

GridCOMP – ProActive/GCM tutorial and and Hands-On Grid Programming

Cédric Dalmaso, Antonio Cansado and Denis Caromel

INRIA - OASIS Team

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

IV Grid@Work, Tsinghua University, Beijing



General agenda

- Talk: A short introduction to ProActive middleware
- Practical session: ProActive fundamentals
- Talk: ProActive / GCM
- Practical session: ProActive / GCM
- Talk: IDE
- Talk: Autonomic
- Practical session: Autonomic



Short Introduction to *ProActive*



Agenda (Update It)

- Overview
- Programming
- Deploying
- What else?
- GUIs and tools
- Applications
- Conclusion



Overview



The team : 30+ members

- OASIS Team at INRIA in Nice, France
- Joint team INRIA / CNRS / Univ. Nice
- Team leader: Denis Caromel
 - 3 professors
 - 2 researchers
 - 1 postdoc
 - 7 engineers
 - 7 PhD students
 - + Interns, visiting researchers...
- Collaborations: ObjectWeb, CoreGRID, **GridCOMP** etc..



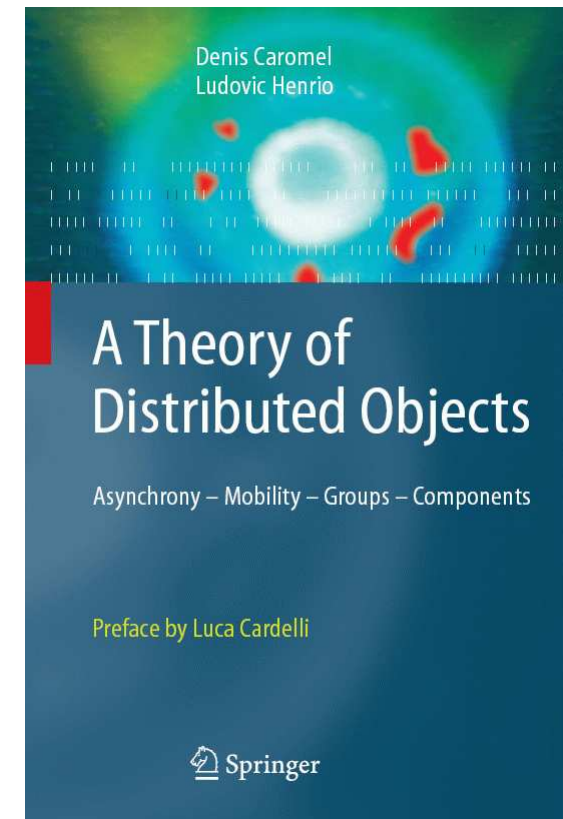
The library

- Originates from work on Eiffel //
 - Started in 1999
- Official releases ~ every 6 months
 - ProActive 3.2.1 released in April 2007
 - (Version 3.9 soon)
- Metrics:
 - 2000 classes, ~ 300.000 LOC (160.000 NCLOC)
- Compliant with several (*de facto*) standards
- ProActive startup: ActiveEon
 - Training, Consulting, Integrating and Support

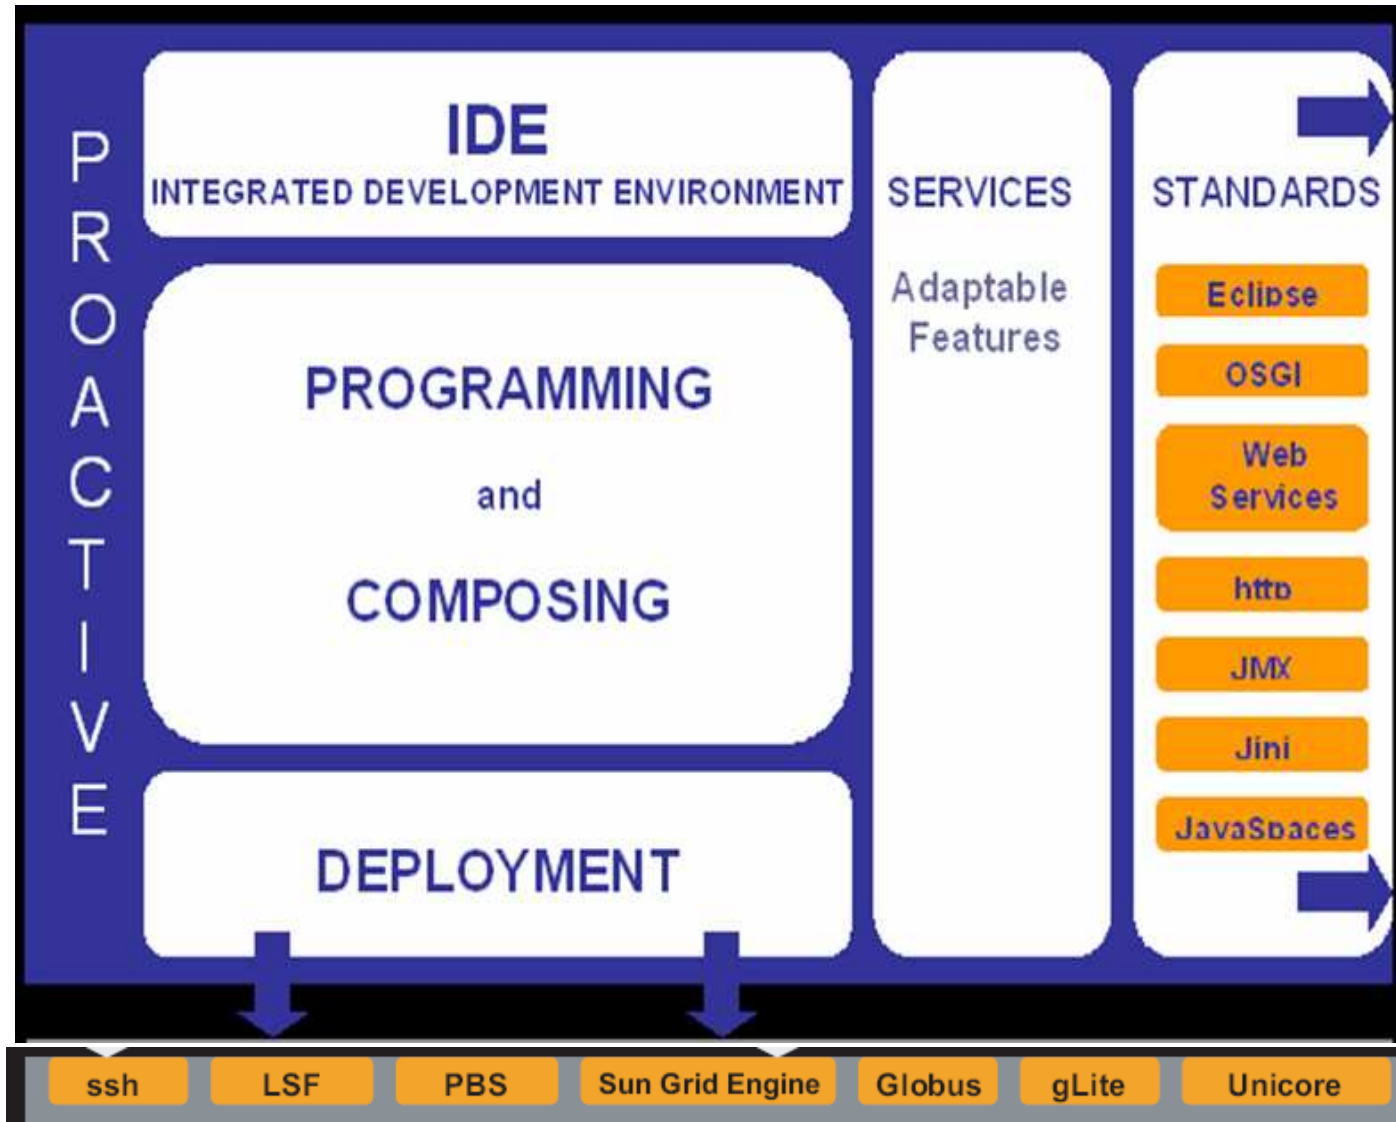


Theory

- Henrio & Caromel
- ASP Calculus:
 - Asynchronous Sequential Processes
 - Based on Sigma-Calculus (Abadi-Cardelli)
- Formal Proofs of **determinism** (in greek)
- Implemented in ProActive



