Introduction & HPC

Ivan Girotto – igirotto@ictp.it

HPC Applications Specialist

International Centre for Theoretical Physics (ICTP)

Mission - An institute run by scientists for scientists

- Foster the growth of advanced studies and research in physical and mathematical sciences, especially in support of excellence in developing countries.
- Develop high-level scientific programmes keeping in mind the needs of developing countries, and provide an international forum of scientific contact for scientists from all countries.
- Conduct research at the highest international standards and maintain a conducive environment of scientific inquiry for the entire ICTP community.
- Thanks to the generous funding from the Italian Government, UNESCO and the IAEA, ICTP has been able to initiate and implement various schemes of support and assistance to scientists from developing countries.

The Abdus Salam
International Centre
for Theoretical Physics

United Nations
Educational, Scientific and
Cultural Organization

IAEA
International Atomic Energy Agency

ICTP from Trieste to the World

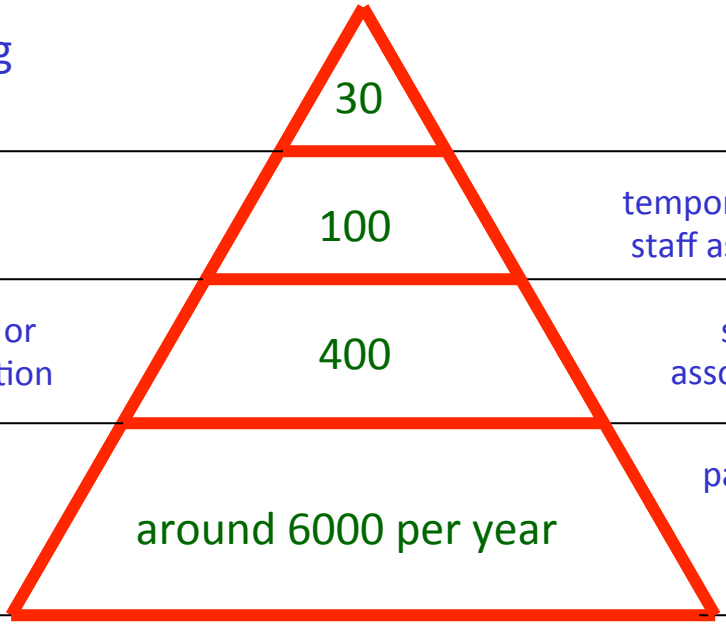

ICTP IN NUMBERS 2016

ICTP VISITORS 2016

5827 VISITORS [25% FEMALE] FROM **135** NATIONS; **56** TRAINING ACTIVITIES ON CAMPUS, **21** IN DEVELOPING COUNTRIES

8 DAYS AVERAGE LENGTH OF VISIT FOR CONFERENCE PARTICIPANTS

57 DAYS AVERAGE FOR RESEARCH VISITORS

59 POSTDOCS ON CAMPUS [47% FROM DEVELOPING COUNTRIES]

232 STUDENTS ENROLLED IN PRE-PHD EDUCATIONAL PROGRAMMES

367 SCIENTISTS ENGAGED IN CAREER DEVELOPMENT PROGRAMMES

TOP 10 DEVELOPING COUNTRIES, BY REGION

Africa	N. of visitors	Asia	N. of visitors	Latin America	N. of visitors
South Africa	85	India	378	Brazil	170
Ghana	79	Iran	324	Argentina	160
Nigeria	78	China	175	Mexico	77
Egypt	69	Pakistan	72	Colombia	51
Cameroon	48	Malaysia	45	Cuba	49
Morocco	41	Singapore	44	Chile	33
Algeria	35	Turkey	43	Venezuela	18
Tunisia	32	Korea Rep.	39	Costa Rica	17
Kenya	30	Philippines	28	Peru	17
Senegal	30	Viet Nam	24	Guatemala	10

FLOW OF INCOMING VISITORS BY MONTH AT ICTP

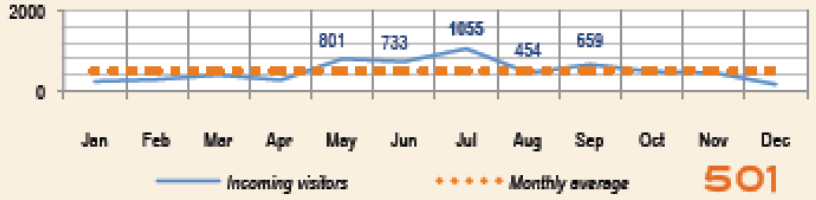

501

22 REGIONAL TRAINING ACTIVITIES

6 - MATHEMATICS (MATH)
5- APPLIED PHYSICS (AP)
5 - EARTH SYSTEM PHYSICS (ESP)
3 - CONDENSED MATTER PHYSICS (CMSP)
2 - HIGH ENERGY PHYSICS (HECAP)
1- QUANTITATIVE LIFE SCIENCES (QLS)

DISTRIBUTION OF FEMALE VISITORS

60% OF FEMALE VISITORS ARE FROM DEVELOPING REGIONS:

Asia-Pacific	416
Africa	184
Latin America	168
Eastern Europe	119

COURSE PARTICIPANTS BY RESEARCH AREA

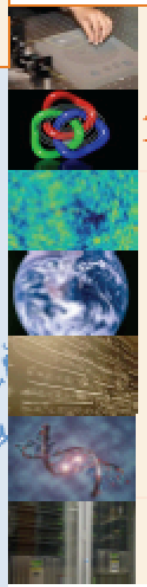

1324 AP

1168 CMSP

814 HECAP

697 ESP

575 Math

359 QLS

281 HPC (High Performance Computing)

1,500 MONTHS OF TRAINING TO COURSE PARTICIPANTS LECTURED BY MORE THAN 1,300 EXPERTS

COUNTRIES REPRESENTED

20 LDCs in Africa

17 from rest of Africa

6 LDCs in Asia

30 from rest of Asia

19 from Latin America

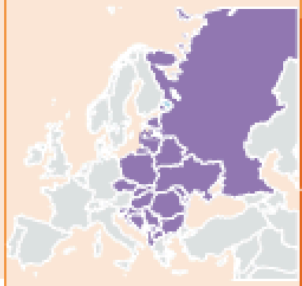

19 from Eastern Europe

The Abdus Salam International Centre for Theoretical Physics

- More than 140,000 visits since 1970
- 190 countries represented
- 20% of ICTP visiting scientists are women

