

The **POWER** *of* **MEASURING**

How performance measurements impact
the modeling of computing systems' power

L. Fontoura Cupertino, G. Da Costa and J.-M. Pierson

IRIT - Institut de Recherche en Informatique de Toulouse
UPS - Université Paul Sabatier

GreenDays, Toulouse, Mars 2015



Outline

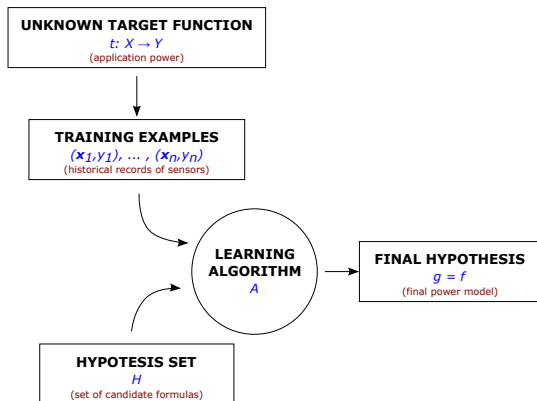
- ① Introduction
- ② Measurement and Uncertainty
- ③ System-wide power models
- ④ Conclusions

- ① Introduction
- ② Measurement and Uncertainty
- ③ System-wide power models
- ④ Conclusions

Introduction


- ▶ Power models can be used in a wide range of applications
 - ▶ Power estimation of applications
 - ▶ Energy efficient scheduling
- ▶ Current models' limitations
 - ▶ Application specific or lack accuracy
 - ▶ Hardware dependent
- ▶ Machine learning
 - ▶ Simple models (calibration)
 - ▶ Noise may be modeled (large number of inputs)
- ▶ Measurements
 - ▶ Accuracy

Machine Learning for power estimation




- ▶ Definition of the Training Set
- ▶ Performance indicators
- ▶ Accurate power measurements


Data Acquisition Infrastructure



Data Acquisition Infrastructure




Data Acquisition Infrastructure



Outline


- ① Introduction
- ② Measurement and Uncertainty
- ③ System-wide power models
- ④ Conclusions

Infrastructure's issues



- ▶ Acquisition response time (latency)
- ▶ Power Conversion Losses
- ▶ Time shifted data
- ▶ Repeated measurements
- ▶ Steady state


Acquisition latency



	Min.	\bar{x}	σ	Max.	Samples
Power	0.3189	0.6162	0.1894	1.646	3810
KPI	0.0021	0.0060	0.0047	0.052	2537


Acquisition response times

KPI are measured onsite presenting a small latency to be acquired, while the power requires a Bluetooth connection.



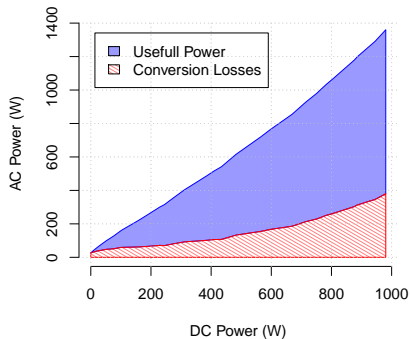
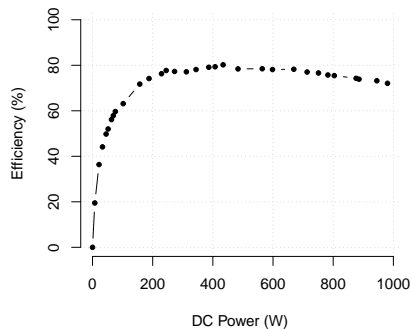
Acquisition response times

KPI are measured onsite presenting a small latency to be acquired, while the power requires a Bluetooth connection.





Power Conversion Losses

- ▶ PSU Input/Output power data provided by PSU's vendors



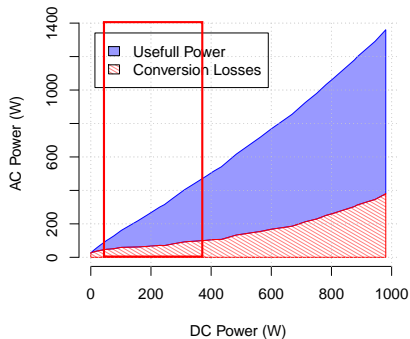
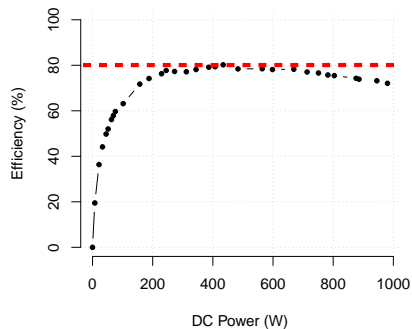
Power Conversion Losses

