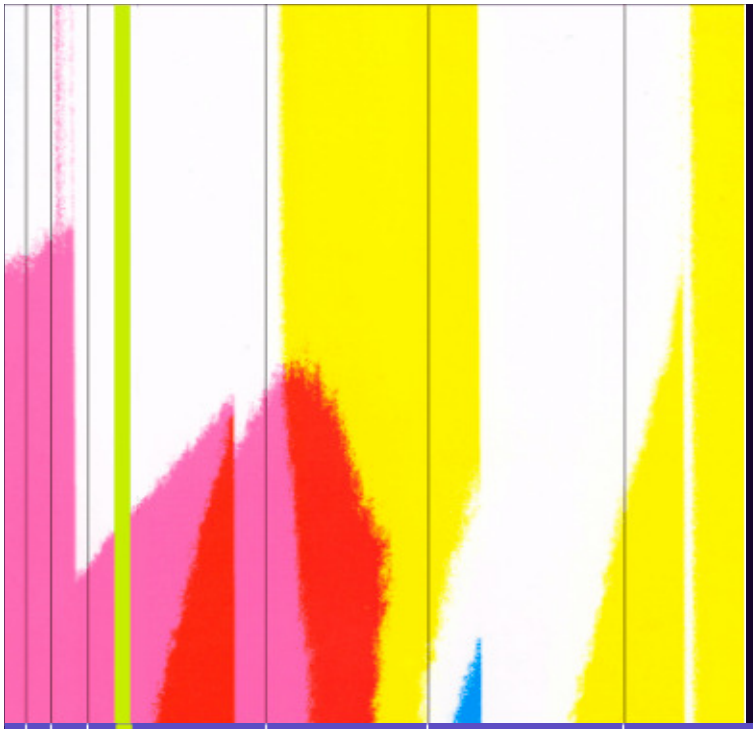

A PERSPECTIVE ON HIGH-PERFORMANCE COLLABORATIVE PLATFORMS FOR MULTIDISCIPLINE SIMULATION & OPTIMIZATION

Toàn NGUYÊN

Harbin (China), March 30, 2009

40 ans la révolution de l'information

INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET EN AUTOMATIQUE

INRIA RHÔNE-ALPES

2nd China-EU Workshop on
Multiphysics Modelling, Simulation, Validation and Optimization

Outline

INRIA

Example Scenario

Multidiscipline challenges

HPC Contributions

Workflows

Conclusions

Outline

INRIA

Example Scenario

Multidiscipline challenges

HPC Contributions

Workflows

Conclusions

INRIA

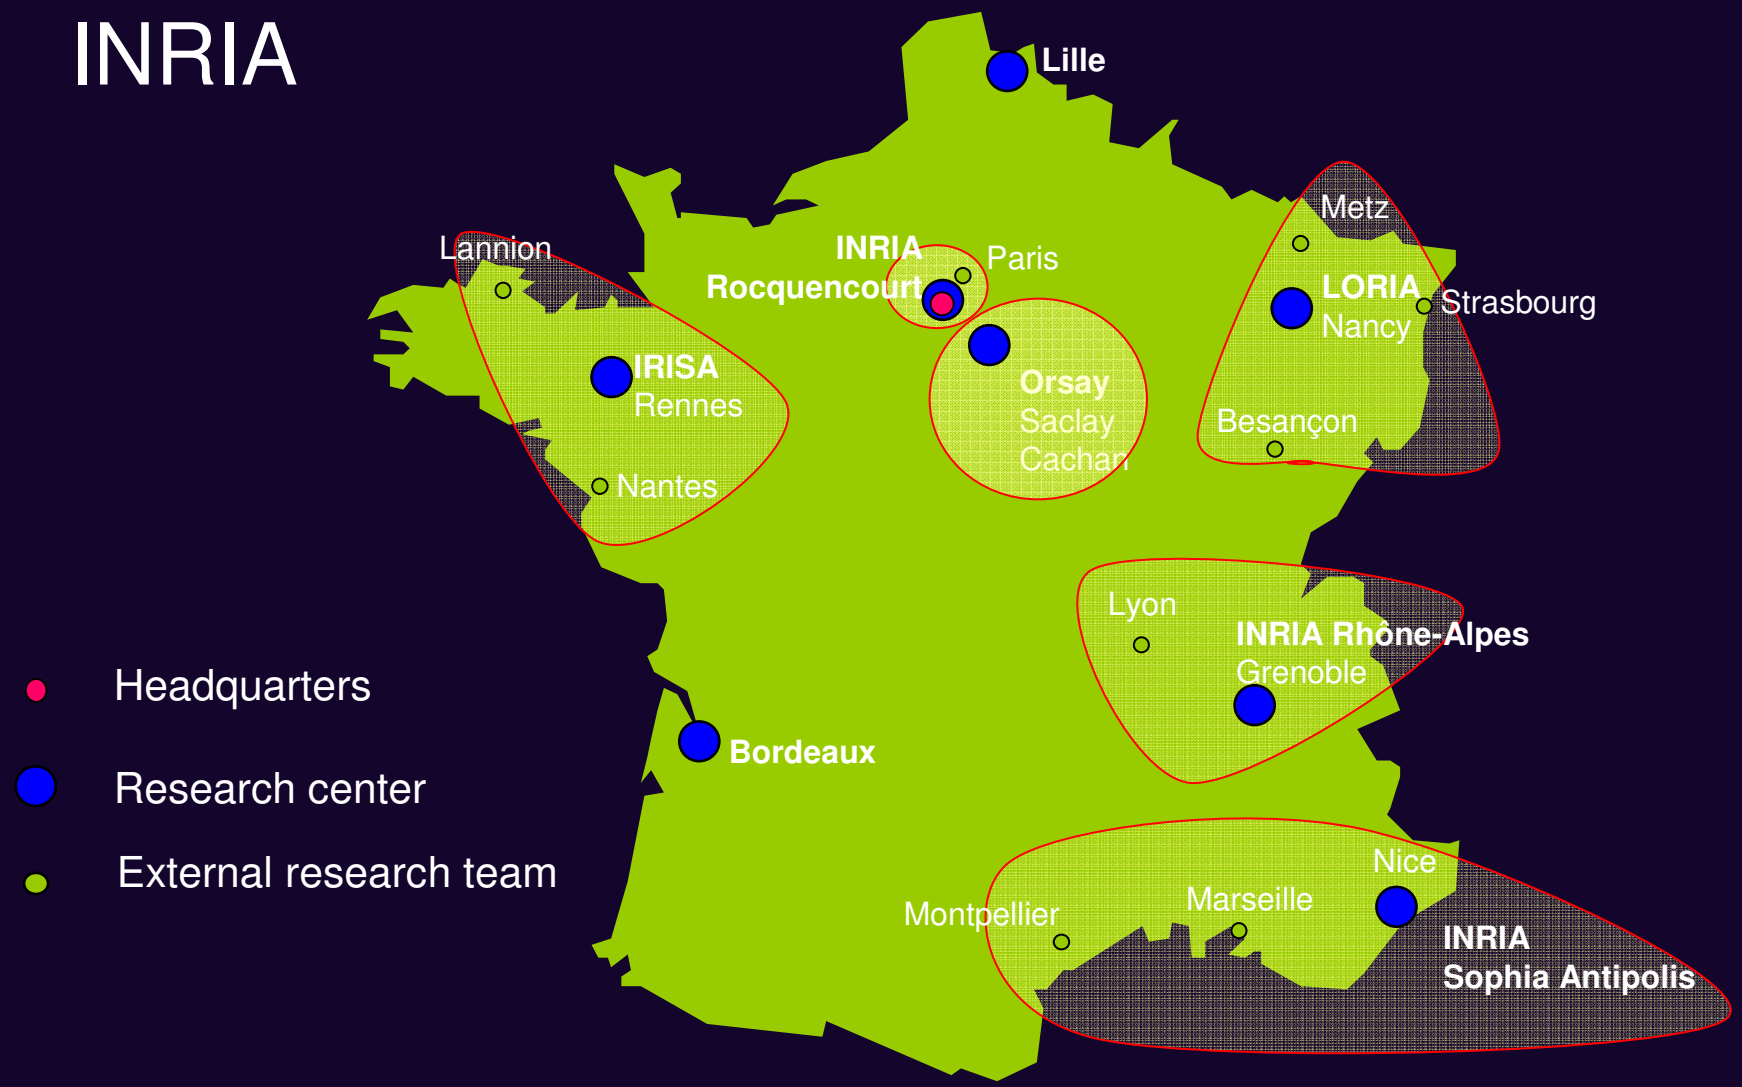
Lille
Metz
Lannion
INRIA
Rocquencourt
Paris
LORIA
Nancy
Strasbourg
IRISA
Rennes
Orsay
Saclay
Cachan
Besançon
Nantes
Lyon
INRIA Rhône-Alpes
Grenoble
Bordeaux
Nice
Marseille
Montpellier
INRIA
Sophia Antipolis
Headquarters
Research center
External research team

INRIA Key figures

Jan. 2008

A workforce of 3,800

including 2,000 paid by INRIA

2,800 scientists
- 950 PhD candidates (400 from abroad)
- 500 post-docs and engineers

1050 staff
- 468 research scientists
- 560 ETA
- 300 interns

120 projects in FP6 and 30+ in FP7

Budget: 180 M€ (tax not incl.)

