Fisica Teorica e Calcolo HPC nell'INFN

Leonardo Cosmai
INFN Bari

"Large Scale Computing at INFN" , 13 Febbraio 2017

Dipartimento di Fisica Sapienza Università di Roma

Roma Tre, 9 Settembre 2014

Large scale computing at INFN?

Tuesday, 9 September 2014 from **10:00** to **16:30** (Europe/Rome)
at **Universita' Roma III (Aula 57)**
via della Vasca Navale 84 - Roma

The case for a renewed support of computational theoretical physics at INFN

Roma Sapienza, 13 Febbraio 2017

Large scale computing at INFN

13 February 2017 *Dipartimento di Fisica Sapienza Univ. di Roma*
Europe/Rome timezone

Gruppo Teorico e calcolo HPC: una lunga tradizione

Le risorse di calcolo:

- ~1985-1995: Cray 1M, Cray X-MP/12-48, Cray Y-MP, Cray T3D/64 (CINECA)

- ~1985-2010: APE, APE100, APEmille, apeNEXT

- ~2010-…: cluster locali, cluster CSN4 (Zephiro)

- ~2012: nuova convenzione CINECA: BlueGene/Q (FERMI), IBM cluster (GALILEO), Lenovo cluster (MARCONI)

- …

apeNEXT (fino a ~ 2010)

apeNEXT:
Computational Challenges
and First Physics Results

Arcetri, Firenze, February 8 -10, 2007

The Galileo Galilei Institute for Theoretical Physics

Talks

Past

10-02-2007	10:30	**Norman Christ**	Chiral Symmetry, Petaflops and Solving Low Energy QCD
10-02-2007	09:30	**Ulrich H. E. Hansmann**	In Silico Folding of Small Proteins
09-02-2007	16:35	**Alessandro Papa**	Analytic continuation from an imaginary chemical potential
09-02-2007	16:10	**Leonardo Cosmai**	QCD in external fields
09-02-2007	14:50	**Massimo D'Elia**	Vacuum structure and deconfinement
09-02-2007	14:25	**Edwin Laermann**	Bulk QCD thermodynamics at small quark masses
09-02-2007	14:00	**Olaf Kaczmarek**	Heavy quark bindings at high temperature
09-02-2007	11:45	**Maria Paola Lombardo**	Topics in QCD at Finite T and/or mu
09-02-2007	11:00	**Federico Toschi**	Turbulent channel flow on apeNEXT
09-02-2007	10:00	**Francesco Di Renzo**	NSPT@apeNEXT
09-02-2007	09:15	**Hubert Simma**	Fundamental parameters of QCD from the lattice
08-02-2007	17:45	**Nazario Tantalo**	B ---> Dlnu form factors through the Step Scaling Method
08-02-2007	16:45	**Frithjof Karsch**	QCD at finite temperature and density
08-02-2007	15:45	**Gerrit Schierholz**	QCD simulations at realistic quark masses: probing the chiral limit
08-02-2007	15:15	**Dirk Pleiter**	Nucleon structure from lattice QCD
08-02-2007	14:30	**Karl Jansen**	Twisted mass fermions
08-02-2007	12:30	**Giancarlo Rossi**	From theoretical models to simulations: a computational challenge
08-02-2007	12:05	**Vittorio Lubicz**	Lattice QCD and flavor physics (II): Impact of lattice on the Unitarity Triangle Analysis
08-02-2007	11:40	**Silvano Simula**	Lattice QCD and flavor physics (I): Determination of Vus
08-02-2007	10:30	**Olivier Pene**	Physics with the French apeNEXT
08-02-2007	10:00	**Nicola Cabibbo**	The APE experience

http://www.ggi.fi.infn.it//index.php?page=events.inc&id=10

I cluster (dal 2007)

- "Test FARM per HPC" (24 biprocessori Xeon @3.2GHz, Infiniband)
  —> CNAF (~2006) ~100 GFlops di picco

- I cluster locali (Bari (BA-CS-GE-LE-NA-TO), Catania (CT-PD-LNS-RM1), MilanoBicocca (MIB-PR-RM2-TO), Pisa(PI-MI-RM1)) (~ 2007-2008)
  ~ 500 cores di calcolo ~ 4 Mcorehours

- cluster CSN4 "Zefiro" (~2010) O(2000) cores 20 TFlops
  ~16 Mcorehours

Il progetto Premiale SUMA e la nuova convenzione INFN-CINECA (~2012)

- SUMA 1900 Keuro, BA / FE / MIB / PI / PR / RM1 / RM2 / RM3

Updated proposal and workplan of the SUMA project

The INFN-SUMA team: R. Alfieri, R. Ammendola, S. Arezzini, D. Badoni, M. Brambilla, R. Benzi, L. Biferale, A. Carboni, A. Ciampa, L. Cosmai, R. Covati, G. De Divitiis M. D' Elia, R. De Pietri, C. Destri, F. Di Renzo, A. Feo, R. Frezzotti P. Giannozzi, L. Giusti, G. La Penna, V. Lubicz, A. Lonardo, E. Mazzoni, V. Minicozzi, S. Morante, P. S. Paolucci, A. Papa M. Papinutto, M. Pepe, R. Petronzio M. Pivanti, S. Piras, F. Rapuano, D. Rossetti, G. C. Rossi, A. Salamon, G. Salina, M. Sbragaglia, S. F. Schifano, S. Simula, N. Tantalo C. Tarantino, R. Tripiccione, E. Vicari, P. Vicini, A. Vladikas

January 14, 2013

Project meetings and workshops

Meetings

April 1-2, 2014 (Ferrara): **Agenda**

November 8, 2013 (Roma, La Sapienza): **Agenda**

July 26, 2013 (Roma, La Sapienza): **Agenda**

December 11, 2012 (CINECA) : **Agenda**

Workshops

February 11-13, 2015 (Trento) : **Agenda**

~ 22 annualità postdoc
35 journal papers
38 conference papers

cluster Zefiro

GALILEO

Il progetto Premiale SUMA e la nuova convenzione CINECA (~2012)

- GALILEO

(~ Feb 2015 - Oct 2017 (?))

Model: IBM NeXtScale

Architecture: Linux Infiniband Cluster

Nodes: 516
Processors: 2 8-cores Intel Haswell 2.40 GHz per node
Cores: 16 cores/node, 8256 cores in total
Accelerators: 2 Intel Phi 7120p per node on 384 nodes (768 in total); 2 NVIDIA K80 per node on 40 nodes (80 in total, 20 available for scientific research)
RAM: 128 GB/node, 8 GB/core
Internal Network: Infiniband with 4x QDR switches
Disk Space: 2.000 TB of local scratch
Peak Performance: 1.000 TFlop/s (to be defined)

~ 70 Mcorehours / anno

(INFN: 15 Mcorehours/anno)

- FERMI

(~ Sep 2012 - Jul 2016)

Architecture: IBM BlueGene/Q
Model: 10 racks
Processor Type: IBM PowerA2, 1.6 GHz
Computing Cores: 163840
Computing Nodes: 10240, 16 core each
RAM: 16 GB/node, 1GB/core
Internal Network: custom with 11 links -> 5D Torus
Disk Space: 2.6 PB of scratch space
Peak Performance: 2PFlop/s

~ 1300 Mcorehours / anno

(INFN: ~500 Mcorehours in 4 anni

Le risorse di calcolo negli ultimi ~20 anni (*)

