Quo Vadis HPC?

HPC Co-Design as Service for Science

Prof. Dr. Martin Schulz

Technical University of Munich (TUM)

Chair for Computer Architecture and Parallel Systems

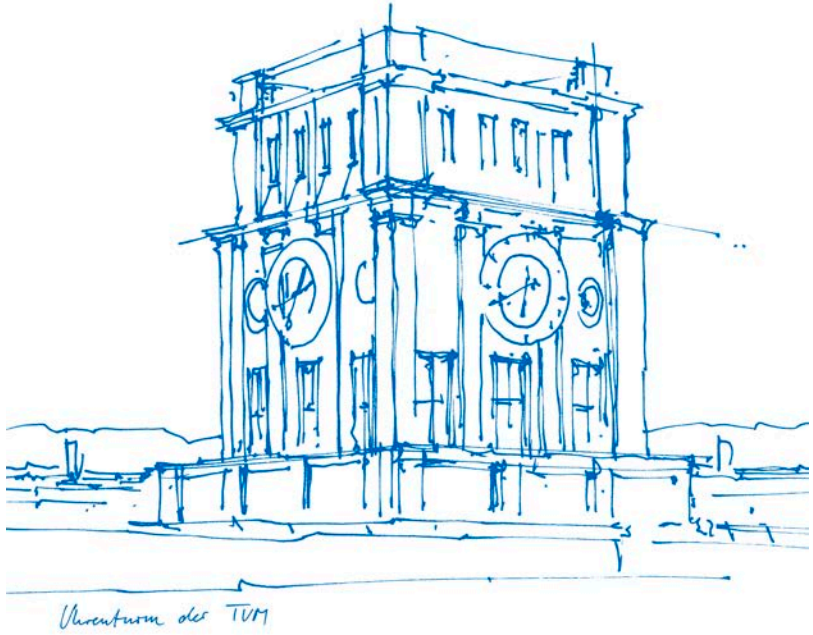

SuperMUC-NG
Next-Gen Science Symposium
Leibniz Supercomputing Centre, Garching
November 22, 2018

HPC is the Backbone for Modern Science

Modeling and simulation is the third pillar of science.

It complements experiments and theory
and enables new scientific understanding.

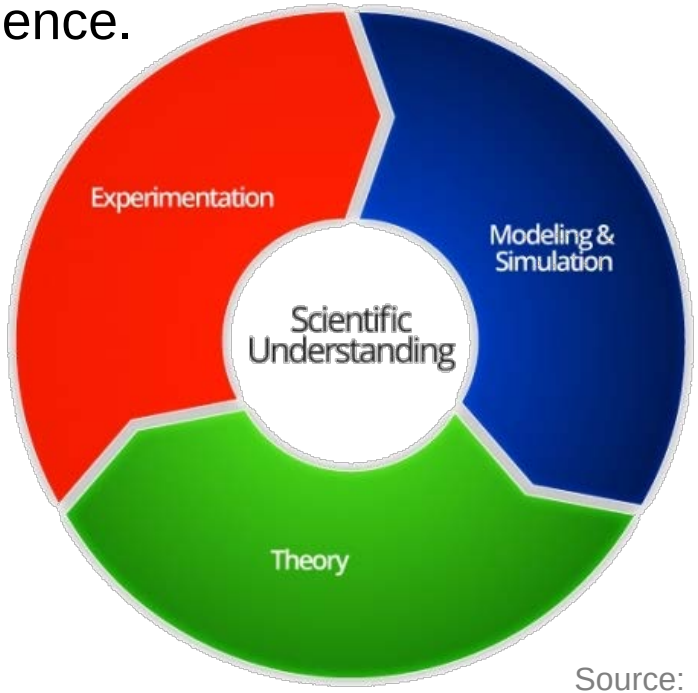

Source:
US Department of Energy

Example: Material Solidification Process

Molecular dynamics at LLNL with ddcMD: 2 Million atom run (2005)

- Used all of Blue Gene/L (128K cores) with close to perfect scaling
- Led to new scientific observations