- ▶ PSU Input/Output power data provided by PSU's vendors



Power Conversion Losses


- ▶ PSU Input/Output power data provided by PSU's vendors



Modeling PSU's conversion losses

- ▶ Univariate polynomial model:


$$P_{DC} = \mathbf{w} \cdot [1 P_{AC} P_{AC}^2 P_{AC}^3 P_{AC}^4 P_{AC}^6 P_{AC}^7 P_{AC}^8]$$



Modeling PSU's conversion losses


- Univariate polynomial model:

$$P_{DC} = \mathbf{w} \cdot [1 P_{AC} P_{AC}^2 P_{AC}^3 P_{AC}^4 P_{AC}^6 P_{AC}^7 P_{AC}^8]$$




Time shifted data

Plogg's measurements take some time to be updated???




Time shifted data

Plogg's measurements take some time to be updated???




Time shifted data

Plogg's measurements take some time to be updated???




Repeated measurements data

Plogg's measurements reply repeated values.




Repeated measurements data

Plogg's measurements reply repeated values.




Steady state data

A simple approach may be just to consider steady state data.



Steady state data

A simple approach may be just to consider steady state data.



Outline

- ① Introduction
- ② Measurement and Uncertainty
- ③ System-wide power models**
- ④ Conclusions

Learning Workload

Maximal impact of each module's device on the system's power consumption


RECS Module	Load	Temp	Processor			RAM		NIC	
			Cache	Cstate	Pstate	Access	Alloc	Down	Up
i7	103.0	2.2	6.8	72.7	48.9	10.2	0.0	8.8	8.2
Atom	3.6	N/A	2.1	3.0	2.4	4.5	0.0	3.0	3.3

Learning Workload

Maximal impact of each module's device on the system's power consumption

RECS Module	Processor					RAM		NIC	
	Load	Temp	Cache	Cstate	Pstate	Access	Alloc	Down	Up
i7	103.0	2.2	6.8	72.7	48.9	10.2	0.0	8.8	8.2
Atom	3.6	N/A	2.1	3.0	2.4	4.5	0.0	3.0	3.3

- ▶ Generic and synthetic
 - ▶ Training set: CPU, memory and network stressed
 - ▶ 3 frequencies: min (1.2GHz), max (2.3GHz) and boost (3.1GHz)

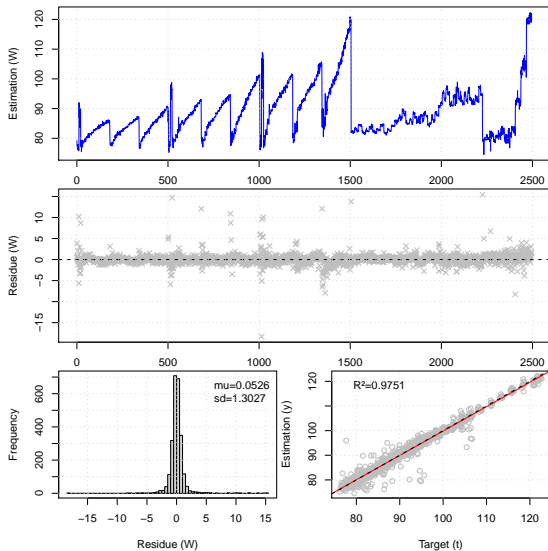


- ▶ 1 workload execution (run) to learn the model
- ▶ (4 + 10) workload runs to validate the model


- ▶ Artificial Neural Network
 - ▶ Multi-layer perceptron network
 - ▶ Training, validation and test subsetting
 - ▶ Topology definition

