Summary of HPC Prospects across ATLAS

- Why ?
- Where ?
- How ?

Eric Lançon

Some concern about LHC computing

CPU consumption above pledges both at T1s and T2s

Sites provide un-pledged resources (thank you!)

Experiments needs are larger than official requests

source : EGI accounting portal

T1 CPU pledge usage [%]

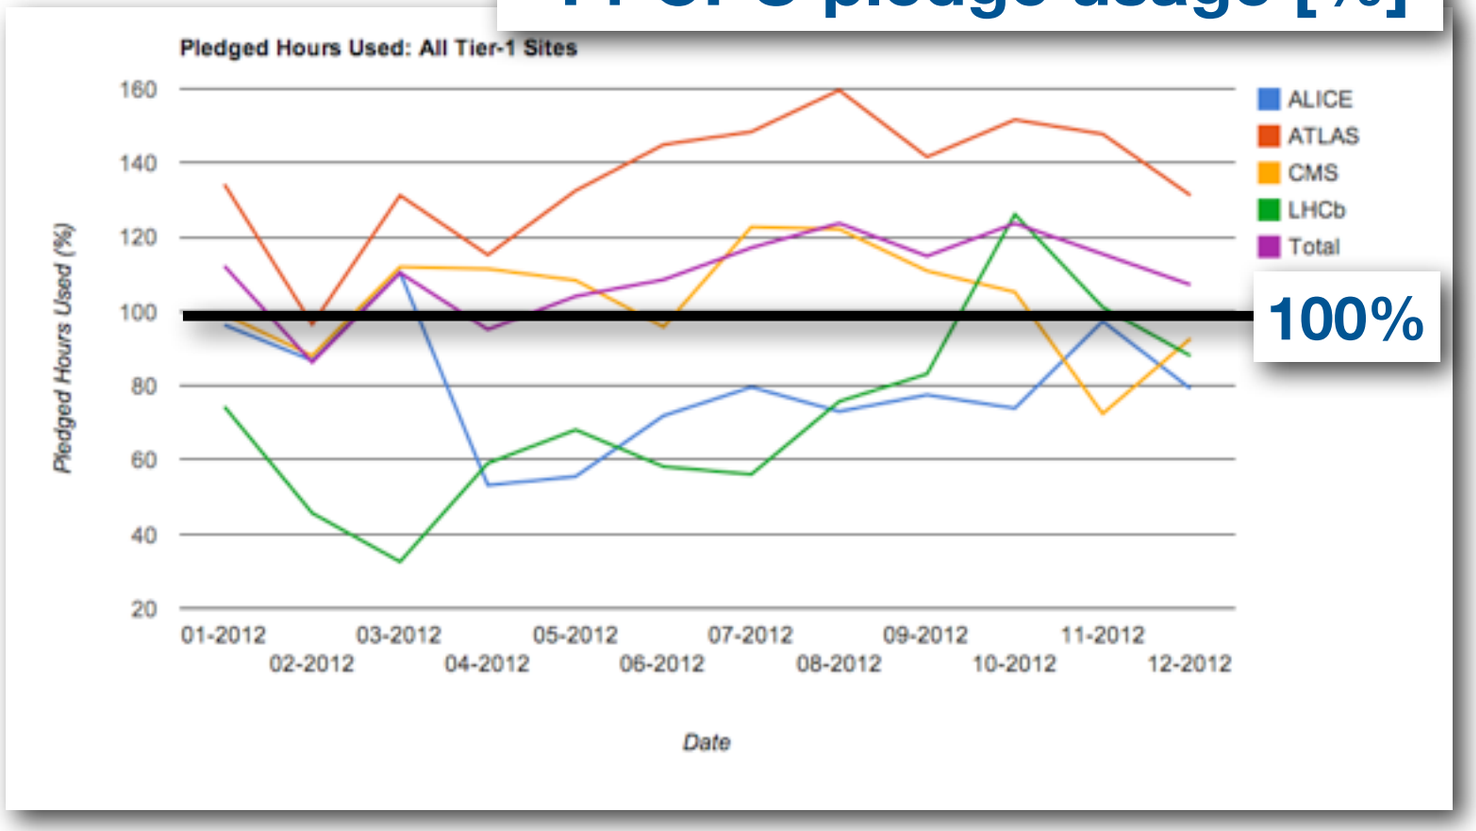

T2 CPU pledge usage [%]