20% from contracts, licenses, 8% from Europe

INRIA organization

Research projects

- The key organization: 180 project-teams (10-20 persons, 4 years, 10-90% funding)

Research centers

- Decentralized, self-supporting structure with strong connections with the regional environment, no intermediate administrative level

Scientific and functional departments

- Coordination and organization role

Discussion structures

- Center Committee in the research centers and discussion committee at the national level

Evaluation bodies

- Research Projects Committee in the research centers, Evaluation Committee at the national level

Steering bodies

- Research centers management teams, national management committees

KEY-WORDS

- What: COLLABORATIVE INTEGRATION PLATFORMS
- How: DISTRIBUTED HIGH-PERFORMANCE ENVIRONMENTS

Programming & Communication

- Where: INNOVATIVE MULTIDISCIPLINE APPLICATIONS

Modeling & Simulation

GOALS

- PROVIDE SEAMLESS & CONFIGURABLE PLATFORMS

- SUPPORTED BY:
  1 - HIGH-PERFORMANCE DISTRIBUTED SYSTEMS
  2 - CONFIGURABLE & RESILIENT SYSTEMS
  3 - PROGRAMMING MODELS & TOOLS for DISTRIBUTED COMPUTING

Technology Push

- APPLIED TO LARGE-SCALE INNOVATIVE MULTIDISCIPLINE APPLICATIONS

Application Pull

Outline

INRIA

Example Scenario

Multidiscipline challenges

HPC Contributions

