INFOCOMP 2012

Future High End Systems: Chances and Challenges for Intelligent Applications and Infrastructures

October 22, 2012 - Venice, Italy

International Conference on Advanced Communications and Computation (INFOCOMP 2012)

INFOCOMP/DataSys
October 21–26, 2012 - Venice, Italy

INFOCOMP 2012 International Panel: Future High End Systems

Panelists

- *Claus-Peter Rückemann* (Moderator), Leibniz Universität Hannover / Westfälische Wilhelms-Universität Münster (WWU) / North-German Supercomputing Alliance (HLRN), Germany
- *Igor Melatti*, Sapienza University of Rome, Italy
- *Brian A. Worley*, Oak Ridge National Laboratory (ORNL), USA
- *Lutz Schubert*, High Performance Computing Centre Stuttgart (HLRS), Germany
- *Wolfgang Hommel*, Leibniz Supercomputing Centre (LRZ), Munich, Germany
- *Udo Inden*, Cologne University of Applied Sciences, Cologne, Germany

INFOCOMP 2012: `http://www.iaria.org/conferences2012/INFOCOMP12.html`

INFOCOMP 2012 International Panel: Future High End Systems

Chances and Challenges for Intelligent Applications and Infrastructures

- **Scale-rush:** Which limitations do we encounter?
- **System and hardware architectures:** Roadmap to use?
- **Infrastructures, frameworks, applications:** What do we need?
- **Intelligent software:** What are the suggested benefits?
- **Knowledge and information:** Status in complex systems?
- **Automation and programming:** Next steps ahead?
- **Parallelisation and optimization:** Key issues?
- **Technology/software:** Is there an integration breaktrough?
- **Your ideas:** Who and what are we operating high end systems for and why does general progress take so long?

INFOCOMP 2012 International Panel: Post-Discussion Summary

Post-Discussion Summary:

- *Scientific and discipline view*: For predictable time, the top High Performance Computing segment will stay "individual" (configuration, customisation, ...).
- *Complementary business view*: Markets start to enter the high end segment.
- There is *neither a general hardware nor a general software solution* for HPC (physical, hardware, and software limitations, high end bottlenecks).
- There is *no general* solution for HPC from what-so-ever intelligent applications.
- *Automation*, *verification*, *integrated systems*, *intelligent systems*, and *knowledge discovery* become important when dealing with increasing complexity.
- Energy consumption / economical deployment force to optimise wherever possible.
- HPC scalability, usability, and ergonomical use are *not* out-of-the-box.
- *Intelligence* and *simplicity* are essential for *balanced* high end solutions.
- Further performance increase will need technological innovations *in addition* to implementing the available software and hardware approaches.
- Essential documentation, content, and knowledge resources have to be made *long-term* persistent (international collaboration, education, classification, re-use).
- High End Computing will need strong and sustainable long-term support for and from *education and teaching* in order to create a wide range of communities and platforms for common usage scenarios and to accelerate the exploitation.

INFOCOMP 2012 International Panel: Table of Presentations

Panelist Presentations:

- **Wrapup Experiences for Future Advanced Scientific Supercomputing** (*Rückemann*)
- **Formal Verification and Automatic Synthesis for Future High-End Systems** (*Melatti*)
- **Knowledge Discovery from Data** (*Worley*)
- **The End of High Performance Computing?** (*Schubert*)
- **Challenges from the HPC service provider perspective** (*Hommel*)
- **Markets will determine the future of HPC and HEC-Systems** (*Inden*)

International Panel INFOCOMP 2012

Future High End Systems:

Chances and Challenges for Intelligent Applications and Infrastructures

Wrapup Experiences for Future Advanced Scientific Supercomputing

The International Conference on Advanced Communications and Computation (INFOCOMP 2012)
October 21–26, 2012, Venice, Italy

Dr. rer. nat. Claus-Peter Rückemann[1,2,3]

[1] Leibniz Universität Hannover, Hannover, Germany
[2] Westfälische Wilhelms-Universität Münster (WWU), Münster, Germany
[3] North-German Supercomputing Alliance (HLRN), Germany

ruckema(at)uni-muenster.de

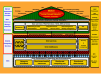

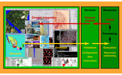

Status on High End Systems

Commonly understood physics and terminology

- In High Performance Computing, supercomputers -i.e. computer systems at the upper performance limit of currently feasible processing capacity- are employed to solve challenging scientific problems.
  A "distributed High Performance Computing" cannot exist.
- Grid and Cloud Computing are distributed computing.
- "High End Computing" summarises Grid, Cloud, and HPC.

Supercomputing / High Performance Computing / Cloud and Distributed Computing

Facing system related challenges regarding disciplines, services, resources, and development:

- Complexity (application-system scenarios, resources, services)?
- Data Size (Big Data – Volume, Velocity, Variability)?
- Compute (computation, memory, and benchmarks)?
- Data locality?
- Nature of numerics and algorithms?
- Exascale – Out-of-Scale?

Status on High End Systems

Advanced Scientific Computing

- Thinking in efficient "Advanced Scientific Computing algorithms".
- Architectures (Shared memory, Graphics Processing Units / GPU)?
- New language concepts beyond MPI (PGAS and others)?
- What about parallel runtimes, compilation, and optimization?
- Portability of optimisation?
- Higher level of automation and programming?
- Complex systems operation and costs?
- Licensing (many many cores)?
- Individual solutions.
- Last 30 years experiences: Any significantly complex High Performance Computing / Supercomputing development does take at least 5 years before productive in large scale.
- Subjective trend: 90 percent of completeness can be considered weak.

Status on High End Systems

Business and industry computing

- Thinking in "business model and concept" perspective.
- Algorithms are developed on contract base.
- Legal regulations and security?
- Standardised industry support?
- Solutions and developments are expected to take place just in time (1-2 years).
- It is easier to rename a task "solved" than creating a solution.
- Subjective trend: 10 percent of (scientific) completeness can be considered economic.