++NEW++
The ICTP Partner Institutes

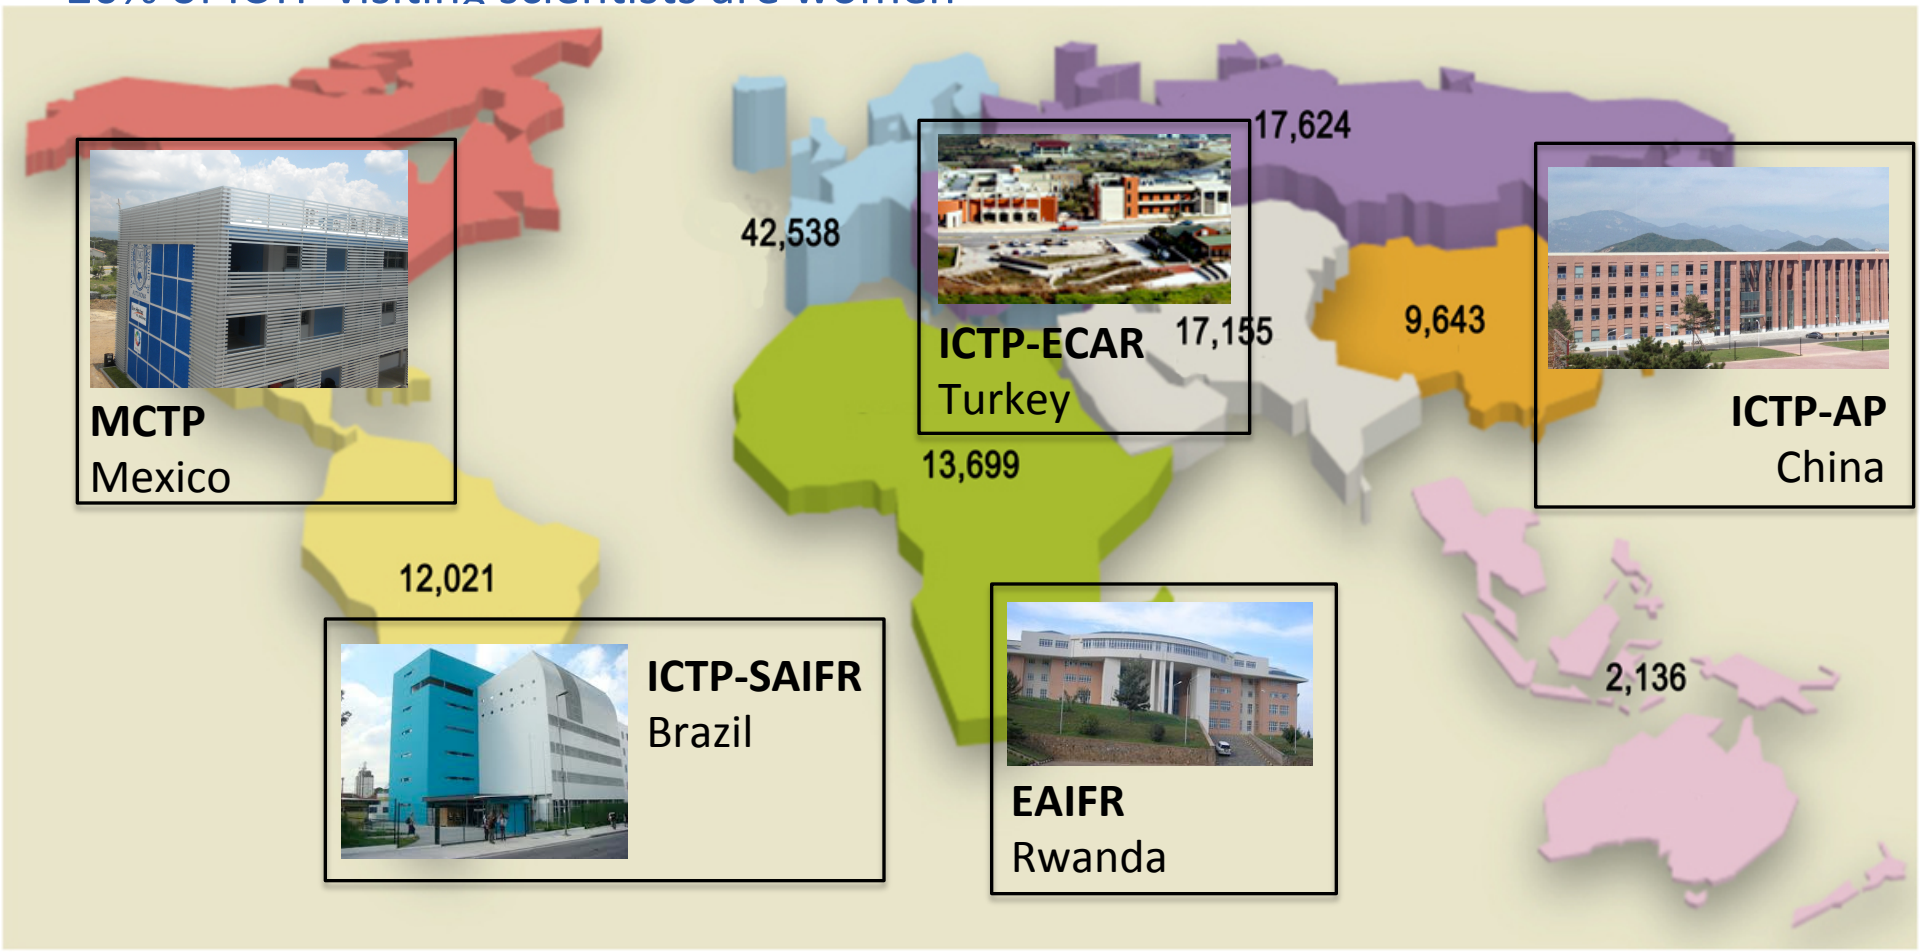

The Abdus Salam International Centre for Theoretical Physics

PRE-PHD PROGRAMMES	DEGREE PROGRAMMES	CAREER DEVELOPMENT	LABORATORY OPPORTUNITIES	SCIENTIFIC OUTREACH
ICTP Postgraduate Diploma Programme	Joint ICTP/SISSA PhD Programme in Physics and Mathematics	Conferences, workshops and schools	Training and Research in Italian Laboratories	Office of External Activities
ICTP/IAEA Sandwich Training Education Programme	Joint PhD Programme, Earth Science and Fluid Mechanics	Junior Associates	ICTP-ELETTRA Users Programme	ICTP Partner Institutes
	Physics PhD Program	Regular Associates	ICTP Laboratories	Science Dissemination Unit
	Joint Masters in Physics	Senior Associates		African Review of Physics
	Joint ICTP/Collegio Carlo Alberto Program in Economics	Federated Institutes		ICTP in East Africa
	International Master, Physics of Complex Systems	OFID Postgraduate Fellowship		Physics Without Frontiers
	Master of Advanced Studies in Medical Physics	The Kuwait Programme at ICTP		
	Masters in High Performance Computing			

ICTP Scientific Calendar

- Schools, Conferences, Workshops around the year
- Half of them on subjects related to main research areas (core)
- The rest on many subjects:
  medical physics, optics, nano physics, plasma physics, electronics, high performance computing, biophysics, satellite navigation, science dissemination and e-learning, m-science, entrepreneurship, nuclear physics (IAEA), teacher training, 3-D Printing, etc...
- http://www.ictp.it/scientific-calendar.aspx