Workflows

Conclusions

Aeroelastic Simulation: M=0.8, Nz=1, Z=11277 m

AleniaAeronautica

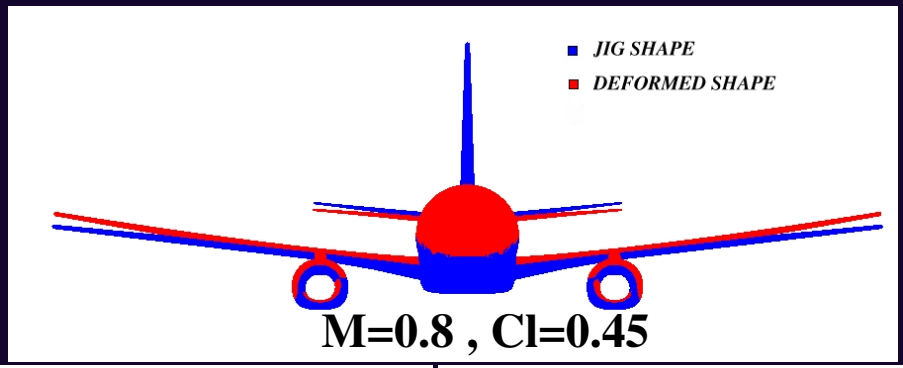

M=0.8 , Cl=0.45

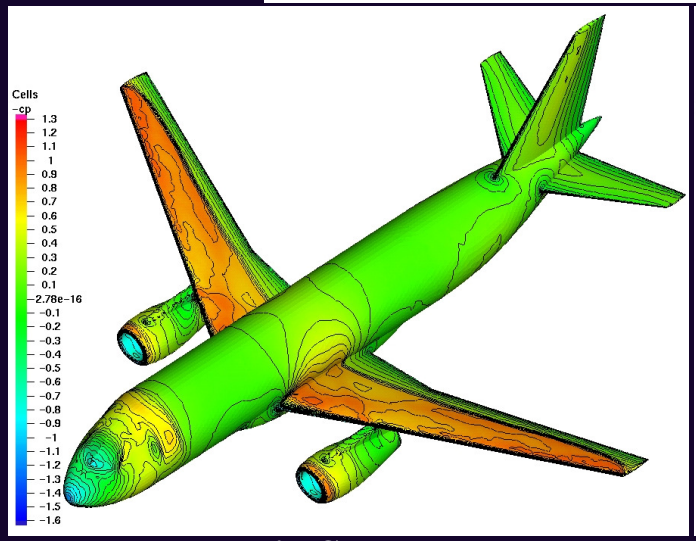

Jig Shape

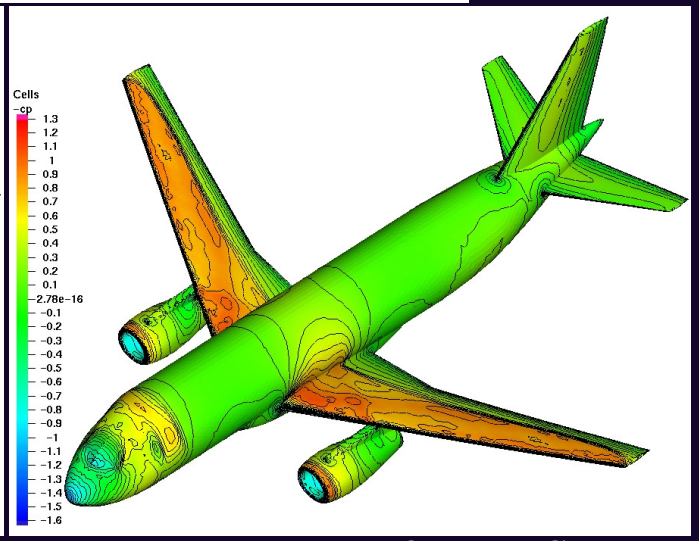

Deformed Shape

Fluid Structure Interaction: Process Overview

Aero-structural Design Process

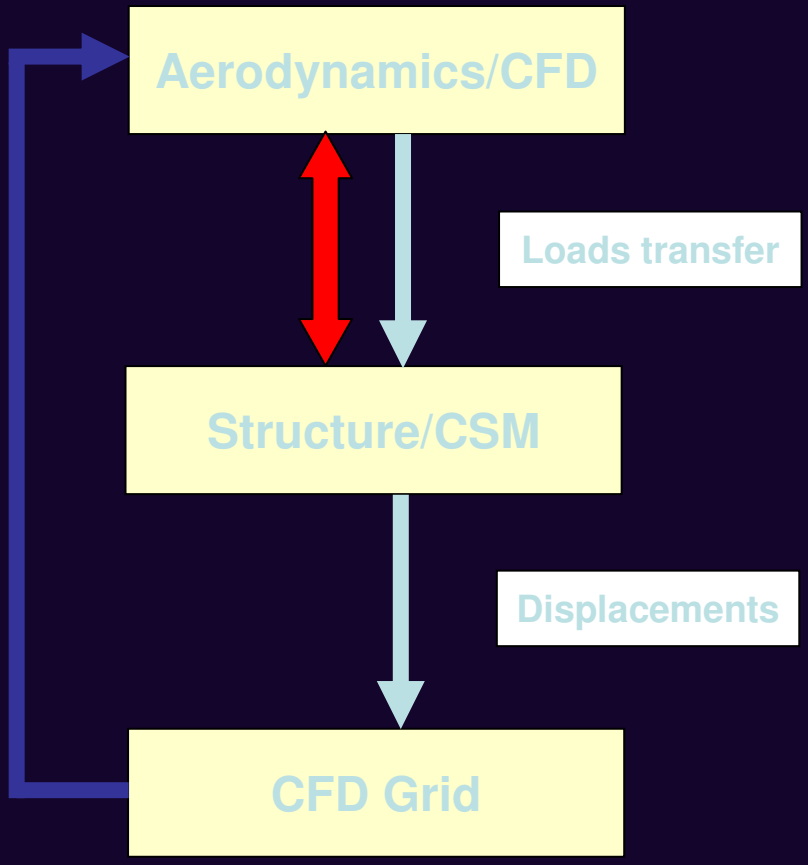

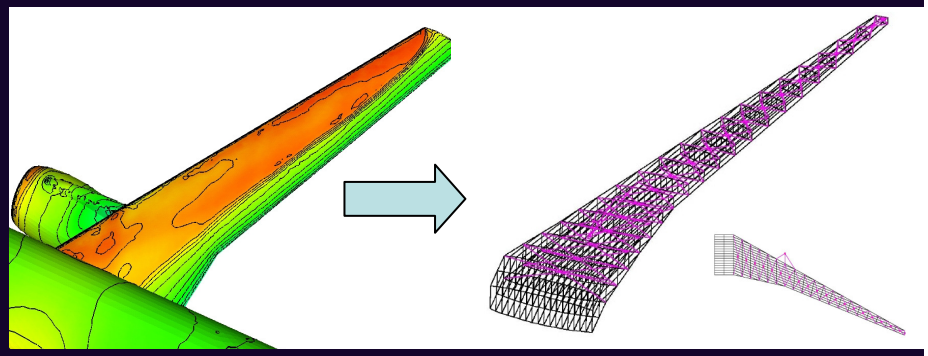

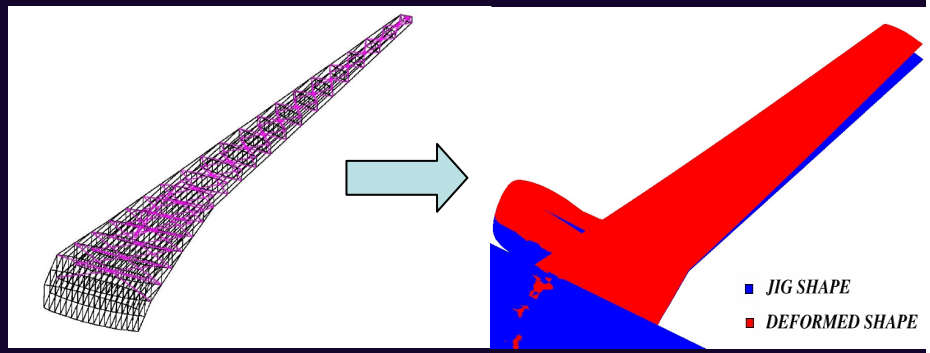

Process definition

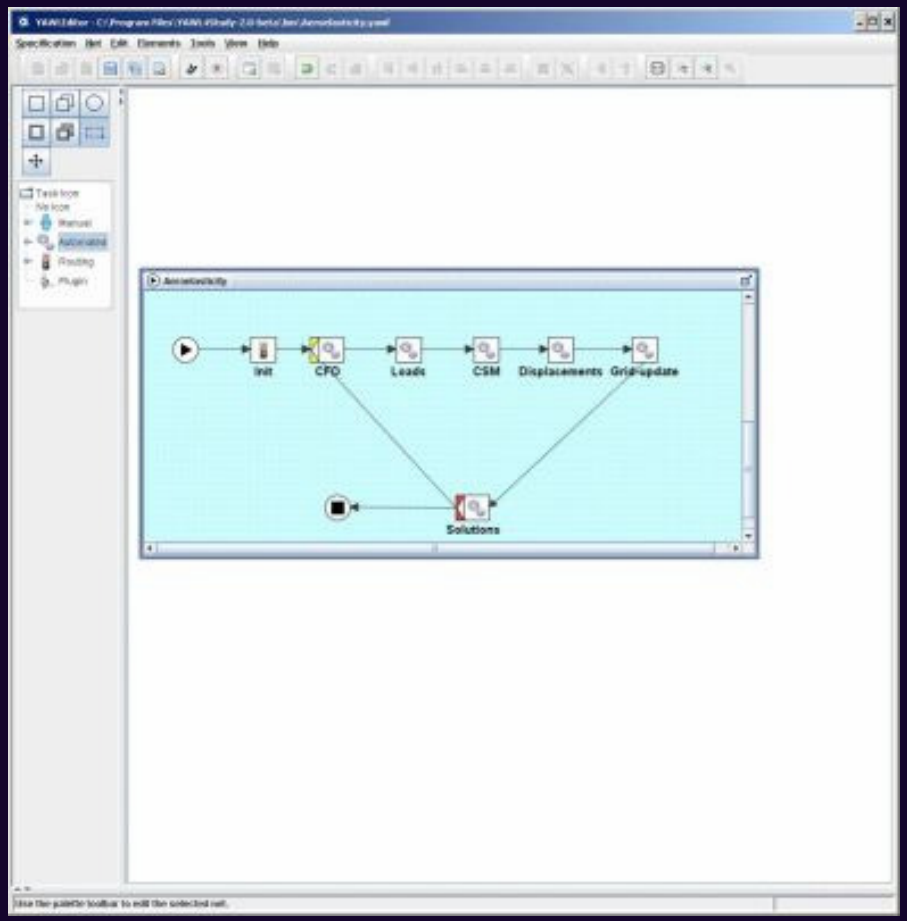

Innovative multidiscipline Applications

Combine algorithms and methodologies

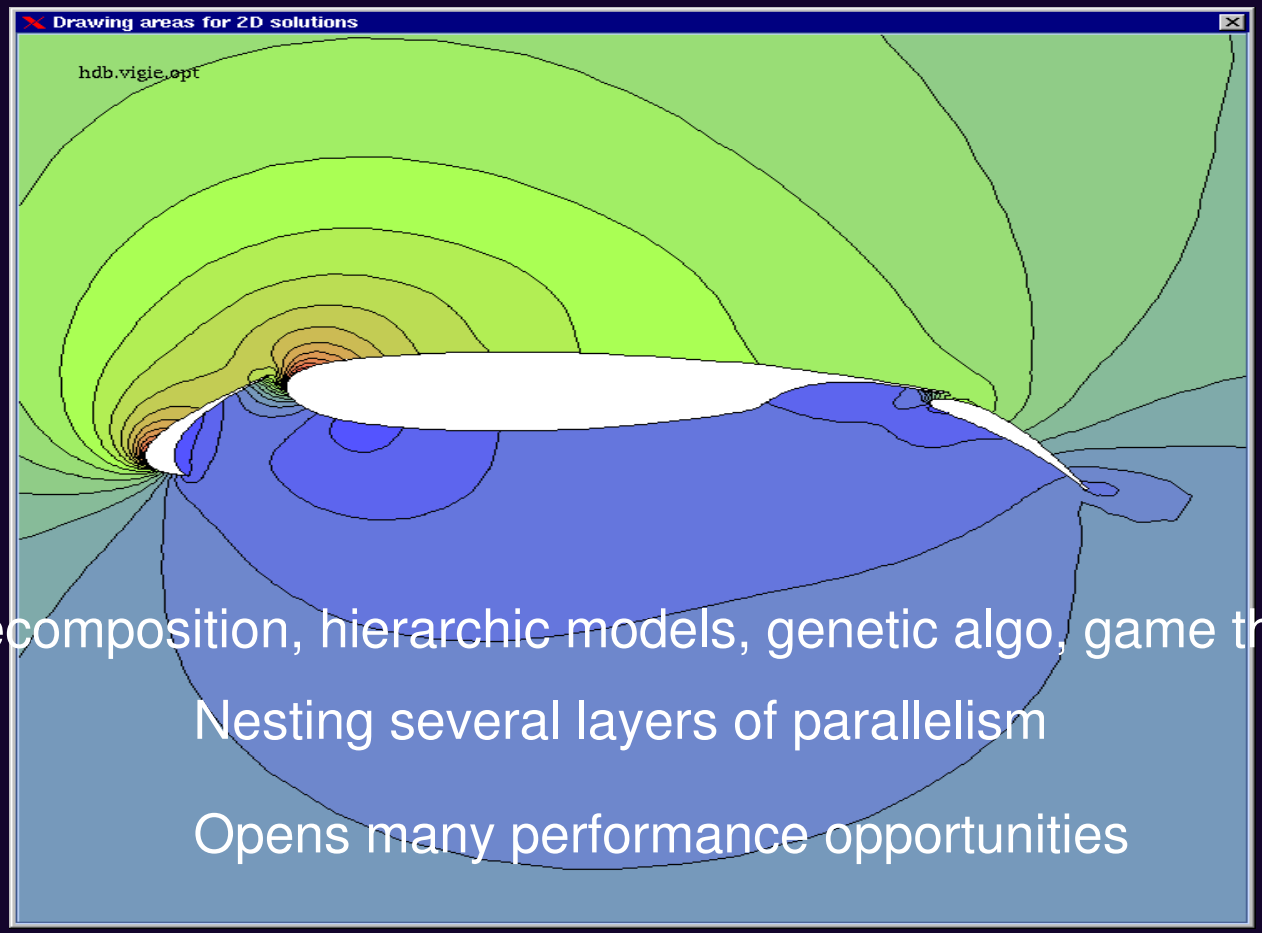

Domain decomposition, hierarchic models, genetic algo, game theory…

Nesting several layers of parallelism

Opens many performance opportunities

Outline

INRIA

Example Scenario

Multidiscipline challenges

HPC Contributions

Workflows

Conclusions

Multiphysics challenges

- Reduce time to market, design better & faster
- Use efficient, robust & validated methodologies & tools
- Set-up collaborative teams, procedures & workplans
- Deploy & take-up HPC resources •
- Design large-scale multiphysics simulation & optimization •
- Contribute to knowledge for safer, cleaner, quieter designs
- Contribute to virtual flight tests, virtual certification by 2020

