Trends in Computing Power used by Various BOINC Communities

Yuri Gordienko

G.V.Kurdyumov Institute for Metal Physics, National Academy of Sciences (Kyiv, Ukraine)

10th BOINC Workshop – 30 September 2014 – Budapest

Under partial support and in cooperation with EU FP7 SCI-BUS project, contract no. RI-283481

State of the Art – CPU

Moore' Law:
CPU transistors versus dates of introduction. The line corresponds to exponential growth with transistor count doubling every two years

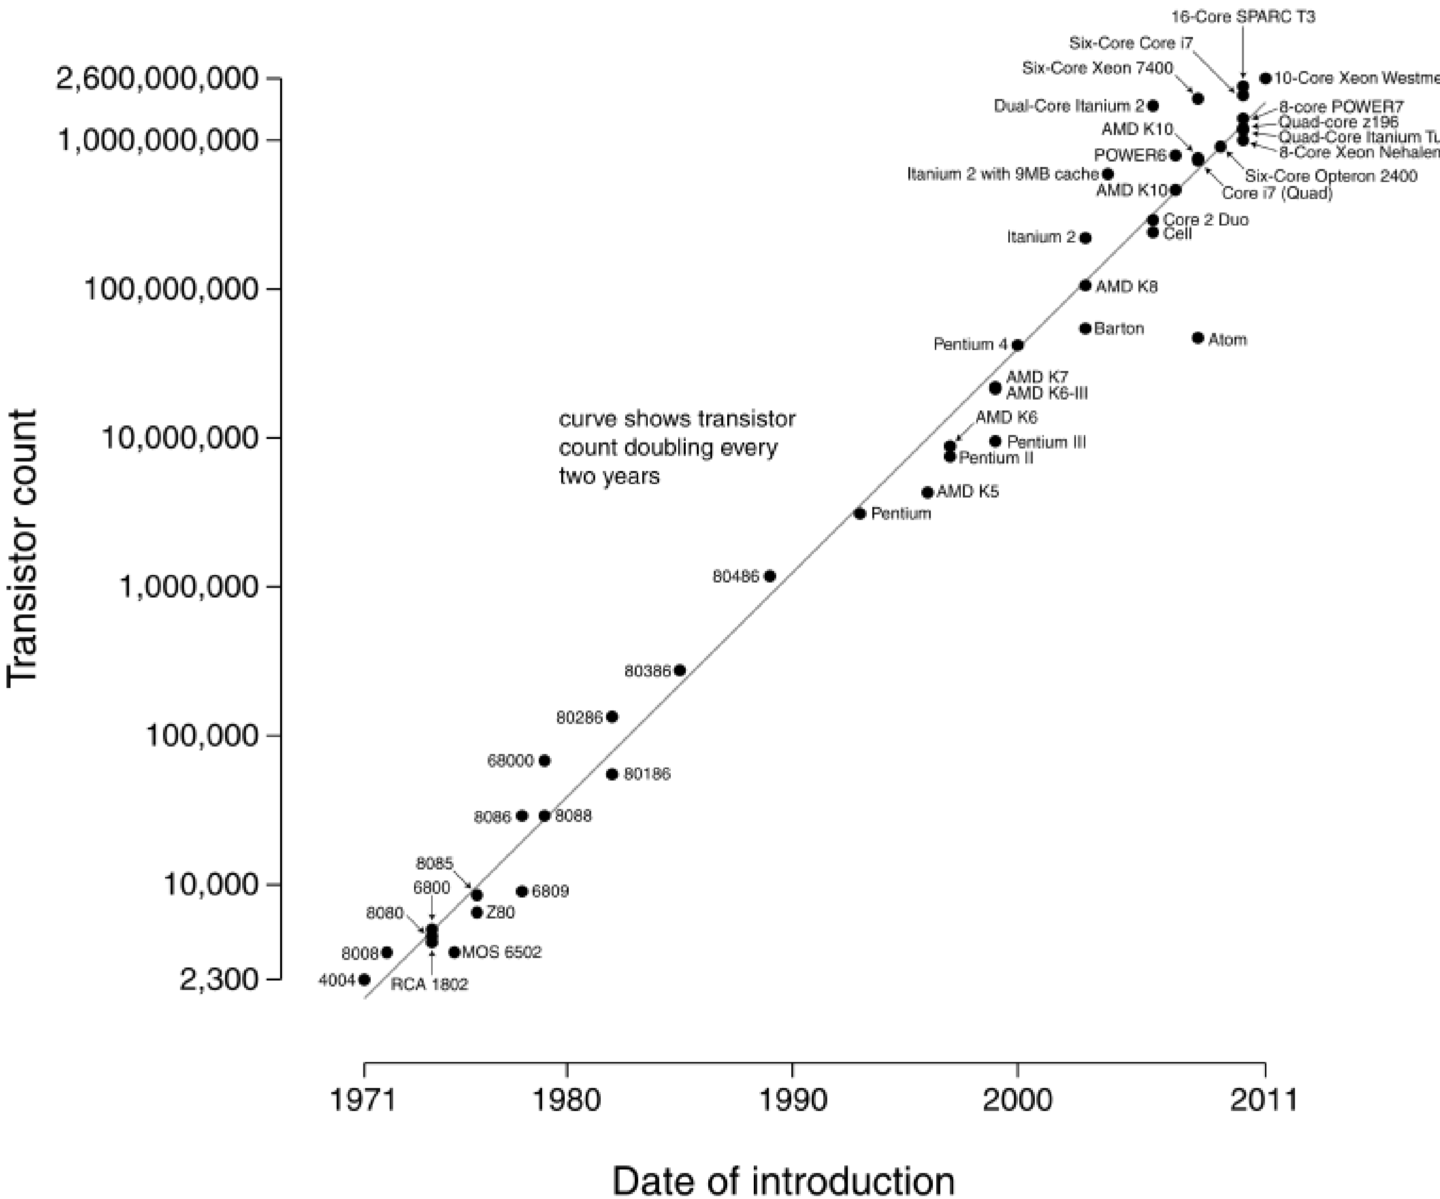

State of the Art – High-tech

Kurzweil's extension of Moore's law: **calculations per second versus time** - from integrated circuits to earlier transistors, vacuum tubes, relays and electromechanical computers

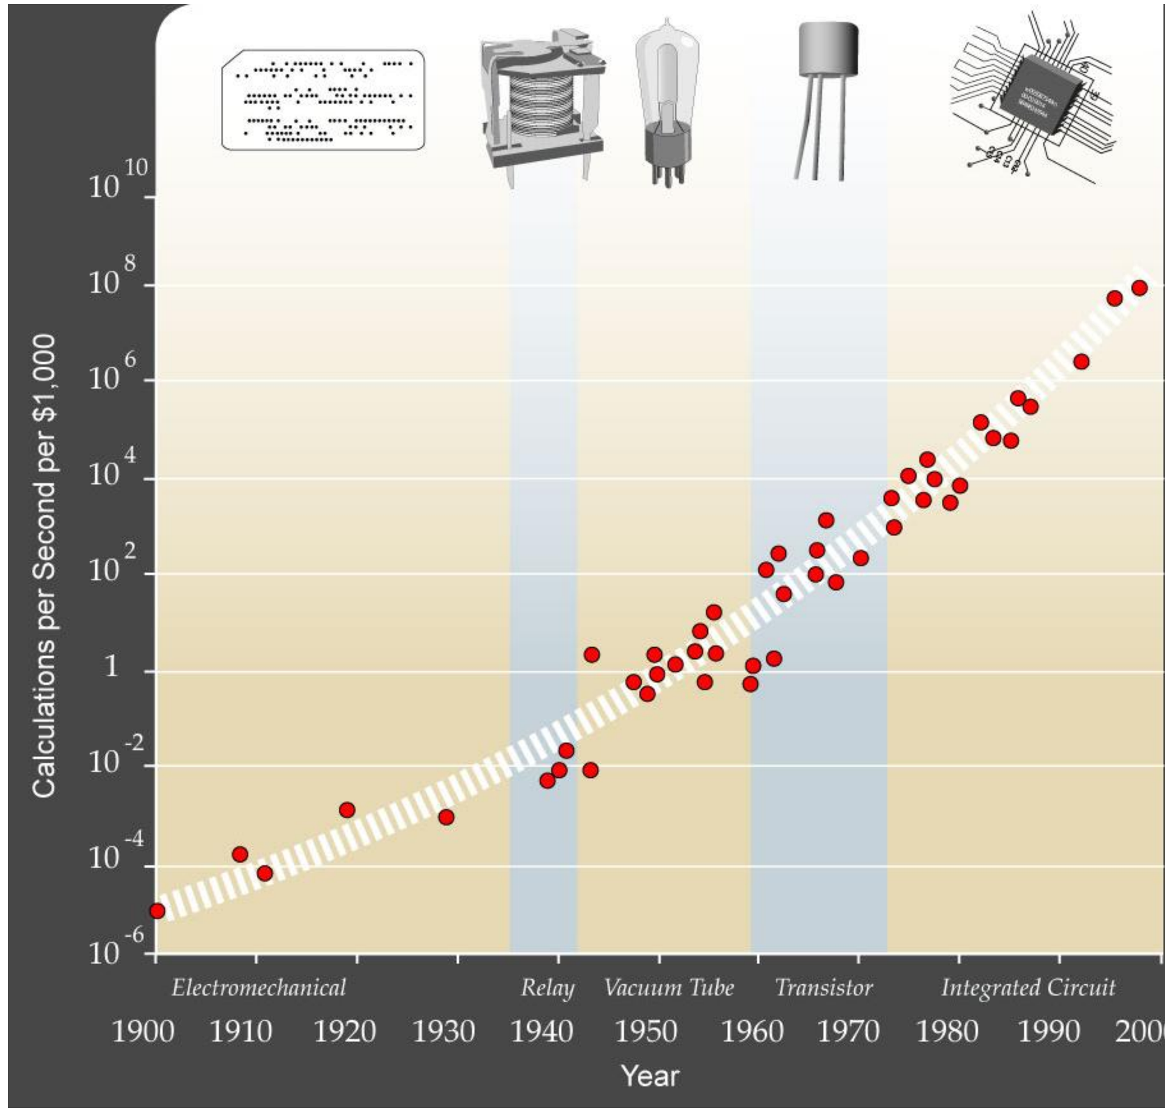

(C) Kurzweil

Let's Measure Progress of Computing Power in BOINC Projects

Motivation:

- Can we measure the actual (not theoretical like in Moore's law) increase of global computing power using …
  … BOINC projects are some subsets of the global PC computing community?
- Can we consider them as "statistically representative samples"?

Aims:

- to explore the progress of computing power in BOINC projects
- to compare it among various BOINC projects
- to check and extend Moore's Law?

Methods:

- use available open statistics on hosts in BOINC-projects
- use closed statistics on actually working machines (not available at the moment)

Data Used

- FLOPs (FPOPs) – the FLoating Point OPerations per second
  "p_fpops"
- IOPs – the Integer OPerations per second
  "p_iops"
- CPUs – the number of CPUs in the host
  "ncpus“

  …

  …
- Others like:
  “p_membw”
- or
  “d_total”
- and more …

Example 1:
SLinCA@Home IMP Desktop Grid

size: small (even **nano**)
history: short (2009-2014-...)
hosts: CPU

Scaling Laws in Cluster Aggregation - SLinCA@Home

Our team in G.V.Kurdyumov Institute for Metal Physics (IMP) , National Academy of Sciences (Kyiv, Ukraine) maintains DG BOINC infrastructure **on the basis of BOINC SZTAKI Desktop Grid technology** at the premises of SLinCA@Home IMP Desktop Grid (http://dg.imp.kiev.ua/slinca).

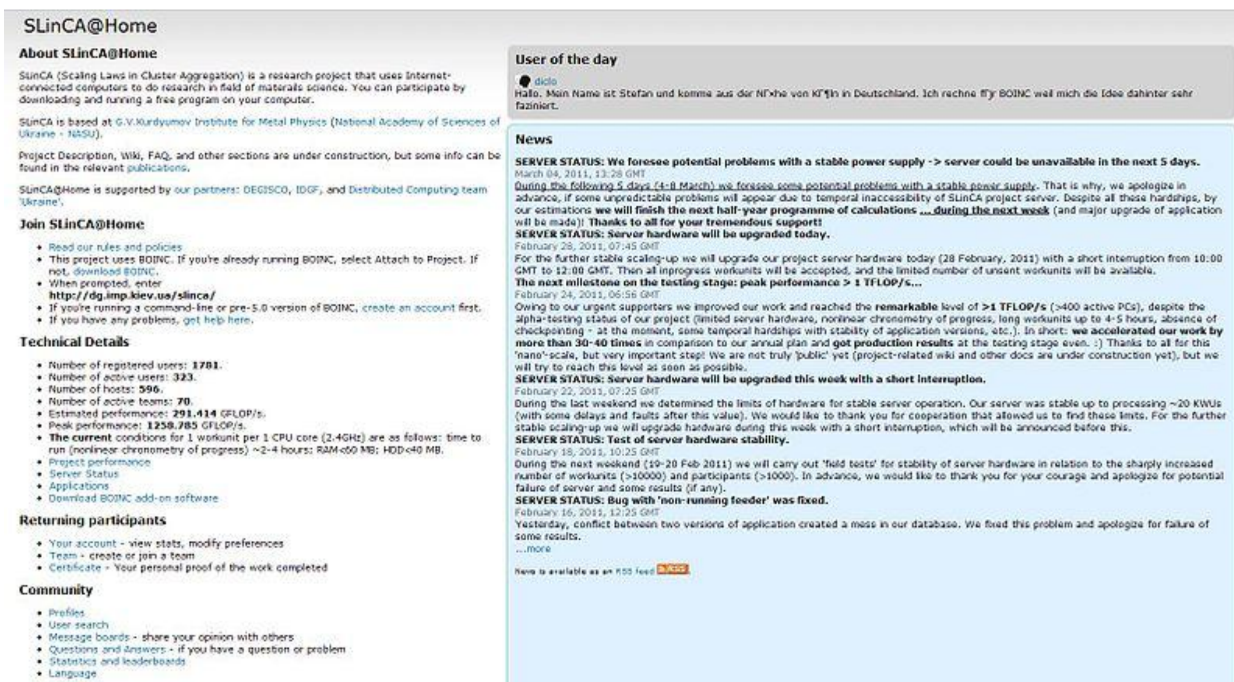

SLinCA@Home – Overview

The current status of IMP Desktop Grid infrastructure:
- ~4000 workers;
- ~ 20 000 in-progress workunits;
- ~150 GFLOPs (average performance)
- ~300-550 GFLOPs (weekly peak performance) and 1.3 TFLOPs (max).

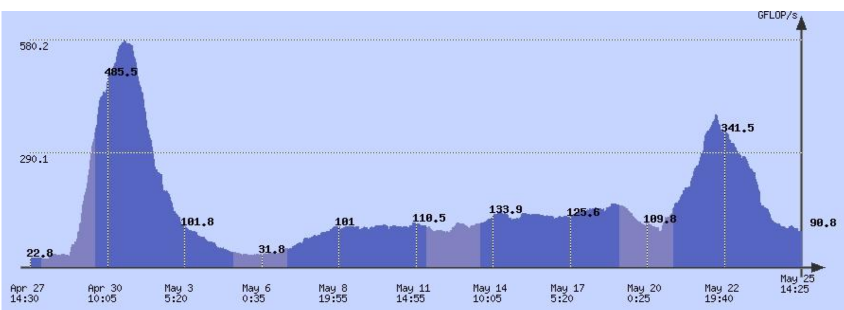

Typical performance timeline

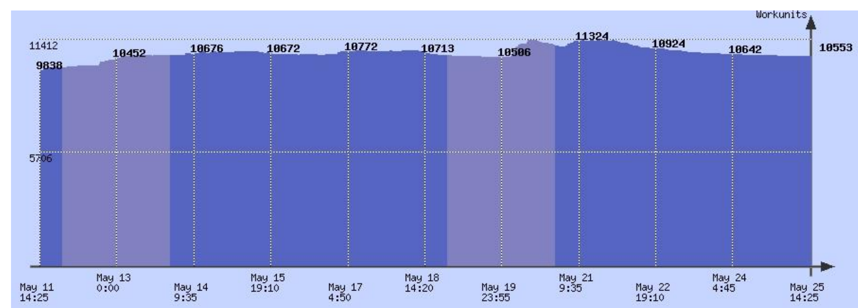

Number of in-progress workunits

SLinCA@Home – Volunteer Community

IMP team mostly cooperate with a national-wide Ukrainian user community on the premises of the public site “Distributed Computing Team ‘Ukraine’” (http://distributed.org.ua), discussed and contributed the best practices in DCI operations at their special fora.