ProActive's Framework in a nutshell

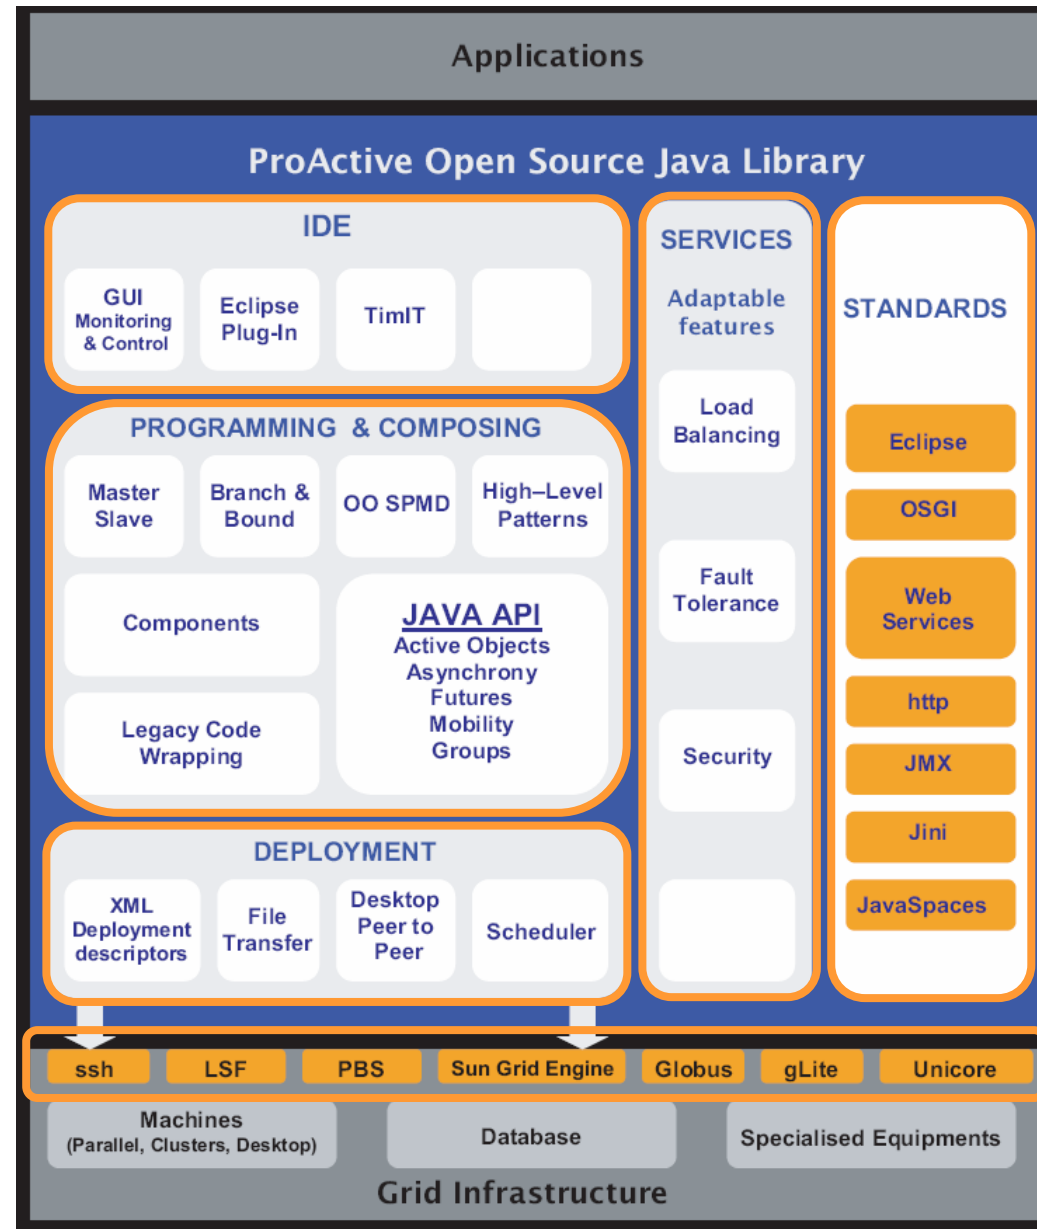


Open
Source
+
PROFESSIONAL
SUPPORT



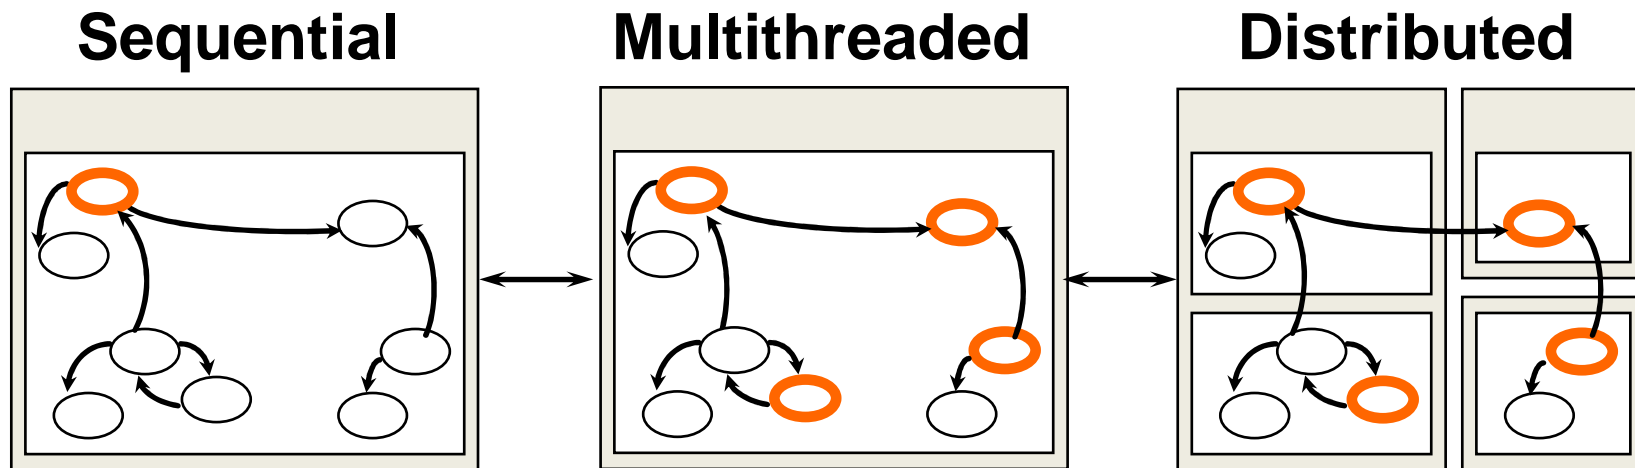
Inside ProActive

- IDE
- PROGRAMMING & COMPOSING
- DEPLOYMENT

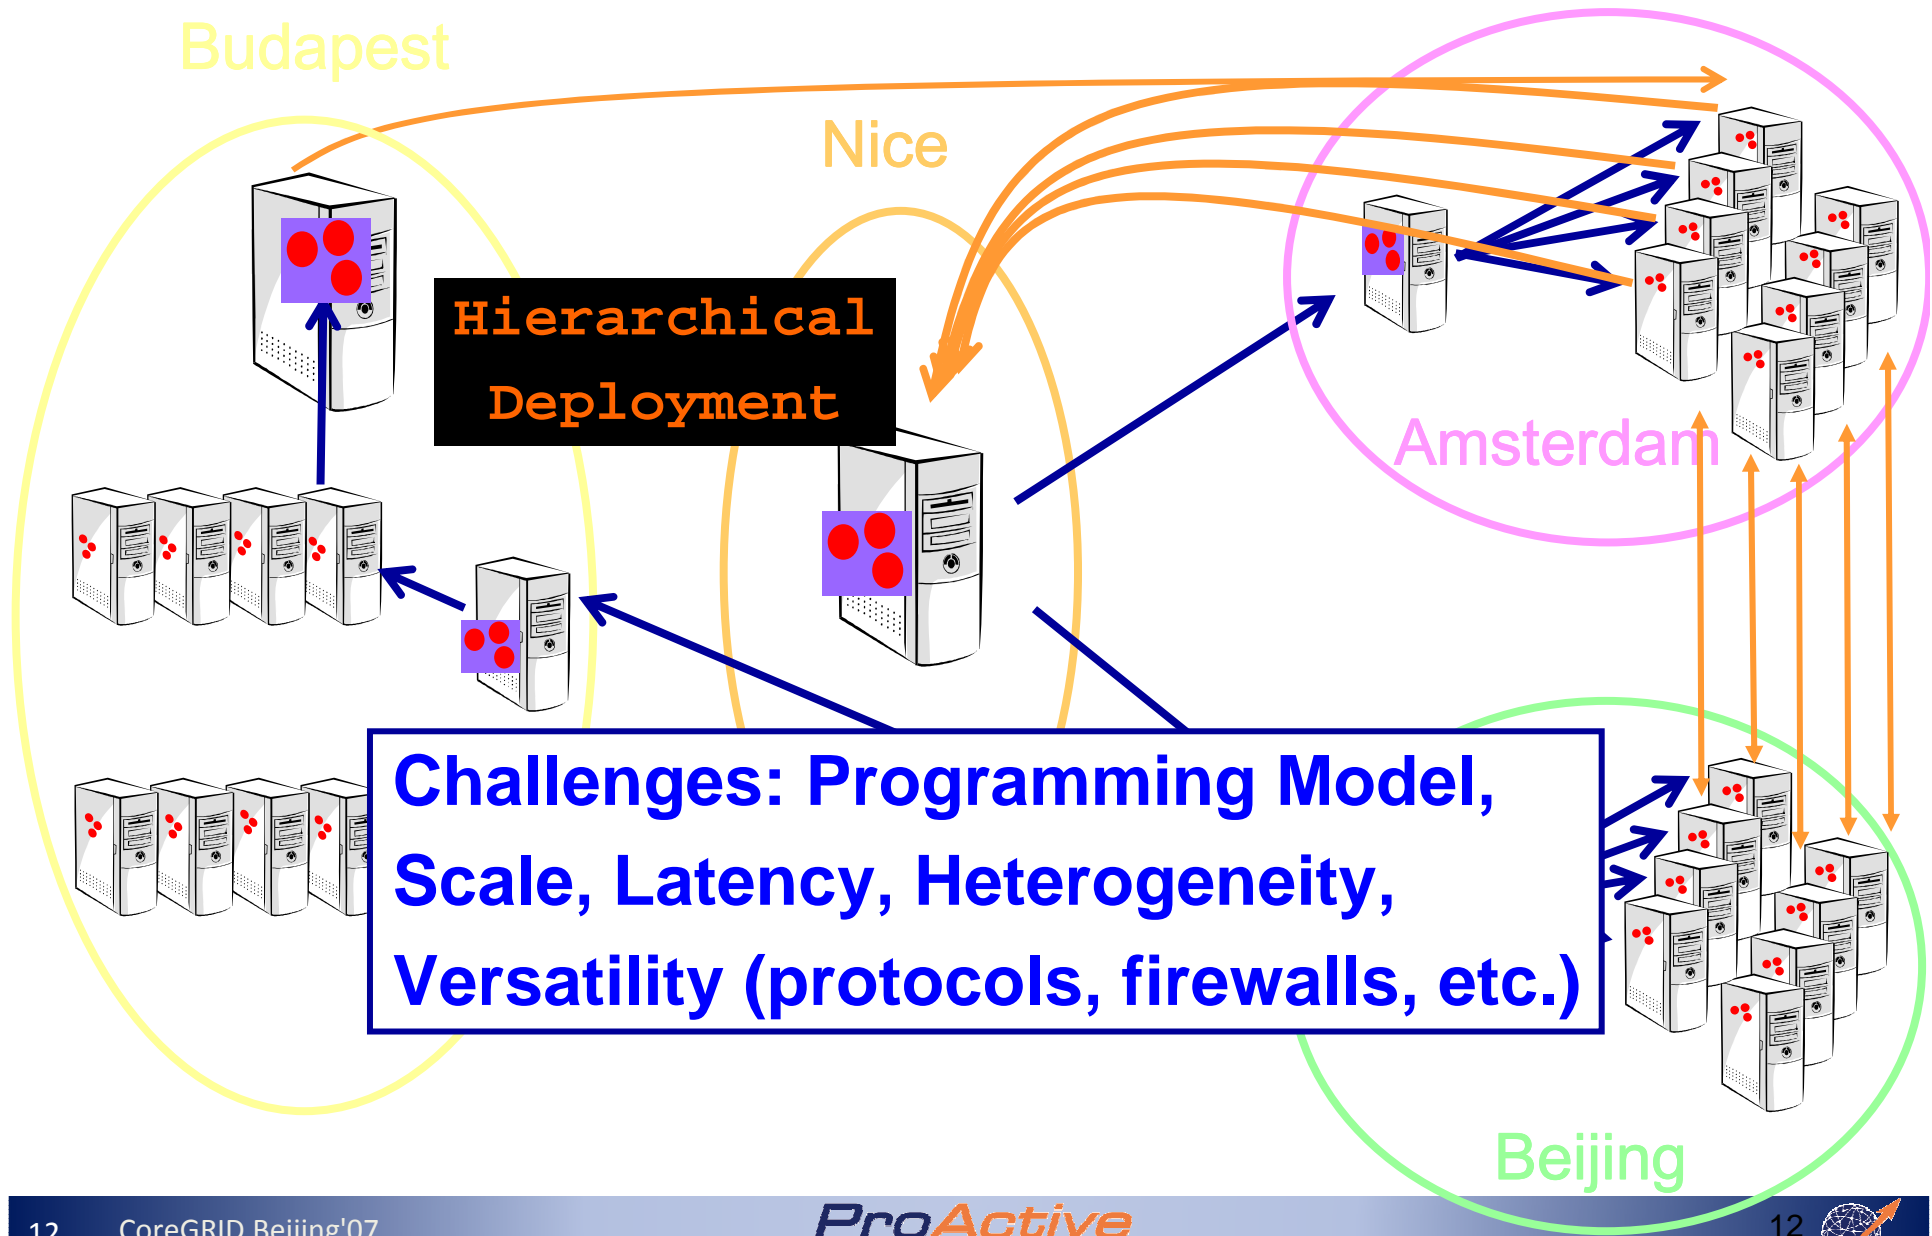


Rationale

- Distributed programming entities
- Parallel processes
- Asynchronism
- Synchronization facilities



Grid Computing with *ProActive*



Programming



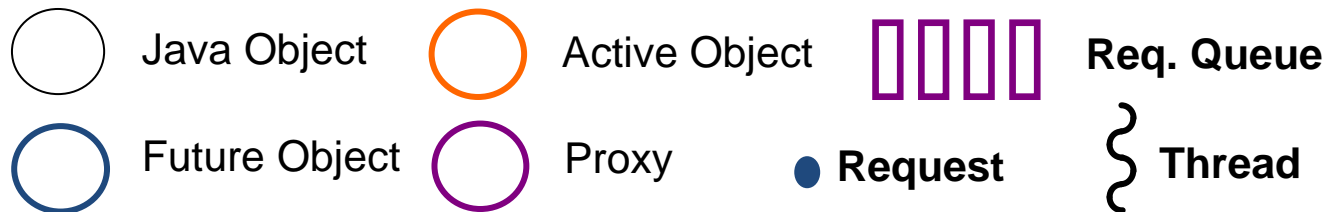
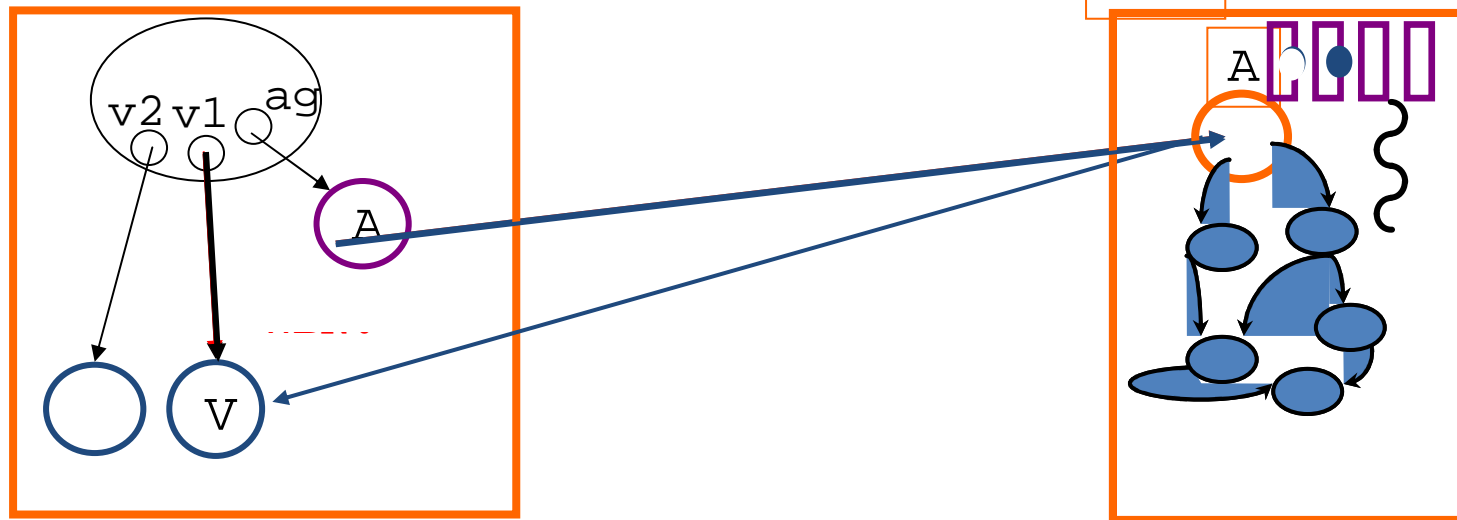
ProActive: Model