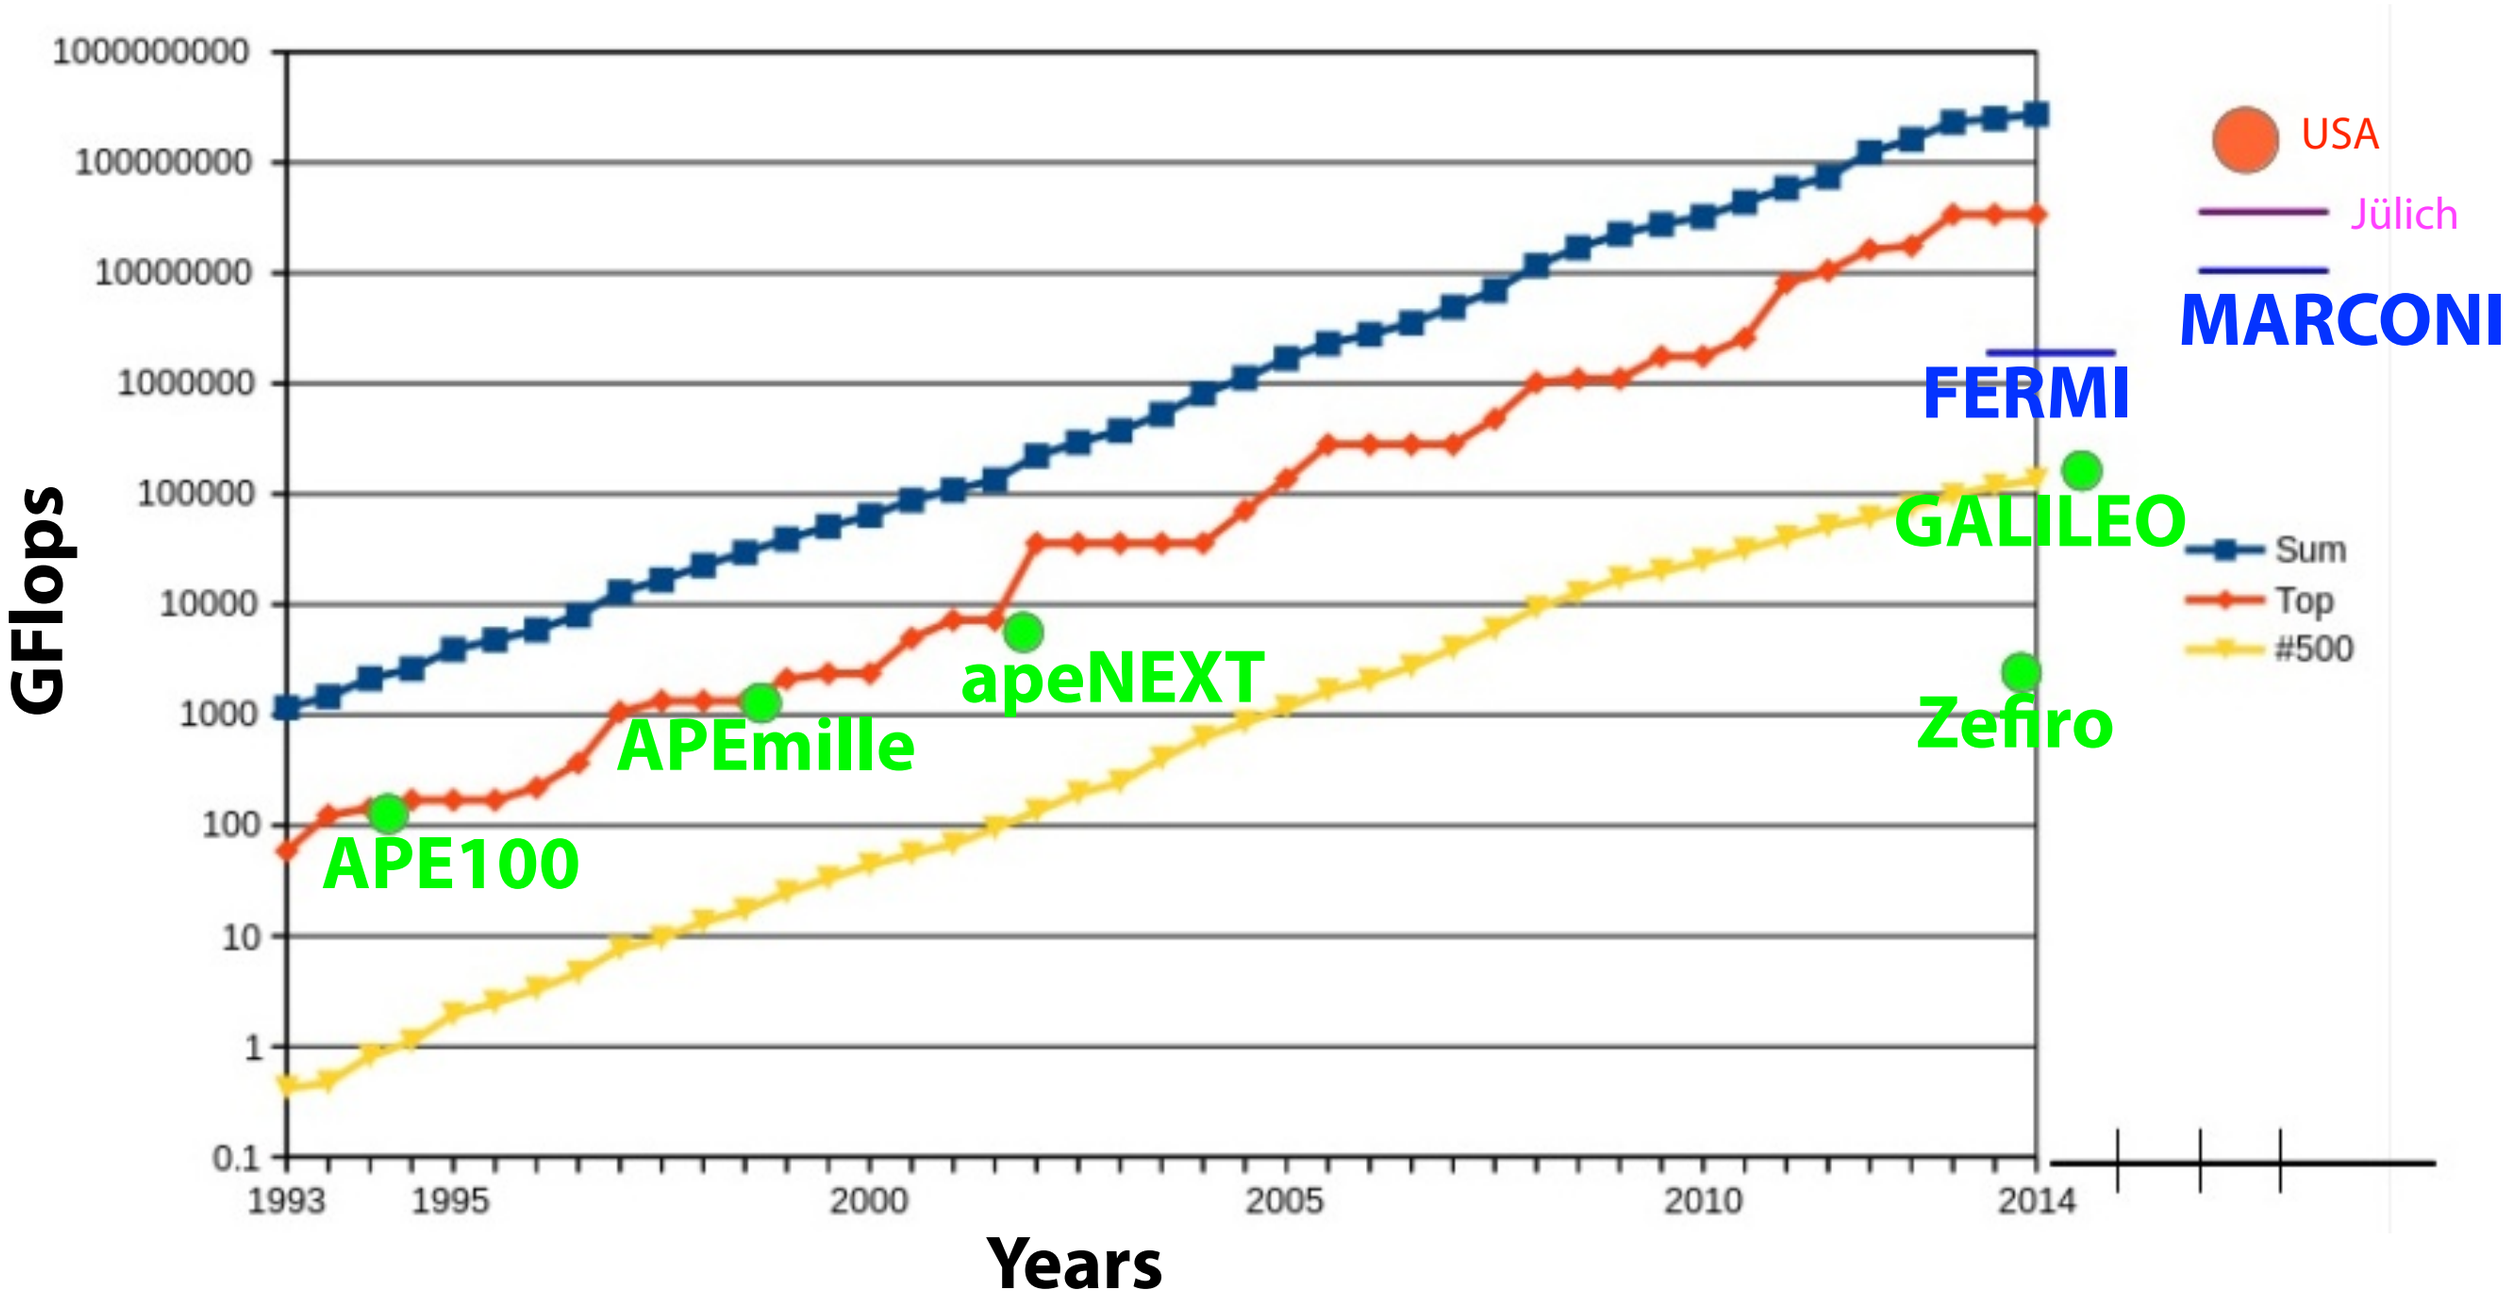

(*) Fig.1 in "The case for a renewed support of computational theoretical physics at INFN"

TOP 500 (Nov 1996)

