

# Energy-Aware Computing Systems

*Energiebewusste Rechensysteme*

## III. Energy Demand Analysis

Timo Hönig

2018-10-31



Preface

Terminology

System Activities and Analysis

Principle of Causality, System Activities

Measurement-Based Analysis

Energy Demand Analysis

Dimensions, Resolution, and Accuracy


Physical Methods

Logical Methods


Summary



# Preface: Importance of Measurement-Based Analysis



System Model (Leonardo da Vinci, 1493)  
429 years before the first flying prototype




Physical System (Pescara, 1922)  
first flying prototype

- analysis: measure physical system to refine model **and** improve system
- measure to answer: fitting, progressing, and comparison
- by the way, the first prototype did not really fly very long...



- energy demand **analysis**
  - originates from the Ancient Greek:  
*analysis*, "a breaking up"; from *ana-* "up, throughout" and *lysis* "a loosening".
  - is the process of breaking a complex topic or substance into smaller parts in order to gain a deeper understanding of it
  - resolution of anything complex into simple elements




- **energy demand analysis**
  - dissecting a computing system and its components with regards to energy demand
  - quantifying energy demand by applying **physical** and **logical** analysis methods
  - disclosure of cause and effect



# Principle of Causality

- causal chain of events related to some other event  $e_i$ :



- $A$ ,  $B$  and  $C$  are system activities
- cause and effect relationship of system activities
- activities ("cause") influence the system state  $\rightarrow$  impact on the energy demand of sequel activities ("effect")



- system level activity → effects are subject to analysis
  - component-based energy demand analysis
  - correlation with software-level activities (i.e. process execution)
- quantify energy demand on the background of different aspects
  - level of activity (i.e. timers)
  - impact on the system (i.e. cache trashing)
    - ↔ first-level analysis (in isolation)
- determine overall relevance of system activity
  - identify (and rule out) cause-and-effect relationships
    - ↔ second-level analysis (with side activities)
- **measurement-based** analyses to investigate consequences with regards to the energy demand is the focus of today's lecture



# Measurement-Based Analysis

- measurement-based analysis of a physical system




Figure 2: Block Diagram of System Modeling Procedure

Wang '93 [2]



# Measurement-Based Analysis

- measurement-based analysis of a physical system

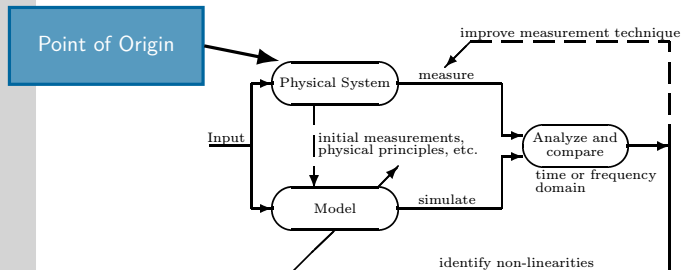


Figure 2: Block Diagram of System Modeling Procedure

Wang '93 [2]

- point of origin: **physical system** (e.g., prototype)
- apply measurements to analyze system properties and impact of external stimuli (i.e., input)



# Measurement-Based Analysis

- measurement-based analysis of a physical system




Figure 2: Block Diagram of System Modeling Procedure

Wang '93 [2]

- goals: improve physical **system**, build **model**, improve **measurement**, identify **input dependencies**
- pre-post comparison, analyze: **actual vs. target**



# Measurement-Based Analysis

- measurement-based analysis of a physical system

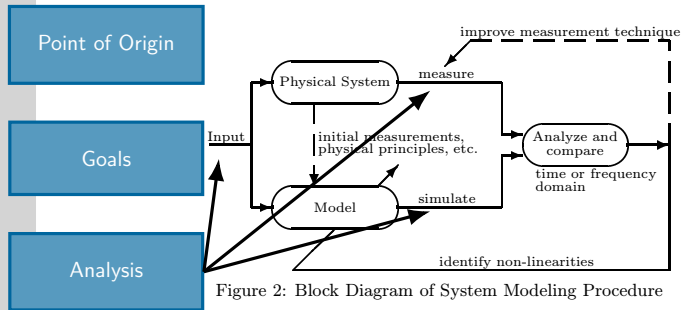


Figure 2: Block Diagram of System Modeling Procedure

Wang '93 [2]

