,81	38,90	43,00
1 1 1 0	45,14	46,35	44,70	37,04
1 1 1 1	47,48	48,53	48,57	46,81

TM: Data Fitting

$$T = T_{abs} - T_{allOFF} = \begin{cases} 0, & M_{coreON} = '0000' \\ SG \cdot P, & otherwise \end{cases}$$

Measured cores temperatures (°C)

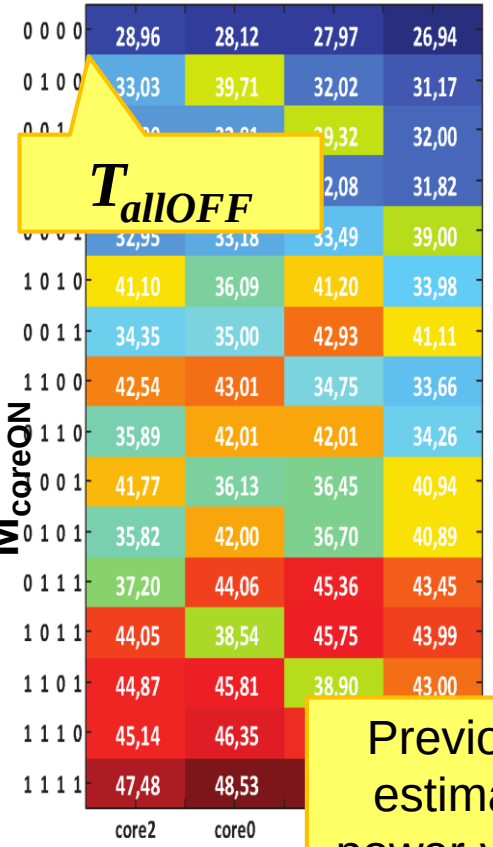

Previously estimated power values

$$\underbrace{[T]_{4\times16}}_{y} = \underbrace{[SG]_{4\times5}}_{\beta} \cdot \underbrace{[P]_{5\times16}}_{X}$$

$$y = \beta X + \varepsilon$$

Standard linear model fitting

where:

$$P = \begin{bmatrix} P_{core} & 0 & . & . & 0 \\ 0 & . & & & . \\ . & & . & & . \\ . & & & P_{core} & 0 \\ 0 & & & 0 & P_{sys} \end{bmatrix} \cdot \begin{bmatrix} 0 & 1 & . & . & 1 \\ 0 & 1 & . & . & 1 \\ 0 & 0 & . & . & 1 \\ 0 & 0 & . & . & 1 \\ 0 & 0 & . & . & 1 \end{bmatrix}_{5\times16}$$

Results

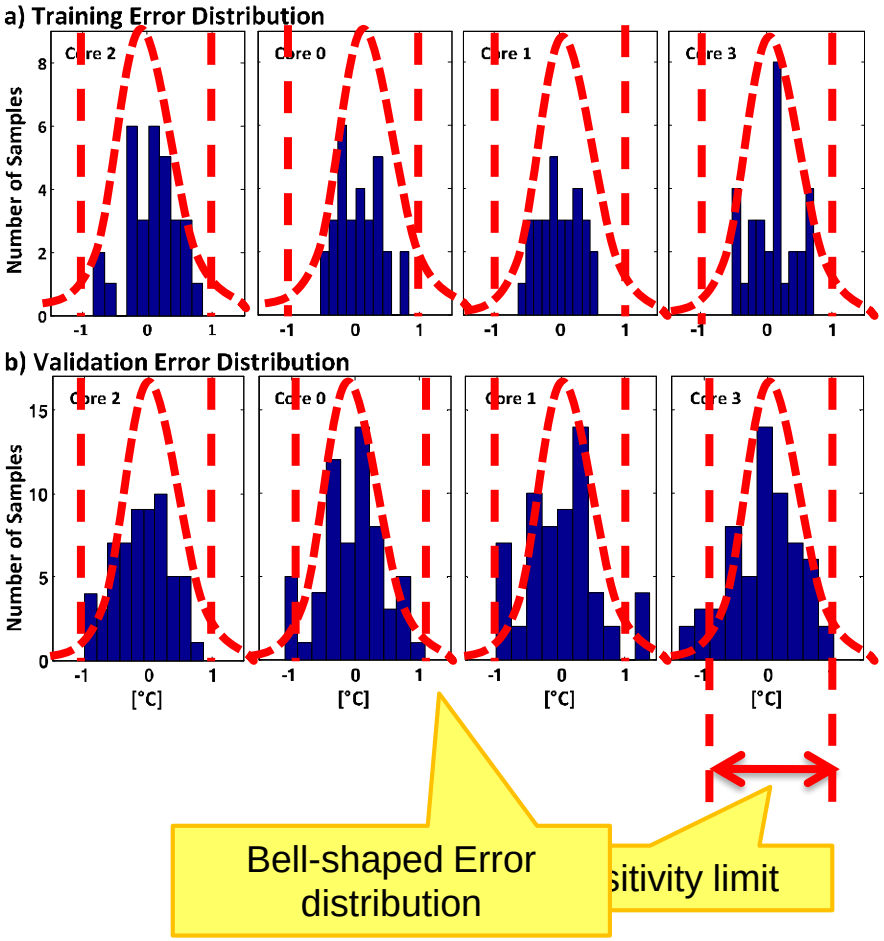

- Extremely accurate in-sample fitting
- Out of sample error below 1.4°C

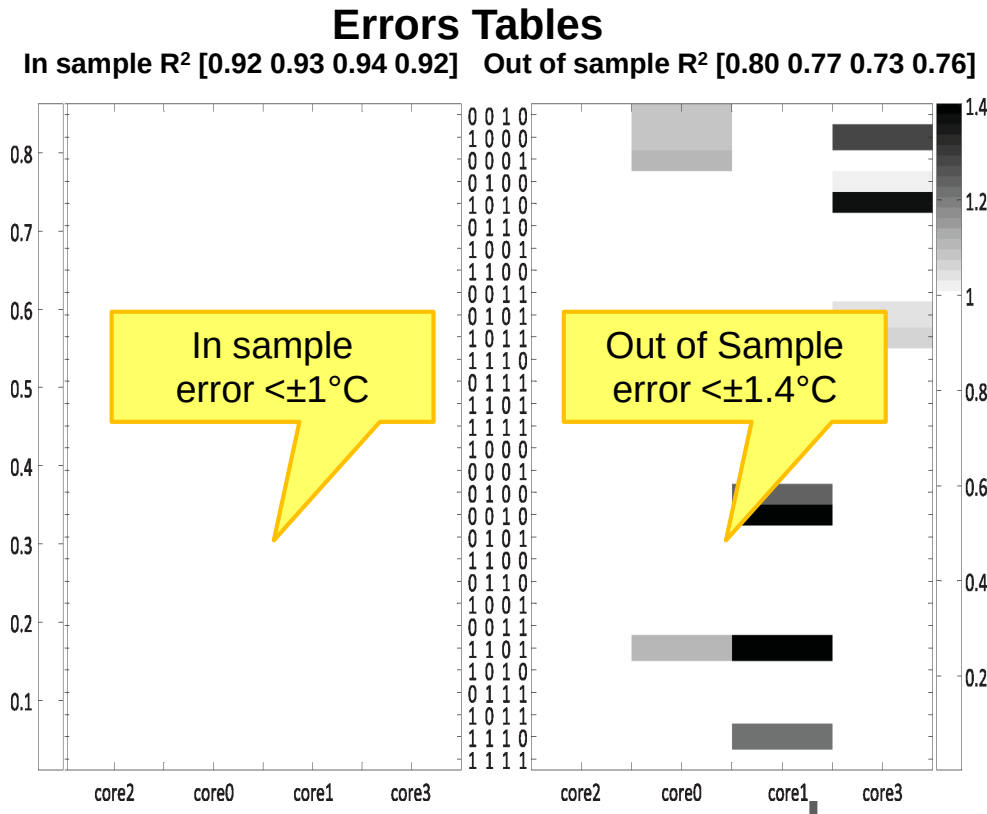

Layout

- Motivation
- Static Thermal Modelling
- **Dynamic Thermal Modelling**
- Scalable Thermal Modelling
- Thermal Management
- Datacenter automation

Model Structure

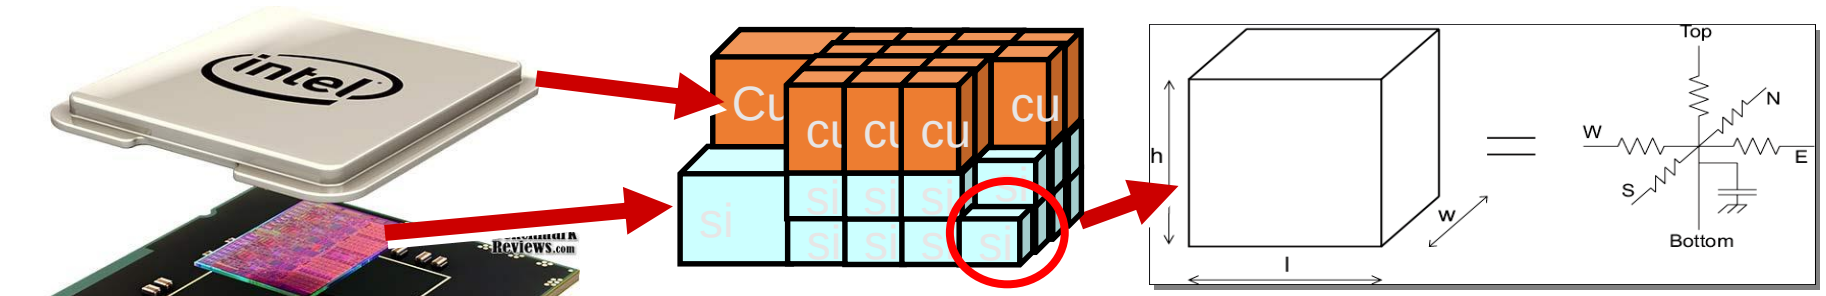

$$x_{n+1} = Ax_n + B\,P_{n,j}$$

$$T_{n,j} = Cx_n$$

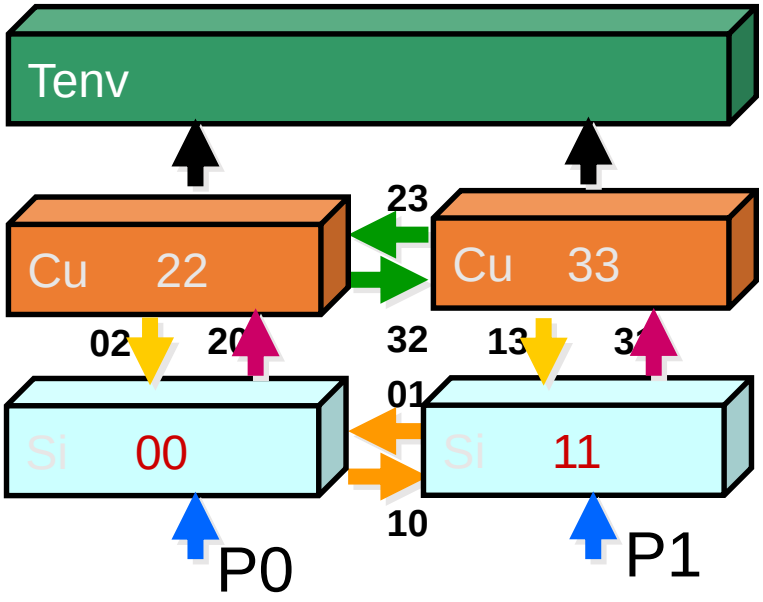

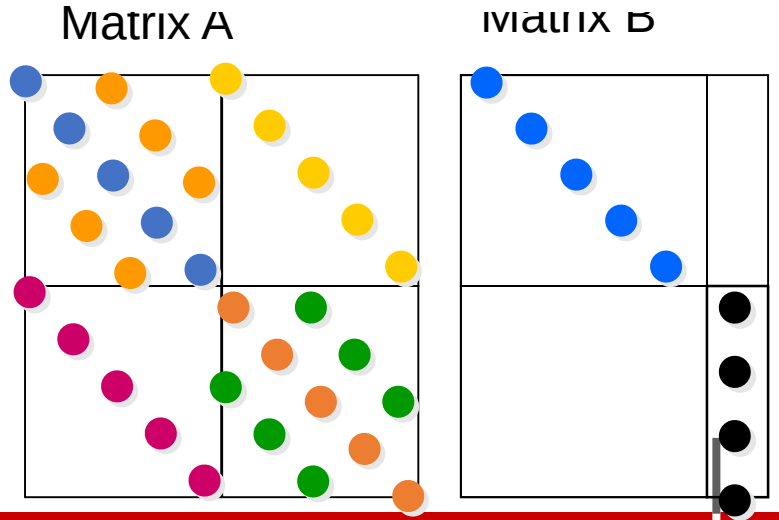

LS System Identification

Test pattern

System response

Data → **Model least square optimization** → **Optimal parameters**

A,B

$$\theta_{opt} \rightarrow \min_{\theta} I(\theta)$$ Parametric optimization

$$\theta = \{A, B\}$$ Parameters

$$I(\theta) = \frac{1}{N_s} \sum_{i=1}^{N_s} e_i^2$$ Cost function

$$e = T_{model} - T_{real}$$ Error Function

Results - Quantization

Thermal Modeling

i7 Server Platform – 4 cores

Ts = **1ms** - **Quantizzation noise**

Step response: $y(t) = K_0 - \sum_{i=1}^{n} k_i \cdot e^{-p_i \cdot t}$

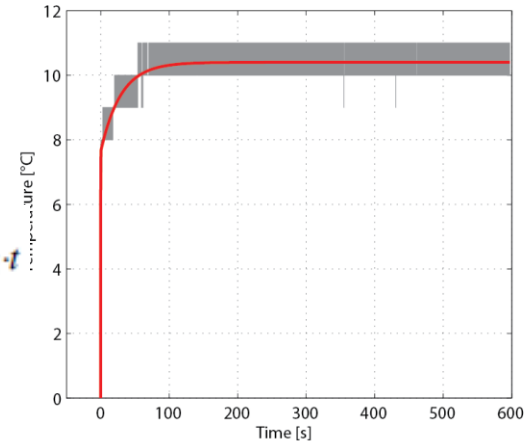

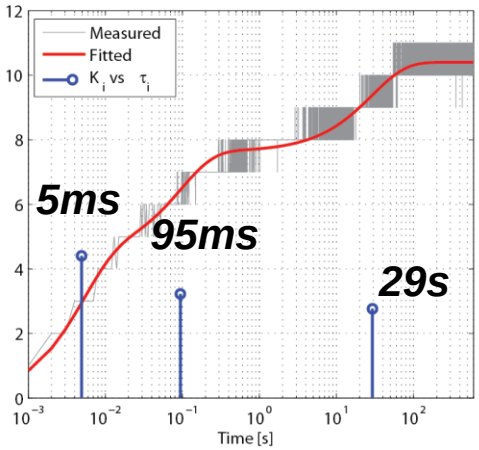

Results - Quantization

Thermal Modeling

i7 Server Platform – 4 cores

Ts = **1ms** - **Quantizzation noise**

Step response: $y(t) = K_0 - \sum^{n} k_i \cdot e^{-p_i \cdot t}$

Black Box

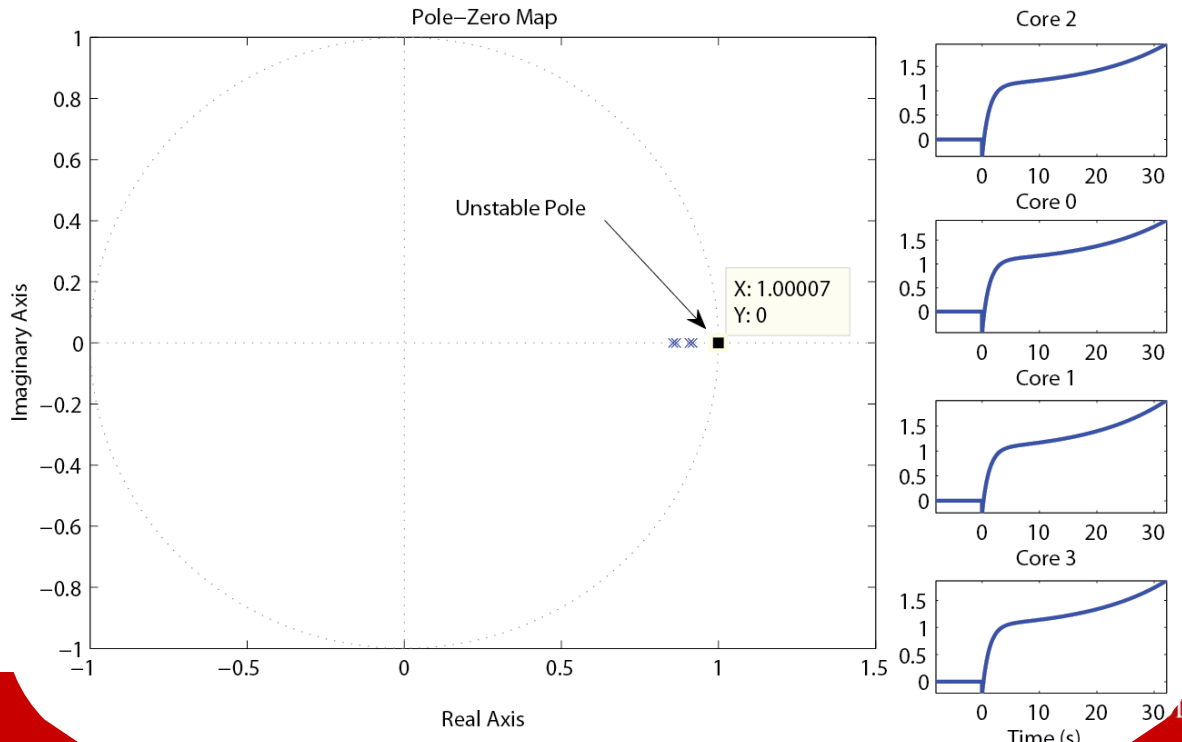

Results - Quantization

Thermal Modeling

i7 Server Platform – 4 cores

Ts = **1ms** - **Quantizzation noise**

Step response: $y(t) = K_0 - \sum^{n} k_i \cdot e^{-p_i \cdot t}$

Black Box Model Estimation Performance

Order	1	2	3
SSE	1,02E+05	2,00E+04	2,59E+04
Time elapsed (s)	300,15	2434,24	12681,57