Transport & Aeronautics Vision 2020

Multiphysics Design, simultation & optimization

nesting complex multidiscipline processes

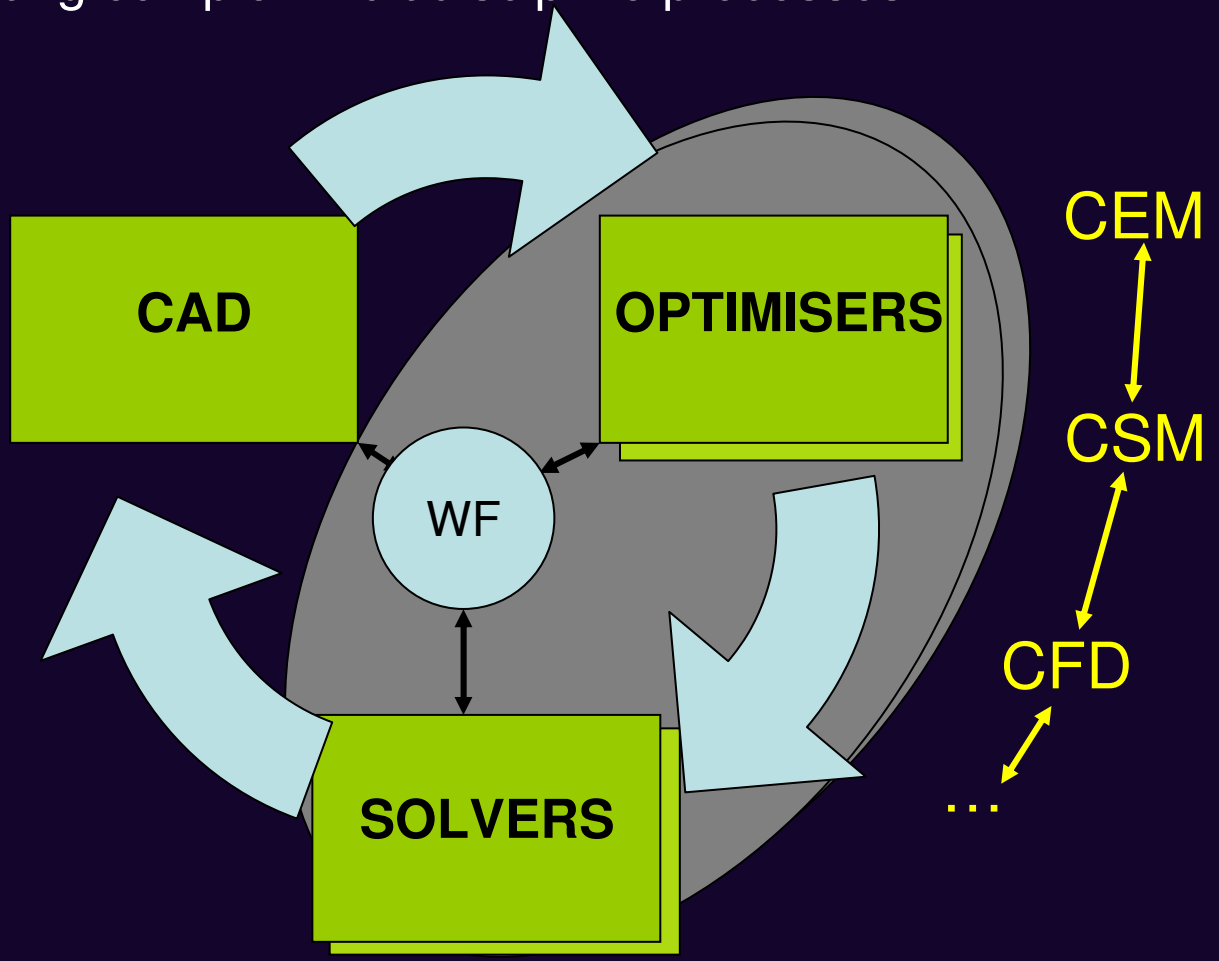

HPC Contributions: ongoing efforts

- High-Performance Computing era: multi-core, GPU,…
- Heterogeneous, parallel, grid and distributed computing
- Sustainable, affordable, mature, HPC & wide-area computing
- Collaborative environments: //viz+HPC, VR workbenches, …
- Disseminate in many domains: genomic, seismic,aerospace, medical, chemical, automotive, astrophysics, energy…

The e-science era

Outline

INRIA

Example Scenario

Multidiscipline challenges

HPC Contributions

Workflows

Conclusions

HPC: from Mega-Flops to Exa-Flops…

CDC6600	Cray2	ASCI Red	Blue Gene /L	Roadrunner	
1MFlops	1GFlops	1TFlops	.5PFlops	1PFlops	1EFlops
1964	1985	1997	2007	2008	…2018
10**6 flops	10**9	10**12	#10**15	10**15	10**18

HPC: the road to Exascale computing…

- Exascale 2018 plan at Intel: computing power >> Moore's law
- Blue Waters (NCSA), INCITE (US): free petascale for industry
- RIKEN KEISOKU (Japan), DEISA, PRACE (EU) INITIATIVES
- MOSART, CRESCENDO, NURESIM (CEA nuclear plant sim)
- ITER Tokamak fusion facility, Argonne, Oak Ridge, MIT Labs

Blue Gene /L (.5 Pflops), /P (3 Pflops), /Q (10 Pflops), …

ISSUES

- COMPUTING POWER IS NOT: WHAT ARE THE ISSUES ?
- TECHNOLOGY BARRIER PLAGUES HPC DISSEMINATION
- HETEROGENEITIES PLAGUE MULTIPHYSICS SIMULATION
- CULTURE & IPR PLAGUE COLLABORATIVE WORK…
- NEED TO DEVELOP COMMON APPROACHES
- • • • PETASCALE DATA MANAGEMENT ISSUE

FOCUS

- NEED TO SOFTEN ACCESS TO HPC TECHNOLOGY
- METHODS & TOOLS TO SUPPORT HETEROGENEITIES
- EXPLORE, DESIGN, COMBINE NEW APPROACHES
- EXPERIMENT ON LARGE-SCALE TESTCASES
- NEED TO DEVELOP COMMON COLLABORATIVE TOOLS
- • • • AUTOMATE SCENARIO DEPLOYMENT & EXECUTION

Outline

INRIA

Example Scenario

Multidiscipline challenges

HPC Contributions

Workflows

Conclusions

OPEN ISSUES

- USERS NOT AWARE OF UNDERLYING HW, OS, GRIDS…
- DISTRIBUTION, FAILURE, SECURITY TRANSPARENT
- COMPOSE & REUSE LEGACY TOOLS, CODES, APPs
- NO USE OF EXOTIC LANGUAGES & COMMANDS : GUI++
- COMPLIANT WITH CURRENT & UPCOMING STANDARDS