Countries Performance Share

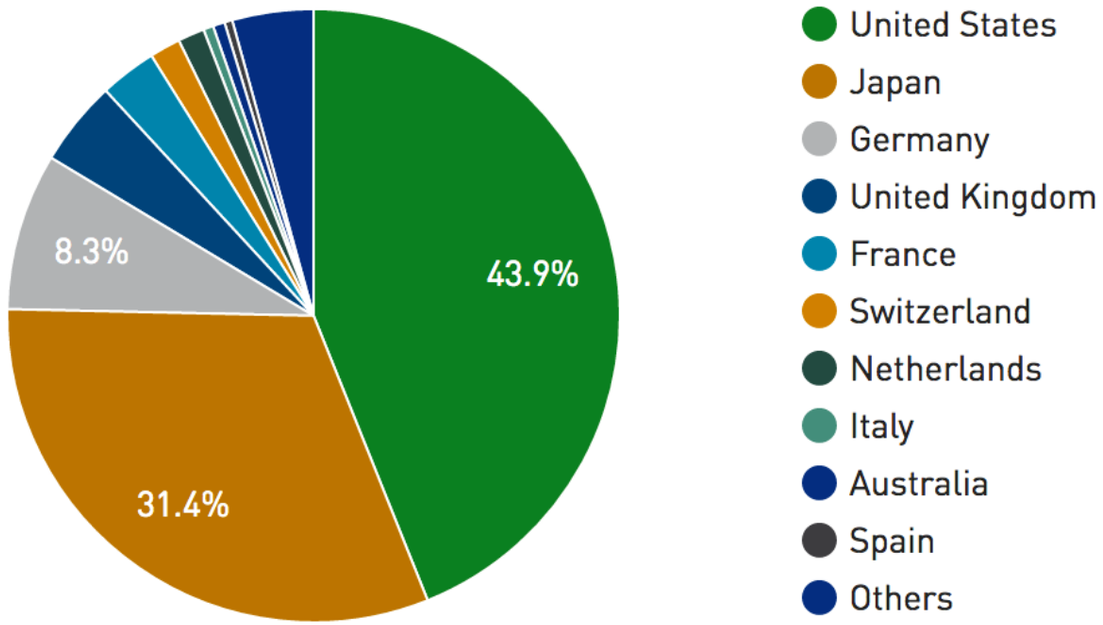

Countries	Count	System Share (%)	Rmax (GFlops)	Rpeak (GFlops)	Cores
United States	267	53.4	3,506	5,281	41,298
Japan	80	16	2,507	3,250	8,534
Germany	51	10.2	660	912	2,585
United Kingdom	18	3.6	360	472	1,488
France	17	3.4	239	357	1,402
Switzerland	9	1.8	130	182	1,015
Netherlands	9	1.8	112	144	363
Italy	6	1.2	42	59	302
China	1	0.2	7	9	32

6 entries found.

Rank	Site	System	Cores	Rmax (GFlop/s)	Rpeak (GFlop/s)	Power (kW)
155	CINECA (/site/47495) Italy	T3D MC128-8 (/system/170950) Cray Inc.	128	12.8	19.2	
299	CINECA (/site/47495) Italy	SP2/32 (/system/169570) IBM	32	6.6	8.5	
307	Centro Di Calcolo Interuniversiatrio Dell Italia (/site/47538) Italy	T3D MCA64-8 (/system/170995) Cray Inc.	64	6.4	9.6	
327	CRS4 (/site/47506) Italy	SP2/30 (/system/169553) IBM	30	6.2	8.0	
407	CILEA (/site/47494) Italy	SPP1600/XA-32 (/system/170723) HPE	32	5.5	7.7	
466	Centro Italiano Ricerche Aerospaziali (CIRA) (/site/47540) Italy	POWER CHALLENGE 10000 (/system/168538) HPE/SGI	16	4.9	6.2	

TOP 500 (Nov 2016)

Countries Performance Share

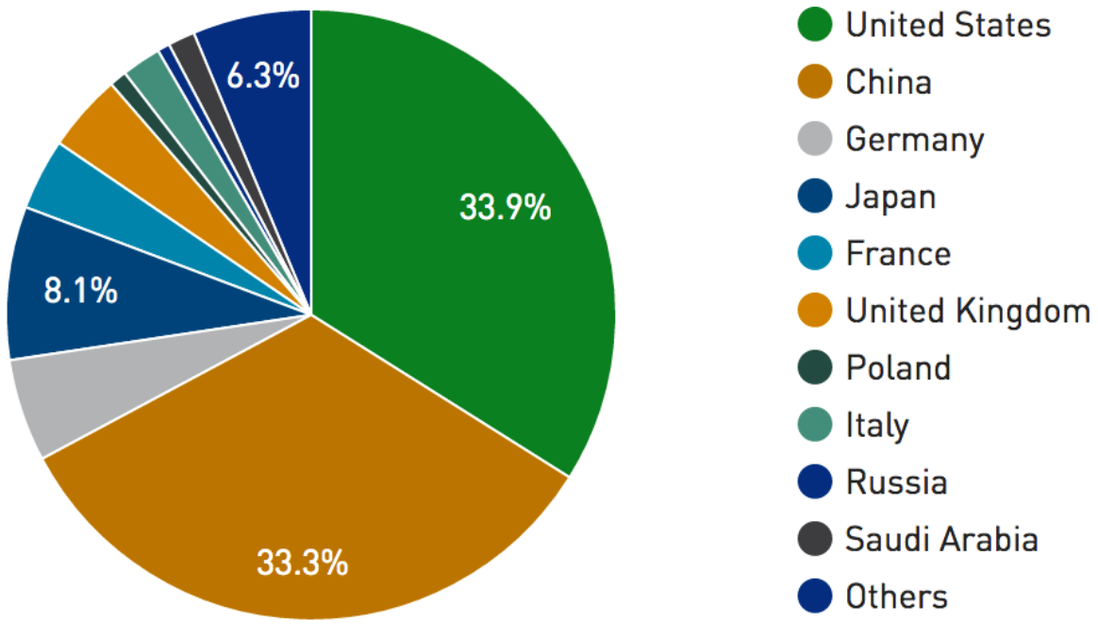

Countries	Count	System Share (%)	Rmax (GFlops)	Rpeak (GFlops)	Cores
United States	171	34.2	228,032,809	327,303,955	11,660,816
China	171	34.2	223,571,136	394,013,392	21,546,512
Germany	31	6.2	36,501,435	45,628,388	1,600,240
Japan	27	5.4	54,486,820	77,371,577	3,946,560
France	20	4	25,398,803	31,727,765	1,158,428
United Kingdom	13	2.6	27,602,596	31,682,369	1,148,968
Poland	7	1.4	6,162,214	8,157,370	208,284
Italy	6	1.2	14,062,113	21,140,514	606,312
Russia	5	1	4,411,812	6,515,928	181,070
Saudi Arabia	5	1	9,577,664	12,798,147	376,336
India	5	1	3,092,368	4,456,051	133,172
Sweden	4	0.8	3,378,143	4,291,609	123,992
Korea, South	4	0.8	5,679,725	7,597,851	169,696
Switzerland	4	0.8	12,273,082	18,811,648	328,608

6 entries found.

Rank	Site	System	Cores	Rmax (TFlop/s)	Rpeak (TFlop/s)	Power (kW)
12	CINECA (/site /47495) Italy	**Marconi Intel Xeon Phi** - CINECA Cluster, Intel Xeon Phi 7250 68C 1.4GHz, Intel Omni-Path (/system /178937) Lenovo	241,808	6,223.0	10,833.0	
29	Exploration & Production - Eni S.p.A. (/site/50489) Italy	**HPC2** - iDataPlex DX360M4, Intel Xeon E5-2680v2 10C 2.8GHz, Infiniband FDR, NVIDIA K20x (/system/178425) IBM	72,000	3,188.0	4,605.0	1,227
56	CINECA (/site /47495) Italy	**Fermi** - BlueGene/Q, Power BQC 16C 1.60GHz, Custom (/system /177720) IBM	163,840	1,788.9	2,097.2	821.9
57	CINECA (/site /47495) Italy	**Marconi Intel Xeon** - Lenovo NeXtScale nx360M5, Xeon E5-2697v4 18C 2.3GHz, Omni-Path (/system/178755) Lenovo	54,432	1,723.9	2,003.1	1,360.8
212	CINECA (/site /47495) Italy	**GALILEO** - IBM NeXtScale nx360M4, Xeon E5-2630v3 8C 2.4GHz, Infiniband QDR, Intel Xeon Phi 7120P (/system/178549) IBM/Lenovo	50,232	684.3	1,103.1	2,825.6
371	Exploration & Production - Eni S.p.A. (/site/50489) Italy	**HPCC1** - iDataPlex DX360M4, Xeon E5-2670 8C 2.600GHz, Infiniband FDR14 (/system/178200) IBM	24,000	454.1	499.2	498.5

