My Background to explain my bias :

Distributed and Parallel Objects

Denis Caromel

OASIS Team, France

INRIA -- CNRS - I3S -- Univ. of Nice Sophia-Antipolis, IUF

SOS 10, Maui, March 6-9

ProActive
Programming, Composing, Deploying on the Grid

A Java API + Tools for Parallel, Distributed Computing

A uniform framework: **An Active Object pattern**

A formal model behind: **Determinism (POPL'04)**

Programming Model:

- **Asynchronous Remote Objects, Wait-By-Necessity**
- **Groups, Mobility, Components, Security, Fault-Tolerance: Checkpoints**

Environment:

- **XML Deployment Descriptors, File Transfers**
- **Interfaced with: `rsh, ssh, LSF, PBS, Globus, Jini, SUN Grid Engine`**
- **Graphical Visualization and monitoring: IC2D**

In the www. ObjectWeb .org Consortium

(Open Source **LGPL**)

ObjectWeb
Open Source Middleware

JEM 3D : Java 3D Electromagnetism

together with Said El Kasmi, Stéphane Lanteri (caiman)

Maxwell 3D equation solver, Finite Volume Method (FVM)

Pre-existing Fortran MPI version: EM3D (CAIMAN team @ INRIA)

Up to 294 machines at the same time (Intranet and cluster)

Large data sets: 150x150x150 (100 million facets)