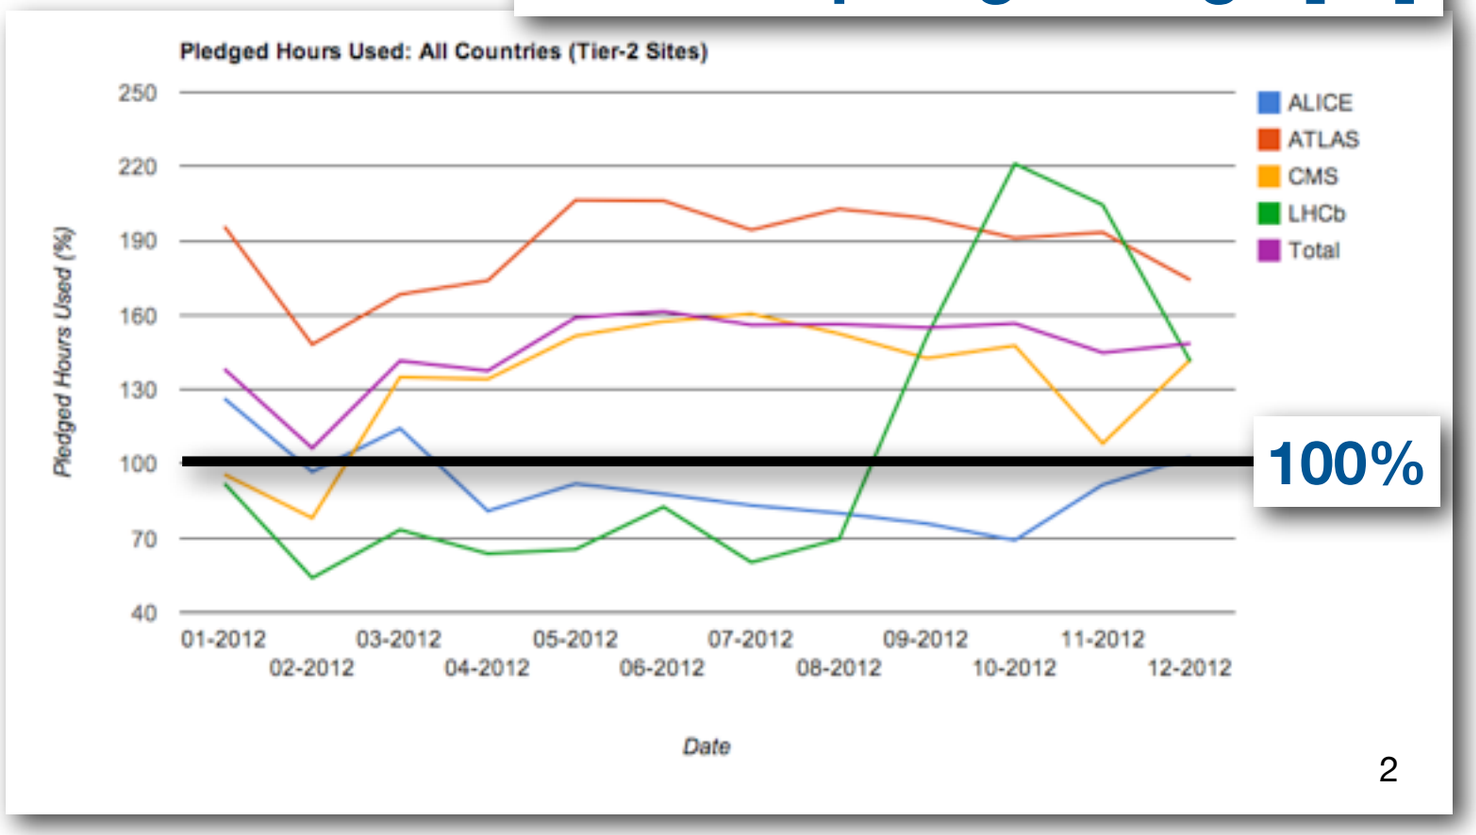

Eric Lançon
irfu
saclay
International Center for Elementary Particle Physics
素粒子物理国際研究センター
ICEPP

HPC for HEP ?

- HPC (High Performance Computing) aka Supercomputers
- High number of processors in close proximity and interconnected
- The opposite of Grid computing…
- Why are they interesting for us ?

Large investments in HPC infrastructures in many countries

From Peta to Exa scales initiatives[1]

Europe (PRACE)