- Active objects : structuring entities (subsystems)
 - Passive objects (fields)
 - 1 thread / AO
 - Request queue
- Full control to serve incoming requests (**reification**)
- Sequential processing
- Typed entities (safe)
- **Asynchronous Communication** between active objects
- **No shared** passive **objects** - deep-copy of parameters

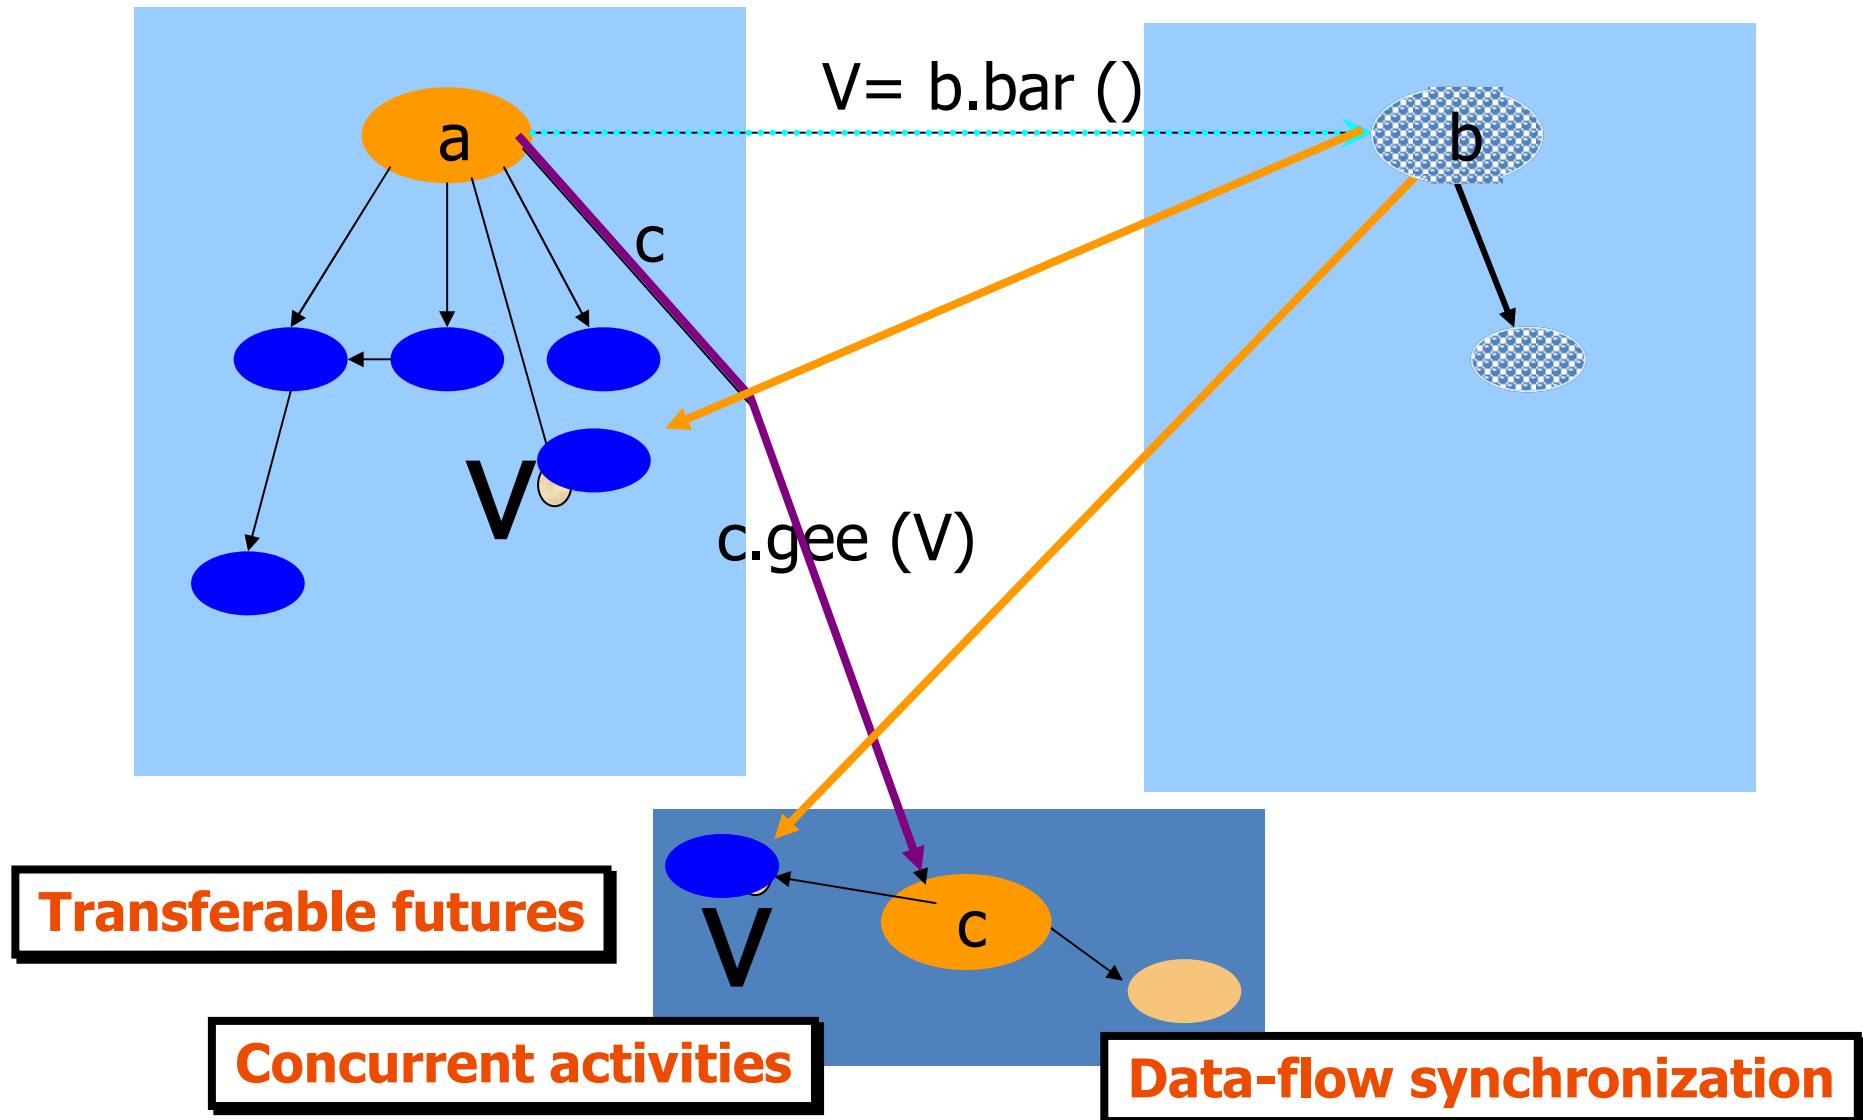


Creation, Invocation and Sync.

- A ag = **newActive** ("A", [...], VirtualNode)
- V v1 = ag.foo (param);
- V v2 = ag.bar (param);
- ...
- JVM** ● v1.bar(); //Wait-By-Necessity



Automatic Continuations



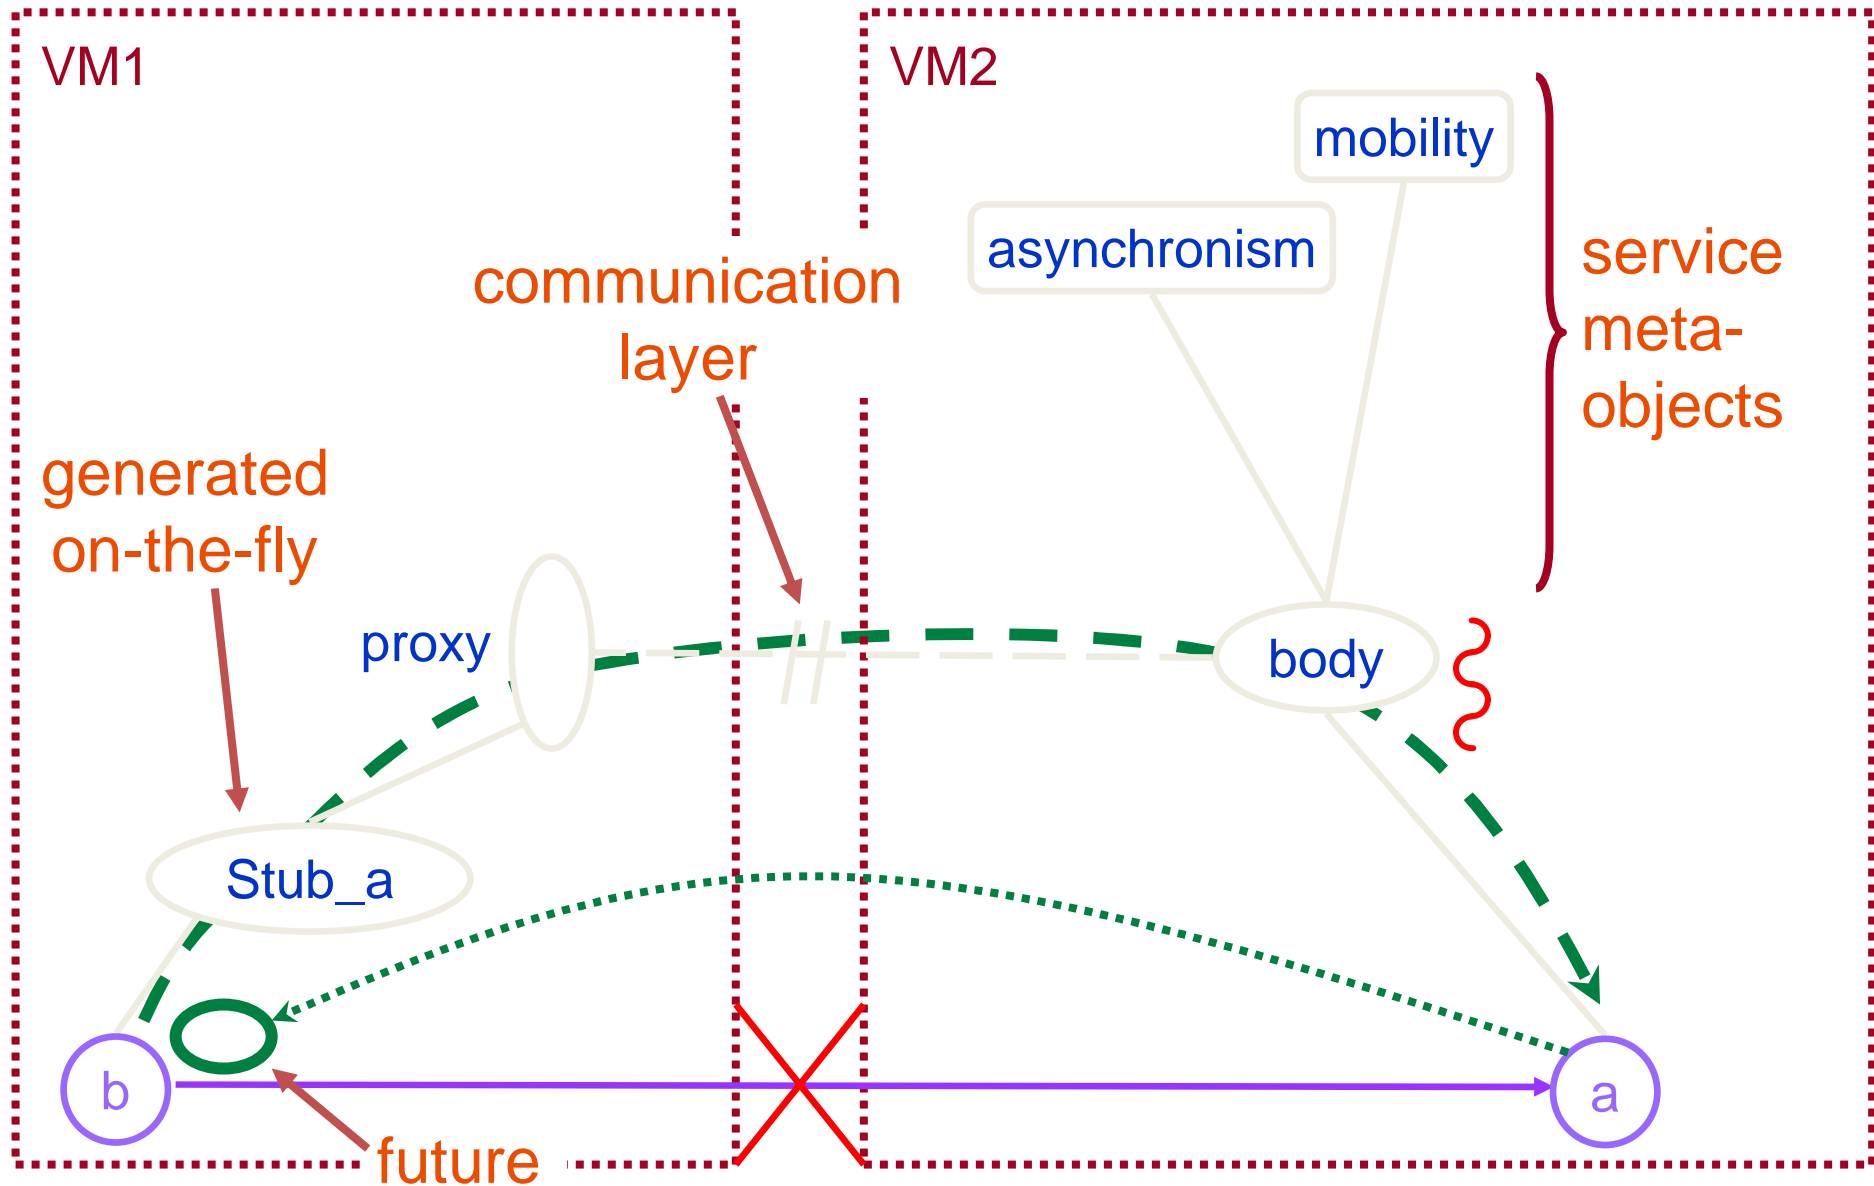
Explicit Synchronizations

```
A ag = newActive ("A", [...], VirtualNode)
V v = ag.foo(param);
...
v.bar(); //Wait-by-necessity
```