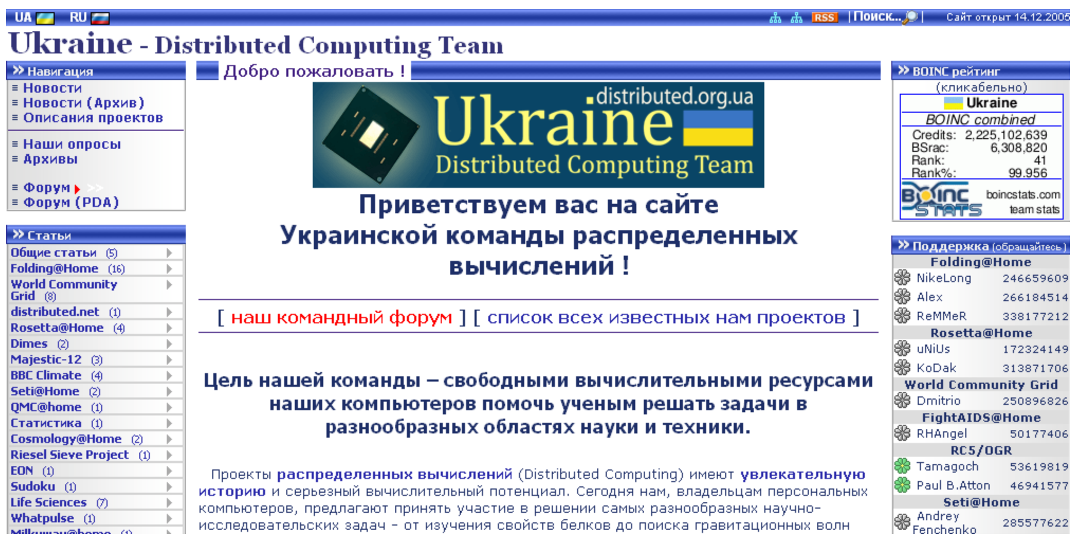

Use Case 1: Stress relaxation phenomena in Al/Cu/Si crystals

<u>Application:</u> LAMMPS-over-DCI (LAMMPS wrapped in DC-API by SZTAKI)
<u>Science Community:</u> physics, materials science – 2 user groups (6 end users): IMP + KNU (Taras Shevchenko Kyiv National University, Kyiv, Ukraine).
<u>Scientific Aim</u>: MD simulation of relaxation behavior of stresses in nanocrystals.

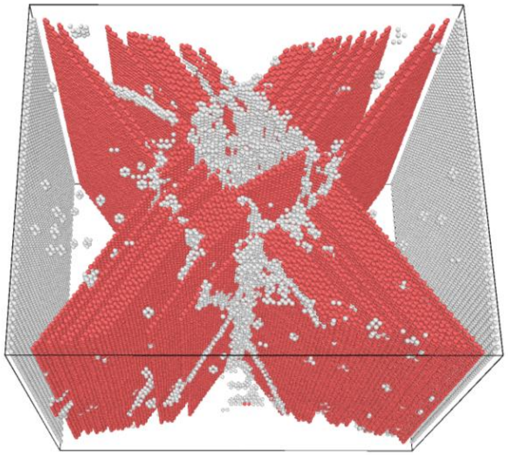

The evolution of the defect substructure in Al nanocrystal (after 15 ps, defects only are shown: gray color — point defects, red color — stacking faults) .

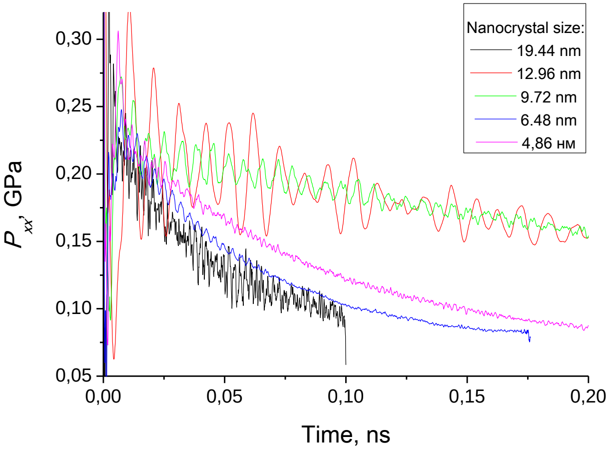

Oscillation of the internal stress P_{xx} for different nanocrystal sizes.

Use Case 2: Nanoindentation of graphene membrane

Application: LAMMPS-over-DCI (LAMMPS wrapped in DC-API by SZTAKI)
Science Community: nanotechnologies – 2 user groups (5 end users): IMP + SPM&RS-Centre (Centre of scanning probe microscopy and resonance spectroscopy, Kyiv, Ukraine).
Scientific Aim: MD simulation of nanoindentation was performed for monolayer graphene membrane in an atomic force microscope.

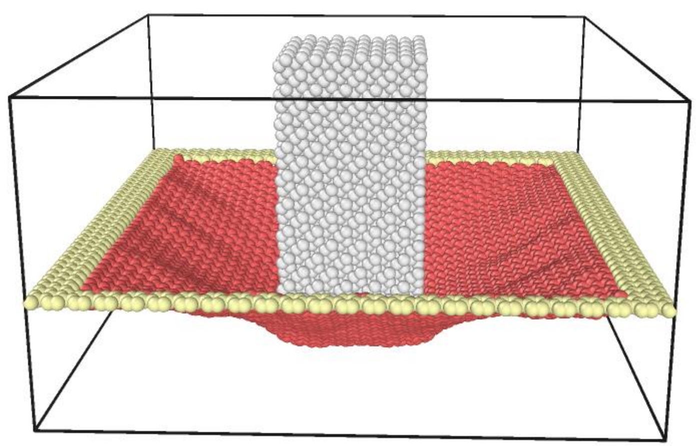

Visualization of nanoindentation:
red color — mobile graphene atoms,
yellow color — fixed graphene atoms,
gray color – nanoindentor atoms.

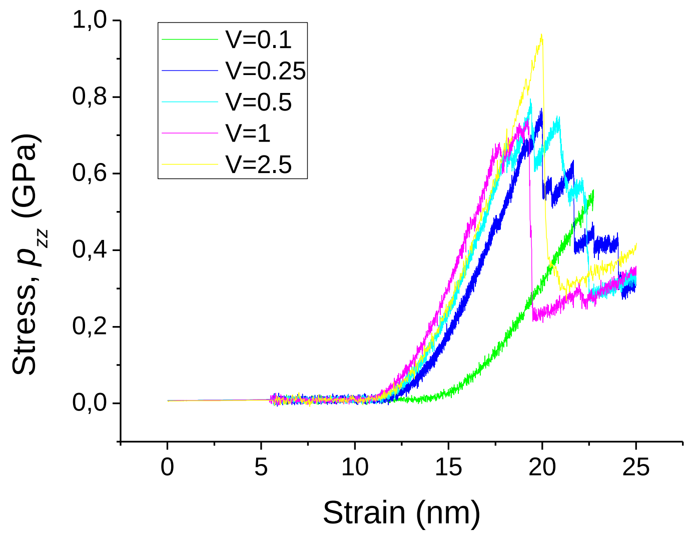

Stress-strain dependencies for different speeds (A/pm) of the nanoindentor.

Use Case 3: Thermal stability of boron nitride nanotubes

Application: LAMMPS-over-DCI (LAMMPS wrapped in DC-API by SZTAKI)
Science Community: nanotechnologies, materials science – 2 user groups (4 end users): IMP + IPMS (Frantsevich Institute for Problems in Materials Science, Kiev, Ukraine).
Scientific Aim: MD simulation of boron nitride nanotubes (BNNT) with exceptional physical properties, which are a prerequisite for their wide practical applications in the future.

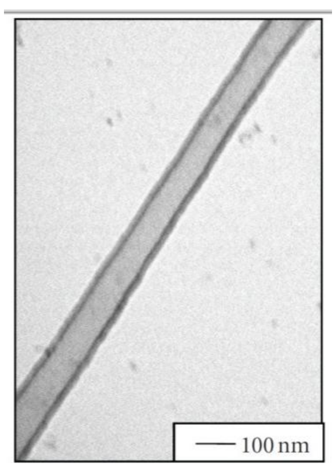

example of boron nitride nanotube (TEM-image)[3]

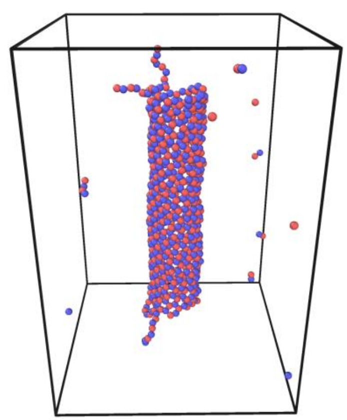

Collapse of NTNB (after 0.1 ns, nitrogen atoms are shown by red color, boron atoms — by blue color).

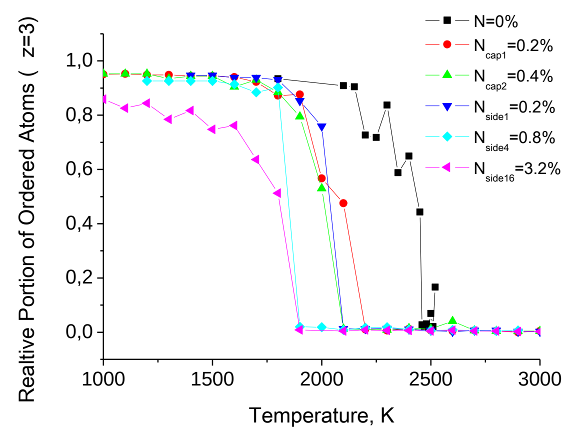

Dependence of decay temperature on the vacancies concentration and location (cap or side) in BNNT

Use Case 4: Manipulations with carbon nanotubes

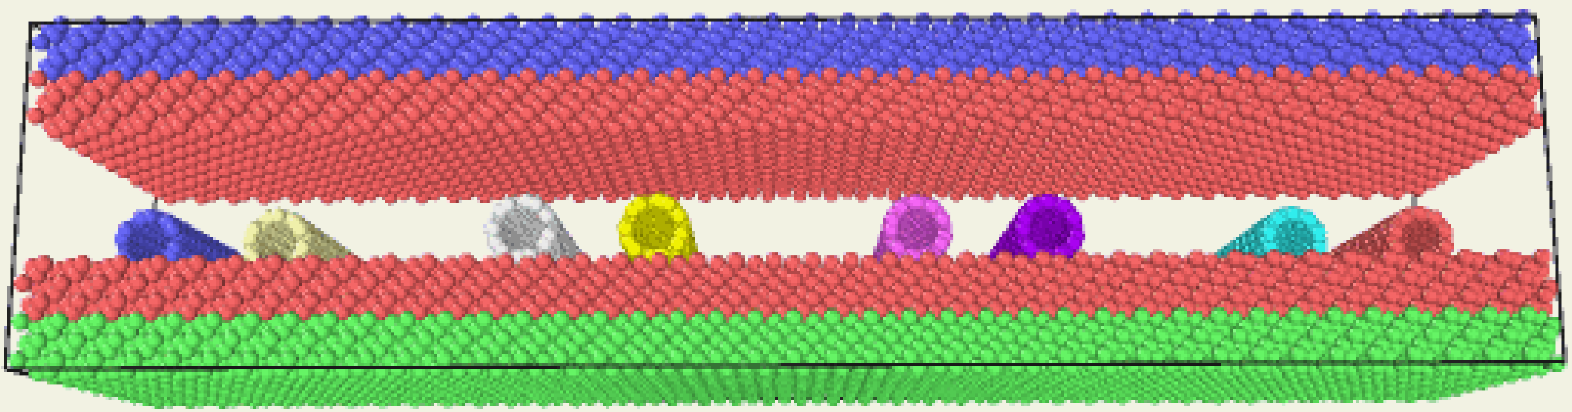

Detachment of m-CNTs after application of driving force per atom **F=0.17 eV/A** and usage of the second Si-substrate ("stamp") in the presence of s-CNTs:
two m-CNT c(6,6); two s-CNT c(7,5), two s-CNT c(9,2), and two m-CNT c(10,0) (from left to right).

Publications

1.Change of Scaling and Appearance of Scale-Free Size Distribution in Aggregation Kinetics by Additive Rules, *Physica A*: Statistical Mechanics and its Applications (2014) DOI: 10.1016/j.physa.2014.06.028, prepint at arxiv.org.

2.IMP Science Gateway: from Portal to Hub of Virtual Experimental Labs in Materials Science, International Workshop on Science Gateways (IWSG-2014) (Dublin, Ireland), preprint at arxiv.org.

3.The Oscillatory Nature of Internal Stress Relaxation due to Self-Organization of the Defect Substructure in Nanocrystals, IV Int. Conf. "Nanoscale Systems: Structure, Properties, Technologies" ("Наноразмерные системы: строение, свойства, технологии") (Kyiv, Ukraine, 2013/11/22), poster (PDF).

4.Thermal Stability of Boron Nitride Nanotubes, IV Int. Conf. "Nanoscale Systems: Structure, Properties, Technologies" ("Наноразмерные системы: строение, свойства, технологии") (Kyiv, Ukraine, 2013/11/22), poster (PDF).

5.Elastic Properties and Strength of Graphene under Nanoindentation Conditions, IV Int. Conf. "Nanoscale Systems: Structure, Properties, Technologies" ("Наноразмерные системы: строение, свойства, технологии") (Kyiv, Ukraine, 2013/11/22), poster (PDF).

6.Complex Workflow Management and Integration of Distributed Computing Resources by Science Gateway Portal for Molecular Dynamics Simulations in Materials Science, Proc. Third International Conference "High Performance Computing" HPC-UA 2013 (Ukraine, Kyiv, October 7-11, 2013), 148-155, PDF.

7.Usage of IMP Science Gateway for Molecular Dynamics Simulations of Various Metal-Organic Nanostructures, Computer Science and Information Technologies, CSIT 2013 (Armenia, Yerevan, September 23-27, 2013), 399-402, color electronic version (high quality); black-white printed version (low quality).

8.Application of the Science Gateway Portal on the Basis of WS-PGRADE Technology for Simulation of Aggregation Kinetics and Molecular Dynamics Simulations of Metal-Organic Nanostructure, International Workshop on Science Gateways (Zurich, Switzerland), 03-05 June, 2013, PDF.

9.From Quantity To Quality: Massive Molecular Dynamics Simulation of Nanostructures under Plastic Deformation in Desktop and Service Grid Distributed Computing Infrastructure, Computer Science, 14, No.1 (2013) DOI: 10.7494/csci.2013.14.1.27.

10.Application of Desktop Grid Technology in Material Science, Proc. of 5th International Conference "Distributed Computing and Grid-technologies in Science and Education" GRID'2012 (Dubna, Russia) pp.123-129, color electronic version (high quality); black-white printed version (low quality).

Example 1:
SLinCA@Home IMP Desktop Grid

size: small (even nano)

history: short (2009-2014-...)

hosts: CPU

Let’s measure
its averaged (per host)
performance characteristics...

