

SEPTEMBER 21-23, BORDEAUX

2015

EURO



MPI



HPC and combustion

G. Staffelbach

A. Dauptain, E. Riber, O. Vermorel, L. Gicquel, B. Cuenot, F.
Duchaine, J. Dombard
and lots of Phds...

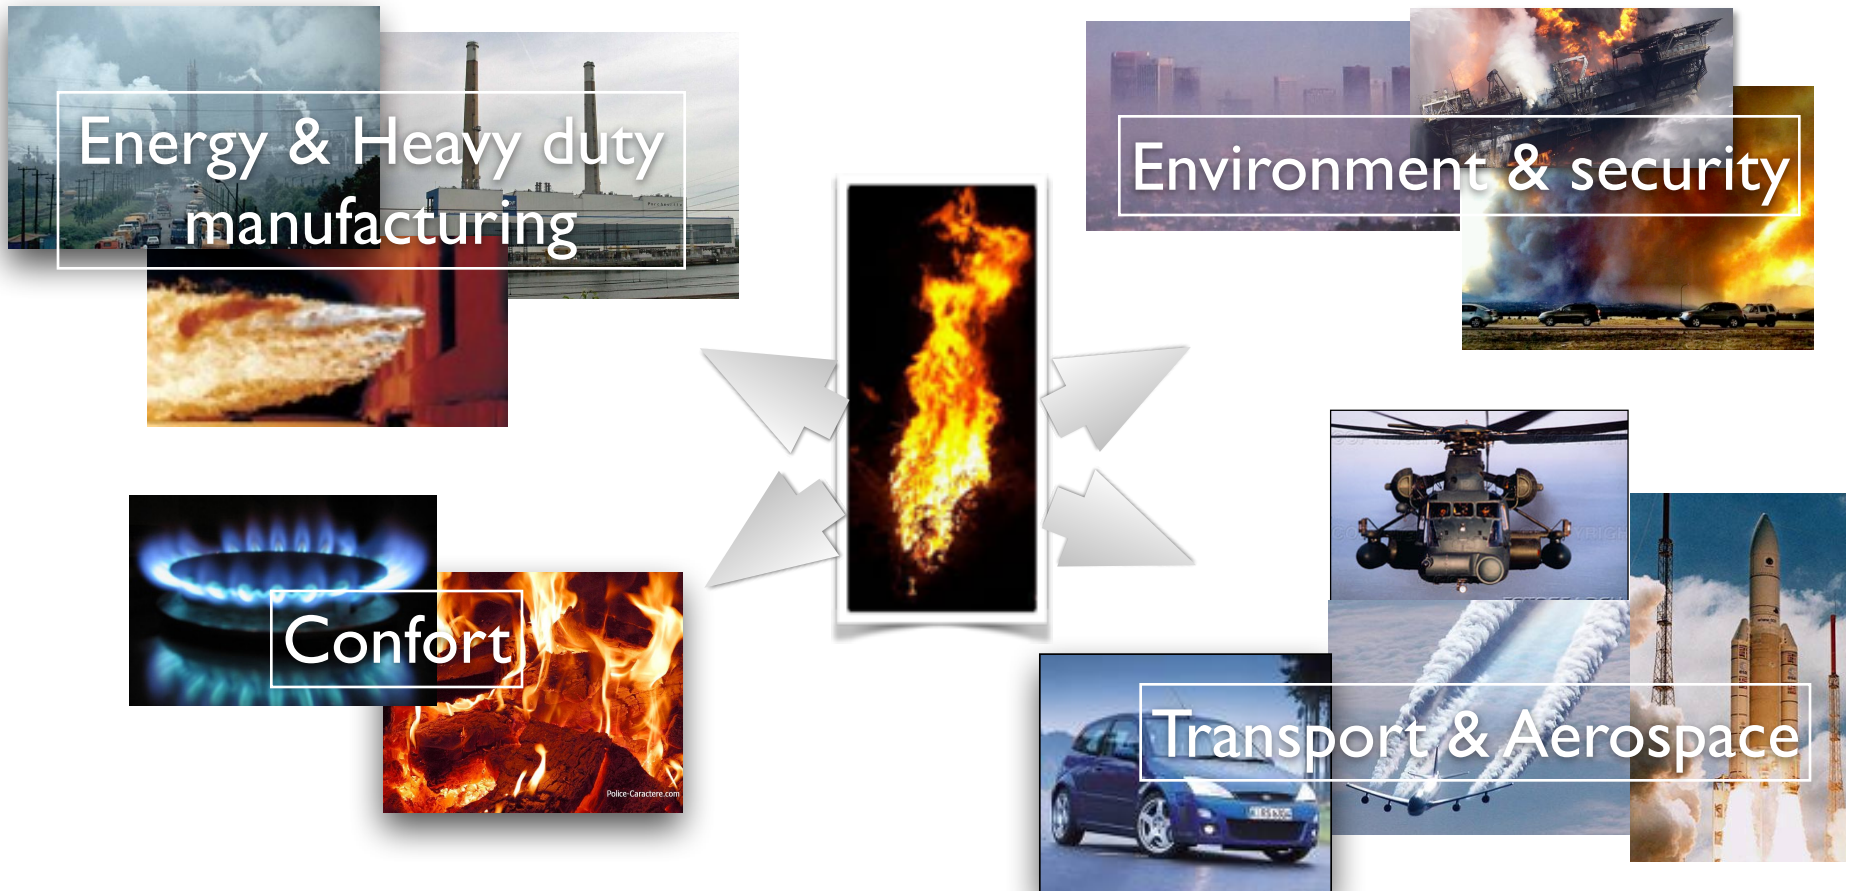


CERFACS

- ➔ non-profit research center in Toulouse France
- ➔ 150 people :
 - Climate change
 - Sparse matrix algorithms
 - Atmospheric pollution
 - Computational Fluid Dynamics
 - Aerodynamics
 - Turbomachinery
 - Combustion

Context

Combustion: An engineering science at the cross-road between chemistry & fluid mechanics with strong technological / industrial and societal implications





Context

➔ Pollution and climate change definitely will ...





Context




- ➔ Definite need to optimise combustion processes

- ➔ Via new technologies:
 - New fuels ?
 - New materials ?
 - New operating conditions ?


- ➔ Optimizing combustion is a priority but ..

Context

- ➔ New conditions can lead to new problems ...
- ◆ One of the usual problems encountered during optimization of combustion systems is combustion ***instabilities*** (« thermoacoustics » in combustion)



Context




Liquid rocket engine (NASA 1957)



Liquid rocket engine (NASA 1963)

What is a combustion instability ?

→ Very simple experiment (40\$)



+



What is a combustion instability ?



What is a combustion instability ?



What is a combustion instability ?



What is a combustion instability ?



What is a combustion instability ?




What is a combustion instability ?






What is a combustion instability ?

- Same experiment with glass walls (Dr Durox, EM2C Paris): 2000 \$ experiment (quartz)



STABLE




UNSTABLE




What is a combustion instability ?

- Same experiment with glass walls (Dr Durox, EM2C Paris): 2000 \$ experiment (quartz)



STABLE




UNSTABLE




What is a combustion instability ?

- Same experiment with glass walls (Dr Durox, EM2C Paris): 2000 \$ experiment (quartz)




STABLE



UNSTABLE


What is a combustion instability ?

EM2C Laboratory
ECP/CNRS



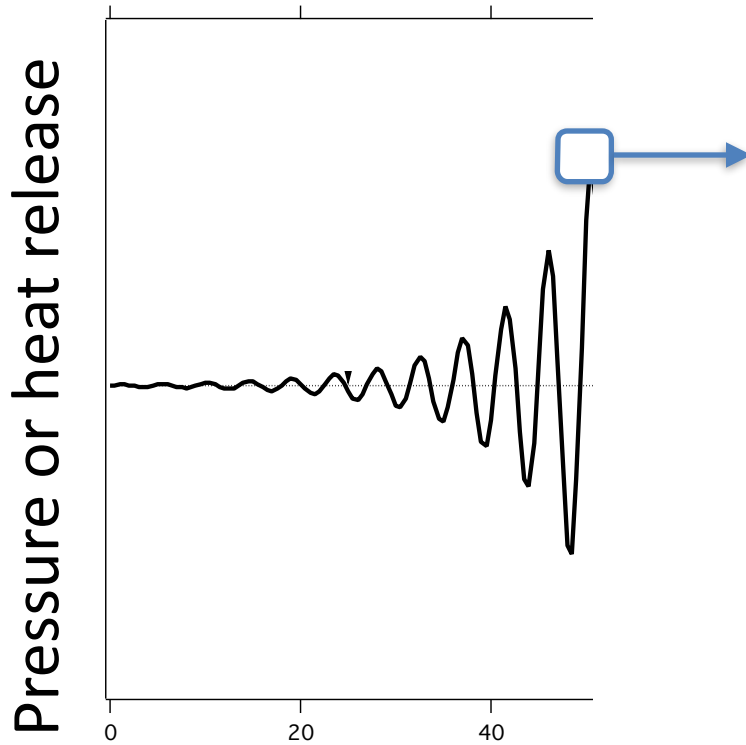
What is a combustion instability ?

EM2C Laboratory
ECP/CNRS



What is a combustion instability ?

➔ If oscillation attain a high enough level ...






CERFAC Tools

THEORY


ACOUSTIC TOOLS

HIGH FIDELITY SIMULATIONS

$$f = \frac{c^0}{2L_c} + \frac{c^{0T0}}{\pi L_c}$$




AVSP



DAYS

AVBP



YEARS

COMPLEXITY & CPU COST

THEORY

ATACAMAC

**NETWORK
MODEL**

**HELMHOLTZ
SIMULATION**

SECTOR LES

LES 360°

ANALYTICAL

1,5 D

3 D Fourier space

3 D LES Time domain




CERFAC Tools

THEORY

ACOUSTIC TOOLS

HIGH FIDELITY SIMULATIONS


$$f = \frac{c^0}{2L_c} + \frac{c^0 \Gamma_0}{\pi L_c}$$



PARALLEL TOOLS

AVSP


AVBP



DAYS



YEARS



COMPLEXITY & CPU COST

THEORY

ATACAMAC

**NETWORK
MODEL**

**HELMHOLTZ
SIMULATION**

SECTOR LES

LES 360°

ANALYTICAL

1,5 D

3 D Fourier space







3 D LES Time domain



The AVBP code



- ➔ Developed by CERFACS and IFP-EN,
- ➔ External/internal flows,
- ➔ Fully compressible turbulent reacting flows (ideal & real gases),
- ➔ DNS/LES approaches,
- ➔ Unstructured hexahedral, tetrahedral, prisms & hybrid meshes,
- ➔ Massively parallel,
- ➔ C/Fortran languages,
- ➔ SPMD approach.
- ➔ Multi-phase solvers (Lagrangian & Eulerian)



2013 PRACE Scientific Annual Report 'success story'

'Most innovative industrial HPC solution in Europe' in ISC'2013

An open science project

RESEARCH ENTITIES

EM2C (Centrale Paris)
IMFT (Toulouse)
CORIA (Rouen)
ONERA
TU Munich
U. Twente
Von Karman institute
U. Sherbrooke
CIEMAT Madrid
ETH Zurich
Gent University

INDUSTRY

SNECMA
TURBOMECA
TOTAL
SNECMA Vernon
RENAULT
PSA
ALSTOM
ANSALDO
HONEYWELL
SIEMENS
AIR LIQUIDE
GDF

CERFACS



HPC centers
(IDRIS, CINES, BSC, GENCI,
CEA TGCC,
ARNL, ORNL, JSC)


INTEL IPCC
DEEP FP7
EU project

- G. Statterbach



CPU capacity
[Flop/s]

Problem
size [DoF]