Black Box

Our approach
LS + physical constraints

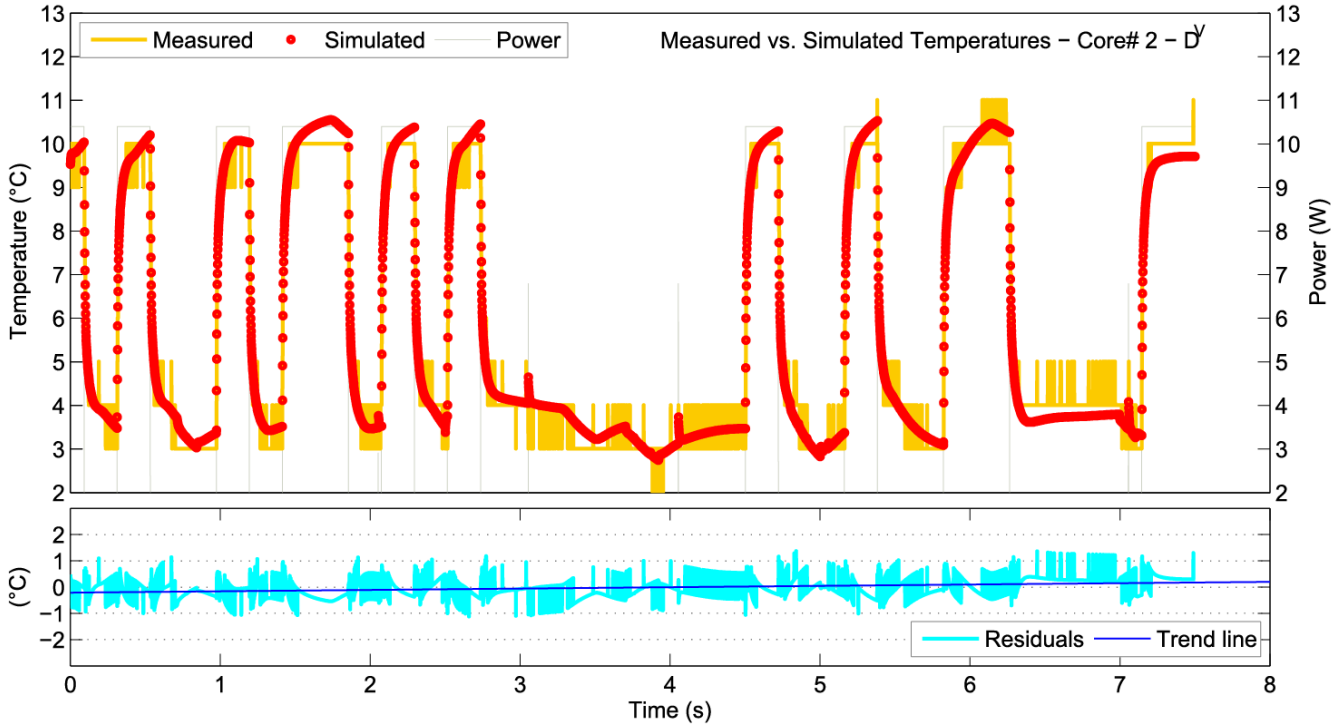

Thermal Model Identification

Single chip cloud computer SCC – 48 cores – 24 tiles

Ts = **100ms** - **Measurment noise**

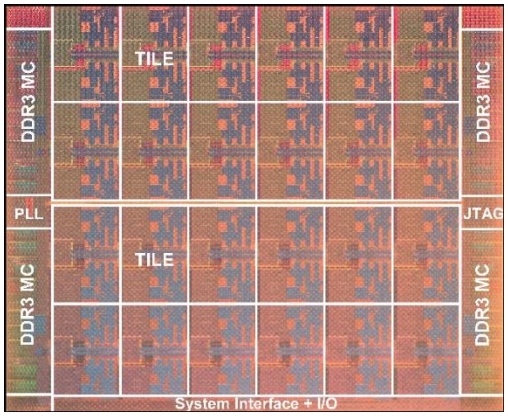

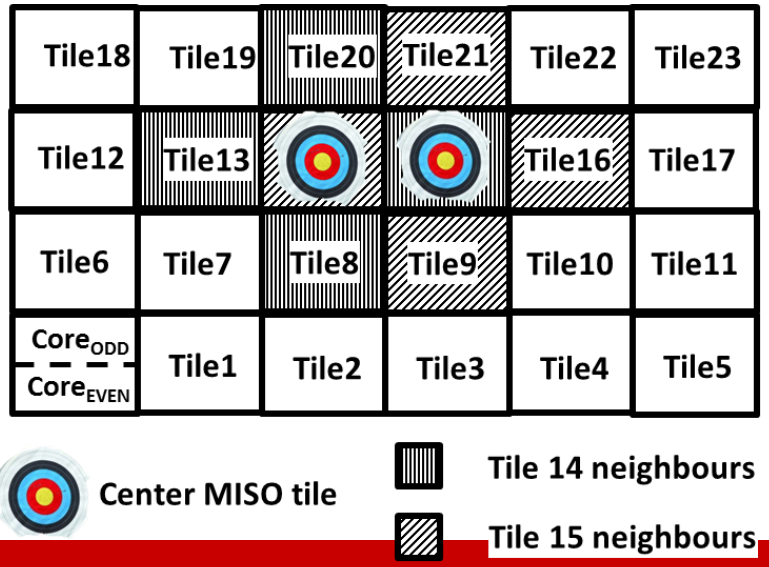

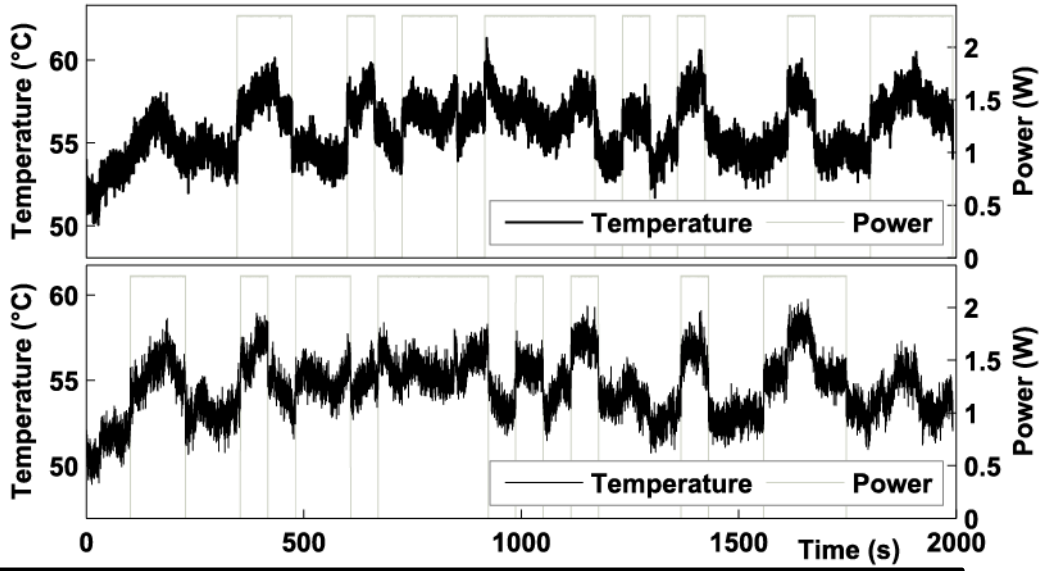

Intel SCC – tile based architecture:
- 2 cores per tile
- thermal sensors per tile (ring oscillators)
- Frequency can be scaled per tile.

MISO Thermal Model Identification:
In: Tile's Power, xy neighbour Tile's temperatures
Out: Tile's Temperature

SID Standard ARX

ARX Solution:

$$A(z^{-1})T(t) = B(z$$

White noise

Uncorrelated

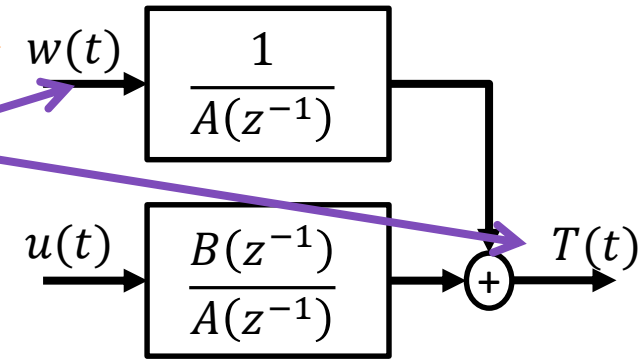

$$T(t) = -a_1 T(t-1) - \ldots - a_n T(t-n) + b_{11} u_1(t-1) + \ldots + b_{1n} u_1(t-n) + \ldots + b_{r1} u_r(t-1) + \ldots + b_{rn} u_r(t-n)$$

Measured

Unknown

$$T(t) = \varphi^T(t)\theta + w(t)$$

Linear composition of the model parameters and measured data

Best Prediction => White residual

$$w(t) = T(t) - \varphi^T(t)\theta$$

Thermal Model Identification

Single chip cloud computer SCC – 48 cores – 24 tiles

Ts = **100ms** - **Measurment noise**

Standard ARX:

- **Designed only for process noise!**
- **Measurement noise induces biases !**

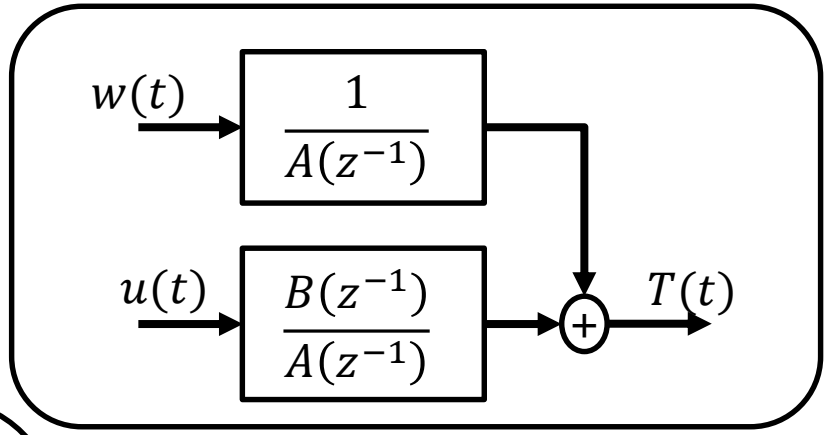

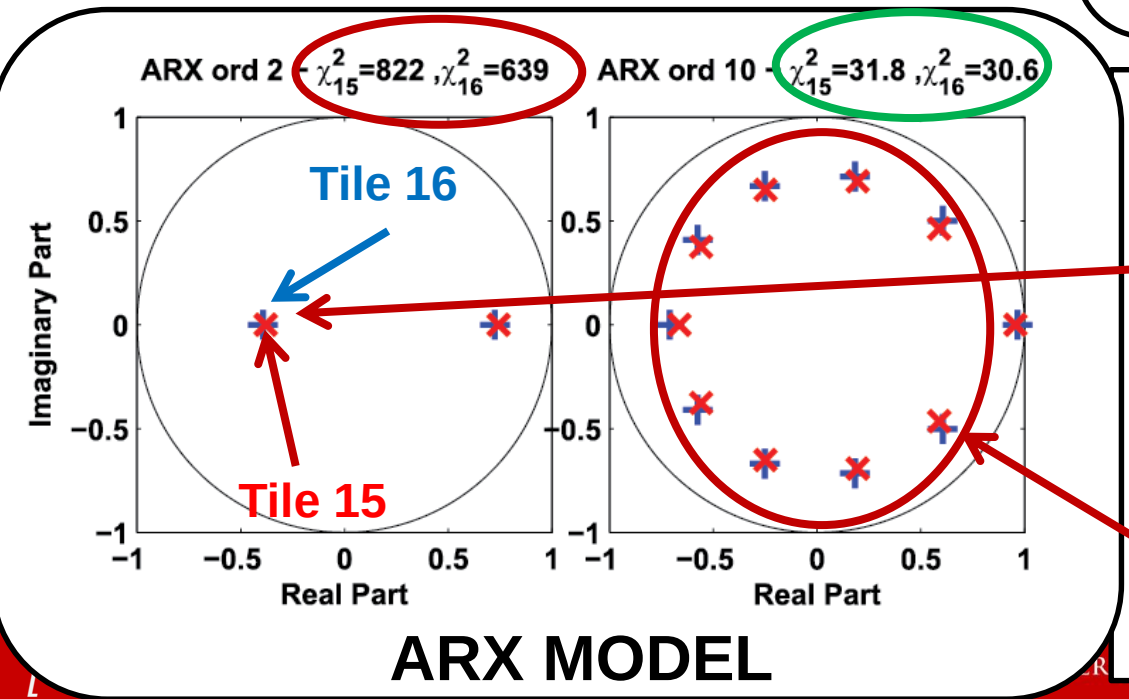

ARX MODEL

ARX - 2° Order :

- χ^2 *failed*
- **Negative pole:**
  - **Not realistic**

ARX - 10° Order :

- χ^2 *passed*
- **Negative & Complex poles**
  - **Not realistic**

Thermal Model Identification

Why ARX is failing?

- **Standard ARX considers white innovation**
- **Its impact on the output temperature get coloured by the model poles**
- **In our system we know we have a white additive noise in the temperature sensors**
- **Two different noise sources**
- **Trying to model it with just the innovation noise source biases the model parameters**
- **Noise in the output cannot just be filtered out as it is unknown – cannot power down a CPU and access the internal temperature sensors**

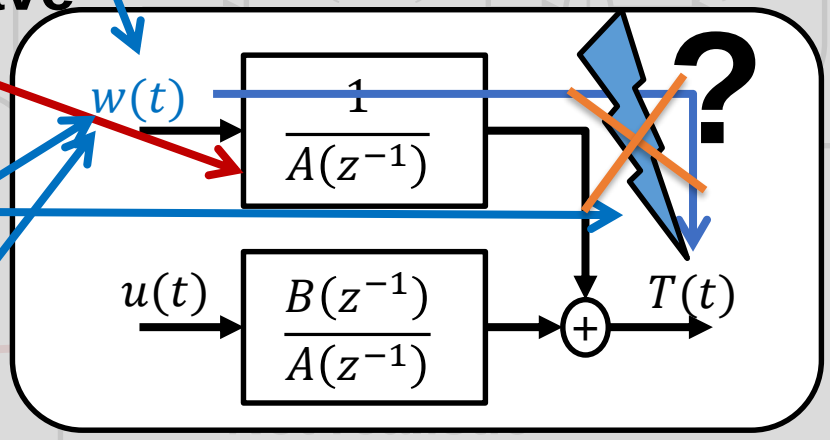

Thermal Model Identification

Why ARX is fa

- **St**
- **I** **by th**
- **In our sy** **a white** **additive** **the temperature** **sensors**
- **Two different noise sources**
- **Trying to model it with just the innovation noise source biases the model parameters**
- **Noise in the output cannot just be filtered out as it is unknown – cannot power down a CPU and access the internal temperature sensors**

OUR SOLUTION:
Directly model the output noise in the ARX problem formulation
By modelling it we can estimate its impact during parameter learning