IMP SLinCA@Home DG - **CPUs/host**

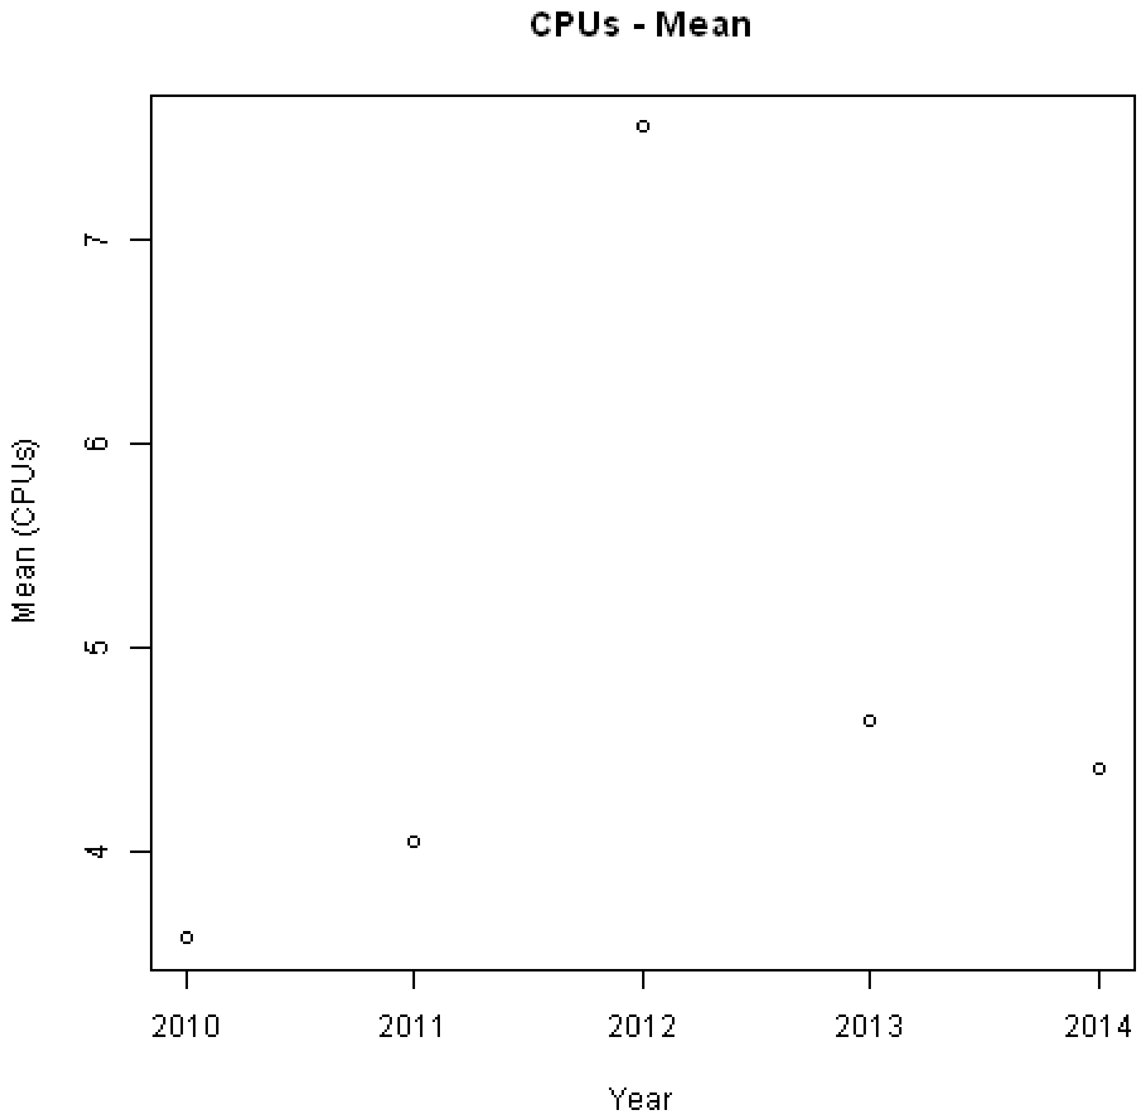

IMP DG - Performance - **IOPs**

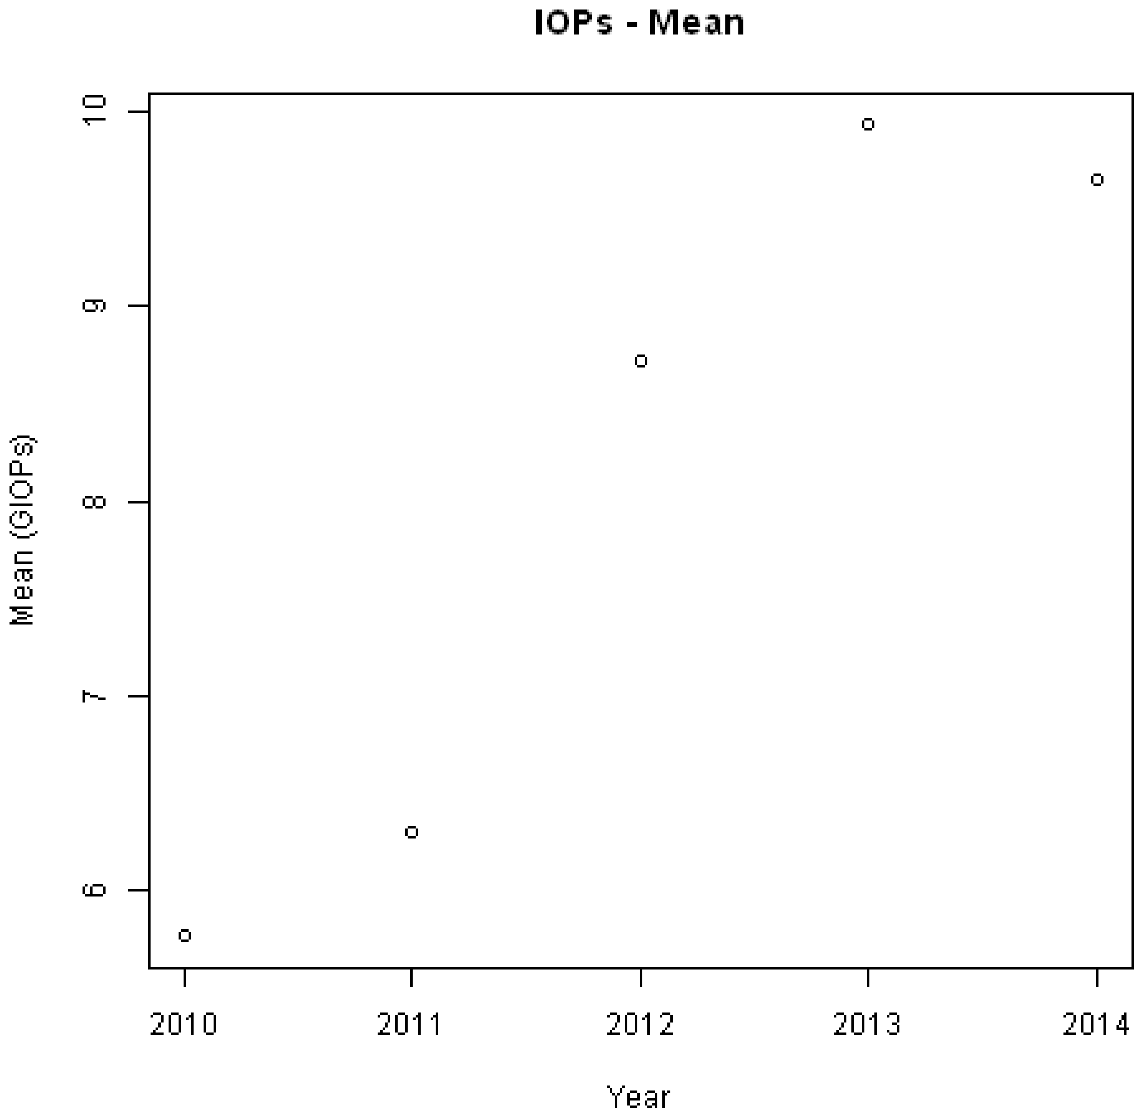

IMP DG - Performance - **FLOPs**

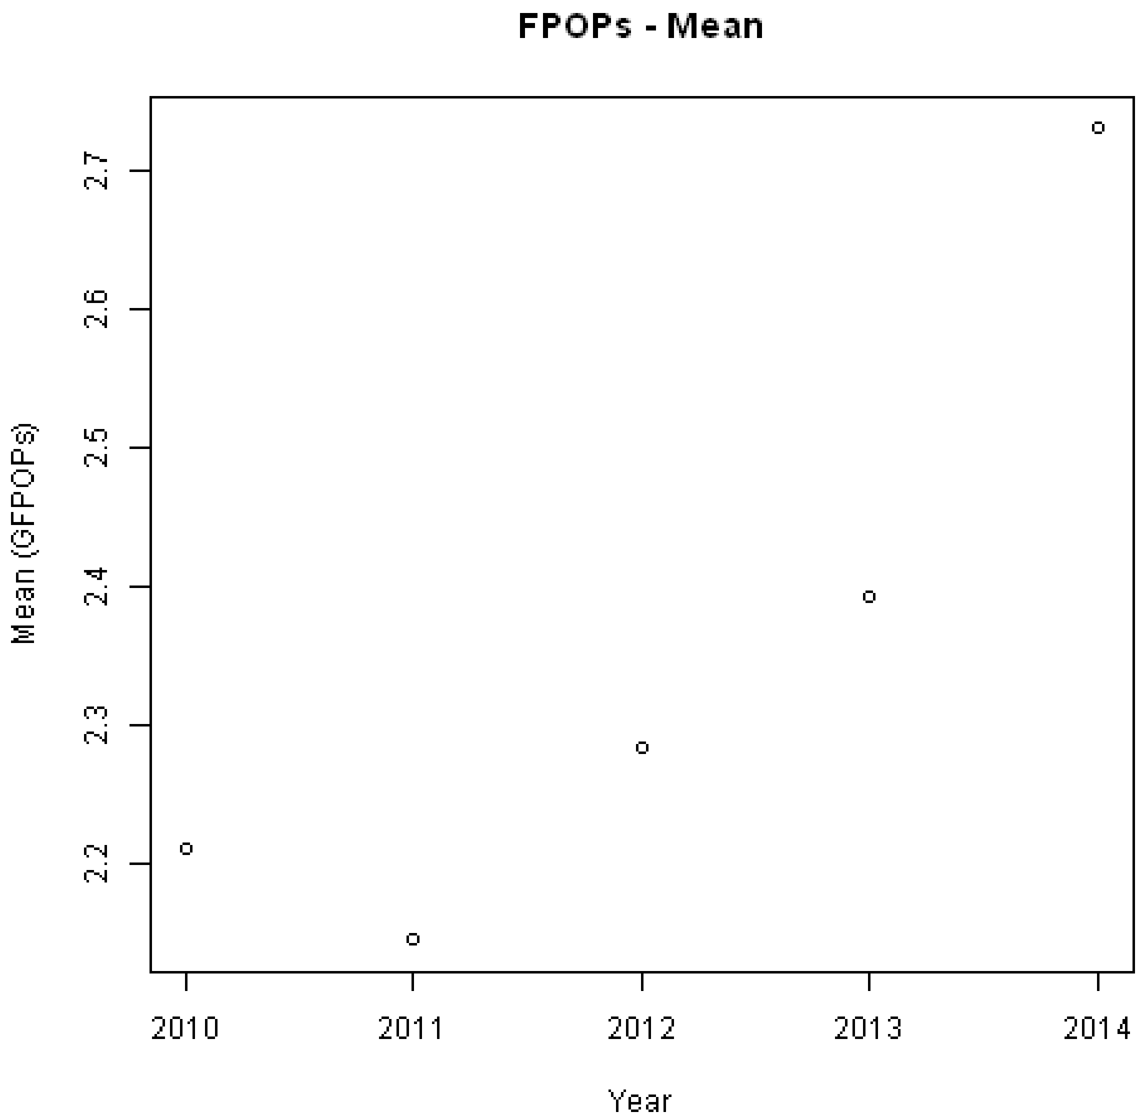

Distribution of Performances

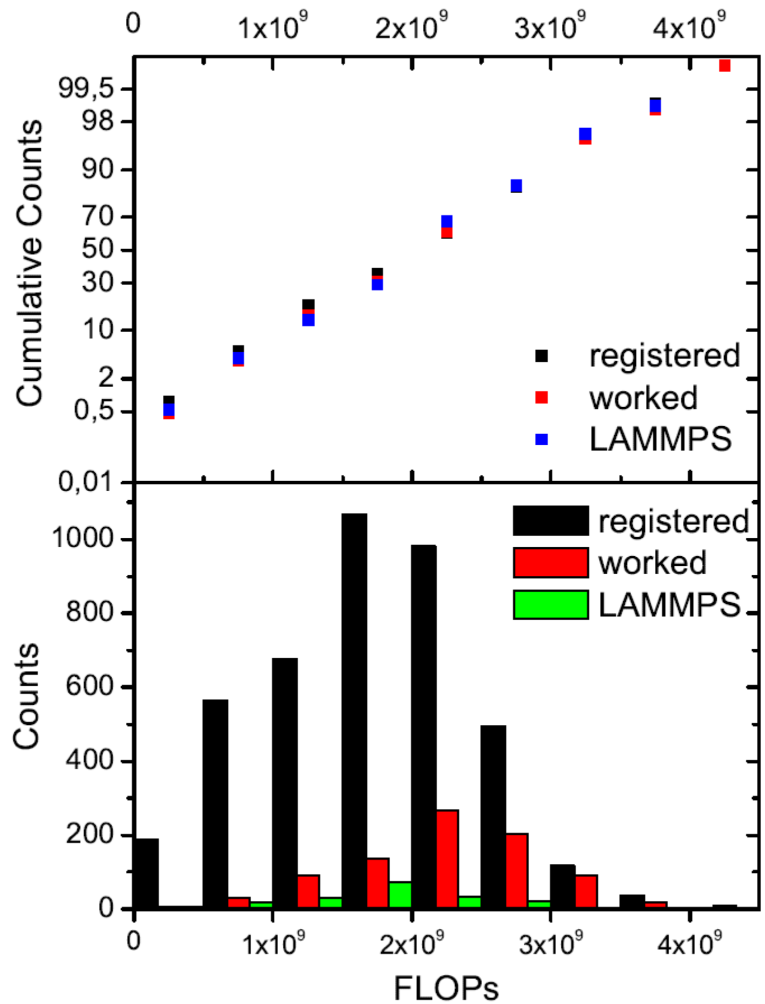

It seems to be **normal**,
BUT
it is more close to **Weibull**!

And Weibull was observed also at Ibercivis BOINC project

From private communication with Alejandro Rivero (in 2012), manager of Volunteer Computing at Ibercivis (2007 to 2011)

Floating Point Operations Per second (FPOPs) in linear-probabilistic coordinates among the registered hosts (black), the worked hosts (red), and the hosts used for LAMMPS-application (green)

Gordienko et al, *Computer Science*, 14 (1), 27-44 (2013)

Distribution of CPUs per Host

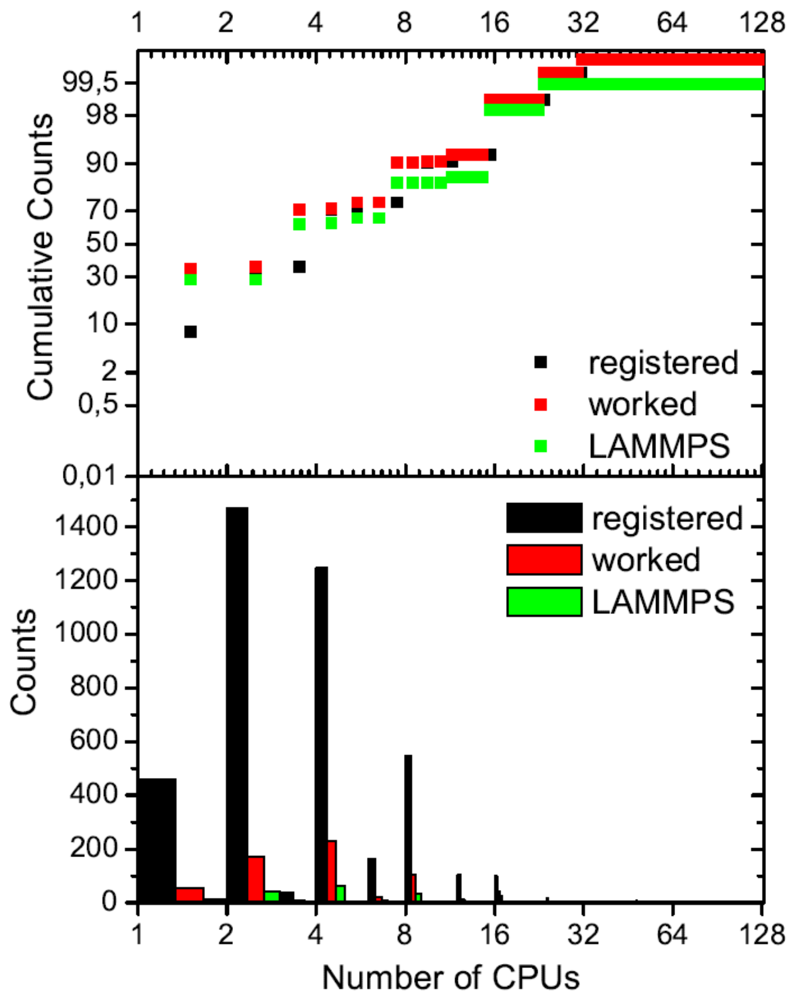

It seems to be **log-normal**

CPUs per host in the logarithmic-probabilistic coordinates (right) among the registered hosts (black), the worked hosts (red), and the hosts used for LAMMPS-application (green)

Gordienko et al, *Computer Science*, 14 (1), 27-44 (2013)