- **Explicit Synchronization:**
 - - ProActive.isAwaited (v); // Test if available
 - - ProActive.waitFor (v); // Wait until availab.
- **Vectors of Futures:**
 - - ProActive.waitForAll (Vector); // Wait All
 - - ProActive.waitForAny (Vector); // Get First



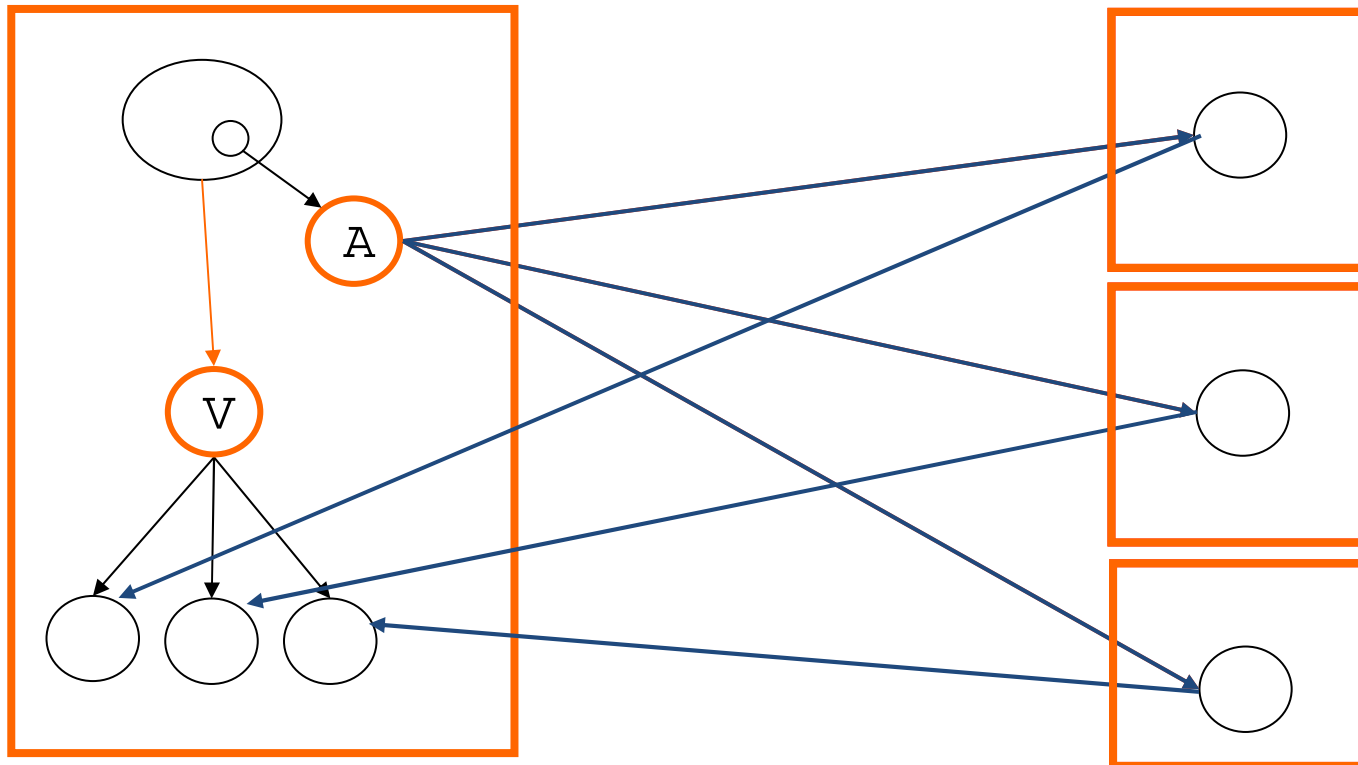
Architecture : a Meta-Object Protocol



Active Objects and Groups

- A ag = `newActiveGroup ("A", [...], VirtualNode)`
- V v = ag.foo(param);
- ...
- v.bar(); //Wait-by-necessity

JVM



○ Typed Group ○ Java or Active Object

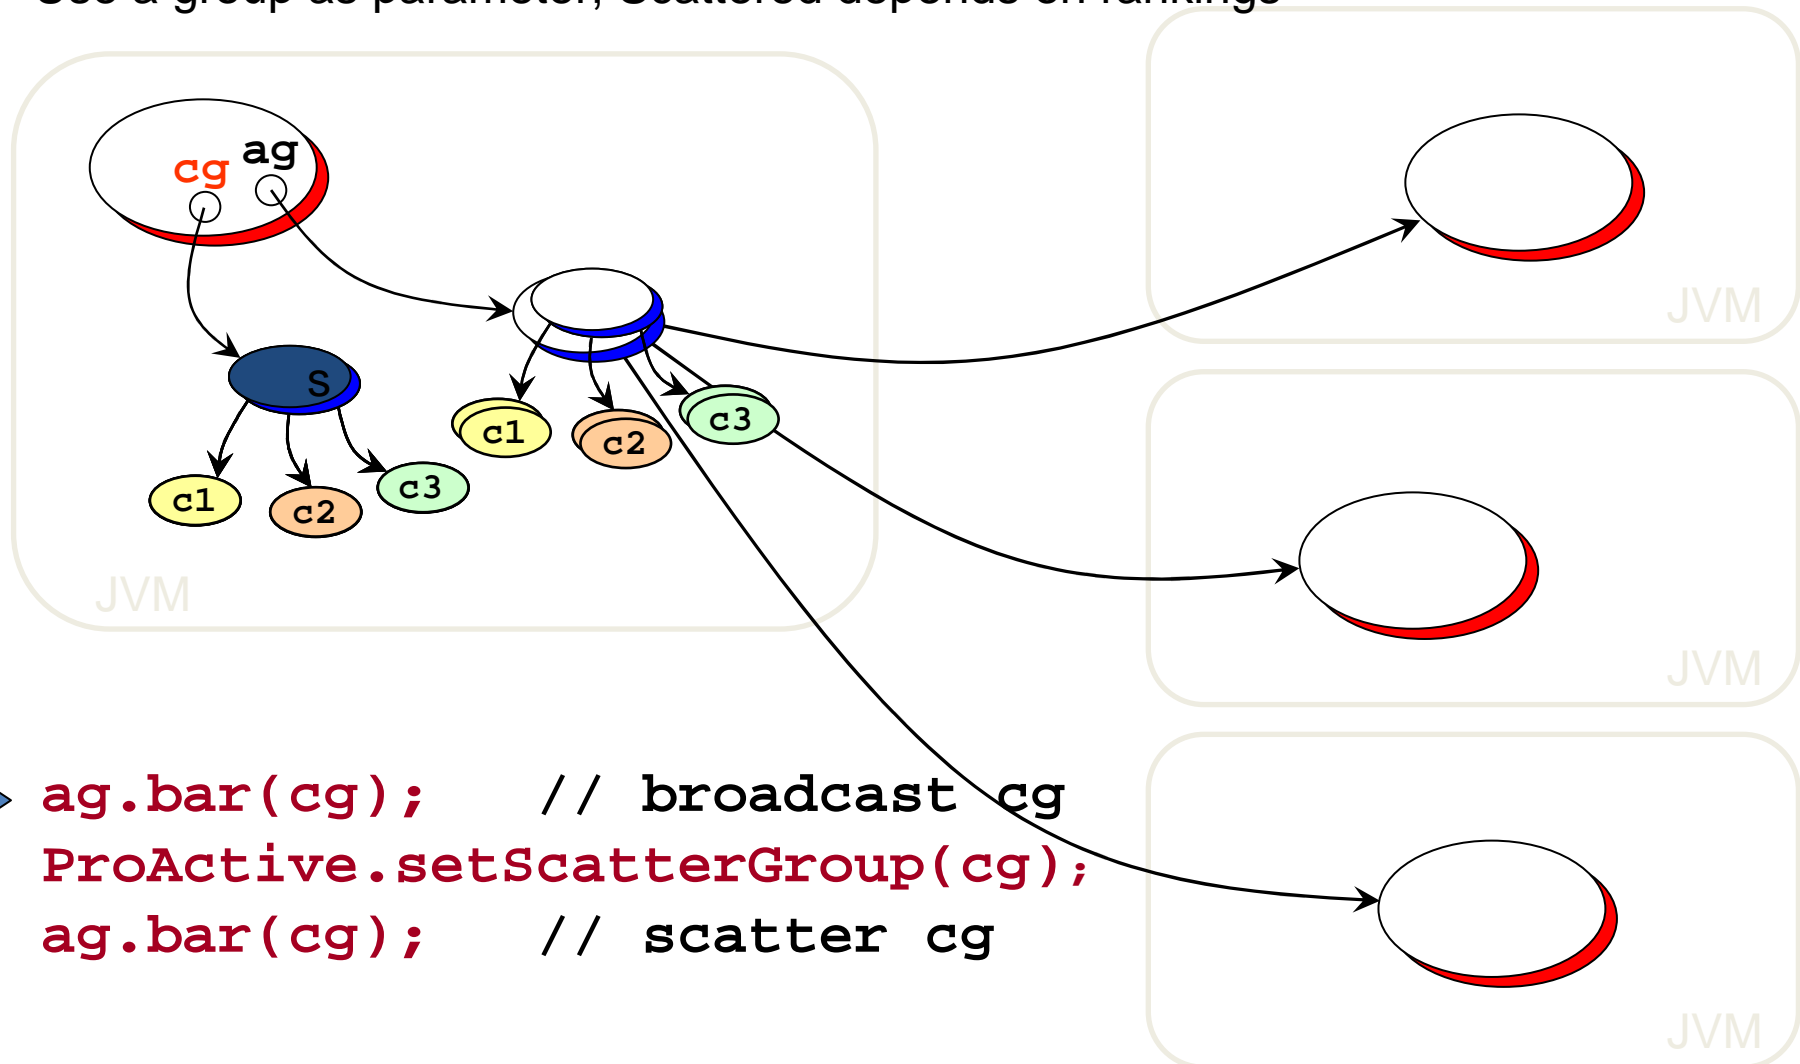
Group, Type, and Asynchrony are crucial for Cpt. and GRID



Broadcast and Scatter

Broadcast is the default behavior

Use a group as parameter, Scattered depends on rankings



➔ `ag.bar(cg); // broadcast cg`
`ProActive.setScatterGroup(cg);`
`ag.bar(cg); // scatter cg`

Deploying



Deployment : an abstract model

- Problem:
 - Heterogeneous environments/protocols
 - Scalability issues (large number of hosts / latency)
- A key principle:
 - Separate design from deployment infrastructure
 - Nothing about infrastructure, protocols or physical resources in the app. code

Creation Protocols

- ssh, gsissh, rsh, rlogin
- lsf, pbs, sun grid engine, oar, prun
- globus(GT2, GT3 and GT4), unicore, glite, arc (nordugrid)

Registry/Lookup and Comm. Protocols

- rmi, http, rmissh, ibis, soap

Files Transfers

- scp, rcp
- unicore, arc (nordugrid)
- other protocols like globus, glite will be supported soon

XML deployment file ⇒

Virtual Node (VN)