CERFACS

2005

EURO MPI 2015 - G. Staffelbach

2010


2015

2020

10^4

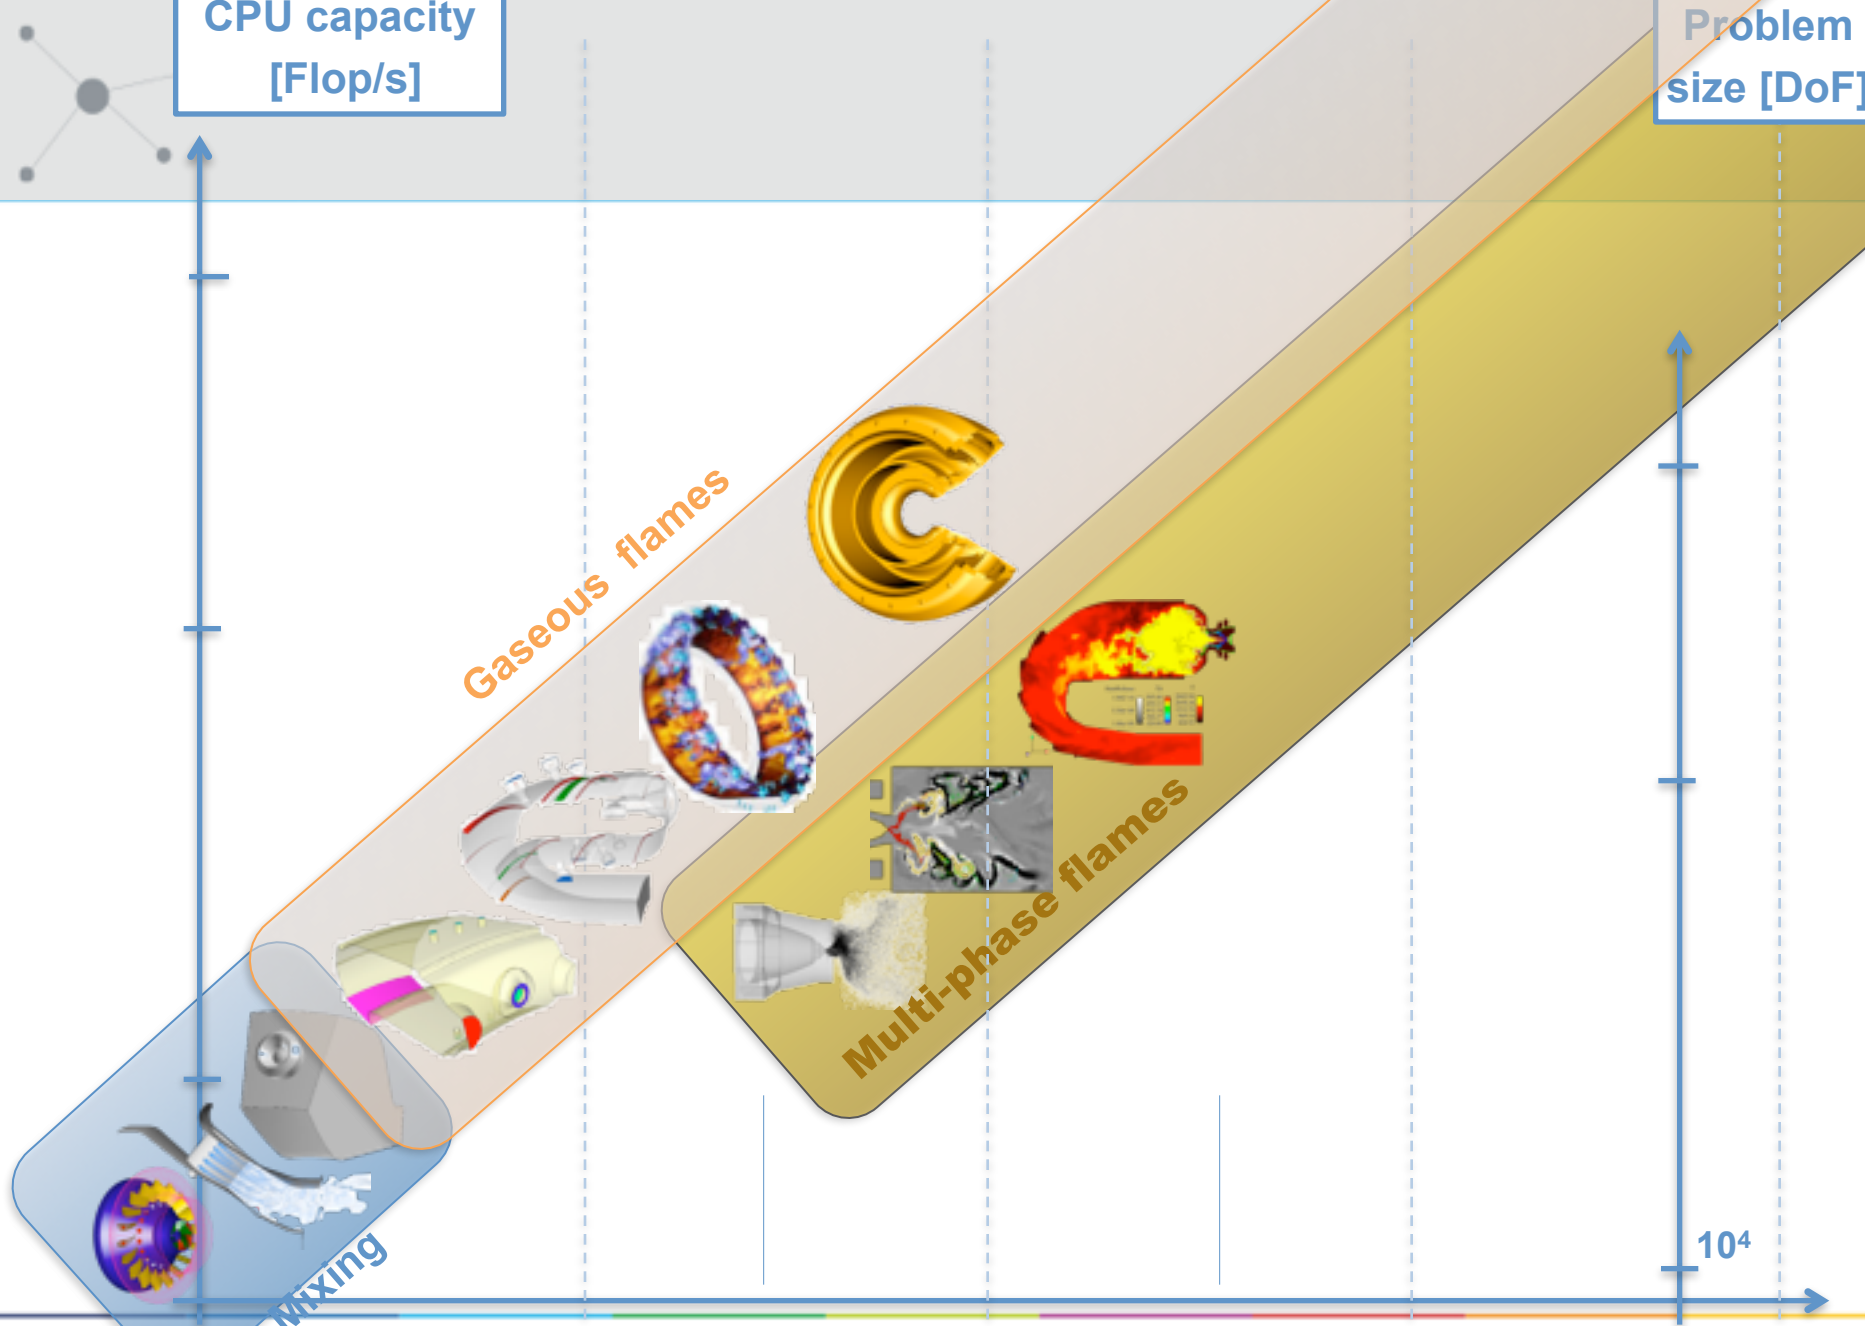
CPU capacity
[Flop/s]

Problem
size [DoF]



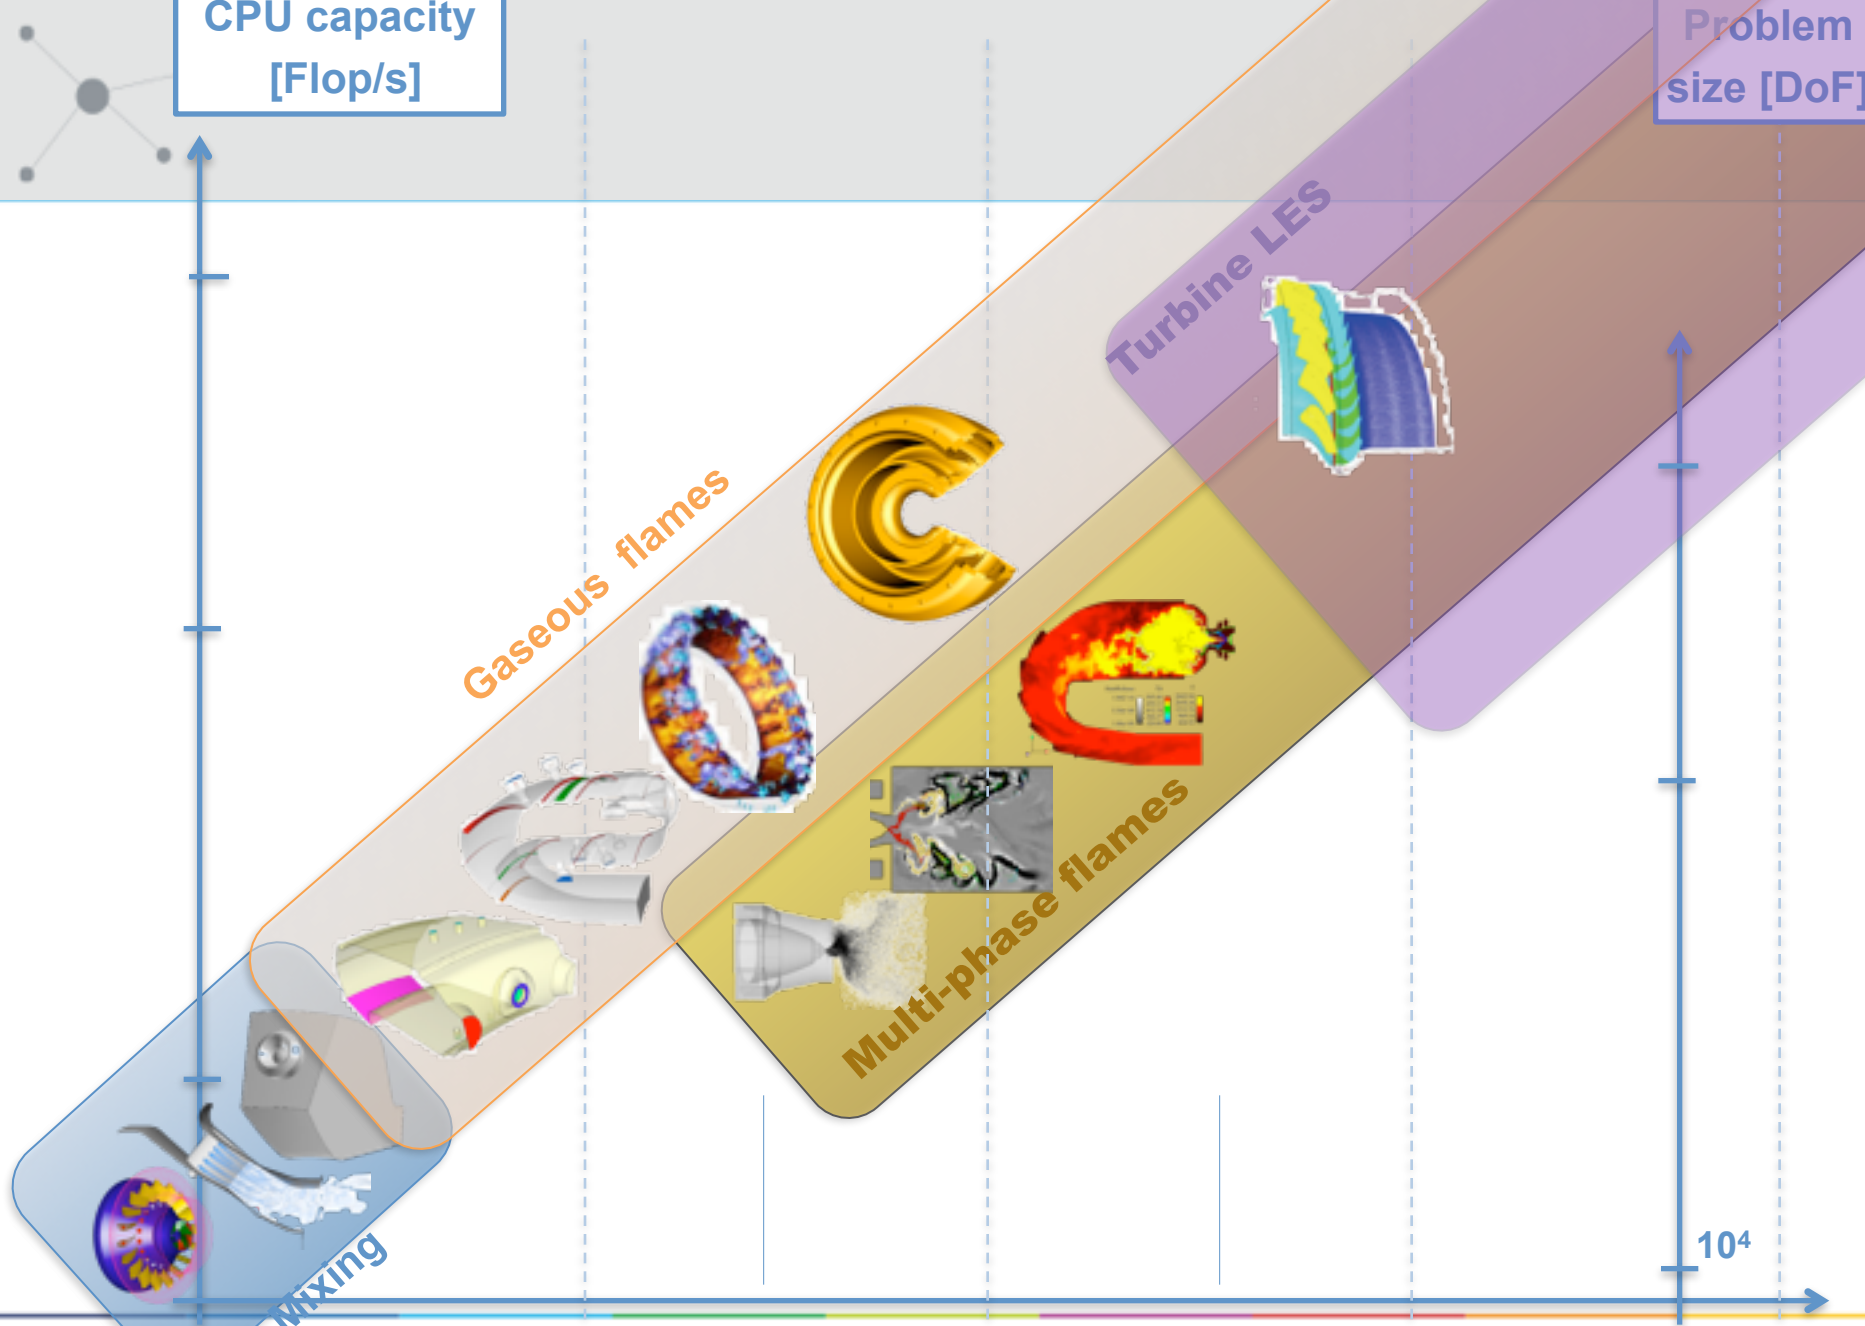
CPU capacity
[Flop/s]