- analysis: iterative approach to build either **refined systems**, **improved models**, or **both** (depending on goal)
- improve **first-order goals**, consider **second-order constraints**



# Dimensions, Resolution, and Accuracy


- energy demand analysis at system level
- *what* system properties to measure?
  - energy demand → requires knowledge of power over time
  - power demand → requires data of circuitry (electric current, voltage)
  - ...
  - time demand → timing **and** time requirements
  - thermal impact and the duality thereof
    - ...minimum/maximum/average of the above system properties
- *how* to measure?
  - physical measurements
  - logical measurements
    - ↔ interlude: measuring errors, wrong statistics, avoiding false conclusion
- consider: measuring resolution and accuracy
  - resolution: ability to differentiate between two similar measured values
  - accuracy: deviation of measured values from the physically ground truth



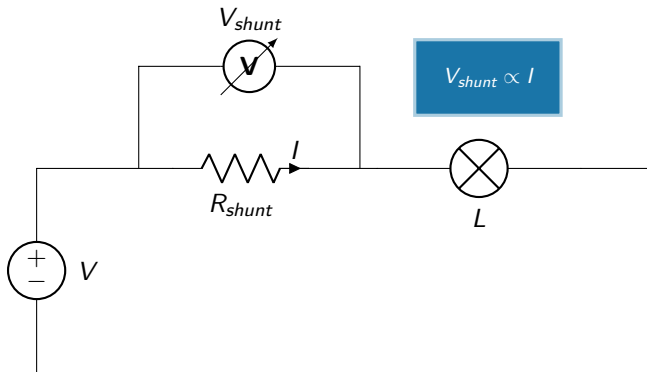
- measuring energy demand: commonly based on measurements which quantify the power demand, integrate power demand over time
- average power · time interval (i.e. execution time)  $\Rightarrow$  energy demand
- different physical measurement methods (direct/indirect)
  - shunt based measurement methods
  - Hall effect measurement methods
  - current mirror based measurement methods
  - ...
- measuring the energy demand is a first step towards energy awareness **and** towards improving energy efficiency



- shunt based measurement; indirection: **current measurement**
- analyze **electric current** in a circuit with constant DC voltage
- idea: build **low resistance path** with a resistor (shunt), measure **voltage drop(s)** across shunt
- determine **electric current** by applying Ohm's law




- known: resistance ( $R$ ); sense: DC voltage ( $V$ )
- calculate: current  $I$  (...and electrical power  $P$ )




- system setup and preparations
  - verify baseline setup: functioning *without* alteration (i.e. without shunt)
  - identify hardware component that is subject to analysis
  - integrate measuring setup (i.e., assemble shunt resistor)
  - determine proper functioning *without* measuring device (i.e. without voltmeter)
- measuring procedure
  - activate and instruct hardware to execute workload → trigger signals
  - analyze setup under varying external conditions
  - reveal non-linearities of measured values
- determining power and energy demand
  - time series
  - calculate current, (average) power, and energy demand
  - apply statistics on measurement data