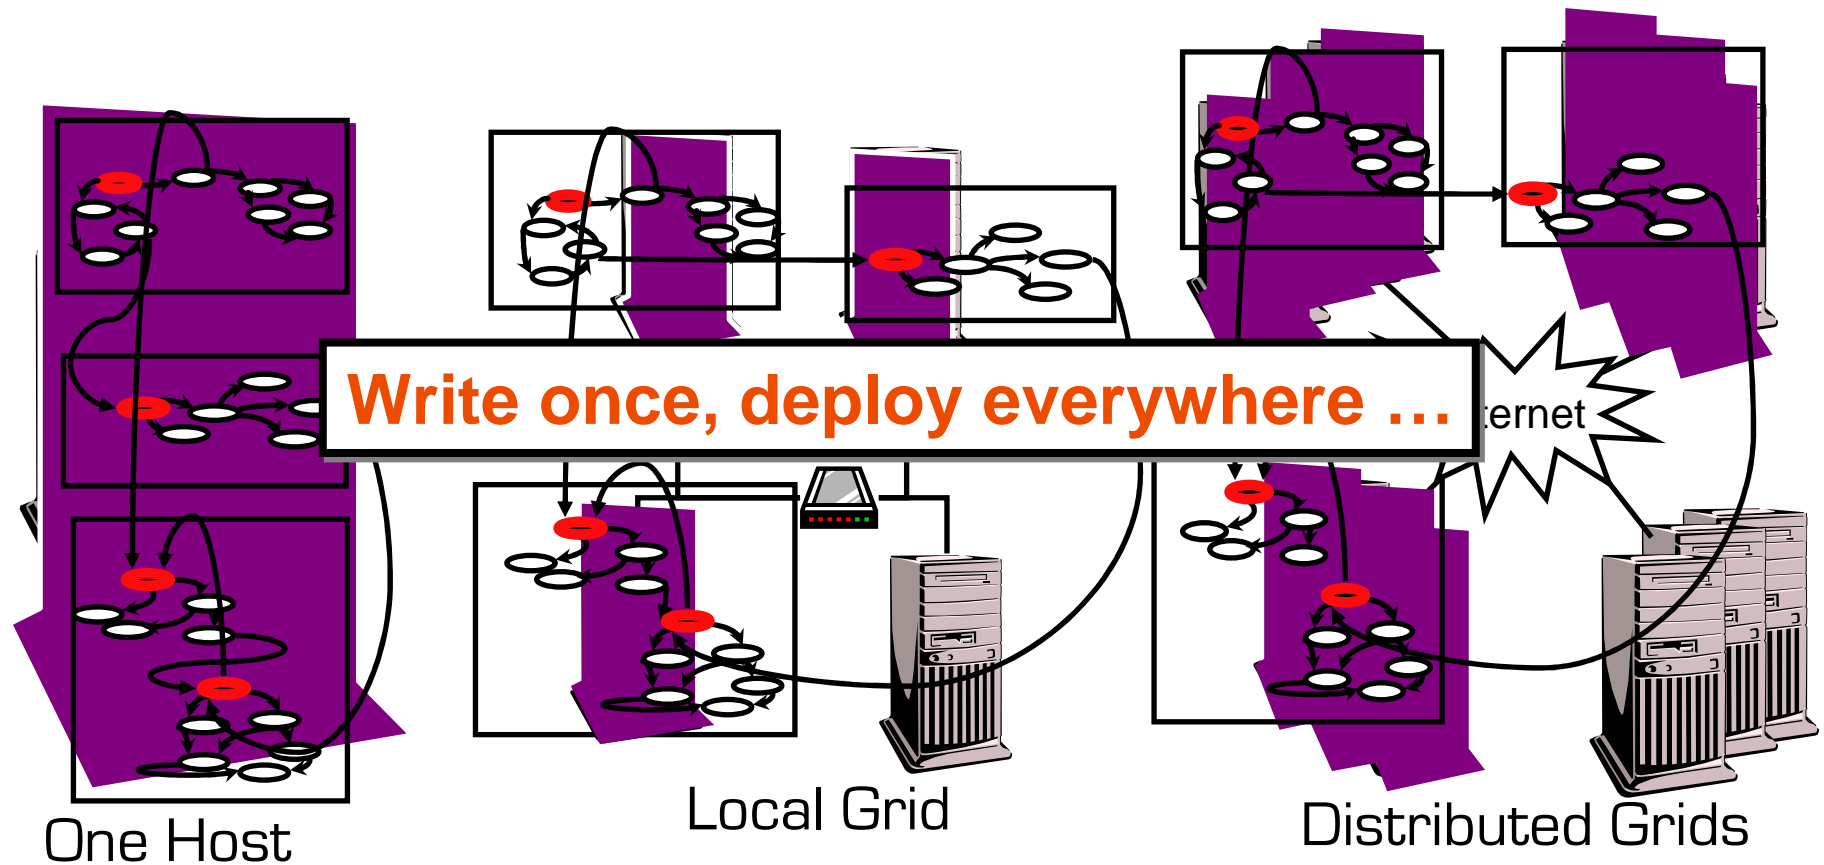
Activating a XML Desc.

-
- <virtualNodesDefinition>
- <virtualNode name=' Dispatcher' />
-

```
•ProActiveDescriptor pad = ProActive.getProactiveDescriptor(String xmlFile);
•// Returns a ProActiveDescriptor object from the xml file
•VirtualNode dispatcher = pad.getVirtualNode(' Dispatcher');
•// Returns the VirtualNode Dispatcher as a java object
•dispatcher.activate();
•// Activates the VirtualNode
•Node node = dispatcher.getNode();
•// Returns the first node available among nodes mapped to the VirtualNode
•
•C3DDispatcher c3dDispatcher = newActive('C3DDispatcher',param, node);
```



Same Application, Many Deployments



What else?



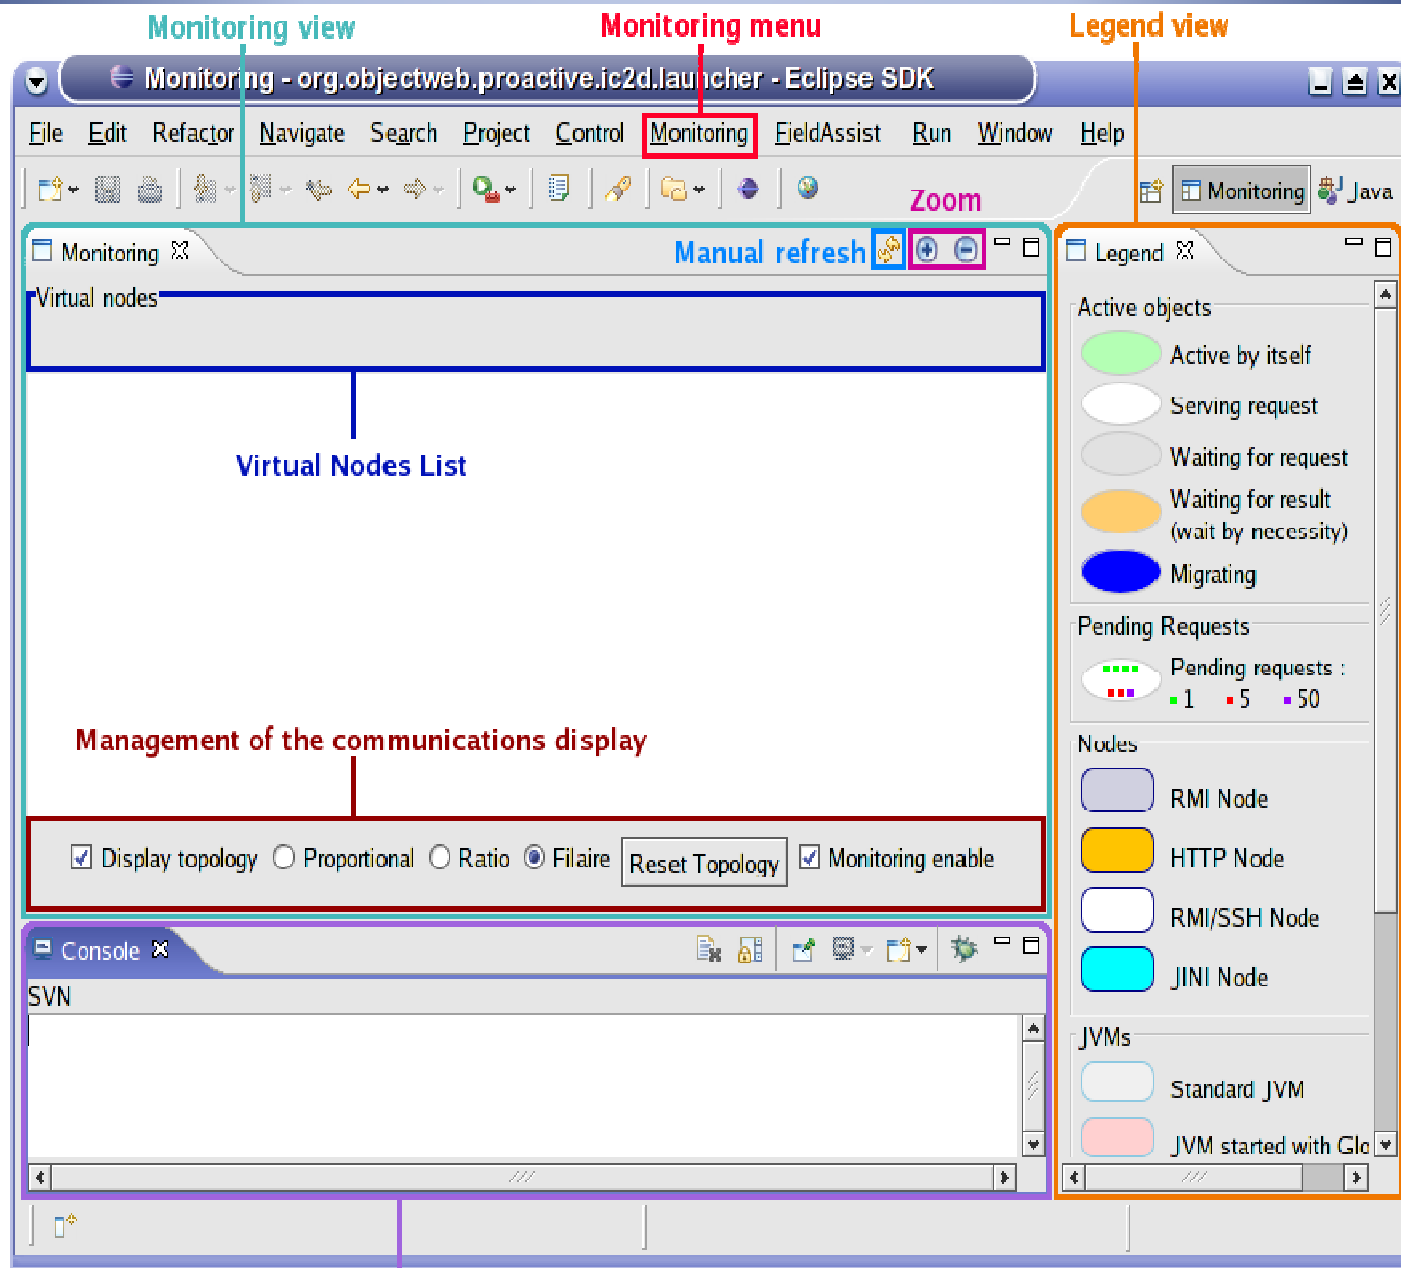
Other important features

- Component framework, the ProActive/GCM (just later!)
- P2P environment + Branch & Bound API
- Legacy code wrapping (MPI!)
- Middleware services:
 - Fault Tolerance
 - SOA integration (OSGI compliancy)
 - Migration
 - Load Balancing
 - Security