Projected Performance Development

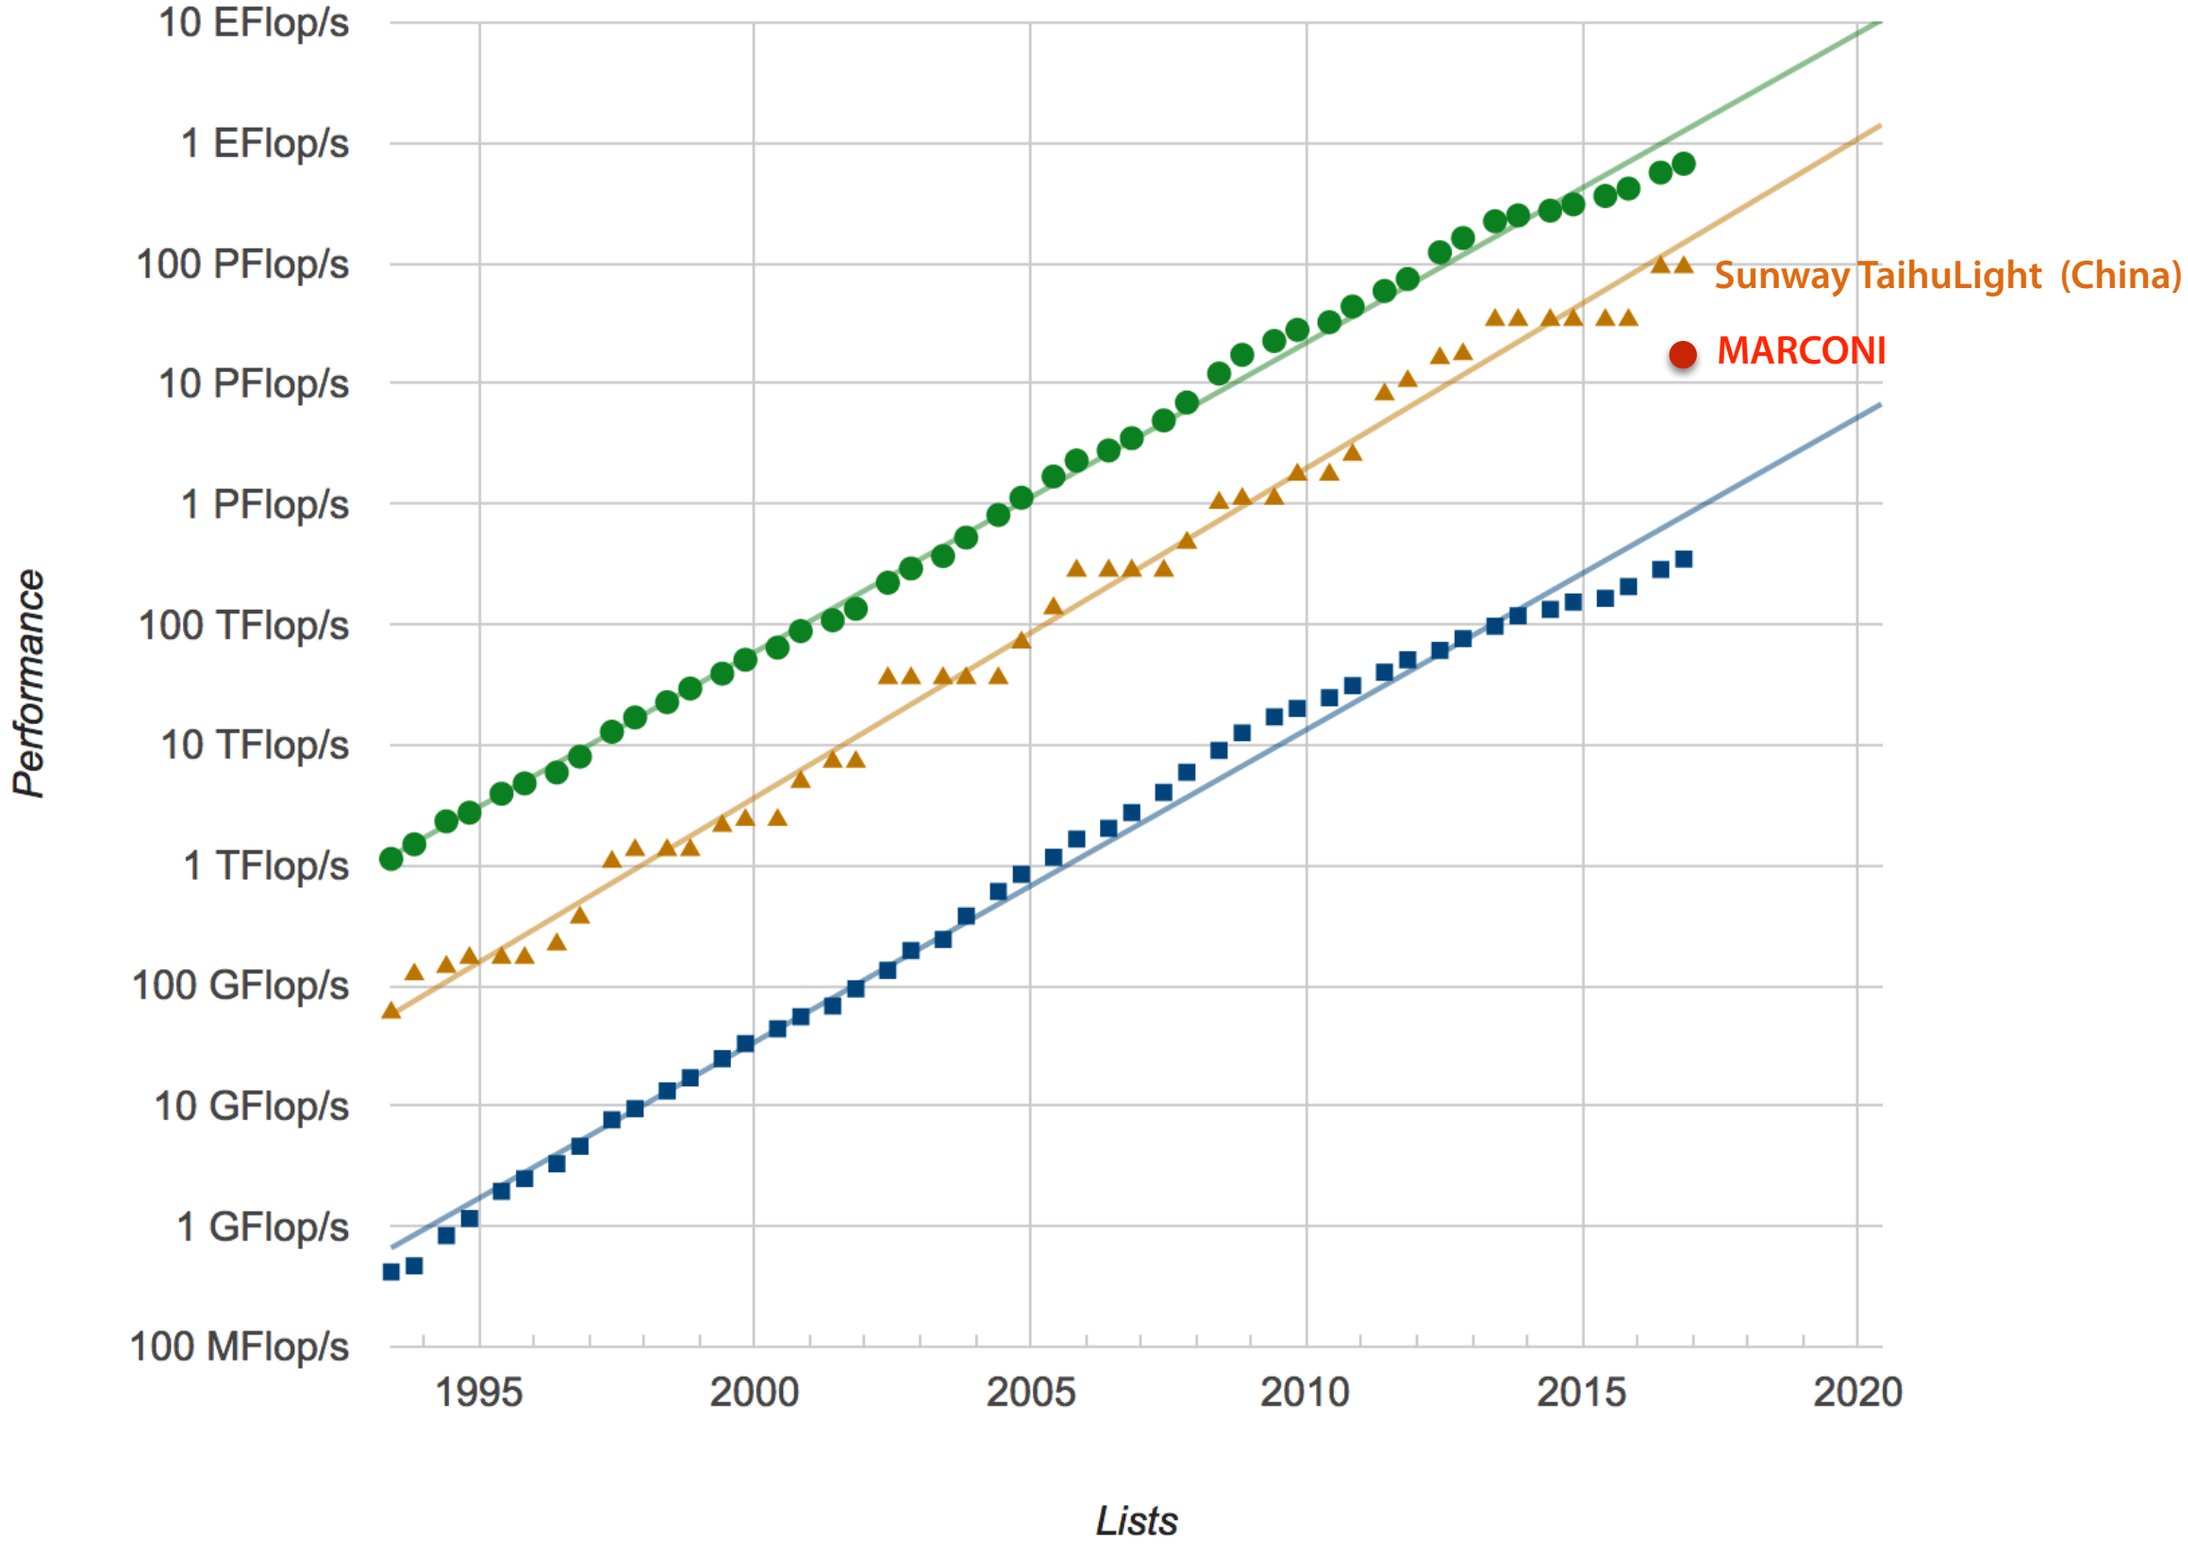

La convenzione con il CINECA

FERMI (SET 2012- JUL 2016)

i.s.	corehours	percentuale
FBS	3,909,987	0.93
FTeCP	36,062,944	8.59
biophys	7,513,439	1.79
fldturb	23,968,009	5.71
lqcd123	106,970,718	25.48
manybody	2,988,883	0.71
npqcd	101,896,796	24.27
qcdlat	96,954,092	23.09
teongrav	39,618,431	9.44
TOTALE	**419,883,299**	**100.00**