Lawrence Livermore National Laboratory

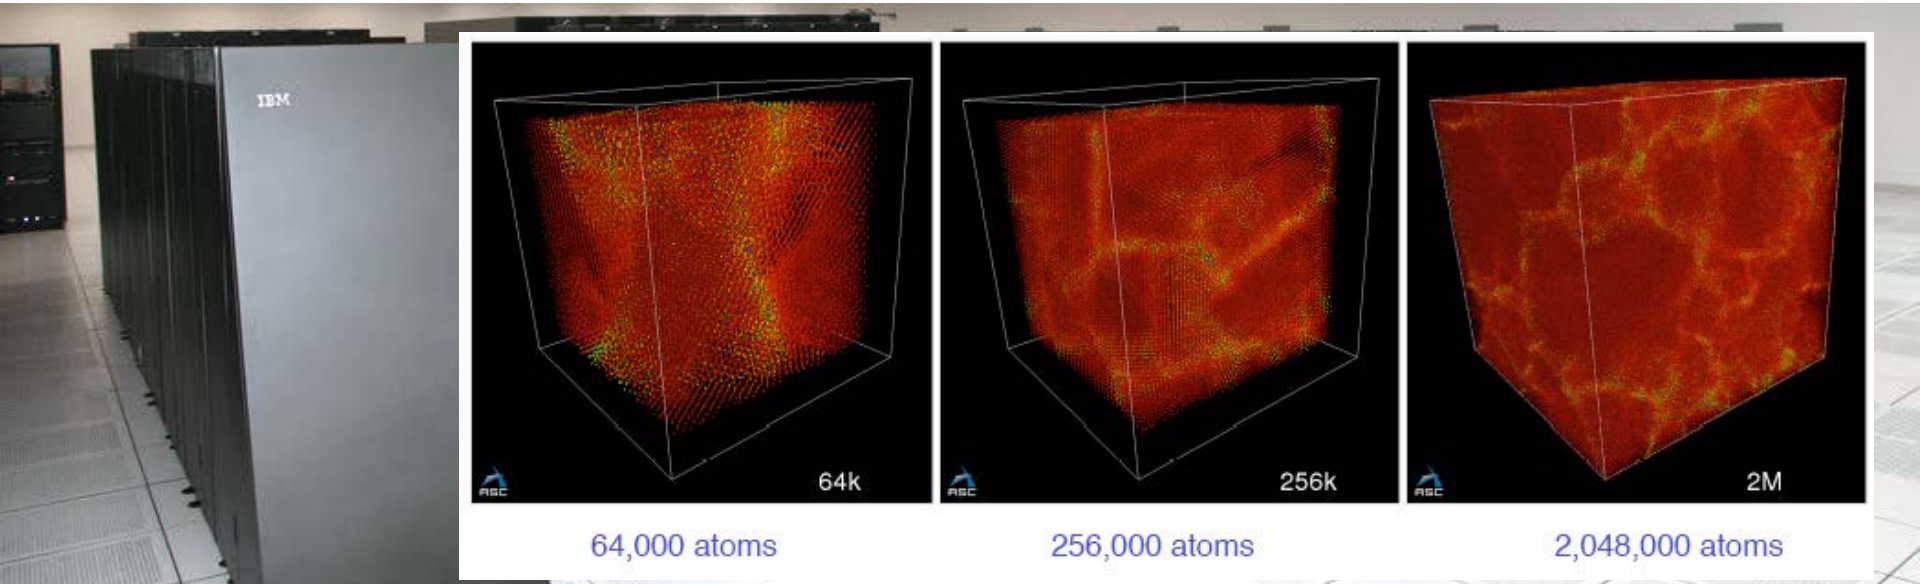

https://e-reports-ext.llnl.gov/pdf/336469.pdf

HPC is the Backbone for Modern Science

Modeling and simulation is the third pillar of science.

It complements experiments and theory and enables new scientific understanding.

Modern Science is data driven.

We need to support workflows from data acquisition to analysis and visualization.

Big Data is a new critical field.

The demands on HPC resources keep growing.

This requires **strong HPC centers** to lead.

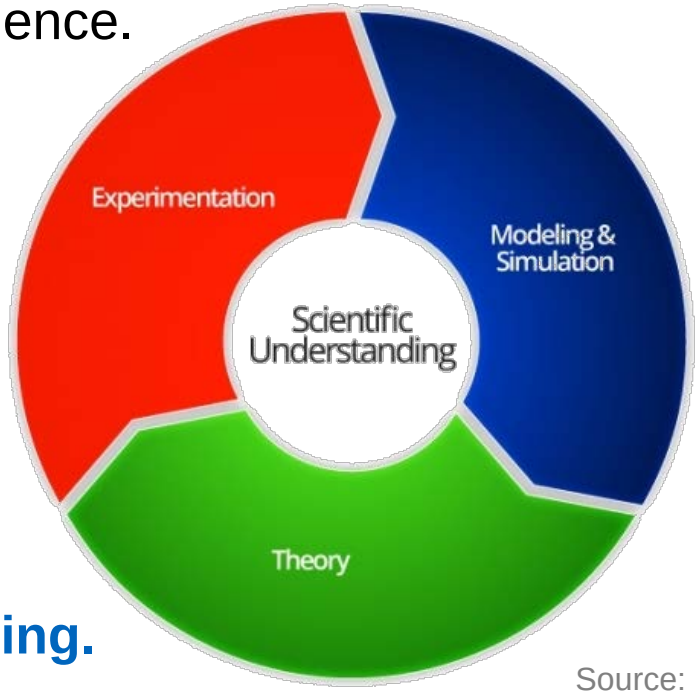

Source:
US Department of Energy

Top 500 / Nov. 2018

USA: Summit (ORNL)

China: Sunway TaihuLight (NSC Wuxi)