  CORBA, SOA, WSDL, WSRF, OGSA, Globus GT3 & GT4, UNICORE…

Distributed Workflows

DISTRIBUTED INTERDEPENDENT COMPUTATIONS

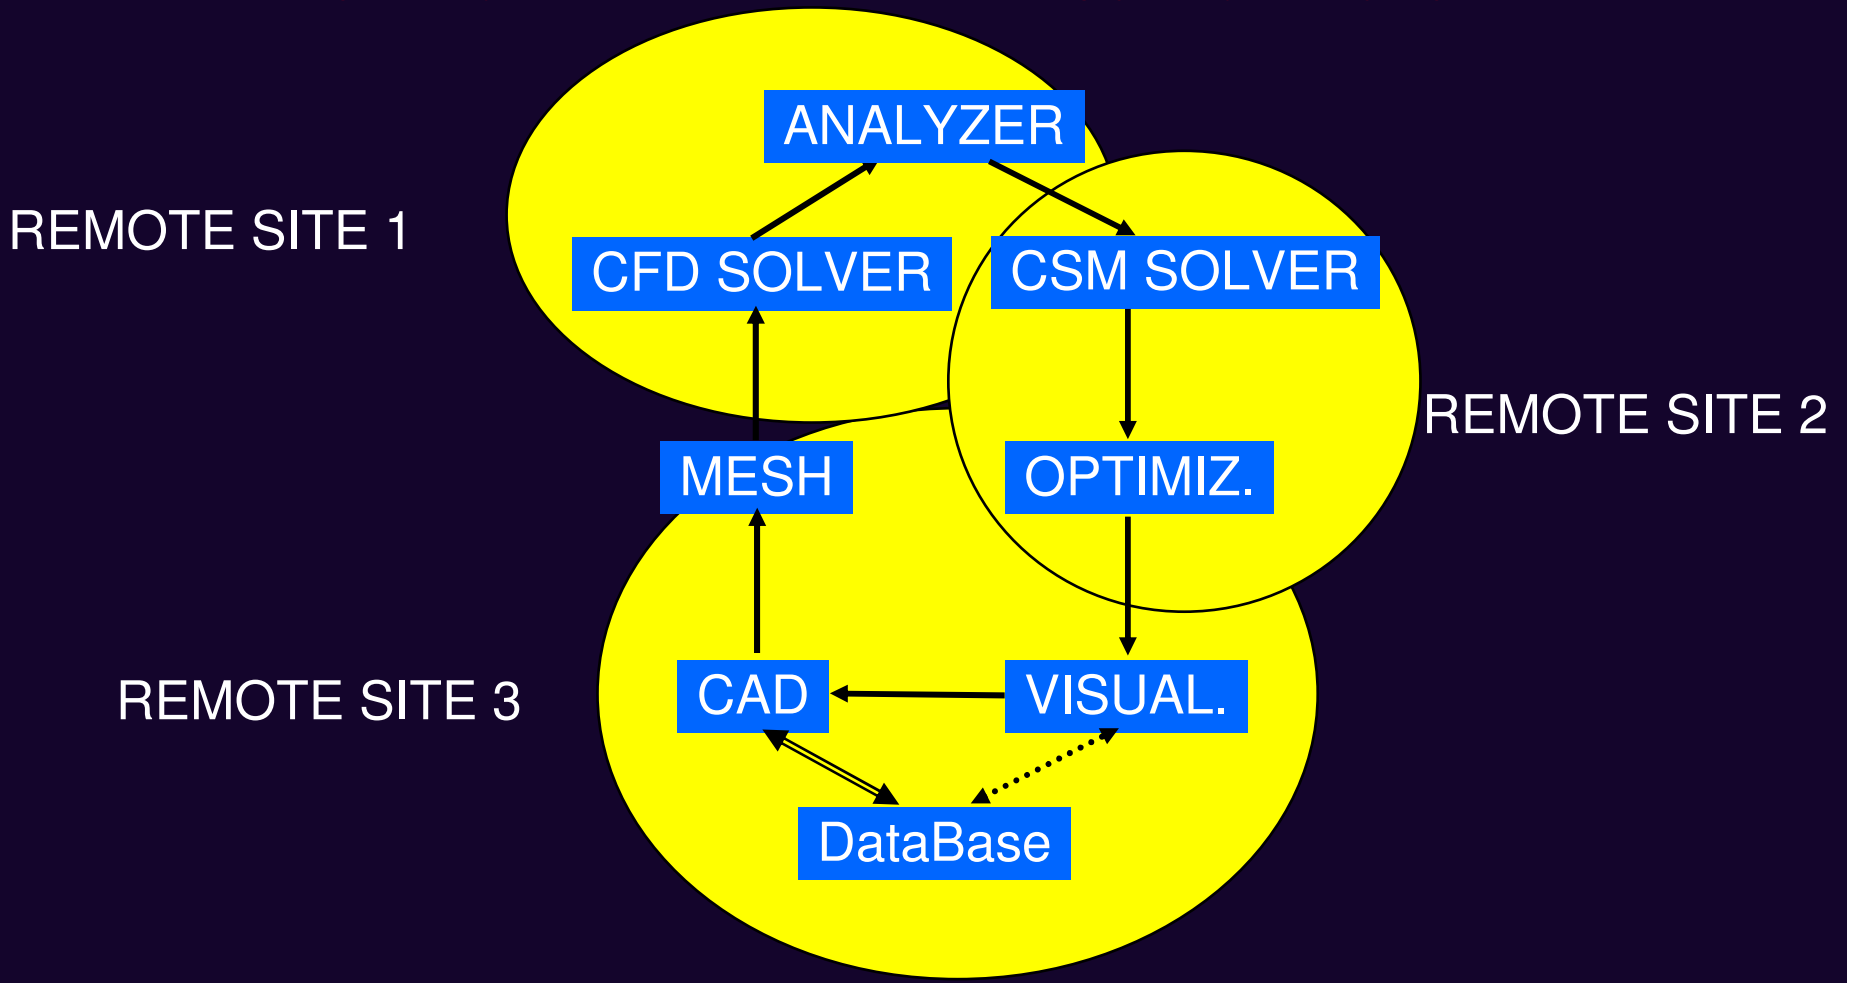

UNCHANGED WHATEVER THE USERS AND COMPUTING INFRASTRUCTURE

Distributed Workflows