Problem
size [DoF]



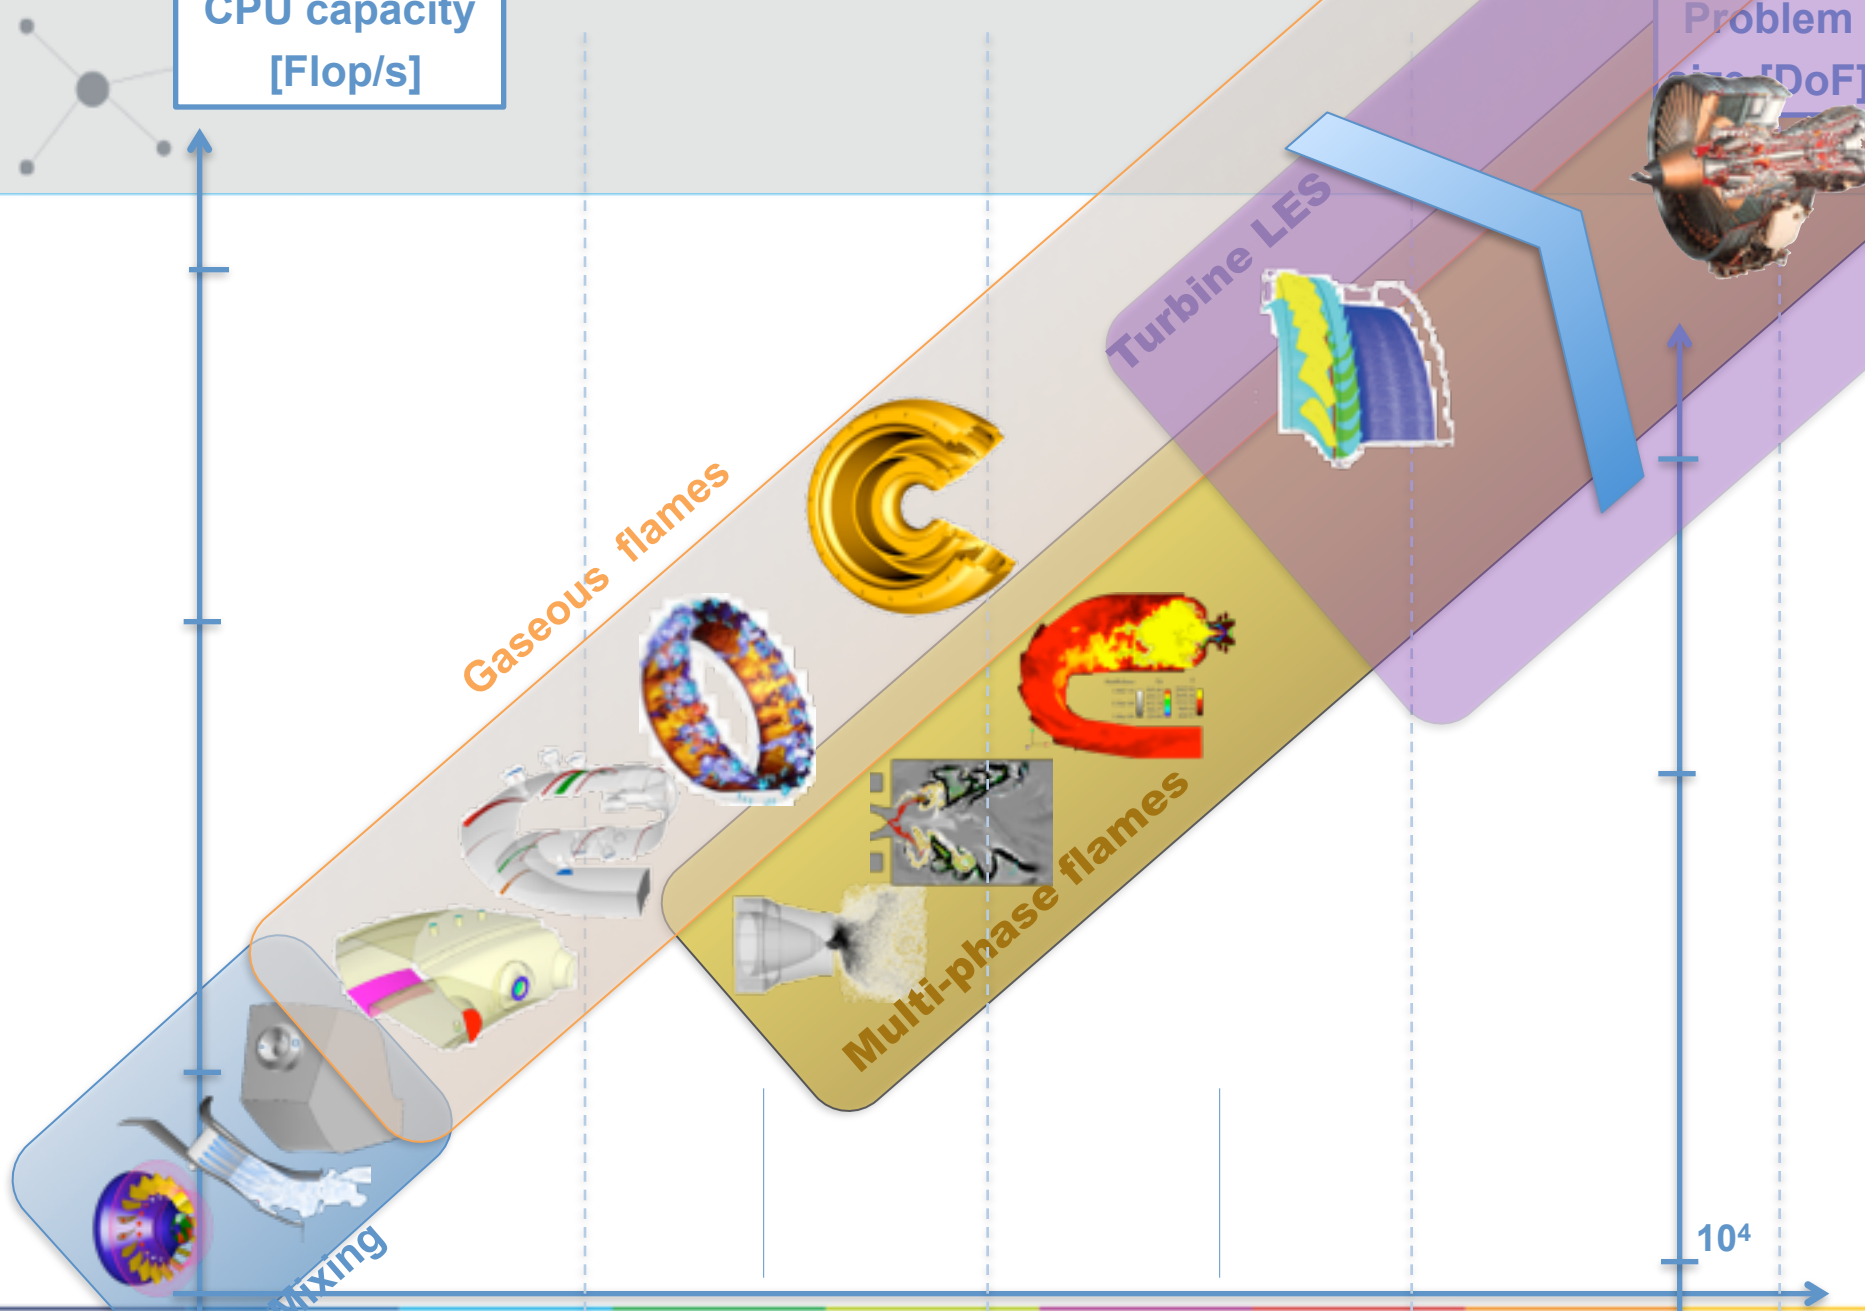
CPU capacity
[Flop/s]

Problem
size [DoF]



CPU capacity
[Flop/s]

Problem
size [DoF]




CPU capacity
[Flop/s]

Problem
size [DoF]




Flames

Turbine LES



Multi-phase flames



Mixing

CERFACS

2005

2010

2015

2020

10⁴



CPU capacity
[Flop/s]

Problem
size [DoF]





Flames

Turbine LES



Multi-phase flames

Multi-physics



Mixing

10⁴

CERFACS

2005

EURO MPI 2015 - G. Staffelbach

2010

2015


2020

CPU capacity
[Flop/s]


Problem
size [DoF]




Flames



Multi-phase flames



Multi-physics



Turbine LES



Transfer time
to INDUSTRY


10⁴

CPU capacity
[Flop/s]


Problem
size [DoF]



Flames



Multi-phase flames



Multi-physics

Turbine LES



Transfer time
to INDUSTRY

4 MPI tasks

2005

EURO MPI 2015 - G. Staffelb

2010

262144 MPI tasks

2020

10⁴

18

"We acknowledge Intel for its support through the Intel Parallel Computing Center program"

This research used resources of the Argonne Leadership Computing Facility at Argonne National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under contract DE-AC02-06CH11357.

"We acknowledge PRACE for awarding us access to resource CURIE based in France at *Très Grand Centre de calcul du CEA (TGCC)*."


"We acknowledge PRACE for awarding us access to resource JUQUEEN based in Germany at *Jülich Supercomputing Centre (JSC)*."

"We acknowledge PRACE for awarding us access to resource HERMIT based in Germany at *High Performance Computing Center Stuttgart (HLRS)*."

"We acknowledge PRACE for awarding us access to resource FERMI based in Italy at *CINECA*."

Combustion instabilities in Gas turbines


- ➔ From experiment to real turbines
- ➔ Use LES to 'predict' instabilities



Wolf et al Comb. Flame, 159: 3398-3413, 2012


Gicquel et al Progress in Energy and Combustion Science. 38, 782-817. 2012.

Parmentier et al Combustion and Flame 159, 7, July 2012, 2374-2387.



Combustion instabilities in Gas turbines


- ➔ From experiment to real turbines
- ➔ Use LES to 'predict' instabilities



Wolf et al Comb. Flame, 159: 3398-3413, 2012

Gicquel et al Progress in Energy and Combustion Science. 38, 782-817. 2012.


Parmentier et al Combustion and Flame 159, 7, July 2012, 2374-2387.