Status on High End Systems

No requirements, no High Performance Computing!

HPC resources providers are dealing with specifiable requirements, not with the challenges of application algorithms.

Perception and solutions?

- Complementary needs and perception on research and business sides?
- Integrated collaboration frameworks and concepts?
- Validation, verification, error correction?
- Redundancy and criticality management?
- Knowledge?
- Classification (Universal Decimal Classification, UDC)?
- Content and context?
- Georeferencing?
- Intelligent software / Multi-Agent Systems (MAS)?
- Is software a solution for every problem?
- Isolated user groups, no holistic view on context and content.
- Who will be the recipients of YOUR work/results?

Status on High End Systems

Distributed Computing

Disciplines:

- Interactive communication requirements (quantity and quality).
- Data transfer to/from distributed resources (interactive and batch).

Services and Resources:

- Communication: Insufficient data transfer for large data volumes.
- Storage: Insufficient backup capacities for large data volumes.

Development of Applications and Methods:

- ... depending on funding, physical resources, consulting.
- ... depending on reliability, high availability, security.

Status on High End Systems

Integrated Information and Computing Systems

- Dynamical tasks,
- Interactive tasks,
- Batch job tasks,
- Requirements range

from	$>$ 100000 requests per interactive application per operation
to	$>$ 100 GB per second, per application; $>$ 1 GB per second per core,
from	single TB datasets
to	millions of small dataset per node,
from	loosely
to	highly parallel.

- We need all these technologies and concepts but without appropriate funding today what is the essence?
- Integrating disciplines
  - Seismics and seismology (SEG-Y), geosciences, natural sciences, ...
  - Geophysics, archaeology, global environment, climatology (NetCDF), ...
  - Spatial information, georeferencing, and search, visualization, streaming data, ...
  - Data and information structure, documentation, classification (UDC), ...

  with Services, and Resources (transfer, storage, compute).
- **Separate the knowledge and describe the context for a long-term effect!**
- **We need to make the overall essence imperishable!**

Vision

Goals for future high end systems

Next generation: The next generation of complex high end systems will not be just only the next sizeof-factor multiplied with all the items required or just another project. The next generation step can neither be reached by hardware nor by software alone.

Holistic view: For reaching a general state of the art there is no efficient solution available other than an holistic handling, integrating user applications, services, and resources.

Common understanding: An overall understanding of integrated use, operation, and provisioning regarding software and hardware is neccessary for appropriate funding.

Real life case studies: Real life large data case studies with keyplayer participants from academia and industry representing disciplines, services, and resources providers.

Challenges

Funding (and) Complexity

Challenges and deficits:

(as identified in last years' *INFOCOMP Partnership Session and Open Call* and *Open Discussion* in Barcelona and this years' *DigitalWorld International Panel on Big Data* in Valencia)

- ((Funding)),
- There is no wearout operation with High End Computing,
- Common availability of (broadband) data transfer capacities,
- Fast and massive I/O and communication solutions,
- Archiving,
- Storage, file systems,
- Reliable and secure data and resources access (homomorphic not practical),
- Data staging with HW and SW,
- Architectures (storage, memory, cores) and energy efficiency.

Issues of future high end systems:

- **Complexity** (holistic view, frameworks, legal aspects).
- **Technological limitations and distinctive features** (physics, scale).
- **Multi-disciplinary work** (documentation, knowledge).
- **Hardware and software integration**.
- **Portability** of algorithms and solutions (automation, optimisation).

Formal Verification and Automatic Synthesis for Future High-End Systems

Igor Melatti

Sapienza Università di Roma

Model Checking Group http://mclab.di.uniroma1.it/
Computer Science Department
Sapienza University of Rome

Oct 22nd, 2012
INFOCOMP 2012
Venice, Italy

Outline

1. Status

2. Vision

Testing or Verification?

- Many complex systems have been developed in the last decades
  - software, hardware, or both
  - hundreds of modules, millions lines of code, ...

Testing or Verification?

- Many complex systems have been developed in the last decades
  - software, hardware, or both
  - hundreds of modules, millions lines of code, ...
- How to *check* correctness for all this?
  - does the system behave as it was specified?

Testing or Verification?

- Many complex systems have been developed in the last decades
  - software, hardware, or both
  - hundreds of modules, millions lines of code, ...
- How to *check* correctness for all this?
  - does the system behave as it was specified?
- Most used approach, by far, is *testing*
  - feed the system (or a prototype) with some input, and check the output

Testing or Verification?

- Testing has many known advantages
- It is easy to set up
  - developers always at least the final system, often several prototypes and *simulators*
- If the system is not critical for some reason, it is ok

Testing or Verification?

- Testing has many known advantages
- It is easy to set up
  - developers always at least the final system, often several prototypes and *simulators*
- If the system is not critical for some reason, it is ok
- For mission-critical (e.g., software running on satellites) or safety-critical systems (e.g., software used for diagnosys) *formal verification* must be used
  - testing may formally prove that errors are present
  - formal verification may also formally prove that errors are *absent*

Testing or Verification?

- Formal verification disadvantages: difficult to use
  - often needs some mathematical background
  - a model of the system must be built
- Requires many computational resources
  - months of computation
  - TB of RAM
- Formal verification big advantage: provides *mathematical certification of correctness*

SAPIENZA
UNIVERSITÀ DI ROMA

Outline

Formal Verification on High End Systems

- Formal verification problems are however mitigated by new approaches
  - e.g., our research group has been able to couple verification and testing (by reusing simulators) and verify operational procedures for actual satellites
  - many examples exist of "real-world" systems verified to date
- Future high-end systems may help as well
- Parallel approaches for formal verification are widely studied
  - e.g., an entire annual conference, PDMC (Parallel and Distributed Model Checking), is dedicated to this issue

Synthesis

- Software for some (sub)systems may be *automatically synthesized* from formal specifications
- Such software will be correct by construction
- Also in this field, big jumps ahead have been done in the last decade
- Since synthesis is often based on formal verification technique, high end systems may help also here