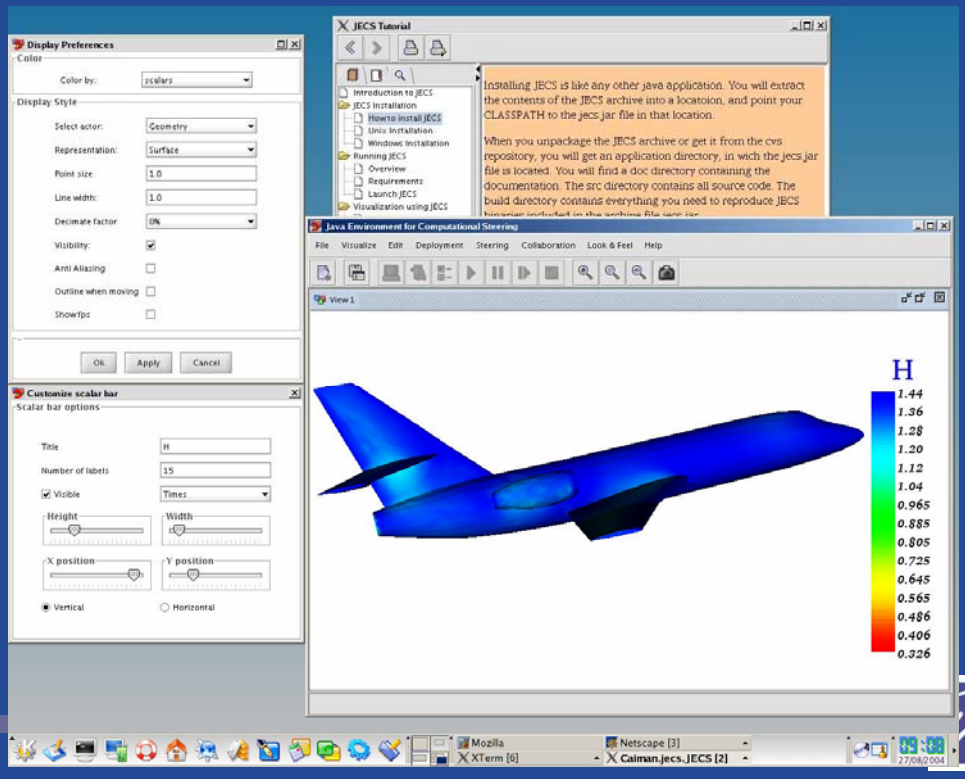

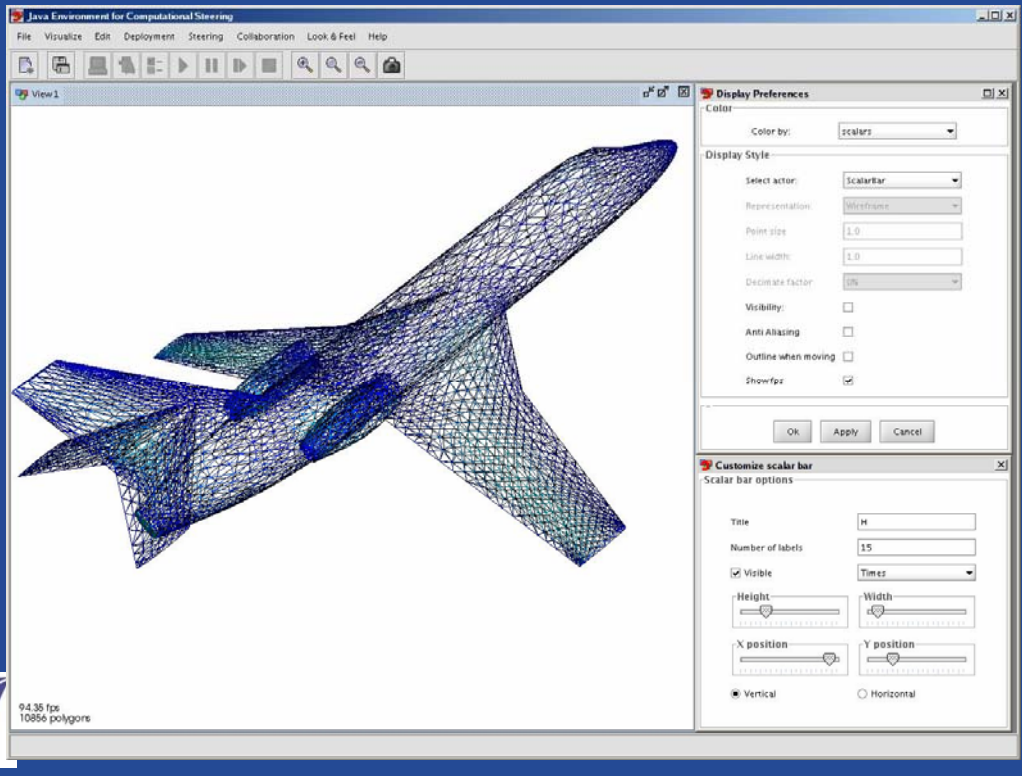

Beating Fortran-MPI ?

Facts:

- Sequential Java vs. Fortran code: 2 times slower
- Large data sets in Java ProActive: 150x150x150 (100 million facets)
- Large number of machines: up to 294 machines in Desktop P2P
- Speed up on **`16 machines:`**
  - Fortran: 13.8
  - ProActive/Ibis: 12
  - ProActive/RMI: 8.8

Grid on 5 clusters (DAS 2): Speed up of 100 on 150 machines

Fortran: no more than 40 proc. …

Beating Fortran MPI with Java ProActive? X/40 (14/16) = 2X/ n (100/150)

Yes, starting at n=105 machines !

Data Intensive: What are the Driving Applications ?

Mainly:

- Sciences and Military, etc.
- Bio Technologies
- Business Data-Mining, Finance-Banking, Insurance

➔ Where is the money ?

➔ Where will be the investment ?

- Consequence:
  - Should fit in standard business IT
  - Portability is the key (not dedicated systems)

Relation Intensive Computation and I/O: One-Sided Operations

Remotely triggered operations:

- MPI one-sided
- X10 asynch { … }
- ProActive Immediate Services
- Active Messages

In the context of Multi-Core Processors:

- might be very important for fast implementations of I/O

Data Intensive Computing should define in the future as:

Using:

- Large scale machines and I/O capabilities
- Dedicated implementation of standard I/O interfaces

But:

- No specific paradigms for application programmer
- No specific user-interfaces

ProActive:
A Java API + Tools for Parallel, Distributed Computing

- A uniform framework: **An Active Object pattern**
- A formal model behind: **Prop. Determinism, insensitivity to deploy.**

Main features:

- **Remotely accessible Objects (Classes, not only Interfaces, Dynamic)**
- **Asynchronous Communications with synchro: automatic Futures**
- **Group Communications, Migration (mobile computations)**
- **XML Deployment Descriptors**
- **Interfaced with various protocols:** `rsh,ssh,LSF,Globus,Jini,RMIregistry`
- **Visualization and monitoring: IC2D**

In the www. ObjectWeb .org Consortium (Open Source middleware)
since April 2002 (**LGPL license**)

What is the status of uptake and exploitation ?