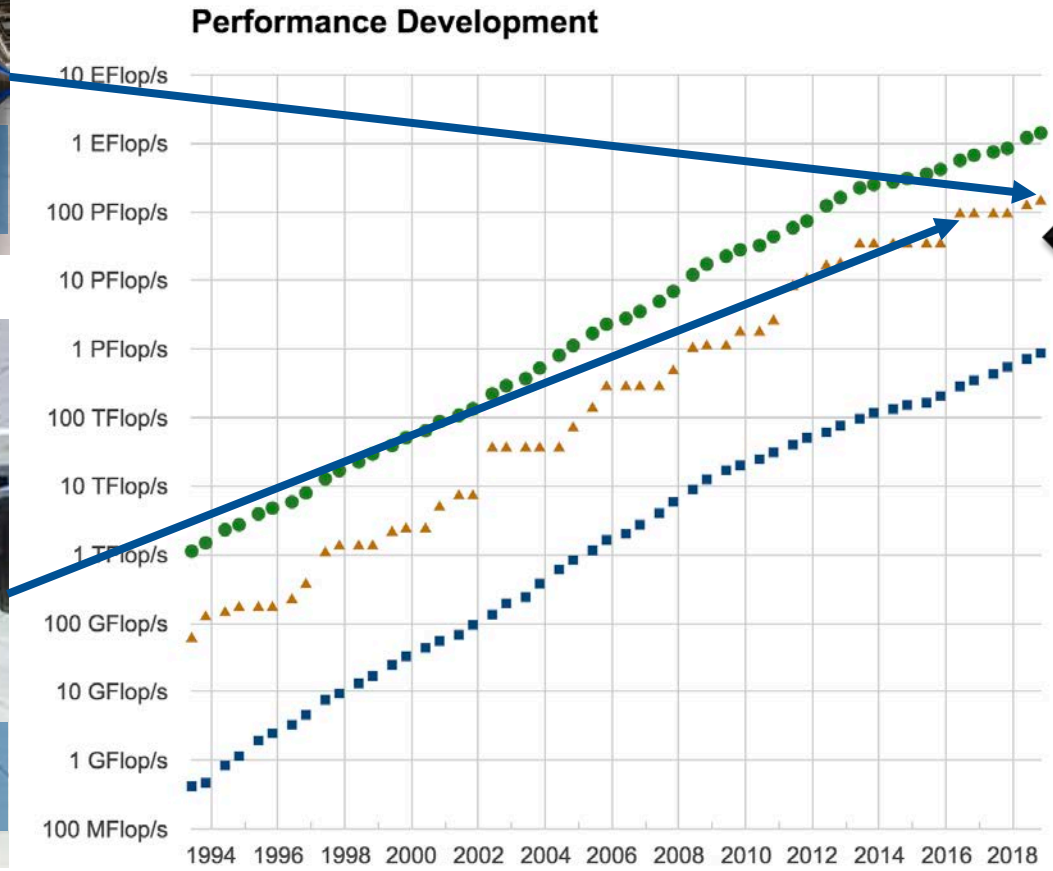

Germany: SuperMUC-NG 19.5 Pflop/s Linpack

Source: Top500, 11.2018

SuperMUC-NG

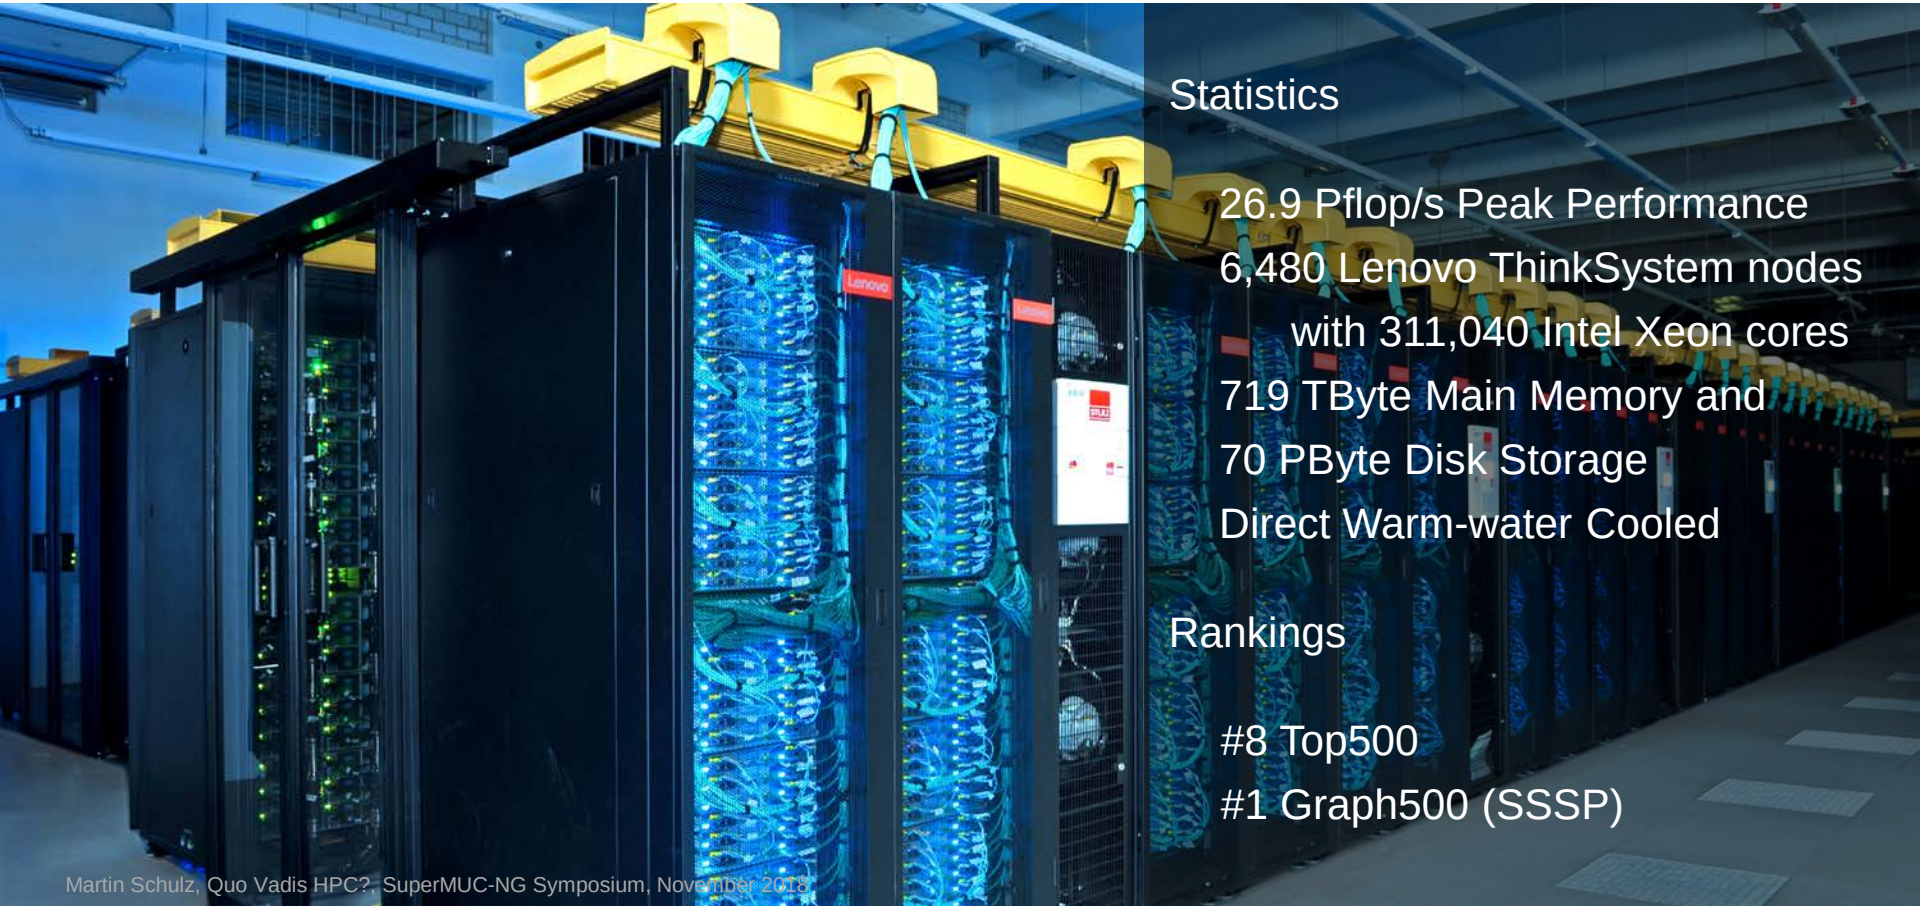

Statistics

26.9 Pflop/s Peak Performance
6,480 Lenovo ThinkSystem nodes
with 311,040 Intel Xeon cores
719 TByte Main Memory and
70 PByte Disk Storage
Direct Warm-water Cooled

Rankings

#8 Top500
#1 Graph500 (SSSP)