Knowledge Discovery from Data

Brian Worley
Director, Computational Sciences and Engineering Division

Oak Ridge National Laboratory, USA

Oak Ridge National Laboratory

The Data Explosion

Experiments

Simulations

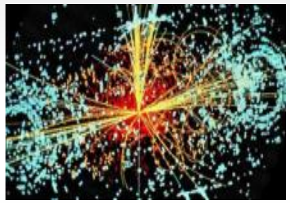

Archives

Social Media

Sensors

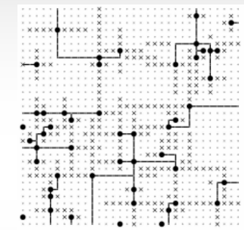

Information Technology

Petabytes Doubling & Doubling

The Challenge
Enable Discovery

Deliver the capability to mine, search and analyze this data in near real time

The Response

Science itself is evolving

The Changing Nature Of Research

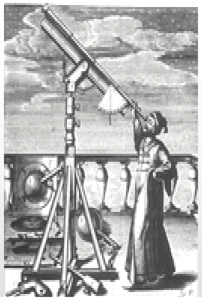

$$\left(\frac{\dot{a}}{a}\right)^2 = \frac{4\pi G\rho}{3} - K\frac{c^2}{a^2}$$

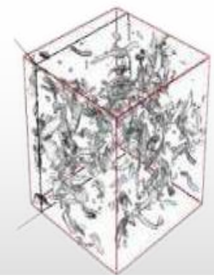

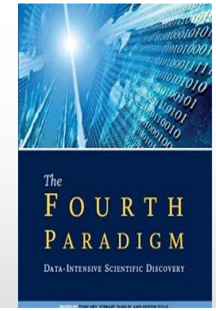

Experiment

Thousand years ago

Description of natural phenomena

Theory

Last few hundred years

Newton's laws, Maxwell's equations…

Computation

Last few decades

Simulation of complex phenomena

Data

Today and the Future

Unify theory, experiment, and simulation with large multidisciplinary data

Using data exploration and data mining (from instruments, sensors, humans…)

Distributed Communities

Knowledge Discovery Science

- Actionable insights from massive, dynamic, disparate data
- Ability to detect, understand, and predict processes underlying the data

Systems

- Sharing and trust
- Social media
- Streaming
- Architecture
- Sensors
- Mobile
- Workflow

Data Analytics

- Text Analysis
- Multi-modality fusion
- Clock-constrained
- Large-scale
- Geo-temporal
- Social networks

Modeling & Simulation

- Large-scale infrastructures
- Physics-based
- Discrete-event & Agent-based
- Real-time/predictive

Cyber Security

- KD-based
- Imbedded systems
- Use of HPC
- M&S-based
- Quantum systems

Mobile platforms (Cell phones, iPad, PDA, UAV)
Desktops, Clusters, and Cloud Computing (Homogeneous and Heterogeneous)
Distributed (Sensor networks, computational platform mixes)
Extreme-scale computing (HPC)

OAK RIDGE National Laboratory

Information Platforms Based on Social Media Features

- Reference by URL
- Contributed content
- Reputation system
- Tagging + search
- User-defined mashups
- Publish-subscribe

VERDE Electric Grid Status
(Real-Time Grid Awareness)

Tracking 2.0
(Cradle-to-grave tracking)

Sensorpedia
(The “Wikipedia of Sensors”)

Knowledge Discovery Framework
(National Biomass Distributions)

OAK RIDGE
National Laboratory

Architecture Concept to Accommodate Disparate, Dynamic, Unstructured Data

Sensor Networks

Distributed Global Databases

Social Media

Sources

Tagging

Faceted Classification Systems

Classify by TYPE
Type A
Type B
Type C

Classify by ZONE
Zone A
Zone B
Zone C

Classify by TIME
Time A
Time B
Time C

Facet

Value

Architectures

Analytics

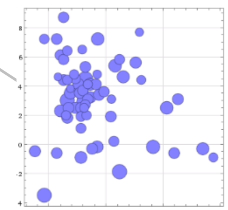

Visualization

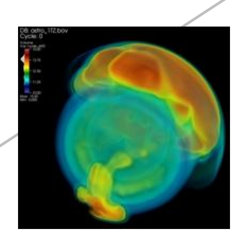

Tools

End User

Thank You

Brian A. Worley

worleyba@ornl.gov

The End of High Performance Computing?

Lutz Schubert {schubert@hlrs.de}

Development Trends

- Power Consumption
  - For every Processor Watt, roughly 1.2 Watt for cooling and provisioning
  - Current Petascale systems: ~1,25 MW
  - Low power CPUs would lead to >490 MW
  - ➢ Projected consumption for Exascale (with current technology): ~1.25 GW without cooling
  - ➢ In comparison, a medium scale power reactor creates around 300 MW
- Stability / Resilience
  - MTBF for one CPU ~10 years
  - Exascale system ~200.000 CPUs
  - ➢ MTBF for a full-scale application on an exascale system: ~26 minutes
- Performance
  - Clock speed is not really increasing anymore
  - ➢ Scalability of SOFTWARE becomes the crucial factor