GUIs and tools





Monitoring View

Job Monitoring View

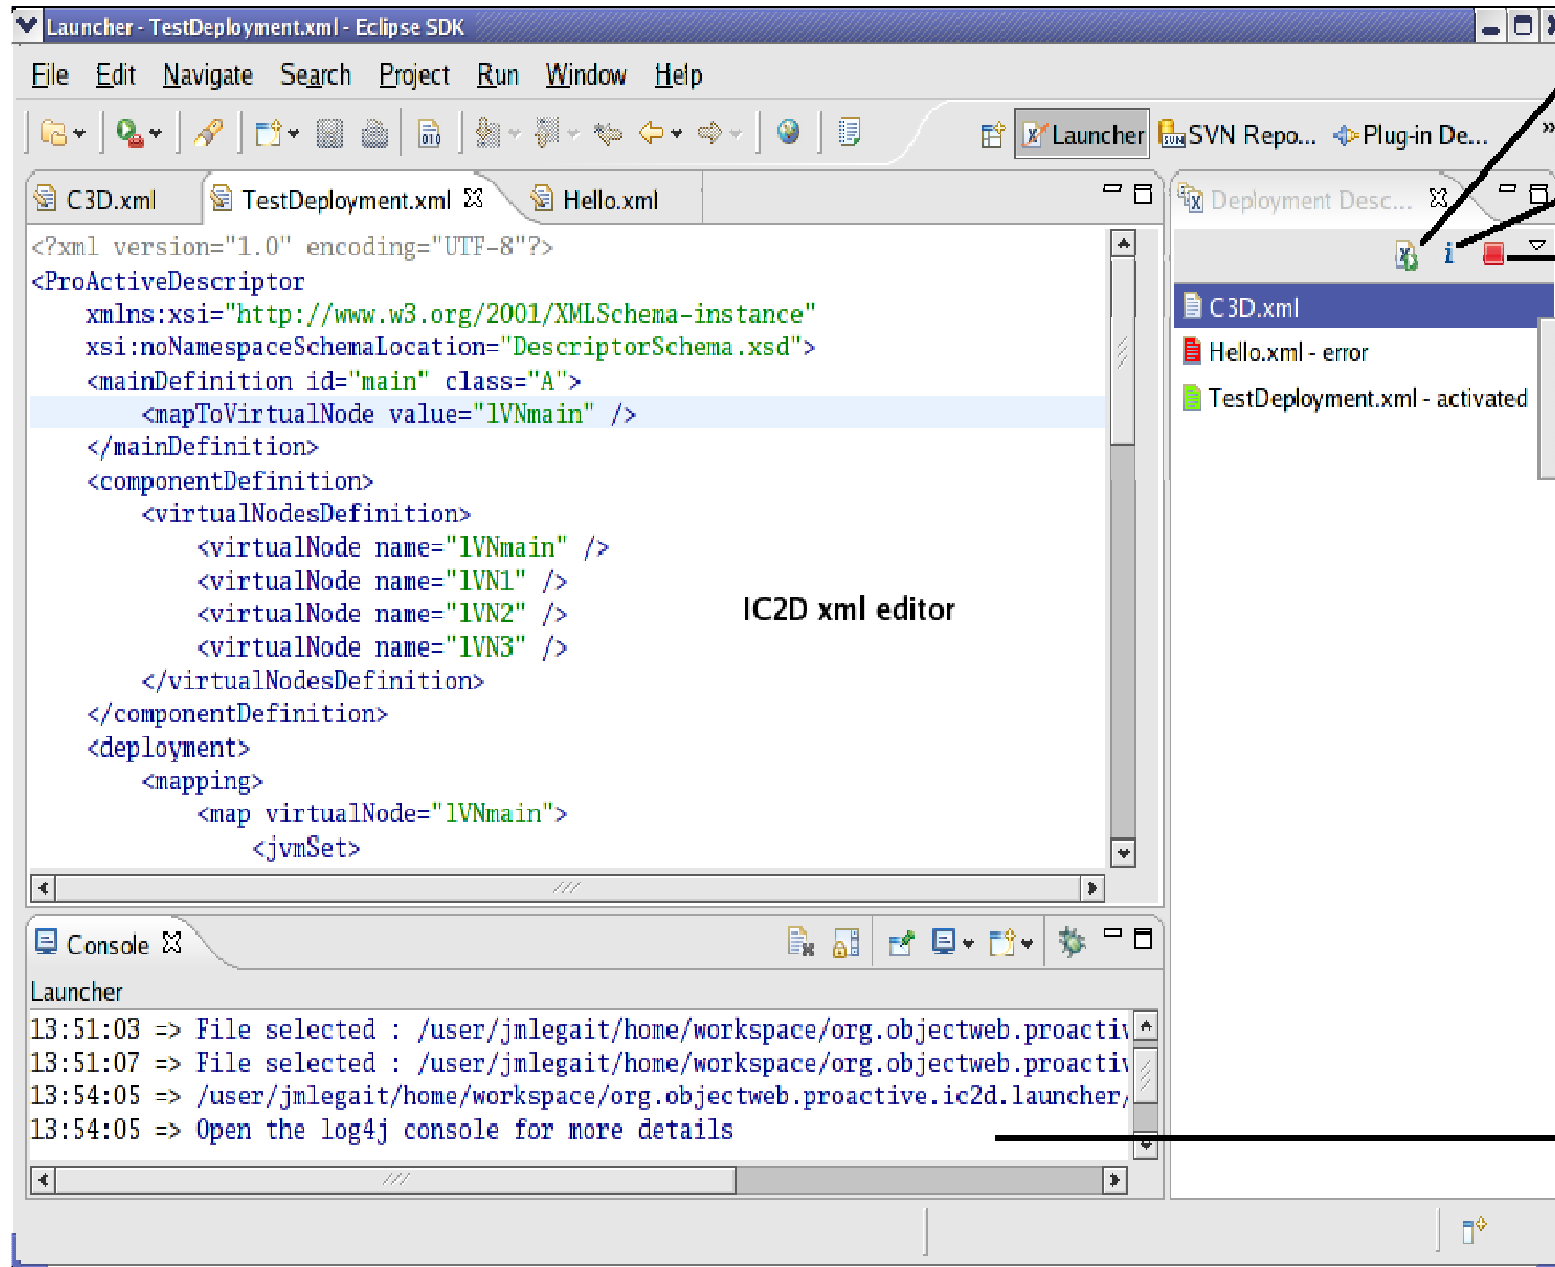
The screenshot displays the Eclipse SDK Monitoring interface, divided into two main sections: the Monitoring View and the Job Monitoring View.

Monitoring View: This view shows a network diagram of virtual nodes. At the top, there are checkboxes for 'Renderer', 'DefaultVN', 'Dispatcher', and 'User'. The main diagram shows a central node 'bebita.inria.fr:1099:OS u...' connected to several other nodes. On the left, a group of nodes includes 'DinnerLayout#2', 'Table#3', 'Philosopher#4', 'Philosopher#5', 'Philosopher#6', 'Philosopher#7', and 'Philosopher#8'. In the center, there are nodes for 'Node Renderer#127...', 'Node Dispatcher#5...', 'Node User#1602644...', and 'Node Renderer#1307...'. Below these are three more nodes: 'stuff.inria.fr:1099:OS und...', 'sidonie.inria.fr:1099:OS ...', and a red-bordered node 'Node Node-4551863...'. At the bottom, there are checkboxes for 'Display topology', 'Proportional', 'Ratio', 'Filaire', and 'Monitoring enable', along with a 'Reset Topology' button.

Job Monitoring View: This view shows a tree structure of jobs. The root is 'DefaultVN (JOB-135745762...)'. Underneath, there are several job entries, including 'bebita.inria.fr:1099:OS un', 'sidonie.inria.fr:1099:OS u', 'Dispatcher (JOB--167207649...', 'User (JOB--294719007)', and 'bebita.inria.fr:1099:OS un'. Each job entry has a sub-tree of nodes, such as 'Node Node60562496...', 'Node User1602644', and 'Node Node-4551863...'. The tree is expanded to show the details of the 'Node Node-4551863...' job.

Console: The console at the bottom shows a message: '15:09:15 => NodeObject id=Node-455186381 already monitored, ccheck for new active objects'.





Launches the deployment descriptor

Gives informations

Kills the application

IC2D xml editor

Launcher console



TimIt



Application Level Timer

The screenshot displays the IC2D monitoring interface. On the left, the 'Timlt View' shows performance metrics for four workers (Worker#1 to Worker#4) at a snapshot time of 18/06/07 15:27:18. Each worker's performance is shown as a horizontal bar chart with categories: computePI_rank, WaitForRequest, SendReply, AfterSerialization, Serialization, BeforeSerialization, SendRequest, Serve, and Total. Worker#1's chart is highlighted with a red box and an arrow pointing to it from the 'Application Level Timer' label.

On the right, the 'Monitoring#1' window shows a 'Virtual nodes' diagram. It lists several nodes: amda.inria.fr:1099:Linux, PA_JVM1530790716_am..., Node matrixNode19..., Worker#4, PA_JVM704475267_am..., Node matrixNode18..., Worker#2, PA_JVM1714913173_am..., PA_JVM1122637657_am..., Node matrixNode57..., Worker#1, PA_JVM2010164699_am..., and Node matrixNode15..., Worker#3. Arrows indicate connections between these nodes.

At the bottom, the 'Console' window shows a log of monitoring events:

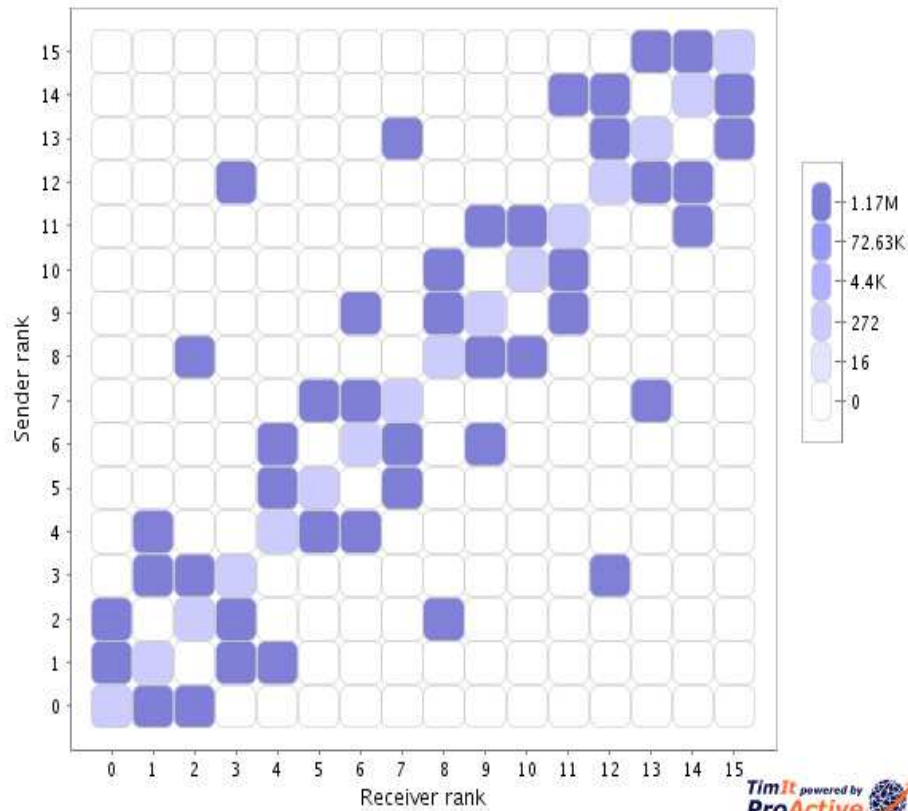
```

Monitoring
15:27:15 => VMobject id=PA_JVM2010164699_ama.inria.fr already monitored, check
15:27:15 => VMobject id=PA_JVM1530790716_ama.inria.fr already monitored, check
15:27:45 => Exploring Host amda.inria.fr with RMI on port 1099
15:27:45 => VMobject id=PA_JVM1714913173_ama.inria.fr already monitored, check
15:27:45 => VMobject id=PA_JVM1122637657_ama.inria.fr already monitored, check
15:27:45 => VMobject id=PA_JVM704475267_ama.inria.fr already monitored, check
15:27:45 => VMobject id=PA_JVM2010164699_ama.inria.fr already monitored, check
15:27:45 => VMobject id=PA_JVM1530790716_ama.inria.fr already monitored, check
  
```

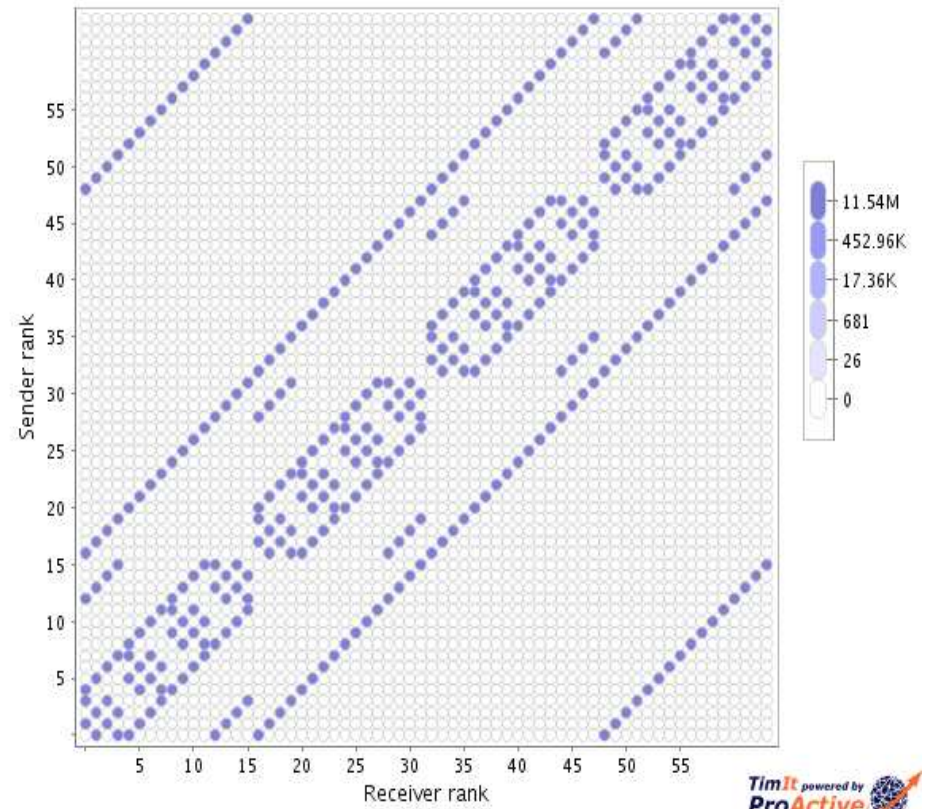
The interface also includes a 'Legend' on the right with categories: Active objects (Active by itself, Serving reqs, Waiting for req, Waiting for res, Migrating, Secure and Ac), Pending Requests (Pending reqs), Nodes (RMI Node, HTTP Node, RMI/SSH Nod), JVMs (Standard JVM, JVM started w), Hosts (Standard Host), and Not Responding (Active Object, JVM). At the bottom, there are controls for 'Auto Reset' (Enable, 7 seconds), 'Drawing style' (Proportional, Ratio, Fixed), and 'Topology' (Display, Reset).