Raw Data

Learning

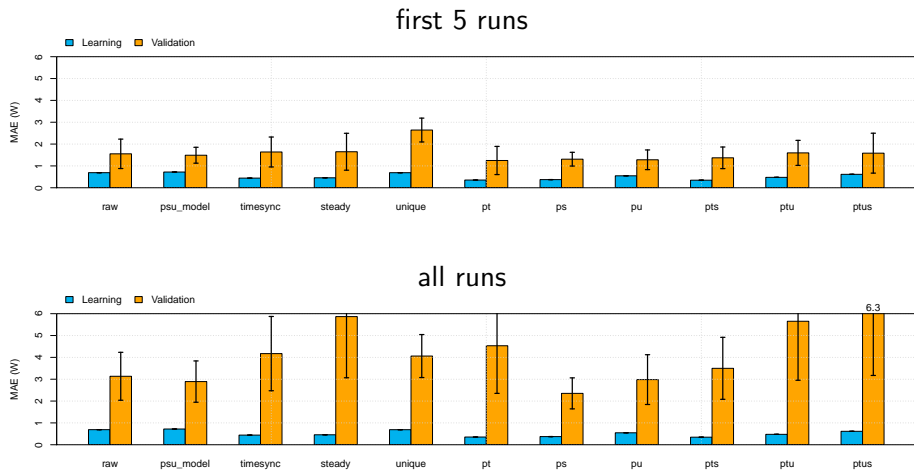


System-wide power models



The results seems fair, but the variances are quite high .. let's run some more times to reduce it.

System-wide power models

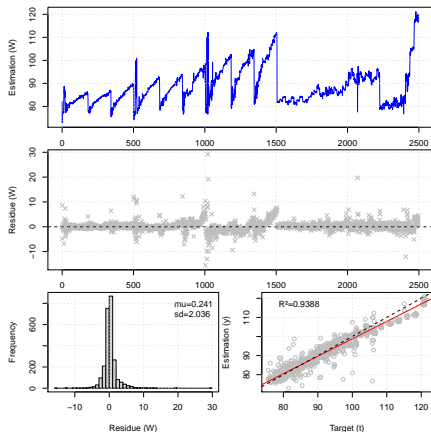


Oh ohhh ... let's have a further look.

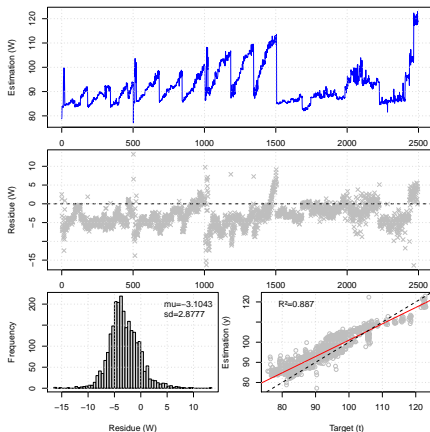
Raw Data

Validation

run 2




run 12



First 5 runs executed after turning the node on. 10 others it was already on. Why are the 5 first good and the 10 others not? Could it be a temperature issue?

Temperature




It seems that all workloads have similar temperature profile ..
Timestamp????