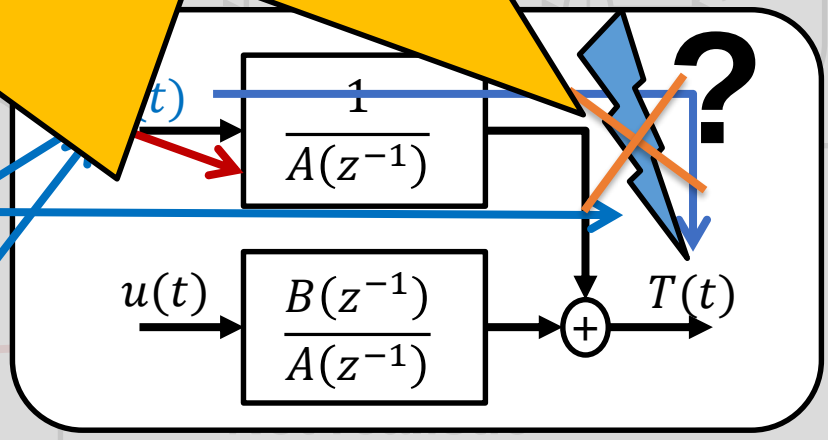

SID ARX+NOISE

Noisy Temperature Sensors

- $T(t) = \bar{T}(t) + v(t)$

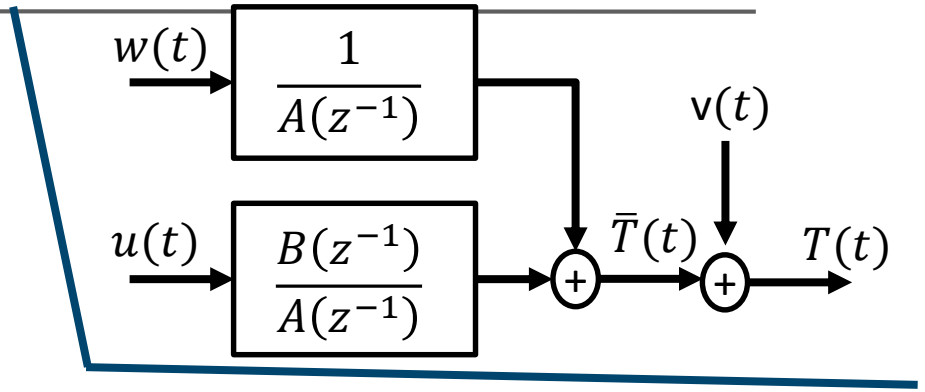

SID ARX+NOISE

Noisy Temperature Sensors

- $T(t) = \bar{T}(t) + v(t)$

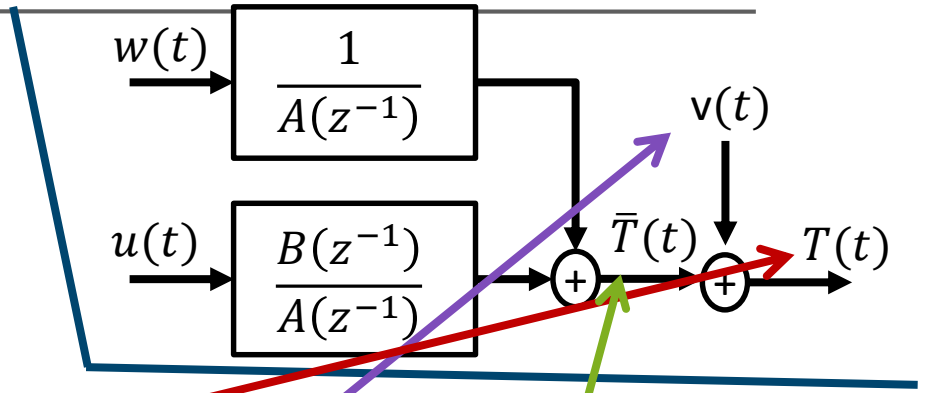

Measured Temperature

Measurement Noise

Real Temperature

SID ARX+NOISE

Noisy Temperature Sensors

- $T(t) = \bar{T}(t) + v(t)$

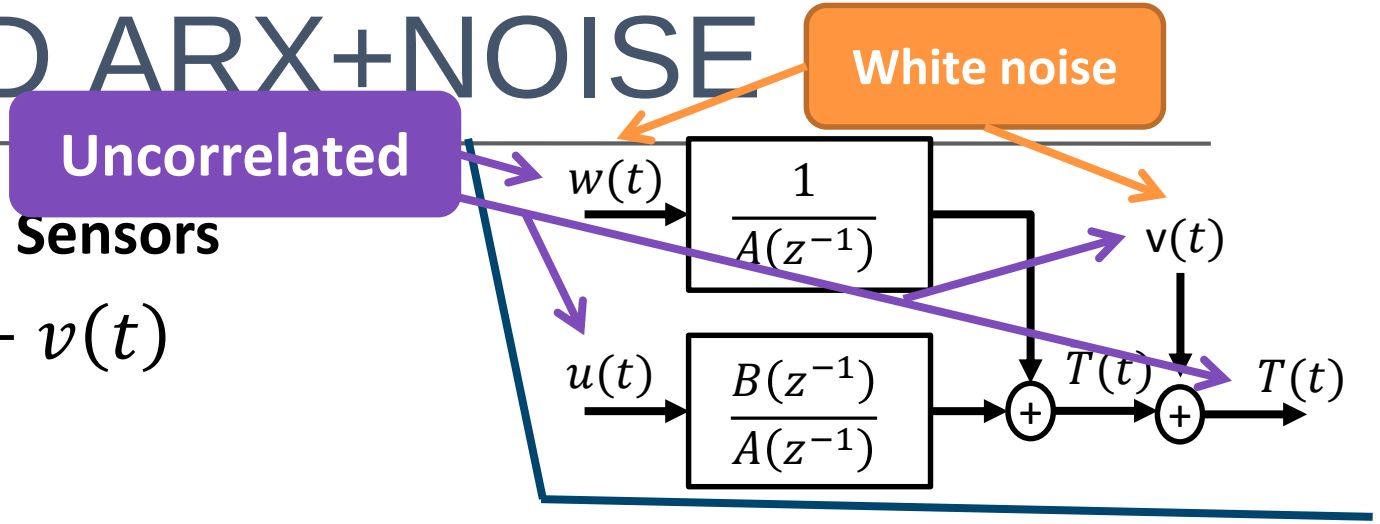

SID ARX+NOISE

Noisy Temperature Sensors

- $T(t) = \bar{T}(t) + v(t)$

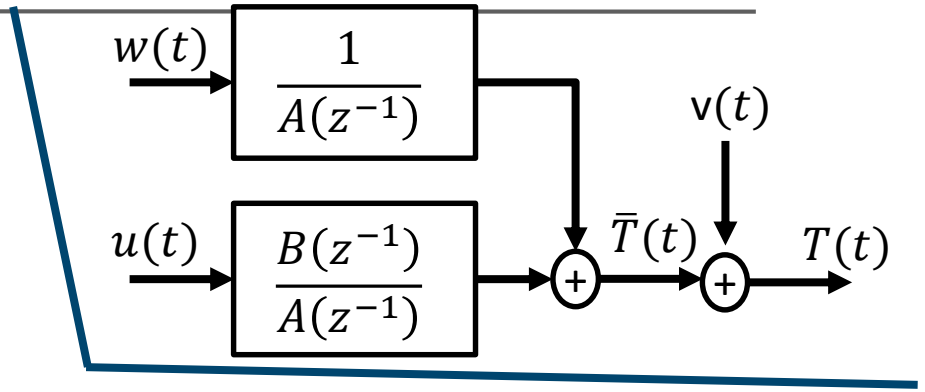

$$\bar{T}(t) = -a_1\bar{T}(t-1) - \ldots - a_n\bar{T}(t-n) + b_{11}u_1(t-1) + \ldots + b_{1n}u_1(t-n) + \ldots + b_{r1}u_r(t-1) + \ldots + b_{rn}u_r(t-n)$$

Real **Unknown**

$$\bar{T}(t) = \bar{\varphi}^T(t)\theta + w(t)$$

$$\varphi(t) = \bar{\varphi}(t) + \tilde{\varphi}(t)$$

Linear composition of the model parameters and real temperature/input

SID ARX+NOISE

Noisy Temperature Sensors

- $T(t) = \bar{T}(t) + v(t)$

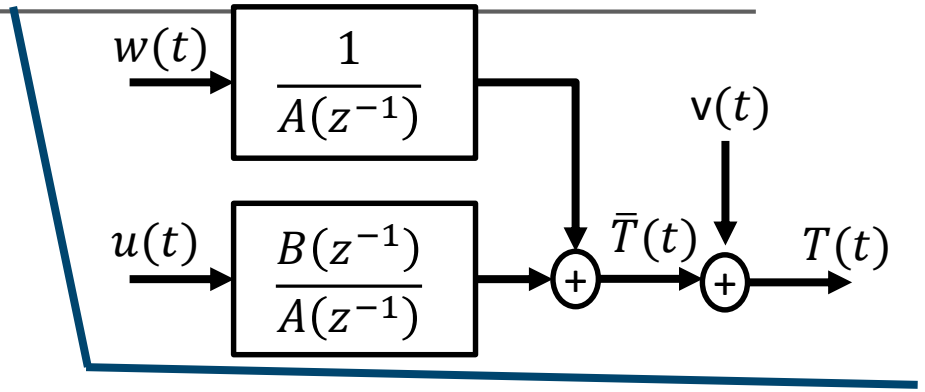

$$\bar{T}(t) = -a_1\bar{T}(t-1) - \ldots - a_n\bar{T}(t-n) + b_{11}u_1(t-1) + \ldots + b_{1n}u_1(t-n) + \ldots + b_{r1}u_r(t-1) + \ldots + b_{rn}u_r(t-n)$$

$$\bar{T}(t) = \bar{\varphi}^T(t)\theta + w(t)$$

$$\varphi(t) = \bar{\varphi}(t) + \tilde{\varphi}(t)$$

$$\bar{\varphi}(t) = [-\bar{T}(t-1)\ldots-\bar{T}(t-n)u_1(t-1)\ldots u_r(t-n)]^T$$

$$\varphi(t) = [-T(t-1)\ldots-T(t-n)u_1(t-1)\ldots u_r(t-n)]^T$$

$$\tilde{\varphi}(t) = [-v(t-1)\ldots-v(t-n)0\ldots0]^T$$

rn

$$T(t) = \varphi^T(t)\theta + e(t)$$

$$e(t) = v(t) - \tilde{\varphi}^T(t)\theta + w(t)$$

Best Prediction => Residual not white $e(t) = T(t) - \varphi^T(t)\theta$

Thermal Model Identification

Single chip cloud computer SCC – 48 cores – 24 tiles

Ts = **100ms** - **Measurment noise**

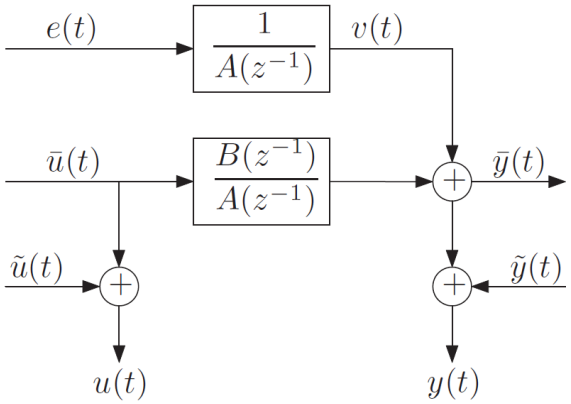

Standard ARX:

- **Designed only for process noise !**
- **Measurement noise induces biases !**

Bias Compensated ARX

(Diversi et al., 2013a, 2014)

- ***Iteratively estimate the noise variance and compensate it in the LS***

Dynamic Frisch

(Diversi et al., 2013b)

- ***Searches solutions compatible with the covariance matrix of the noisy data***
- ***Selection based on a set of low-order and high-order Yule-Walker equations***

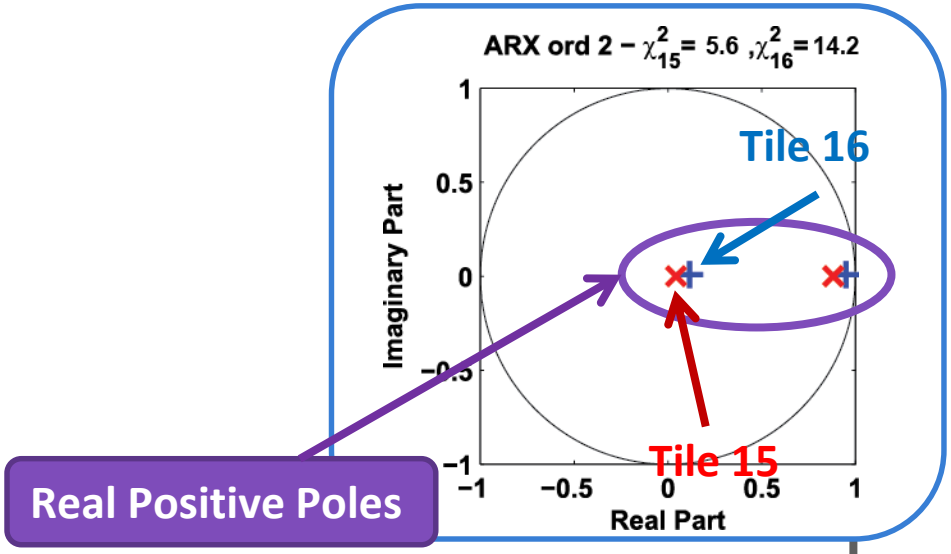

Thermal Model Identification

Sin

Ts

Sta

- ***D***
- ***M***

Bia

(Div

- ***I***
- ***v***

Dyna

(Diversi et al., 2013b)

- ***Searches solutions compatible with the covariance matrix of the noisy data***
- ***Selection based on a set of low-order and high-order Yule-Walker equations***

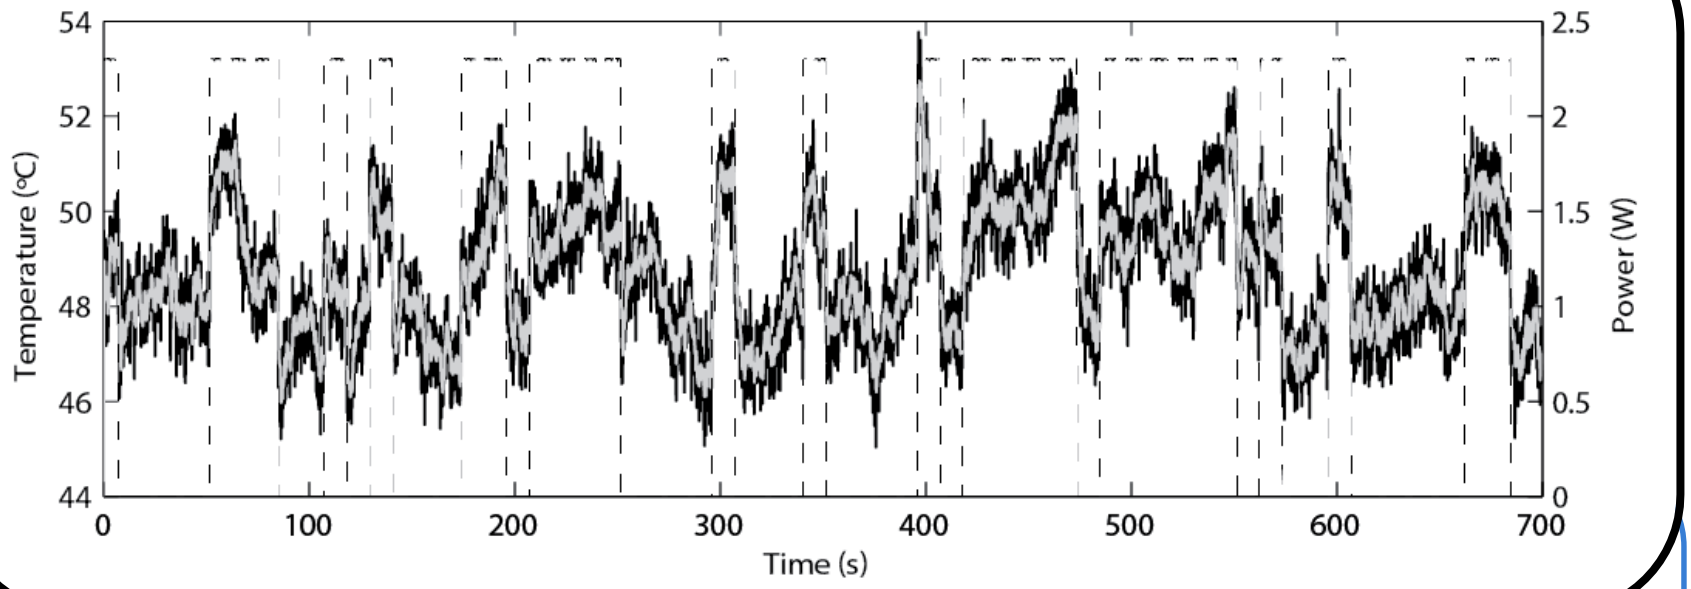

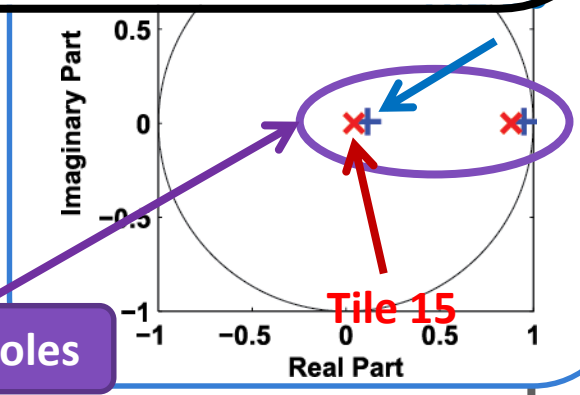

Real Positive Poles

Galileo Modelling

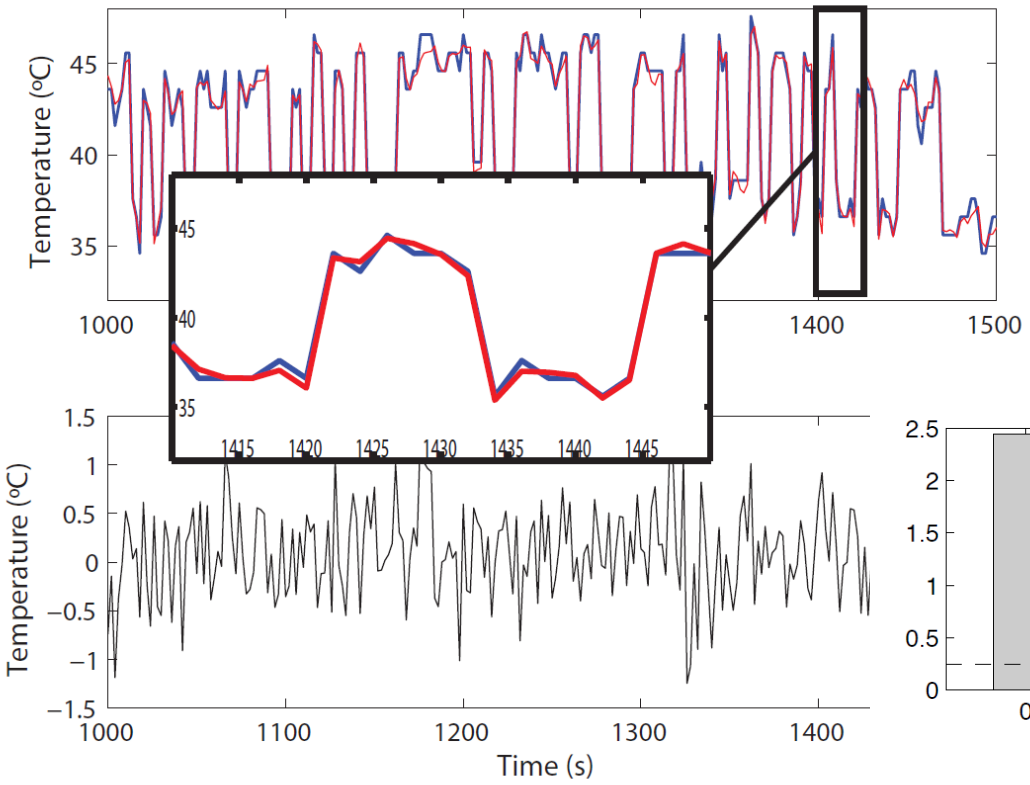

Very accurate prediction <1°C also with quantization noise (±1°C)

Gain w.r.t. core Id

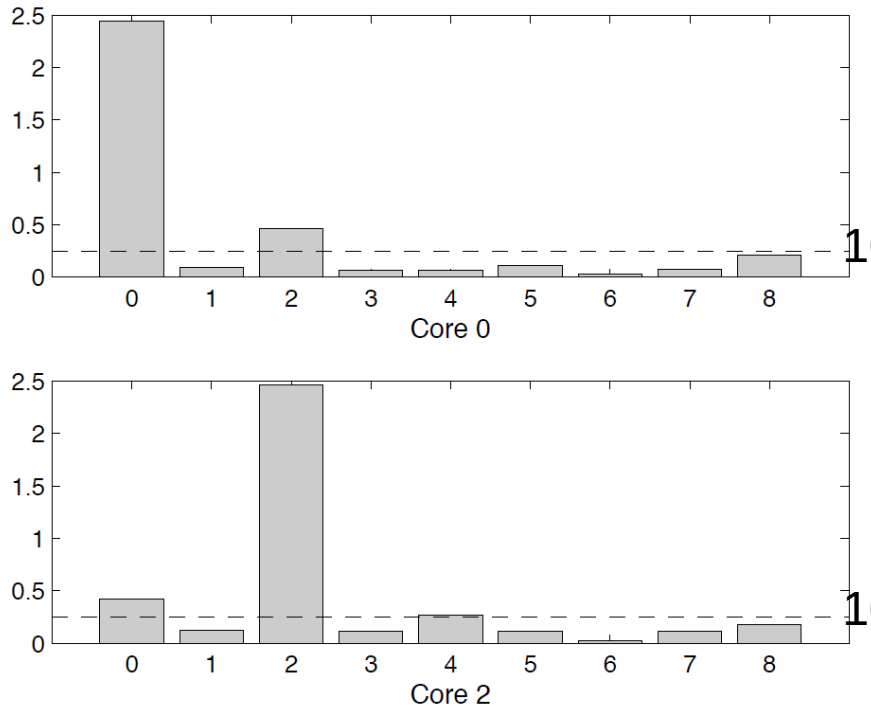

Thermal locality !!