US

Japan

China

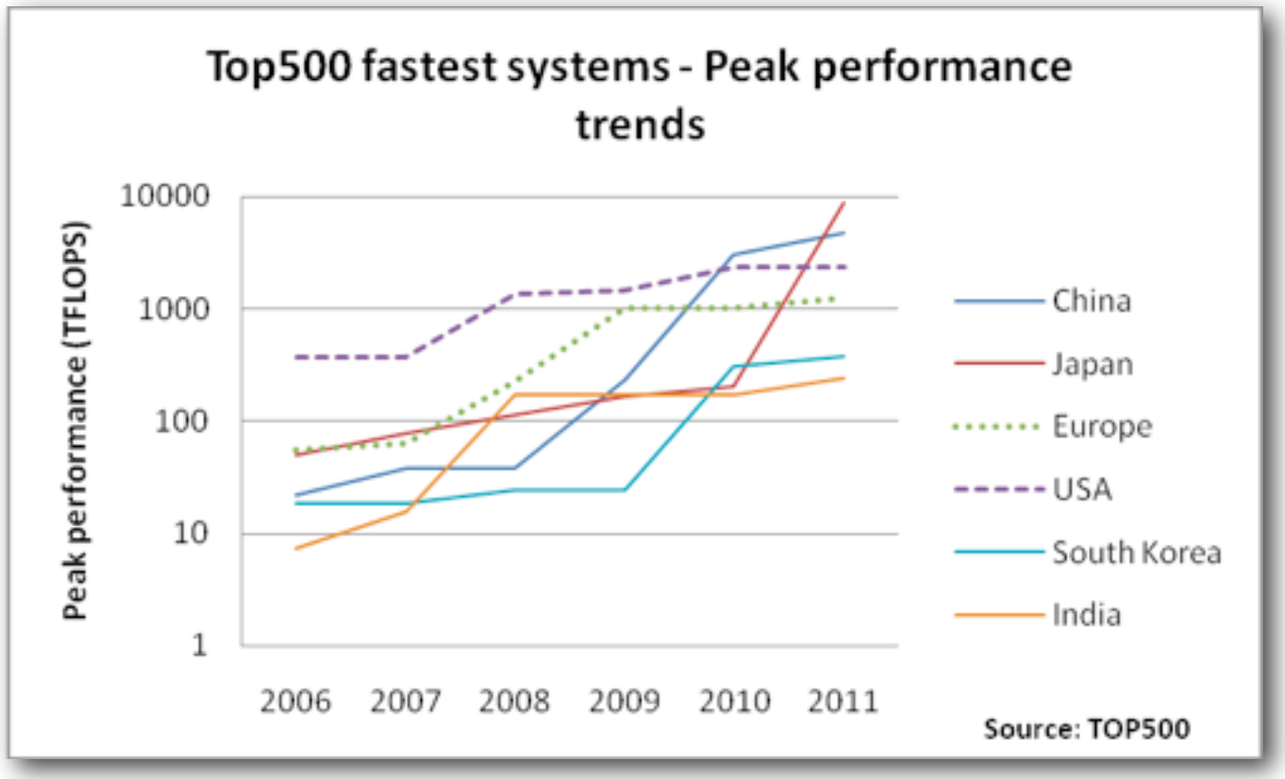

No1 systems (peak performance) in China, Japan, South Korea, India, compared with Europe and the U.S. from 2006 to 2011 (Source TOP500 June lists 2006-2011)

[1] : http://www.eesi-project.eu/pages/menu/eesi-1/publications/investigation-of-hpc-initiatives.php

HPCs have changed a lot in the past 10 years

- Going away from mainframes
- Now : interconnected (InfiniBand,…) clusters of ‘standard’ (to us) processors