$$V_{shunt} \propto I$$

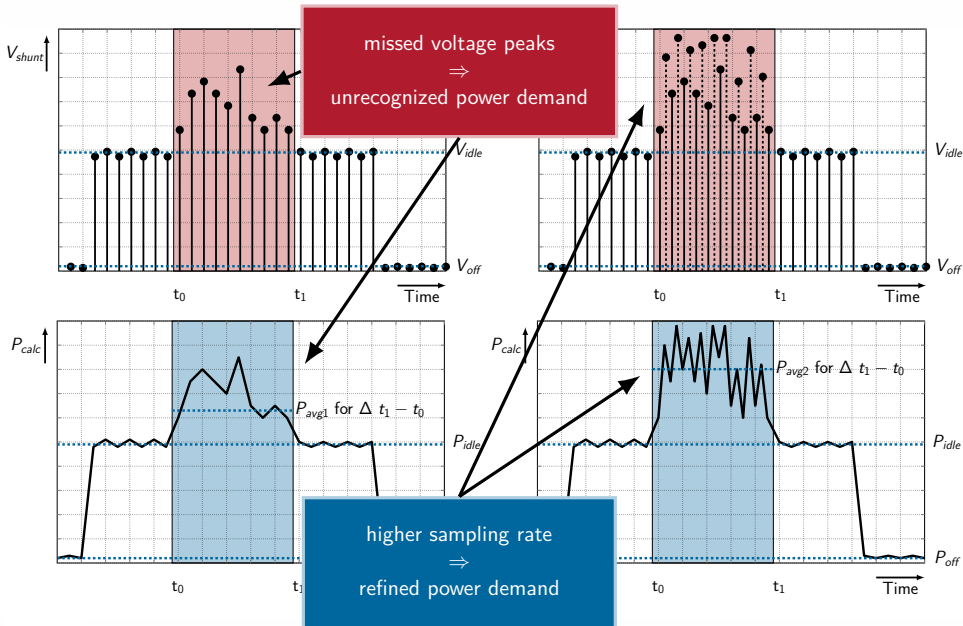


- voltage drop ( $V_{shunt}$ ) across the resistor ( $R_{shunt}$ ) is proportional to the electric current ( $I$ )
- sampling  $V_{shunt}$  is subject to special focus...



# Measuring Energy with a Resistor (Shunt)

Physical



- physical energy demand measurement methods
  - isolated: little overhead 😊, several systems 😞
  - alteration and/or extension of original baseline setup
    - difficult access to power rails in designs using ball grid arrays (e.g., package-on-package, PoP)
    - internal wires inaccessible, for example, when measuring the energy demand of peripherals of a system on chip (SoC)
  - extra efforts, experience in electrical engineering necessary
- logic energy demand measurement methods
  - integrated: overhead-prone 😞, single system 😊
  - energy models, event-based analyses, maintaining of the hardware setup
    - qualitative statements are sufficient for simple analyses
    - correlation of occurrence and frequency of logic (software) events with power or energy demand of hardware components
  - energy models are, despite their limits, often the first choice



- build logical energy model on empirical knowledge as obtained by measurements or calculations based on measurements
  - ↔ requires physical measurements beforehand to establish the model
- static/dynamic components of logical energy demand measurements
  - cost model for specific events in the system
    - execution of a (specific) instruction (within a given system state)
    - transmission of a network packet
    - ...
  - presence of event(s), time and frequency of occurrence
- measurement accuracy depends on quality of energy model
  - consideration of non-deterministic effects (e.g., thermal aspects, impact of unconsidered system activities, state of caches)
  - logical energy models must adapt to the hardware platform and to its specific usage scenario (i.e., system complexity)



- measuring energy demand: identify occurrence of events using hardware performance monitoring counters (PMCs)
- each PMC is configured to measure the occurrence of a particular event (e.g. retired instructions, data cache misses, TLB misses etc.)
- intended use of PMCs: performance analysis
  - ▶ Intel Corporation  
**Intel 64 and IA-32 Architectures Software Developer's Manual Volume 3B: System Programming Guide, Part 2.**
- energy modeling using performance counters
  - event  $\Rightarrow$  demand of a certain amount of energy
  - register relevant events and their frequency/total number of occurrences
  - ▶ Frank Belossa  
**The Benefits of Event-Driven Energy Accounting in Power-Sensitive Systems**  
*Proceedings of the 2000 ACM SIGOPS European Workshop „Beyond the PC: New Challenges for the Operating System”, 2000.*