Layout

- Motivation
- Static Thermal Modelling
- Dynamic Thermal Modelling
- **Thermal Management**
- Datacenter automation

Multitherman Holistic Approach

FP7 ERC Advance MULTITHERMAN:
Multiscale Thermal Management of Computing Systems
PI: Prof Luca Benini

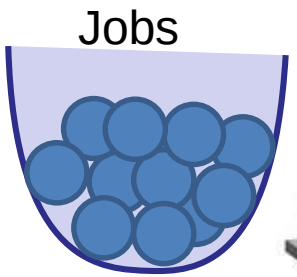

Timescale: ~msec, ~sec

Timescale: ~hours

Multitherman Holistic Approach

FP7 ERC Advance MULTITHERMAN:
Multiscale Thermal Management of Computing Systems
PI: Prof Luca Benini

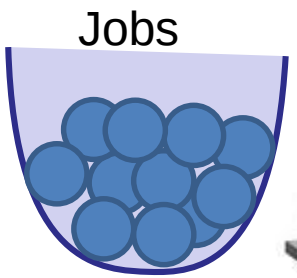

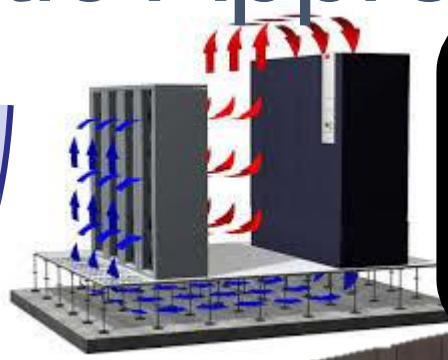

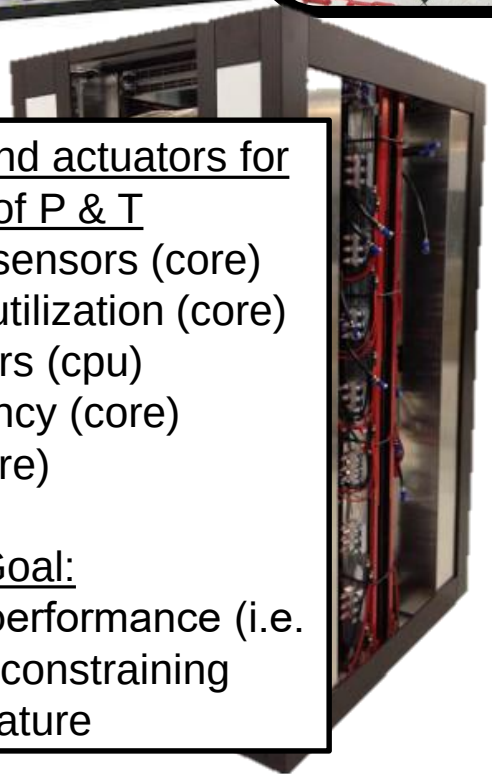

Optimal Set-Point

Chip Level Control

Knowledge

Timescale: ~msec, ~sec

Built-in sensors and actuators for feedback control of P & T

1. Temperature sensors (core)
2. Architectural utilization (core)
3. Power monitors (cpu)
4. Clock Frequency (core)
5. Shutdown (core)

Optimal Control Goal:
Maximize core's performance (i.e. frequency) while constraining maximum temperature

Timescale: ~hours

Multitherman Holistic Approach

FP7 ERC Advance MULTITHERMAN:
Multiscale Thermal Management of Computing Systems
PI: Prof Luca Benini

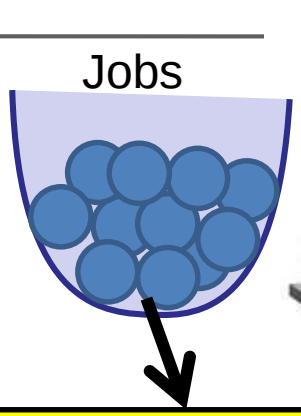

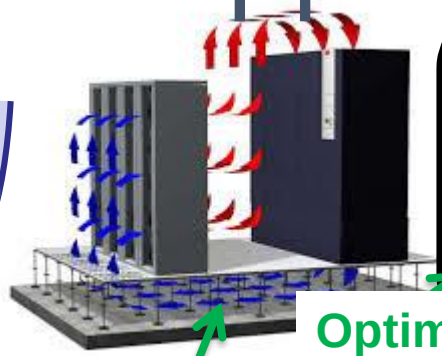

Optimal Set-Point

System Level Control

Knowledge

Timescale: ~hours

Room and chassis sensors and actuators:
1. Machine, Room, Coolant temperatures
2. Blowers, Pumps flows
3. Air conditioner, Chillers activation

Optimal Control Goal:
Global controller which modulates the cooling effort to prevent over-temperatures or excessive local throttling of nodes

Timescale: ~msec, ~sec

Multitherman Holistic Approach

FP7 ERC Advance MULTITHERMAN:
Multiscale Thermal Management of Computing Systems
PI: Prof Luca Benini

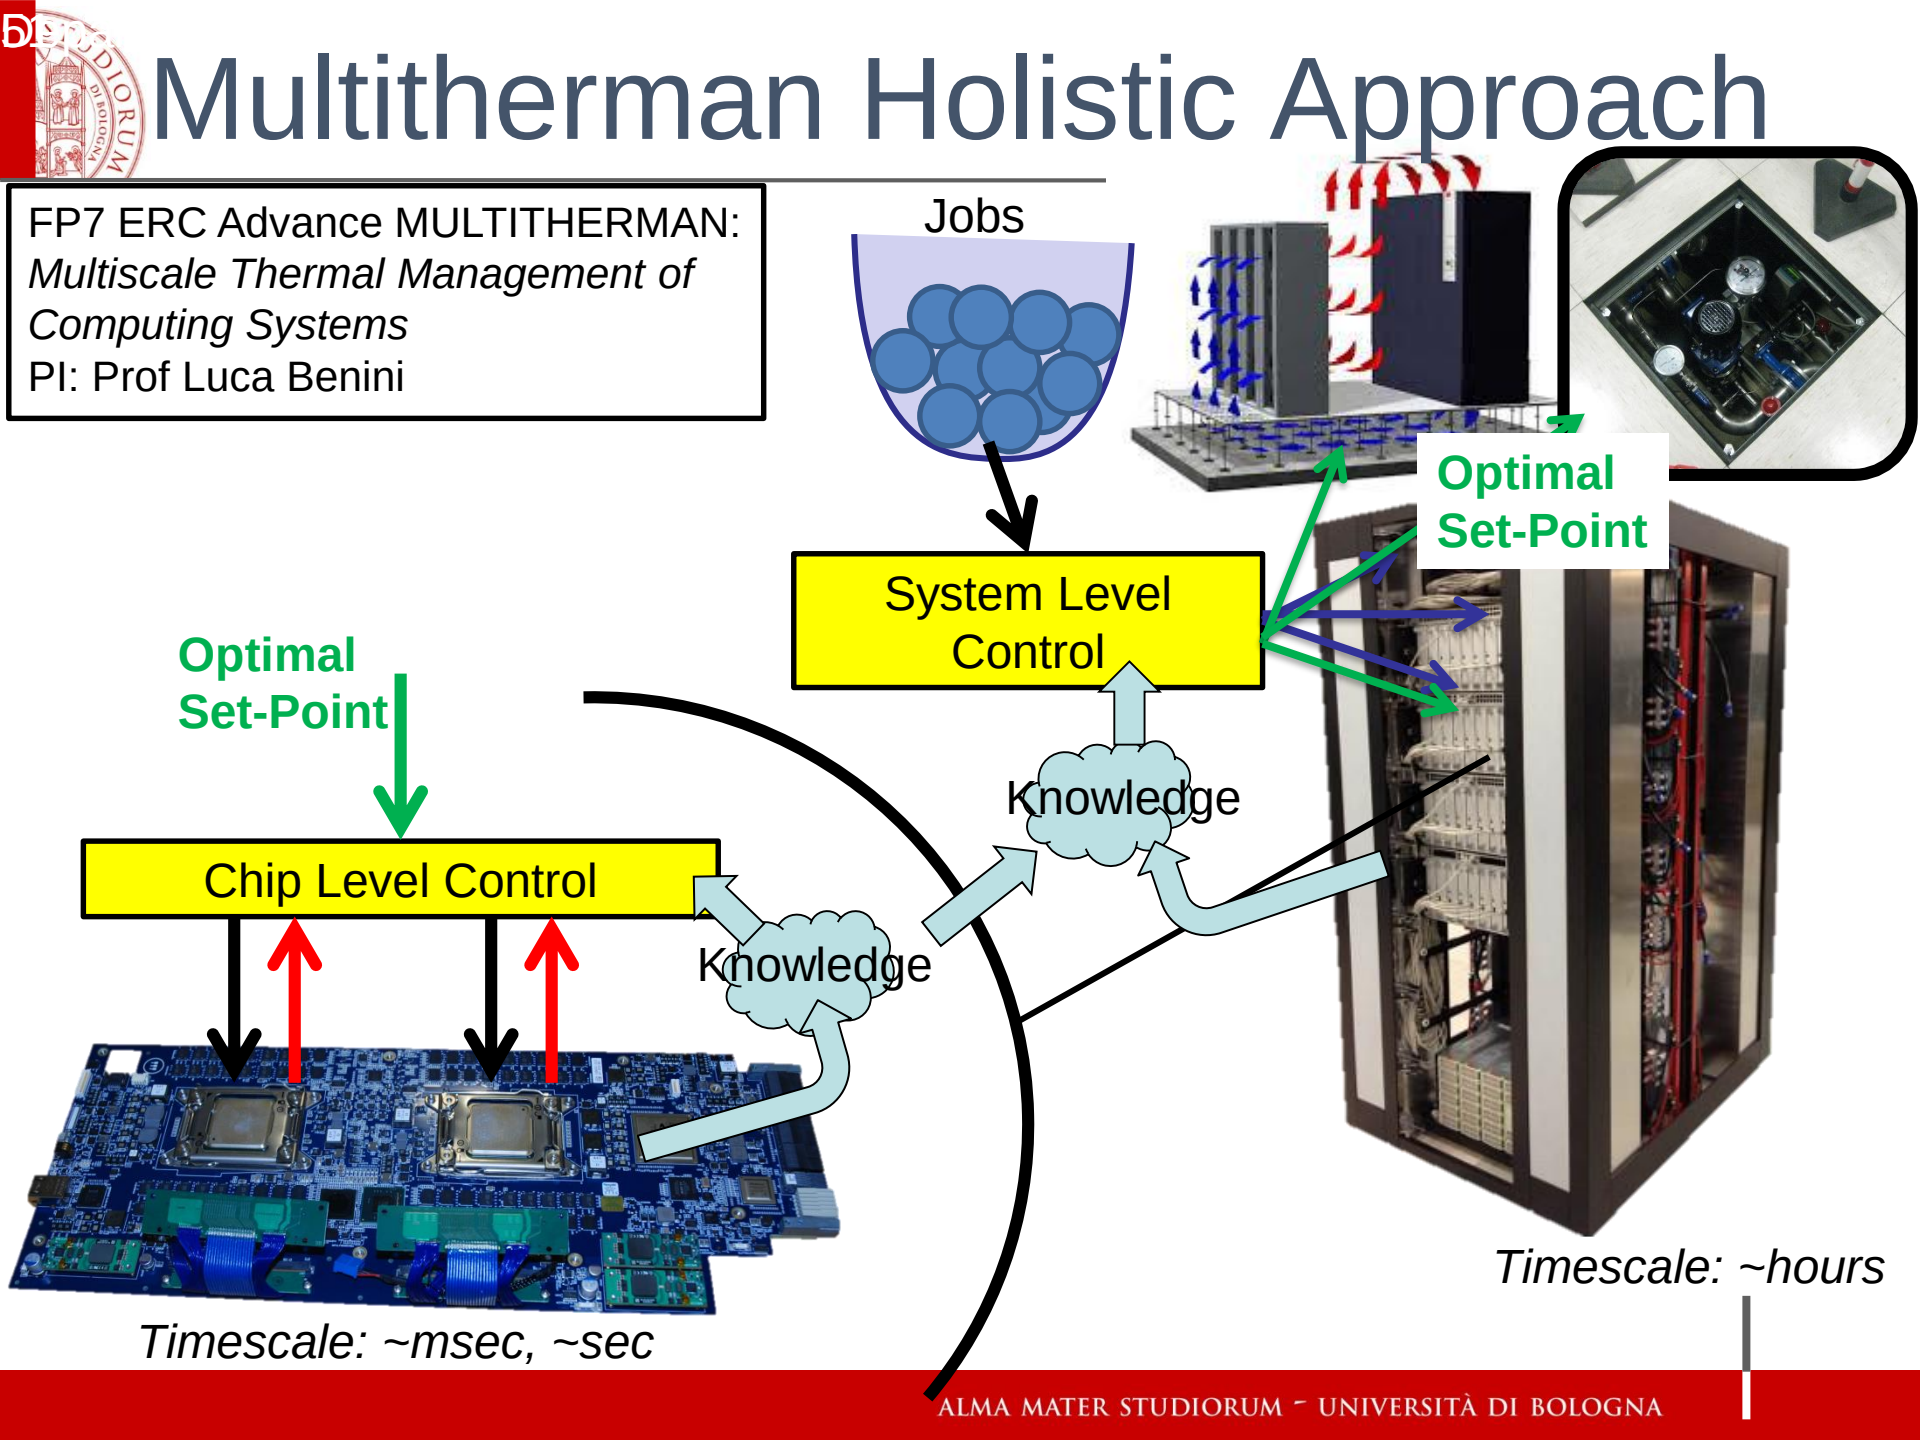

Chip Level Control

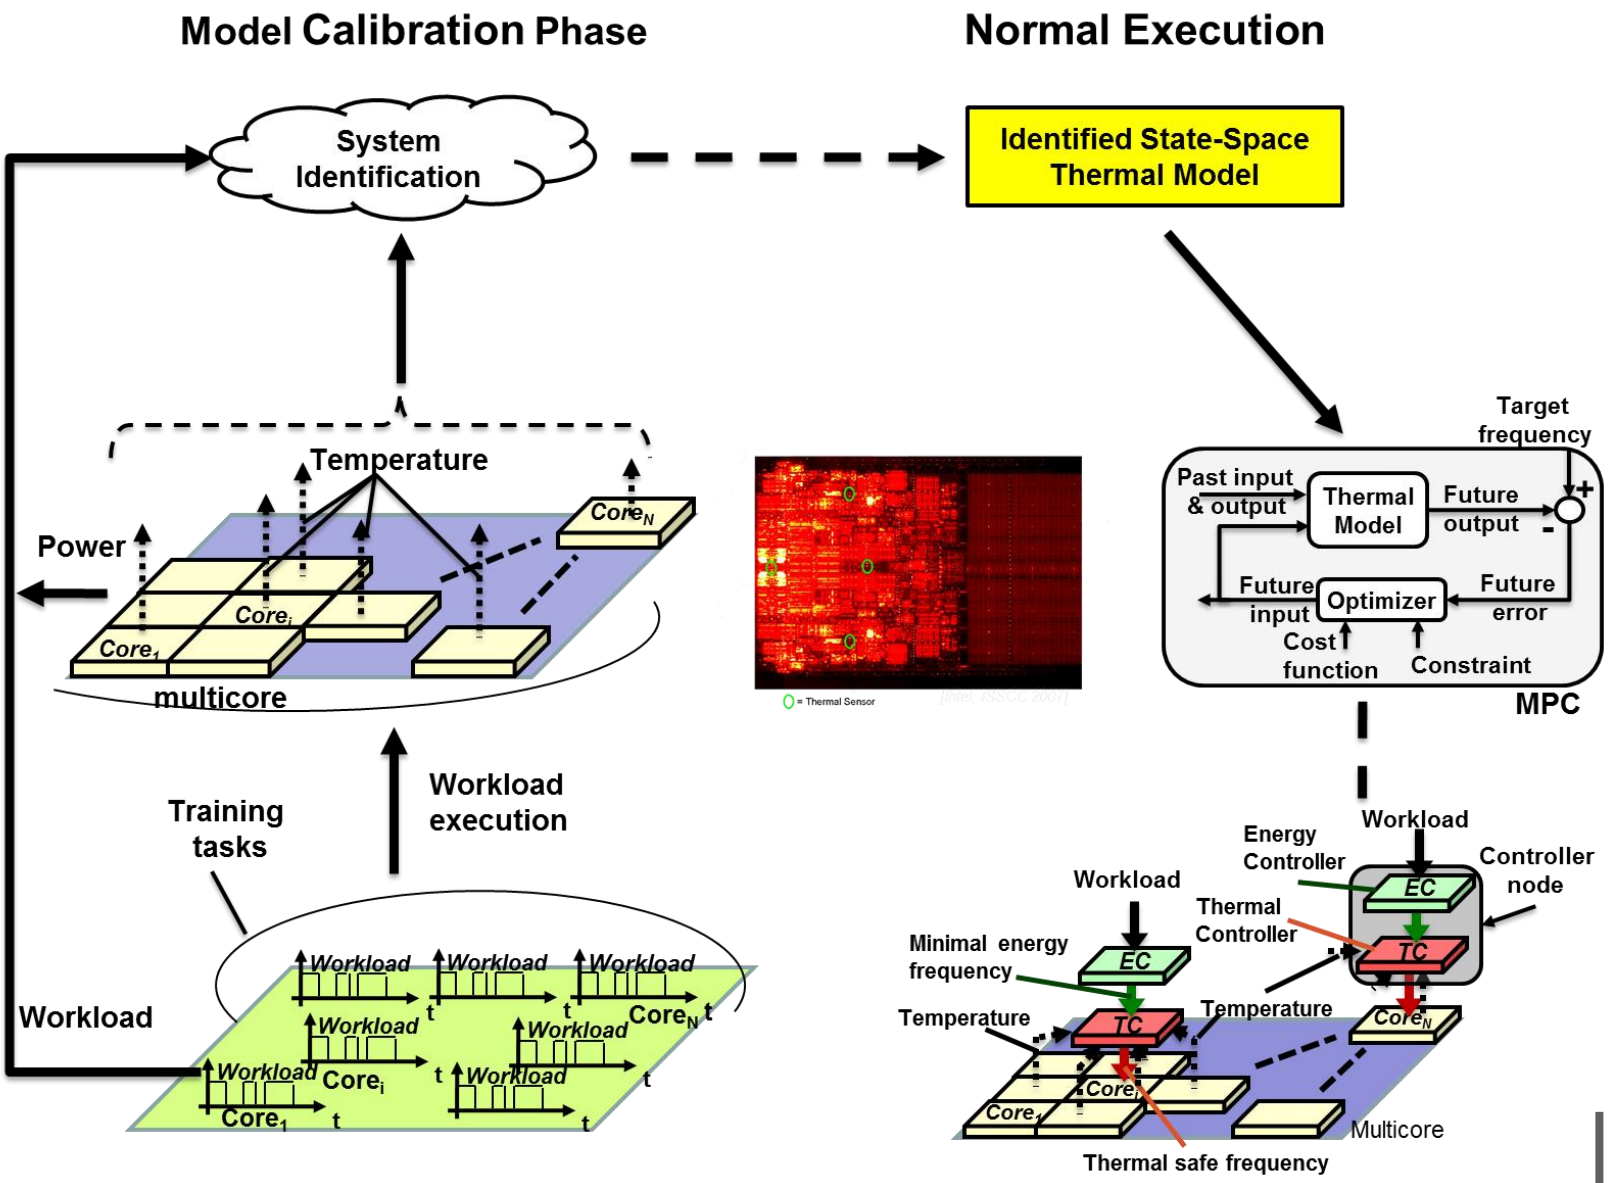

Thermal Controller

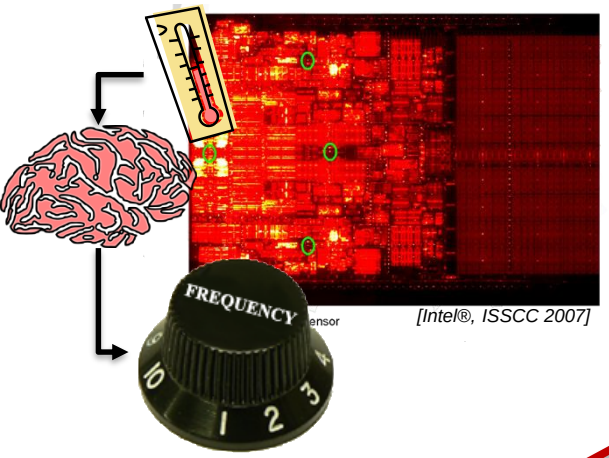

[Intel®, ISSCC 2007]

COMPLEXITY

Threshold based controller

- T > Tmax → low freq
- T < Tmin → high freq
- cannot prevent overshoot
- thermal cycle

Classical feed-back controller

- PID controllers
- Better than threshold based approach
- Cannot prevent overshoot

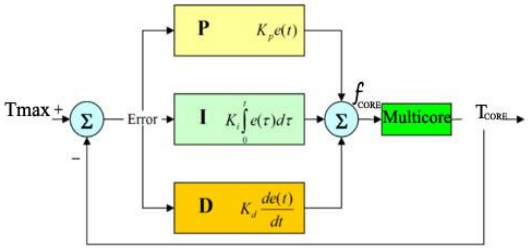

Model Predictive Controller

- Internal prediction: avoid overshoot
- Optimization: maximizes performance

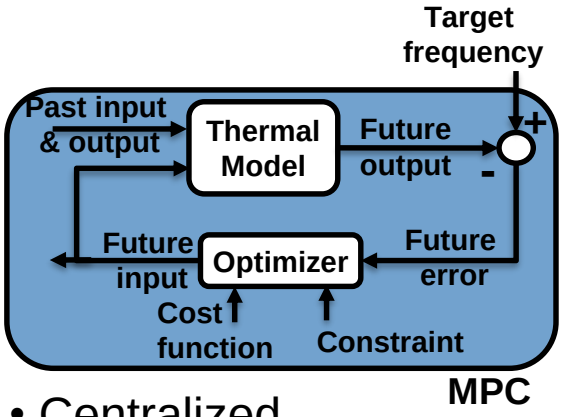

- Centralized
  - aware of neighbor cores thermal influence
  - All at once – MIMO controller
  - Complexity !!!