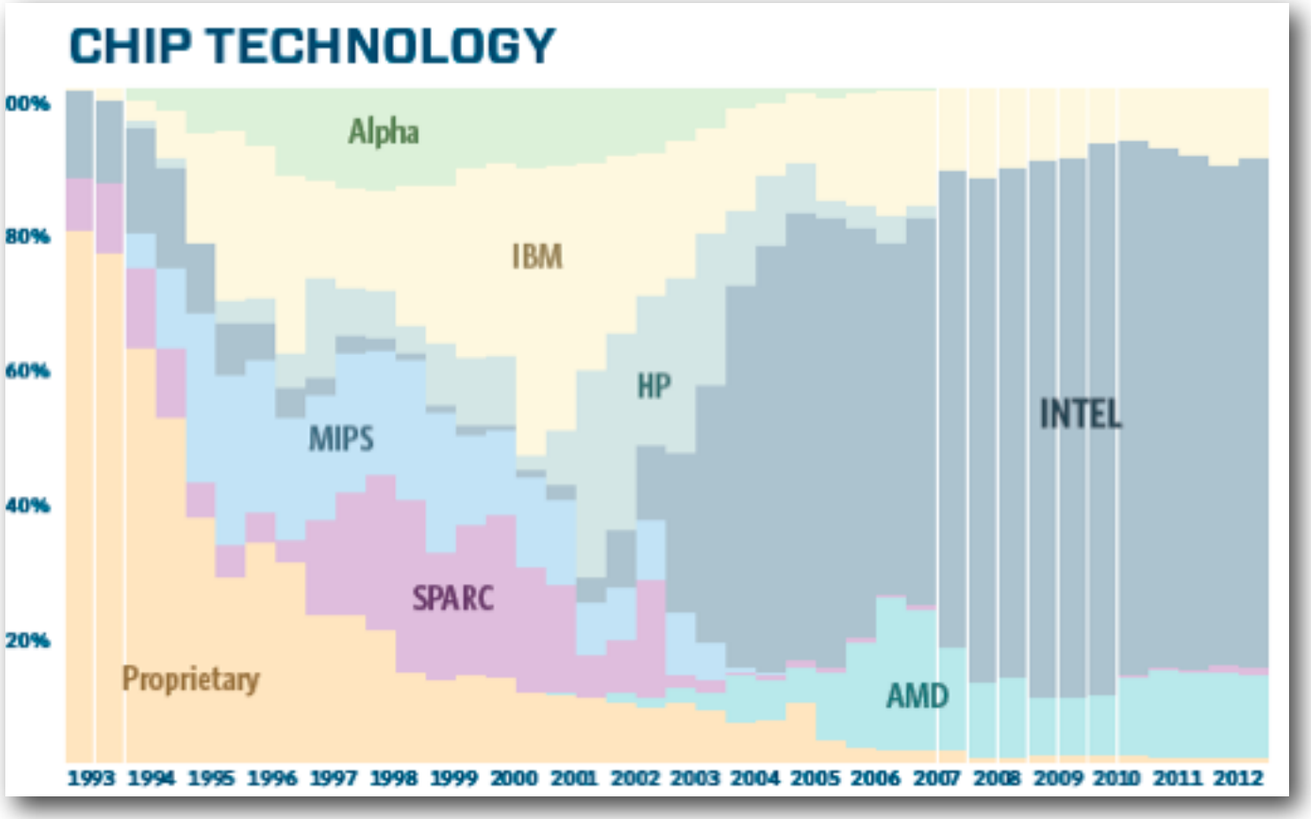

http://www.top500.org/lists/2012/11/

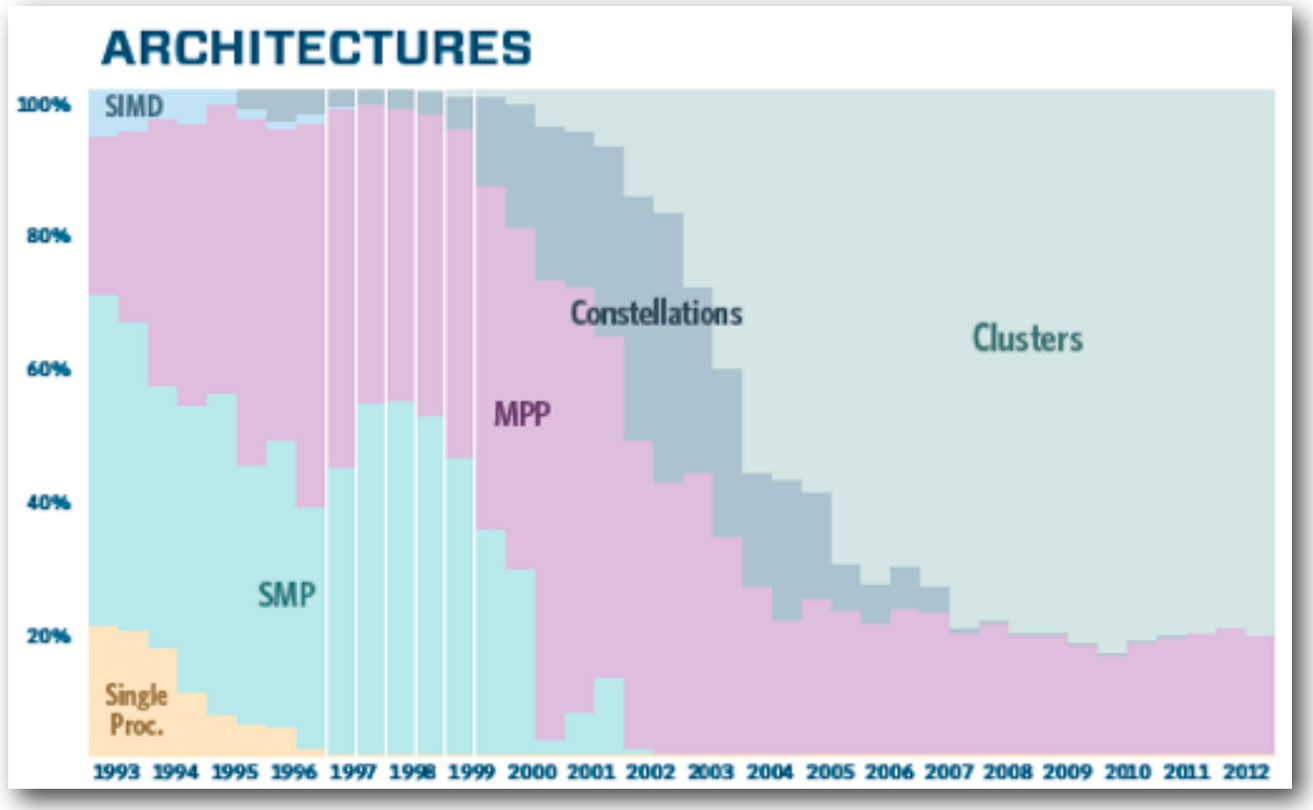

Where those equipments are located

Strong correlation with HEP (ATLAS) geographical locations

Most from energy related institutions (also funding ATLAS computing)