About 10 000 hits / month

About 400 Downloads / month

About 4 000 recorded Downloaders

Among which 400 Industrials

IC2D: Interactive Control and Debugging of Distribution

Main Features:

- Hosts, JVM,
- Nodes
- Active Objects
- Topology
- Migration

- Logical Clock

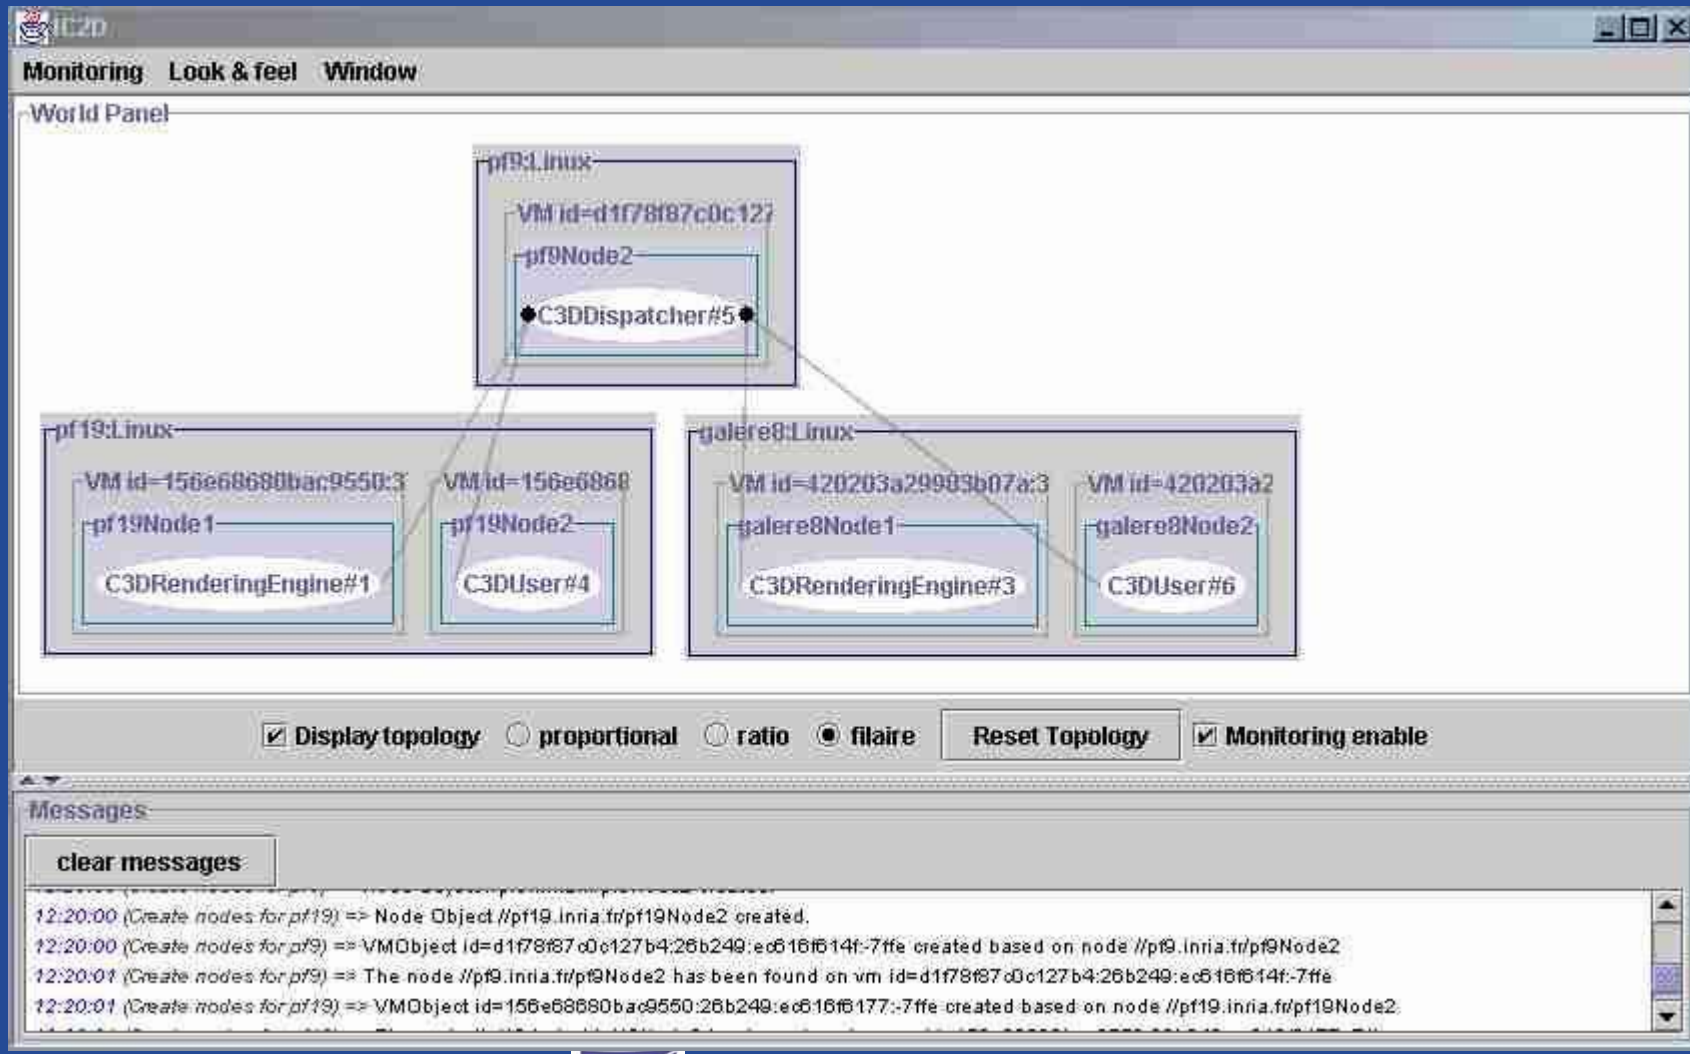

Mobile Application executing on 7 JVMs