PC_USE' SC0FREQ' SC1FREQ' SC2FREQ' SC3FREQ' SC0CSETC0' SC1CSETC0' SC2CSETC0' SC3CSETC0'
SC0CSETC1.IVB' SC1CSETC1.IVB' SC2CSETC1.IVB' SC3CSETC1.IVB' SC0CSETC3.IVB' SC1CSETC3.IVB'
SC2CSETC3.IVB' SC3CSETC3.IVB' SC0CSETC6.IVB' SC1CSETC6.IVB' SC2CSETC6.IVB' SC3CSETC6.IVB'
SC0CSETC7.IVB' SC1CSETC7.IVB' SC2CSETC7.IVB' SC3CSETC7.IVB' PM_USE' SN_SB' SN_RB' SC0FREQ_MSR'
SC1FREQ_MSR' SC2FREQ_MSR' SC3FREQ_MSR' SC0C0_MSR' SC1C0_MSR' SC2C0_MSR' SC3C0_MSR'
PC0PCSCpuClk' PC1PCSCpuClk' PC2PCSCpuClk' PC3PCSCpuClk' PC0PCSTskClk' PC1PCSTskClk' PC2PCSTskClk'
PC3PCSTskClk' PC0PCSPgFlt' PC1PCSPgFlt' PC2PCSPgFlt' PC3PCSPgFlt' PC0PCSCntSw' PC1PCSCntSw' PC2PCSCntSw'
PC3PCSCntSw' PC0PCSCpuMig' PC1PCSCpuMig' PC2PCSCpuMig' PC3PCSCpuMig' PC0PCSPgFltMn' PC1PCSPgFltMn'
PC2PCSPgFltMn' PC3PCSPgFltMn' PC0PCSPgFltMj' PC1PCSPgFltMj' PC2PCSPgFltMj' PC3PCSPgFltMj' PC0PCSAligFlt'
PC1PCSAligFlt' PC2PCSAligFlt' PC3PCSAligFlt' PC0PCSEmFlt' PC1PCSEmFlt' PC2PCSEmFlt' PC3PCSEmFlt'
PC0PCHCpulnst' PC1PCHCpulnst' PC2PCHCpulnst' PC3PCHCpulnst' PC0PCHBraInst' PC1PCHBraInst' PC2PCHBraInst'
PC3PCHBraInst' PC0PCHBraMiss' PC1PCHBraMiss' PC2PCHBraMiss' PC3PCHBraMiss' PC0PCHCacRefs' PC1PCHCacRefs'
PC2PCHCacRefs' PC3PCHCacRefs' PC0PCHCacMisses' PC1PCHCacMisses' PC2PCHCacMisses' PC3PCHCacMisses'
PC0PCHCpuCycl' PC1PCHCpuCycl' PC2PCHCpuCycl' PC3PCHCpuCycl' PC0PCHBusCycl' PC1PCHBusCycl'
PC2PCHBusCycl' PC3PCHBusCycl' PC0PCHStaCycF' PC1PCHStaCycF' PC2PCHStaCycF' PC3PCHStaCycF'
PC0PCHStaCycB' PC1PCHStaCycB' PC2PCHStaCycB' PC3PCHStaCycB' PC0PCHCL1dRA' PC1PCHCL1dRA'
PC2PCHCL1dRA' PC3PCHCL1dRA' PC0PCHCL1dRM' PC1PCHCL1dRM' PC2PCHCL1dRM' PC3PCHCL1dRM'
PC0PCHCL1dWA' PC1PCHCL1dWA' PC2PCHCL1dWA' PC3PCHCL1dWA' PC0PCHCL1dWM' PC1PCHCL1dWM'
PC2PCHCL1dWM' PC3PCHCL1dWM' PC0PCHCL1dPM' PC1PCHCL1dPM' PC2PCHCL1dPM' PC3PCHCL1dPM'
PC0PCHCL1iRM' PC1PCHCL1iRM' PC2PCHCL1iRM' PC3PCHCL1iRM' PC0PCHCDTLBRA' PC1PCHCDTLBRA'
PC2PCHCDTLBRA' PC3PCHCDTLBRA' PC0PCHCDTLBRM' PC1PCHCDTLBRM' PC2PCHCDTLBRM'
PC3PCHCDTLBRM' PC0PCHCDTLBWA' PC1PCHCDTLBWA' PC2PCHCDTLBWA' PC3PCHCDTLBWA'
PC0PCHCDTLBWM' PC1PCHCDTLBWM' PC2PCHCDTLBWM' PC3PCHCDTLBWM' PC0PCHCITLBRA'
PC1PCHCITLBRA' PC2PCHCITLBRA' PC3PCHCITLBRA' PC0PCHCITLBRM' PC1PCHCITLBRM' PC2PCHCITLBRM'
PC3PCHCITLBRM' PC0PCHCLLRA' PC1PCHCLLRA' PC2PCHCLLRA' PC3PCHCLLRA' PC0PCHCLLRM' PC1PCHCLLRM'
PC2PCHCLLRM' PC3PCHCLLRM' PC0PCHCLLWA' PC1PCHCLLWA' PC2PCHCLLWA' PC3PCHCLLWA' PC0PCHCLLWM'
PC1PCHCLLWM' PC2PCHCLLWM' PC3PCHCLLWM' PC0PCHCLLPA' PC1PCHCLLPA' PC2PCHCLLPA' PC3PCHCLLPA'
PC0PCHCLLPM' PC1PCHCLLPM' PC2PCHCLLPM' PC3PCHCLLPM' PC0PCHCBPURA' PC1PCHCBPURA'
PC2PCHCBPURA' PC3PCHCBPURA' PC0PCHCBPURM' PC1PCHCBPURM' PC2PCHCBPURM' PC3PCHCBPURM'
PC0PCHCNodeRA' PC1PCHCNodeRA' PC2PCHCNodeRA' PC3PCHCNodeRA' PC0PCHCNodeRM' PC1PCHCNodeRM'
PC2PCHCNodeRM' PC3PCHCNodeRM' PC0PCHCNodeWA' PC1PCHCNodeWA' PC2PCHCNodeWA' PC3PCHCNodeWA'
PC0PCHCNodeWM' PC1PCHCNodeWM' PC2PCHCNodeWM' PC3PCHCNodeWM' PC0PCHCNodePA' PC1PCHCNodePA'
PC2PCHCNodePA' PC3PCHCNodePA' PC0PCHCNodePM' PC1PCHCNodePM' PC2PCHCNodePM' PC3PCHCNodePM'


No Timestamp... are the power profiles alike?