LES of full 360 Gas Turbine

→ LES of a gas turbine is able to predict the azimuthal combustion instability

10 million CPU hours Bluegene P
32768 MPI tasks



Pressure




Temperature


LES of full 360 Gas Turbine

→ LES of a gas turbine is able to predict the azimuthal combustion instability

10 million CPU hours Bluegene P
32768 MPI tasks



Pressure





Temperature

38.36000 ms

Illustration in the lab

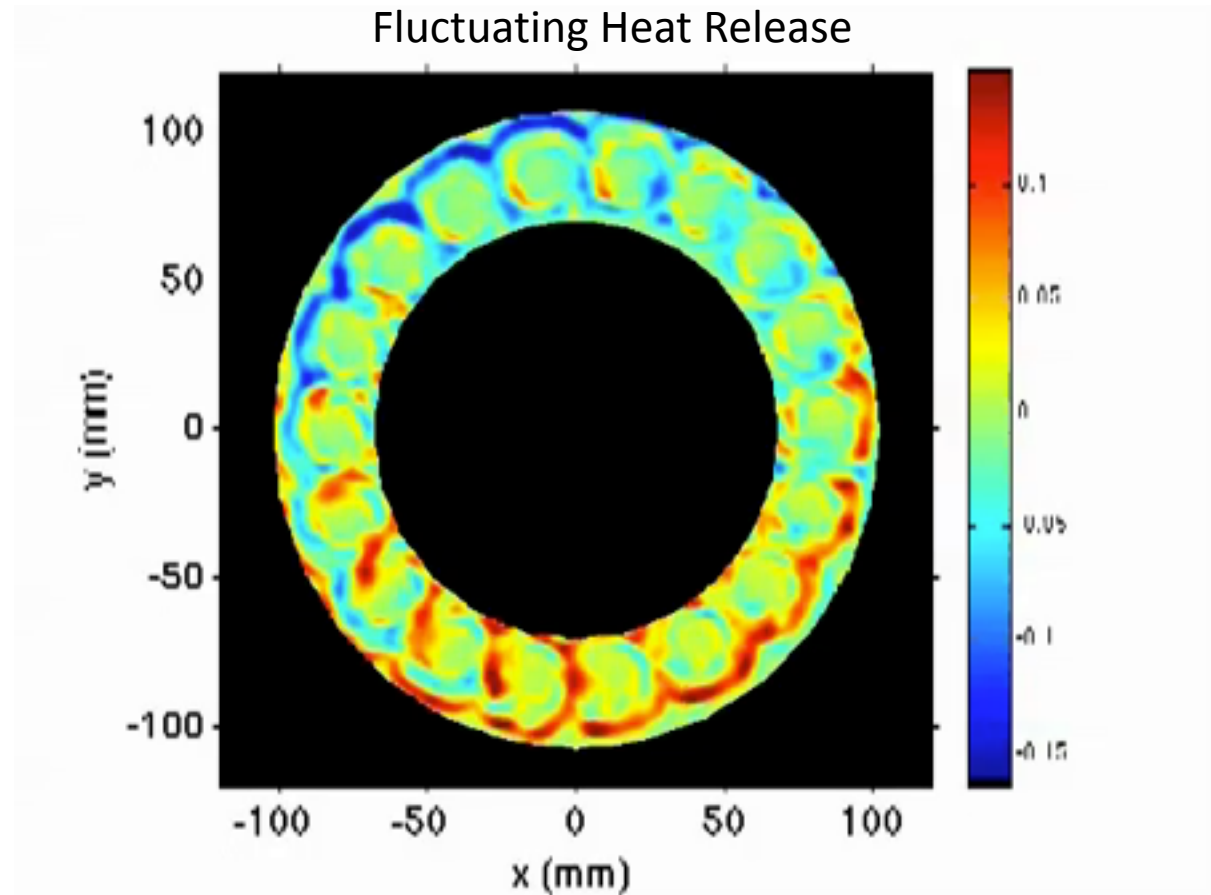

→ DAWSON Experiment (Cambridge 2011)



N.A. Worth, J.R. Dawson, Proc. Combust. Inst. (2012), [http:// dx.doi.org/10.1016/j.proci.2012.05.061](http://dx.doi.org/10.1016/j.proci.2012.05.061)


Illustration in the lab

→ DAWSON Experiment (Cambridge 2011)



N.A. Worth, J.R. Dawson, Proc. Combust. Inst. (2012), [http:// dx.doi.org/10.1016/j.proci.2012.05.061](http://dx.doi.org/10.1016/j.proci.2012.05.061)


Stability prediction



- ➔ Increasing the fuel consumption rate reduces the delay and stabilised the system

wolf et al.


Stability prediction




- ➔ Increasing the fuel consumption rate reduces the delay and stabilised the system

wolf et al.

Stability prediction



- ➔ Increasing the fuel consumption rate reduces the delay and stabilised the system





wolf et al.

Using simulation to assist combustor design : Ignition

Efficient ignition is paramount for safety and economic reasons

Where to add the energy to start the flame?

How to ensure burner to burner flame propagation ?





Using simulation to assist combustor design : Ignition

Efficient ignition is paramount for safety and economic reasons






Where to add the energy to start the flame?

How to ensure burner to burner flame propagation ?



Using simulation to assist combustor design : Ignition

How many burner ? Ignition spark position ?



SP9: L = 90mm






SP16: L = 160mm

SP26: L = 260mm

EXPERIMENTS ! B. Renoud et al CORIA

Using simulation to assist combustor design : Ignition

How many burner ? Ignition spark position ?



SP9: L = 90mm






SP16: L = 160mm

SP26: L = 260mm

EXPERIMENTS ! B. Renoud et al CORIA

Using simulation to assist combustor design : Ignition

How many burner ? Ignition spark position ?



SP9: L = 90mm






SP16: L = 160mm

SP26: L = 260mm

EXPERIMENTS ! B. Renoud et al CORIA

Using simulation to assist combustor design : Ignition

How many burner ? Ignition spark position ?



SP9: L = 90mm

SP16: L = 160mm

SP26: L = 260mm

EXPERIMENTS ! B. Renoud et al CORIA

Using simulation to assist combustor design : Ignition

SP9: L = 90mm

SP26: L = 260mm



Time = 36.8 ms



Time = 36.8 ms



Radial flame propagation

Axial flame propagation

D. barre L. Esclapez

8192 MPI tasks BG Q IDRIS - 10M cpu hours

Using simulation to assist combustor design : Ignition

SP9: L = 90mm

SP26: L = 260mm



Time = 36.8 ms



Time = 36.8 ms



Radial flame propagation

Axial flame propagation

D. barre L. Esclapez



8192 MPI tasks BG Q IDRIS - 10M cpu hours

Ignition on an annular burner

15 million hours CURIE TTGC GENCI


E. Riber, B. Cuenot, F. Duchaine (CERFACS), R. Vicquelin, M. Boileau, M. Philip, T. Schmitt, S. Candel (EM2C)

Simulation of Ignition in a Multiple Annular Combustor injector and comparison with experiments







E. Riber, B. C...
Boileau, M. P...
Simulation of...
injector and c...



CNRS, ONERA, EDF, TOTAL, AIRBUS, STGC, GENCI

elin, M.
stor




Using simulation for safety applications

From LES to DNS of explosions in semi-confined domains

Understand the physical phenomena involved in confined space explosions and validate the methodology for multiple scales

Sydney experiment, Masri et al

25cm length
LES and DNS (1Be elements)



P. Quillatre et al

D. Barré et al



2013 - 20M BG P


2014 - 86M BG Q

CERFACS

EURO MPI 2015 - G. Staffelbach

Using simulation for safety applications

Time: 11.91 ms



P. Quillatre et al
From LES to DNS of explosions in semi-confined domains
D. Barré et al

2013 - 20M BG P

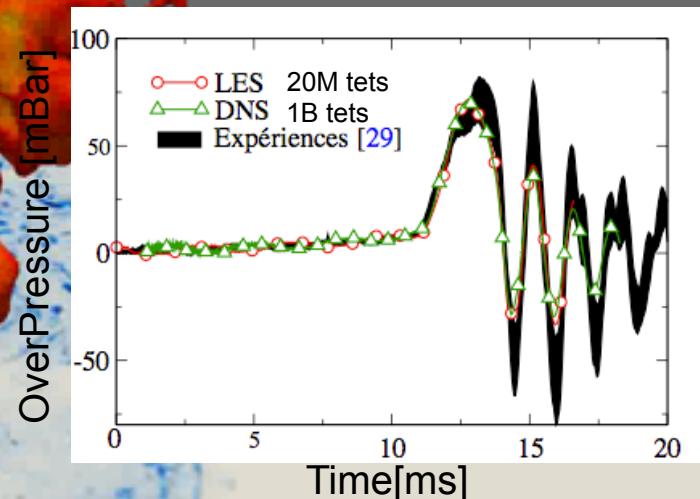
2014 - 86M BG Q



EURO MPI 2015 - G. Staffelbach

Using simulation for safety applications

Time: 11.91 ms



P. Quillatre et al
From LES to DNS of explosions in semi-confined domains

D. Barré et al

2013 - 20M BG P


2014 - 86M BG Q




EURO MPI 2015 - G. Staffelbach

Using simulation for safety applications

Large Eddy Simulation of the 1.5m configuration versus experiment



Experiment performed by Gexcon




Large Eddy Simulation

P. Quillatre et al


Time: 58.0

Using simulation for safety applications

Large Eddy Simulation of the 1.5m configuration versus experiment



Experiment performed by Gexcon




Large Eddy Simulation


P. Quillatre et al

Time: 58.0

The 10 M\$ failure(s)



Liquid rocket engine (NASA 1957)




Liquid rocket engine (NASA 1963)



Combustion instabilities in Rocket engines

→ DLR LAMPOLDSHAUSEN HF7 BKD

Full engine: 42 coaxial injectors
Cryogenic O₂/H₂ propellants
Pressure range: 50-80 bar





sizes in mm



A. Urbano et al. IMFT / EM2C / CERFACS



power ~
100 MW

80 M hours on FERMI (cineca IT) , 9th Prace CALL 40k Euros Simulation !



A. Urbano et al. IMFT / EM2C / CERFACS

80 M hours on FERMI (cineca IT) , 9th Prace CALL 40k Euros Simulation !



A. Urbano et al. IMFT / EM2C / CERFACS

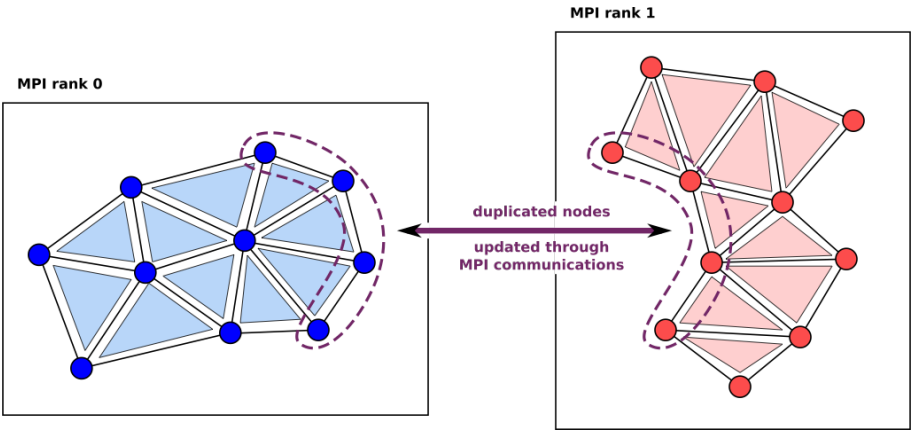





HPC and Combustion


- ➔ These highly complex systems require large amount of computing power and large simulation times.
- ➔ At CERFACS we strive to keep the code as simple and portable as possible while trying to optimise the code as best we can on multiple architecture ..
- ➔ However sometimes physics and HPC do not get along ...

Parallelisation

→ AVBP parallelism relies on MPI non blocking ISEND IRECV with a single node overlapping domain decomposition

