

High Performance Computing for Science & Engineering

'Bridges between Experiment and Theory'

Prof. Dr. - Ing. Morris Riedel

Head of Research Group 'High Productivity Data Processing'
Juelich Supercomputing Centre, Germany

Adjunct Associated Professor
School of Engineering and Natural Sciences, University of Iceland




UNIVERSITY OF ICELAND
SCHOOL OF ENGINEERING AND NATURAL SCIENCES

FACULTY OF INDUSTRIAL ENGINEERING,
MECHANICAL ENGINEERING AND COMPUTER SCIENCE

High Performance Computing

Why do we need fast computing?



*Numerical calculations... Model
...simulation over time*

Experiment
'we observe
the nature'

Theory
'we create
a model
of nature'



Fast = n Floating Point Operations (FLOP) per one second

1 FLOP/s
Example:
 $93386.25 * 10^8$
= ???

>5.000.000.000.000.000
FLOP/s
~ 500.000 cores
~ 2015

1.000.000.000.000.000 FLOP/s
~295.000 cores~2008 (JUGENE)

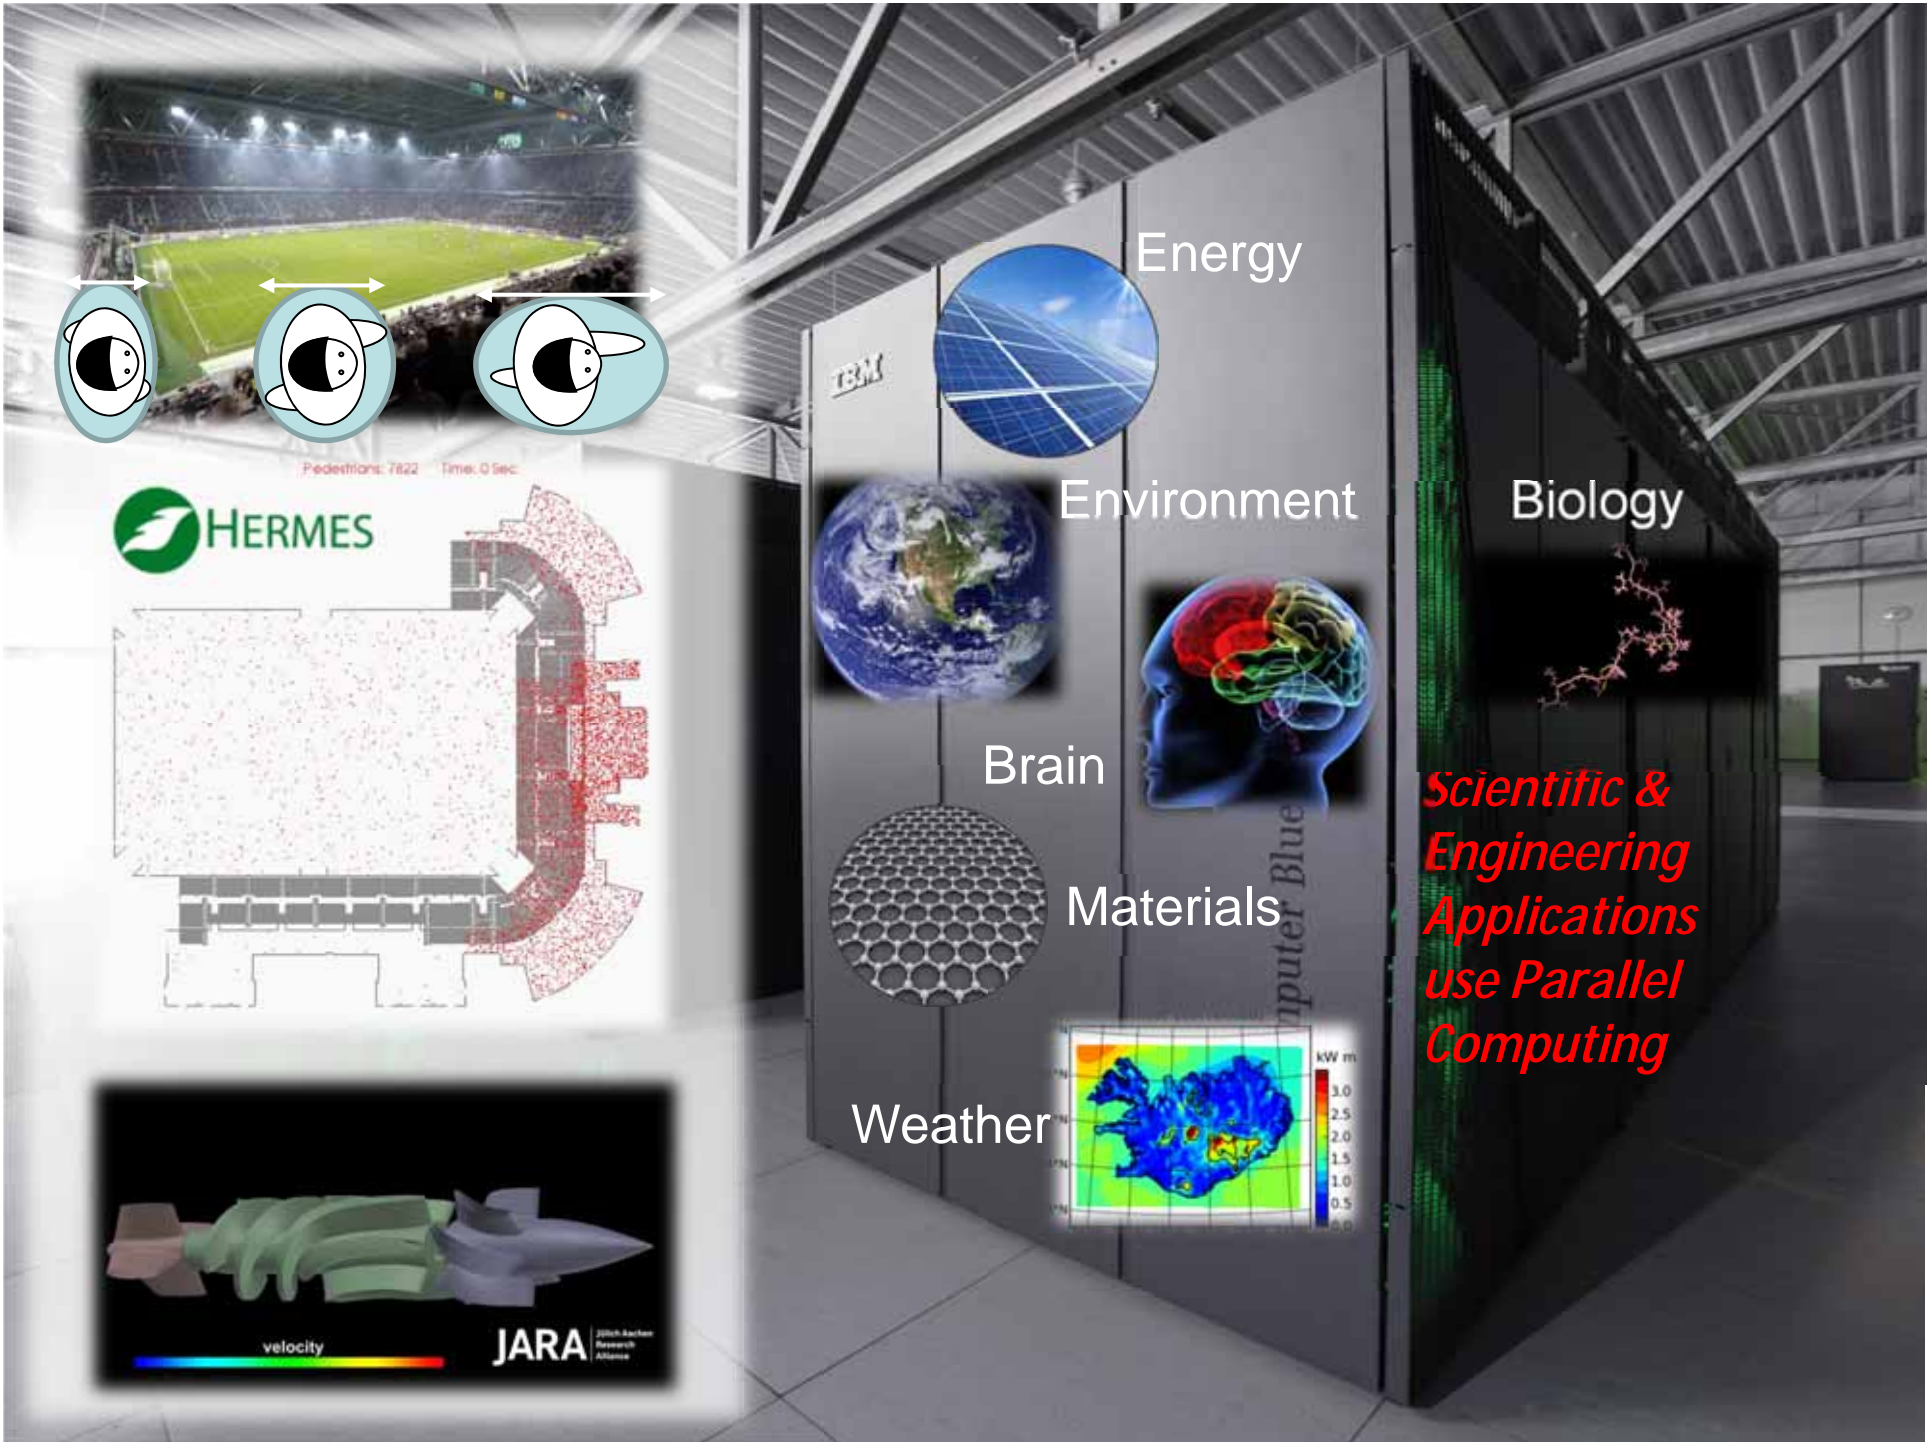
1.000.000 FLOP/s
~1984





Upgrade JUGENE to JUQUEEN

supercomputer Blue Gene





Pedestrians: 7822 Time: 0 Sec




Energy




Environment



Brain



Materials



Weather

Biology

Scientific & Engineering Applications use Parallel Computing

Research on 'Big Data Waves' in Science & Engineering

How can we manage the rising tide of ever-increasing data?

Unsolved questions:

- Scale
- Heterogeneity
- Stewardship
- Curation
- Long-Term Access and Storage

Challenges:

- Collection, Trust, Usability
- Interoperability, Diversity
- Security
- Education and training
- Data publication and access
- Commercial exploitation
- New social paradigms
- Preservation and sustainability




How do we efficiently analyse these large quantities of data?

Making use of Big Data


Applying 'smart data analytics' techniques

Classification




- ✓ *Groups of data exist*
- ✓ *New data classified to existing groups*

Clustering




- ✓ *No groups of data exist*
- ✓ *Create groups from data close to each other*

Regression



- ✓ *Identify a line with a certain slope describing the data*

➤ *Serial algorithms for large volumes of data exist since decades, 'big data' needs their parallelization*



Can I do this in Iceland? – Yes!

Where do I get fast computing?

66°N

14°W

© Illustration by map-europe.net

Nordic HPC

Joint Nordic Supercomputer in Iceland



Contact:

nhpc@hi.is



HPC
System
Garðar



Where do I get the skills? Get involved! Contact Us!

Selected Teaching Activities @ HI → morris@hi.is



Statistical Data Mining Course & Seminars

Big Data Analytics Techniques, Data Analysis, Parallel & Scalable Machine Learning

High Performance Computing A/B Course & Seminars

Parallel Programming, Simulation Sciences, MPI/OpenMP, Map-Reduce/Hadoop/Spark