Energy-Aware Computing Systems

Energiebewusste Rechensysteme

XI. Infrastructure

Timo Hönig

2019-01-23

Agenda

Preface and Terminology

Linking (with) the Physical World

- Dependencies and Impact
- Considerations and Metrics

Infrastructure

- Temperature-Aware Workload Placement
- Building Operating System Services (BOSS)
- Runtime System for Heterogeneous HPC Clusters

Call for Participation

Uncharted Lecture next Wednesday

Summary

Infrastructure

- **motivation**
  - indirect resource demand → costs
  - „many a little makes a mickle"
- **operational goals**
  - identify necessary operating conditions
  - ...but invest reasonable efforts, only

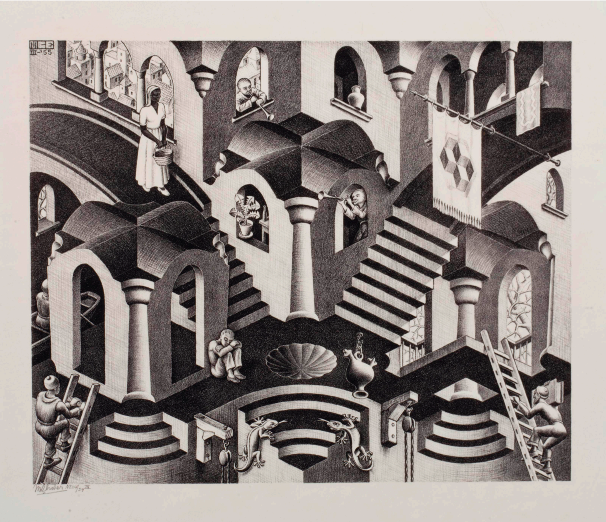

Dependencies and Impact

- dependency on external factors
  - device (e.g., power supplies)
    - $\rightarrow$ *moving parts* (e.g., fans)
      - $\rightarrow$ kinetic energy (e.g., heat)
  - physical properties (ambient air)
  - not all environmental conditions can be controlled

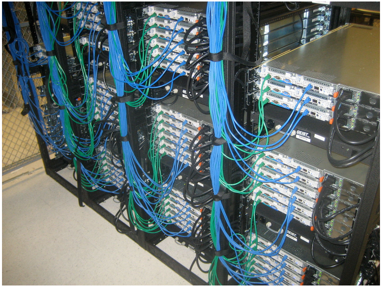

- impact on external systems
  - control resource demand
    - temporal delay
    - workload (re)positioning
  - control types
    - implicit: activation of additional resources
    - explicit: system control

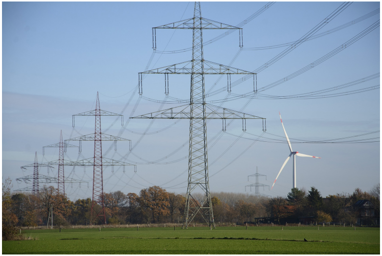

Considerations and Metrics

- common infrastructure considerations
  - power demand of computing systems → secondary energy (i.e., heat)
  - heating, ventilation, and air conditioning (HVAC) **required**
    ↪ especially heat distribution
- common metric: power usage effectiveness (PUE)
  - PUE := total energy demand / energy demand of computing systems
  - total energy demand includes infrastructure, HVAC etc.
  - perfect system: PUE <= 1.0
    ↪ electricity generation
- PUE refinements
  - mixed use of renewable and non-renewable energy
  - reuse of secondary energy
  - various additions and alternatives
    - Green Energy Coefficient (GEC)
    - Energy Reuse Factor (ERF)
    - Carbon Usage Effectiveness (CUE)

Temperature-Aware Workload Placement

- Motivation
  - reduce cooling costs in data centers
  - workload placement → hot spots and cold spots in data centers
  - cf. thermal-aware task scheduling [2], Lecture 7
- Temperature-Aware Workload Placement [5]
  - temperature-aware workload placement algorithms
  - reduce cooling costs over the worst-case placement by almost 50 %

▶ J. Moore et al.
Making Scheduling "Cool": Temperature-Aware Workload Placement in Data Centers
Proceedings of the USENIX Annual Technical Conference (ATC '05), 2005.