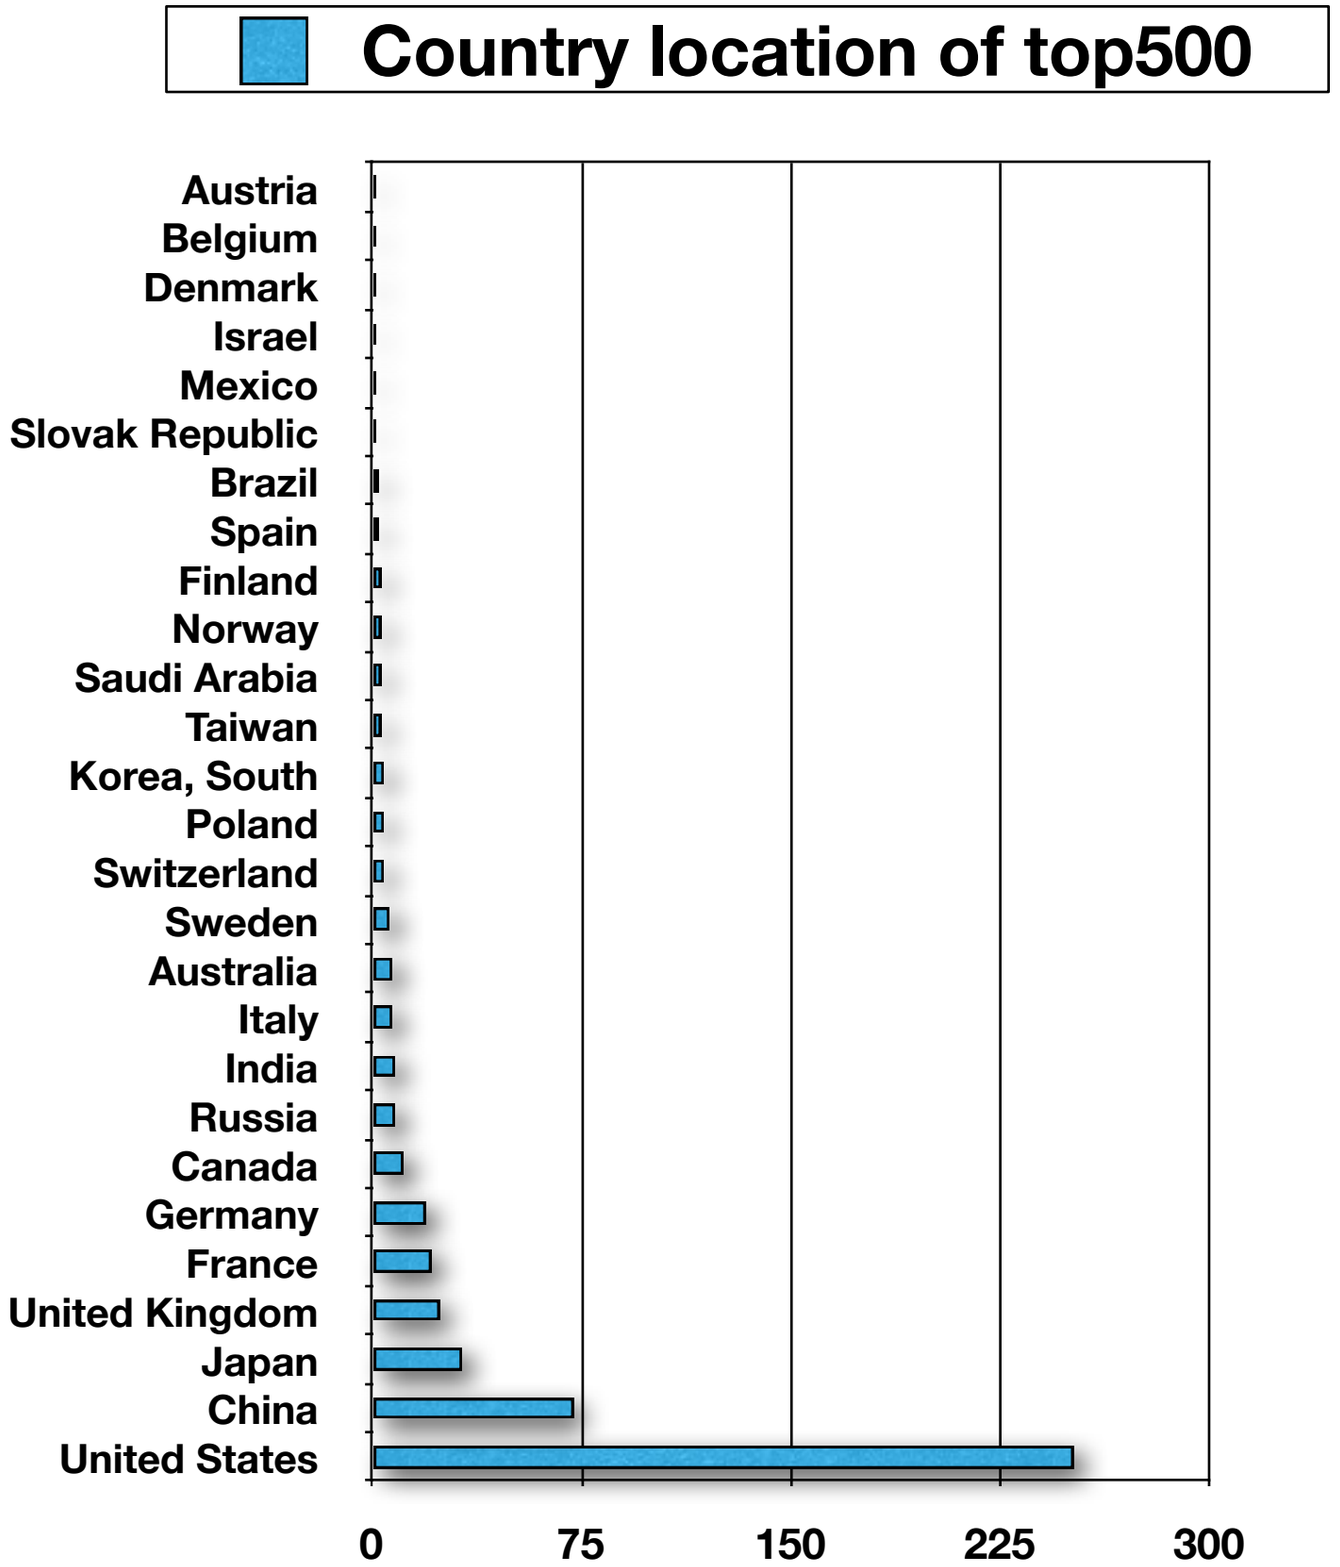

http://www.top500.org/lists/2012/11/

HPC operation

Capability computing oriented (using the maximum computing power to solve a single large problem in the shortest amount of time)

more than

Capacity computing oriented (using efficient cost-effective computing power to solve a small number of somewhat large problems or a large number of small problems)

But situation is evolving

The Majority of TeraGrid/XD Projects Have Modest-Scale Resource Needs

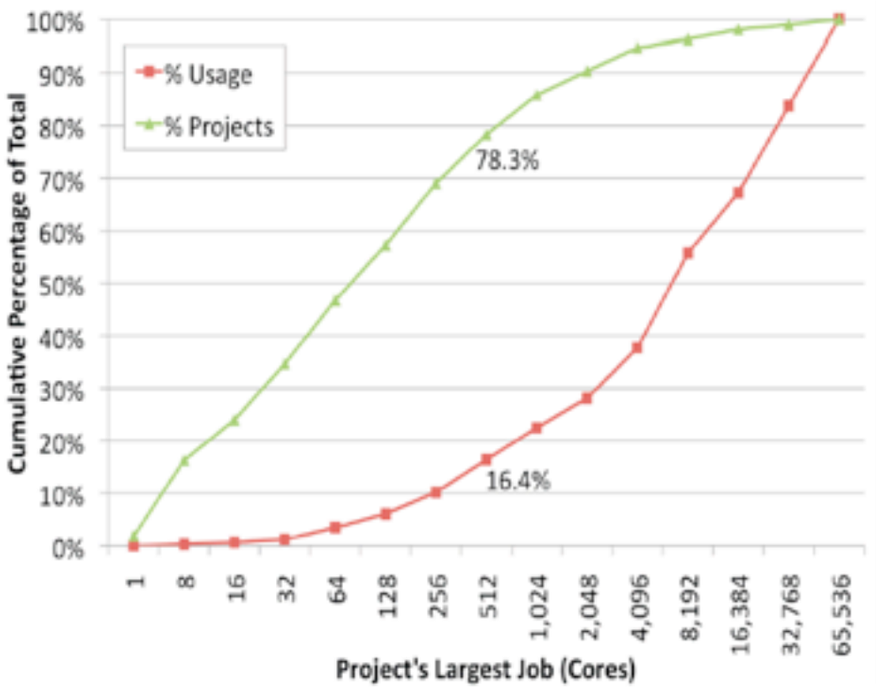

Exceedance distributions of projects and usage as a function of the largest job (core count) run by a project over a full year (FY2009)

- "80/20" @ ~512 cores (FY09)
  - ~80% of projects never run a job larger than this
  - And those projects use <20% of resources
- Only ~1% of projects run jobs as large as 16K cores and those consume >30% of TG resources
- Many projects/users only need modest-scale jobs/resources
- And a modest-size resource can provide the resources for a large number of these projects/users

SDSC SAN DIEGO SUPERCOMPUTER CENTER
at the UNIVERSITY OF CALIFORNIA; SAN DIEGO
UCSD

HEP computation is capacity oriented

What it means for us?

- Most HPC hardware are not unknown to us

- **Large** number of spare CPU cycles are available at HPCs which cannot be used by ‘standard’ HPC applications

- An opportunity for HEP ?