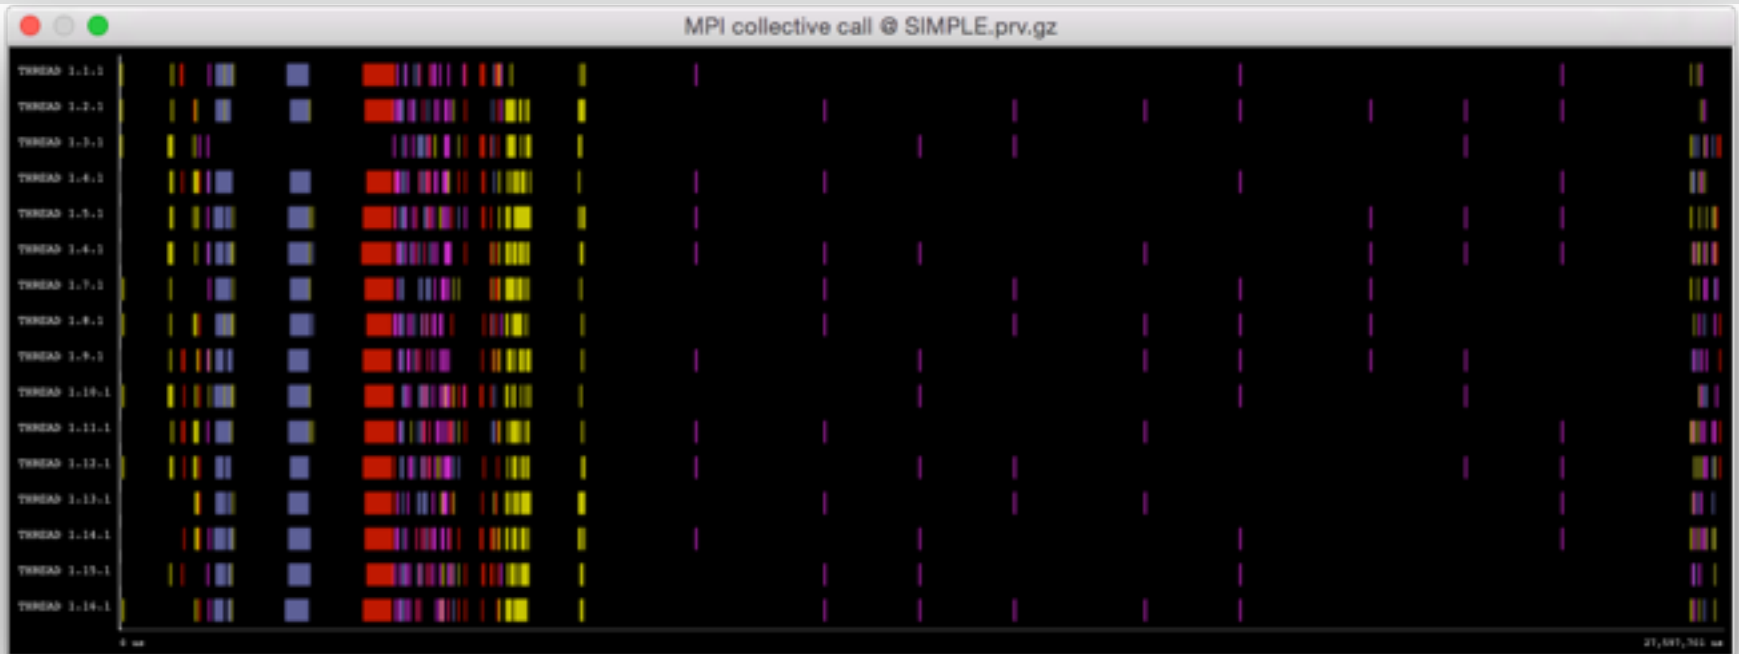


Parallel performance








Collective Calls

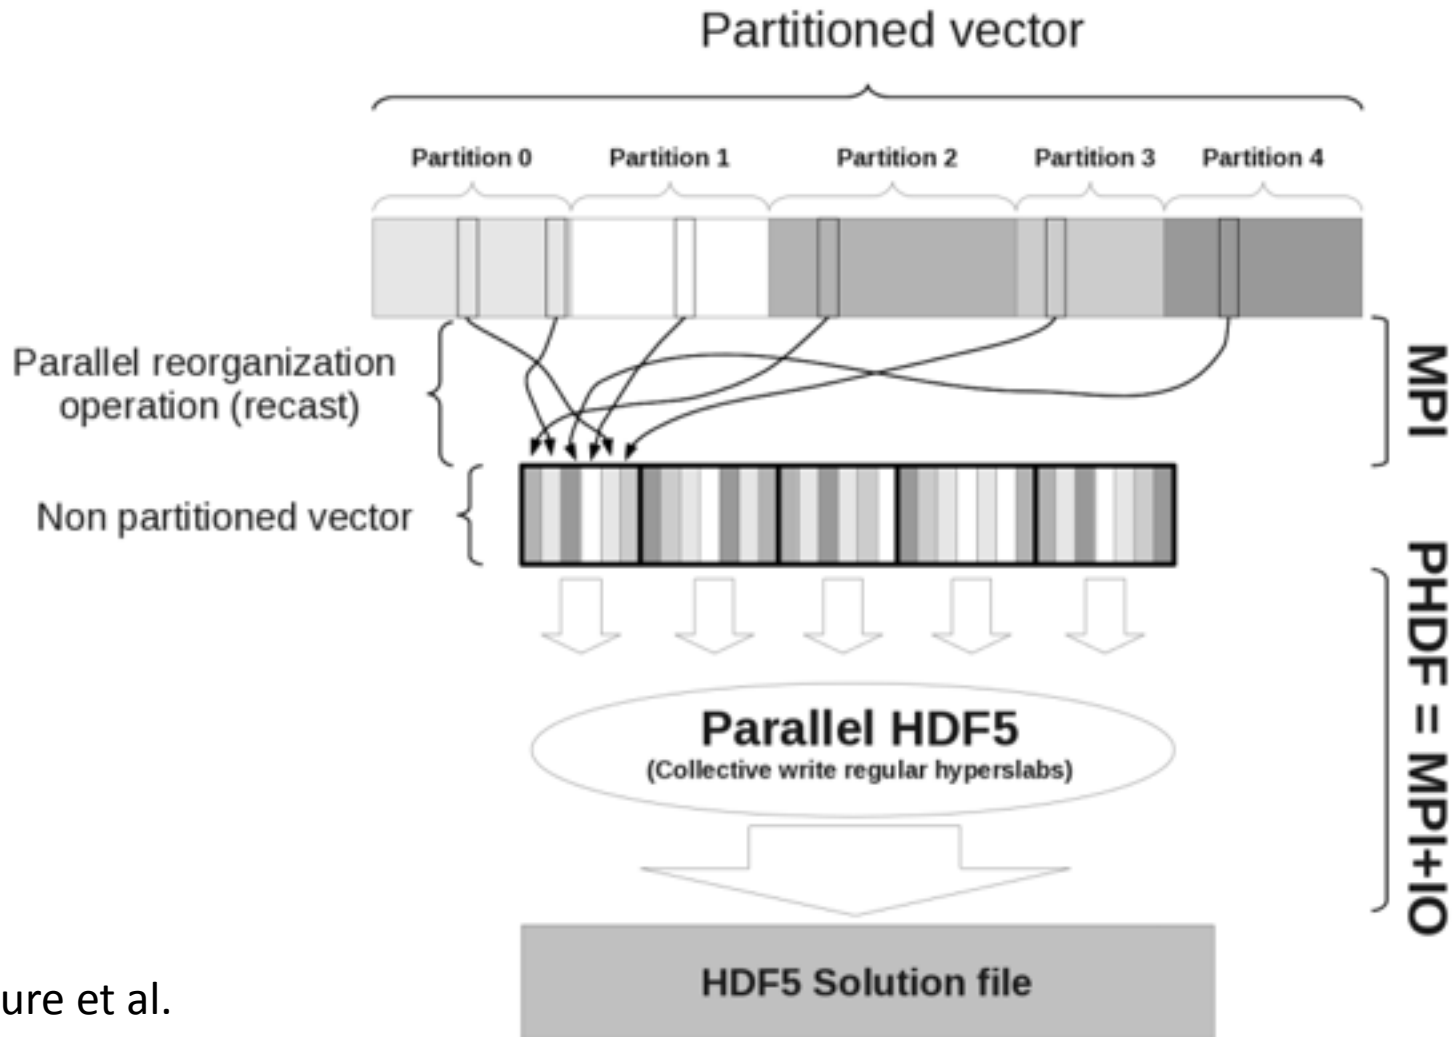


-  MPI_BCAST
-  MPI_BARRIER
-  MPI_ALLTOALL
-  MPI_ALLREDUCE

→ High usage of collectives

-  Physics
-  Monitoring
-  I/O

Parallel I/O



Jaure et al.