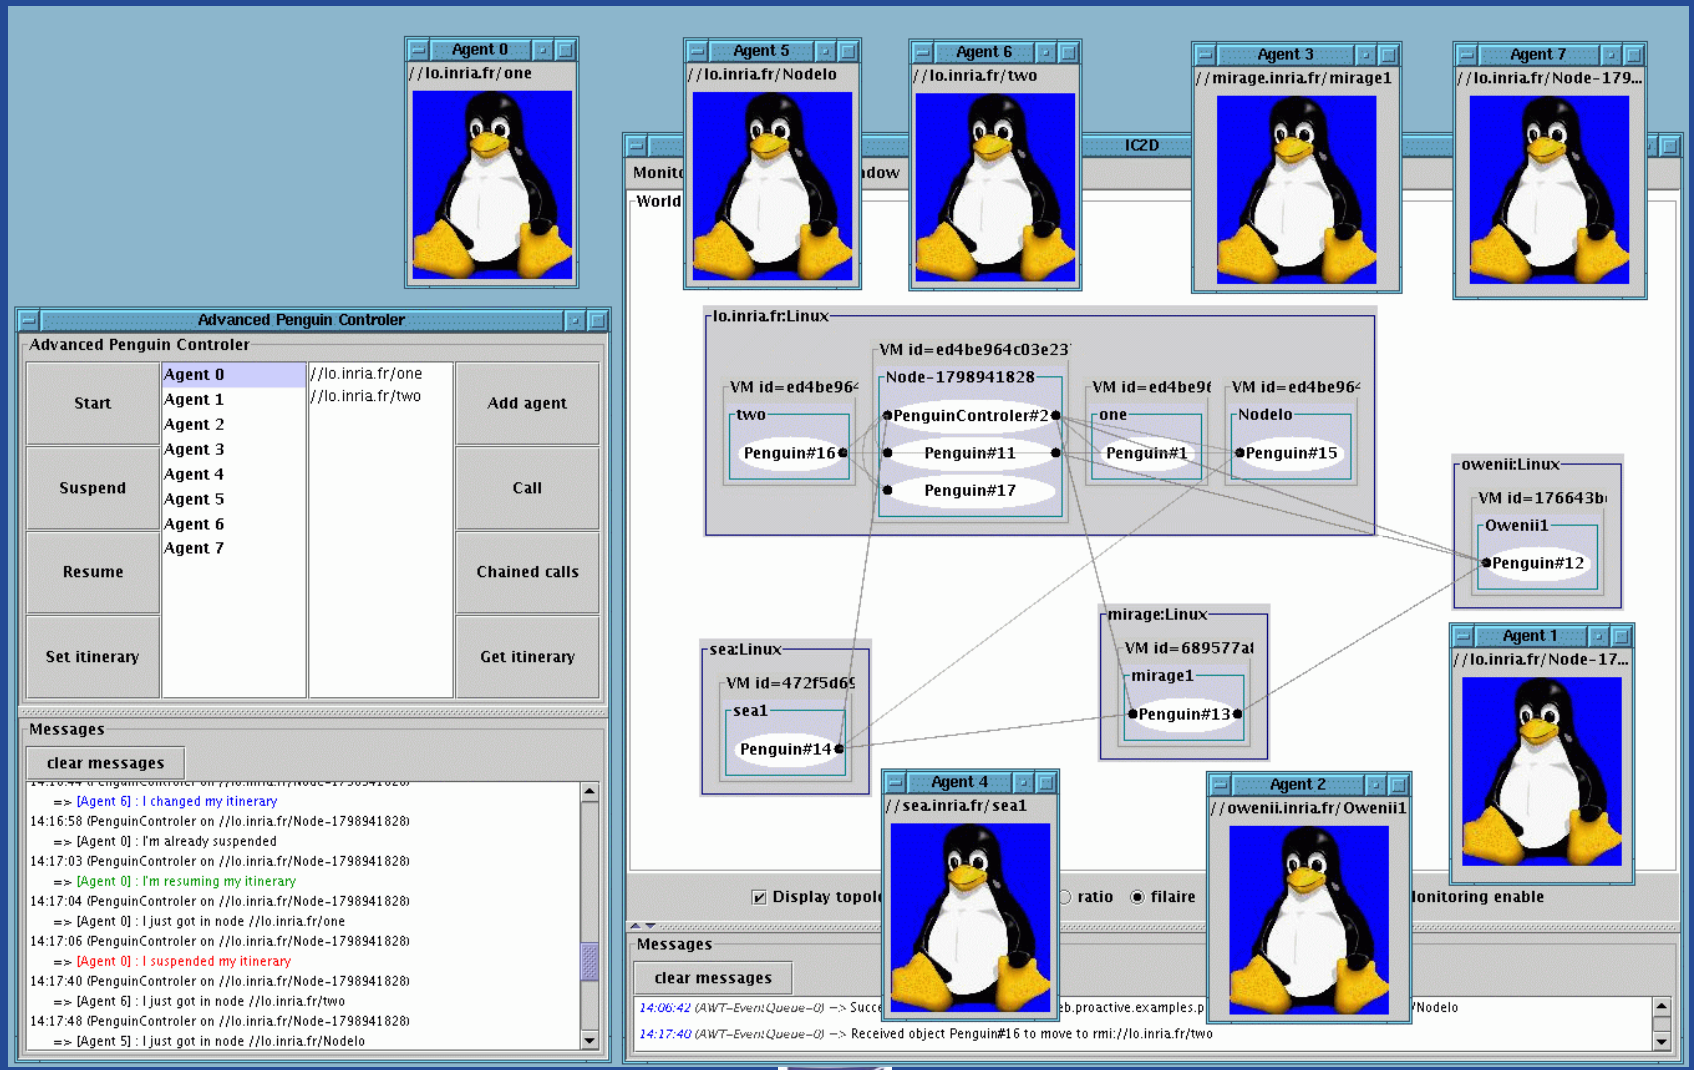

Monitoring of RMI, Globus, Jini, LSF cluster Nice -- Baltimore at SC'02

Width of links proportional to the number of communications

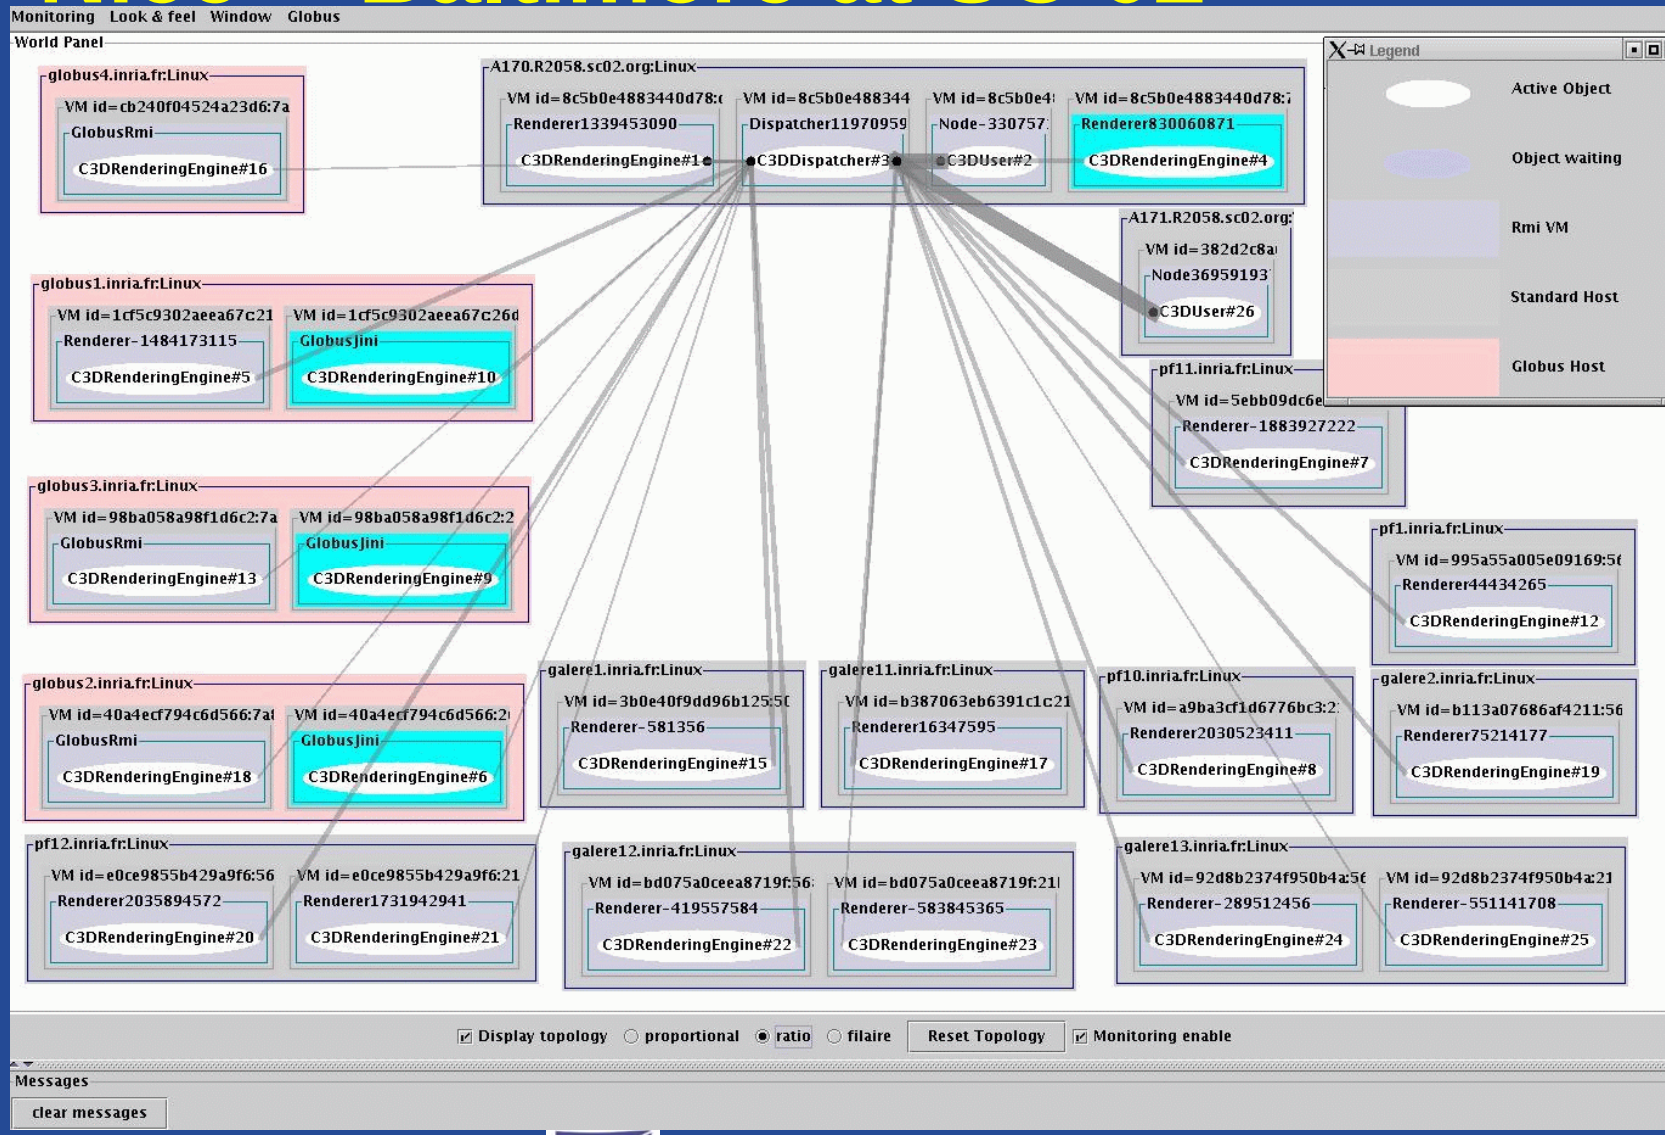

Which success stories do we have?

World record : 53 years of computations in 6 months

Distributed Financial Applications