Building Operating System Services (BOSS)

- Motivation
  - energy demand of buildings: 73 % (2011) in the US (data centers: 3 %)
  - reduce costs of commercial buildings using operating systems for buildings
  - cf. occupancy-driven energy management [1], Lecture 6
- Building Operating System Services (BOSS) [3]
  - distributed operating system plus services
  - query language and control infrastructure

▸ S. Dawson-Haggerty et al.
BOSS: Building Operating System Services
Proceedings of the 10th USENIX Conference on Networked Systems Design and Implementation (NSDI '13), 2013.

Runtime System for Heterogeneous HPC Clusters

- Motivation
  - variable power pricing becomes increasingly common
  - exploit dynamic pricing to reduce operating costs of HPC systems
- Runtime System for Heterogeneous HPC Clusters (Albatross) [4]
  - implement different operation modes for varying electricity pricing
  - use heterogeneous compute components to execute workloads

- T. Hönig et al.
  How to Make Profit:
  Exploiting Fluctuating Electricity Prices with Albatross, A Runtime System for Heterogeneous HPC Clusters
  Proceedings of the 8th International Workshop on Runtime and Operating Systems for Supercomputers (ROSS 2018), 2018.

Subject Matter

- considerations on **infrastructure** are necessary during design and operation of energy-aware systems
- **metrics** help to analyse the efficiency → use different metrics depending on **actual systems** and **infrastructure**
- **exploit external factors** for **own advantage** at system level

- Evaluation: please **participate** until end of this week
- Uncharted Lecture by Ralph Schlenk (Nokia) on **energy-efficient optical networks**
  - Wednesday, 30. January 2019, 14:30 s.t.
  - room 02.037 („e-Studio"), RRZE (second floor)

Reference List I

[1] AGARWAL, Y. ; BALAJI, B. ; GUPTA, R. ; LYLES, J. ; WEI, M. ; WENG, T. :
Occupancy-driven Energy Management for Smart Building Automation.
In: *Proceedings of the 2nd Workshop on Embedded Sensing Systems for Energy-Efficiency in Building (BuildSys '10)*, 2010, S. 1–6

[2] CHOI, J. ; CHER, C.-Y. ; FRANKE, H. ; HAMANN, H. ; WEGER, A. ; BOSE, P. :
Thermal-aware Task Scheduling at the System Software Level.
In: *Proceedings of the 2007 International Symposium on Low Power Electronics and Design (ISLPED'07)*, 2007, S. 213–218

[3] DAWSON-HAGGERTY, S. ; KRIOUKOV, A. ; TANEJA, J. ; KARANDIKAR, S. ; FIERRO, G. ; KITAEV, N. ; CULLER, D. :
BOSS: Building Operating System Services.
In: *Proceedings of the 10th USENIX Conference on Networked Systems Design and Implementation (NSDI '13)*, 2013, S. 443–458

[4] HÖNIG, T. ; EIBEL, C. ; WAGENHÄUSER, A. ; WAGNER, M. ; SCHRÖDER-PREIKSCHAT, W. :
How to make profit: Exploiting fluctuating electricity prices with Albatross, a runtime system for heterogeneous HPC clusters.
In: *Proceedings of the 8th International Workshop on Runtime and Operating Systems for Supercomputers (ROSS'18)* ACM, 2018, S. 1–8

Reference List II

[5] Moore, J. ; Chase, J. ; Ranganathan, P. ; Sharma, R. :
Making Scheduling "Cool": Temperature-aware Workload Placement in Data Centers.
In: *Proceedings of the USENIX Annual Technical Conference (ATC '05)*, 2005, S. 61–75