HPC Top500 Developments

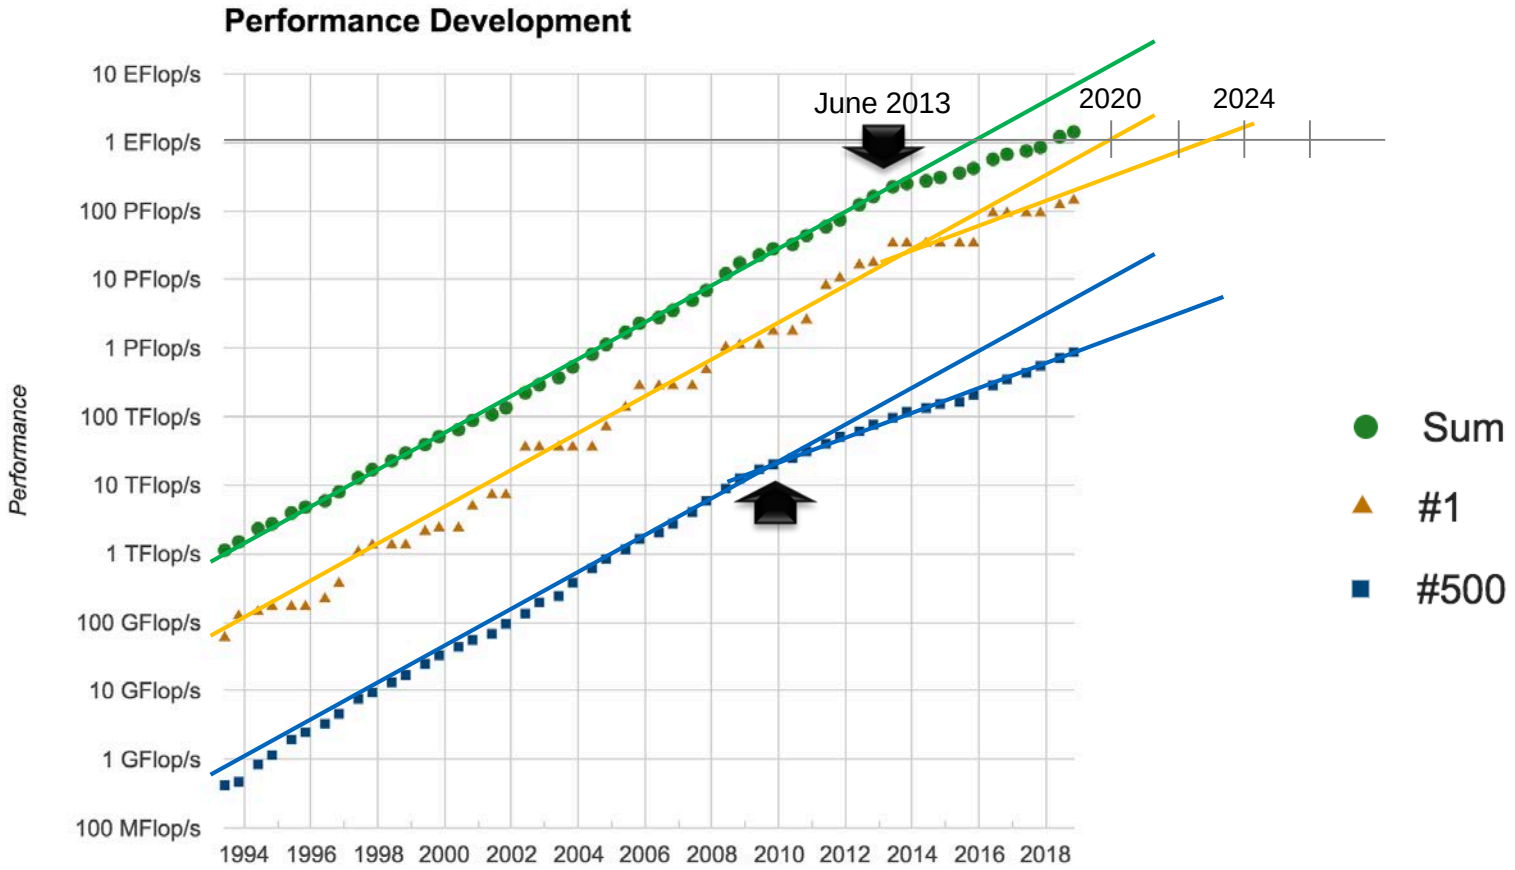

Source: Top500

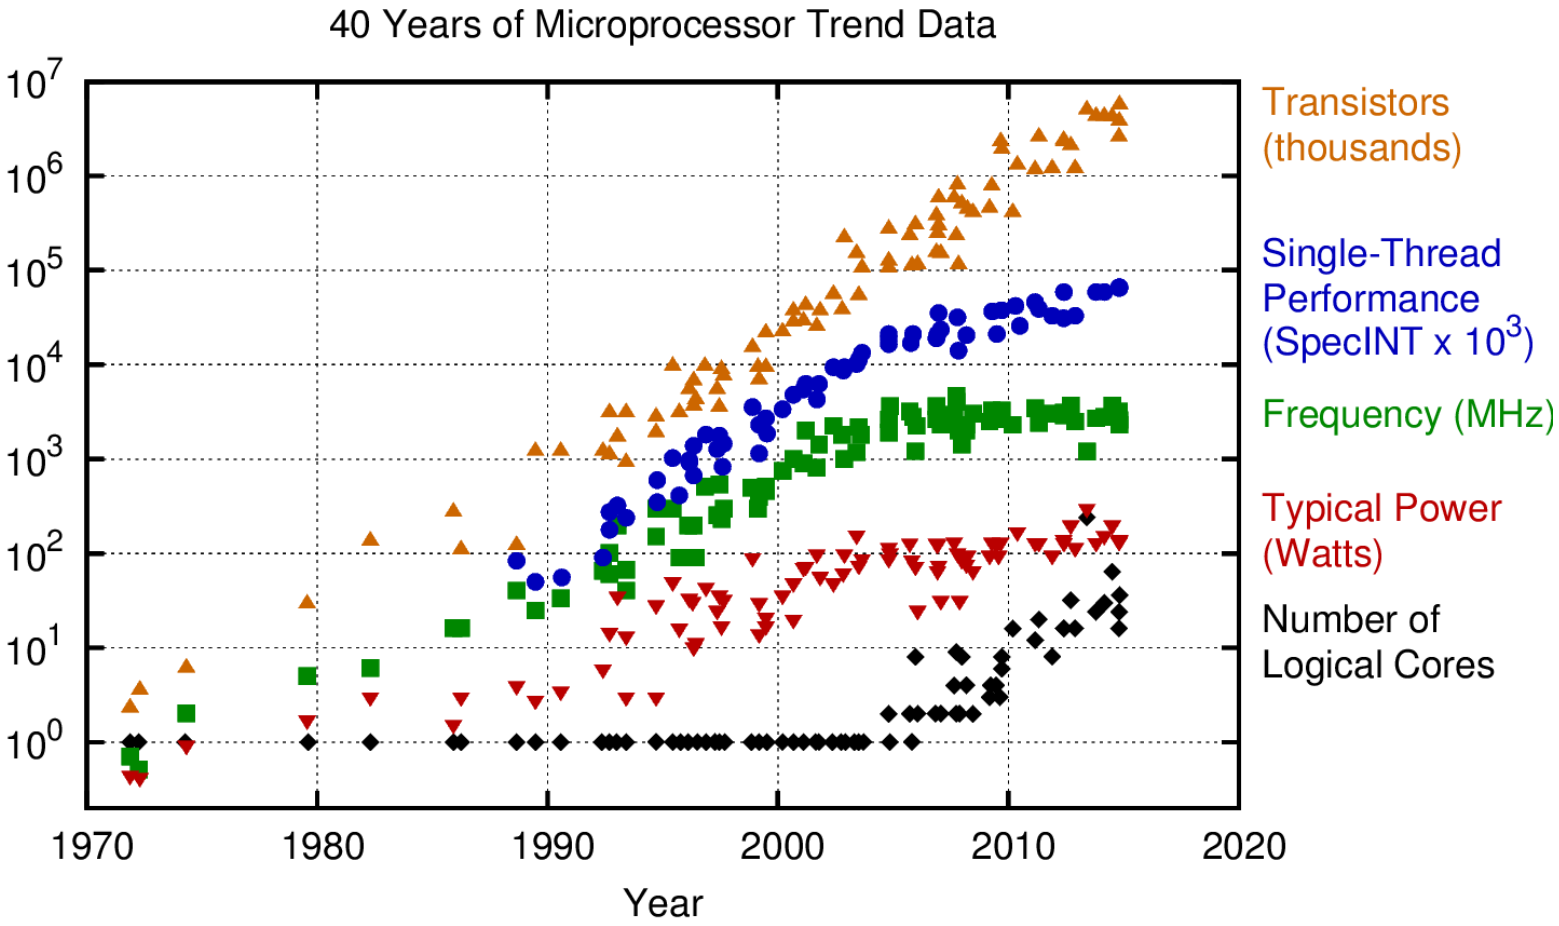

Original data up to the year 2010 collected and plotted by M. Horowitz, F. Labonte, O. Shacham, K. Olukotun, L. Hammond, and C. Batten
New plot and data collected for 2010-2015 by K. Rupp

Turning Points

Breakdown of Dennard Scaling!

> “Power density of transistors stays constant”

Already happened around 2007
Consequence: Parallelism

Breakdown of Moore‘s Law?

> “The number of transistors will double every 18 months”