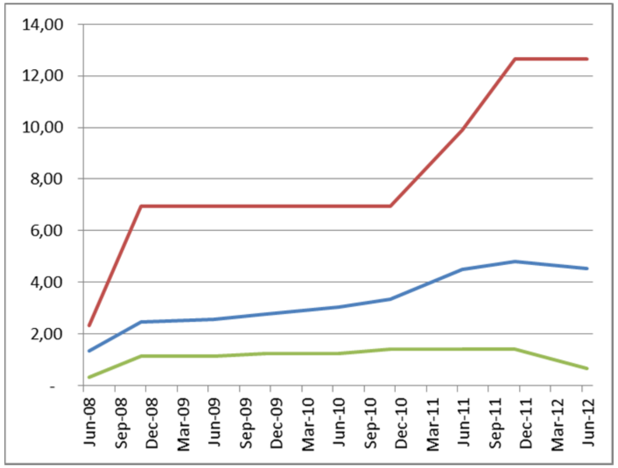

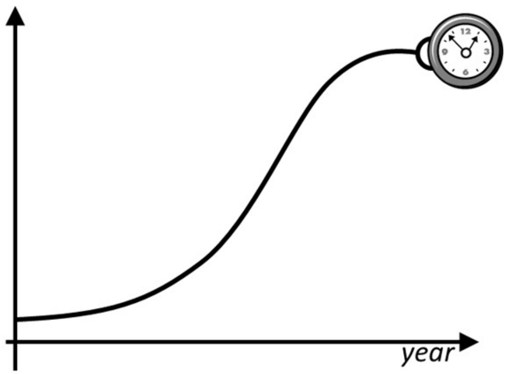

Cost vs. Benefit

- Physical limitations cannot be overcome
  - Memory wall
  - Communication bandwidth & latency
  - Speed of light
  - Power wall
  - Etc.
- Effort to increase performance of soft- and hardware gets higher and higher
- Only applications really benefitting from exascale are those with little to no dependencies (Embarrassingly parallel; Multi-level applications)
  - ➢ Require big memory
  - ➢ Do not require fast interconnects
  - ➢ Not challenging; do not require expensive systems
- How do we serve the essential problem spaces though?