PRACE, the Partnership for Advanced Computing in Europe

SuperMUC a PRACE Tier-0 centre :
155,000 Sandy Bridge cores, 2.8M HS06

WLCG 2013 T0/1/2 pledges ~2.0M HS06

In principle yes, however...

- Not all HPC hardware types easily usable by ATLAS
- Issues with disk space/core and memory/core : ‘nicely’ meet ATLAS software re-engineering effort
- Difficulties to use HPC centers for I/O intensive applications
- Outbound connectivity of HPC centers is an issue
- …

HPC usage & procedures

Application calls (~heavy procedures), selection board

Scheduling of applications

You get what you asked for at most

Little experience in experimental HEP world (Lattice QCD computations, cosmology,...)


PRACE – Partnership for Advanced Computing in Europe

Guide for Applicants to Tier-0 Resources

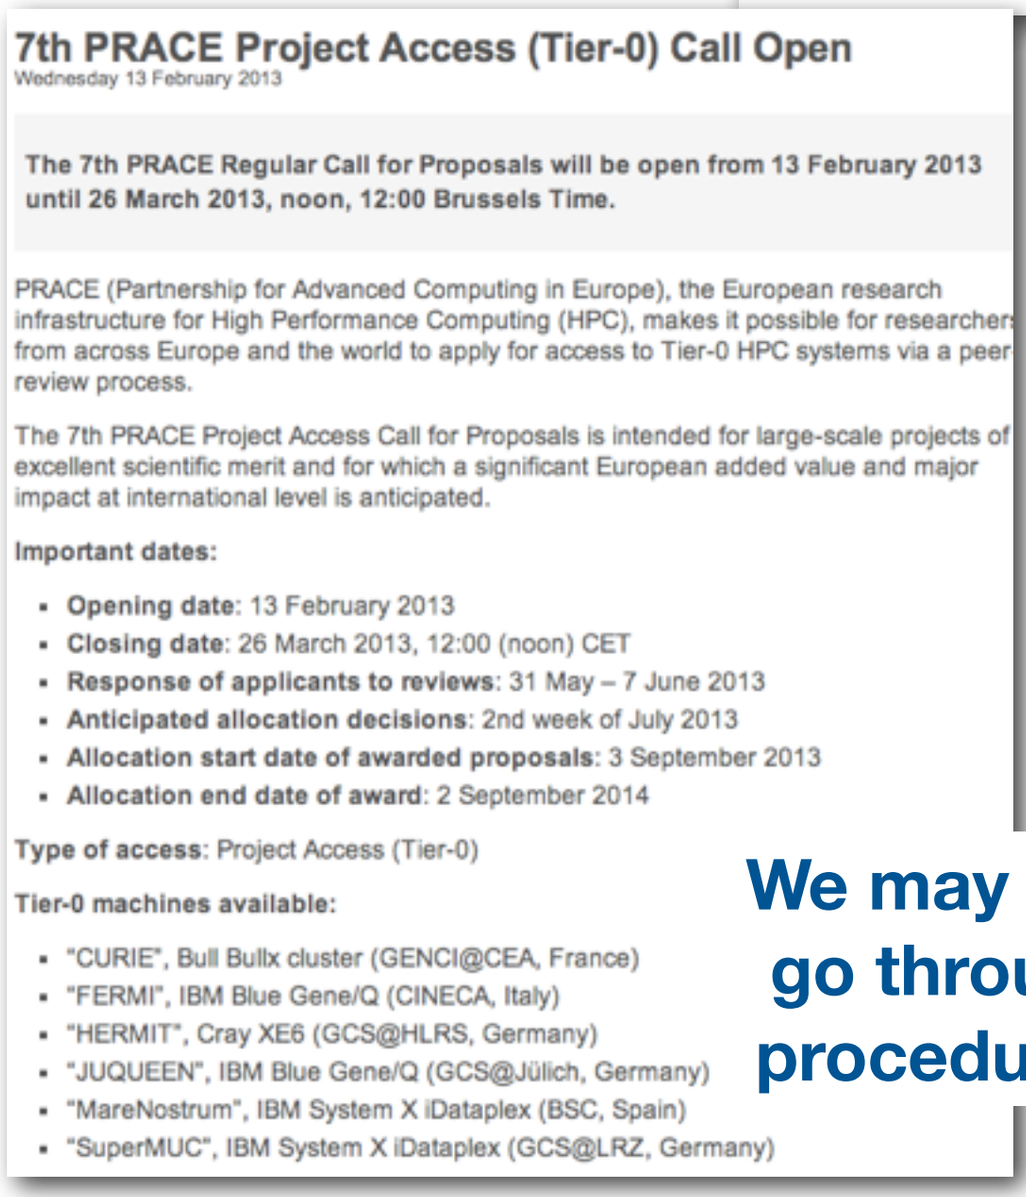
7th PRACE Project Access (Tier-0) Call Open

Wednesday 13 February 2013

The 7th PRACE Regular Call for Proposals will be open from 13 February 2013 until 26 March 2013, noon, 12:00 Brussels Time.