Distribution of RAM sizes

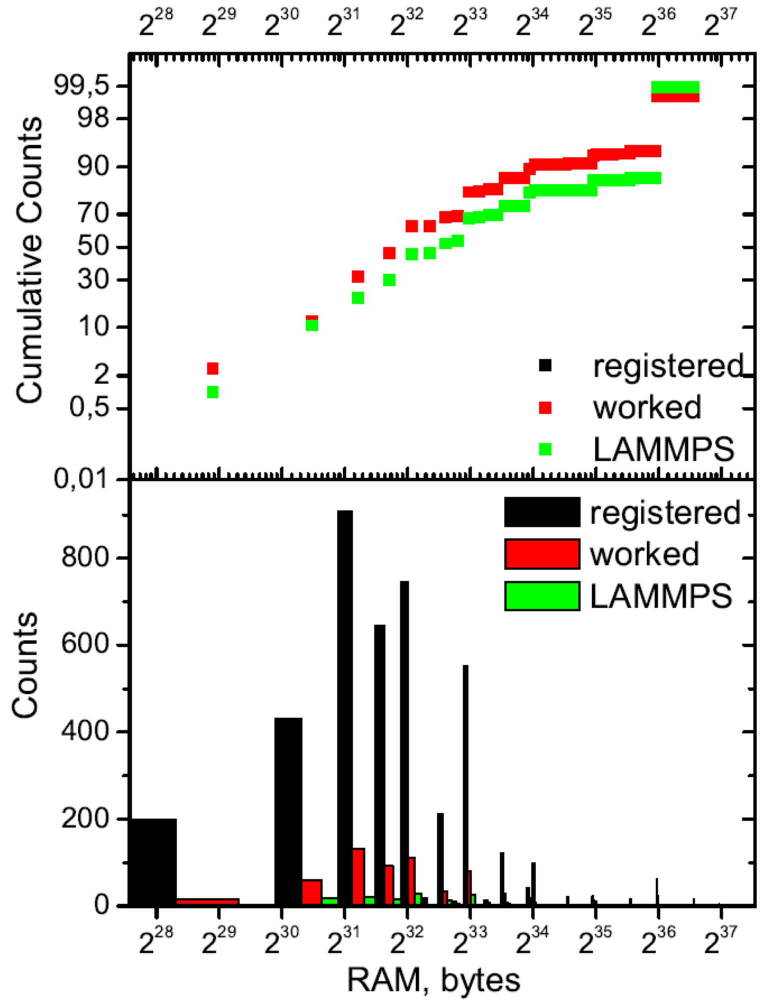

It seems to be **log-normal**… at least for low values

RAM sizes for the registered hosts (black), the worked hosts (red), and the hosts used for LAMMPS-application (green)

Gordienko et al, *Computer Science*, 14 (1), 27-44 (2013)

Distribution of HDD sizes

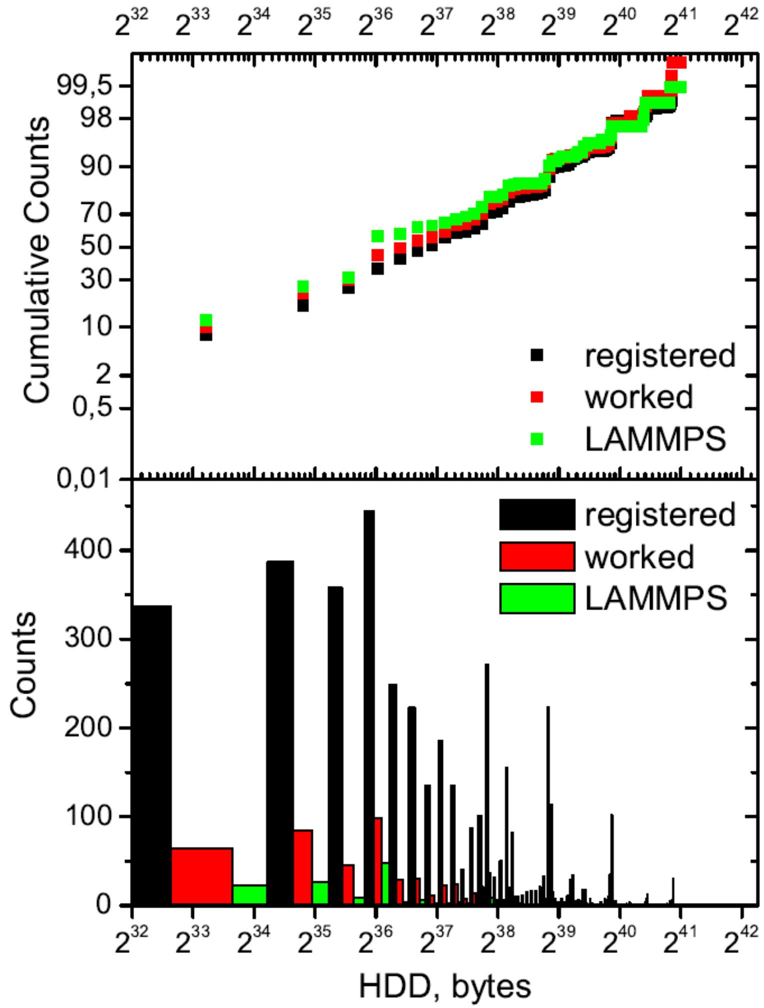

It seems to be **log-normal**… at least for low values

HDD sizes for the registered hosts (black), the worked hosts (red), and the hosts used for LAMMPS-application (green)

Gordienko et al, *Computer Science*, 14 (1), 27-44 (2013)

Example 2:
SZTAKI Desktop Grid

size: large,
history: long,
hosts: CPU

SZTAKI Desktop Grid

MTA SZTAKI

SZTAKI DESKTOP GRID

BOINC

Number of users: **126538** Number of hosts: **105221**
Estimated performance of last 48 hours: **1333.382 GFlop/s**
Peak performance: **3.4 TFlop/s**

Proud participant of IDGF
International Desktop Grid Federation | Supported by IDGF-SP | http://desktopgridfederation.org

About SZTAKI Desktop Grid

SZTAKI Desktop Grid is operated by the Laboratory of Parallel and Distributed Systems at the Hungarian Academy of Sciences, Budapest, Hungary.

The SZTAKI Desktop Grid and its applications are partly supported by the ongoing IDGF-SP project. The work leading to these results has received funding from the European Union Seventh Framework Programme (FP7/2007-2013) under grant agreement nº RI-312297. The experts of the International Desktop Grid Federation provide further support for the SZTAKI Desktop Grid infrastructure, its applications, and its integration into the European Grid Infrastructure.

Join SZTAKI Desktop Grid

Applications being run by SZTAKI Desktop Grid

Application	Field	Information	Launch	Results
Zeta-Search	Mathematics	Описание	Nov, 2013	Results
KOPI beta	Linguistics	Description (EN), Description (HU)	Sep, 2011	
BinSys	Mathematics	Description (EN), Description (HU)	Jun, 2005	Results
UC-Explorer	Physics	Описание	Feb, 2009	