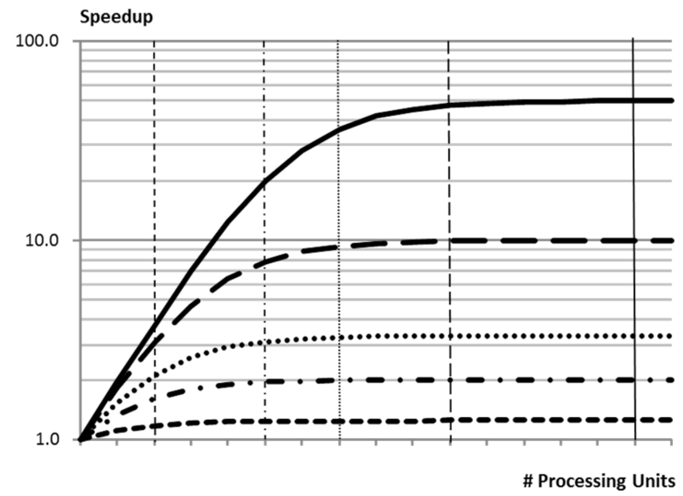

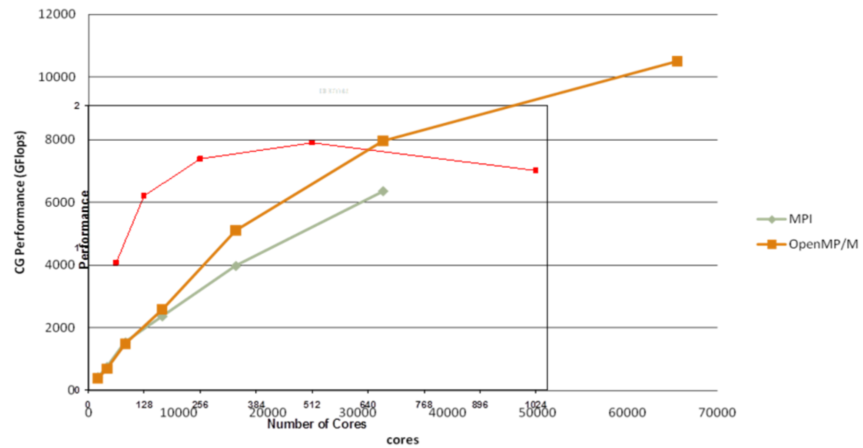

Predicting the Future

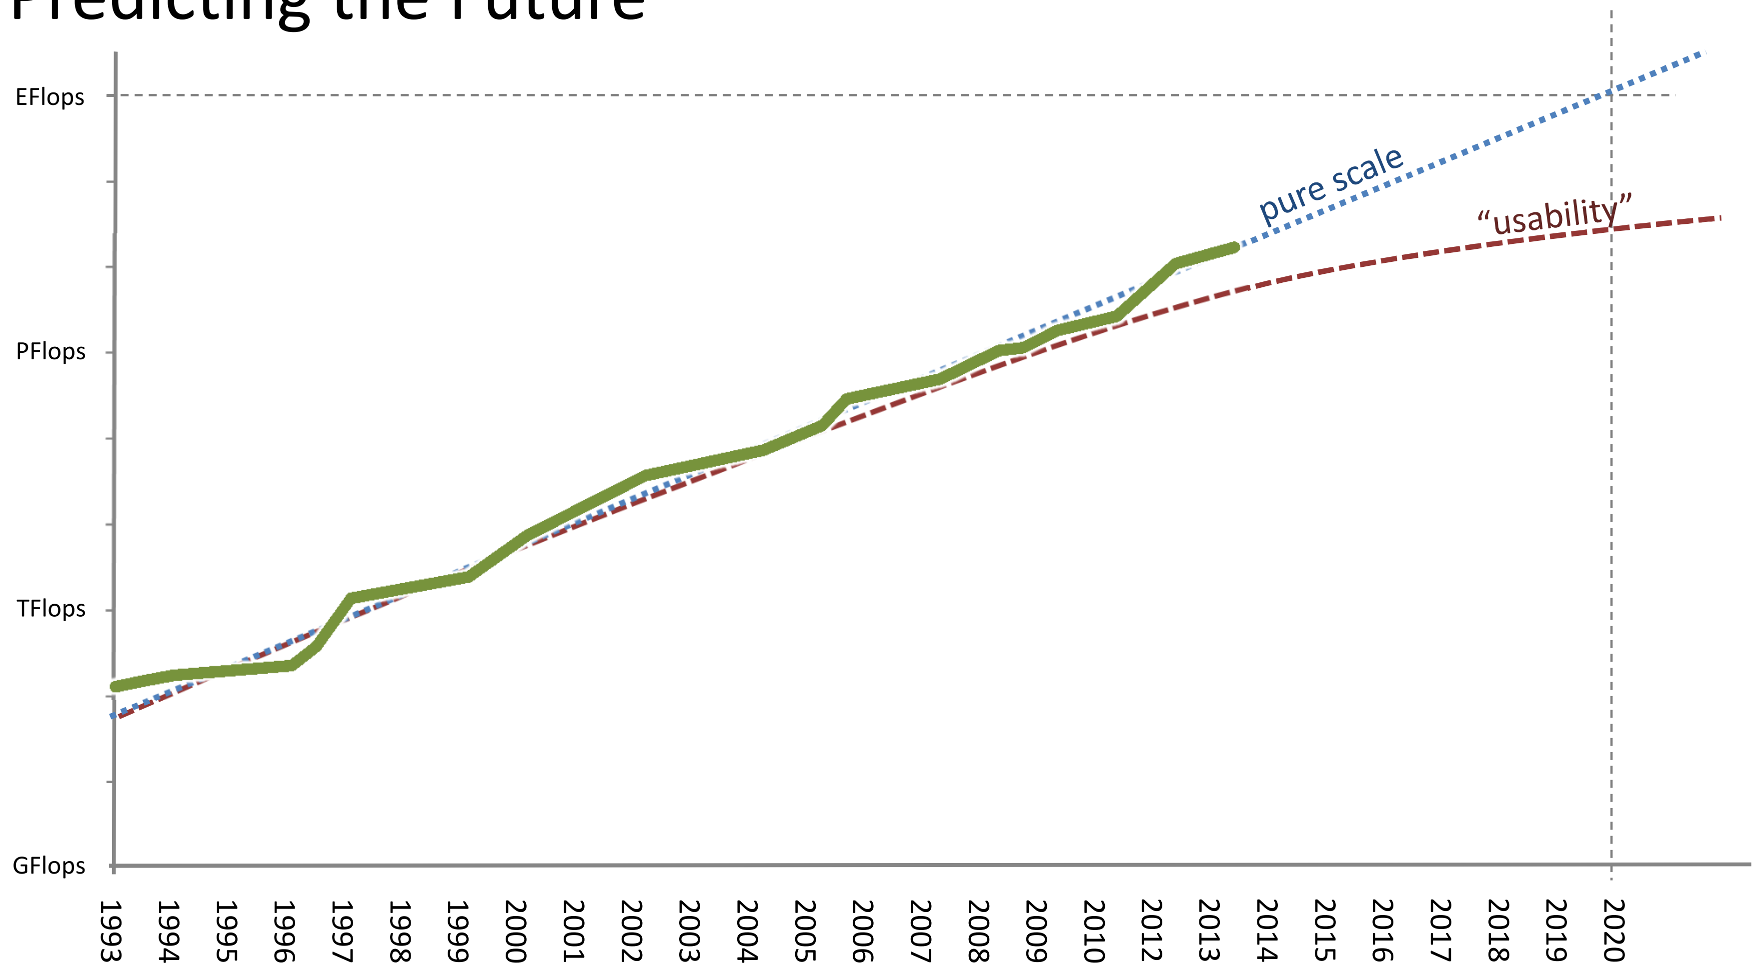

What about future applications?

- Cannot talk about "general" performance increment anymore
- Distinguish roughly:
  - Embarrassingly parallel applications:
    - HPC communication more expensive than necessary
    - ➢ Move to the cloud?
  - Loosely coupled applications
    - Increase the problem scope with minimal impact on communication
    - Can still exploit growth in systems, but the benefit is decreasing
    - ➢ Will this reach exascale?
  - Tightly coupled applications (hard case)
    - The problem scope impacts directly on communication scope
    - Generally already decrease in performance with more resources
    - ➢ Current systems and approaches are insufficient

➢ For the real hard and interesting cases we need a complete new computing approaches!
  - ➢ Multi-level programming
  - ➢ Exploitation of concurrency
  - ➢ Move away from Turing model
  - ➢ Alternatives to Silicone?