Scientific Sections

High Energy Cosmology and Astroparticle Physics

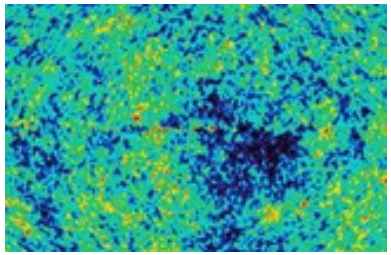

Condensed Matters and Statistical Physics

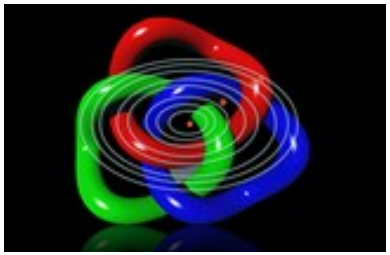

Earth System Physics

Mathematics

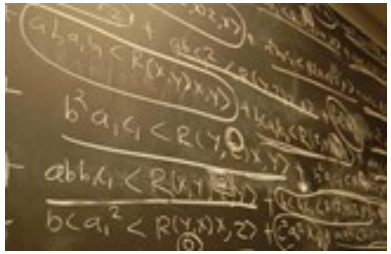

Applied Physics

New areas

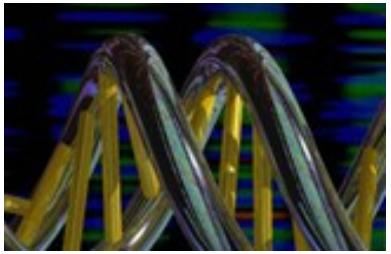

HPC Staff and Collaborators

Dr. David Grellsheid
Herwig Software Manager
@ Durham University

Dr. Axel Kohlmeyer
Full Professor of Research
@ Temple University

Dr. Clement Onime
Responsible IT/HPC Infrastructure
@ ICTP

me
HPC Application Specialist
@ ICTP

High-Performance & Scientific Computing activities at the ICTP

- HPC service and HPC application consulting
  - in house HPC facility (Argo)
  - research enablement on massively parallel systems for HPC on both national service (CINECA) and EU infrastructures (PRACE)
- Training on HPC and Scientific Programming

MHPC in pills: www.mhpc.it

- High-level educational program: not an Ms.C. program!
- Intensive training aimed to build knowledge in solving complex problems with an HPC approach
- Innovative, hands-on based training
- Aimed to people with strong interest in:
  - advanced programming for scientific computing
  - software optimization
  - management of computing platforms
  - data management and data analytics

MHPC
Master in High Performance Computing

Background Requirements

- Candidates must have some experience in programming and a competence in at least one of the languages between C, C++ and/or Fortran
  - Python knowledge is a plus
- A sound knowledge of Linux operating system
- Master level of a scientific degree is required
- No prior HPC knowledge is assumed
- Enthusiasm is a must

1 year program divided in 6-8 months courses and 6 month project (some overlap)

Mandatory

- Scientific Programming Environment
- Introduction to Computer Architectures for HPC
- Object Oriented Programming
- Parallel Programming
- Introduction to Numerical Analysis
- Advanced Computer Architectures and Optimizations
- Parallel Data Management and Data Exchange
- High Performance Computing Technology
- Best Practices in Scientific Computing

Optional Choice

- Data structures, sorting and searching algorithms in serial and parallel
- Lookup tables, cell lists and neighbor lists
- Domain decomposition techniques
- Parallel FFT techniques
- Parallel Linear Algebra
- Multipole expansion, multi-grid methods
- Adaptive Meshes
- Maximum likelihood techniques
- Cluster or network or graph analysis
- Monte Carlo methods
- Agent-based models
- Automatic differentiation
- DFT from source to code

Latin American Introductory School on Parallel Programming and Parallel Architecture for High Performance Computing

- International experience
- Parallel programming for HPC
- HPC approach to parallel codes
- Best practise experiences
- Computer architectures for HPC

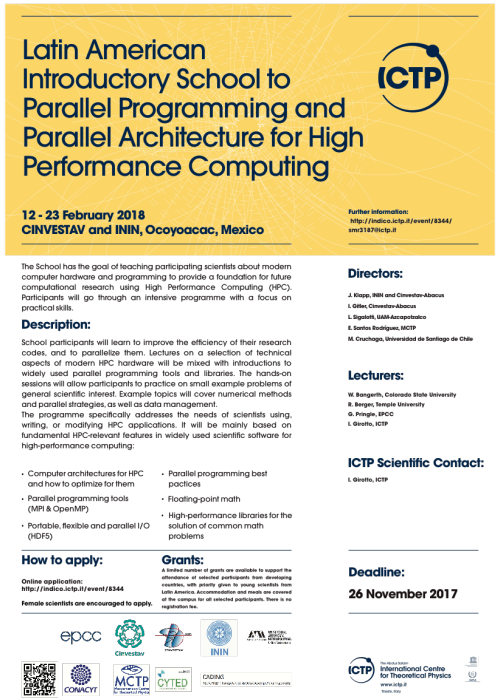

Latin American Introductory School to Parallel Programming and Parallel Architecture for High Performance Computing

ICTP

12 - 23 February 2018
CINVESTAV and ININ, Ocoyoacac, Mexico

Further information:
http://indico.ictp.it/event/8344/sm3187@ictp.it

The School has the goal of teaching participating scientists about modern computer hardware and programming to provide a foundation for future computational research using High Performance Computing (HPC). Participants will go through an intensive programme with a focus on practical skills.

Description:

School participants will learn to improve the efficiency of their research codes, and to parallelize them. Lectures on a selection of technical aspects of modern HPC hardware will be mixed with introductions to widely used parallel programming tools and libraries. The hands-on sessions will allow participants to practice on small example problems of general scientific interest. Example topics will cover numerical methods and parallel strategies, as well as data management.
The programme specifically addresses the needs of scientists using, writing, or modifying HPC applications. It will be mainly based on fundamental HPC-relevant features in widely used scientific software for high-performance computing:

- Computer architectures for HPC and how to optimize for them
- Parallel programming tools (MPI & OpenMP)
- Portable, flexible and parallel I/O (HDF5)
- Parallel programming best pactices
- Floating-point math
- High-performance libraries for the solution of common math problems

How to apply:

Online application:
http://indico.ictp.it/event/8344

Female scientists are encouraged to apply.

Grants:

A limited number of grants are available to support the attendance of selected participants from developing countries, with priority given to young scientists from Latin America. Accommodation and meals are covered at the campus for all selected participants. There is no registration fee.

epcc Cinvestav ININ CONACYT MCTP CYTED CADING

Directors:

J. Kiapp, ININ and Cinvestav-Abacus
I. Gitler, Cinvestav-Abacus
L. Sigalotti, UAM-Azcapotzalco
E. Santos Rodriguez, MCTP
M. Cruchaga, Universidad de Santiago de Chile

Lecturers:

W. Rangerth, Colorado State University
R. Berger, Temple University
G. Pringle, EPCC
I. Girotto, ICTP

ICTP Scientific Contact:

I. Girotto, ICTP

Deadline:

26 November 2017

ICTP The Abdus Salam International Centre for Theoretical Physics www.ictp.it Trieste, Italy

Dr. Gavin Pringle
Applications Consultant
EPCC @ University of of Edinburgh (UK)

Dr. Richard Berger
Research Assistant Professor
Temple University (USA)

Prof. Wolfgang Bangerth
Professor
Colorado State University (USA)

Dr. J. Manuel Solano-Altamirano
Research Assistant Professor
Benemérita Universidad Autónoma de Puebla (Mexico)

Jimmy Aguilar Mena
PhD Student
BSC-CNS (Spain)

Marlon Brenes Navarro
PhD Student
Trinity College Dublin (Ireland)

Dr. William Fernando Oquendo Patiño
Assistant Professor
Universidad de La Sabana (Colombia)

Why use Computers in Science?