FIG. 1. Correlation of retired microinstructions and energy consumption

Bellosa '00 [1]

- workload: integer operations; microinstructions  $\Rightarrow$  energy demand
- calibration: multimeter (integrated power monitor, 1 Watt resolution)






FIG. 3. Correlation of L2 Cache references and energy consumption

Bellosa '00 [1]

- workload: varying integer and memory operations
- second level cache  $\Rightarrow$  energy demand



- Intel Running Average Power Limit (RAPL)
- between the worlds: logical and physical measurements
  - originally, RAPL was using a software power model → **logical** measurements with hardware performance counters and I/O models
  - recent Intel CPUs (i.e. Haswell and onwards) → **physical** measurements
- hybrid approach towards energy-aware systems
  - adjusting performance levels (i.e. dynamic voltage and frequency scaling)  
⇒ impacting power demand
  - adjusting power levels (i.e. power capping)  
⇒ impacting performance



- physical and logical methods to measure the energy demand
  - analysis method strictly depends on system and use case
  - alternatives that augment and complement each other (i.e. verification)
- isolation and partitioning
  - separate measuring device from device under test (physical)
  - determine and quantify the influence of the measurement (logical and physical)
- overhead and side-effects
  - increased resource demand, system slowdown **or** speedup
  - interrelation to otherwise unrelated system components



## Interlude: Measuring Done Right

- performing measurements is good, but...
- common measurement problems
  - measuring the **wrong thing**
  - drawing inappropriate conclusions
  - using bad statistics
  - ignoring **system interaction**
  - ignoring timing granularity
- ...and further considerations
  - comparing apples to oranges
  - comparing end-to-end measurements with the sum of parts
  - using the **wrong metrics**
  - **mistakes**

▶ Margo Seltzer and Aaron Brown

### **Measuring Computer Systems: How to Measure Performar**

*Proceedings of the annual conference on USENIX Annual Technical Conference, 1997.*

## Mars Probe Lost Due to Simple Math Error

By ROBERT LEE HOTZ  
TIMES SCIENCE WRITER

NASA lost its \$125-million Mars Climate Orbiter because spacecraft engineers failed to convert from English to metric measurements when exchanging vital data before the craft was launched, space agency officials said Thursday.

A navigation team at the Jet Propulsion Laboratory used the metric system of millimeters and meters in its calculations, while Lockheed Martin Astronautics in Denver, which designed and built the spacecraft, provided crucial acceleration data in the English system of inches, feet and pounds.

As a result, JPL engineers mistook acceleration readings measured in English units of pound-seconds for a metric measure of force called newton-seconds.

In a sense, the spacecraft was lost in translation:

"That is so dumb," said John  
Please see MARS, A.35

- **complex systems** require thorough understanding of individual system aspects to allow focussed **energy demand analyses**
- **physical** and **logical** energy demand measurements have individual benefits and are often complementary
- available analysis **methods** must be suitable for individual use
- reading list for Lecture 4:
  - ▶ Andreas Weissel and Frank Bellosa  
**Process Cruise Control:  
Event-Driven Clock Scaling for Dynamic Power Management**  
*Proceedings of the International Conference on Compilers,  
Architecture and Synthesis for Embedded Systems (CASES), 2002.*



- [1] BELLOSA, F. :  
The Benefits of Event-Driven Energy Accounting in Power-Sensitive Systems.  
In: *Proceedings of the 2000 ACM SIGOPS European Workshop „Beyond the PC: New Challenges for the Operating System” (EW '00)* ACM, 2000, S. 37–42
- [2] WANG, F. ; ABRAMOVITCH, D. ; FRANKLIN, G. :  
A Method for Verifying Measurements and Models of Linear and Nonlinear Systems.  
In: *Proceedings of the 1993 IEEE American Control Conference*, 1993, S. 93–97