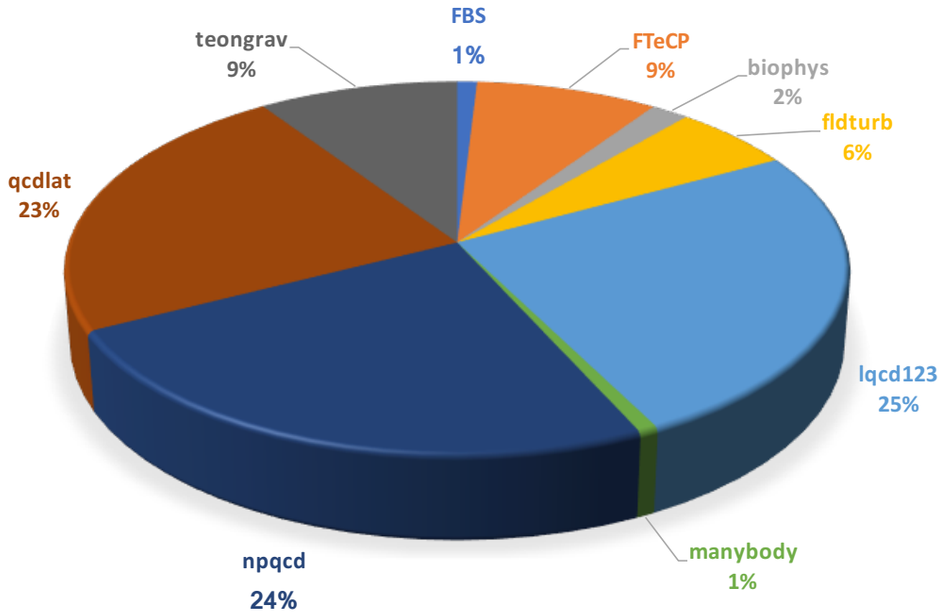

Area	corehours	percentuale
Lattice QCD	341,884,550	81.42
General Relativity	39,618,431	9.44
Quantitative Biology	7,513,439	1.79
Fluid Dynamics	23,968,009	5.71
Nuclear Physics	6,898,870	1.64
Totale	**419,883,299**	**100.00**

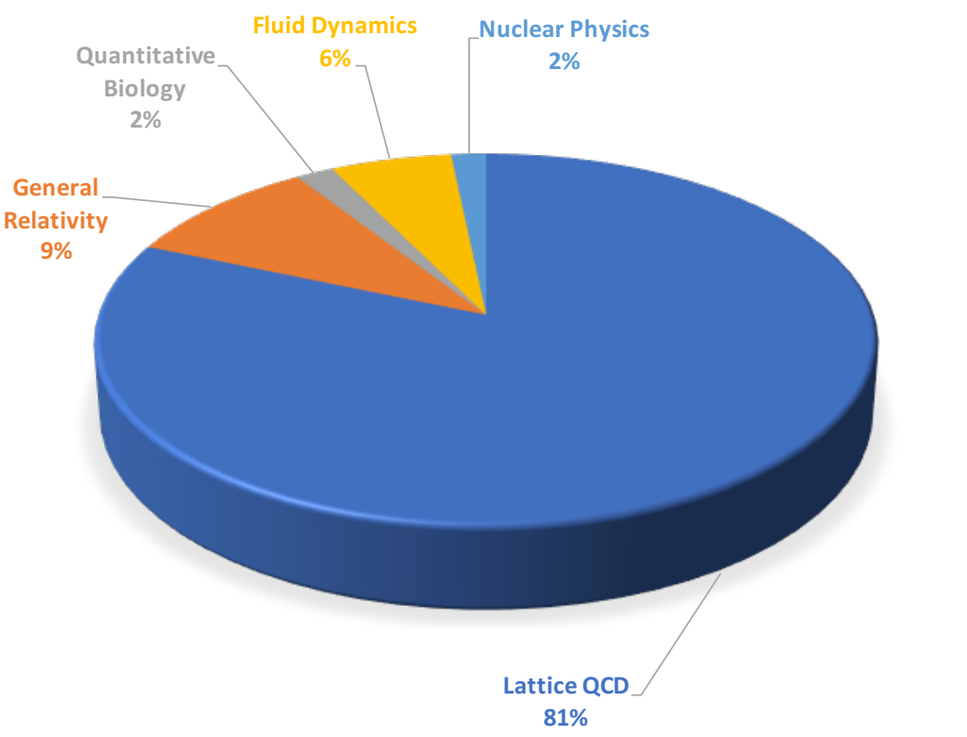

La convenzione con il CINECA

GALILEO (FEB 2015- JAN 2017)

i.s.	corehours	percentuale
FBS	251,737	0.96
FTeCP	558,171	2.12
biophys	3,184,282	12.09
fldturb	2,338,680	8.88
lqcd123	3,064,944	11.64
manybody	1,001,300	3.80
npqcd	4,736,025	17.99
qcdlat	6,201,524	23.55
teongrav	4,112,987	15.62
LHCb	37,156	0.14
test	91,437	0.35
nemesys	502,815	1.91
fusion	249,980	0.95
TOTALE	**26,331,038**	**100.00**

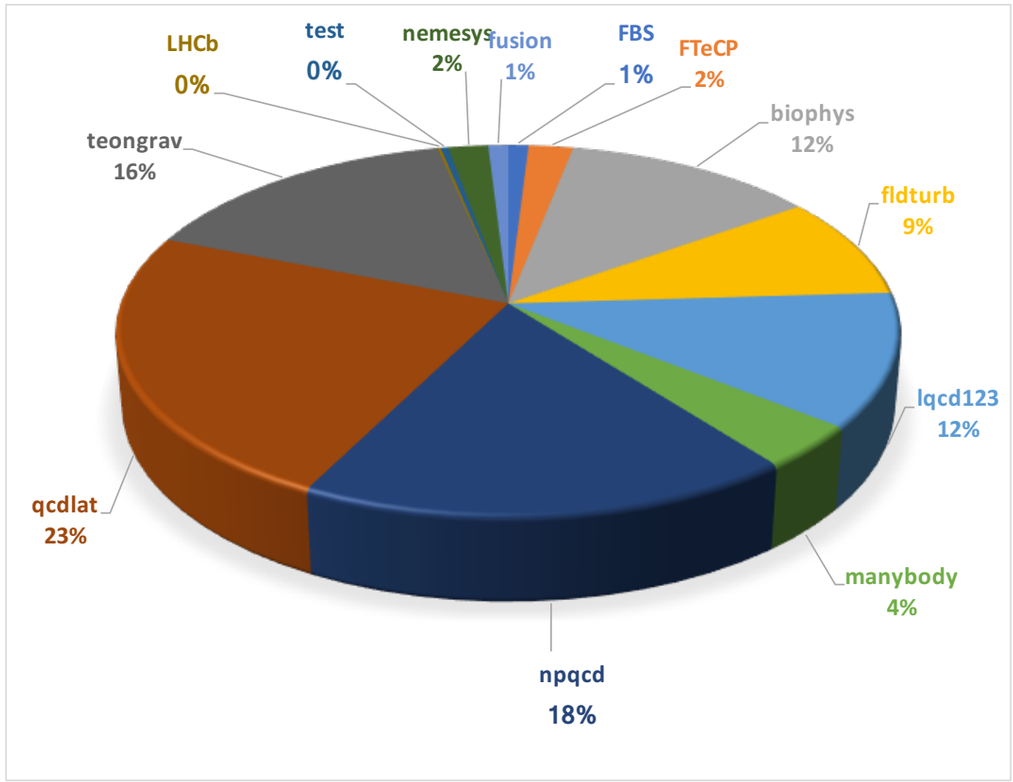

Area	corehours	percentuale
Lattice QCD	14,560,664	55.30
General Relativity	4,112,987	15.62
Quantitative Biology	3,184,282	12.09
Fluid Dynamics	2,338,680	8.88
Nuclear Physics	1,253,037	4.76
Condensed matter	502,815	1.91
Others	378,573	1.44
TOTALE	**26,331,038**	**100.00**