- Use complex theories without a closed solution: solve equations or problems that can only be solved numerically, i.e. by inserting numbers into expressions and analyzing the results
- Do “impossible” experiments: study (virtual) experiments, where the boundary conditions are inaccessible or not controllable
- Benchmark correctness of models and theories: the better a model/theory reproduces known experimental results, the better its predictions

SW in Science

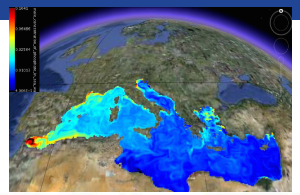

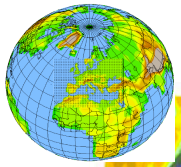

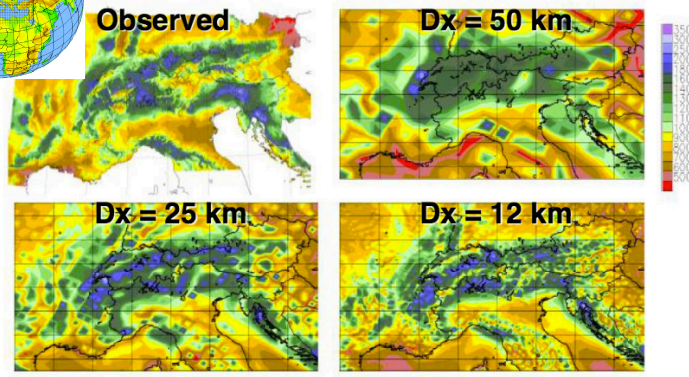

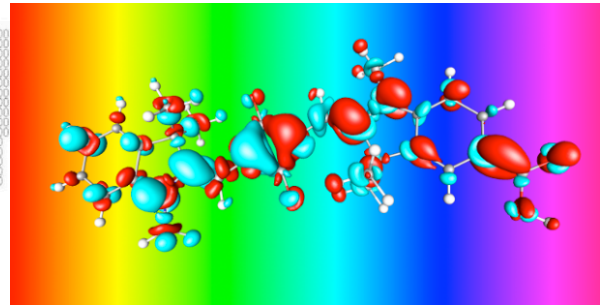

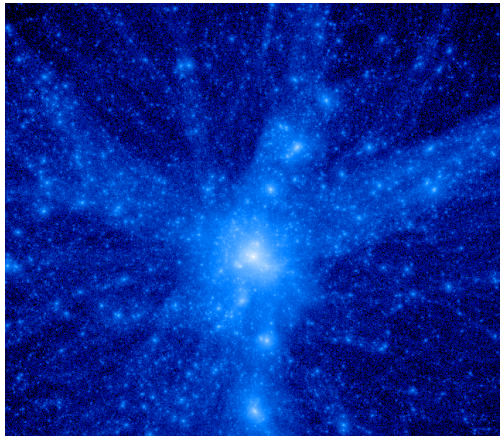

LHCb Event Display

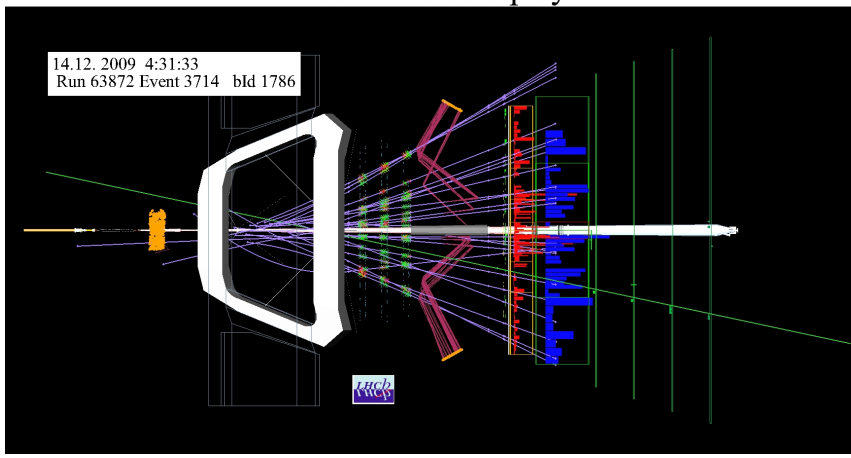

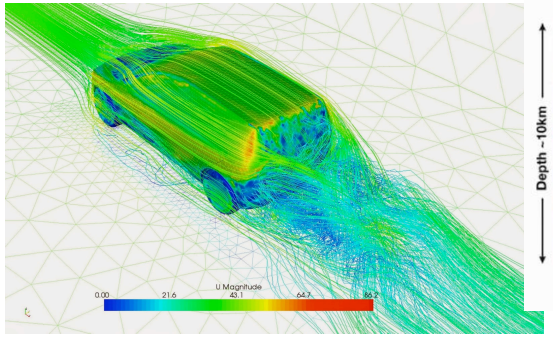

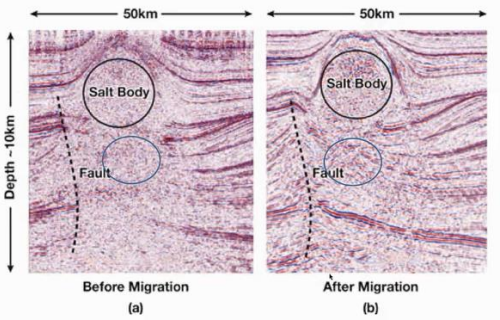

What is High-Performance Computing (HPC)?

- Not a real definition, depends from the prospective:
  - HPC is when I care how fast I get an answer
  - HPC is when I foresee my problem to get bigger and bigger
- Thus HPC can happen on:
  - A workstation, desktop, laptop, smartphone!
  - A supercomputer
  - A Linux Cluster
  - A grid or a cloud
  - Cyberinfrastructure = any combination of the above
- HPC means also High-Productivity Computing

Why would HPC matter to you?

- Scientific computing is becoming more important in many research disciplines
- Problems become more complex, thus need complex software and teams of researchers with diverse expertise working together
- HPC hardware is more complex, application performance depends on many factors
- Technology is also for increasing competitiveness
- HPC knowledge is an opportunity