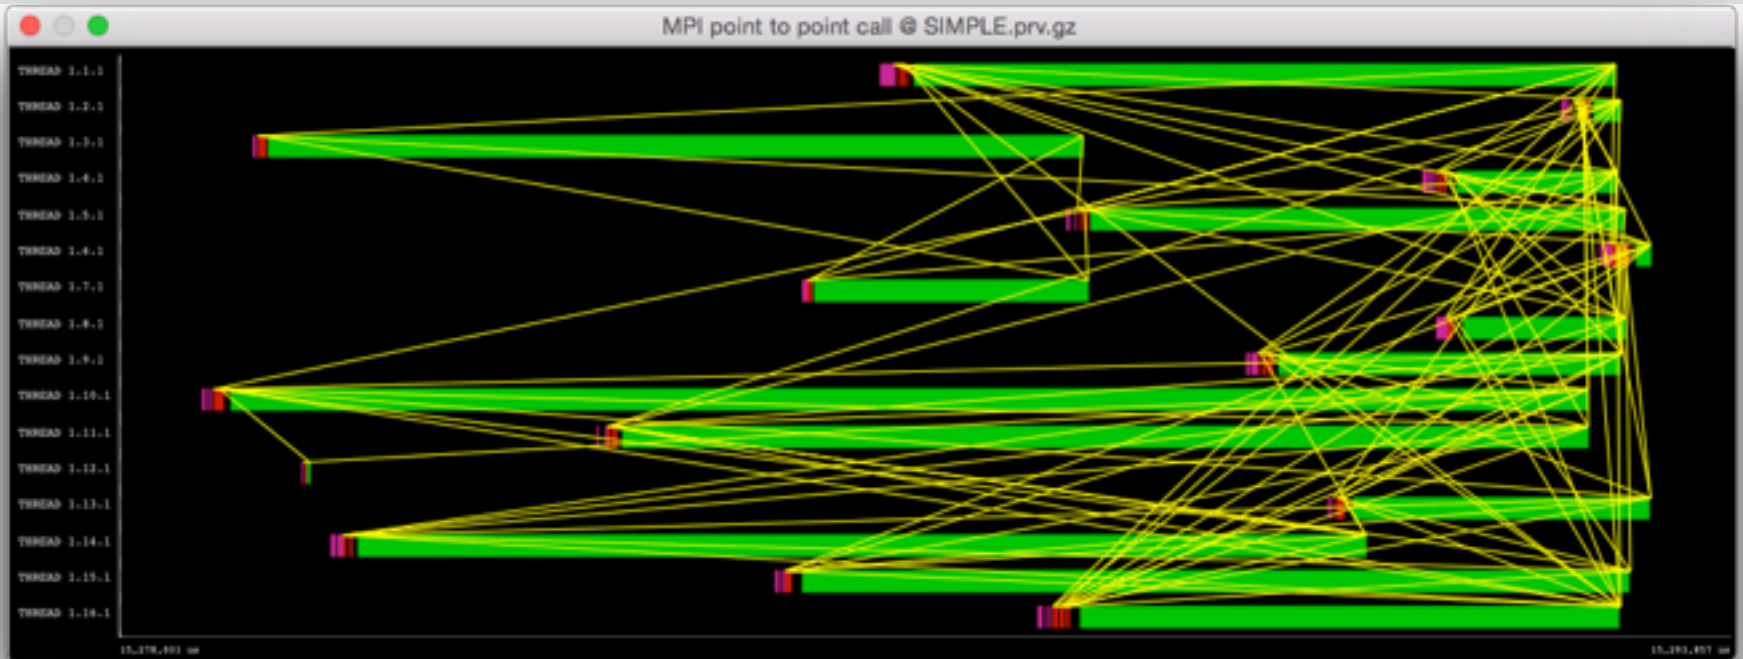
Point to Point



-  MPI_WAITALL
-  MPI_IRecv
-  MPI_Isend



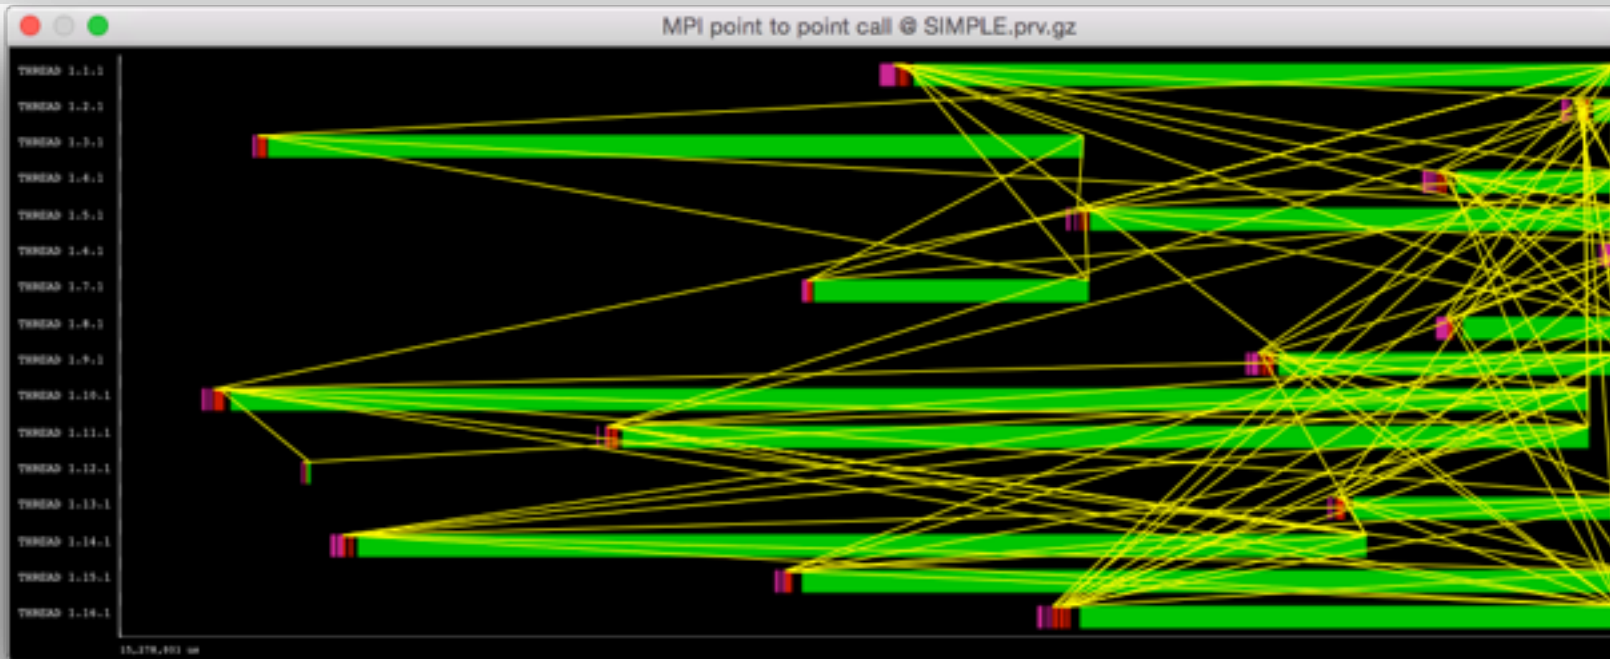
Point to Point



-  MPI_WAITALL
-  MPI_IRECV
-  MPI_ISEND



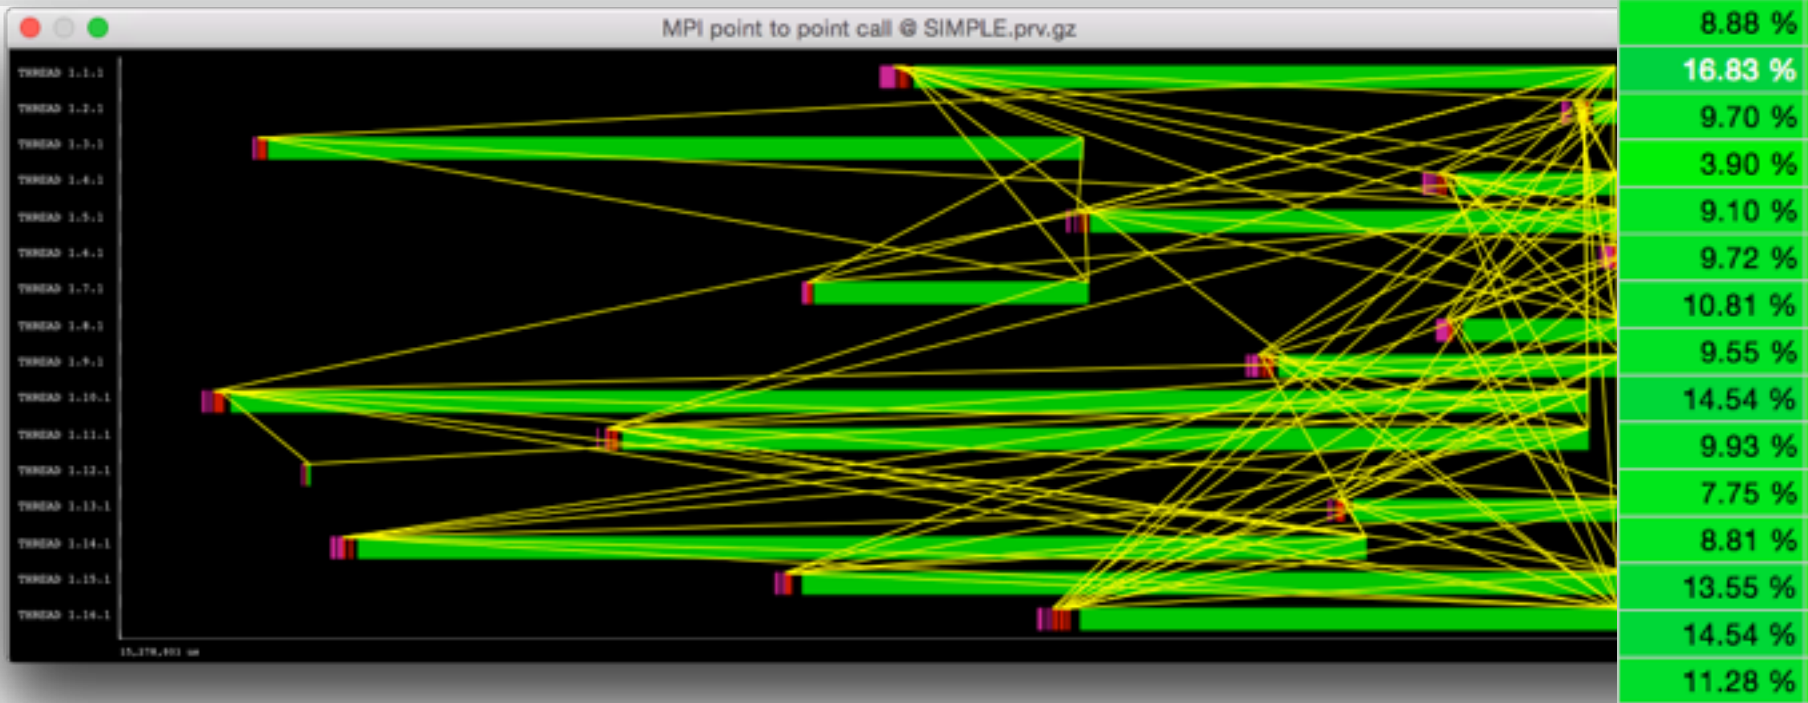
Point to Point



MPI_Waitall
16.37 %
8.88 %
16.83 %
9.70 %
3.90 %
9.10 %
9.72 %
10.81 %
9.55 %
14.54 %
9.93 %
7.75 %
8.81 %
13.55 %
14.54 %
11.28 %

- MPI_WAITALL
- MPI_IRecv
- MPI_Isend


Point to Point



- Small communication pattern
 - High imbalance

Point to Point

- THREAD 1.1.1
- THREAD 1.2.1
- THREAD 1.3.1
- THREAD 1.4.1
- THREAD 1.5.1
- THREAD 1.6.1
- THREAD 1.7.1
- THREAD 1.8.1
- THREAD 1.9.1
- THREAD 1.10.1
- THREAD 1.11.1
- THREAD 1.12.1
- THREAD 1.13.1
- THREAD 1.14.1
- THREAD 1.15.1
- THREAD 1.16.1





MPI_Waitall
16.37 %
8.88 %
16.83 %
9.70 %
3.90 %
9.10 %
9.72 %
10.81 %
9.55 %
14.54 %
9.93 %
7.75 %
8.81 %
13.55 %
14.54 %
11.28 %

- MPI_WAITALL
- MPI_IRecv
- MPI_Isend

→ Small communication pattern
■ High imbalance


Balanced partition does not mean balance communications ...

➔ Simple Cylinder on 4 domains ...



rank	elements
1	18949
2	19011
3	18997
4	19011

➔ Simple C



PREVIOUS	
rank	elements
1	18949
2	19011
3	18997
4	19011

CONTIGUOUS	
rank	elements
1	38050
2	16336
3	16979
4	4603



Current Challenges

- ➔ Imbalance can be somewhat compensated by Communication/Computation overlap
 - Contiguous partitioning
 - Reduces the number of neighbours !
 - Is not 'fully balanced' : imbalanced vertices/task
- ➔ High usage of collective leads to radical differences depending on the implementations.
 - Collectives are extremely important for CFD applications
- ➔ Overlap of (MPI) I/O and computation



Current frontier Science : MPMD coupling/Multiphysics

- ➔ One code all the science ... Not realistic
(Optimisation, physical models , Complexity of maintaining the tool)
- ➔ Solution : MPMD code coupling
- ➔ Use optimize code for each physics and use a coupler to handle the interaction


- ➔ CERFACS and ONERA develop the Open source coupler

Current frontier Science : MPMD coupling/Multiphysics

- ➔ One code all the science ... Not realistic
(Optimisation, physical models , Complexity of maintaining the tool)
- ➔ Solution : MPMD code coupling
- ➔ Use optimize code for each physics and use a coupler to handle the interaction


- ➔ CERFACS and ONERA develop the Open source coupler **O-Palm**

Example of code coupling method for moving parts



- ➔ Characteristics are similar to the sliding mesh approach (Francois, ISABE 2011):
- ➔ Overlapping section width depends on numerical stencil 2nd and 3rd order interpolation available (Linear and Hermite-type)


Example of code coupling method for moving parts



- ➔ Characteristics are similar to the sliding mesh approach (Francois, ISABE 2011):
- ➔ Overlapping section width depends on numerical stencil 2nd and 3rd order interpolation available (Linear and Hermite-type)

Combustion chamber / Turbine Simulation

→ First LES of combustion and first turbine stage




- Possibility to perform each simulation separate before doing the coupled case
- Hand made load balancing between the two instances.

D. Papadogianis et al

Combustion chamber / Turbine Simulation

→ First LES of combustion and first turbine stage




- Possibility to perform each simulation separate before doing the coupled case
- Hand made load balancing between the two instances.

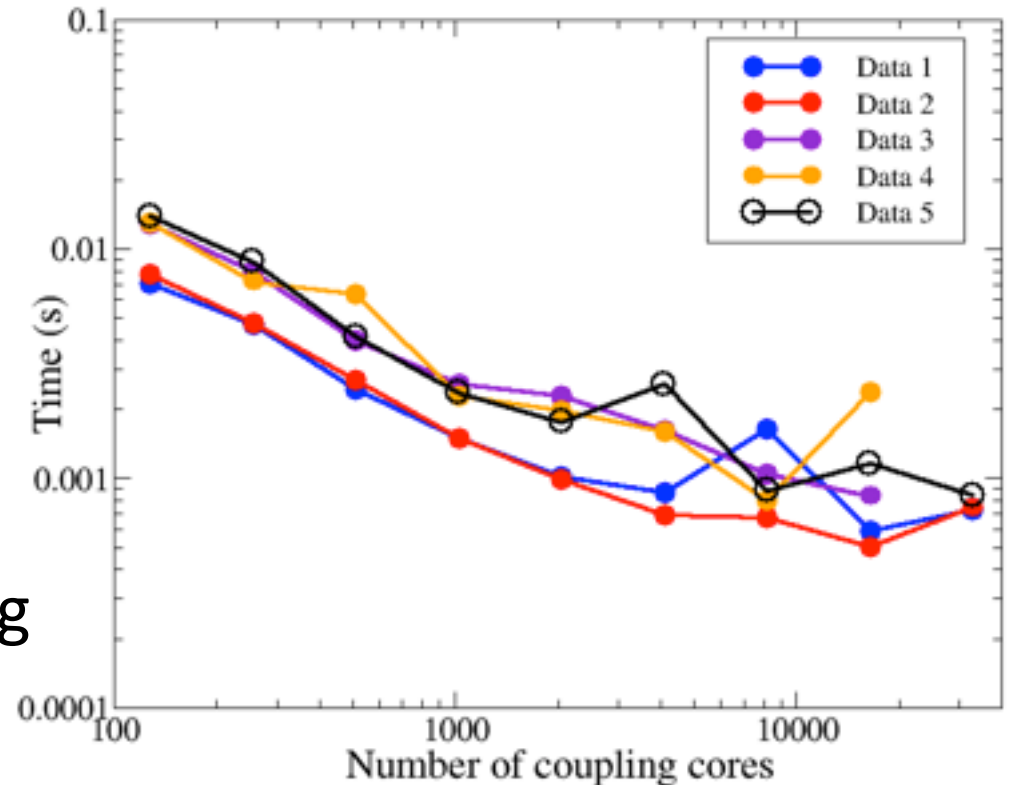
D. Papadogianis et al

Full compressor LES

Duchaine et al.




Exchanges time for TITAN CRAY (ORNL)



- ➔ Multiple stage coupling
- ➔ Two AVBP instances


Full compressor LES

Duchaine et al.



- ➔ Multiple stage coupling
- ➔ Two AVBP instances


Exchanges time for TITAN CRAY (ORNL)



Need for co-partitioning methods :


=> Partition instance 2 based on instance 1 to limit exchanges

Full compressor LES



Multi-Stage Gas Turbine Computation. Grand Challenge OCCIGEN. GENCI/CINES. 7680 Haswell Cores.
IDRIS Grand Challenge 6144 MPI tasks
J. Delaborde, F. Duchaine, L. Gicquel, G. Staffelbach

Full compressor LES




Multi-Stage Gas Turbine Computation. Grand Challenge OCCIGEN. GENCI/CINES. 7680 Haswell Cores.
IDRIS Grand Challenge 6144 MPI tasks
J. Delaborde, F. Duchaine, L. Gicquel, G. Staffelbach



Fluid / Solid thermal conduction

Florent DUCHAINE - CERFACS

AVBP & AVTP Coupling






Fluid / Solid thermal conduction

Florent DUCHAINE - CERFACS


AVBP & AVTP Coupling



Fun application: 3D Large Eddy Simulation of a booster trail.