Thank you

Lutz Schubert {schubert@hlrs.de}

MNM TEAM

MUNICH NETWORK MANAGEMENT TEAM

Challenges from the HPC service provider perspective

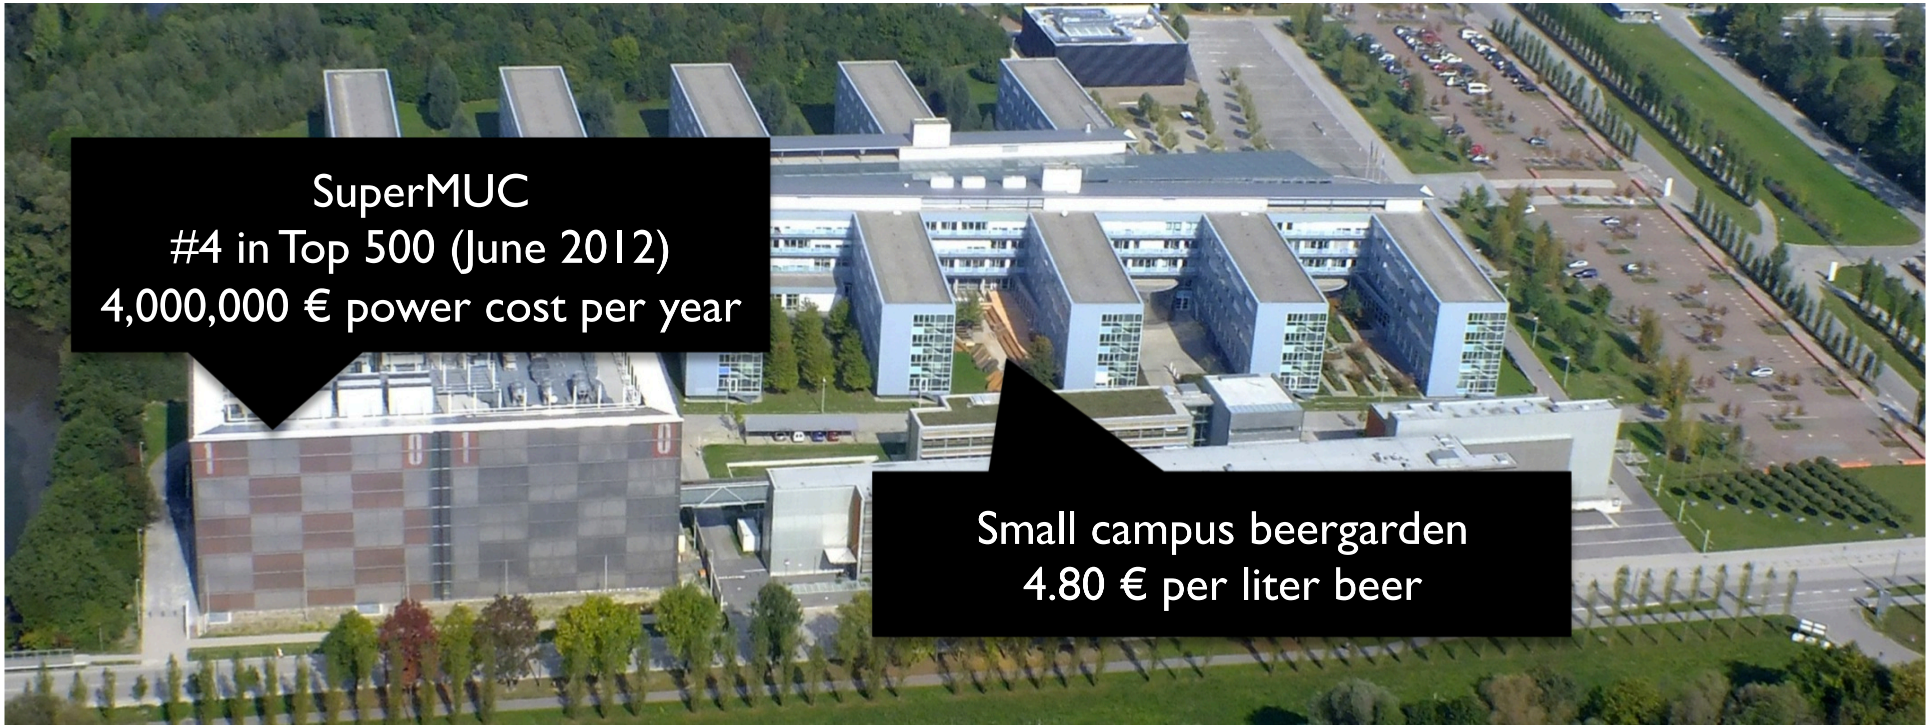

Photo: Ernst A. Graf, 2012

Wolfgang Hommel

Leibniz Supercomputing Centre, Munich, Germany

INFOCOMP 2012 Panel on Future High End Systems: Chances and Challenges

Current status

- Top 10 supercomputers 2012:
  - $O(10^5)$-$O(10^6)$ cores
  - Avg. power consumption << 5 MW
- Legacy code vs. intelligent apps:
  - Hardware has a half-life of years
  - Code has a half-life of decades
  - Many applications run at 5-10% performance only
- Infrastructure: evolution, innovation, intelligence?
  - Bottlenecks as expected: interconnects, file system i/o, WAN data transfers
  - Lack of infrastructure management automation

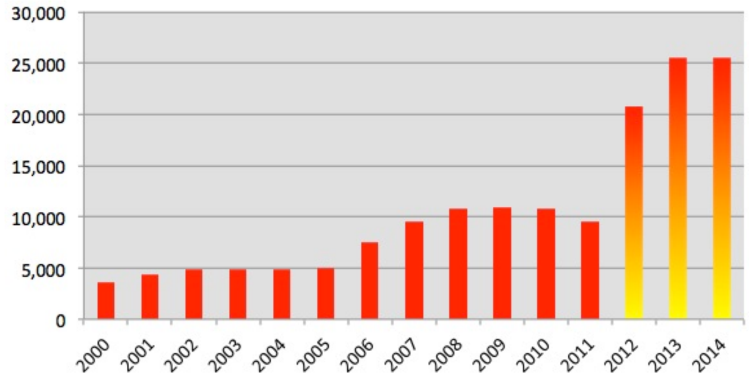

LRZ: Power consumption in MWh

SuperMUC: 147,456 cores / 3.42 MW

Bad today – worse tomorrow?

- High infrastructure costs
  - Raising electricity tariffs increase power costs
  - New systems require power/cooling upgrades
    - Increased air flows for direct air cooling
    - Cold and hot water loops
    - Oil-immersion cooling
- User-friendly application development?
  - User-level checkpoints required
  - Lack of high-performance open source libs
  - Lack of perf. analysis/fault prediction tools
- Long crash recovery duration; example:
  - *11 minutes* blackout at LRZ in Sept. 2012
  - *21 hours* until all nodes were rebooted
  - *5 days* until file system repairs finished

LRZ: Air cooling infrastructure (2006)

LRZ: Hot water cooling infrastructure (2012)

Where are we headed in this decade?

- Software challenges
  - O(10^9), heterogeneous cores
  - In-interconnect processing
  - Power and fault control
- Hardware challenges
  - Reduce bottlenecks and power consumption of RAM, interconnects and file i/o
- Intelligence must be in the right places
  - Intelligent, auto-tuning open source libraries
  - Keep applications simple, avoid extra-smartness
  - Embrace new programming paradigms
    - Communication avoidance
    - Data-driven flows, out-of-order execution
  - Intensive collaboration with vendors and user trainings
  - Automate infrastructure management, e.g., fault recovery

It looks like you have a hot water cooling infrastructure.