DECLARATIVE REPRESENTATIONS OF HETEROGENEOUS DISTRIBUTED INTERDEPENDENT COMPUTATIONS

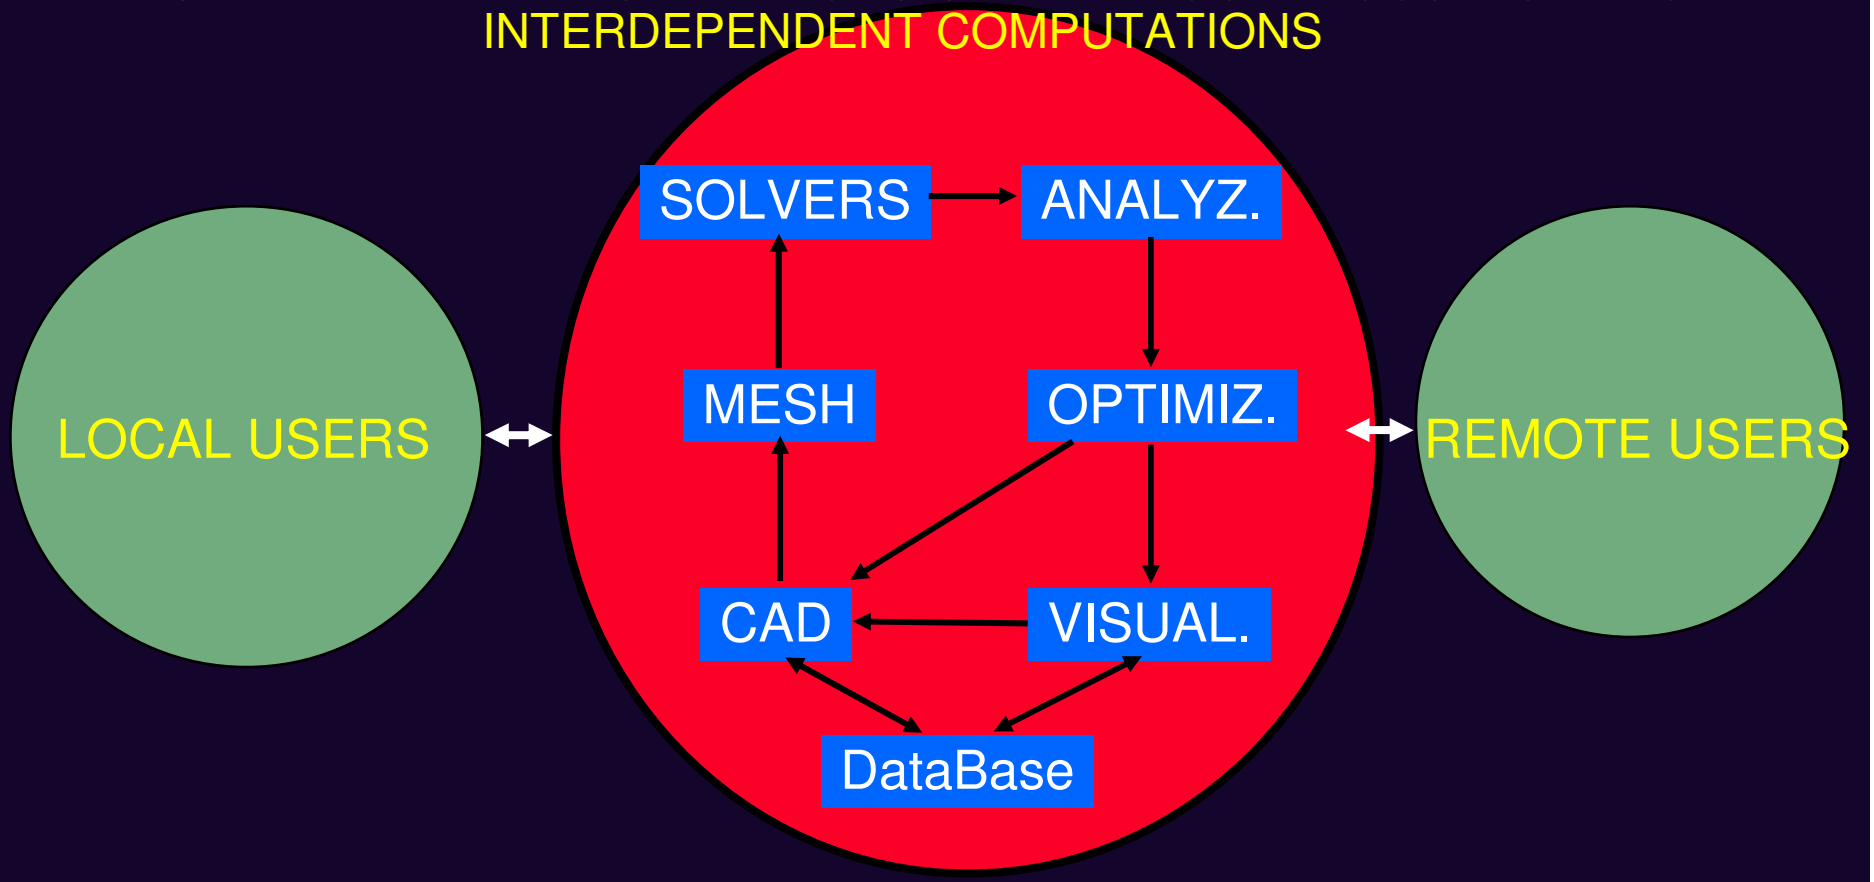

UNCHANGED WHATEVER THE INFRASTRUCTURE AND USERS

Distributed High-Performance Infrastructures

WIDE-AREA BROADBAND NETWORKS

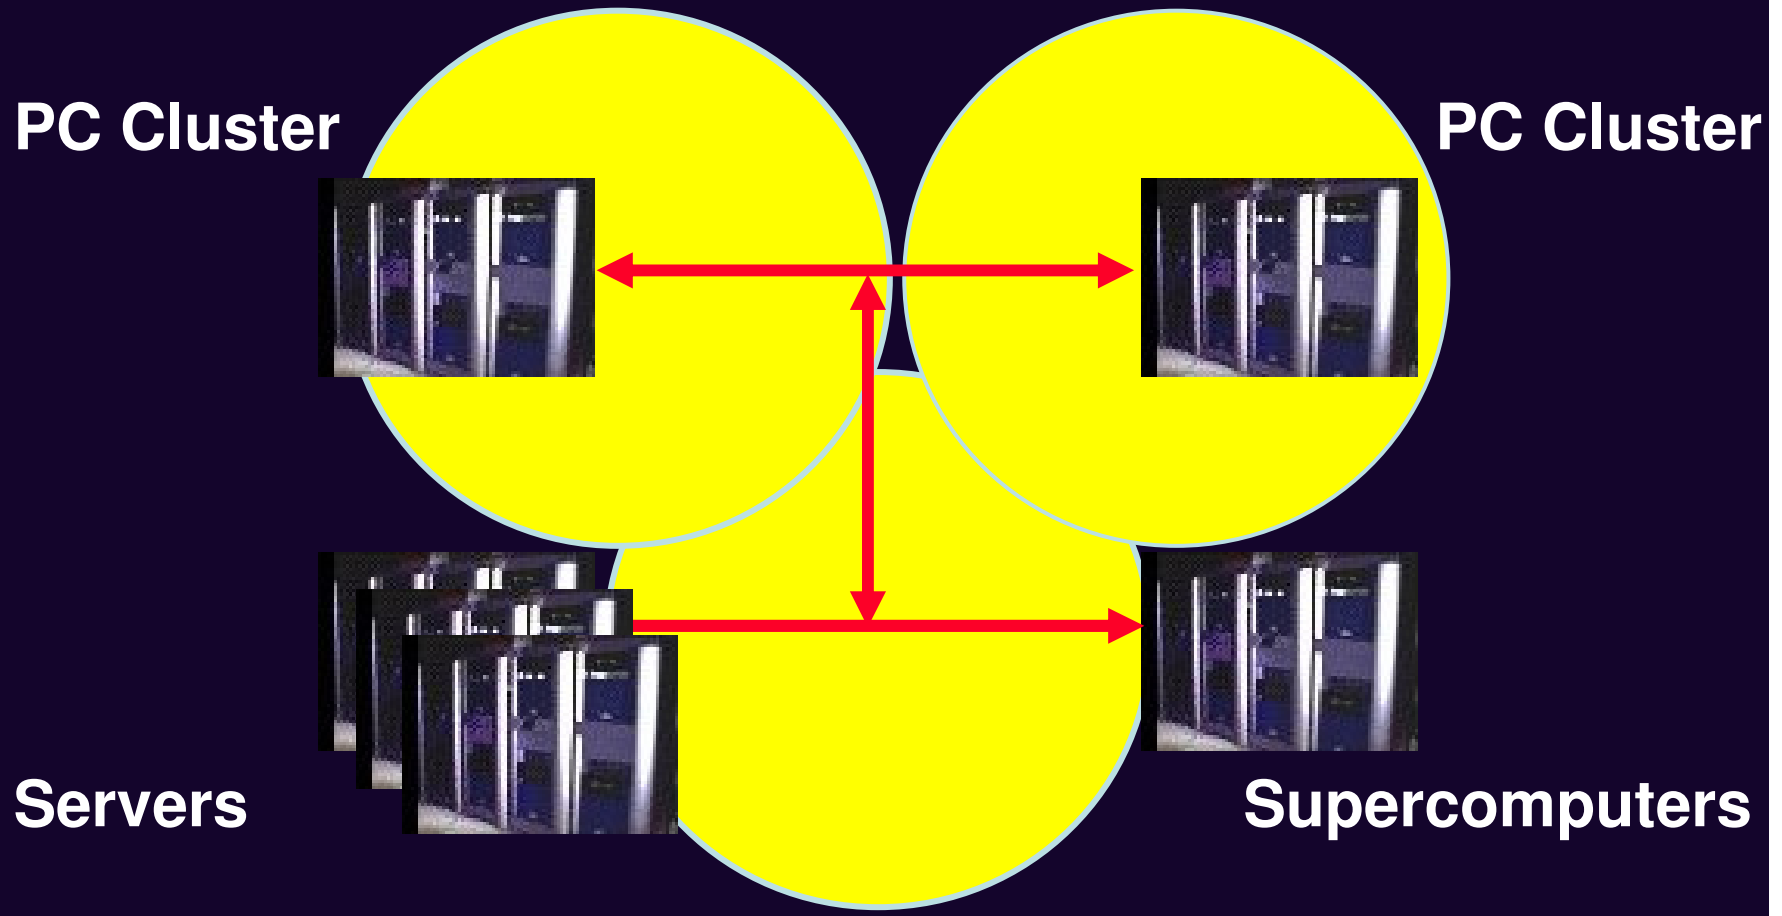

Putting it all together ?