A. Poubeau R. Paoli

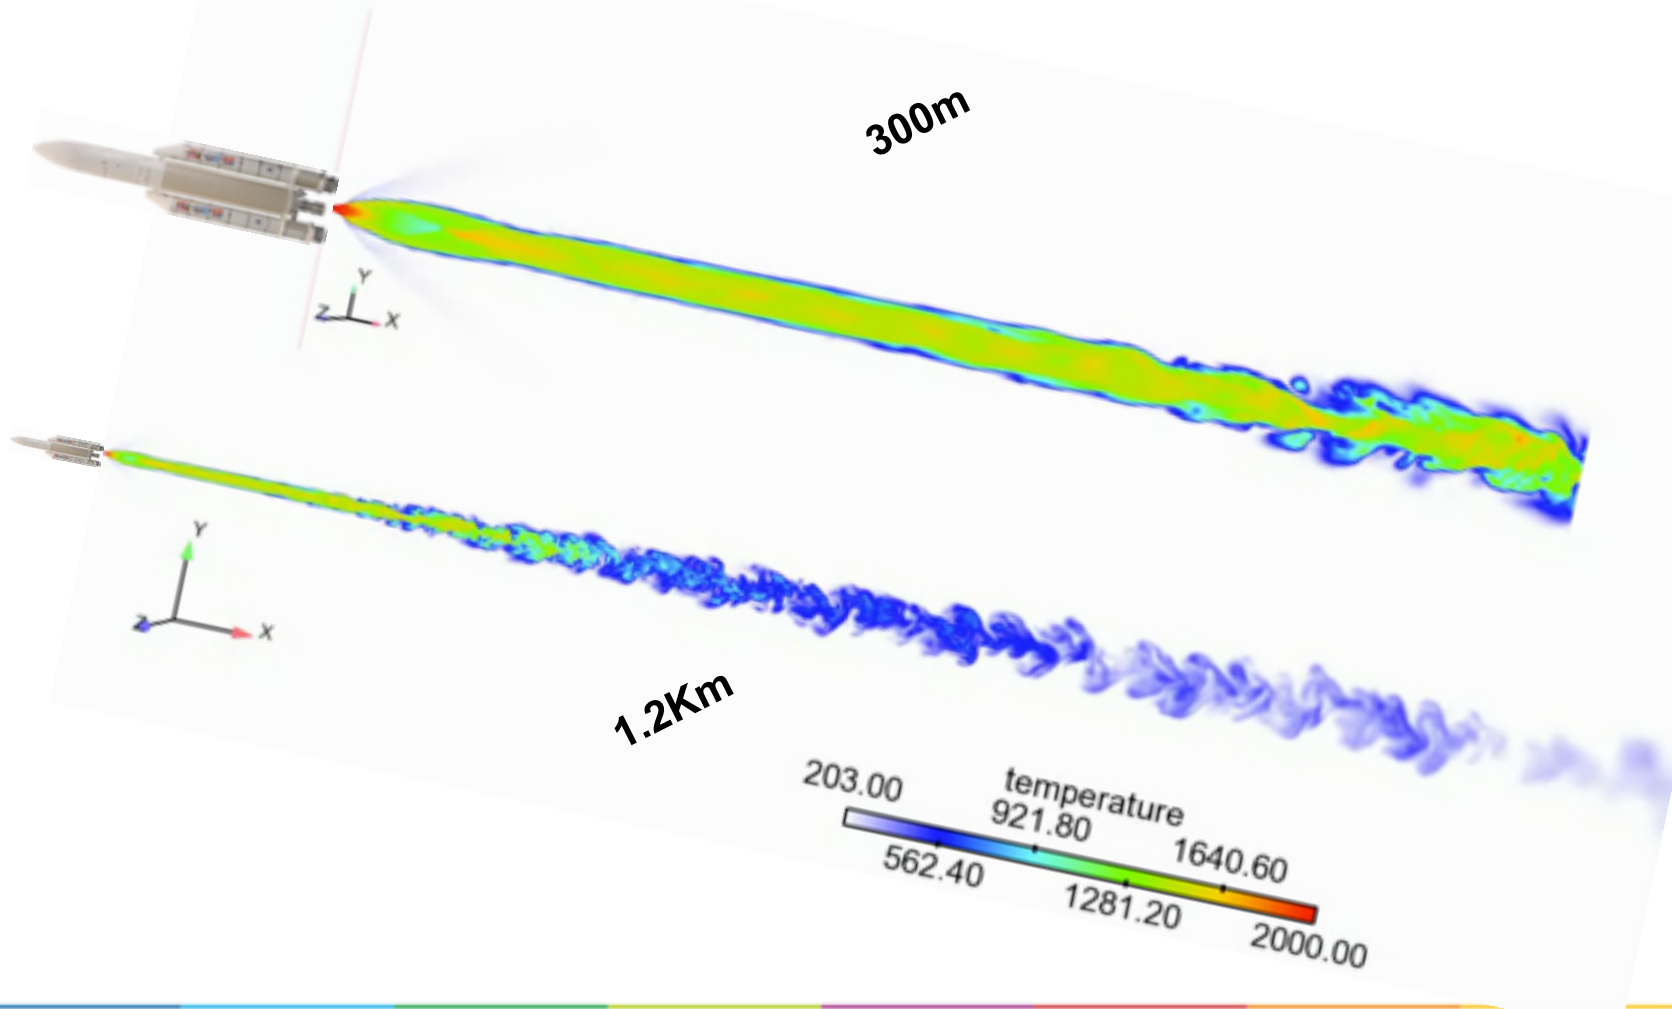
Booster trail simulation



Fun application: 3D Large Eddy Simulation of a booster trail.

A. Poubeau R. Paoli

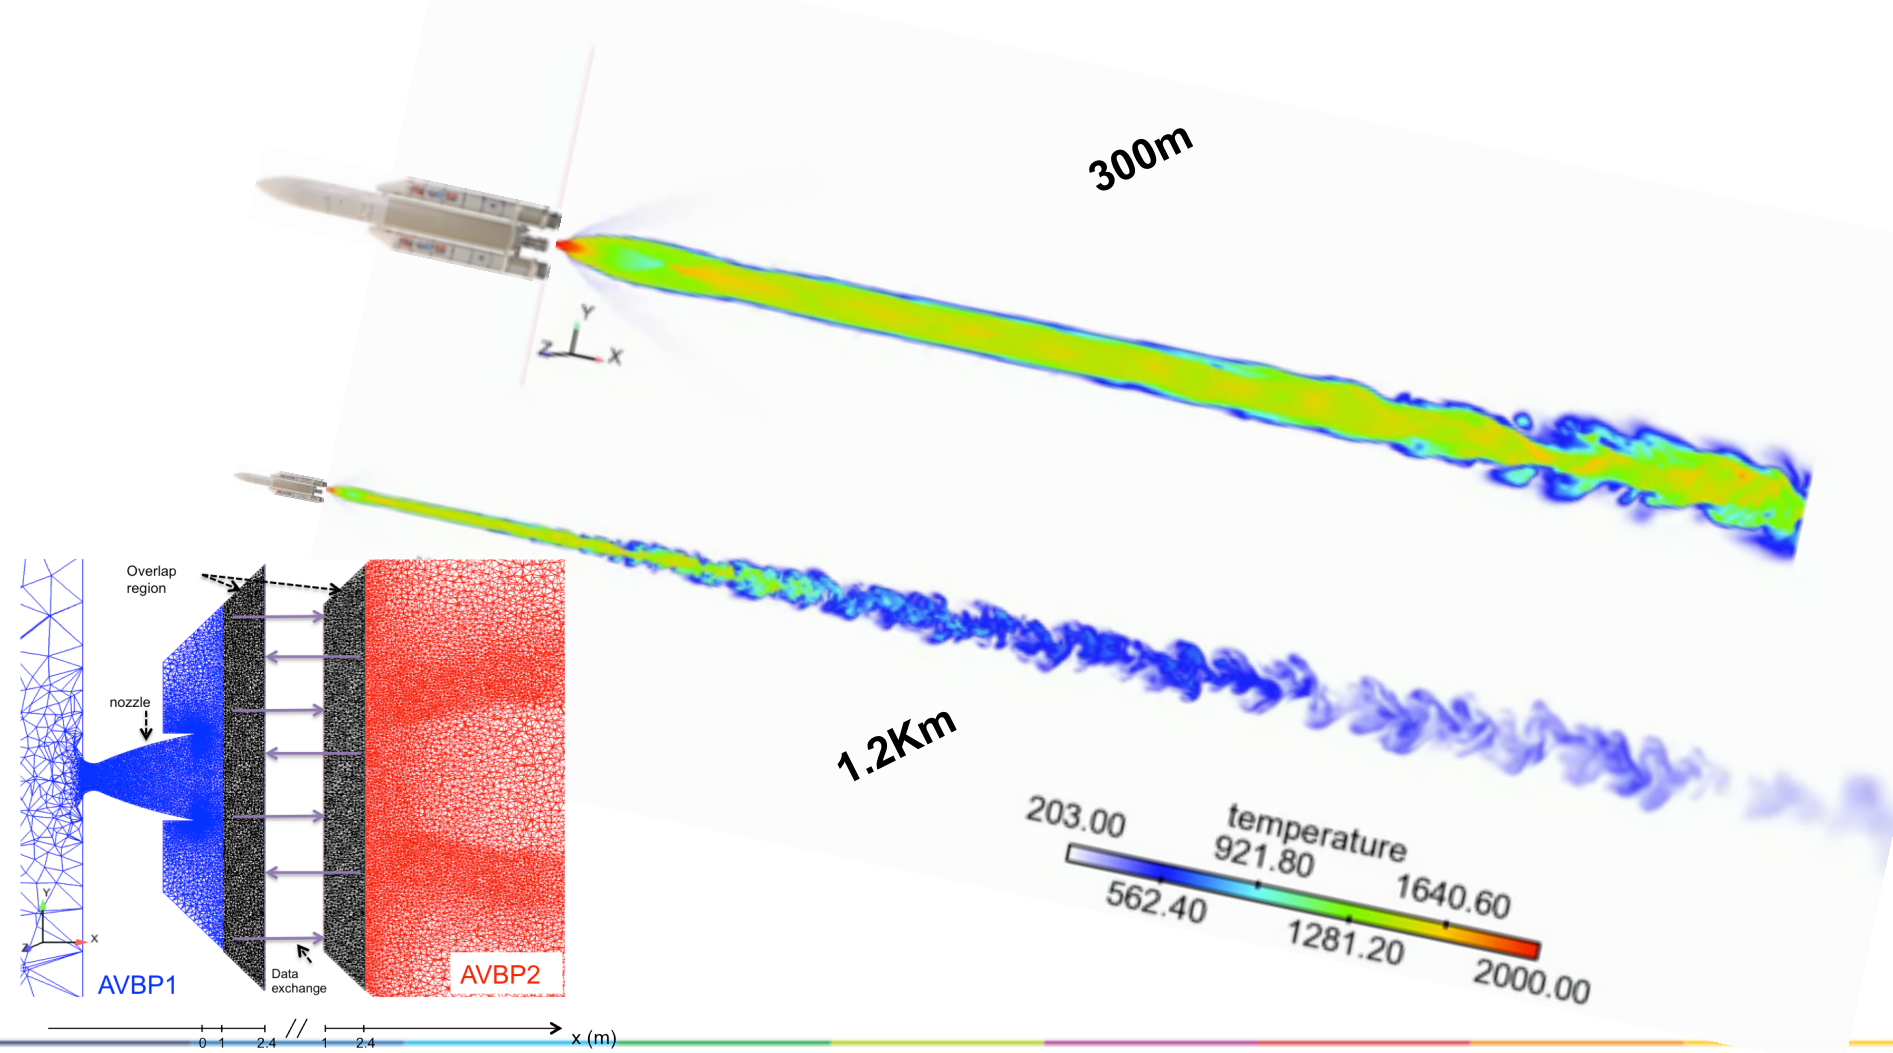
Booster trail simulation



Fun application: 3D Large Eddy Simulation of a booster trail.