Power profiles




Power seems ok as well .. running a full test for all the variables ..

Memory allocation



Look!! The memory presents a significant variation .. should we take it off?


System-wide power models




Humm it seems better than before .. let's take a closer look.

System-wide power models

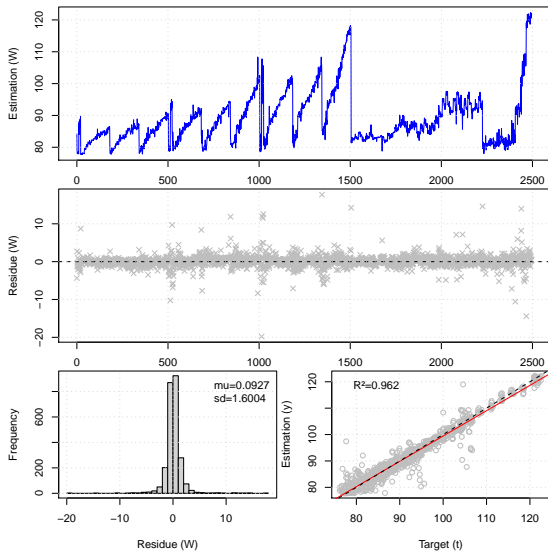
all runs with memory usage



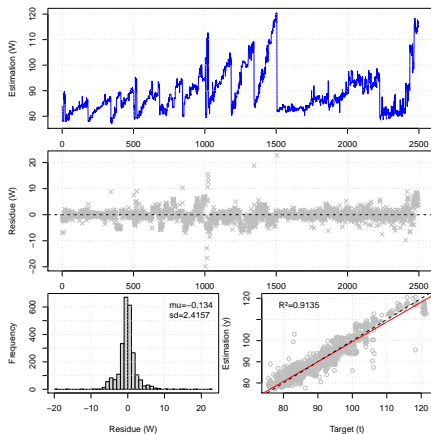
all runs without memory usage



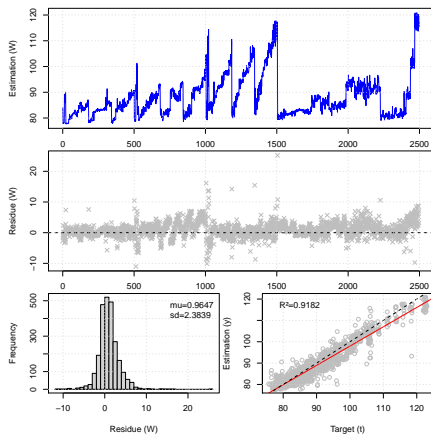
Wow wow wow .. wait! But that is just noisy data, it should have no impact on the model!



run 2



run 12



Outline

- ① Introduction
- ② Measurement and Uncertainty
- ③ System-wide power models
- ④ Conclusions

Conclusions

Conclusions:

- ▶ Machine Learning requires accurate measurements
- ▶ The number of post-processing procedures will vary according to the infrastructure
- ▶ When using external power meter, energy conversion losses need to be modeled
- ▶ Identify the most important variables to avoid modeling noise

Conclusions

Conclusions:

- ▶ Machine Learning requires accurate measurements
- ▶ The number of post-processing procedures will vary according to the infrastructure
- ▶ When using external power meter, energy conversion losses need to be modeled
- ▶ Identify the most important variables to avoid modeling noise

Ongoing research:

- ▶ Variable reduction based on the correlation between simple models and the residuals
- ▶ Validation with real use cases

Conclusions

Conclusions:

- ▶ Machine Learning requires accurate measurements
- ▶ The number of post-processing procedures will vary according to the infrastructure
- ▶ When using external power meter, energy conversion losses need to be modeled
- ▶ Identify the most important variables to avoid modeling noise

Ongoing research:

- ▶ Variable reduction based on the correlation between simple models and the residuals
- ▶ Validation with real use cases

Thank you.