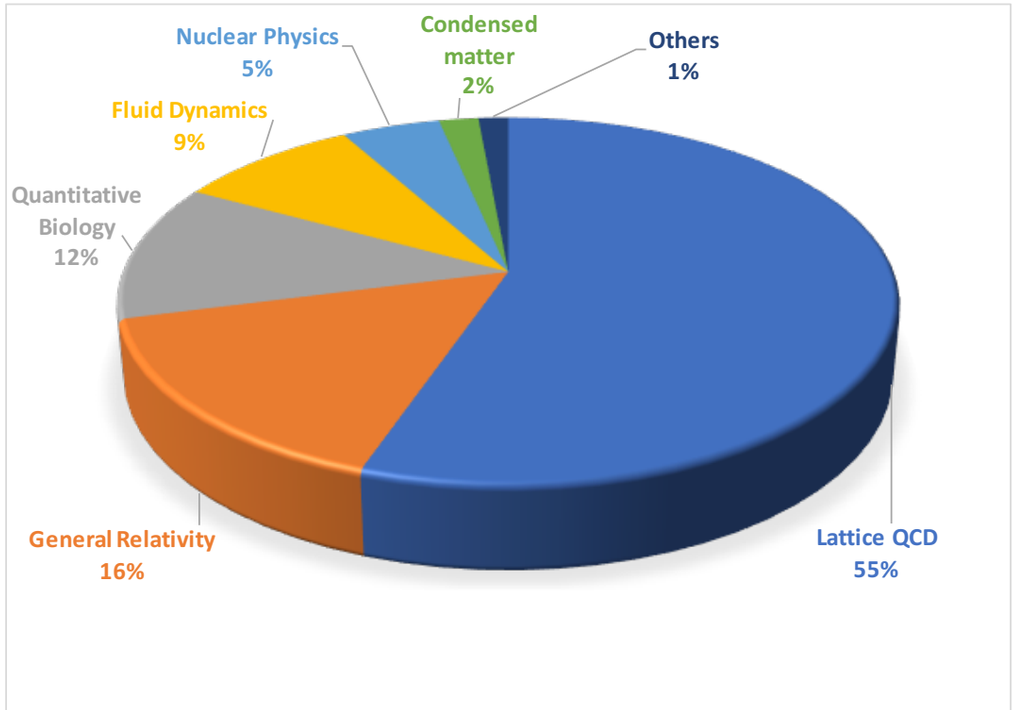

Un confronto con le altre attività

09/02/2017, 11:39

FERMI

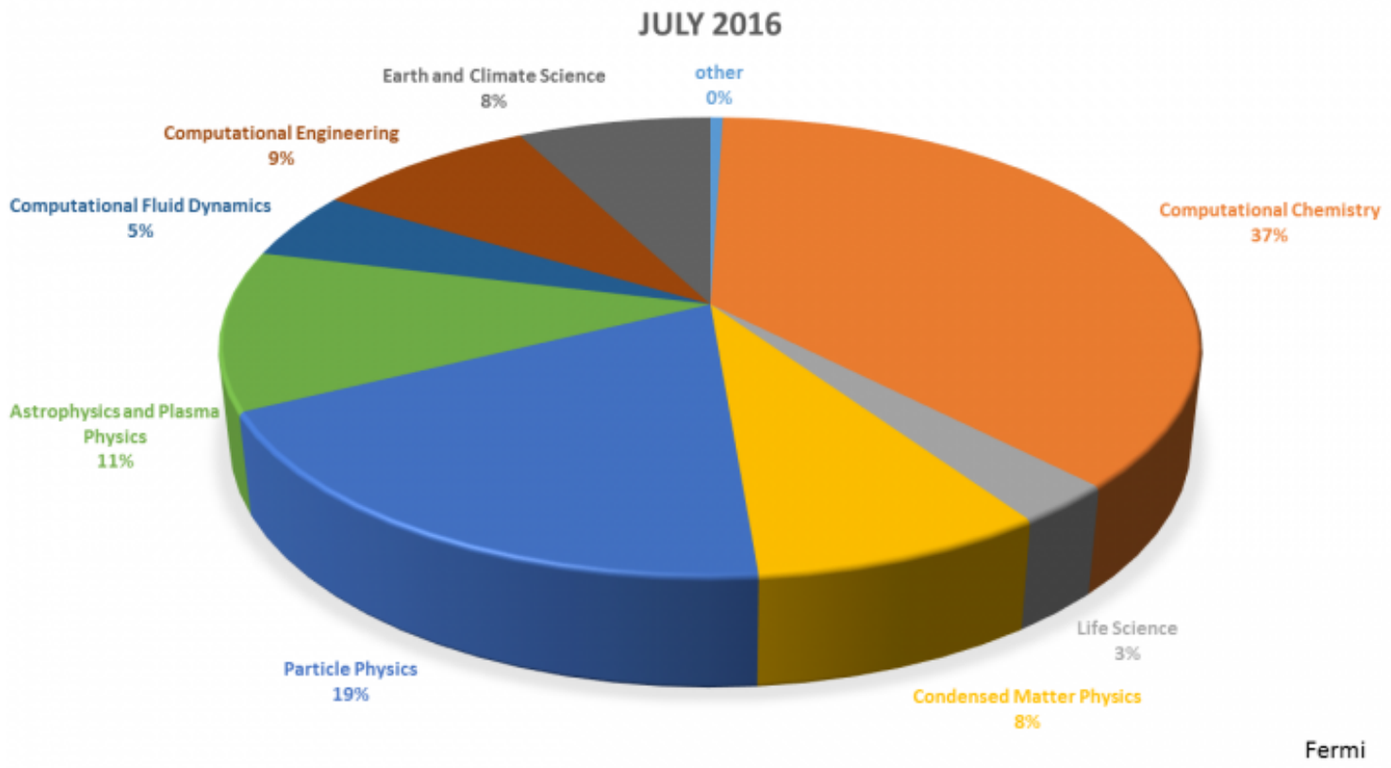

MARCONI

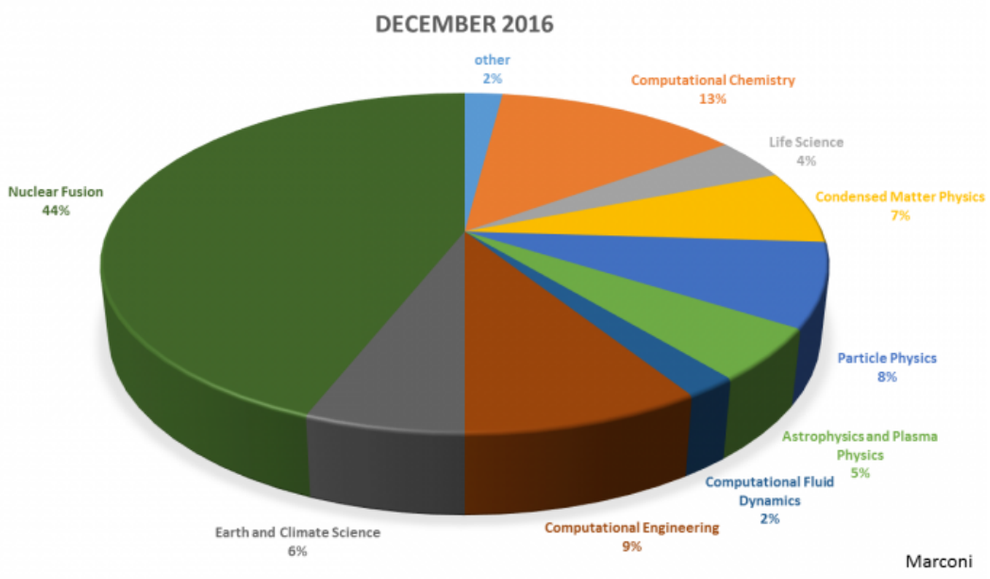

GALILEO

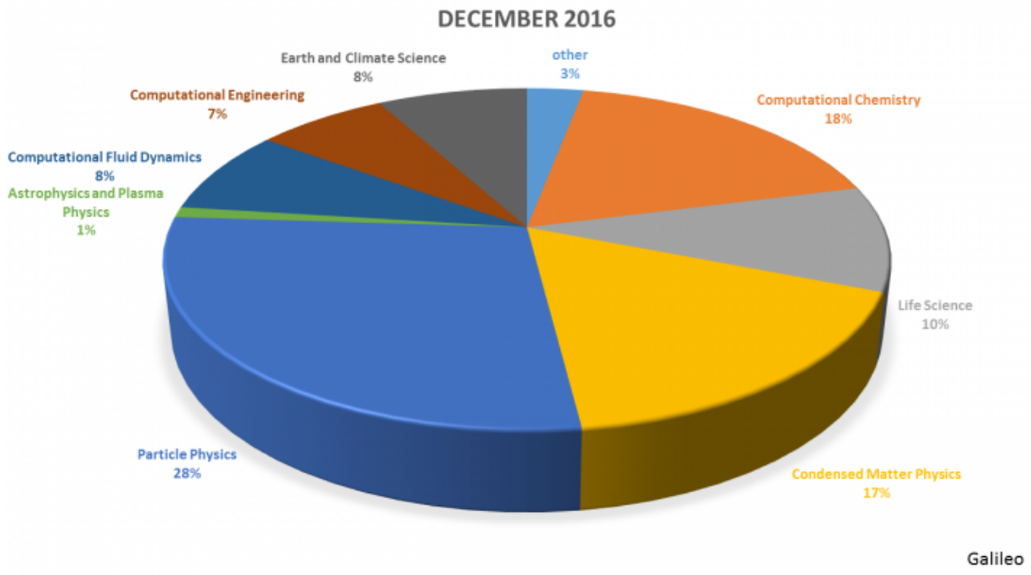

Supporto alle attività HPC e HTC dell'INFN

The case for a renewed support of computational theoretical physics at INFN

S. Arezzini, L. Biferale*, M. Caselle*, A. Cea, A. Ciampa, L. Cosmai, G. de Divitiis, M. D'Elia, R. De Pietri*, C. Destri, P. Dimopoulos, F. Di Renzo, R. Frezzotti, L. Giusti*, G. La Penna, M. P. Lombardo*, V. Lubicz*, E. Marinari, G. Martinelli, V. Minicozzi, S. Morante*, M. Nicodemi*, A. Papa*, M. Papinutto, G. Parisi, F. Pederiva*, A. Pellissetto*, M. Pepe, F. Rapuano, C. Ratti, G. C. Rossi*, G. Salina, F. Schifano, S. Simula, N. Tantalo, C. Tarantino, R. Tripiccione*, E. Vicari, P. Vicini*, M. Viviani*, T. Vladikas

* *Conveners*

(Dated: February 18, 2015)

We outline a proposal for a strong and renewed support by INFN to its research groups active in theoretical computational physics. We argue that theoretical computational physics is important *per se* but it is also an important tool to understand in full the results and implications of present and future experiments. We describe a consistent line of action, review the physics area that would benefit from this line of support, and assess the level of resources needed to put this proposal in practice. We discuss several important advantages that this initiative would bring to all of INFN.

	2014	2016	2017
LGT: hadron physics	100	200	300
LGT: QGP and BSM	90	300	360
LGT: flavor physics	120	240	360
General relativity	60	140	170
Quantitative Biology	5	12	15
Fluid Dynamics	60	120	160
Nuclear Physics	10	15	20
Grand Total (Mcore-h)	**445**	**1027**	**1385**
Grand Total (Eq. Pflops)	**0.640**	**1.49**	**2.00**

TABLE II: Estimate of the yearly computing requirements of several areas of computational physics; units are Mcore-hours. In the last line, the grand total is converted to the computing power of an HPC system fully availab to the INFN community.

Ministero

Istruzione

Ufficio Stampa

Roma, 23 dicembre 2015

Ricerca, da Cipe ok a 5 progetti per oltre 30 milioni