MPC Scalability

MPC *Complexity*

- **Implicit** - *a.k.a. on-line*
  - computational burden
- **Explicit** – *a.k.a. off-line*
  - high memory occupation

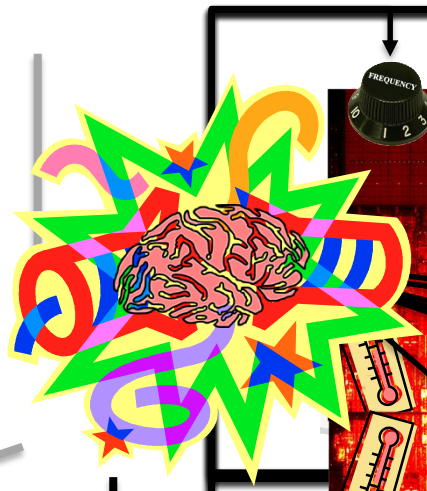

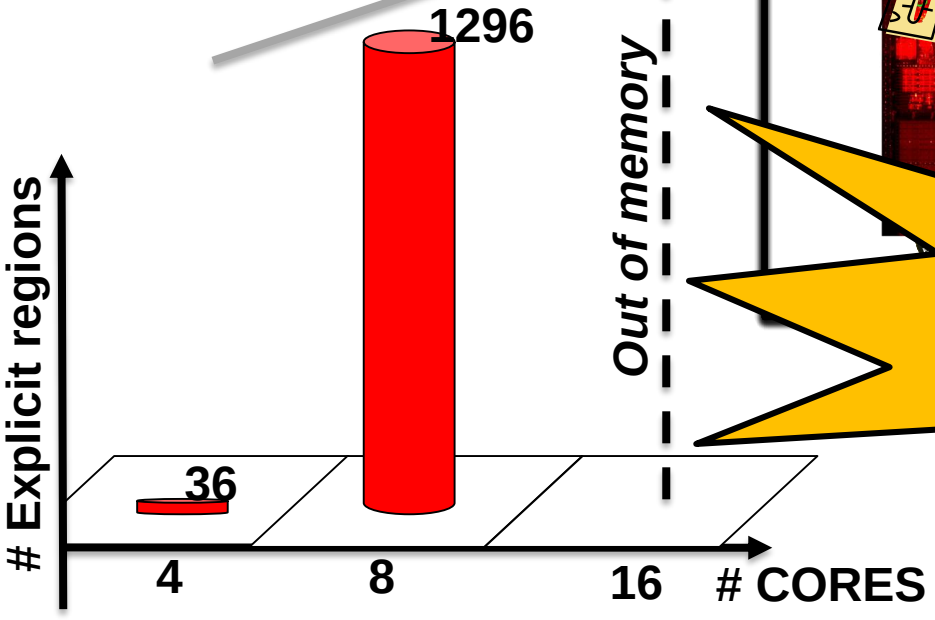

Complexity grows *superlinearly* with number of cores!!

Addressing Scalability

On a short time window, power has a local thermal effect!

One controller for each core
Controller uses:
- local power & thermal model
- neighbor's temperatures

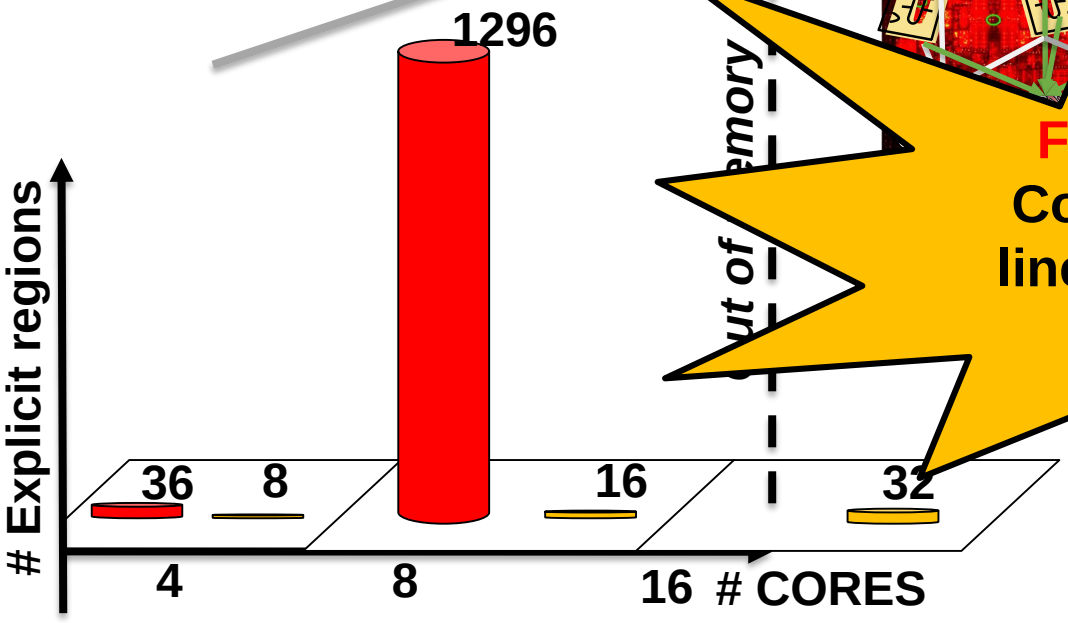

Fully distributed
Complexity scales linearly with #cores

Model Predictive Distributed Control

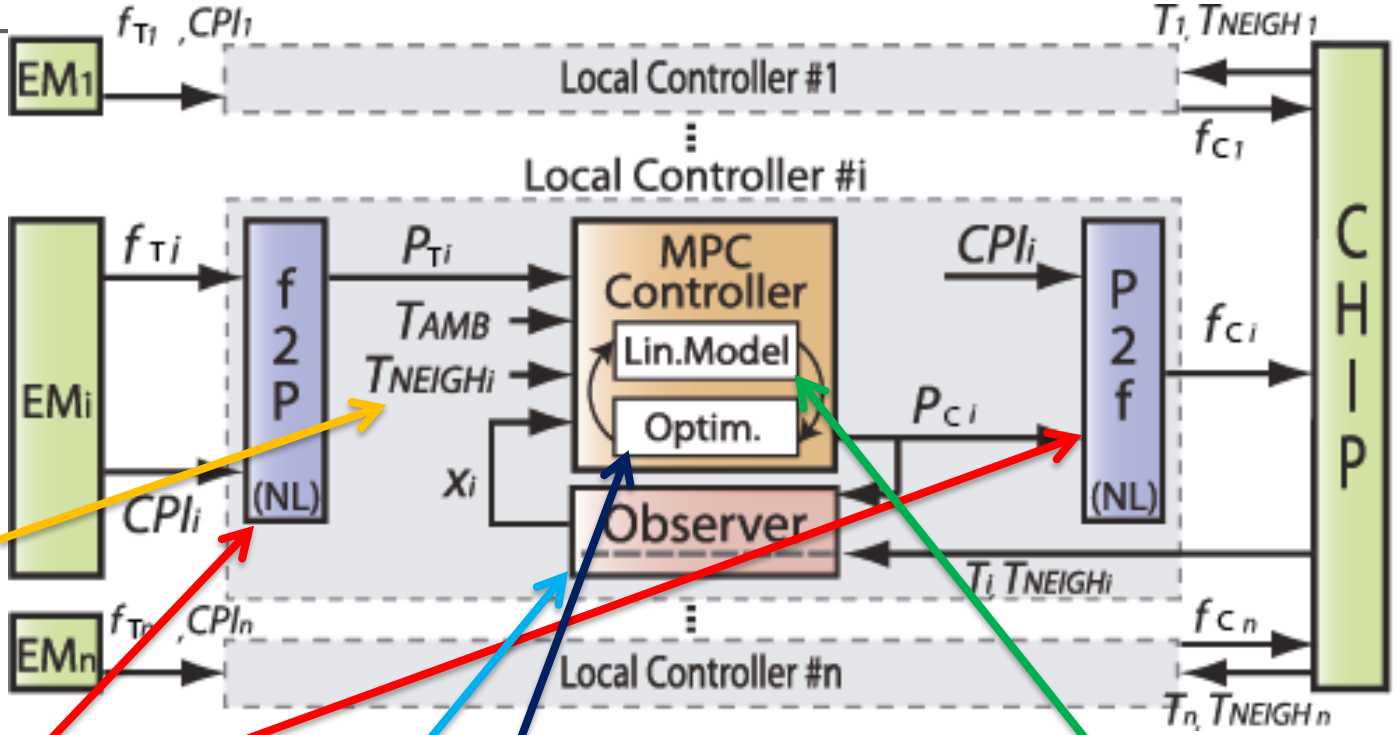

2 states per core

Nonlinear (Frequency to Power)

Implicit formulation

$$\min \sum_{k=0}^{N-1} \|Q \cdot (f_{TC,1}(t+k) - f_{EC,1}(t+k))\|_2$$

s.t

$$0 \le T_1(t+k|t) \le T_{MAX,1}$$

Linear (Power to Temperature)

$$\hat{x}(k+1) = A_1 \cdot \hat{x}(k) + B_1 \cdot u(k) + L_1 \cdot (T_1 - C_1 \cdot \hat{x}(k))$$

Classic Luenberger state observer

Model Predictive Distributed Control

(Bartolini et al., 2013) — (Tilli et al., 2012) — (Tilli et al., 2015)

Distributed
vs
Centralized
vs
PID

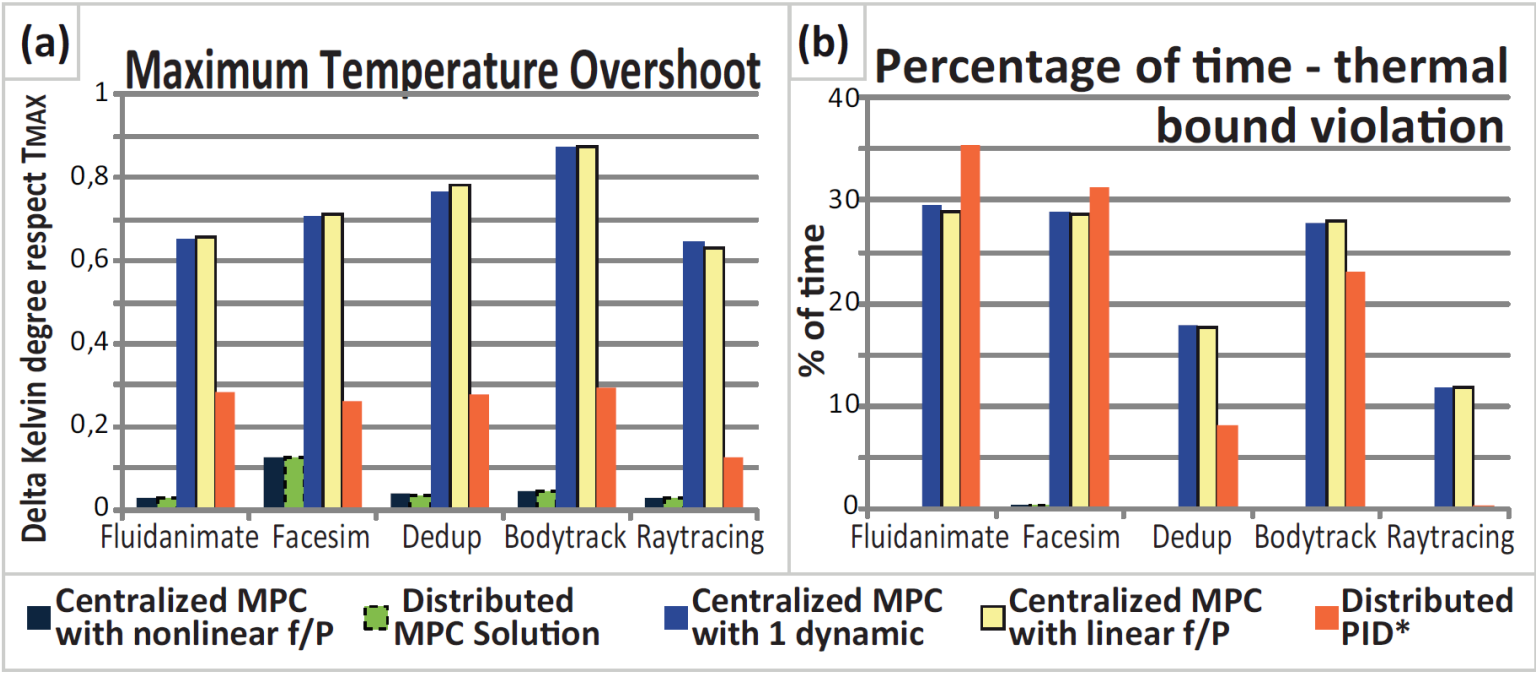

Complexity

a) *Centralized MPC Complexity*

	MPC_explicit	MPC_implicit
4	81	7,70
8	6561	9,00
16	OUT	24,20
48	OUT	85,50
	# regions	time (us)

b) *Distributed MPC Complexity (single core)*

f2P	0,061	time (us)
Observer	0,743	time (us)
MPC (Impl)	4,690	time (us)
MPC (Expl)	2	# regions
P2f	1,188	time (us)

Layout

- Motivation
- Static Thermal Modelling
- Dynamic Thermal Modelling
- Thermal Management
- **Datacenter automation**