URGENT NEED TO SIMPLIFY APPLICATION & USER INTERFACES

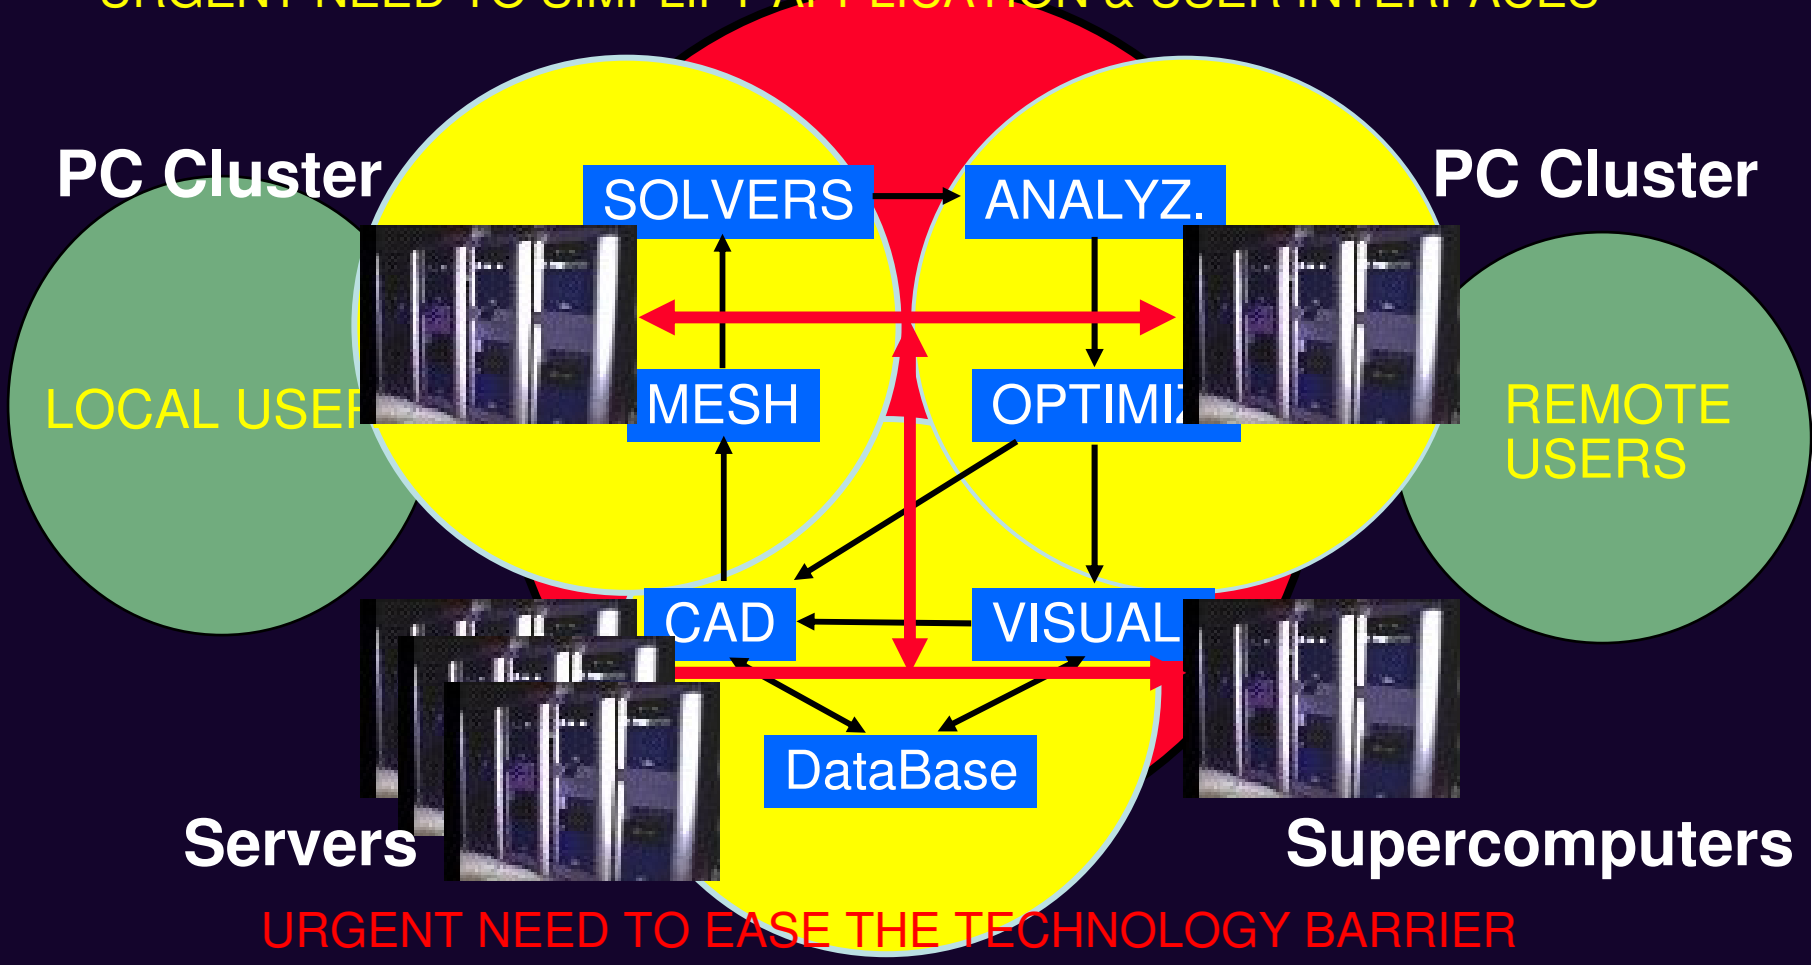

URGENT NEED TO EASE THE TECHNOLOGY BARRIER

Distributed Workflows: customization

FOCUS ON APPLICATION & USERS

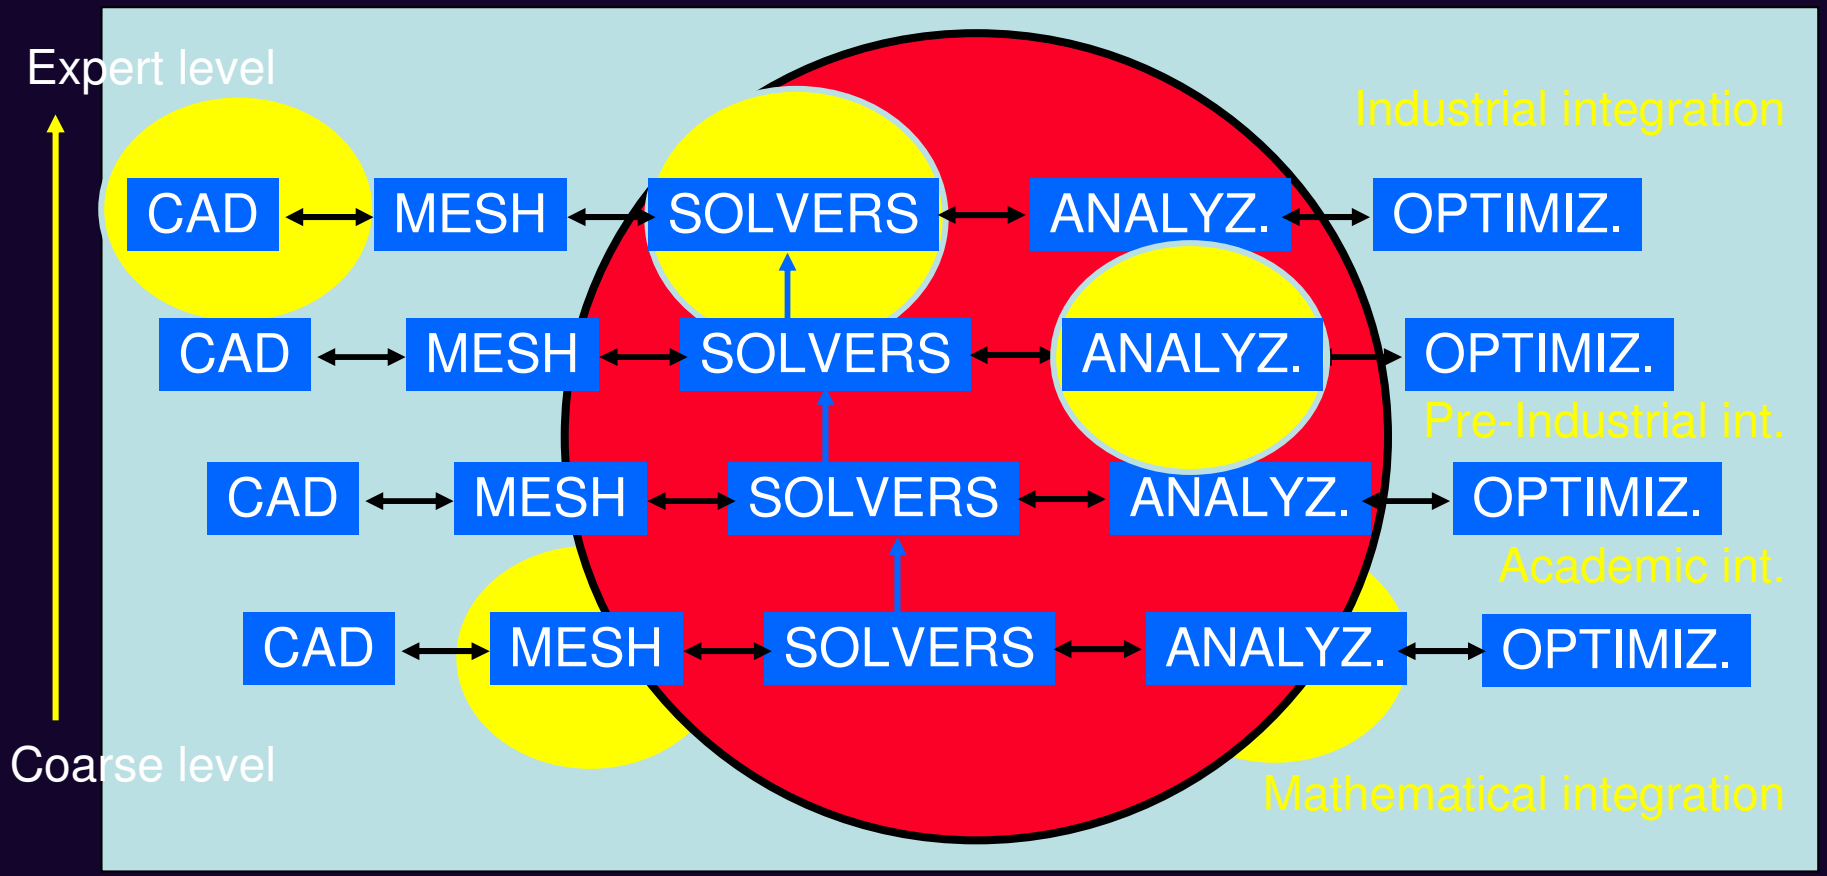

Layered Declarative Application Environments

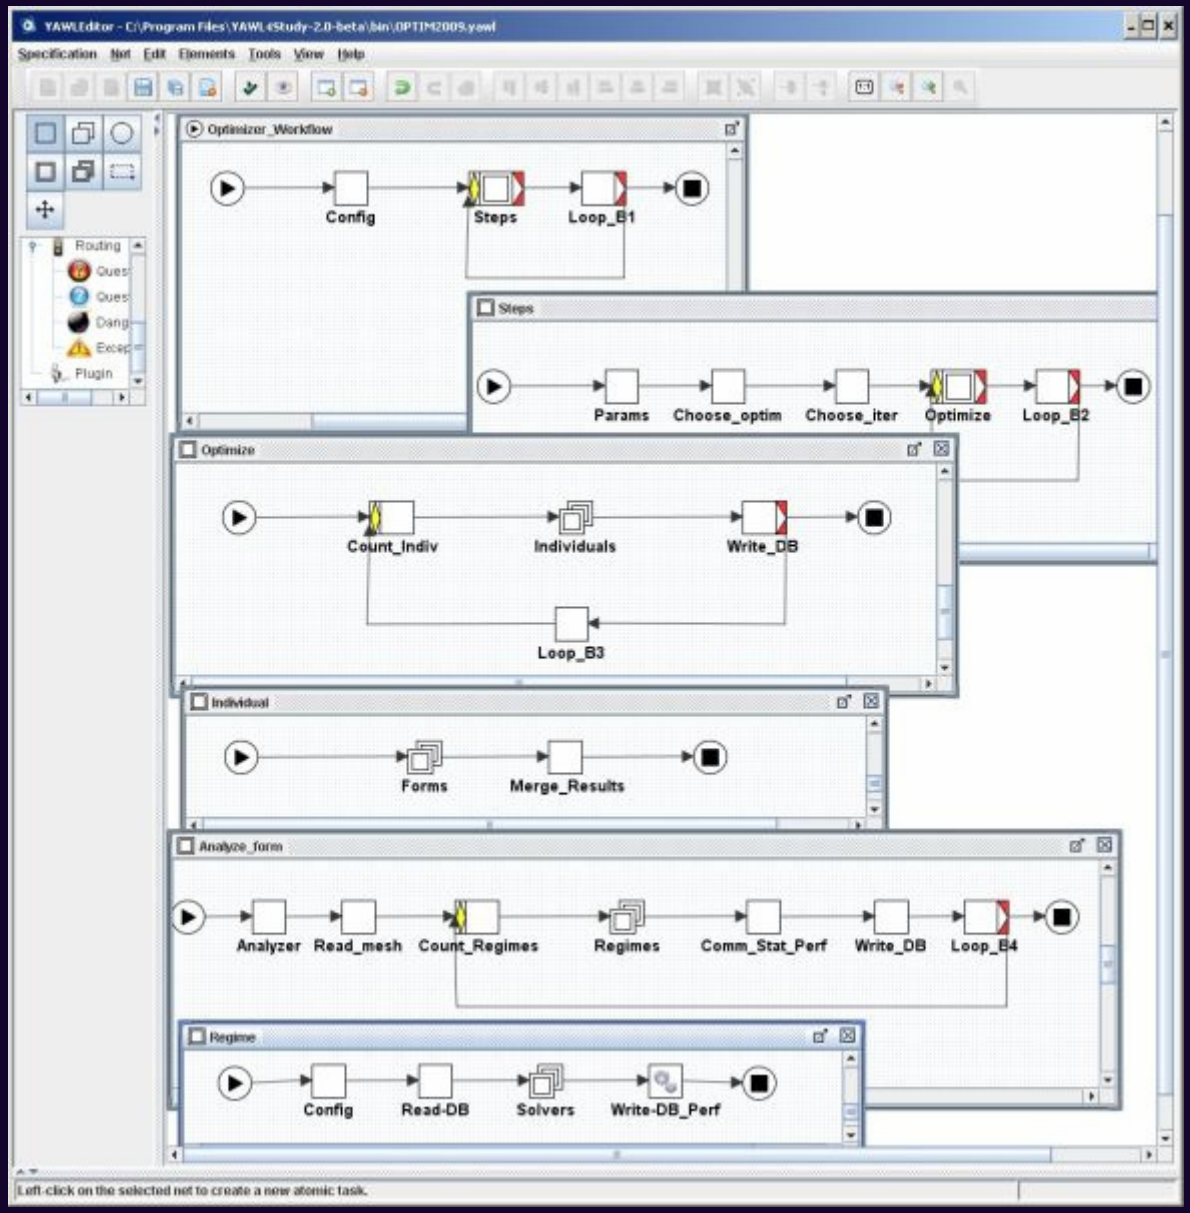
YAWLEditor - C:\Program Files\YAWL4Study-2.0-beta\bin\OPTIM2009.yawl
Specification Net Edit Elements Tools View Help
Routing
Plugin
Optimizer_Workflow
Config
Steps
Loop_B1
Steps
Params
Choose_optim
Choose_iter
Optimize
Loop_B2
Optimize
Count_Indiv
Individuals
Write_DB
Loop_B3
Individual
Forms
Merge_Results
Analyze_form
Analyzer
Read_mesh
Count_Regimes
Regimes
Comm_Stat_Perf
Write_DB
Loop_B4
Regime
Config
Read-DB
Solvers
Write-DB_Perf
Left-click on the selected net to create a new atomic task.