A. Poubeau R. Paoli

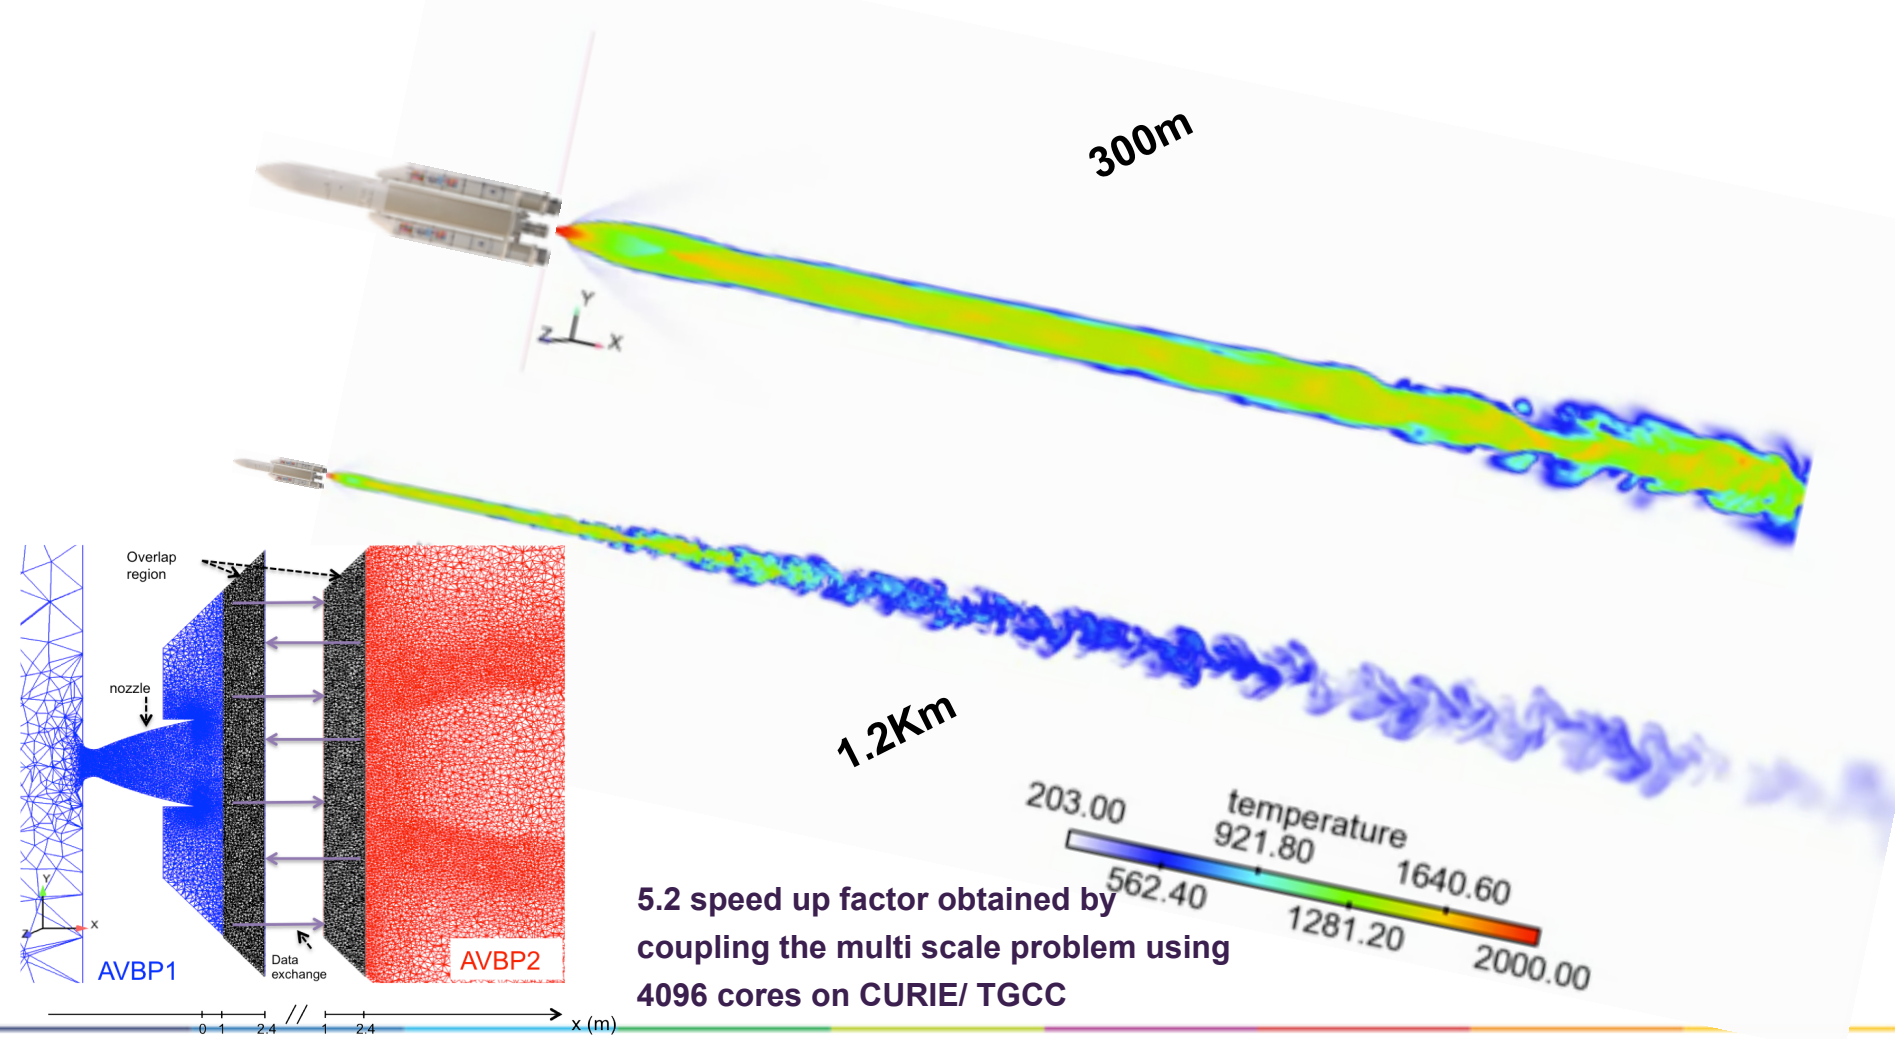
Booster trail simulation



Fun application: 3D Large Eddy Simulation of a booster trail.

A. Poubeau R. Paoli

Booster trail simulation





5.2 speed up factor obtained by coupling the multi scale problem using 4096 cores on CURIE/ TGCC



Fun application



A. Bonhomme et al.





Fun application

A. Bonhomme et al.





CFD MPI challenges



- ➔ High usage of collectives ..
 - Physics likes maxima/Sums ... Not optimised on all implementations.
 - For code coupling .. reductions on sub-communicators is 100 times slower than on MPI_COMM_WORLD (BG \Q, CRAY)
- ➔ Enormous dependency on partitioning .. Most balanced work is not always the best approach ..
- ➔ How to handle heterogeneous machines with only MPI ?

MPI limitations on Heterogeneous MIC systems

- ◆ MPI execution times is highly dominated by MPI



Execution times distributions on 2 Xeon Phi
-- average times for 1 simulation time step --



XEON PHI

- ➔ Possible to avoid hybrid approach via runtime or MPI put/get implementation ?

Code modernisation : MPI+OPENMP4

➔ Introduction of a second level of decomposition

MPI decomposition

Parmetis / Ptscotch

Group decomposition

METIS


all ranks

task decomposition

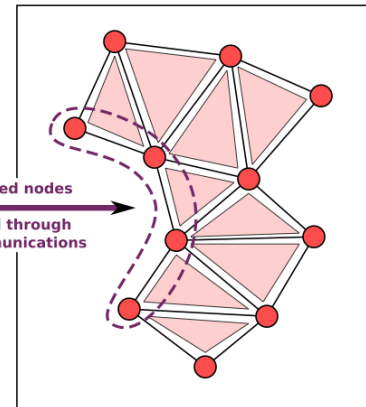
compute node

MPI rank 1

MPI rank 0




duplicated nodes
updated through
MPI communications



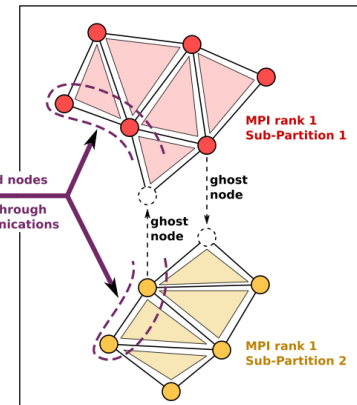
task

MPI rank 0

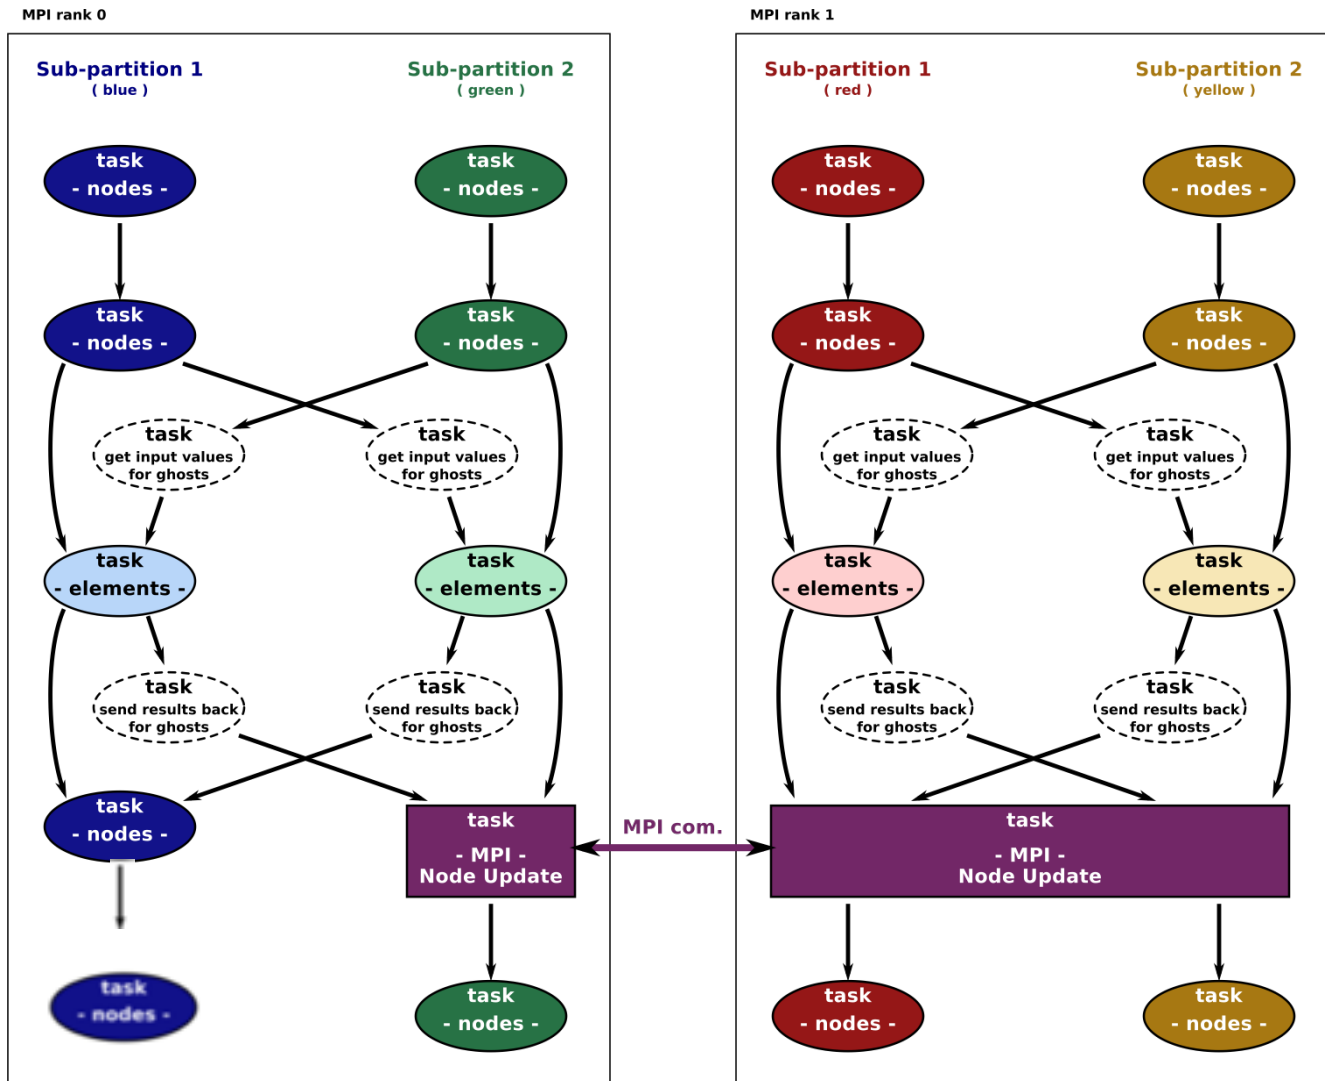


duplicated nodes
updated through
MPI communications


MPI rank 1




MPI+OPENMP4 tasks



Parallel efficiency on a single Xeon with 60 cores on a small test case







Parallel efficiency on a single Xeon with 60 cores on a small test case




THANK YOU

ICAMDAC A. Misdaris
2070 CA


A. Bonhomme



J.M. Senoner



D. Papadogianis




F. Duchaine



THANK YOU

ICAMDAC A. Misdaris
2070 CA


A. Bonhomme




Velocity [m/s]
0 20 40 60 80



J.M. Senoner



D. Papadogiannis



F. Duchaine