More & More Computing …

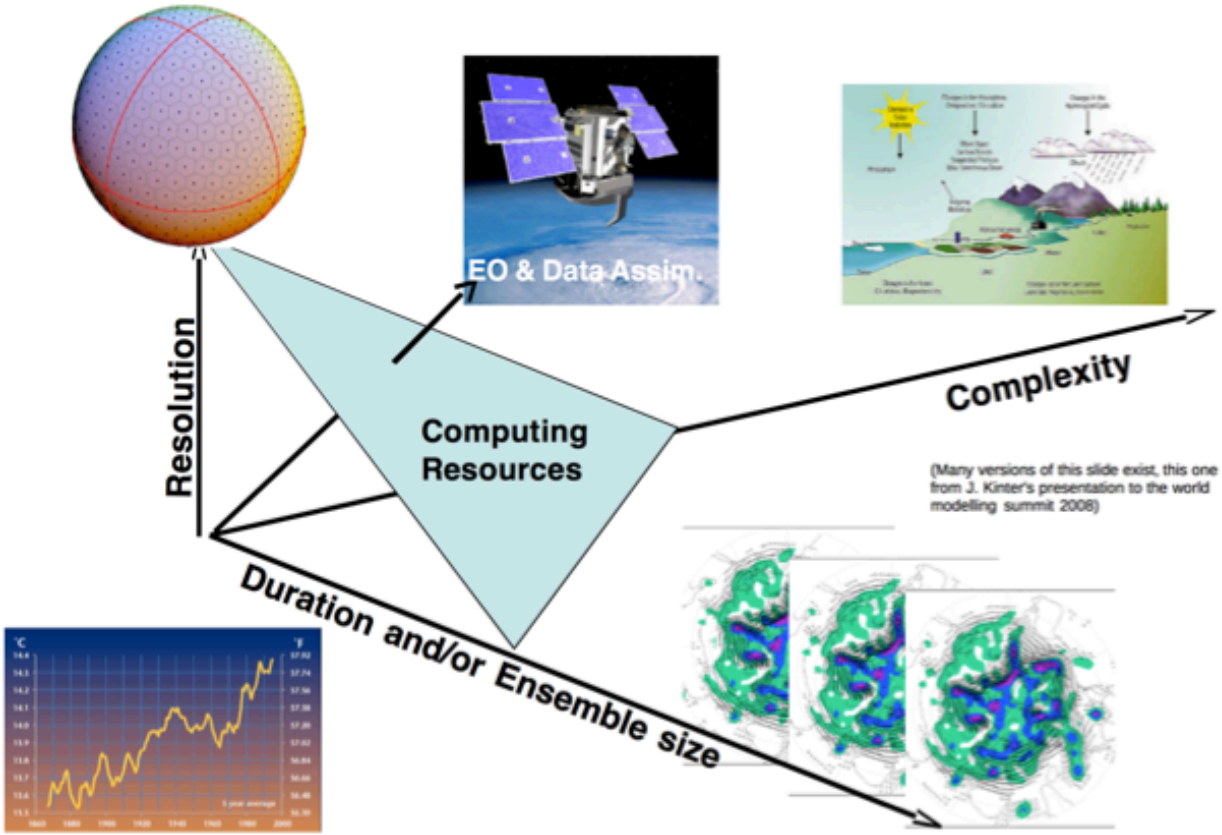

HPC INFRASTRUCTURE

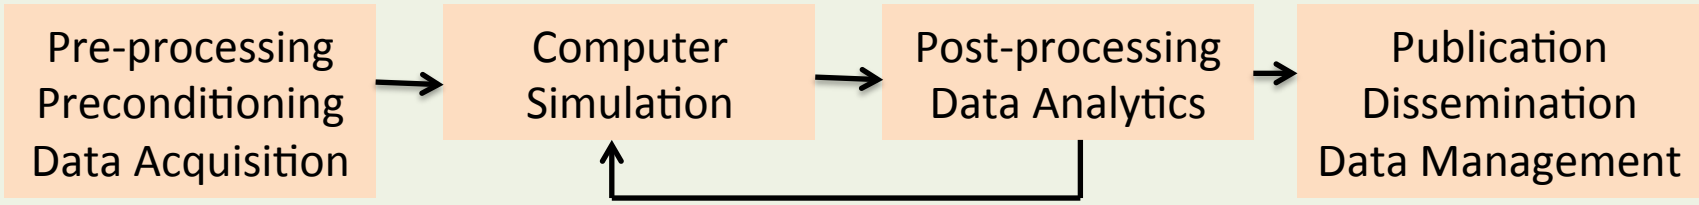

SW Workflow & Parallel Applications

Compilers/Libraries/Debugging & Profiling

HW/Resource Management/File System/...

Parallelism - 101

- there are two main reasons to write a parallel program:
  - access to larger amount of memory (aggregated, going bigger)
  - reduce time to solution (going faster)

Programming Parallel Paradigms

- Are the tools we use to express the parallelism for on a given architecture
- They differ in how programmers can manage and define key features like:
  - parallel regions
  - concurrency
  - process communication
  - synchronism

How do we evaluate the improvement?

- We want estimate the amount of the introduced overhead => $T_o = n_{pes} T_P - T_S$
- But to quantify the improvement we use the term **Speedup**:

$$S_P = \frac{T_S}{T_P}$$

Speedup

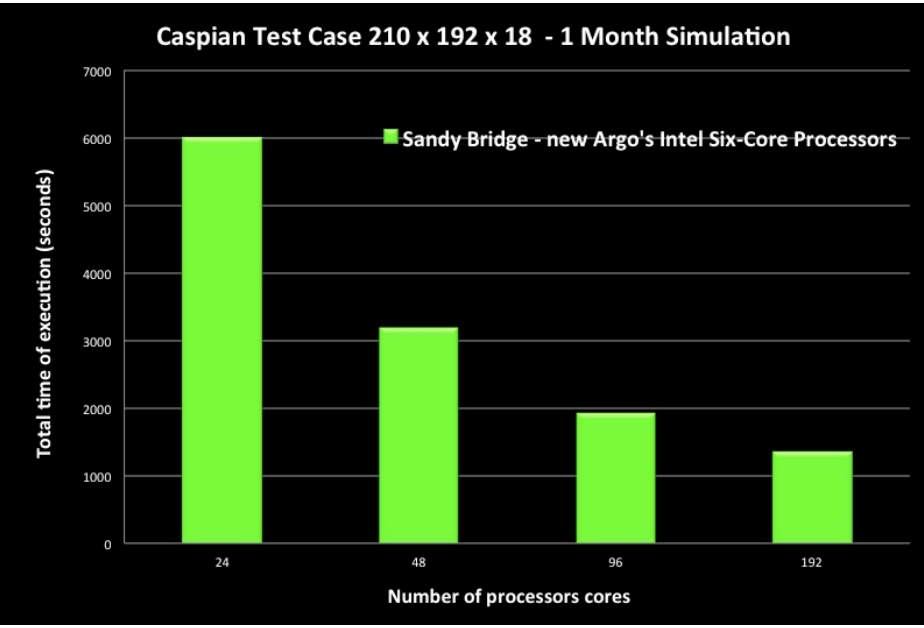
Caspian Test Case 210 x 192 x 18 - 1 Month Simulation
Total time of execution (seconds)
7000
6000
5000
4000
3000
2000
1000
0
Sandy Bridge - new Argo's Intel Six-Core Processors
24
48
96
192
Number of processors cores

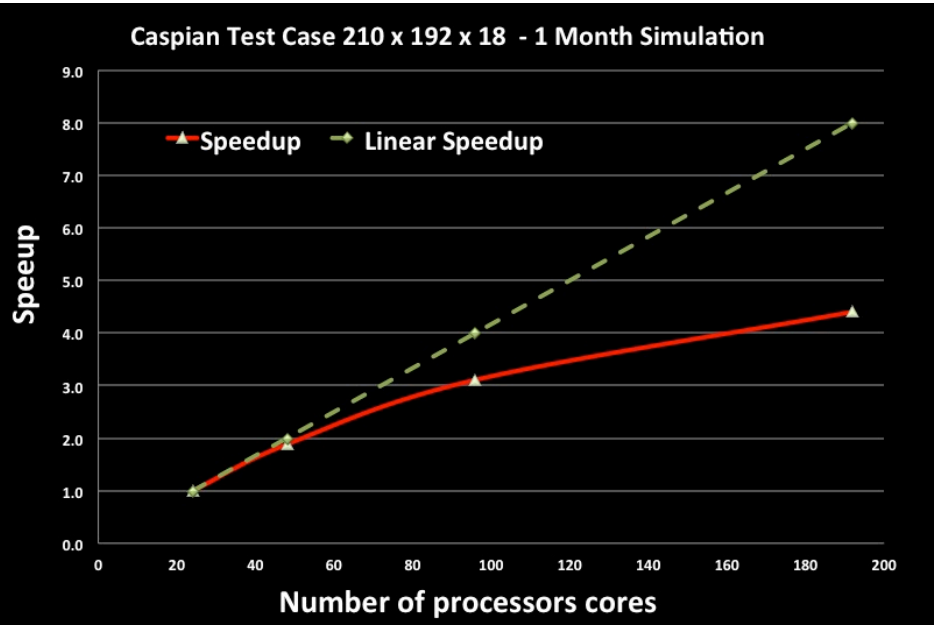
Caspian Test Case 210 x 192 x 18 - 1 Month Simulation
Speeup
9.0
8.0
7.0
6.0
5.0
4.0
3.0
2.0
1.0
0.0
Speedup
Linear Speedup
0
20
40
60
80
100
120
140
160
180
200
Number of processors cores