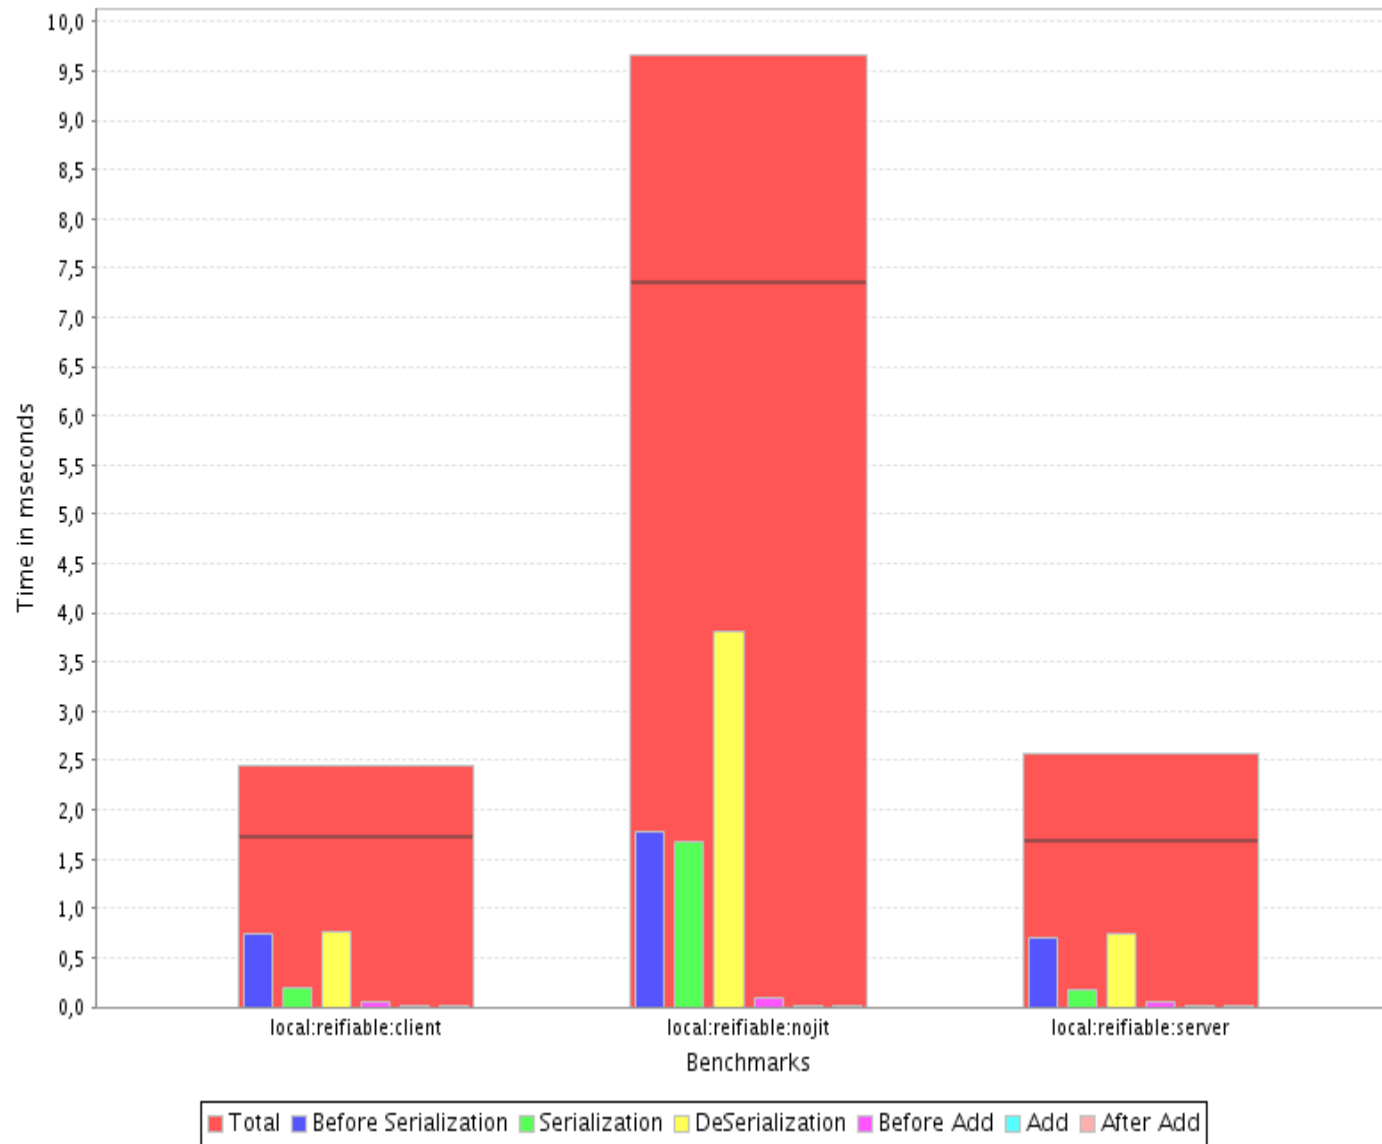
CG Communication pattern S 16
Data density distribution (CG version :1.1)



MG Communication pattern C 64
Data density distribution (MG version :1.0)



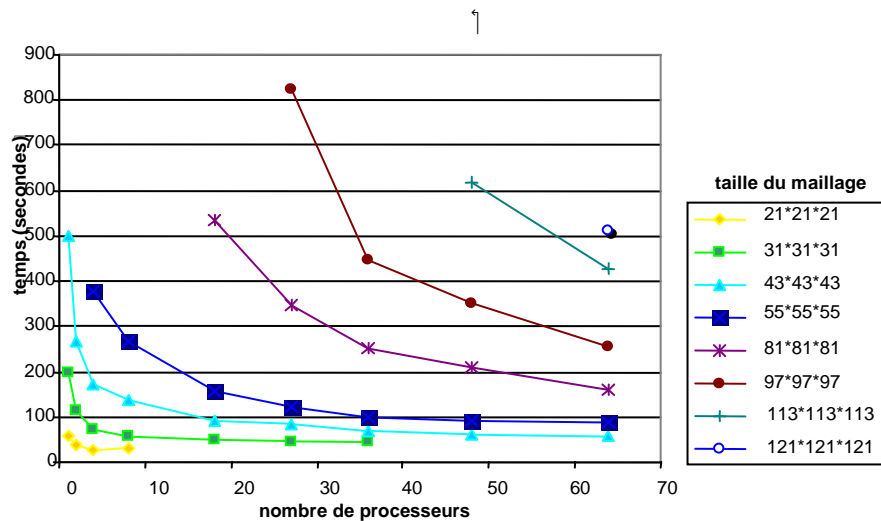
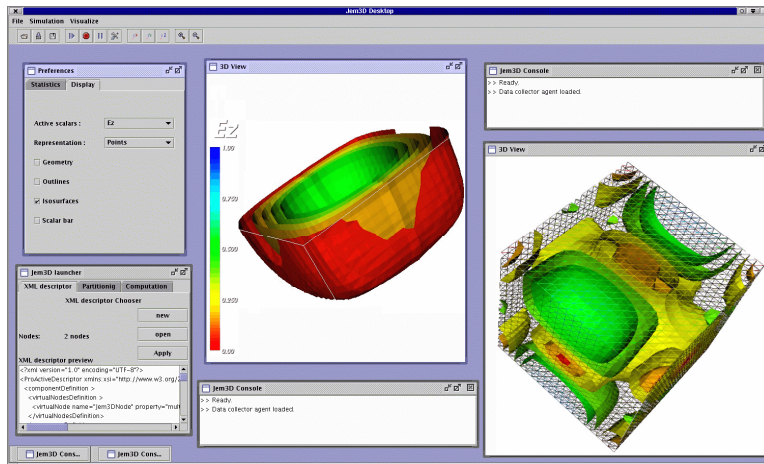
Remote Call – JIT options (1000 iterations, warmup =1000)