Workflow contributions to Multiphysics Design

- DECLARATIVE REPRESENTATION OF HETEROGENEOUS DISTRIBUTED INTERDEPENDENT COMPUTATIONS
- AUTOMATE & MANAGE COMPLEX DISTRIBUTED COMP.
- USED AS COMPLEX APPLICATION FACTORIES
- RUNTIME BINDING
- IN PLACE REFERENCING
- EXHIBIT MULTI-DISCIPLINE INTERACTION PROTOCOLS •
- DATA STREAMING FOR ASYNCHRONOUS // COMPUTAT.

Workflow contribution to Multiphysics Design

- MANY WORKFLOW candidates: Taverna, YAWL, Pegasus…
- WORKFLOW BUS
- SHARE & COMPOSE WORKFLOWS
- SERVICE ORIENTED ARCHITECTURES
- REUSABLE TRACEABLE COMPOSITE WORKFLOWS
- REPRODUCIBILITY, PROVENANCE OF DATA

Workflows for Multiphysics Design

- PLUG-IN COMPONENTS, DATABASES AND SOFTWARE
- INTERACTIVE MULTI-USER PLATFORMS
- MULTI-DISCIPLINE TRIAL/ERROR ENVIRONMENTS
- EXHIBIT APPLICATION DEPENDENCIES
- COMPUTATIONAL STEERING =
  ON-LINE MONITORING + INTERACTIVE STEERING

OPEN ISSUES

- FLEXIBILITY (i.e. dynamic (re)-configuration), USER EXCEPTION HANDLING (i.e. parameter updates)… DISTRIBUTED PAUSE, ROLL-BACK, RESTART… RESILIENCY, PETA-SCALE VOLUMES DATA…
- COLLABORATION on WORKFLOWS
  - China, Finland, …
- VIRTUAL FLIGHT-TESTS IN AERONAUTICS Vision 2020
  - FP7 Transport (Aeronautics) Call 2009…
- FLUID-STRUCTURE APPLICATIONS IN AUTOMOBILE
  - OMD2 (ANR) : RENAULT, TAO, OASIS, ECN, EMSE, ENS Cachan, UTC, SIREHNA, ACTIVEON…

Outline

INRIA

Example Scenario

Multidiscipline challenges

HPC Contributions

Workflows

Conclusions

To summarize

- Workflows are formal specifications of scientific processes: data collection, analysis, computation & publication
- Support sharing scientific processes descriptions across organizational and discipline bounderies
- Ability to evolve processes, compare, identify differences
- Solve heterogeneity in large-scale multiphysics apps
  - implementations: application orchestration & choreography
  - computing environments: supercomputers, large clusters, …

The snowball effect…

CONCLUSION

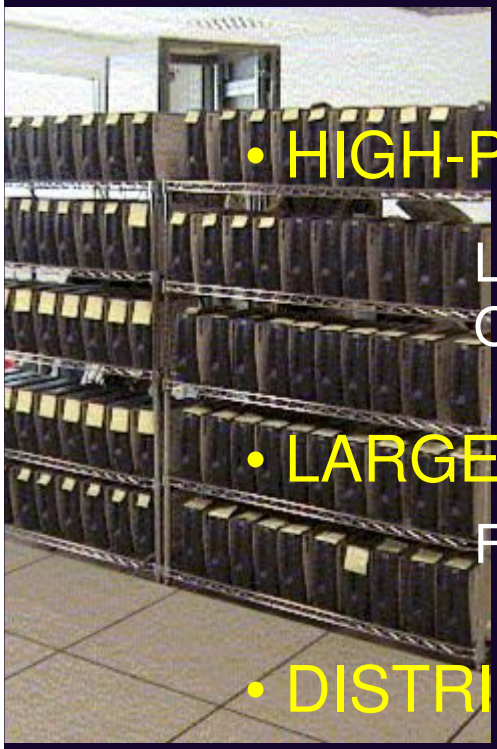

- HIGH-PERFORMANCE COMPUTING PLATFORMS

  LARGE SCALE MULTIDISCIPLINE APPLICATIONS ON COLLABORATIVE ENVIRONMENTS

- LARGE-SCALE MULTI-DISCIPLINE OPTIMIZATION

  FLEXIBLE, RESILIENT, EFFECTIVE

- DISTRIBUTED PLATFORMS

  RE-CONFIGURABLE, FAULT-TOLERANT

CONCLUSION

TOWARDS VIRTUAL FLIGHT-TESTS

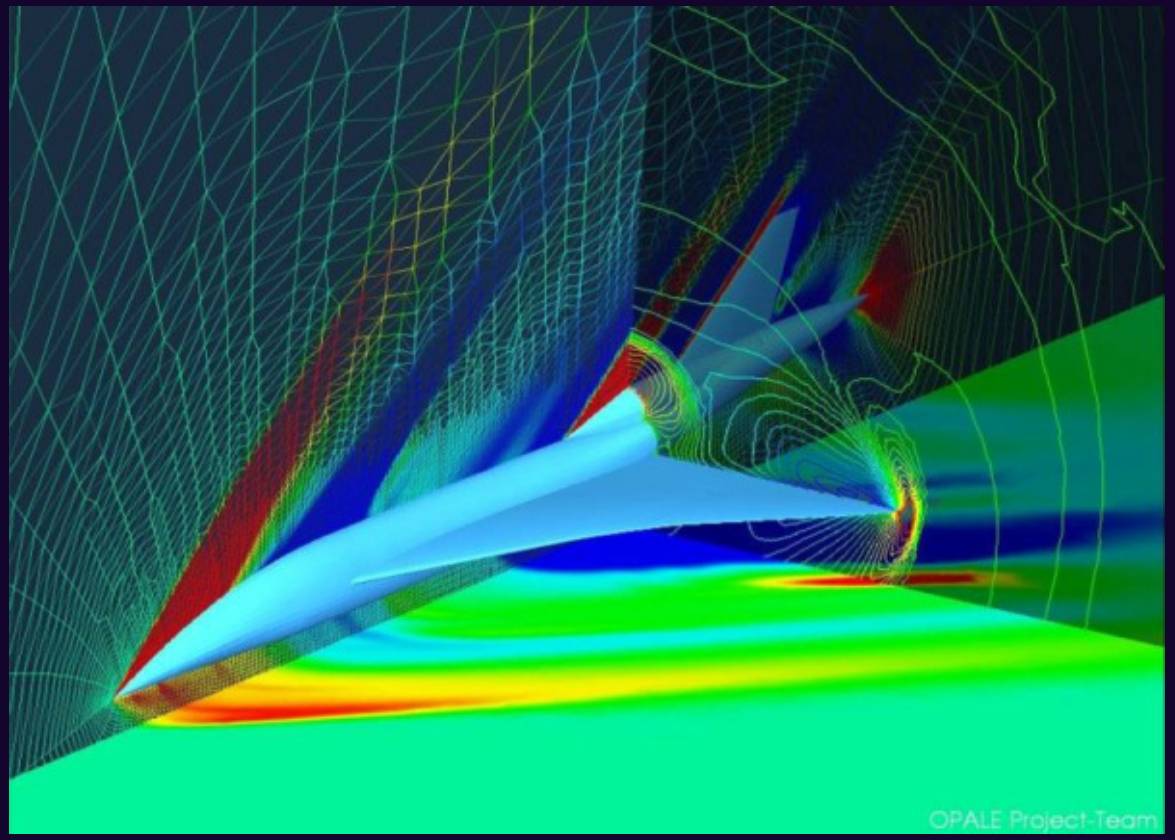

HIGH-PERFORMANCE COLLABORATIVE PLATFORMS

« THE DIGITAL DYNAMIC AIRCRAFT »

REFERENCES

http://www-opale.inrialpes.fr

- **http://www.inrialpes.fr/opale/coupled2009.pdf**
- **http://www.inrialpes.fr/opale/CSCWD2008.pdf**
- **http://www.inrialpes.fr/opale/ECCOMAS2008.pdf**
- **http://www.inrialpes.fr/opale/CDVE2008.pdf**

Phone : +33 4 76 61 52 40

Toan.Nguyen@inrialpes.fr

Links

- IEEE Computer, December 2007,
  Examining the Challenges of Scientific Workflows
- NSF Workshop, Arlington (Va), 1-2 May 2006
  http://vtcpc.isi.edu/wiki/index.php/Main_Page
- Ter@Tec Initiative (France)
  High-Performance Numeric Simulation
  http://www.teratec.fr/
- LARGE-SCALE DISTRIBUTED INFRASTRUCTURES
  Open Science Grid, TeraGrid, NMI, NAREGI, PRACE…
- IEEE Technical Committee on Scalable Computing
  Workflow Management in Scalable Computing Env.
  http://www.swinflow.org/tcsc/wmsce.htm

Links

- vl-e
  University of Amsterdam, http://www.vl-e.nl
- PRACE Partnership for Advanced Computing in Europe
  Persistent HPC service & infrastructure
  http://www.prace-project.eu
- CCGRID'09 Cluster Computing & Grid Conf.
  Lyon (F) May 2008, Shanghai May 2009, Melbourne 2010
  http://www.ens-lyon.fr/LIP/RESO/ccgrid2008/
  http://grid.sjtu.edu.cn/ccgrid09/
- e-Science 2008 4th IEEE Intl Conf.
  University of Indiana, Indianapolis, December 2008
- Grid Applications and Middleware Workshop
  Wrocaw, Poland (GAMW'2009), Sept. 2009