Efficiency

- Only embarrassing parallel algorithm can obtain an ideal Speedup
- The **Efficiency** is a measure of the fraction of time for which a processing element is usefully employed:

$$E_P = \frac{S_P}{p}$$

Efficiency

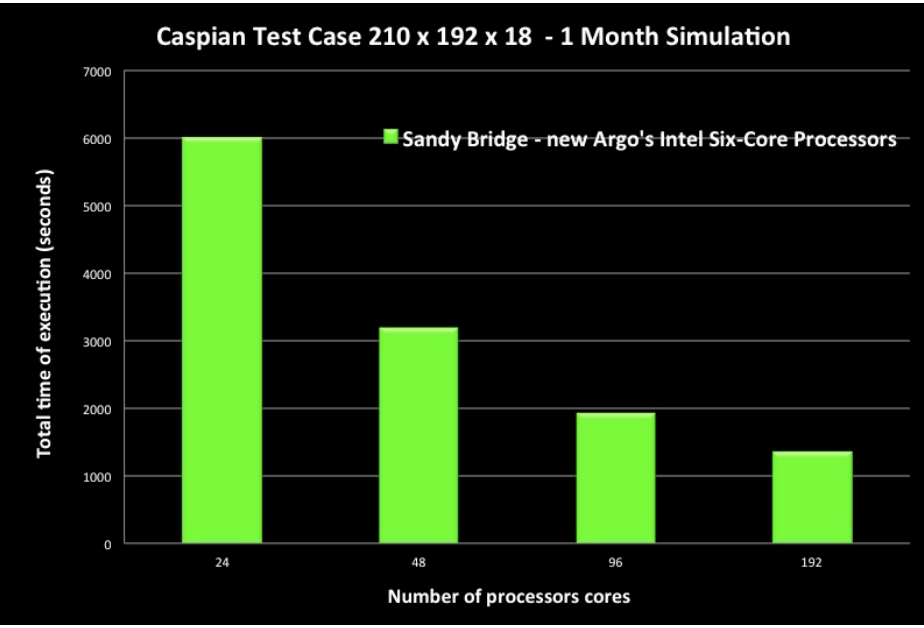
Caspian Test Case 210 x 192 x 18 - 1 Month Simulation
Total time of execution (seconds)
7000
6000
5000
4000
3000
2000
1000
0
Sandy Bridge - new Argo's Intel Six-Core Processors
24
48
96
192
Number of processors cores

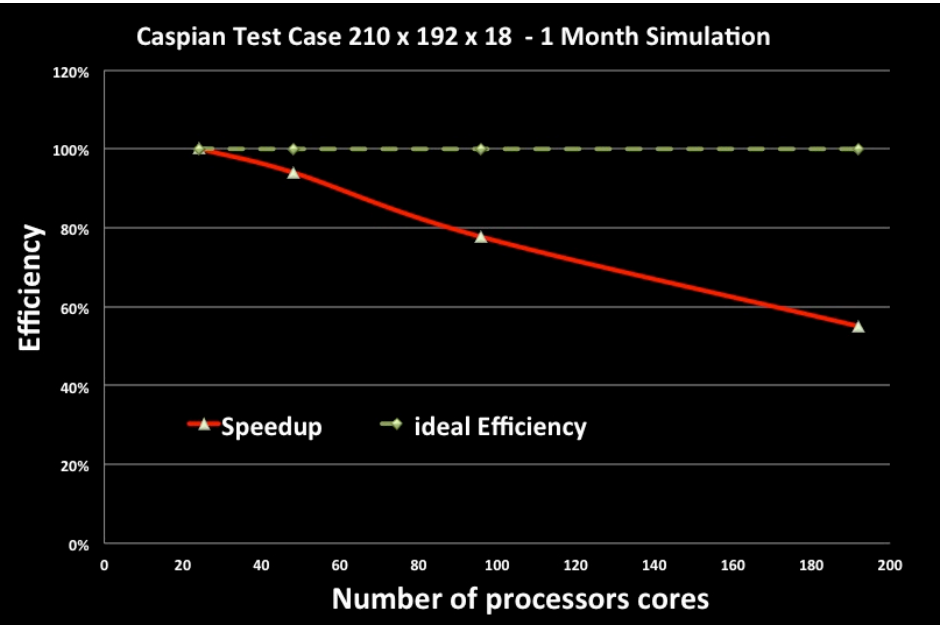
Caspian Test Case 210 x 192 x 18 - 1 Month Simulation
Efficiency
120%
100%
80%
60%
40%
20%
0%
Speedup
ideal Efficiency
0
20
40
60
80
100
120
140
160
180
200
Number of processors cores

Scalability

- When we want consider the scalability of our problem we are interested in two main features:
  - how much faster do we go increasing the number of processes for a fixed problem size (strong scaling)
  - how does the application behave if we increase the problem size keeping the workload fixed per processors