http://szdg.lpds.sztaki.hu/szdg/

SZTAKI DG - CPUs/host

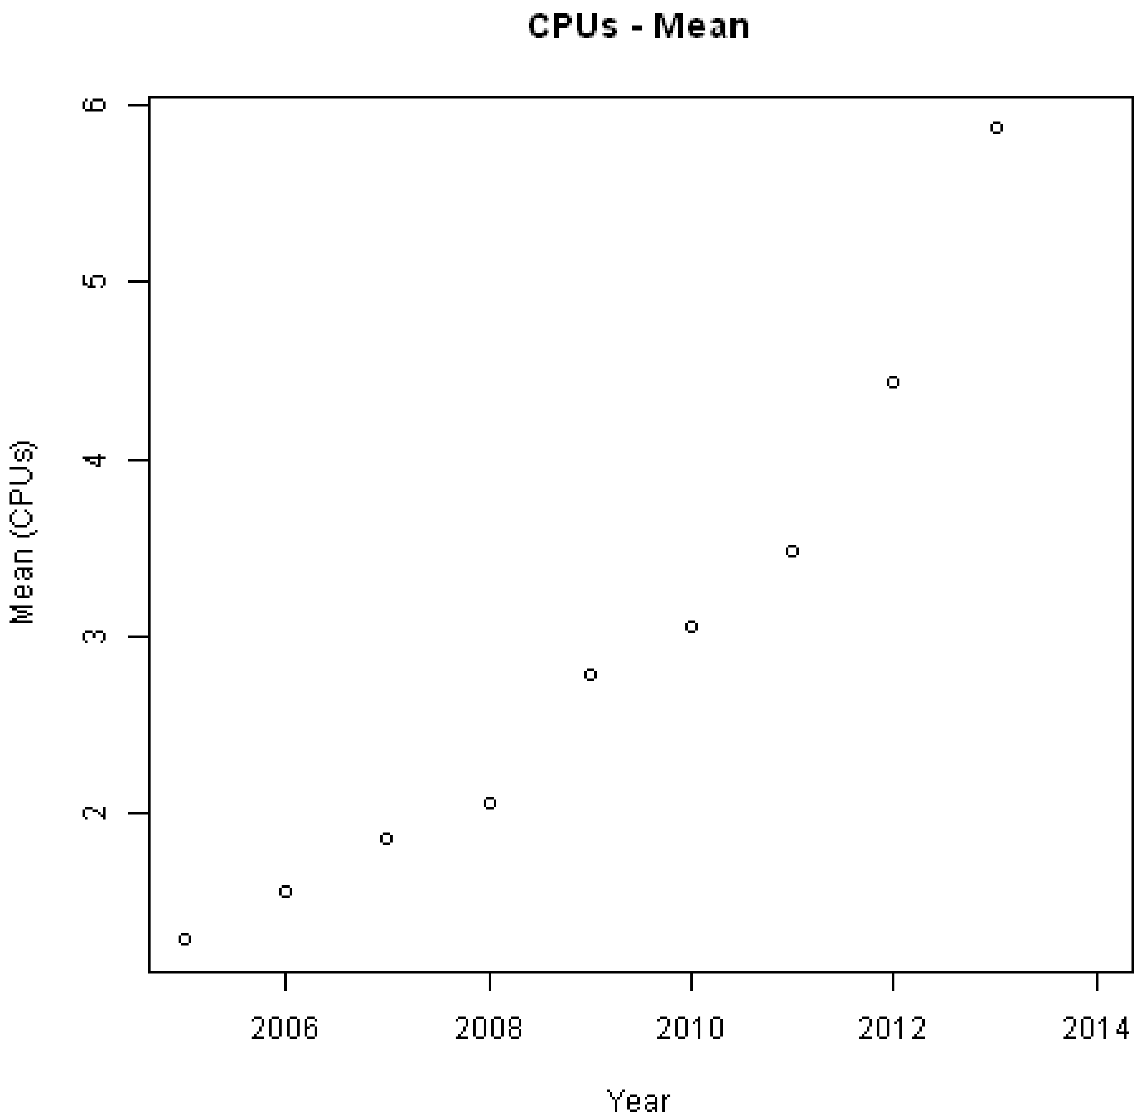

SZTAKI DG - Performance - IOPs

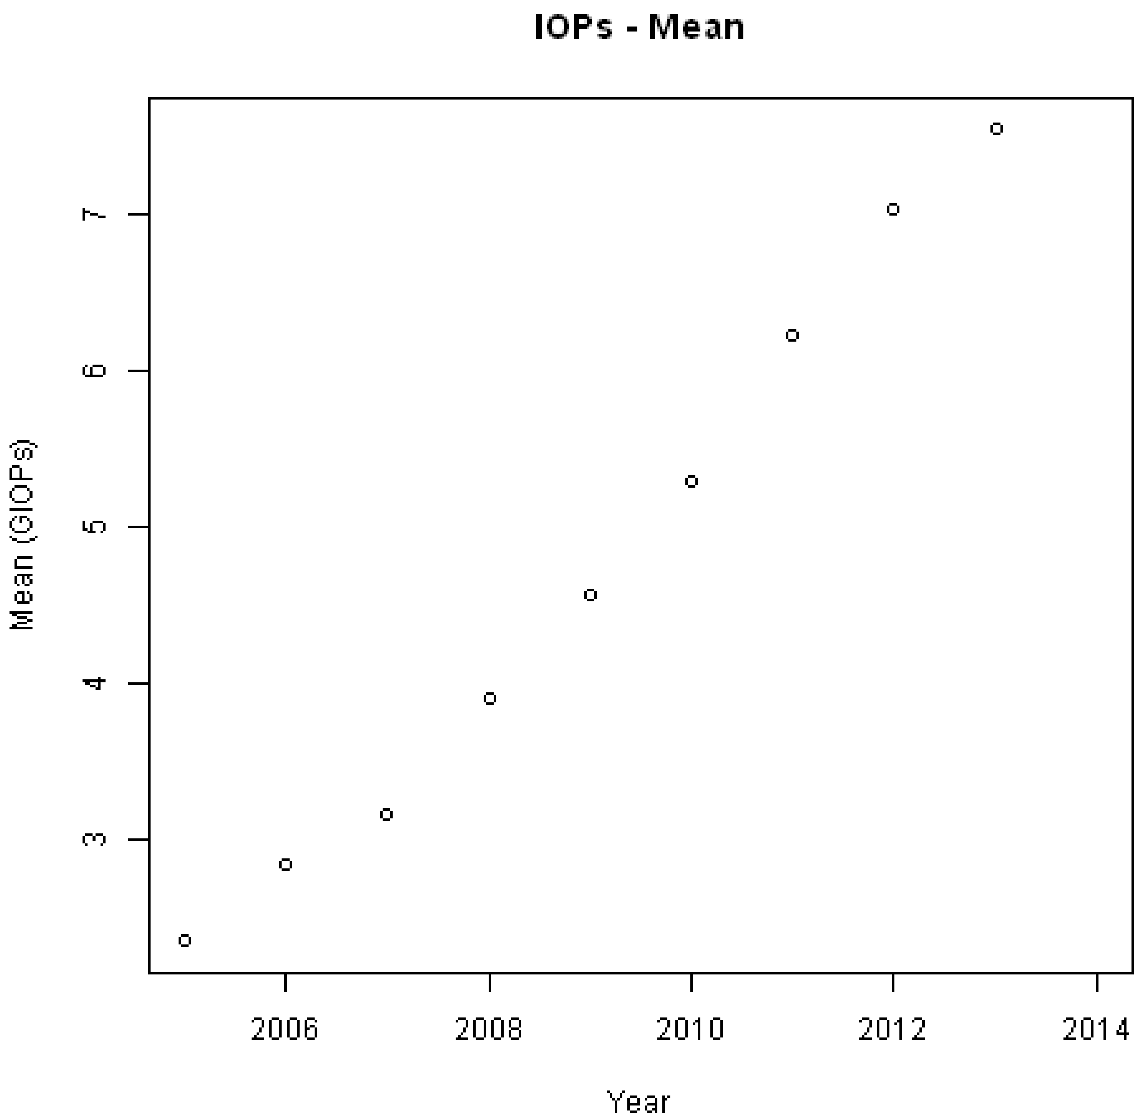

SZTAKI DG - Performance - FLOPs

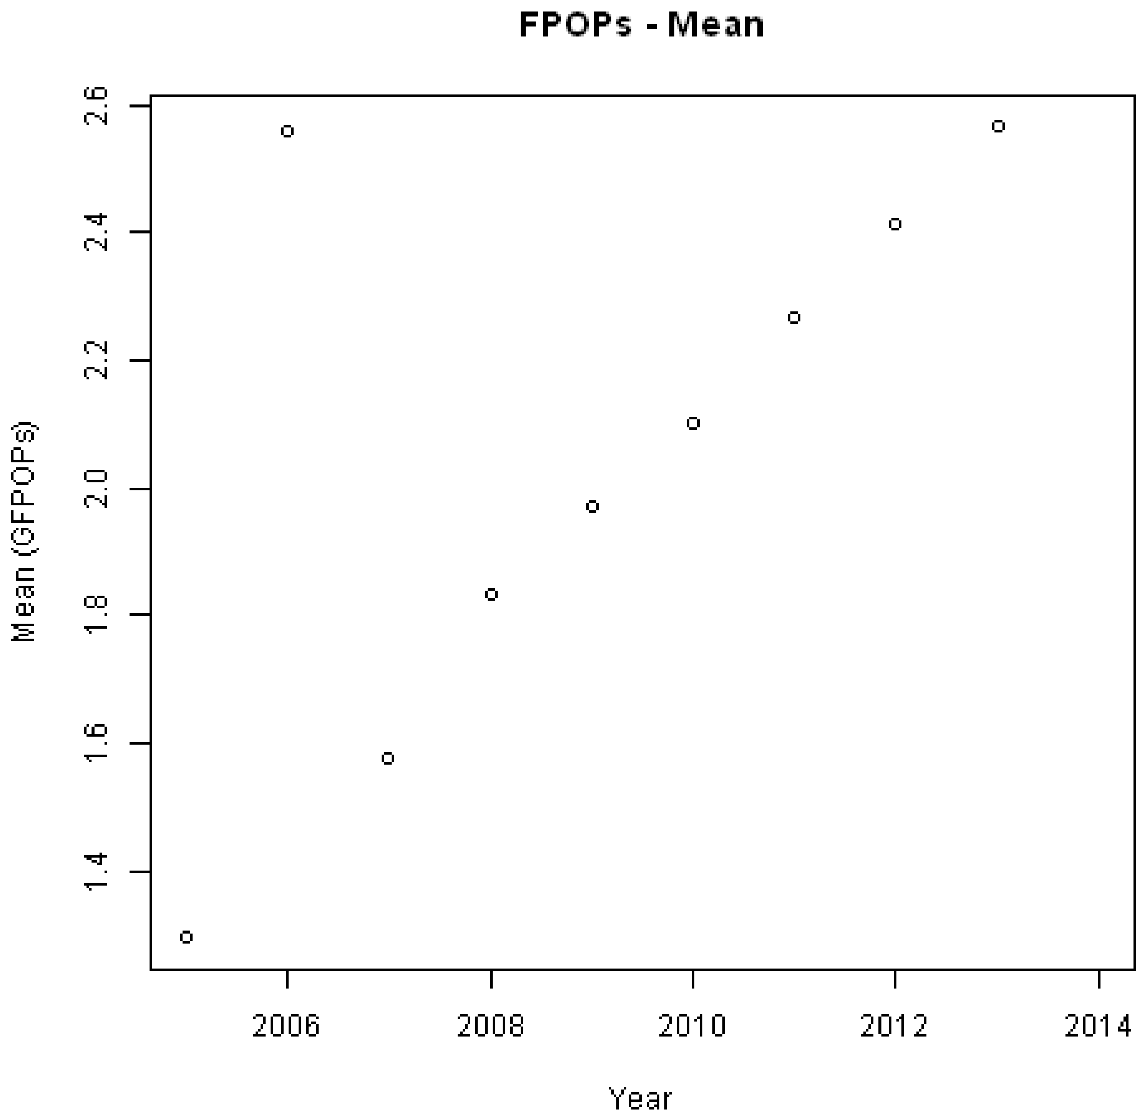

Example 3:
Ibercivis Desktop Grid

size: large,
history: long,
hosts: CPU

Ibercivis

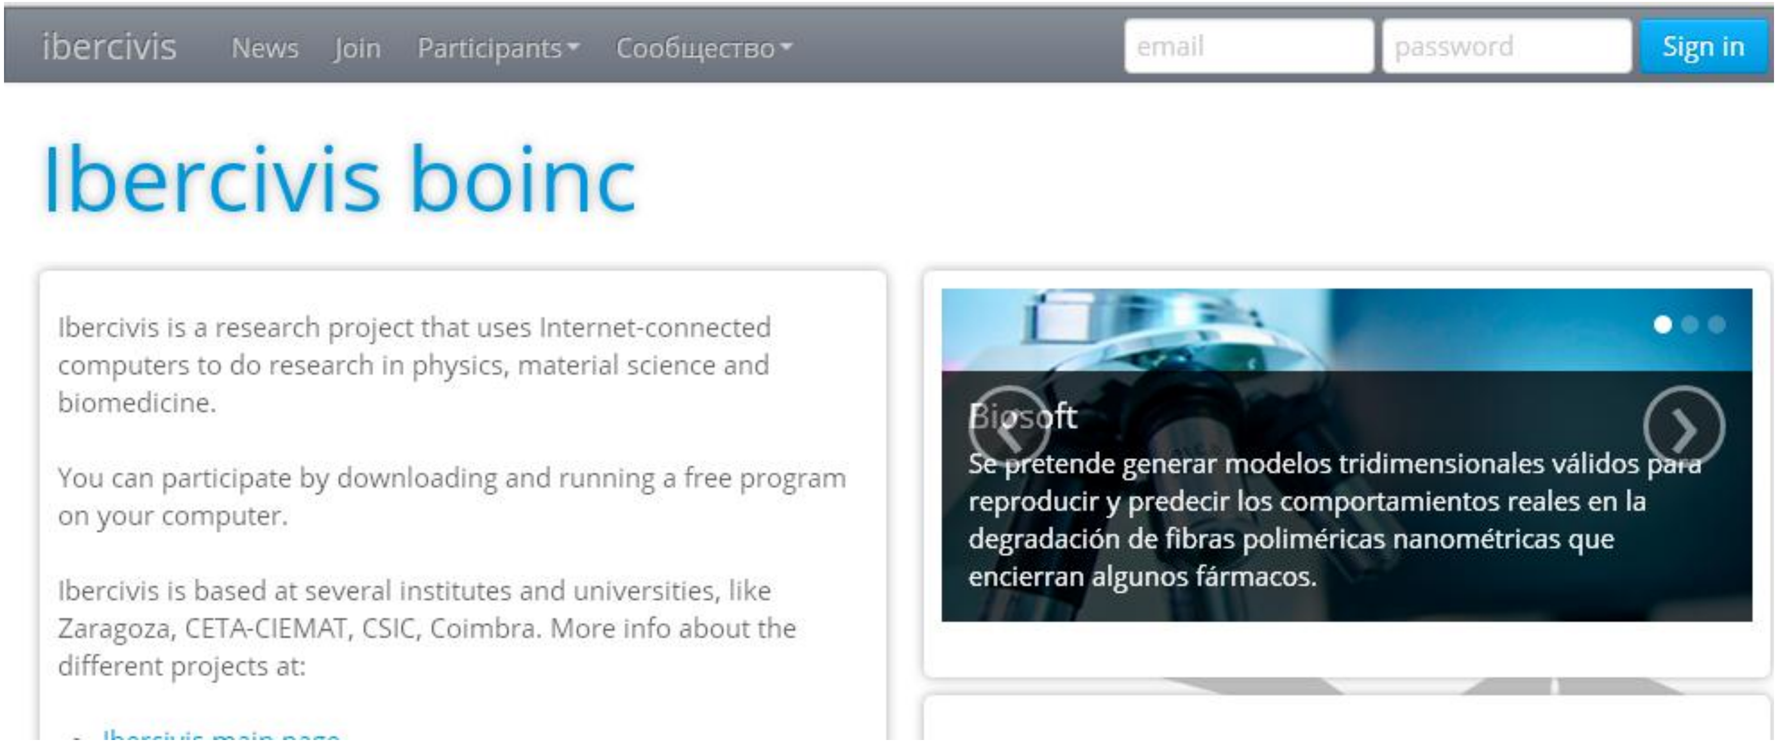

http://registro.ibercivis.es/

Ibercivis - **CPUs/host**

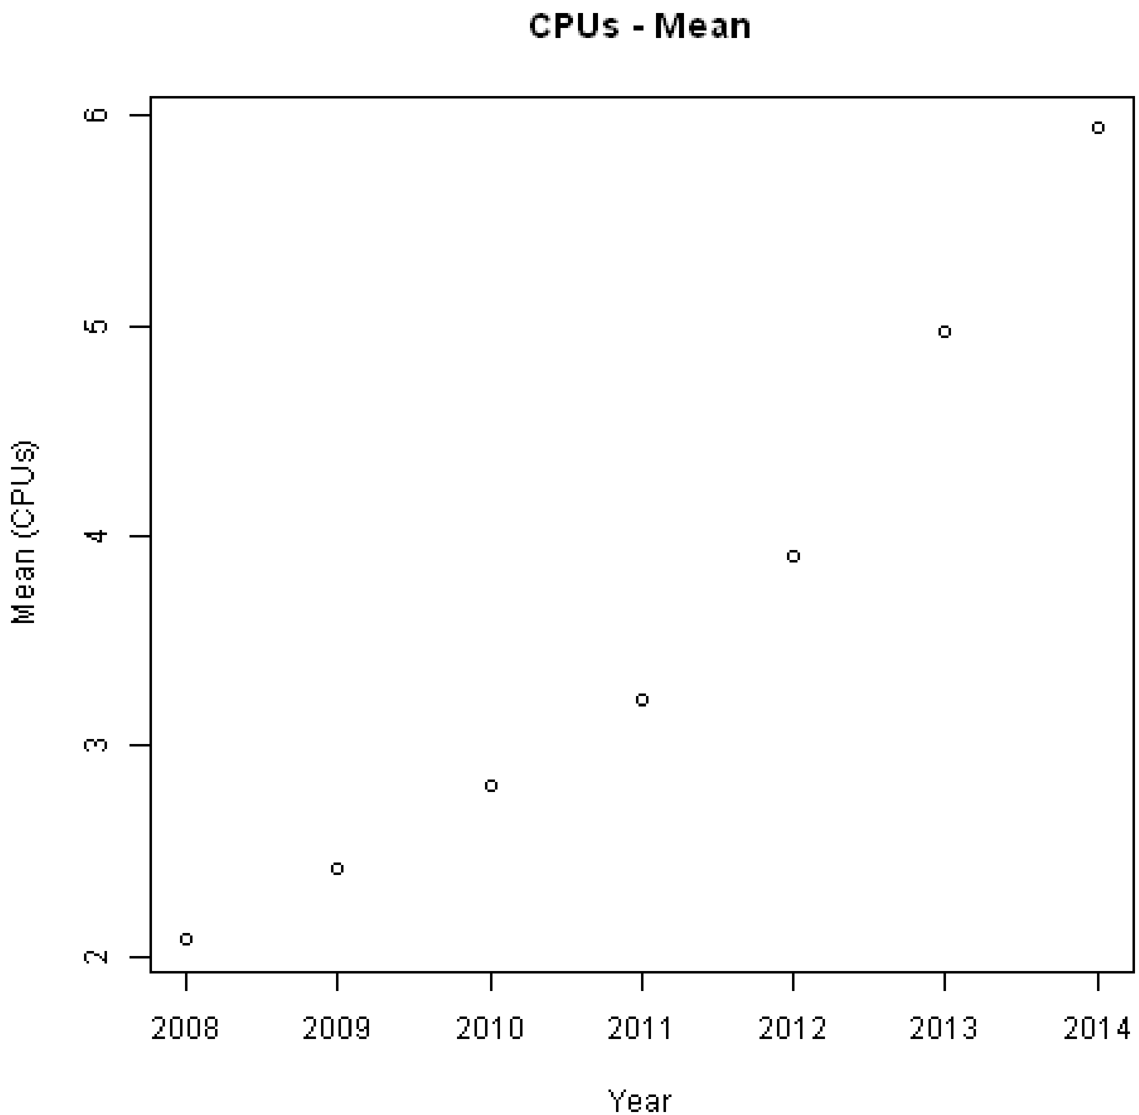

Ibercivis - Performance - **IOPs**

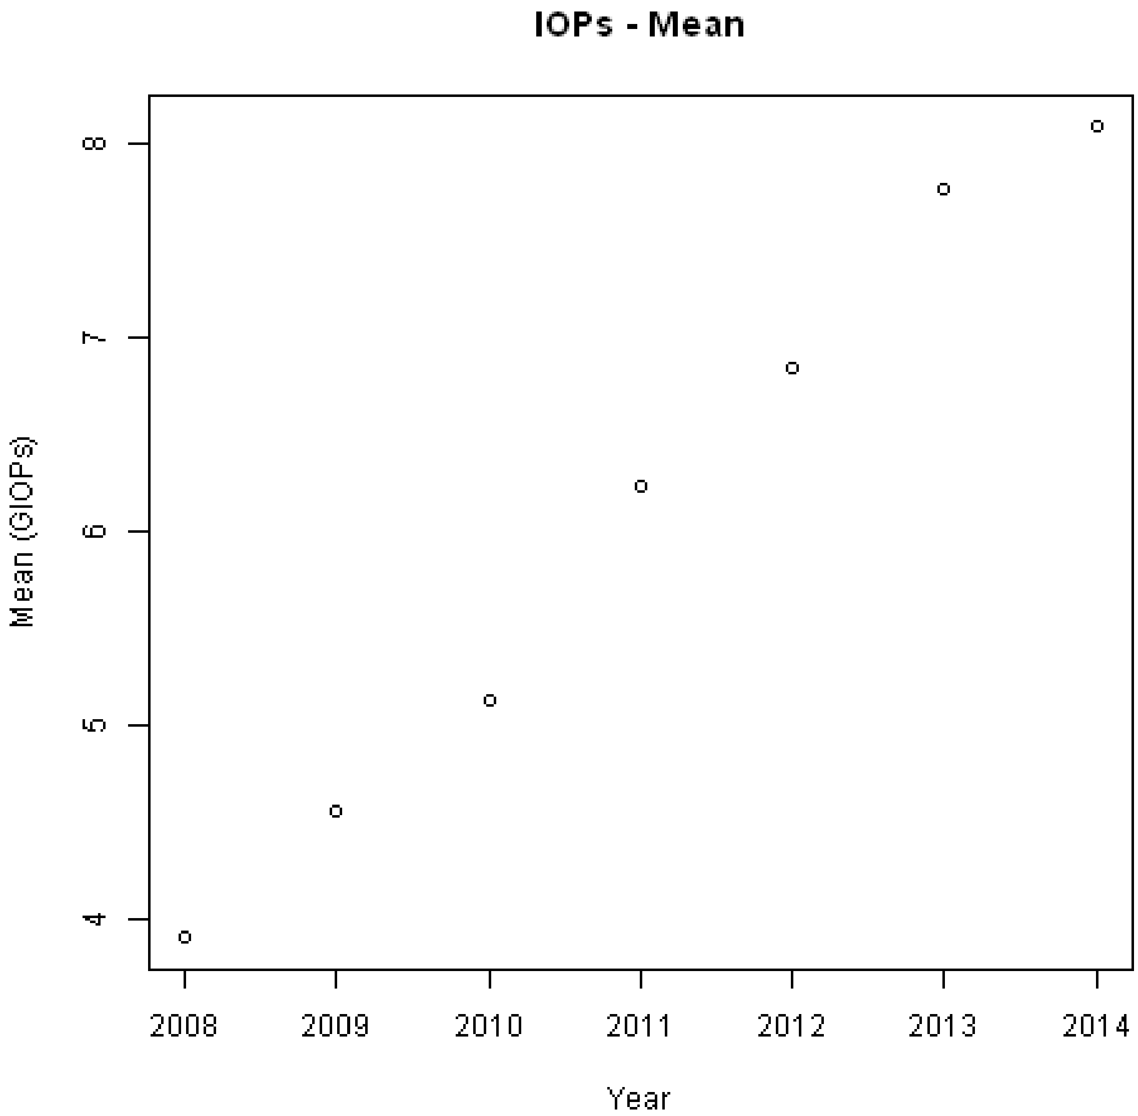

Ibercivis - Performance - **FLOPs**

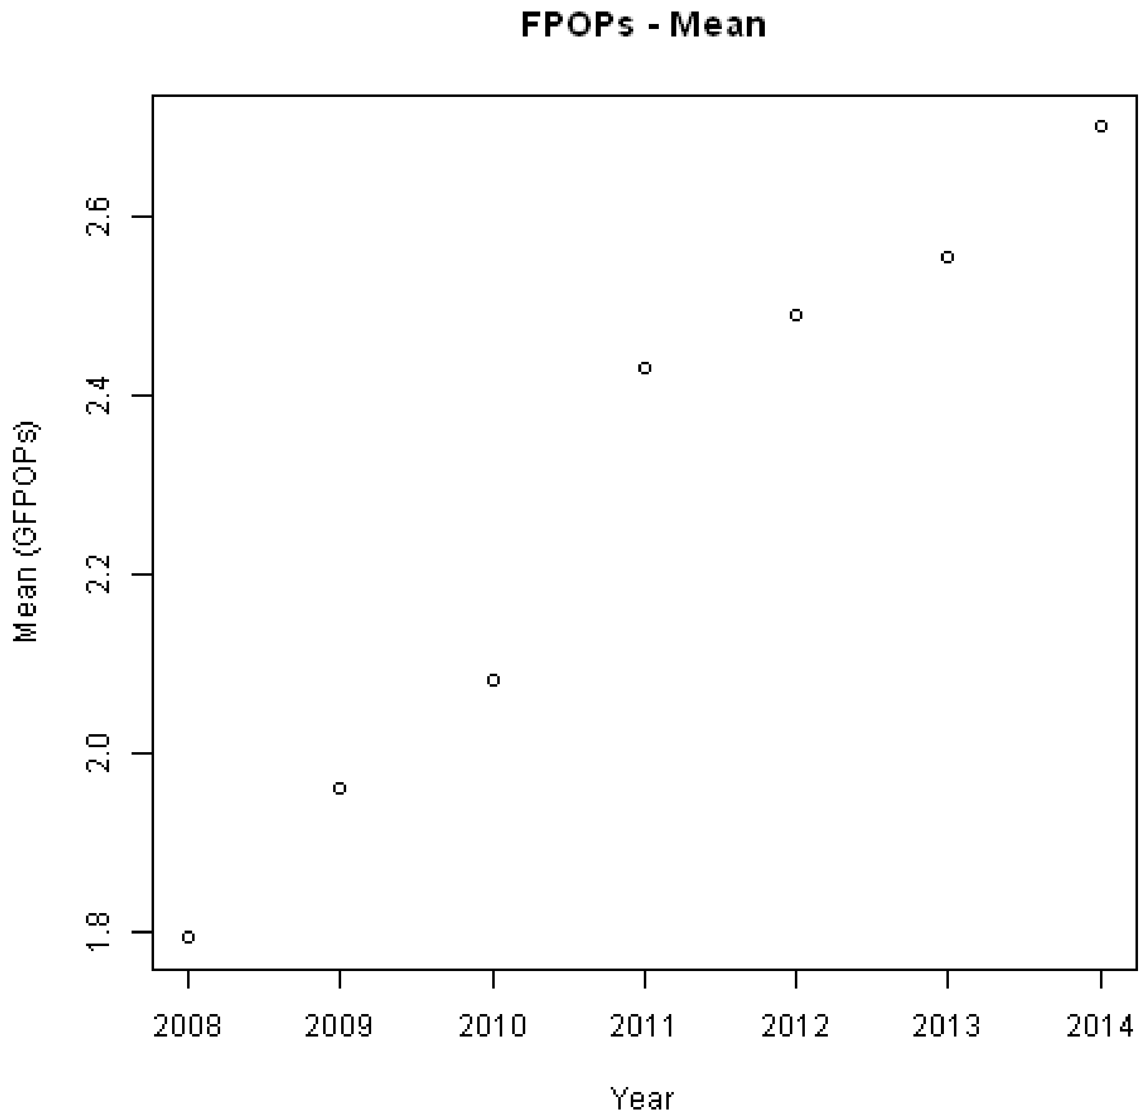

Example 4:
EDGeS Desktop Grid

size: medium,
history: short,
hosts: CPU

EDGeS

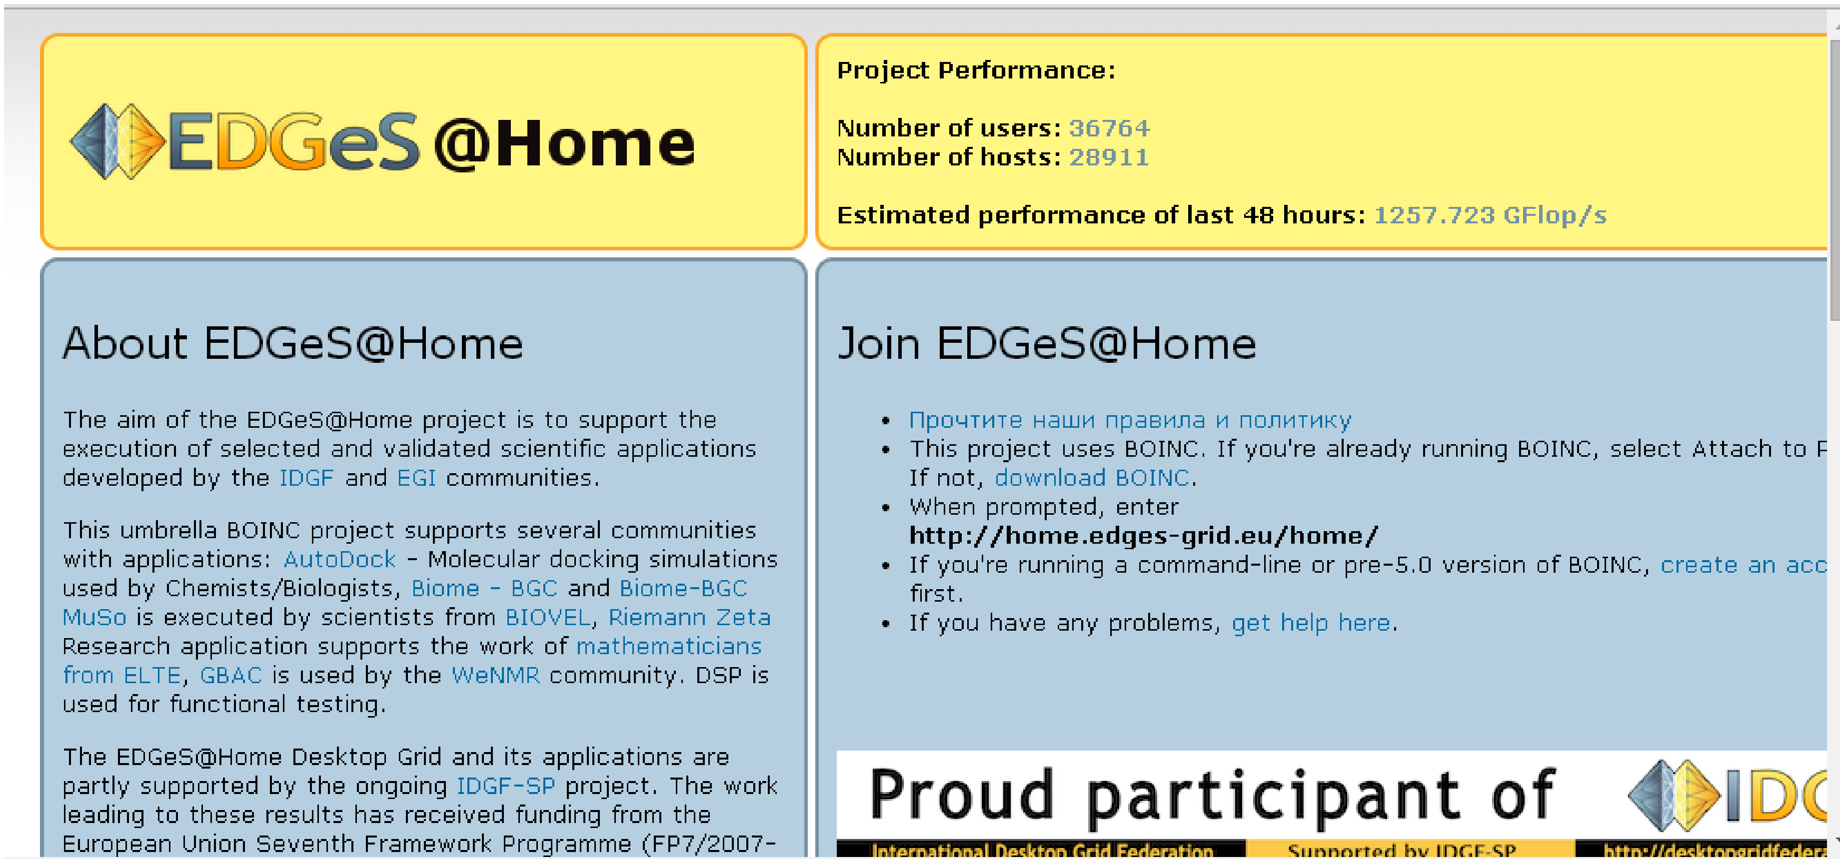

http://home.edges-grid.eu/home/

EDGeS DG - **CPUs/host**

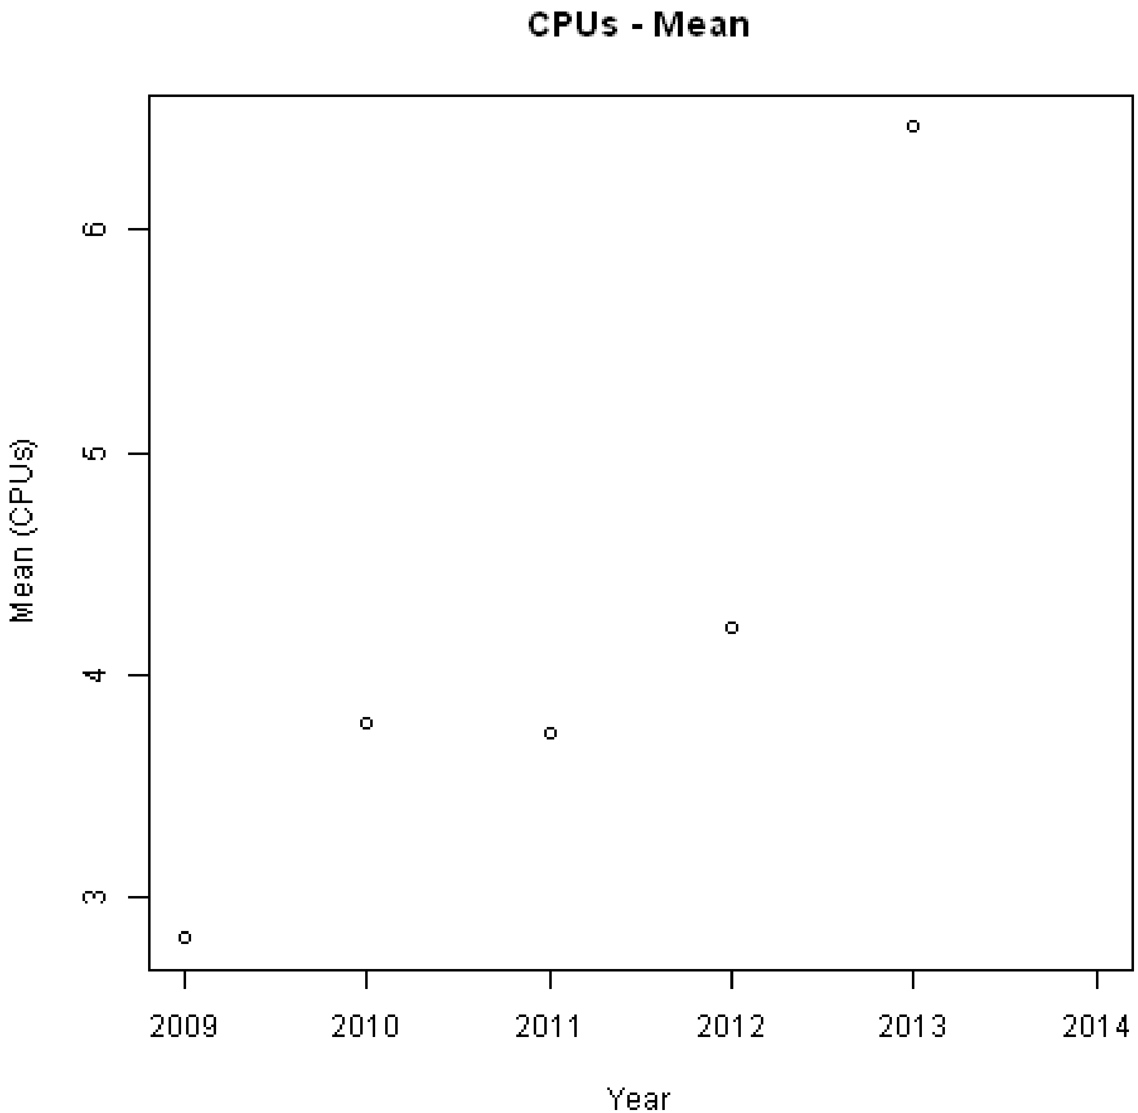

EDGeS DG - Performance - **IOPs**

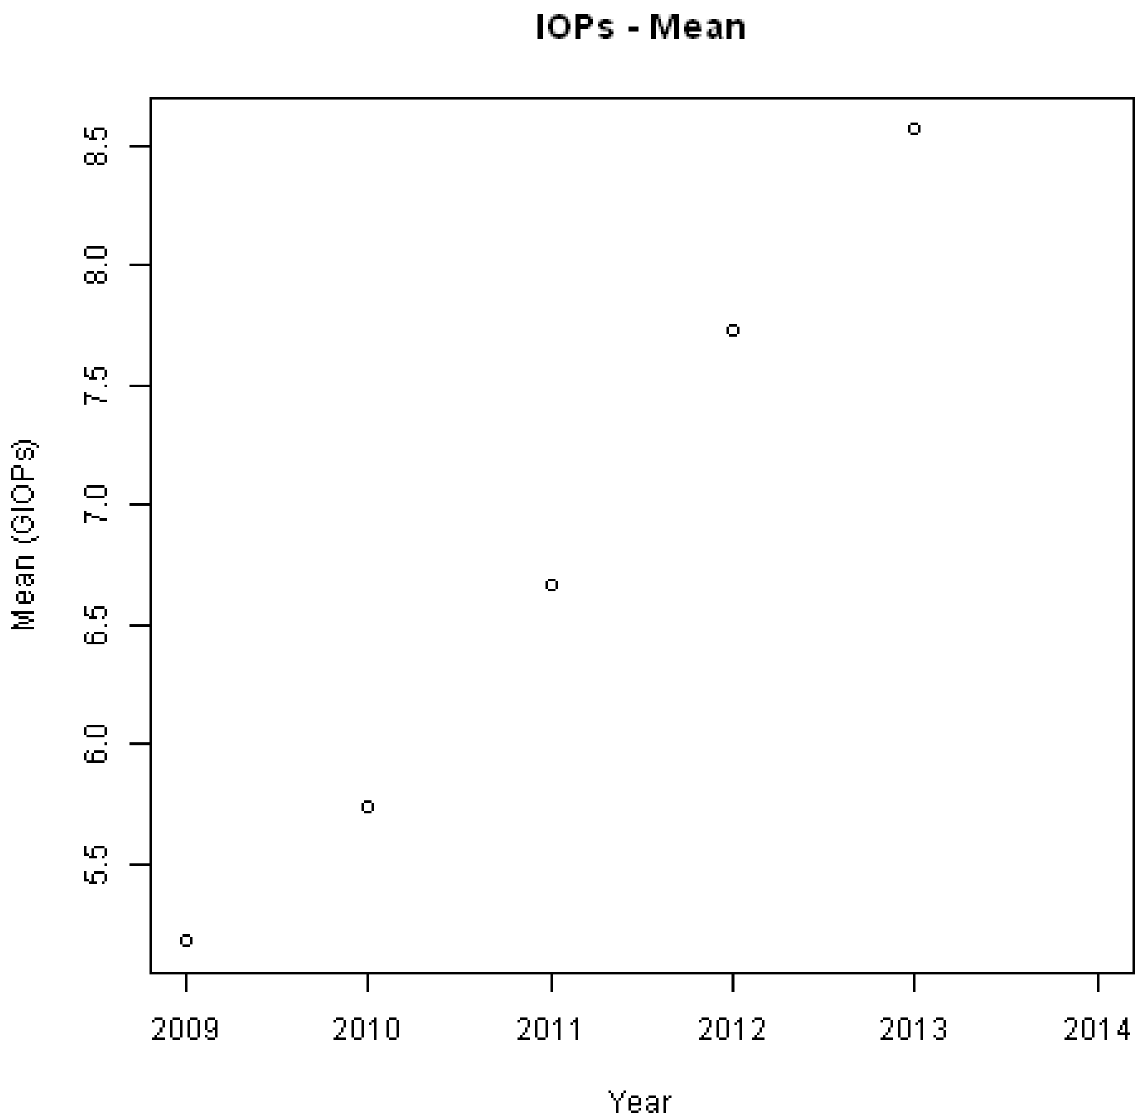

EDGeS DG - Performance - **FLOPs**

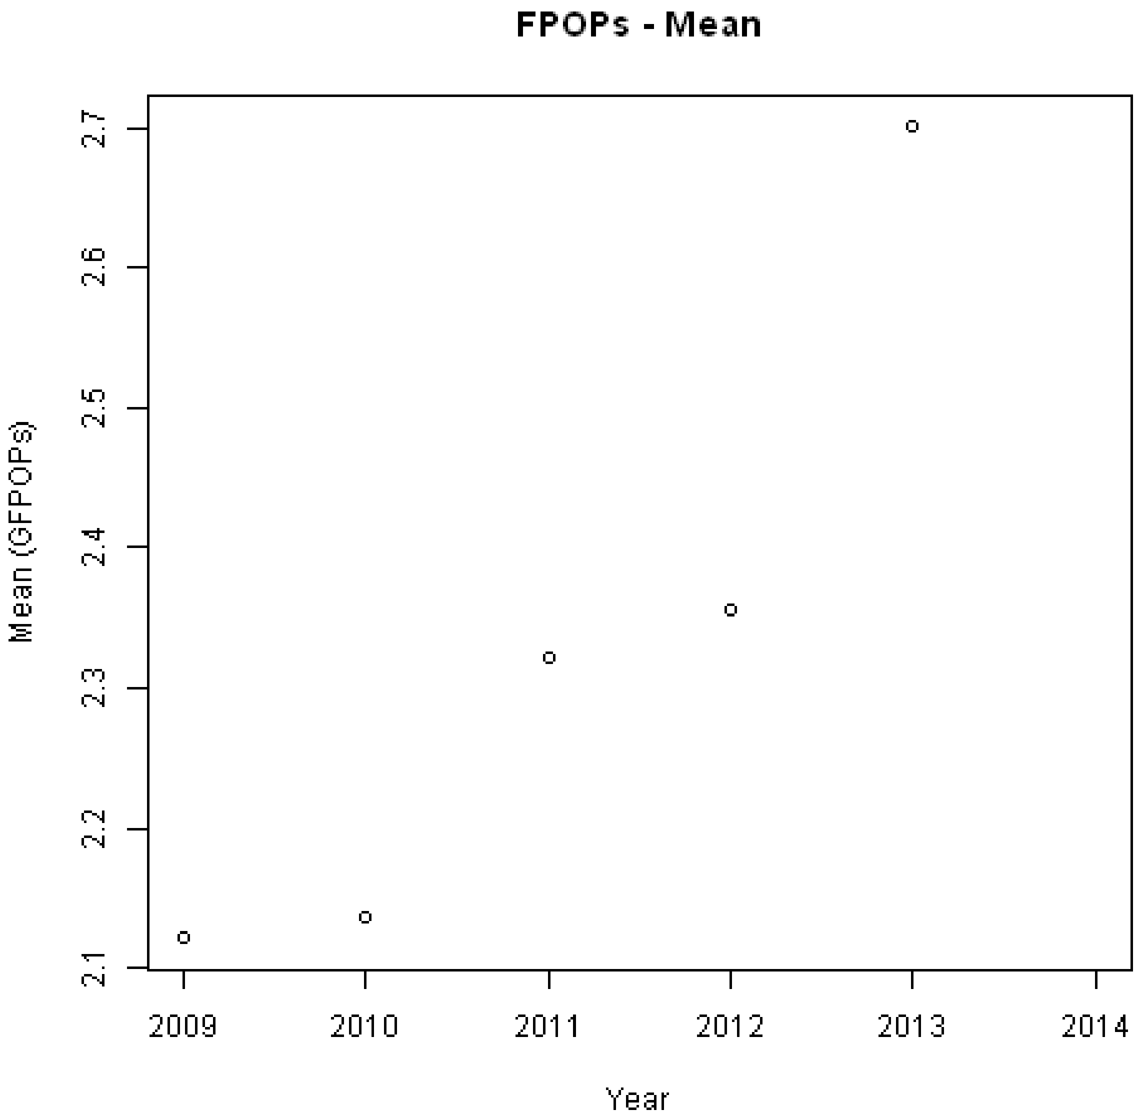