Dall'infrastruttura di calcolo, agli stage internazionali

Oltre 30 milioni di euro per finanziare progetti di ricerca. Il Cipe (Comitato interministeriale di programmazione economica) ha ammesso al finanziamento, a valere sul Fondo integrativo speciale per la ricerca (Fisr), cinque iniziative, dall'infrastruttura di calcolo agli stage per studenti universitari e ricercatori in aziende italiane all'estero.

"High performance data network"

Proposto dall'Istituto Nazionale di Fisica Nucleare (INFN), il progetto vuole contribuire a realizzare un'infrastruttura di calcolo innovativa di punta a livello nazionale, dove fare anche formazione a giovani laureati e dottori di ricerca. Valore complessivo: **€ 13.500.000**.

Le risorse di calcolo 2017 (al CINECA)

- convenzione INFN-CINECA
  —> 6% MARCONI + 15 Mcorehours GALILEO

- MARCONI "A1" **18 Mcorehours**

Model: Lenovo NeXtScale
Architecture: Intel OmniPath Cluster
Nodes: 1.512
Processors: 2 x 18-cores Intel Xeon E5-2697 v4 (Broadwell) at 2.30 GHz
Cores: 36 cores/node, 54.432 cores in total
RAM: 128 GB/node, 3.5 GB/core
Internal Network: Intel OmniPath
Disk Space: 17PB (raw) of local storage
Peak Performance: 2 PFlop/s

- MARCONI "A2" **120 Mcorehours**

Model: Lenovo Adam Pass
Architecture: Intel OmniPath Cluster

Nodes: 3.600
Processors: 1 x 68-cores Intel Xeon Phi 7250 CPU (Knights Landing) at 1.40 GHz
Cores: 68 cores/node (272 with HyperThreading), 244.800 cores in total
RAM: 16 GB/node of MCDRAM and 96 GB/node of DDR4
Internal Network: Intel OmniPath Architecture 2:1
Disk Space: 17PB (raw) of local storage
Peak Performance: 11 PFlop/s

- MARCONI "A3" (2500 nodi Skylake)

INFN: ~ 400 nodi Intel Skylake
nodo Skylake:
2*20 cores @2.3GHz, 192 GB Ram,
2,950 GFlops

~1.2 PFlops ~ JUL 2017

(alcune) CONCLUSIONI

- Collaborazione INFN - CINECA rende possibile affrontare progetti di fisica con utilizzo di significative risorse hardware.
- La comunità scientifica INFN ha a disposizione una buona base di risorse di calcolo con la possibilità di accedere con maggiore probabilità di successo alle risorse ISCRA-PRACE.
- Consolidare e favorire un adeguato upgrade nei prossimi anni delle risorse di calcolo HPC a disposizione della comunità di fisica teorica INFN.