Bands parallelization scaling

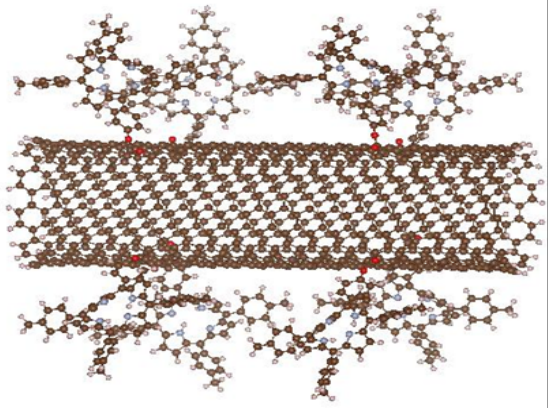

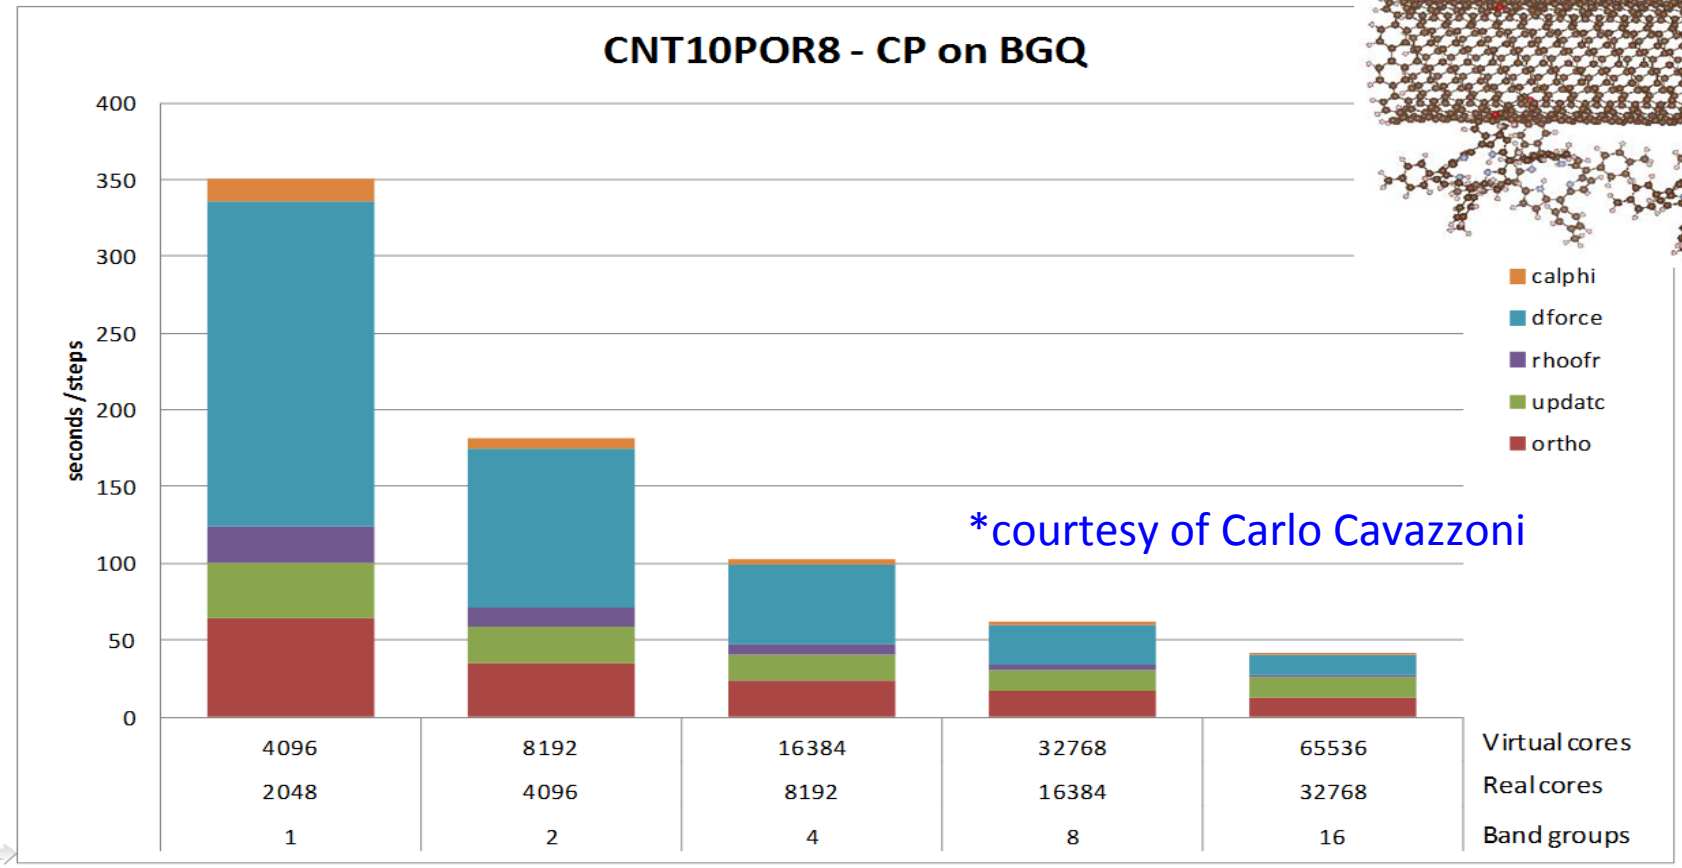

*courtesy of Carlo Cavazzoni

How do we get the profiling? /1

- Code instrumentation

```
double cclock()
/*
      Return the second elapsed
      since Epoch (00:00:00 UTC, January 1, 1970)
*/

{
    struct timeval tmp;
    double sec;
    gettimeofday( &tmp, (struct timezone *)0 );
    sec = tmp.tv_sec + ((double)tmp.tv_usec)/1000000.0;
    return sec;

}
```

- Profiling tools

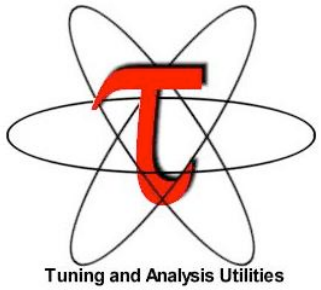

The GNU Profiler (GPROF)

scalasca

How do we get the profiling? /2

- Code instrumentation

```
 Writing output data file c8_atm213_k111.save

 init_run     :     91.65s CPU     91.65s WALL (       1 calls)
 electrons    :   3366.51s CPU   3366.51s WALL (       1 calls)
 forces       :     16.68s CPU     16.68s WALL (       1 calls)
 stress       :    209.17s CPU    209.17s WALL (       1 calls)

 Called by init_run:
 wfcinit      :     68.98s CPU     68.98s WALL (       1 calls)
 potinit      :      4.75s CPU      4.75s WALL (       1 calls)

 Called by electrons:
 c_bands      :   3000.94s CPU   3000.94s WALL (      23 calls)
 sum_band     :    192.26s CPU    192.26s WALL (      23 calls)
 v_of_rho     :      4.41s CPU      4.41s WALL (      24 calls)
 mix_rho      :      6.72s CPU      6.72s WALL (      23 calls)

 Called by c_bands:
 init_us_2    :      2.12s CPU      2.12s WALL (      47 calls)
 cegterg      :   2994.88s CPU   2994.88s WALL (      23 calls)

 Called by *egterg:
 h_psi        :    940.26s CPU    940.26s WALL (      70 calls)
 g_psi        :     30.53s CPU     30.53s WALL (      46 calls)
 cdiaghg      :   1223.83s CPU   1223.83s WALL (      69 calls)

 Called by h_psi:
 add_vuspsi   :     78.78s CPU     78.78s WALL (      70 calls)

 General routines
 calbec       :     65.14s CPU     65.14s WALL (      72 calls)
 fft          :      9.65s CPU      9.65s WALL (     271 calls)
 ffts         :      2.55s CPU      2.55s WALL (     474 calls)
 fftw         :    894.47s CPU    894.51s WALL (   75284 calls)
 davcio       :     32.45s CPU     32.45s WALL (      23 calls)

 Parallel routines
 fft_scatter  :    284.51s CPU    284.65s WALL (   76029 calls)
 ALLTOALL     :     61.81s CPU     61.82s WALL (   75272 calls)
 EXX routines

 PWSCF        :     1h 1m CPU       1h 1m WALL
```