A New Trend: Datacentre Automation

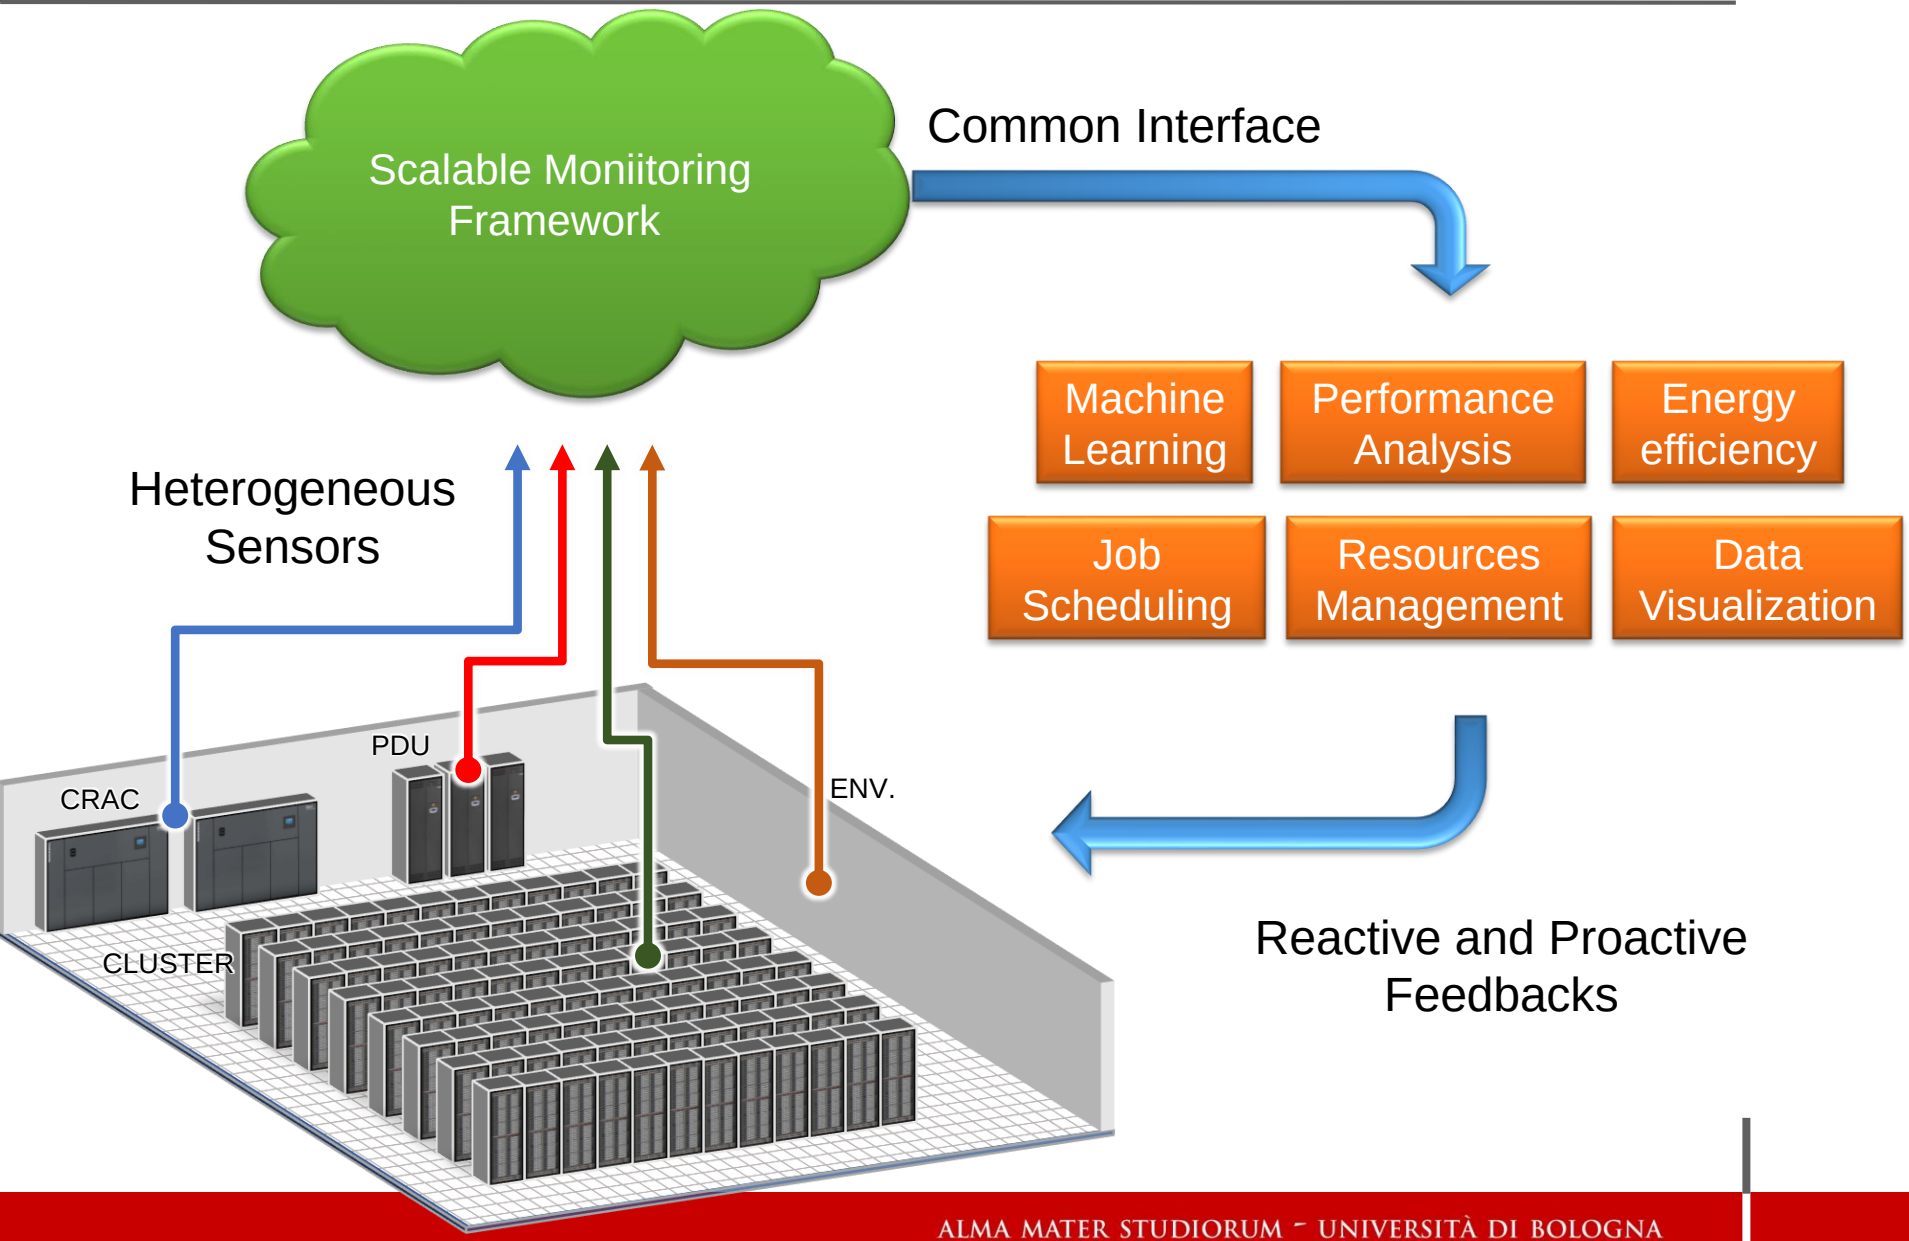

The Big Data & DL backbone

Computing clusters

- Not only the computing engine of Big Data solution
- Also a complex industrial plant and a growing industry in ER
- A compute nodes can produces ~ 100/1000 metrics/s * “peta/exa scale” = Big Data!

Datacenter automation – improve energy/cost efficiency and effectiveness – Industry 4.0 thanks to:

- Live collection and processing of large telemetry data (>100GB/day x cluster)
- On-line generation of “plant models” – a.k.a. digital tweens”, security break detection

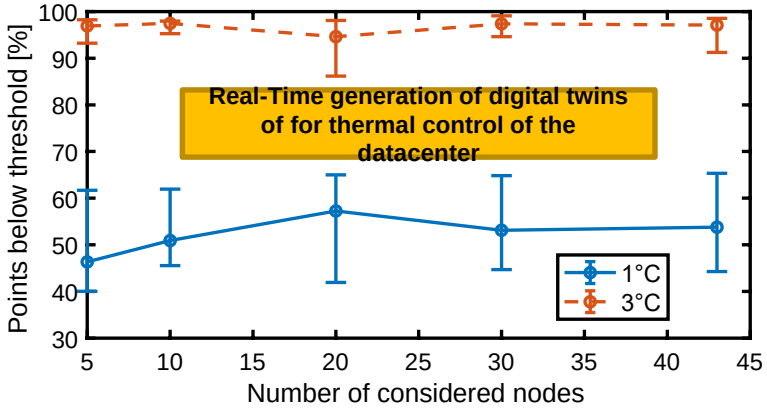

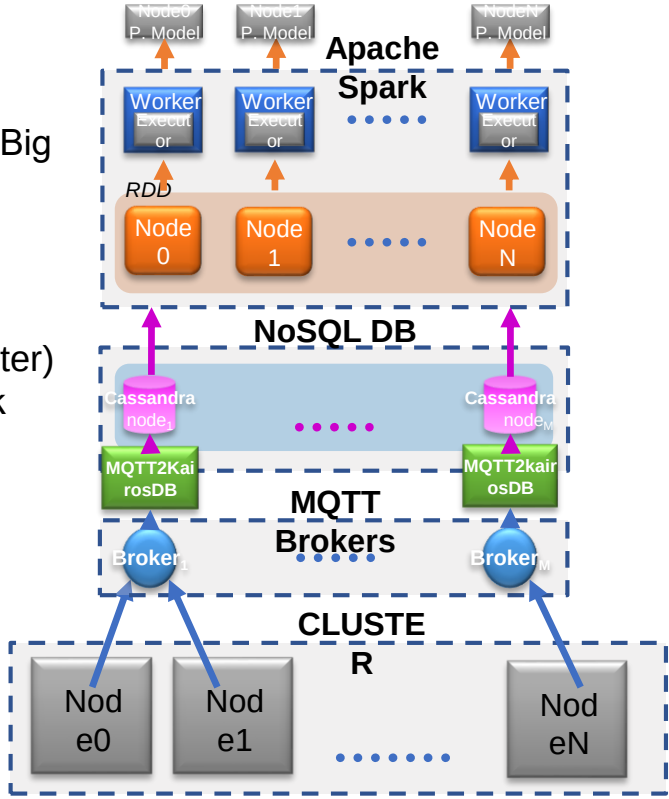

ExaMon: an Industy 4.0 approach to datacenter automation

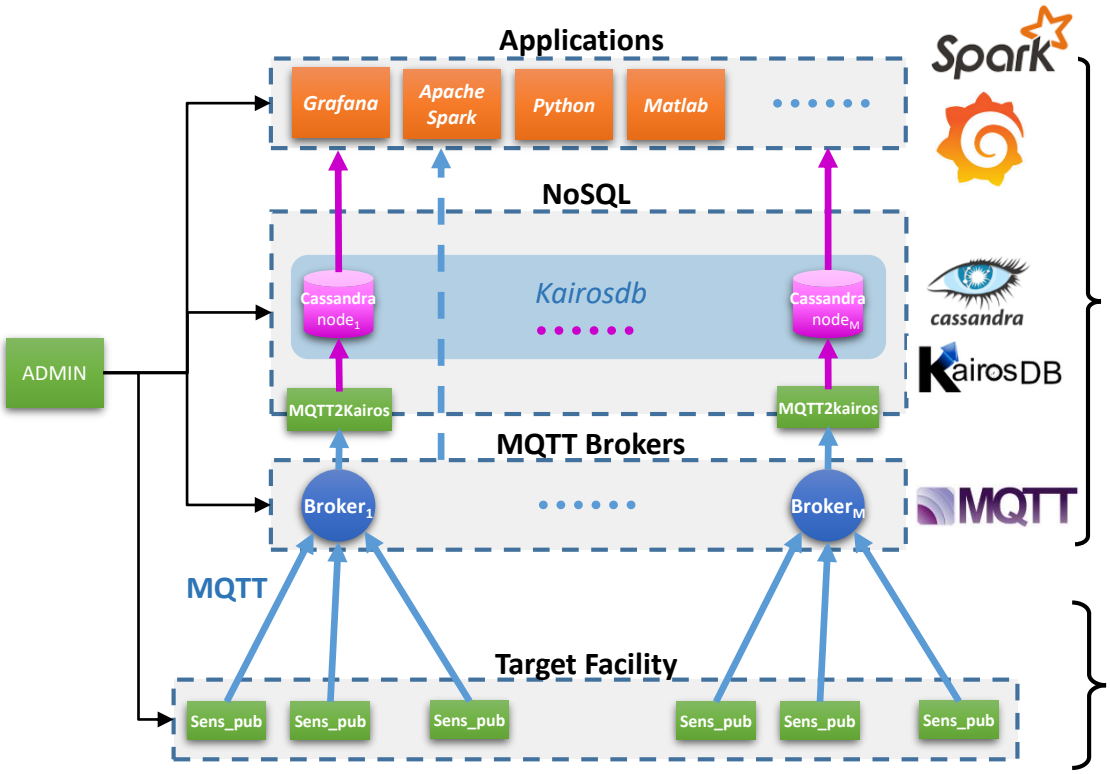

Front-end

- Host: management node
- **Docker containers**
- **~45KS/s**

Back-end

- **MQTT**–enabled monitoring agents (e.g. Dig)

EXAMON

ExaMon: an Industy 4.0 approach to datacenter automation

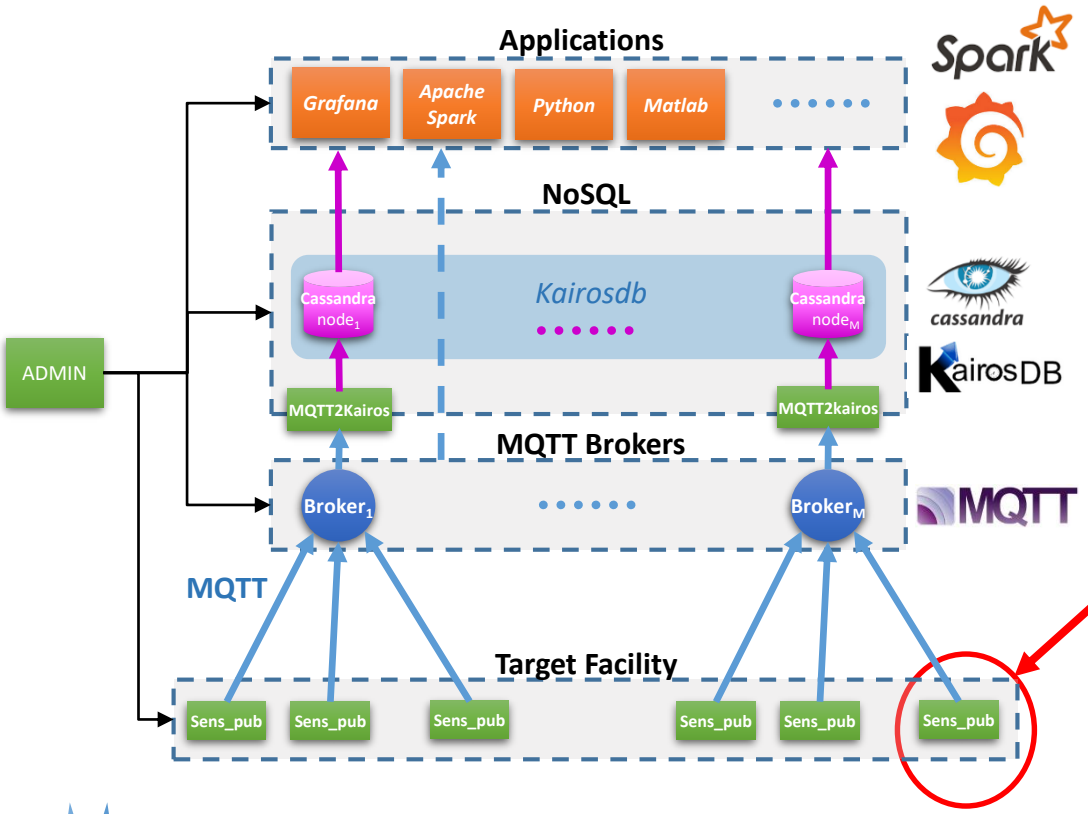

Monitoring Agents:

- Send Power and Performance measurements to the FE

EXAMON

ExaMon: an Industy 4.0 approach to datacenter automation

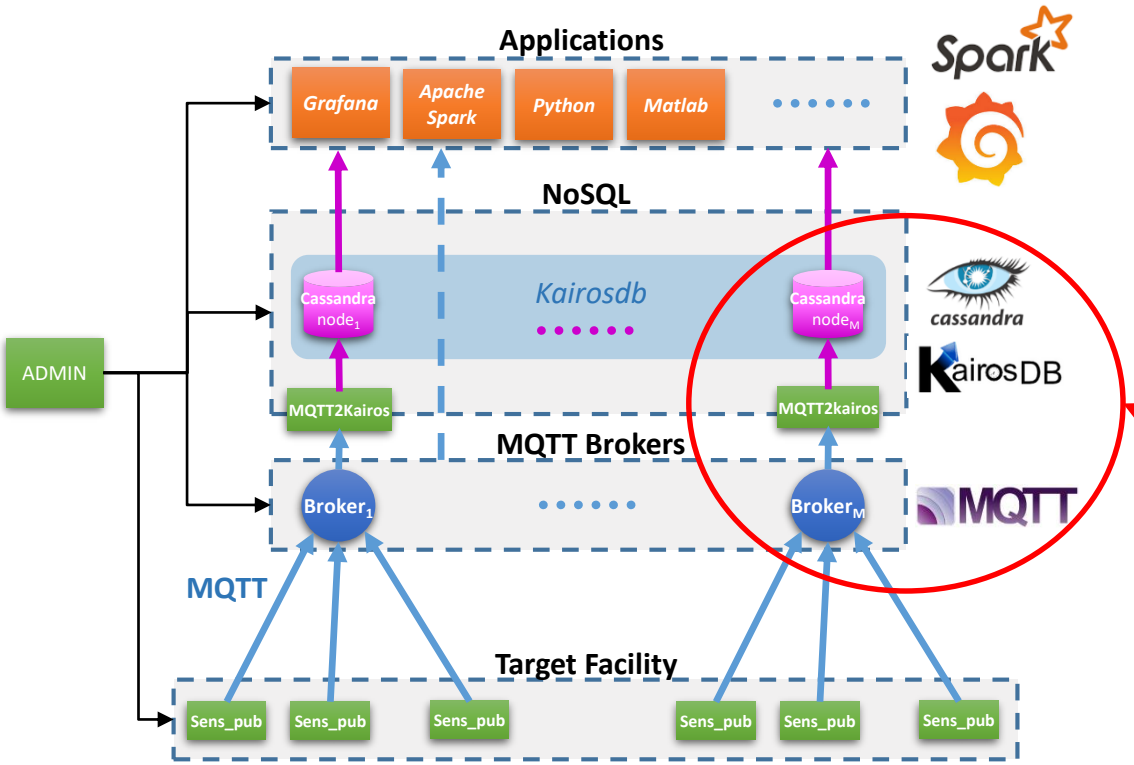

Broker:

- Forward data to the listeners (e.g. kairosDB)

Mqtt2kairosdb:

- Interface between MQTT and KairosDB
- KairosDB is a front-end to handle time series in Cassandra

Cassandra:

- NoSQL database
- Highly scalable
- Optimized to balance the load on multiple nodes

EXAMON

ExaMon: an Industy 4.0 approach to datacenter automation

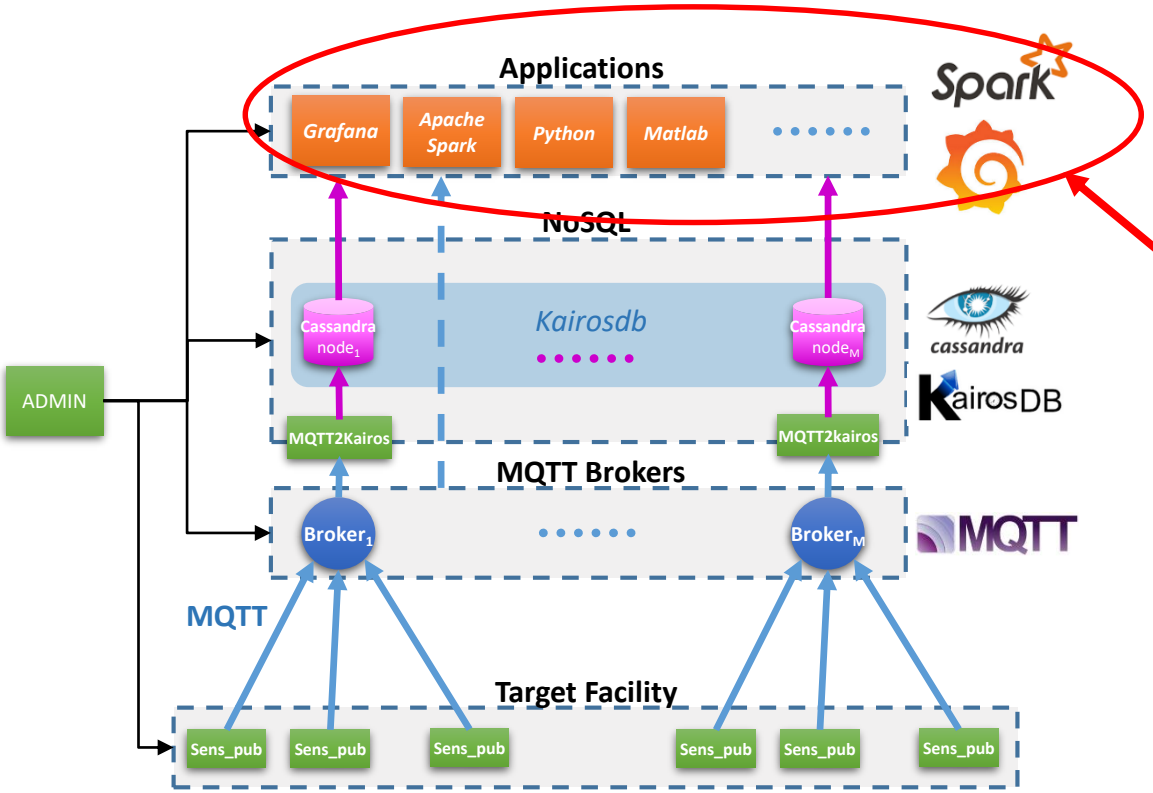

Application layer:

- Grafana, Apache Spark, etc …
- Aggregate metrics for Data Visualization, ML Analysis, Post Processing, etc …

EXAMON

ExaMon: Scalable Data Collection and Analytics

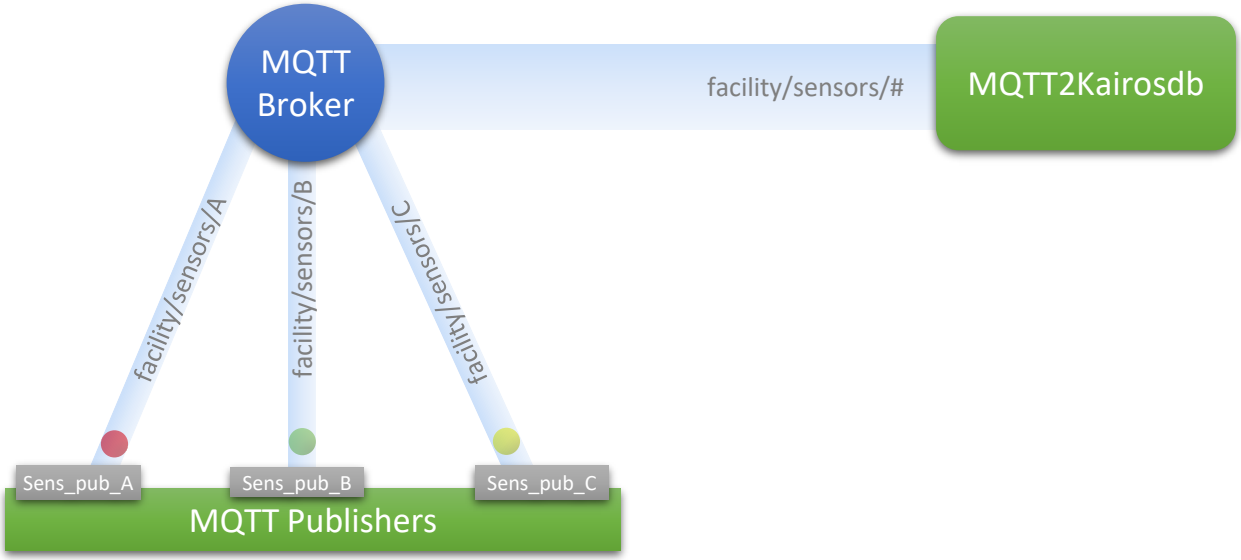

{Key,Value} = TS, Measurement
Topic = /davide/node1/Metric

EXAMON

ExaMon: Batch & Streaming

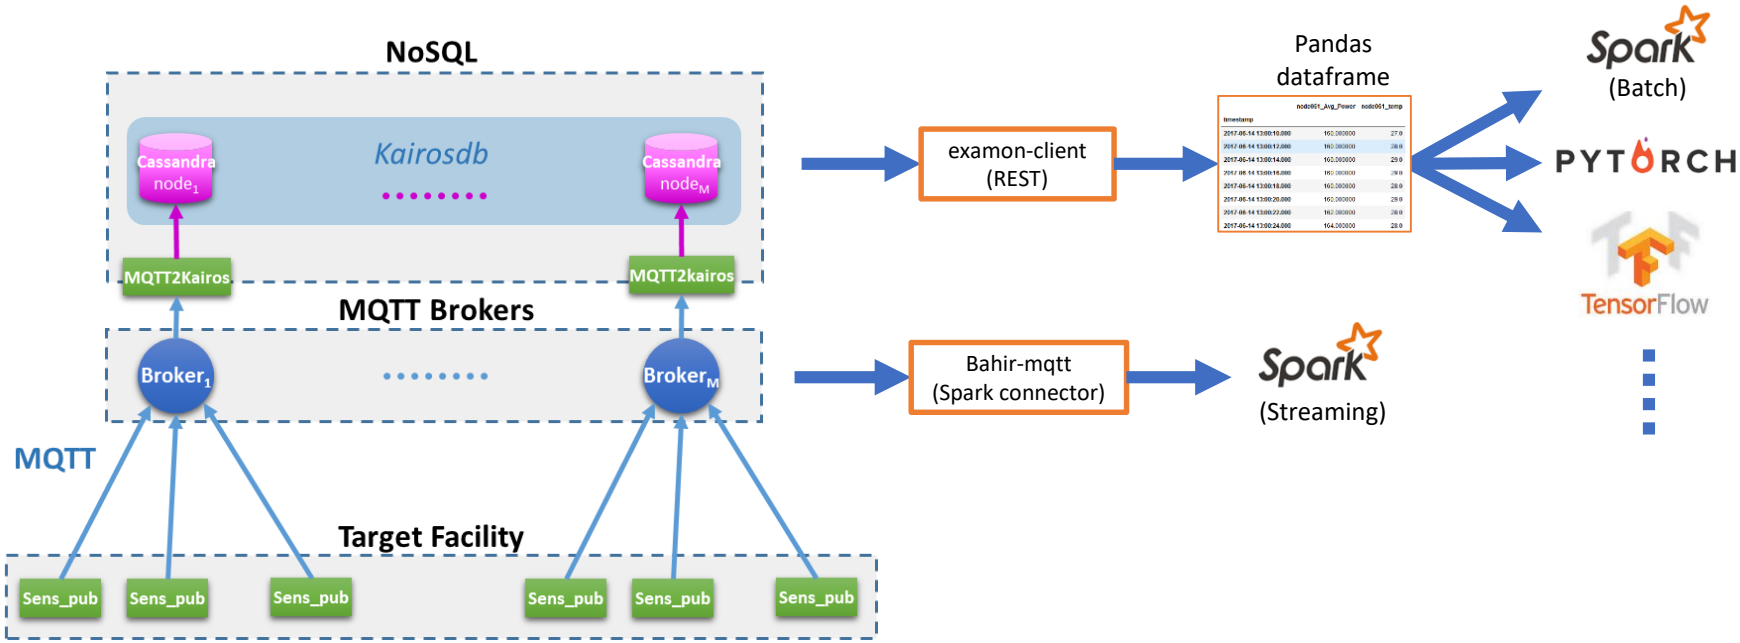

EXAMON

Use Case: Online Power Model Learning

- **Idea: Coupling the monitoring framework with Apache Spark to calculate CPU power model parameters**

Power Model:

$$\mathrm{P}_{pkg0+1} = \sum_{i=0}^{N_c-1} (a_i + b_i \mathrm{IPS}_i)\mathrm{f}_i$$

Spark MLlib: *Streaming linear regression with SGD*

$$\min_{\mathbf{w}\in\mathbb{R}^d} f(\mathbf{w})$$

$$\mathbf{w} = [a_0, \ldots, a_{N_c}, b_0, \ldots, b_{N_c}]$$

$$\mathbf{w}^{(t+1)} := \mathbf{w}^{(t)} - \gamma f'_{\mathbf{w},\mathbf{i}}$$

Spark Program

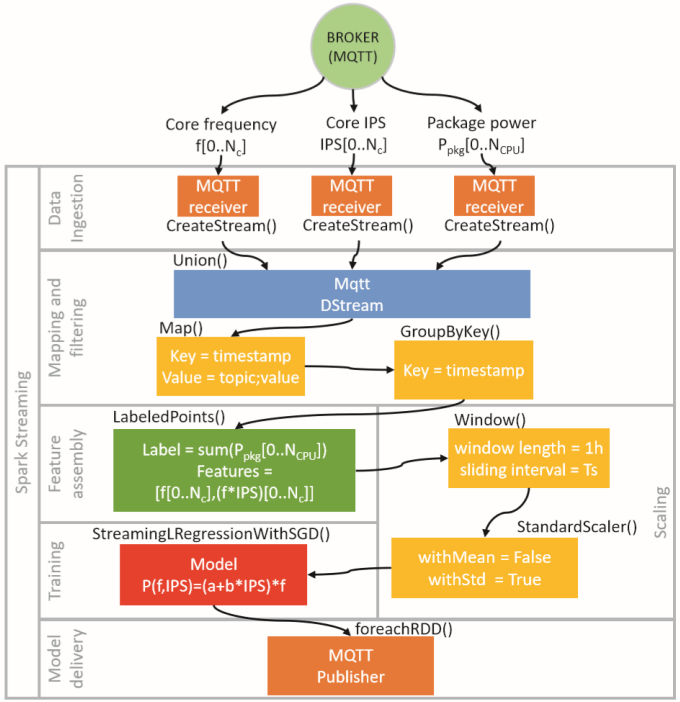

18 September 2018

Use Case: Online Power Model Learning

Topic	Payload
<organization name>/<cluster name>/<node name>/pmu/core/0/IPS	<val>;<timestamp0>
<organization name>/<cluster name>/<node name>/pmu/core/3/freq	<val>;<timestamp0>
<organization name>/<cluster name>/<node name>/pmu/cpu/0/P_pkg	<val>;<timestamp0>
<organization name>/<cluster name>/<node name>/pmu/core/1/IPS	<val>;<timestamp0>
.....	
<organization name>/<cluster name>/<node name>/pmu/core/1/IPS	<val>;<timestampN>

	FEATURES								LABELS			
Metrics	Freq				IPS				P_Pkg			
core/CPU	0	1	..	N	0	1	...	N	0	1	..	M
timestamp0	val	val	val	val	val	val	val	val	val	val	val	val
timestamp1	val	val	val	val	val	val	val	val	val	val	val	val
...	...	...	...	...	...	...	...	...	...	...	...	...
timestampN	val	val	val	val	val	val	val	val	val	val	val	val

Topic	Payload
<organization name>/<cluster name>/<node name>/power_model/a0	<val>;<timestamp0>
<organization name>/<cluster name>/<node name>/power_model/a1	<val>;<timestamp1>
<organization name>/<cluster name>/<node name>/power_model/a2	<val>;<timestamp2>
<organization name>/<cluster name>/<node name>/power_model/a3	<val>;<timestamp3>
.....	
<organization name>/<cluster name>/<node name>/power_model/bN	<val>;<timestampN>

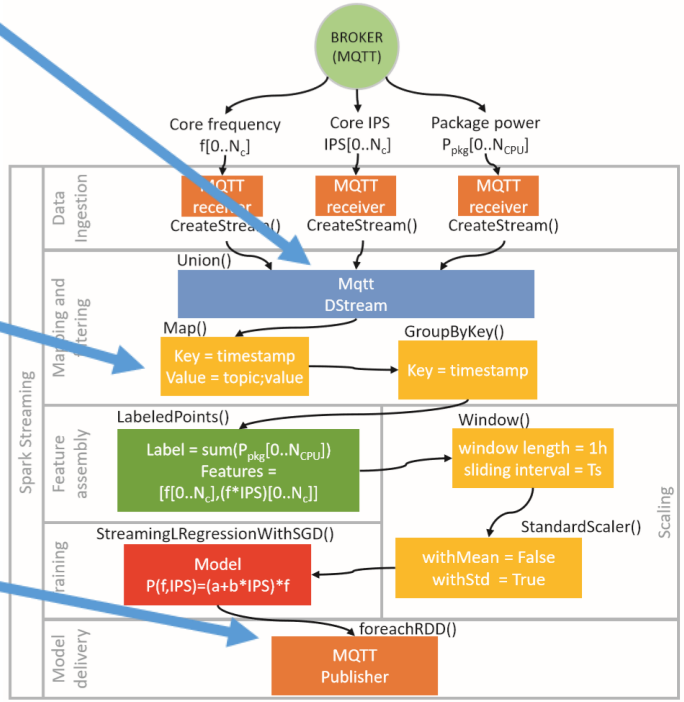

Use Case: Galileo (CINECA) setup

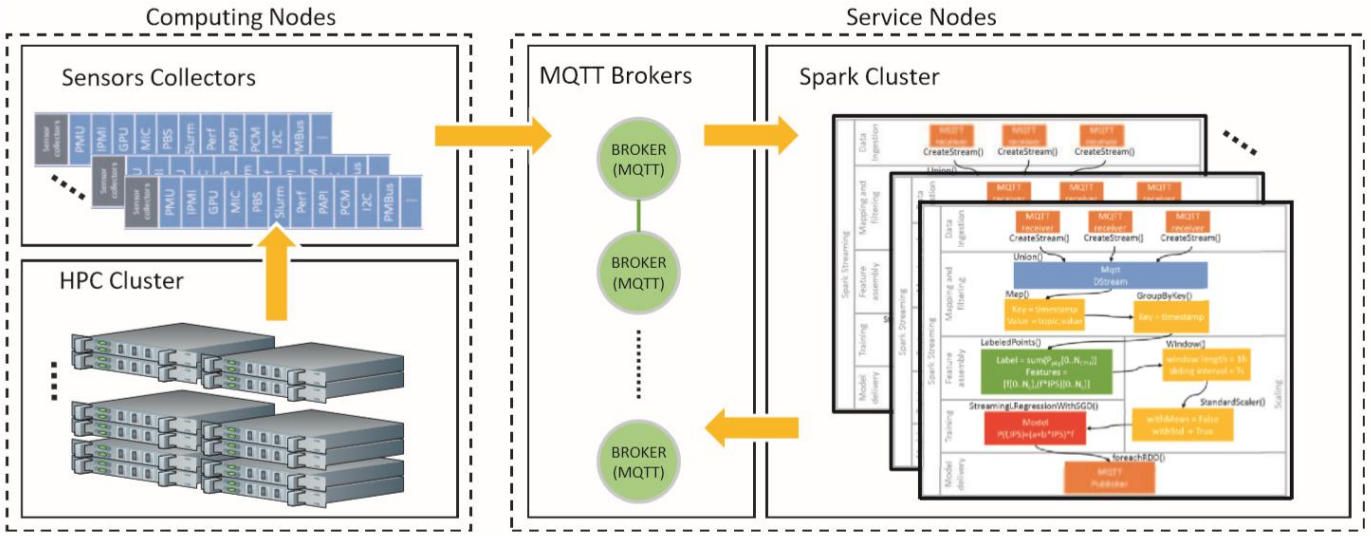

- **516 nodes**
- **430 metrics/node**

- **221880 total metrics**
- **144 Mb/s**
  - **@ 2 seconds sampling time**

Use Case: Results

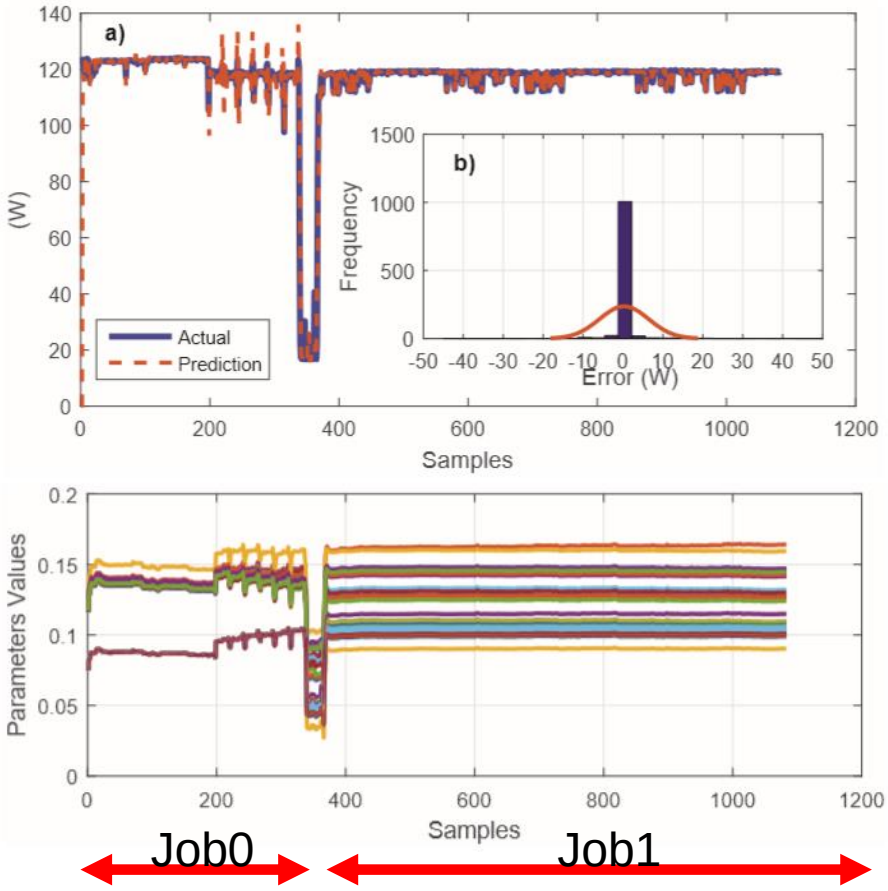

- **Actual Power vs Model prediction (CPU0 + CPU1)**
  - Transition between two different workload phases running on the node
  - Online algorithm promptly recovers, learning a new model
  - 5-6 samples recovery time)
- Average model accuracy:
  - ±25mW
- Model update time (iteration):
  - 600ms

Algorithm returned satisfactory performance!