PRACE (Partnership for Advanced Computing in Europe), the European research infrastructure for High Performance Computing (HPC), makes it possible for researchers from across Europe and the world to apply for access to Tier-0 HPC systems via a peer review process.

The 7th PRACE Project Access Call for Proposals is intended for large-scale projects of excellent scientific merit and for which a significant European added value and major impact at international level is anticipated.

Important dates:

- **Opening date**: 13 February 2013
- **Closing date**: 26 March 2013, 12:00 (noon) CET
- **Response of applicants to reviews**: 31 May – 7 June 2013
- **Anticipated allocation decisions**: 2nd week of July 2013
- **Allocation start date of awarded proposals**: 3 September 2013
- **Allocation end date of award**: 2 September 2014

Type of access: Project Access (Tier-0)

Tier-0 machines available:

- "CURIE", Bull Bullx cluster (GENCI@CEA, France)
- "FERMI", IBM Blue Gene/Q (CINECA, Italy)
- "HERMIT", Cray XE6 (GCS@HLRS, Germany)
- "JUQUEEN", IBM Blue Gene/Q (GCS@Jülich, Germany)
- "MareNostrum", IBM System X iDataplex (BSC, Spain)
- "SuperMUC", IBM System X iDataplex (GCS@LRZ, Germany)

We may not want to go through the full procedures to start

ATLAS survey

Many contacts initiated by ATLAS members with national/local HPC institutions

Early investigation phase

Would be nice to have a global ATLAS view of activities in this areas and some prospects

Richard's suggestion

Questionnaire submitted to ICB members

1. *Is there HPC facilities in your country? which kind of hardware?*
2. *Could they be used by ATLAS?*
3. *if yes do you have an estimate of the amount of CPU that can be delivered to ATLAS?*
4. *is HEP being encouraged to use HPC facilities by 1) the HEP funding source and/or 2) the HPC funding source?*
5. *How keen are the operators of HPC facilities to see HEP computing on their machines?*
6. *What level of help might be available to exploit these machines – could this be a source of non-HEP effort helping to re-engineer our code for future computer architectures?*
7. *Is exploitation of HEP facilities expected to endanger traditional HEP computing budgets?*

No details of answers provided today only some general comments

Answers received

1. All have HPC facilities available (or soon e.g. Greece), mostly x86 based, as expected…

2. Possibly available for ATLAS (some already used e.g. NDGF)

Country
Canada
China
France
Germany
Greece
Italy
Norway
Spain
Sweden
Switzerland
USA

HPC availability for ATLAS ?

3. Amount of CPU that can be available for ATLAS ?

- ranges/country from T2 to T1 size and more
- very few countries provided quantitative answers

4. HEP encouraged to use HPC facilities ?

- No except in US
- But situation may evolve/is evolving in various countries

Interactions with HPC

5. HPC operators views about HEP computing

- Different world, different way of working in most of the countries, dialog just started
- Different kind of applications, but HEP can increase HPC centers usage. Evolving way of managing HPC centers
- Most pro-active in US?

6. Support from non-HEP source to re-engineer ATLAS code ?

- Too early to get firm answer in most of the countries (likely very little)
- Some support for code porting might be available in US

Effect on HEP computing budget?

7. Is exploitation of HEP facilities expected to endanger traditional HEP computing budgets?

- *Definitively yes in most of the countries*
- *Unitarity issues*

- *However :*
  - *The net result for HEP could be positive for CPU balance*
  - *Given momentum of HPC world and their interest in big data volume processing*
  - *and under-published computing needs for HEP (CMS processed parked-data at SDSC[1])*

[1]http://www.sdsc.edu/News%20Items/PR040413_lhc.html

素粒子物理国際研究センター
International Center for Elementary Particle Physics

ICEPP

To summarize :

危机与机遇

危機と機会

crisis & opportunities

irfu

cea

saclay

Eric Lançon

HPC contributes to Higgs boson search

Researchers at CERN drive pioneering research into the origins of the universe with a powerful HPC solution analysing one billion proton collisions per second

"As a result of the number of cores in the Dell C-class servers and the architecture of the processors, we have 50 per cent more capacity to analyse the collisions."

Marc Dobson, Applied Physicist at CMS (System Administrator), CERN

Customer profile

CERN

supplier

Organization	CERN
Industry	Science
Country	Switzerland
Employees	2,400 employees; 11,000 external researchers
Website	public.web.cern.ch/public

Business need

Scientists at CERN needed to increase computing power to process data from the CMS experiment at the world-famous Large Hadron Collider (LHC) as they searched for the Higgs boson.

Solution

http://i.dell.com/sites/doccontent/corporate/case-studies/en/Documents/2013-cern-10011805.pdf

Backup