Would you like this intelligent application to

- ☐ brew coffee
- ☐ heat your building
- ☑ just run

Panel Session

Future High End Systems: Chances and Challenges for Intelligent Applications and Infrastructures

Markets will determine the future of HPC and HEC-Systems

Udo Inden, Cologne University of Applied Sciences
Research Centre Knowledge Management / Centre for Applications of Intelligent Systems

What will future HEC-Systems look like?

What will reach exascale first:

Large super-computing centres or large networks of mobile devices?

*" ... That is to say that by **2020 we can expect mobile devices** (laptops, tablets and smart phones) **with a performance of around 1 TFLOPs.***

*A **standard laptop with a standard GPU** already has a performance of **50 GFLOPs**. Furthermore such a system **can even now be fitted with a performant GPU** card to raise its performance to **1 TFLOPs**. The GPU in a **typical tablet or smart phone** currently performs at **~ 5 GFLOPs**.*

All this clearly supports the conclusion that by 2020 we will have hand-held mobile devices capable of performing several 100 GFLOPs or more. The … question is to what use this performance will be put."

That is only 10^5 – 10^6 less than super-computers

What will 1.000.000 mobile devices coupled via internet mean?

→ INCOCOMP 2011: "CLIPS" – Collaborative Intelligent Applications

Francis Wray, "A Brief Future of Computing", 2012 HPC Planet (HPC Planet is a Support Action, funded by the 7th EU Framework Program, 2009)
http://www.planethpc.eu/index.php?option=com_content&view=article&id=66:a-brief-future-of-computing&catid=1:articles&Itemid=3
Retr. Oct. 2nd 2012

complexity

Major Phases of the Evolution of Operations Complexity

Operations' Models from Henry Ford to the Semantic Web

There will be a massive increase of HPC-applications

- **in technical engineering**
  (product and production technology development)
- **in business operations management**
  - marketing & innovation intelligence
  - supply-chain & production management
  - Integrated risk management

This will have impact on the HPC-Industry!

(which may not know yet that it is an industry)

"Internet Universe" and Big Data

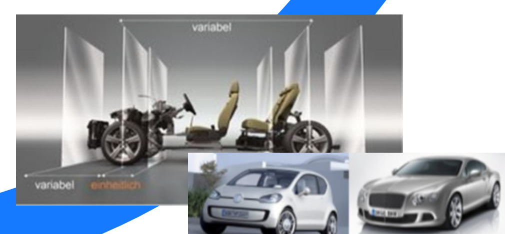

VW Matrix Platform

Tin Lizzy

Taichi Ohno

Udo Inden, Cologne University of Applied Sciences

Panel Session

Future High End Systems: Chances and Challenges for Intelligent Applications and Infrastructures

Massive Parallel Hardware Systems

Thesis:

In 10 years nobody will care about technology but just expect service!