Thermal modeling at scale

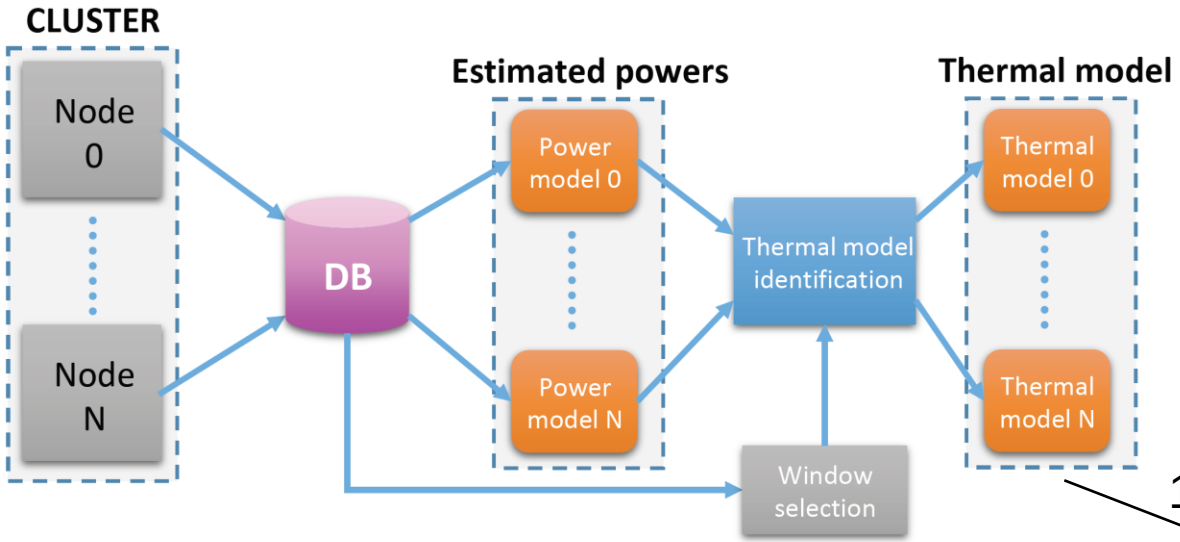

14days divided in 1.5day window

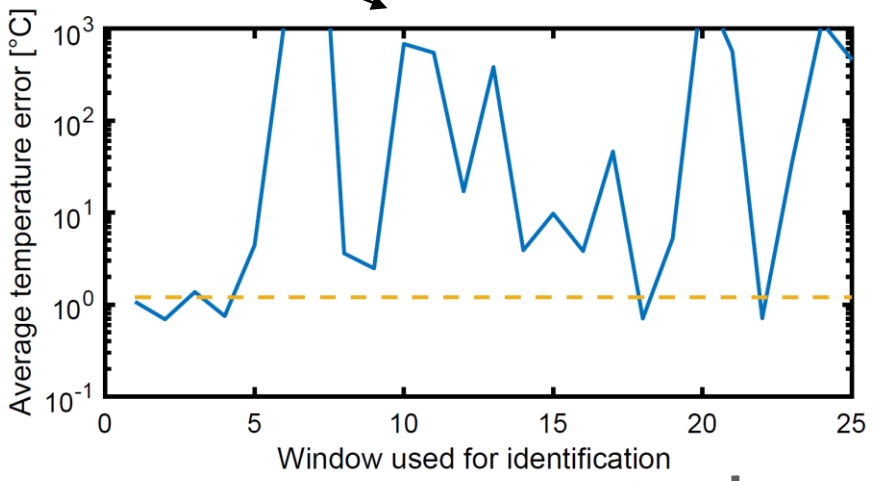

Problem with training

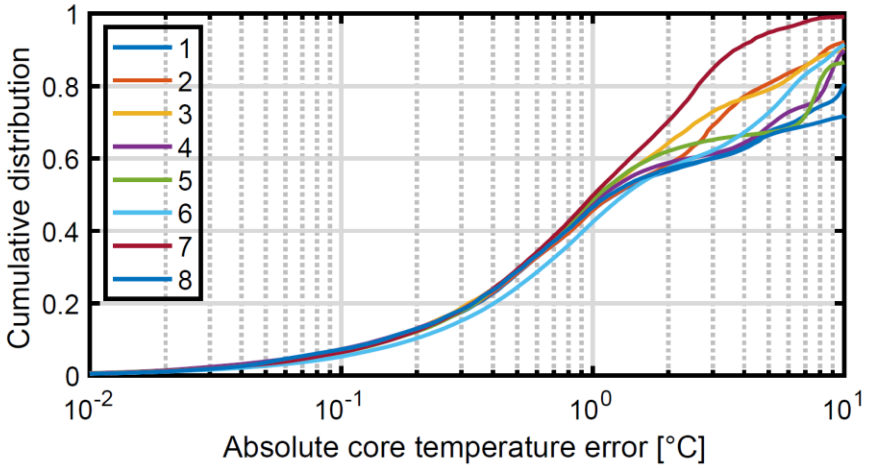

(a) window 16

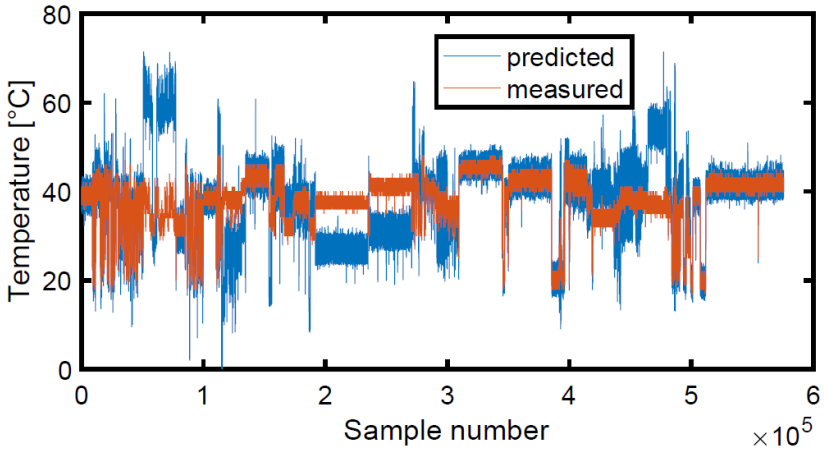

(a) window 16

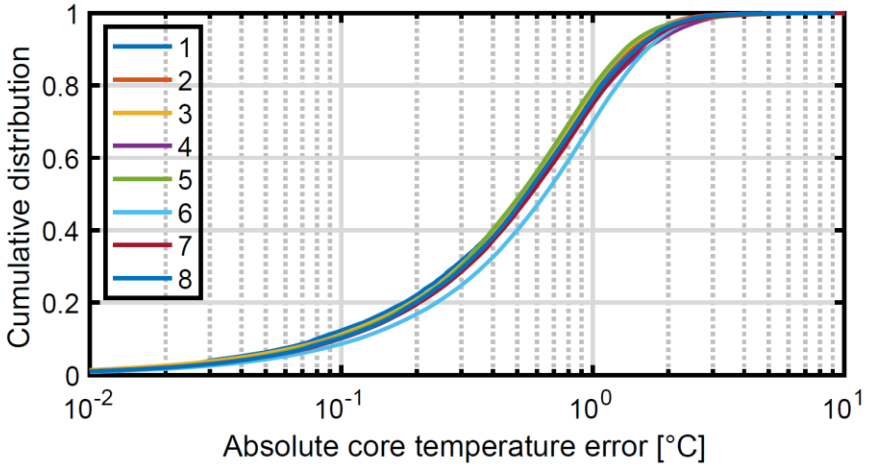

(b) window 22

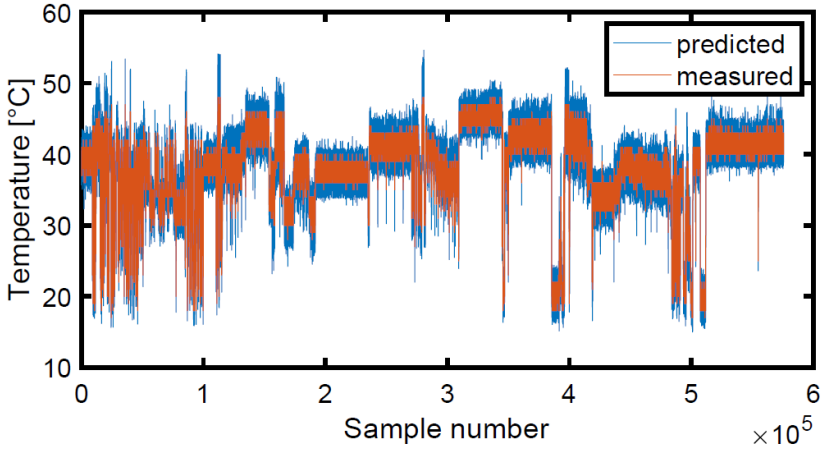

(b) window 22

How to evaluate goodness

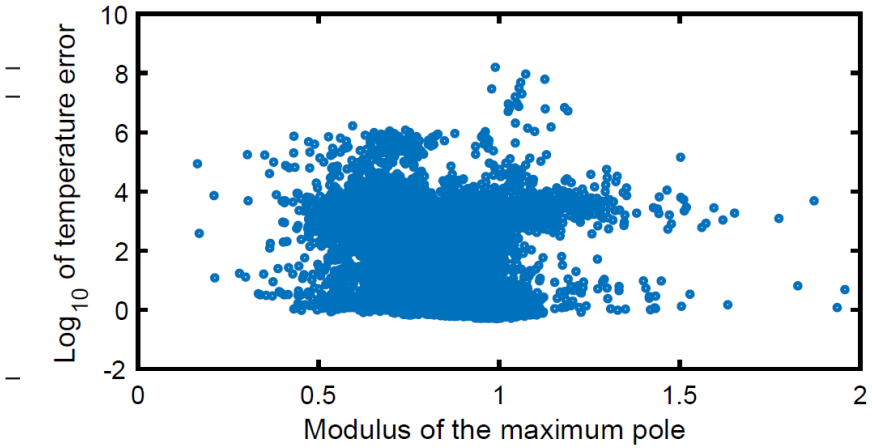

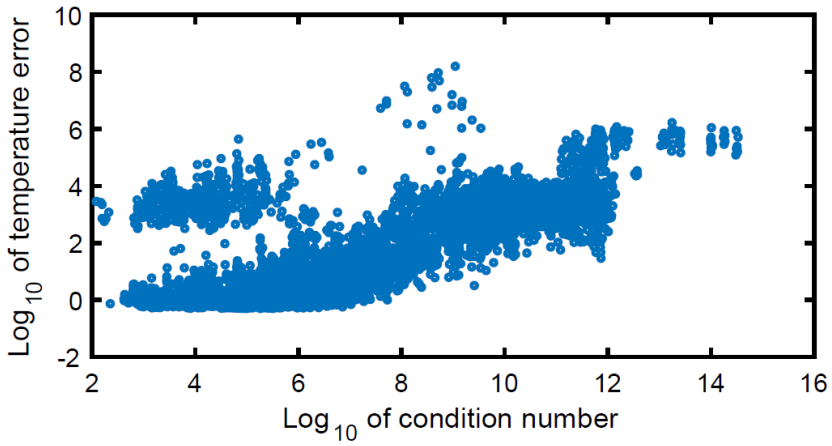

Experts features not good!

Deep Learning approaches beats expert user in determining the «good» windows

D.A.V.I.D.E. (#18 Green500 Nov'17)

E4 COMPUTER ENGINEERING

D.A.V.I.D.E.
SUPERCOMPUTER
(**D**evelopment of an
Added
Value
Infrastructure
Designed in
Europe)

PRACE

TOP500 The List. THE GREEN 500

OpenPOWER

D.A.V.I.D.E. SUPERCOMPUTER

(Development of an Added Value Infrastructure Designed in Europe)

OCP form factor compute node based on IBM Minsky

4x NVIDIA Tesla P100 HSMX2

2 x IBM POWER8 with NVLink

2xIB EDR

BusBar

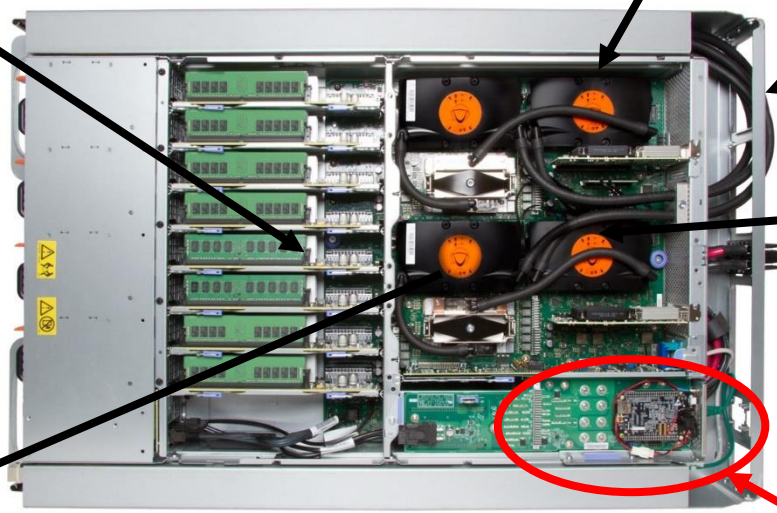

PRACE

LIQUID COOLING

DiG

ETH Zurich / University of Bologna

OUT-OF-BAND HIGH RESOLUTION POWER AND PERFORMANCE MONITORING

DiG: High Resolution Out-of-band Power Monitoring

- Out-of-band → Zero overhead
- Collect more than 1.5 kS/s, 7/7d, 24/24h, for all users
- Architecture independent (i.e. tested on Intel, ARM and IBM)
- Fine grain → down to ms scale (sampling @800 kS/s + avg)
- IoT communication technology (MQTT) → scalable
- Edge Computing!

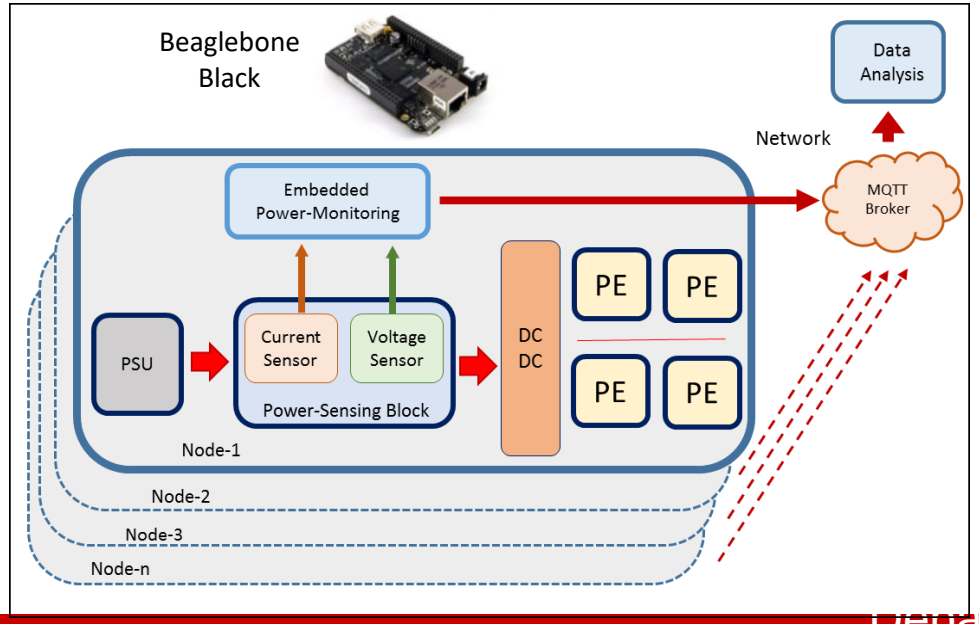

Conclusion

- Machine learning and DL is a key tool for solving engineering problems!
- We are now able to design predictive control loop on supercomputing processors
- Future works: Datacenter automation and anomaly detection!

Pubblications – Naif Modelling

- Davide Rossi, Antonio Pullini, Igor Loi, Michael Gautschi, Frank K. Gürkaynak, Andrea Bartolini, Philippe Flatresse, Luca Benini, **A 60 GOPS/W, −1.8 V to 0.9 V body bias ULP cluster in 28 nm UTBB FD-SOI technology**, JSSE15
- Conficoni C.; Bartolini A.; Tilli A.; Tecchiolli G.; Benini L**.; Energy-Aware Cooling for Hot-Water Cooled Supercomputers**, DATE 15
- Conficoni, Christian, et al. "**HPC Cooling: A Flexible Modeling Tool for Effective Design and Management**." IEEE Transactions on Sustainable Computing (2018).
- Bartolini, Andrea, et al. "**The DAVIDE Big-Data-Powered Fine-Grain Power and Performance Monitoring Support.**" (2018).
- Beneventi, Francesco, et al. "**Continuous learning of HPC infrastructure models using big data analytics and in-memory processing tools**." 2017 Design, Automation & Test in Europe Conference & Exhibition (DATE). IEEE, 2017.
- Bartolini A.; Cacciari M.; Cavazzoni C.; Tecchiolli G.; Benini L.; **Unveiling Eurora — Thermal and power characterization of the most energy-efficient supercomputer in the world**, DATE14
- Fraternali, Francesco, et al. "**Quantifying the impact of variability and heterogeneity on the energy efficiency for a next-generation ultra-green supercomputer**." IEEE Transactions on Parallel and Distributed Systems 29.7 (2018): 1575-1588.
- F. Fraternali, A. Bartolini, C. Cavazzoni, G. Tecchiolli, L. Benini, **Quantifying the impact of variability on the energy efficiency for a next-generation ultra-green supercomputer**; ISLPED14

Pubblications Thermal

- Beneventi F. ; Bartolini A. ; Tilli A. ; Benini L., **An Effective Gray-Box Identification Procedure for Multicore Thermal Modeling** TC14
- Diversi R.; Tilli A.; Bartolini A.; Beneventi F.; Benini L.; **Bias-Compensated Least Squares Identification of Distributed Thermal Models for Many-Core Systems-on-Chip**, TCAS 14
- Diversi, R; Bartolini, A.; Tilli, A.; Beneventi, F.; Benini, L.; ,"**SCC thermal model identification via advanced bias-compensated least-squares**", DATE13 **BPA**
- Bartolini, A., Diversi, R., Cesarini, D., & Beneventi, F.. **Self-Aware Thermal Management for High Performance Computing Processors.** IEEE Design & Test.
- Diversi, Roberto, et al. "**Thermal model identification of supercomputing nodes in production environment**." Industrial Electronics Society, IECON 2016-42nd Annual Conference of the IEEE. IEEE, 2016.
- Bartolini, Andrea, et al. "**Multiscale Thermal Management of Computing Systems-The MULTITHERMAN approach**." IFAC-PapersOnLine 50.1 (2017): 6709-6716.
- Beneventi, Francesco, et al**. "Static thermal model learning for high-performance multicore servers."** Computer Communications and Networks (ICCCN), 2011 Proceedings of 20th International Conference on. IEEE, 2011.
- Beneventi, Francesco, et al. "**Cooling-aware node-level task allocation for next-generation green HPC systems.**" High Performance Computing & Simulation (HPCS), 2016 International Conference on. IEEE, 2016.
- Bartolini, Andrea, et al. "**Thermal and energy management of high-performance multicores: Distributed and self-calibrating model-predictive controller**." IEEE Transactions on Parallel and Distributed Systems 24.1 (2013): 170-183.

Pubblications - NN

- Netti, A., Kiziltan, Z., Babaoglu, O., Sirbu, A., Bartolini, A., & Borghesi, A. (2018). **FINJ: A Fault Injection Tool for HPC Systems**. arXiv preprint arXiv:1807.10056.
- Lombardi, Michele, Michela Milano, and Andrea Bartolini. "**Empirical decision model learning**." Artificial Intelligence 244 (2017): 343-367.
- Bartolini, Andrea, et al. **"Optimization and Controlled Systems: A Case Study on Thermal Aware Workload Dispatching**." AAAI. 2012.

Pubblications – big data

- Bartolini, Andrea, et al. "**The DAVIDE Big-Data-Powered Fine-Grain Power and Performance Monitoring Support.**" (2018).
- Beneventi, Francesco, et al. "**Continuous learning of HPC infrastructure models using big data analytics and in-memory processing tools**." 2017 Design, Automation & Test in Europe Conference & Exhibition (DATE). IEEE, 2017.

Acknowledge

- HPC team:
  - Luca Benini, Francesco Beneventi, Antonio Libri, Andrea Borghesi, Federico Pittino, Alessandro Petrella, Daniele Cesarini, Christian Conficoni, Roberto Diversi, Andrea Tilli, Michele Lombardi, Michela Milano