- Profiling tools

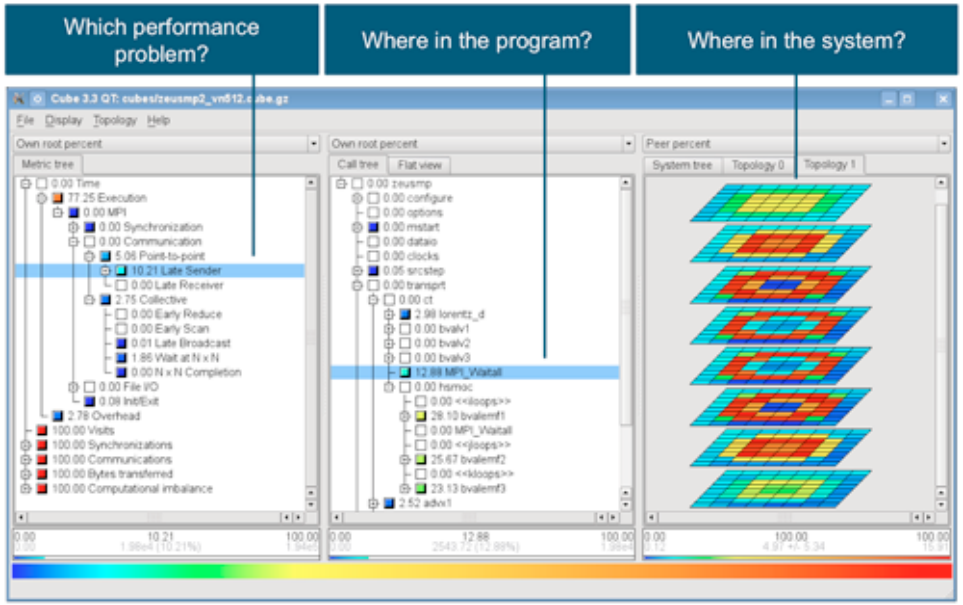

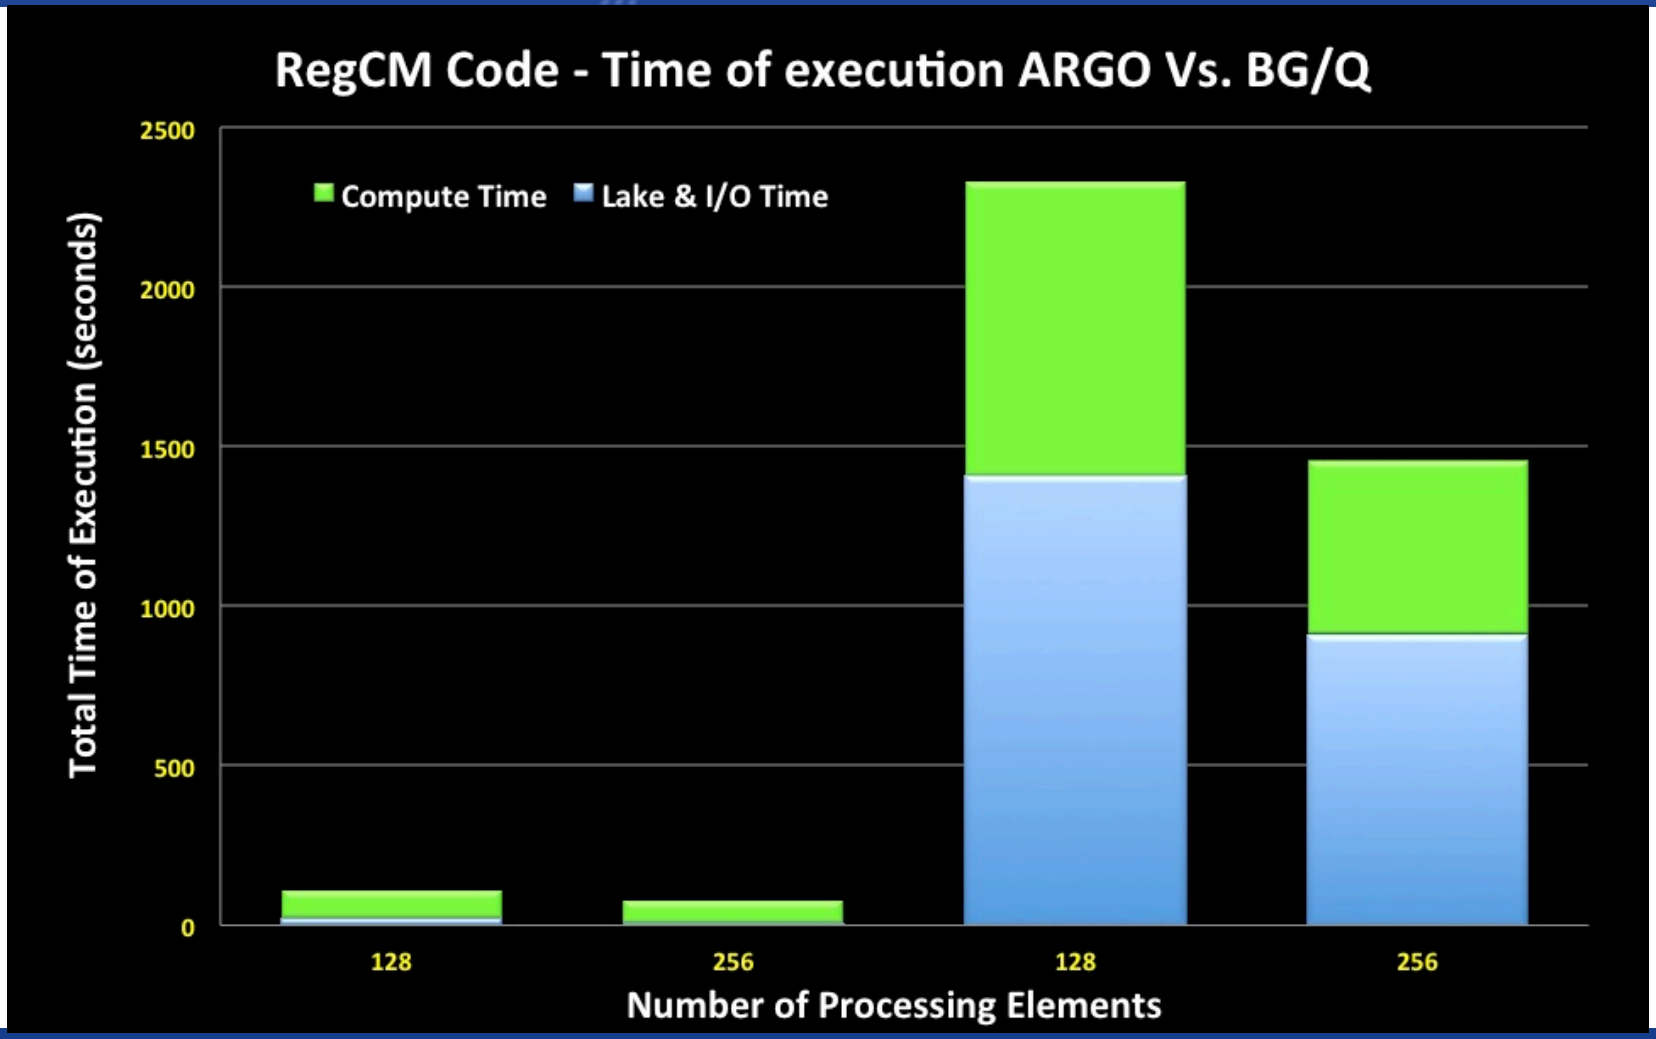
RegCM Code - Time of execution ARGO Vs. BG/Q
Compute Time
Lake & I/O Time
Total Time of Execution (seconds)
2500
2000
1500
1000
500
0
128
256
128
256
Number of Processing Elements

Thanks for Your Attention!

epcc

Cinvestav

CINVESTAV

INININ

UNIVERSIDAD AUTÓNOMA METROPOLITANA
Unidad Azcapotzalco

CONACYT

MCTP
Mesoamerican Centre for Theoretical Physics

CYTED

CADING
RED CYTED | Computación de Alto Desempeño en Ingeniería