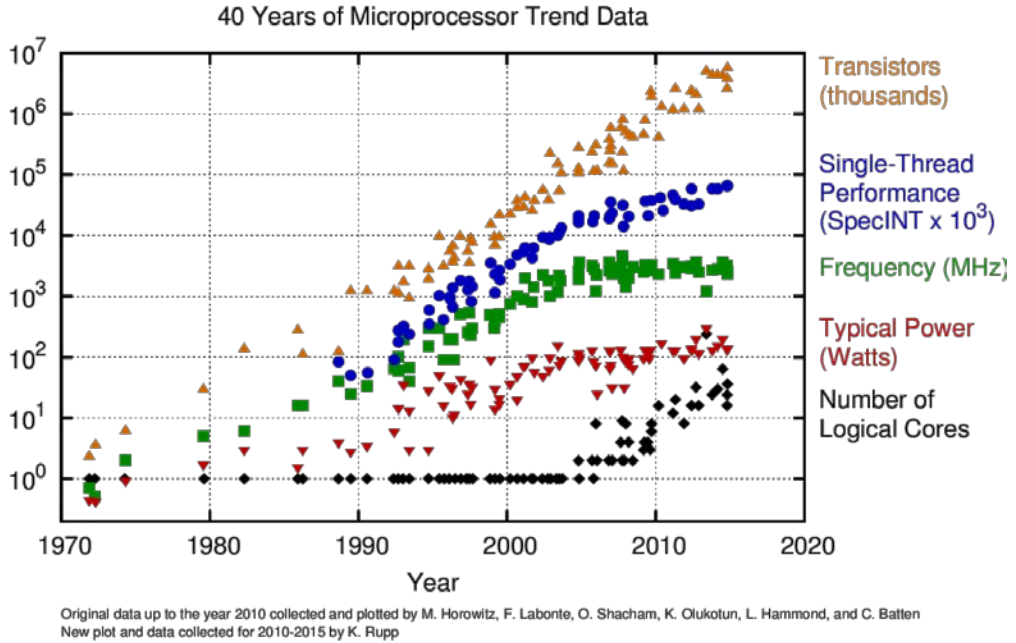

Current push to 7nm process is rocky.
What comes after it?

How Can The HPC Community Bypass These Barriers?

The demands of modeling & simulation keep growing

Strong HPC is imperative to academia and industry.
Stagnated HPC endangers innovation.

Centers need to impact architecture development

Buying off-the-rack is no longer a viable concept.

HPC research must drive development

In hardware and software with a commitment to increased efficiency.

Efficient computing = More science

Application-Driven Co-Design as a Service to Science

Understanding	Influencing	Advancing
Workloads Better	Vendor Designs	Software Co-Design

Understanding Workloads Better

HPC systems must support a large and diverse set of workloads

- Need to capture and understand application workloads.
- Need to capture the impact of workloads on the system.

LRZ DCDB project provides a firm basis

Goal: extension to continuous monitoring

Challenge: analytics of the gathered data

It's its own big data problem!

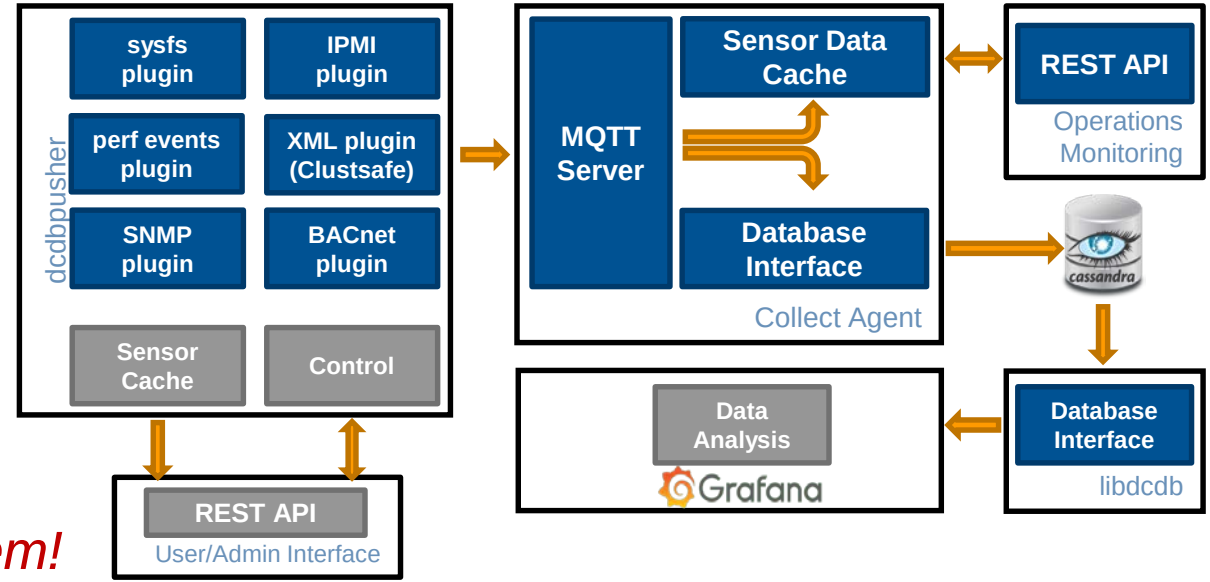

Source: Michael Ott, Daniele Tafani, LRZ

Influencing Vendor Designs

We need to use our position to influence future architecture generations

- Progress on conventional CPUs
- Evolution of GPUs
- Design of accelerators beyond GPUs
  - Leaving von Neumann
  - Quantum Computing
- Advancement of networks and I/O
- Development of new packaging solutions

This requires symbiotic efforts to leap forward

- Research in modeling and analysis
- Early interaction with vendors
- Trusted, long-term relationships

Advancing Software Co-Design

Hardware advances alone are insufficient

Need system software to map applications to hardware.
Need to help applications exploit advanced architectures.