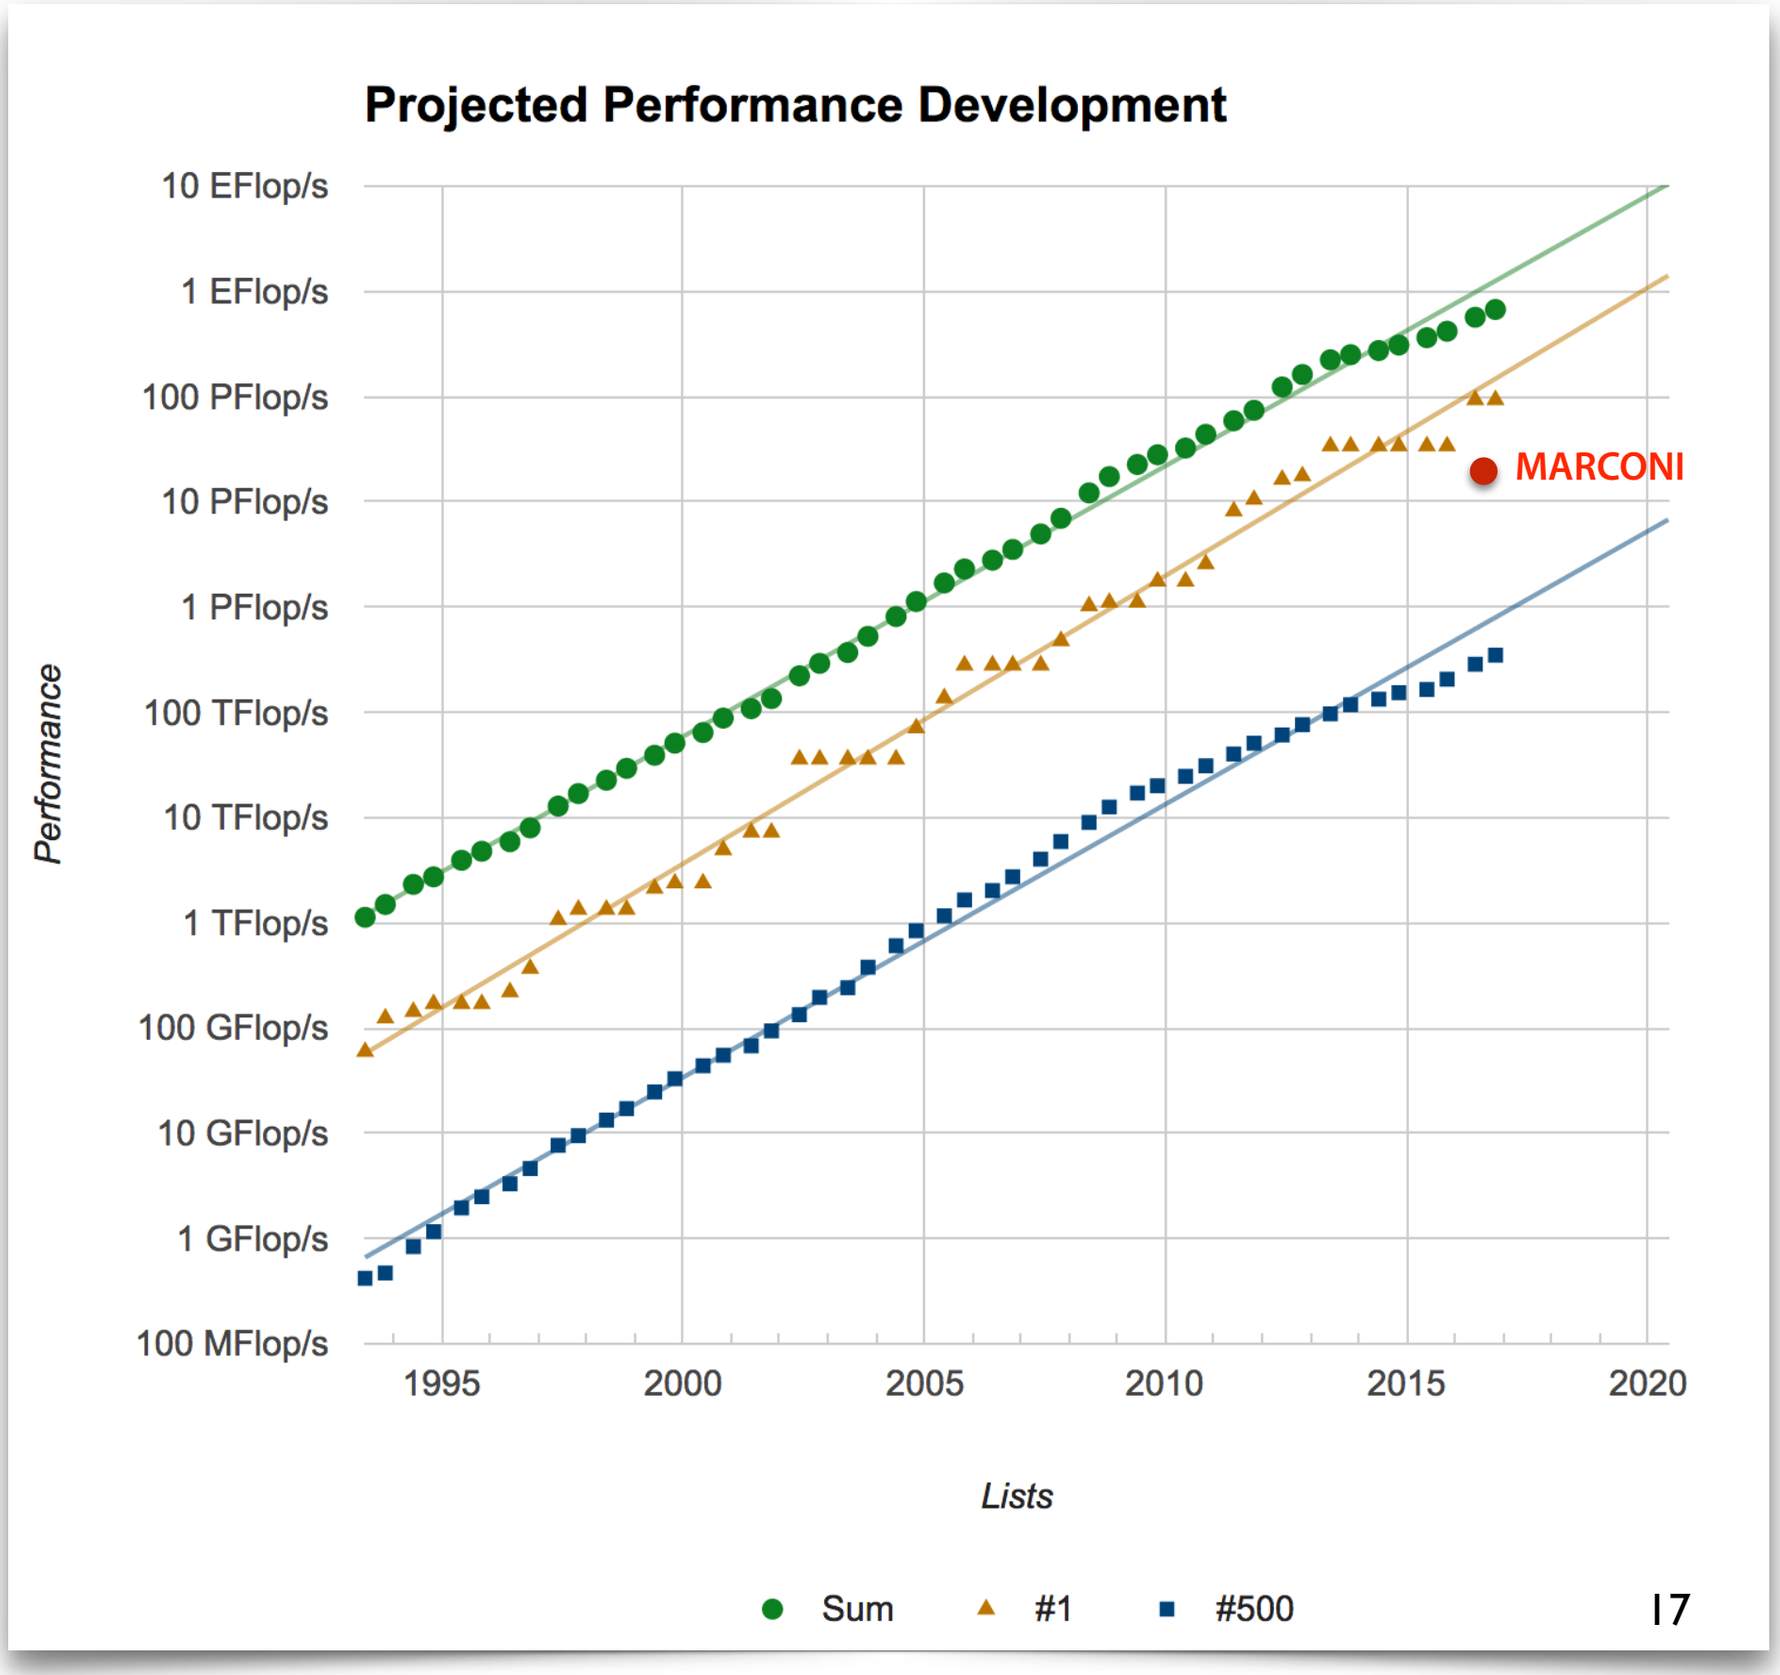

Il Programma di oggi

Presentazioni

presentazioni i.s.

Introduzione	*COSMAI, Leonardo*
aula "Careri" primo piano edificio "Marconi", Dipartimento di Fisica Sapienza Univ. di Roma	10:00 - 10:15
BIOPHYS (1)	*NICODEMI, Mario*
aula "Careri" primo piano edificio "Marconi", Dipartimento di Fisica Sapienza Univ. di Roma	10:15 - 10:20
BIOPHYS (2)	*STELLATO, Francesco*
aula "Careri" primo piano edificio "Marconi", Dipartimento di Fisica Sapienza Univ. di Roma	10:20 - 10:25
FBS	*VIVIANI, Michele*
aula "Careri" primo piano edificio "Marconi", Dipartimento di Fisica Sapienza Univ. di Roma	10:25 - 10:35
FIELDTURB	*LANOTTE, Alessandra S.*
aula "Careri" primo piano edificio "Marconi", Dipartimento di Fisica Sapienza Univ. di Roma	10:35 - 10:45
LQCD123	*SIMULA, Silvano*
aula "Careri" primo piano edificio "Marconi", Dipartimento di Fisica Sapienza Univ. di Roma	10:45 - 10:55
MANYBODY	*PEDERIVA, Francesco*
aula "Careri" primo piano edificio "Marconi", Dipartimento di Fisica Sapienza Univ. di Roma	10:55 - 11:05
NEMESYS	*SINDONA, Antonio*
aula "Careri" primo piano edificio "Marconi", Dipartimento di Fisica Sapienza Univ. di Roma	11:05 - 11:15
NPQCD	*D'ELIA, Massimo*
aula "Careri" primo piano edificio "Marconi", Dipartimento di Fisica Sapienza Univ. di Roma	11:15 - 11:25
QCDLAT	*GIUSTI, Leonardo*
aula "Careri" primo piano edificio "Marconi", Dipartimento di Fisica Sapienza Univ. di Roma	11:25 - 11:35
QFT_HEP	*COLANGELO, Pietro*
aula "Careri" primo piano edificio "Marconi", Dipartimento di Fisica Sapienza Univ. di Roma	11:35 - 11:45
SFT	*CASELLE, Michele*
aula "Careri" primo piano edificio "Marconi", Dipartimento di Fisica Sapienza Univ. di Roma	11:45 - 11:55
STRENGTH	*ITACO, Nunzio*
aula "Careri" primo piano edificio "Marconi", Dipartimento di Fisica Sapienza Univ. di Roma	11:55 - 12:05
TEONGRAV	*DE PIETRI, Roberto*
aula "Careri" primo piano edificio "Marconi", Dipartimento di Fisica Sapienza Univ. di Roma	12:05 - 12:15
Neutron Star Matter (NeuMatt)	*DRAGO, Alessandro et al.*
aula "Careri" primo piano edificio "Marconi", Dipartimento di Fisica Sapienza Univ. di Roma	12:15 - 12:20
Cosmological computing with InDark	*NATOLI, Paolo*
aula "Careri" primo piano edificio "Marconi", Dipartimento di Fisica Sapienza Univ. di Roma	12:20 - 12:30
La macchina sperimentale	*VICINI, Piero*
aula "Careri" primo piano edificio "Marconi", Dipartimento di Fisica Sapienza Univ. di Roma	12:30 - 12:40
Eventuali ulteriori presentazioni e inizio discussione	
aula "Careri" primo piano edificio "Marconi", Dipartimento di Fisica Sapienza Univ. di Roma	12:40 - 13:00

(14:30- 17:00) Discussione

"The case for a renewed support of computational theoretical physics at INFN": 2018-2020