GRID PlugTests (2004, 2005): Deploying on 2700 CPUs worldwide


ObjectWeb
Open Source Middleware
ProActive
Programming, Composing, Deploying on the Grid
PLUGTESTS
Amsterdam
Belfast
Beijing
Benevento
Fribourg
Heraklion
Hong Kong
Juelich
Lille
Los Angeles
Metz
Melbourne
Nancy
Oslo
Pisa
Shanghai
Rio de Janeiro
Zurich

Which success stories do we have?

World record : 53 years of computations in 6 months

Distributed Financial Applications

GRID PlugTests (2004, 2005): Deploying on 2700 CPUs worldwide

Hard: x86_32, x86_64, ia64, AIX, SGIIrix, PPC, Sparc

OS: Linux, AIX, SGIIrix, MacOS, Solaris

Interoperability: EGEE gLite, Unicore, NorduGrid, LSF, PBS, Ibis/… Globus

Open Source Middleware for the GRID

Programming Concurrent, Parallel, Distributed Applications with Active Objects

Denis Caromel, et al.

ProActive.ObjectWeb.org or GOOGLE ProActive

OASIS Team

INRIA -- CNRS - I3S -- Univ. of Nice Sophia-Antipolis, IUF

IEEE ISSADS Symposium and School,
Guadalajara, Mexico, Jan. 23-27 2006

1. GRIDs
2. Distributed Objects – Components
3. Application and Benchmarks
4. Conclusion – Perspective: Adaptive