System Software Co-Design

Evolution of Programming Models via Standards
I/O and Visualization Systems
Resource Manager

International Collaborations

Imperative on the road to exascale
Critical mass for guiding vendor design
Knowledge exchange for risk reduction

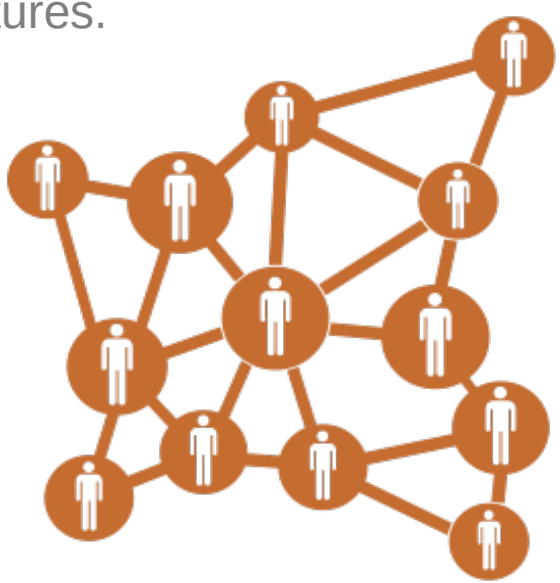

The PowerStack Initiative

Managing power & energy
is critical to achieve exascale

Need hierarchical systems to
cover a wide spread set of scenarios

International initiative to identify

- … Common terminology
- … Compatible components
- … Site-specific policies

To bring research to production

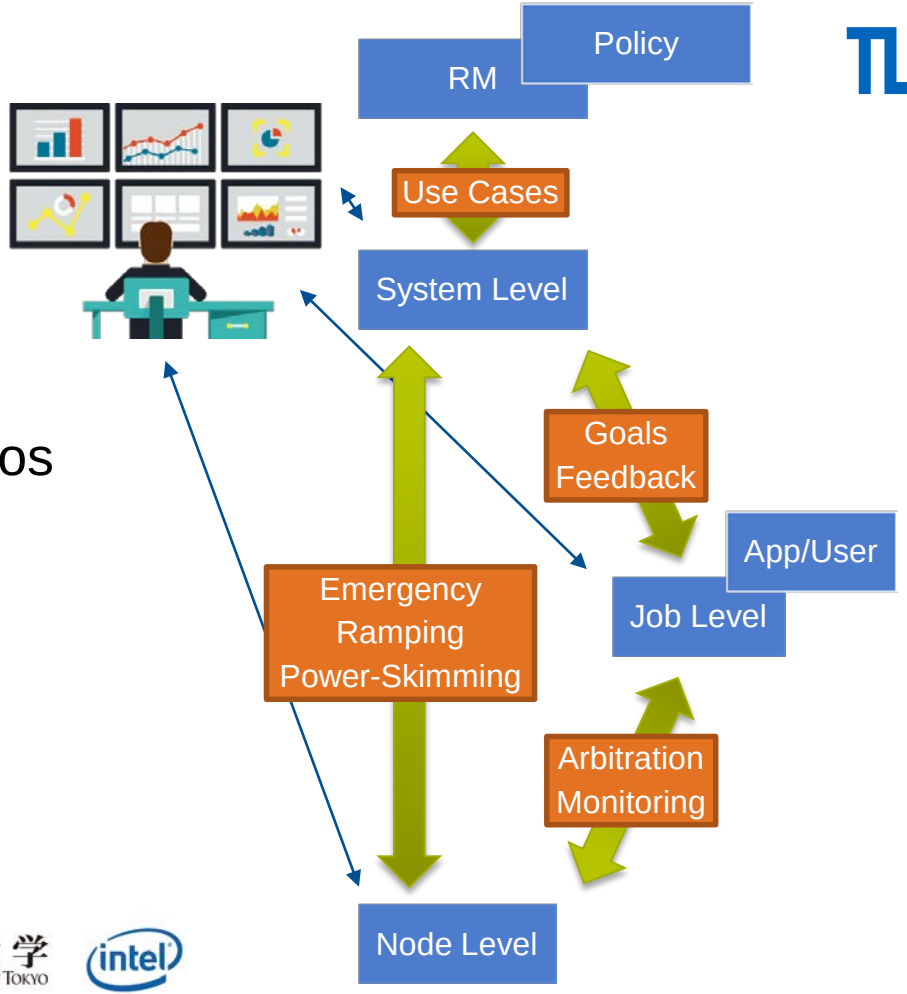

Quo Vadis HPC?

Partnerships with Hardware Vendors
- HPC is (again) no longer a commodity
- Focus on targeted solutions, which are co-developed

Partnerships with Software Communities
- Strive for an open and widely available ecosystem
- Involvement in standards where useful

Partnerships with Application Teams
- Driver must be application requirements
- Continuously capture workloads

Integrated Co-Design research by HPC Centers in close collaboration with universities

Quo Vadis HPC?

HPC centers as leaders driving the …

- ... advancement of HPC architectures
- ... advancement of HPC ecosystems
- ... advancement of HPC industry

All of which
deliver the advancement of science

To do this, we must forge …

- ... earlier engagement with vendors
- ... long-term partnerships
- ... strong ties with universities
- ... global collaborations