Some Applications



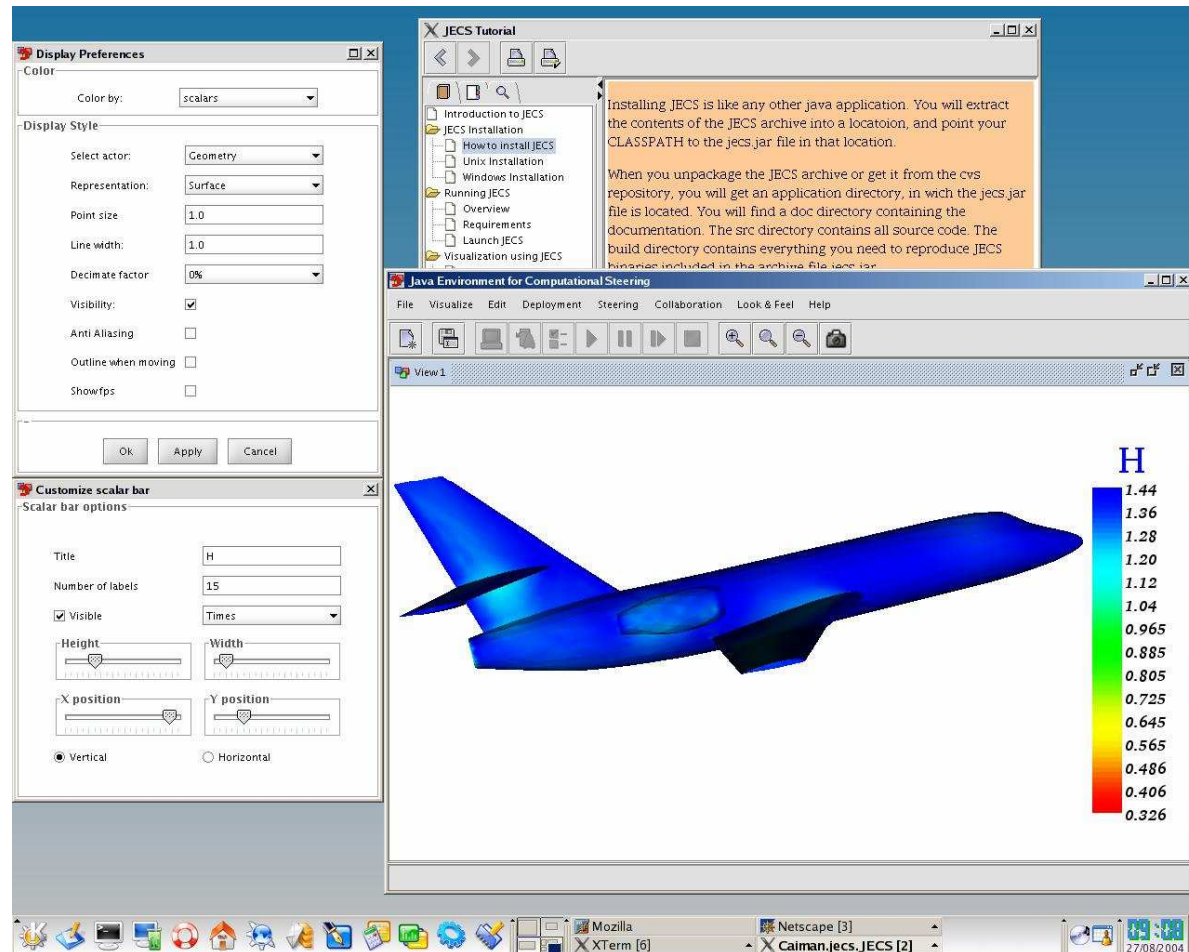
Java 3D Electromagnetism



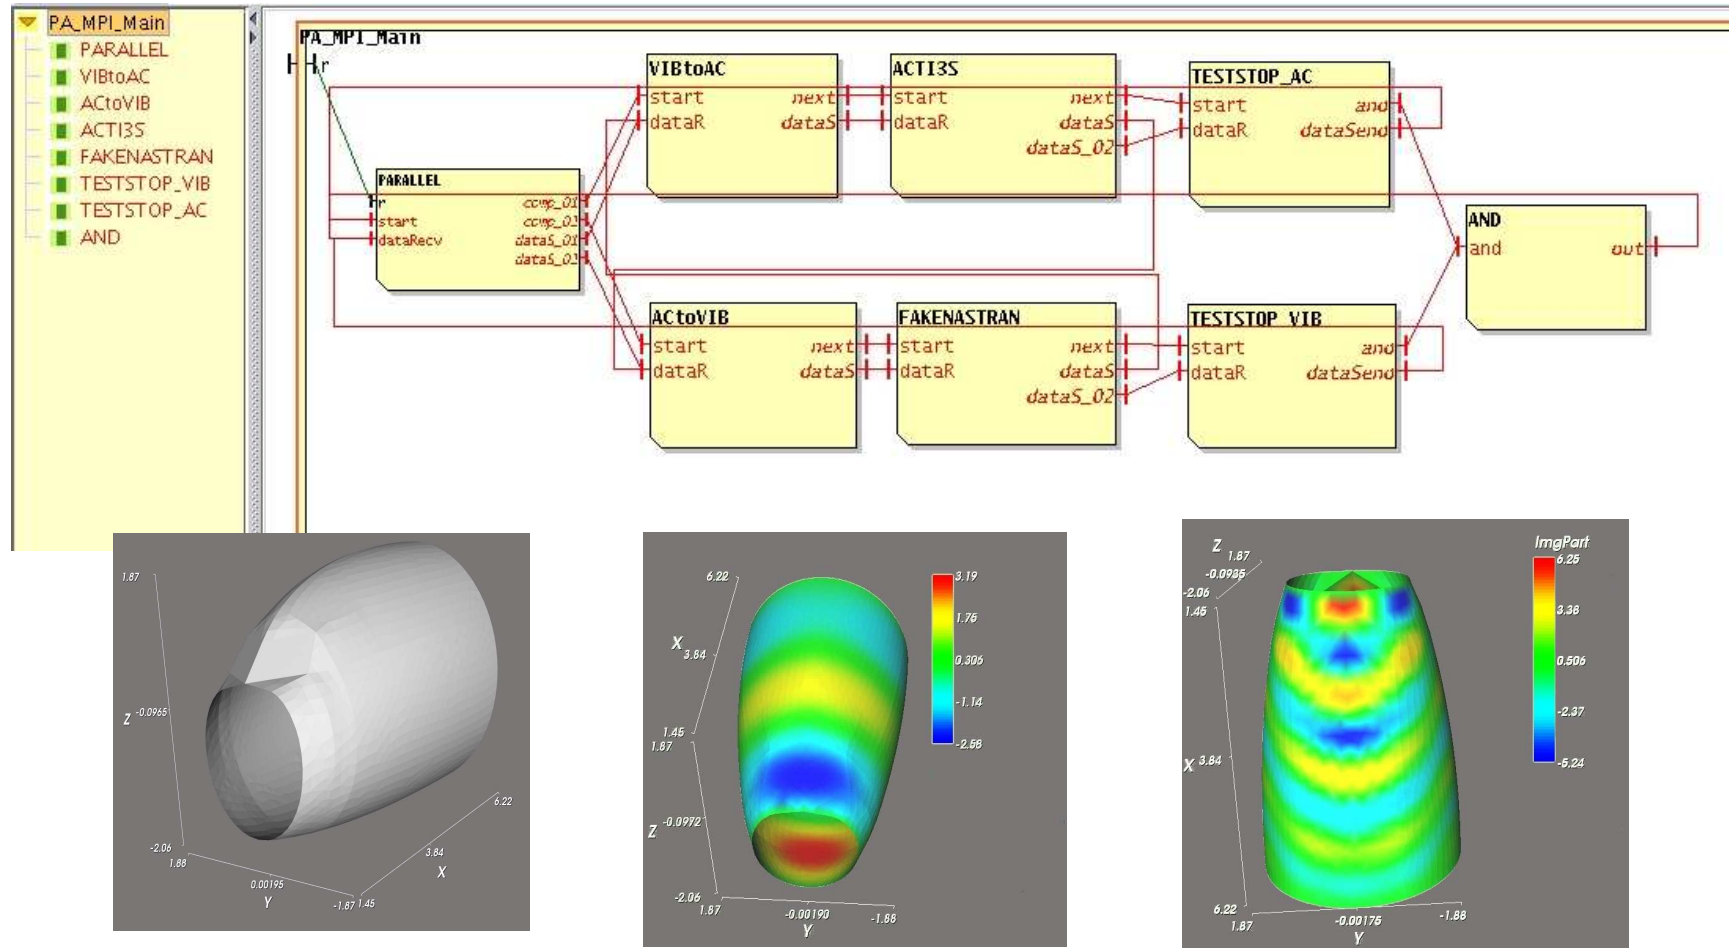
- Maxwell 3D equation solver, Finite Volume Method (FVM)
- Pre-existing **Fortran MPI** version: **EM3D** (CAIMAN team @ INRIA)
- **300+ machines at the same time** (Intranet and cluster)
- Large data sets: 150x150x150 (100 million facets)



J ECS : A Generic Version of Jem3D



Code Coupling : Vibro Acoustic (courtesy of EADS)



Scilab Grid Toolkit

Grid Scilab ToolBox

Command

Scilab Engines

- Engine0
- Engine1
- Engine2
- Engine3
- Engine4
- Engine5

Pending Tasks					
Id Task	Script	Priority	Awaited Time(ms)	State	
Task22	test_scilab.sce	Normal	79000		
Task23	test_scilab.sce	Normal	16000		
Task24	test_scilab.sce	-	-		

Cancel Clear

Executing Tasks					
Id Task	Script	Id Engine	Global Time(ms)	State	
Task14	test_scilab.sce	-	-		
Task15	test_scilab.sce	Engine5	53014		
Task16	test_scilab.sce	-	-		
Task17	test_scilab.sce	Engine0	48063		
Task18	test_scilab.sce	Engine1	45000		

Kill Clear

Terminated Tasks					
Id Task	Script	Execution Time(ms)	Global Time(ms)	State	
Task5	test_scilab.sce	2	1055		
Task7	test_scilab.sce	2	1059		
Task8	test_scilab.sce	40	1042		
Task9	test_scilab.sce	5	1038		
Task10	test_scilab.sce	6	1068		
Task11	test_scilab.sce	1	1068		
Task12	test_scilab.sce	1	1040		

Save Delete

Operations

```

> Deployment is running: /auto/sea/u/sea/0/user/amargin/ProActive_RELEASE/descriptors/examples/ProActiveScilabLocal.xml
-> Deployment is successful: /auto/sea/u/sea/0/user/amargin/ProActive_RELEASE/descriptors/examples/ProActiveScilabLocal.xml
-> Add new Scilab Task :Task0
-> Execute Scilab Task :Task0
-> Terminate Scilab Task :Task0
-> Add new Scilab Task :Task1
-> Execute Scilab Task :Task1
-> Terminate Scilab Task :Task1
    
```

Clear

ProActive Scilab

ToolBox Legend

Pending Tasks

- Pending
- Cancelled

Executing Tasks

- Executing
- Killed

Terminated Tasks

- Succeeded
- Aborted



Post Production movie processing

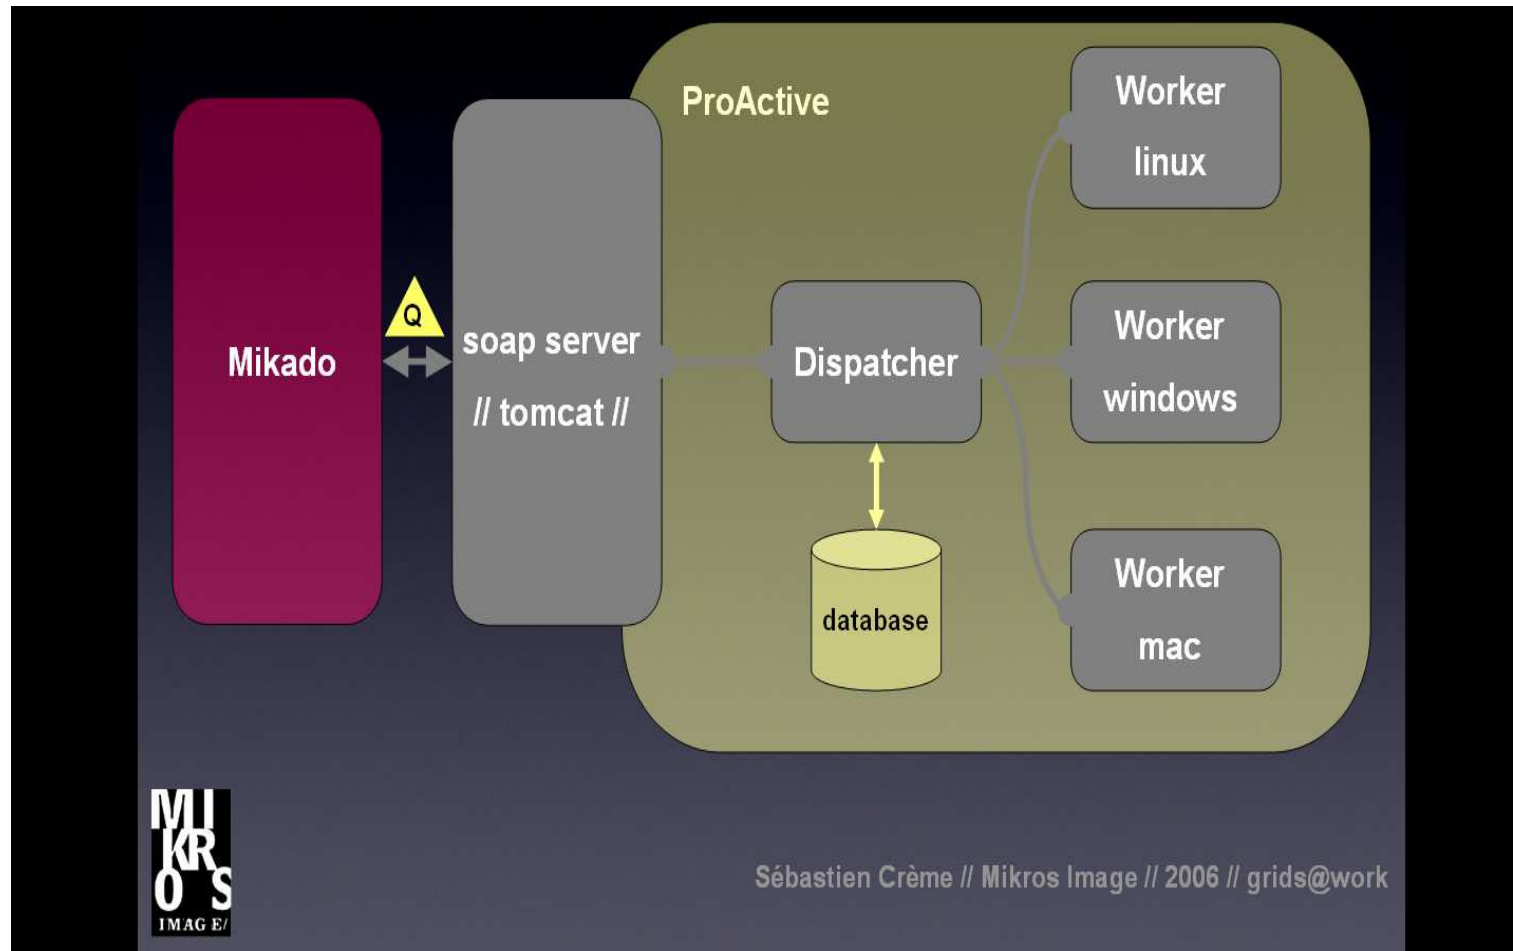
Frames Making of Nissan



Sébastien Crème // Mikros Image // 20



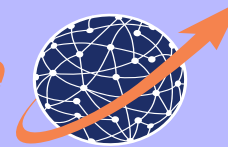
Post Production movie processing



Large Scale Deployments

ObjectWeb
Open Source Middleware

ProActive
Programming, Composing, Deploying on the Grid



PLUGTESTS
THE INTEROPERABILITY SERVICE

Grid Plugtests 2004, 2005 & 2006

20 to 40 sites worldwide

100 GFlops in 2004

1700 Gflops in 2006

4130 cores (2111) in 2006

IBM, Sun, Bull, Apple, x86, 64bits...

Linux, Windows, Solaris, MacOS

ssh, rsh, sshGSI, GRAM

PBS, LSF, SGE, OAR, Globus, Prun

Grid'5000 - DAS etc...

P2P INRIA Infrastructure
53 years computation in 6
months on 200+ machines

On-going activities

- Programming model
 - Grid Component Model, adaptive components
 - Model checking, formal verification of behavioral properties
 - High level parallelism patterns (skeletons)
- Deployment
 - OSGi gateways
 - MPI / native codes wrapping
 - Easier specification, Scheduler
- Middleware services
 - Security at application level
 - Distributed garbage collection
- Industrial strength product
 - Quality development process, Support, Service

**ProActive/GCM User Group and Contest at
GRIDs@Work 2007: IV GRID**

PLUGTESTS, □

Joint European Union/China GRID

28 Oct.-2 Nov. 2007, Beijing, China



Conclusion

- Usability - High Level Abstractions - Latency

Strong programming model

formal model, active objects, groups, components

Versatile deployment framework

Interfaced with Grid & cluster standards

Pluggable middleware services

Mobility, fault tolerance, security etc...





<http://proactive.objectweb.org>

Let's practice !