Example 5:
GPUGRID

size: medium,
history: medium,
hosts: GPU

GPUGRID

http://www.gpugrid.net/

GPUGRID - **CPUs/host**

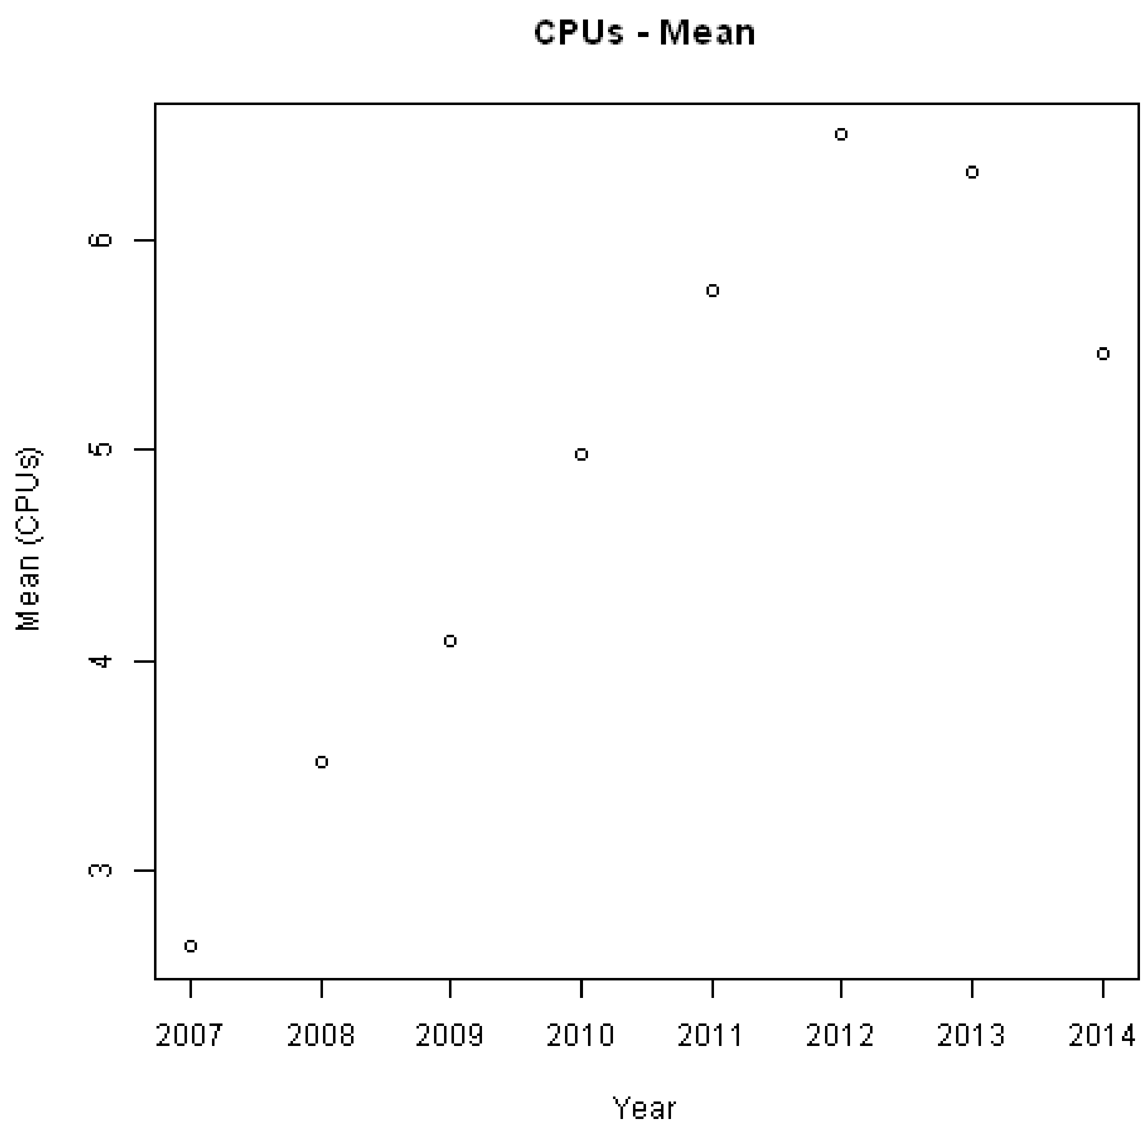

GPUGRID - Performance - **IOPs**

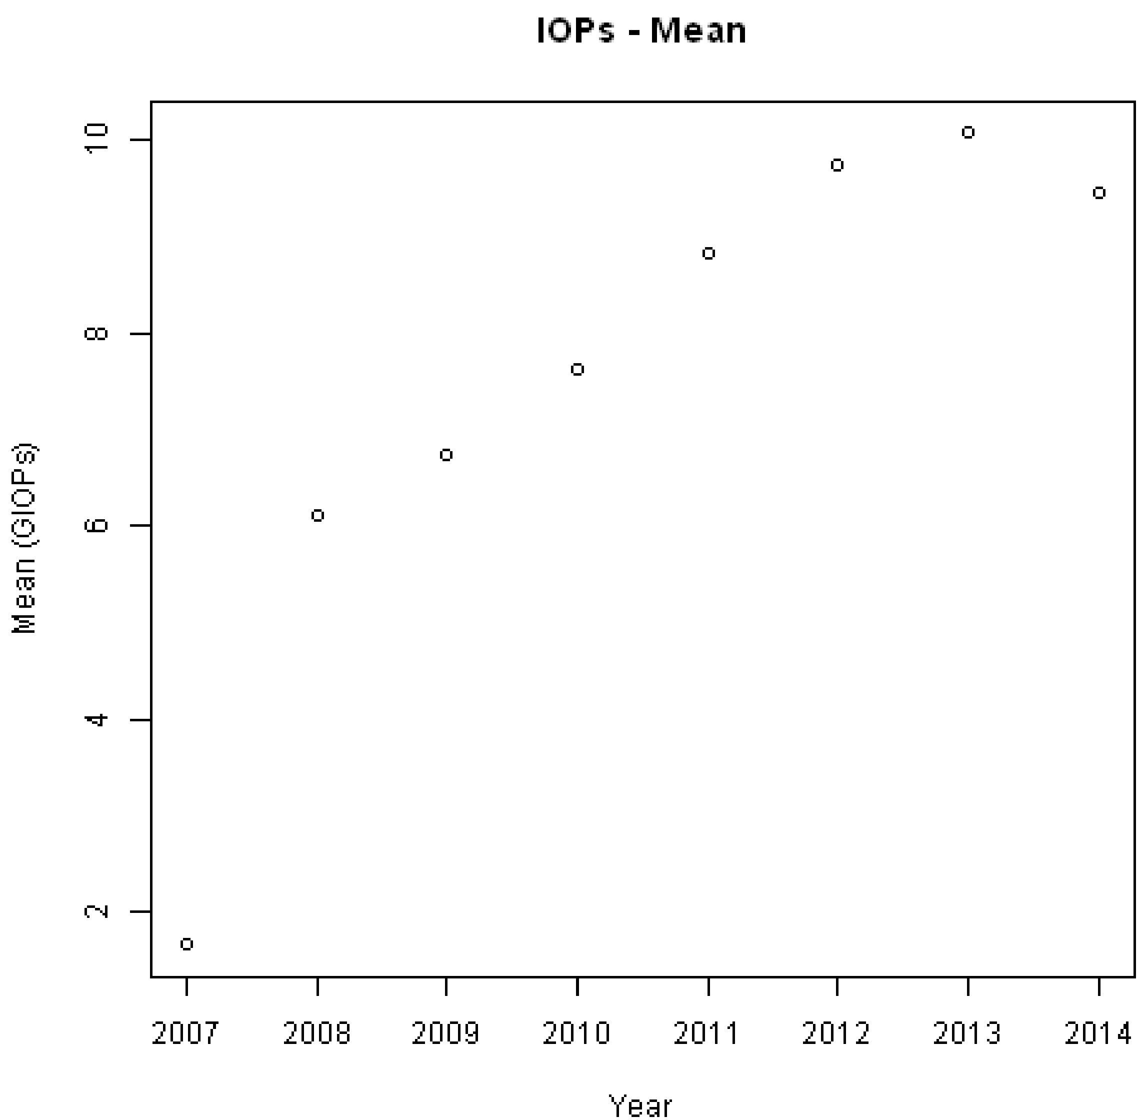

GPUGRID - Performance - **FLOPs**

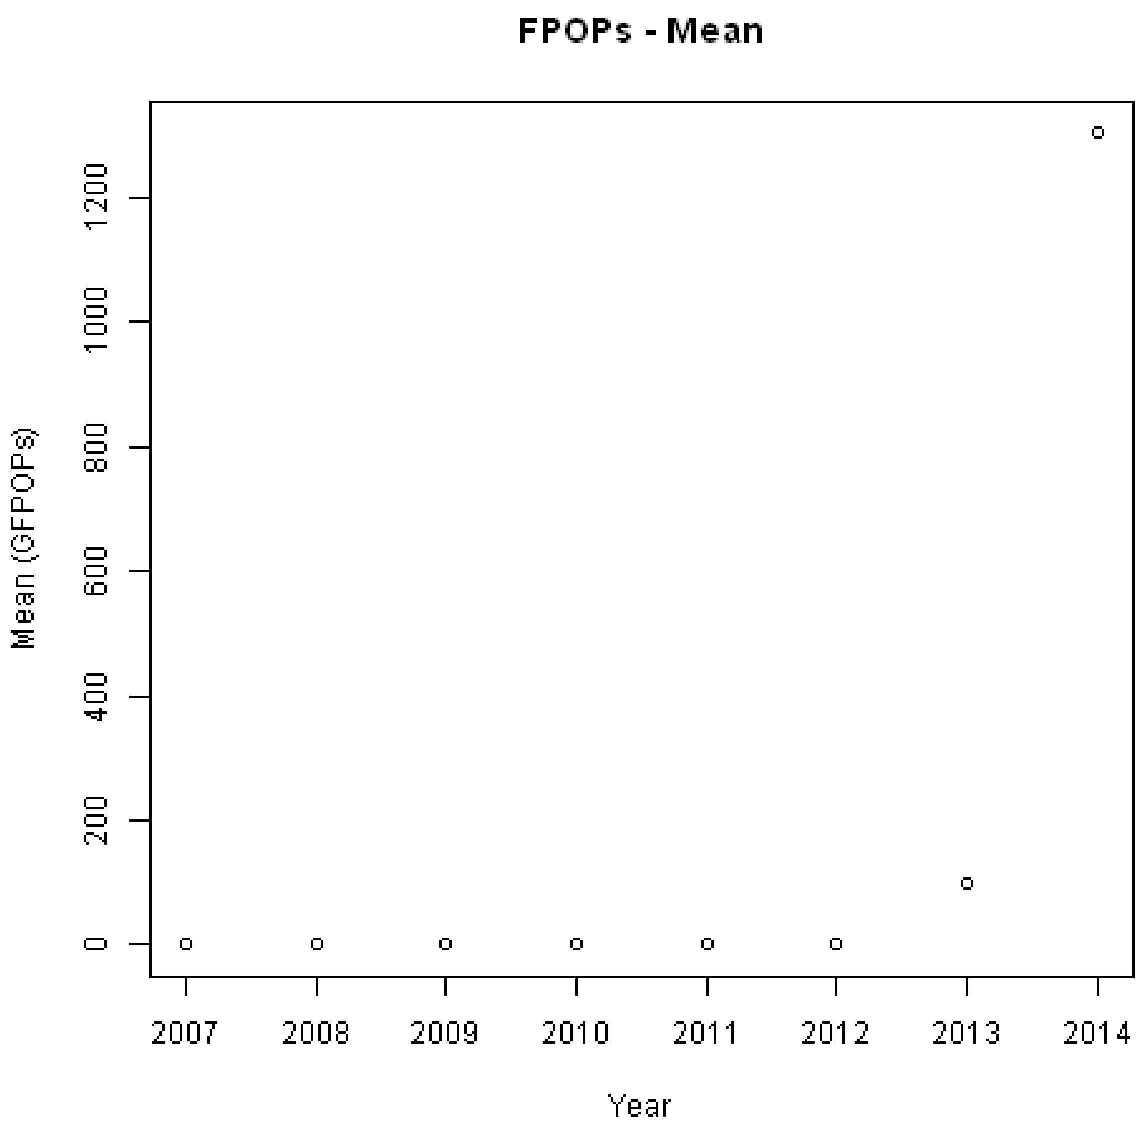

Summary
as to the examples

CPUs/host

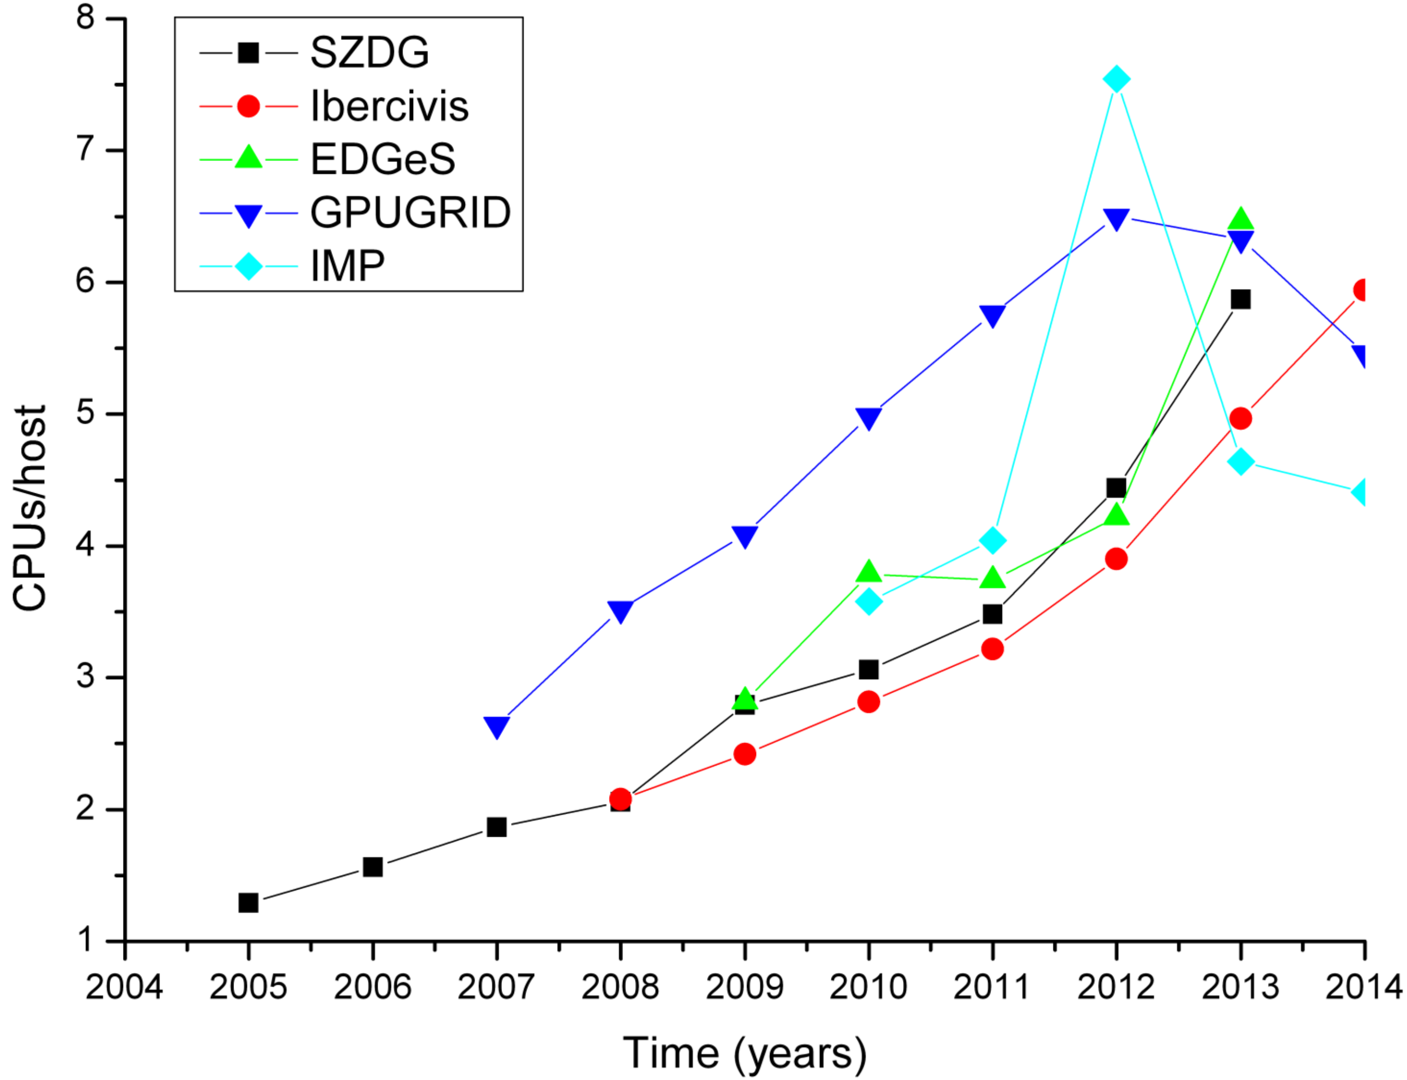
SZDG
Ibercivis
EDGeS
GPUGRID
IMP
CPUs/host
1
2
3
4
5
6
7
8
2004
2005
2006
2007
2008
2009
2010
2011
2012
2013
2014
Time (years)

IOPs

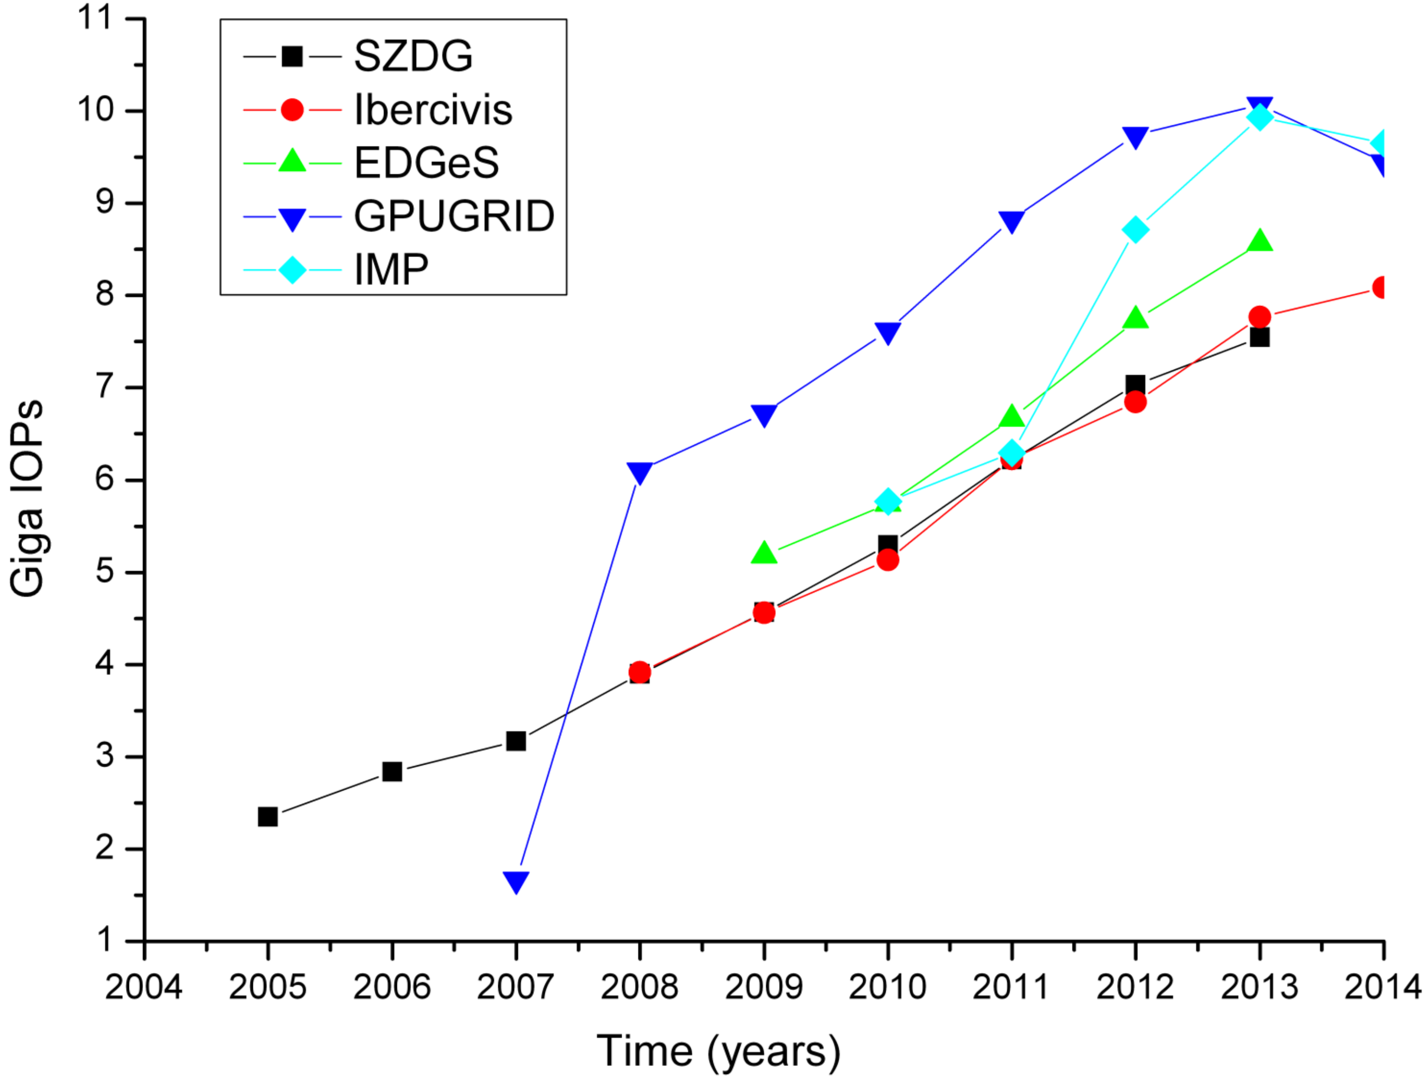

Giga IOPs
Time (years)
SZDG
Ibercivis
EDGeS
GPUGRID
IMP
1
2
3
4
5
6
7
8
9
10
11
2004
2005
2006
2007
2008
2009
2010
2011
2012
2013
2014

FLOPs

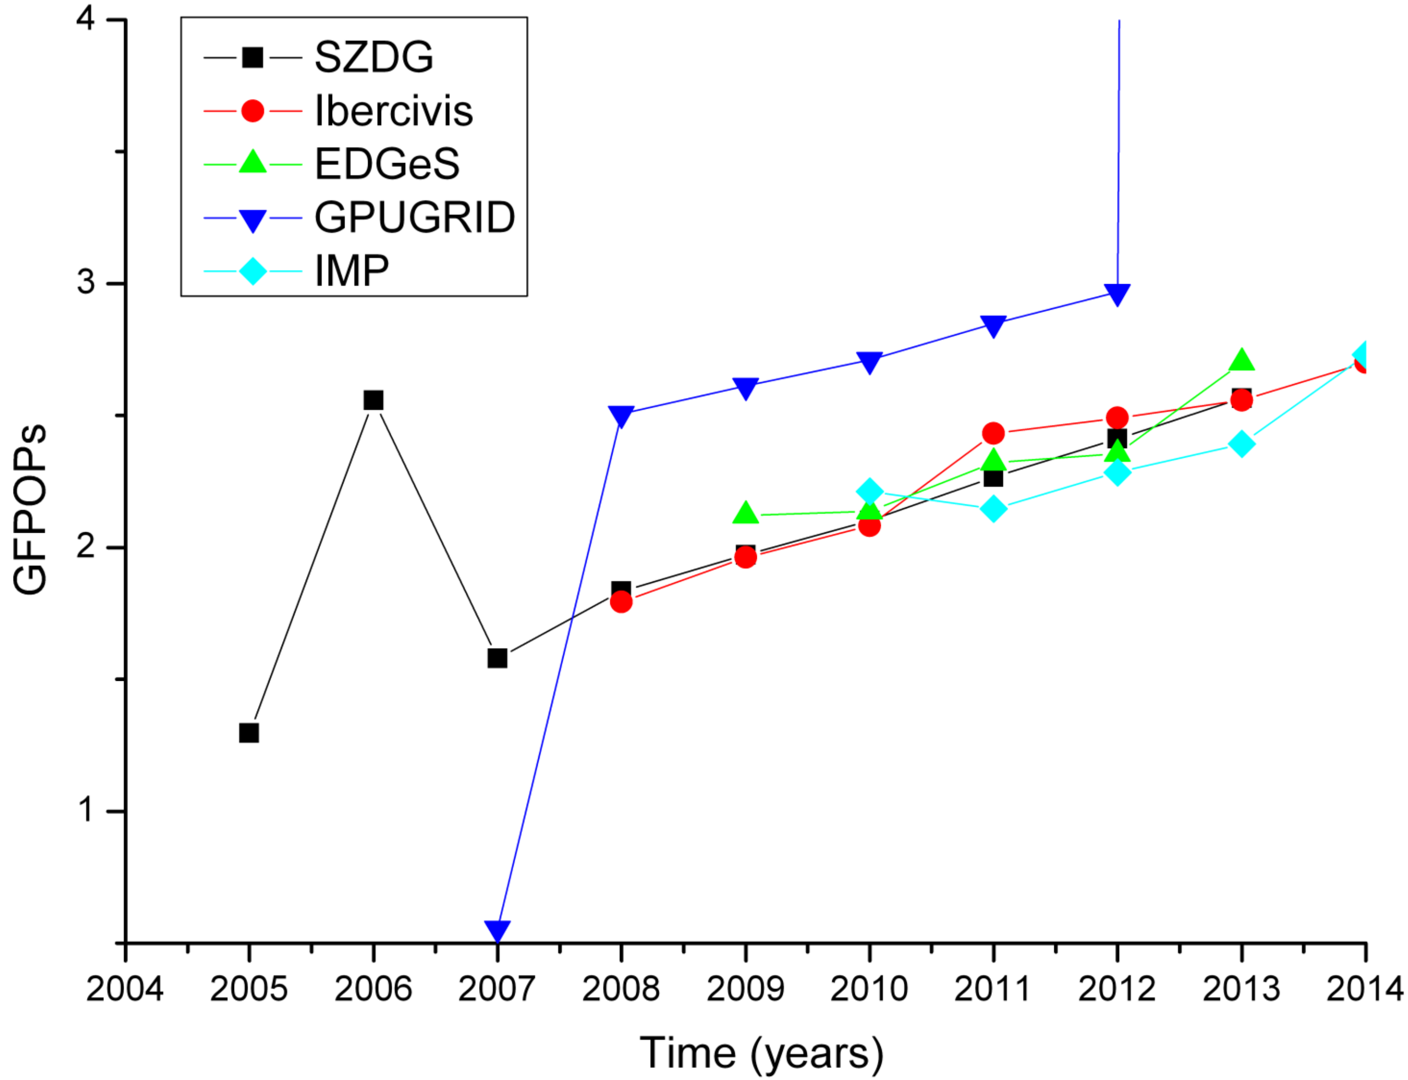

SZDG
Ibercivis
EDGeS
GPUGRID
IMP
GFPOPs
4
3
2
1
2004
2005
2006
2007
2008
2009
2010
2011
2012
2013
2014
Time (years)

Progress of FPOPs and IOPs

FPOPS:

Community	Increase	Time Range	Speedup
SZDG	1,27	8	0,159
EDGeS	0,58	4	0,145