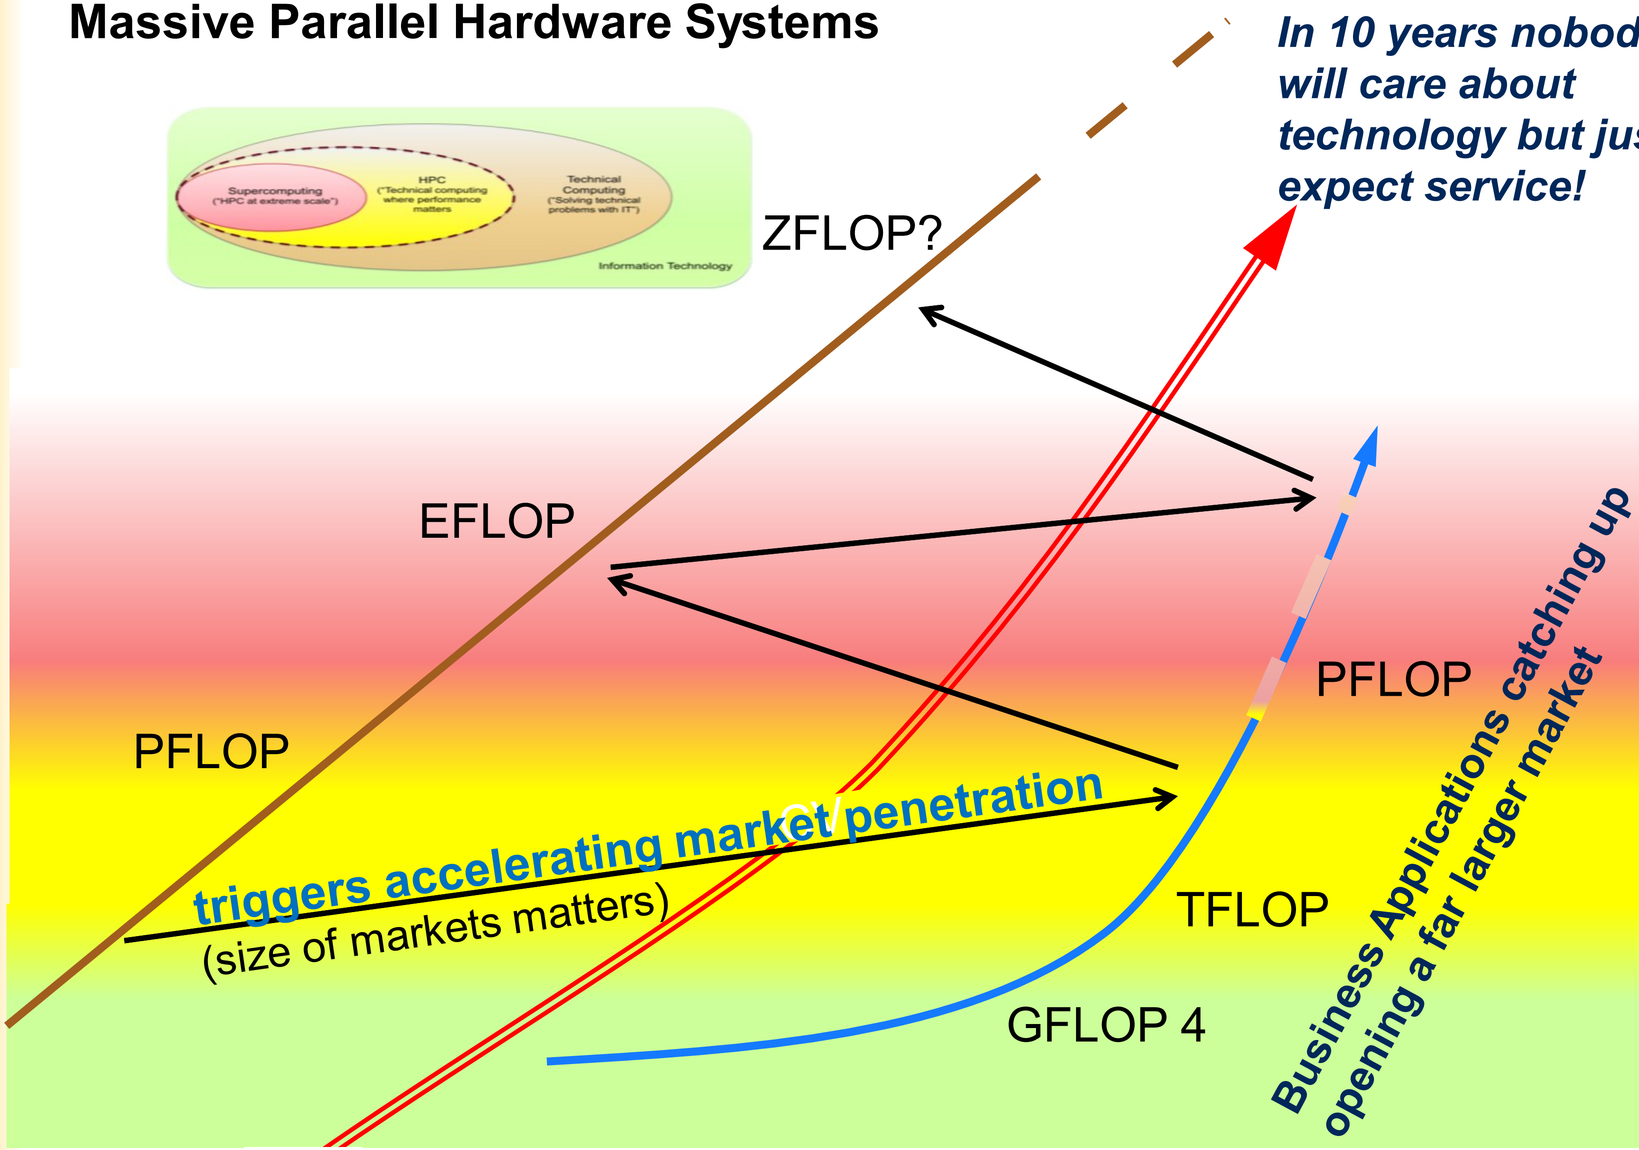

Business-models of providers will change towards more service

Product → **accompanying services** → **services instead product**

Car → car + guarantees … → mobility instead of cars

Commoditization

- **Capabilities offered and demanded** trigger each other (valid for all competitive markets)
- **If there is a market, solutions will develop finally carrying highest level technology into consumer markets …**

  Although **high end systems** are still beyond reach of the majority (*Bugatti Veyron*) more and more HPC technology will penetrate lower market segments (*series cars*)
- **Speed increases** since rate adaptation matters in competition!
- **Variety of paths of penetration** will increase, e.g. include networks of mobile or embedded systems and result in a portfolio of solutions focusing different niches.
- **ICT does not change the logic of markets at all!**

  ICT changes markets by offering a new food to market's.

What will be the Architecture of Future of Computing Services?

High-Performance Computer Centres?

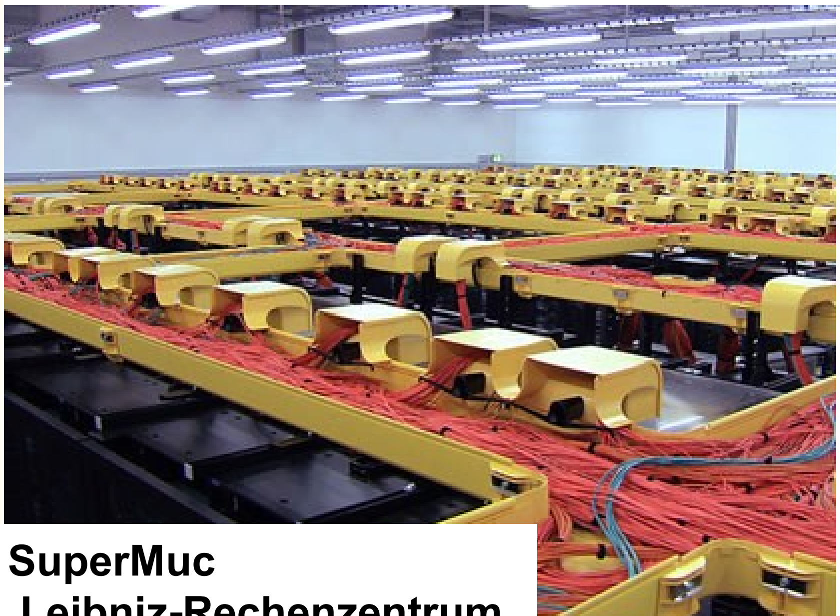

SuperMuc
Leibniz-Rechenzentrum

One customer left?

- **Pentagon, NSA ...**
- **CERN + ...**

Massive Networks of Mobile Devices?

Future Mainstream?
Millions of devices teaming up in solving common problems and self-organising the world or just selling their capacity ...