Ibercivis	0,91	6	0,152
GPUGRID	0,46	4	0,115

MEAN: 0,14±0,02 (GFPOPs/year)

IOPS:

Community	Increase	Time Range	Speedup
SZDG	5,2	8	0,65
EDGeS	3.36	4	0,84
Ibercivis	4,17	6	0,695
GPUGRID	3,96	5	0,792

MEAN: 0,74±0,09 (GFPOPs/year)

Summary on Progress of Computing Power in BOINC Projects

- Several BOINC **projects of different kinds** (big-small, long-short, CPU-GPU, geography) **demonstrate the similar progress** of CPUs/host, FPOPs, IOPs, therefore they can be considered as some subsets of the global PC computing community.
- The **longest** (>6 years) and **largest** (>10000 hosts) **BOINC projects are very similar** in these aspects and **can be considered as "statistically representative samples"**
- The **actual increase of global computing power** (in IOPS/FPOPs) **averaged over communities is very slow** (0,74/0,14 per year) in comparison to Moore's law (2x each 2 years)
- This metrics – let's title it like **BOINC Law** (or **Anderson's ☺ Law**) - can be considered as extension of Moore's Law **(?)**

BUT

- **These results are previous and should be confirmed by future ...**

Future Actions – **other (much bigger) communities**

- e.g. ABC (big size >160 MB)

...

then go to

...

- WCG (huge size >2.5 GB)
- SETI (biggest size >4 GB, longest history >11 years)
- LHC (biggest hype, because of Higgs bozon)...

Future Actions – It should be **better (much deeper) analysis**

To take into account:

- actually **working** hosts,
- host subsets **for some applications** (low sense, but it can be interesting for load balancing, planning, scaling, ...),
- **weights** on credits earned (if it has sense?)
- “finer-grain” time series (from **years to months** and weeks?)

Future Actions – **other (much wider) metrics**

- Higher moment (2^{nd} – std, 3^{rd} – skewness, 4^{th} - kurtosis, ...) analysis,
- Bootstrapping analysis,
- Scaling analysis
- ...

Thank you
for your attention!

Under partial support and in cooperation with EU FP7 SCI-BUS project, contract no. RI-283481