

Strong Programming Model to bridge Distributed & Multi-Core Computing

D. Caromel, et al.

Agenda

1. Background: OASIS, ActiveEon
2. Multi-Cores
3. Programming,
Optimizing
Scheduling
4. Enterprise Parallel Computing



Speed: Application + Development: Productivity

Key Objectives

- ❑ **Parallel Programming Model and Tools**
 - desperately needed
 - for the masses
 - for new architectures (Multi-cores)
- ❑ **As Effective as possible:**
 - Efficient
 - However Programmer Productivity is first KSF
- ❑ **For both Multi-cores and Distributed**
 - Actually the way around
- ❑ **Some Handling of “Large-scale” (Grid, Clouds)**



1. Background

- A joint team, Now about 35 persons
- 2004: First ProActive User Group
- 2009, April: ProActive 4.1, Distributed & Parallel:
From Multi-cores to Enterprise GRIDs



OASIS Team Composition (35)

□ Researchers (5):

- D. Caromel (UNSA, Det. INRIA)
- E. Madelaine (INRIA)
- F. Baude (UNSA)
- F. Huet (UNSA)
- L. Henrio (CNRS)

□ PhDs (11):

- Antonio Cansado (INRIA, Conic)
- Brian Amedro (SCS-Agos)
- Cristian Ruz (INRIA, Conicyt)
- Elton Mathias (INRIA-Cordi)
- Imen Filali (SCS-Agos / FP7 SO)
- Marcela Rivera (INRIA, Conicyt)
- Muhammad Khan (STIC-Asia)
- Paul Naoumenko (INRIA/Région)
- Viet Dung Doan (FP6 Bionets)
- Virginie Contes (SOA4ALL)
- Guilherme Pezzi (AGOS, CIFR)

□ + Visitors + Interns



Located in Sophia Antipolis, between
Nice and Cannes,
Visitors and Students Welcome!

Startup Company Born of INRIA

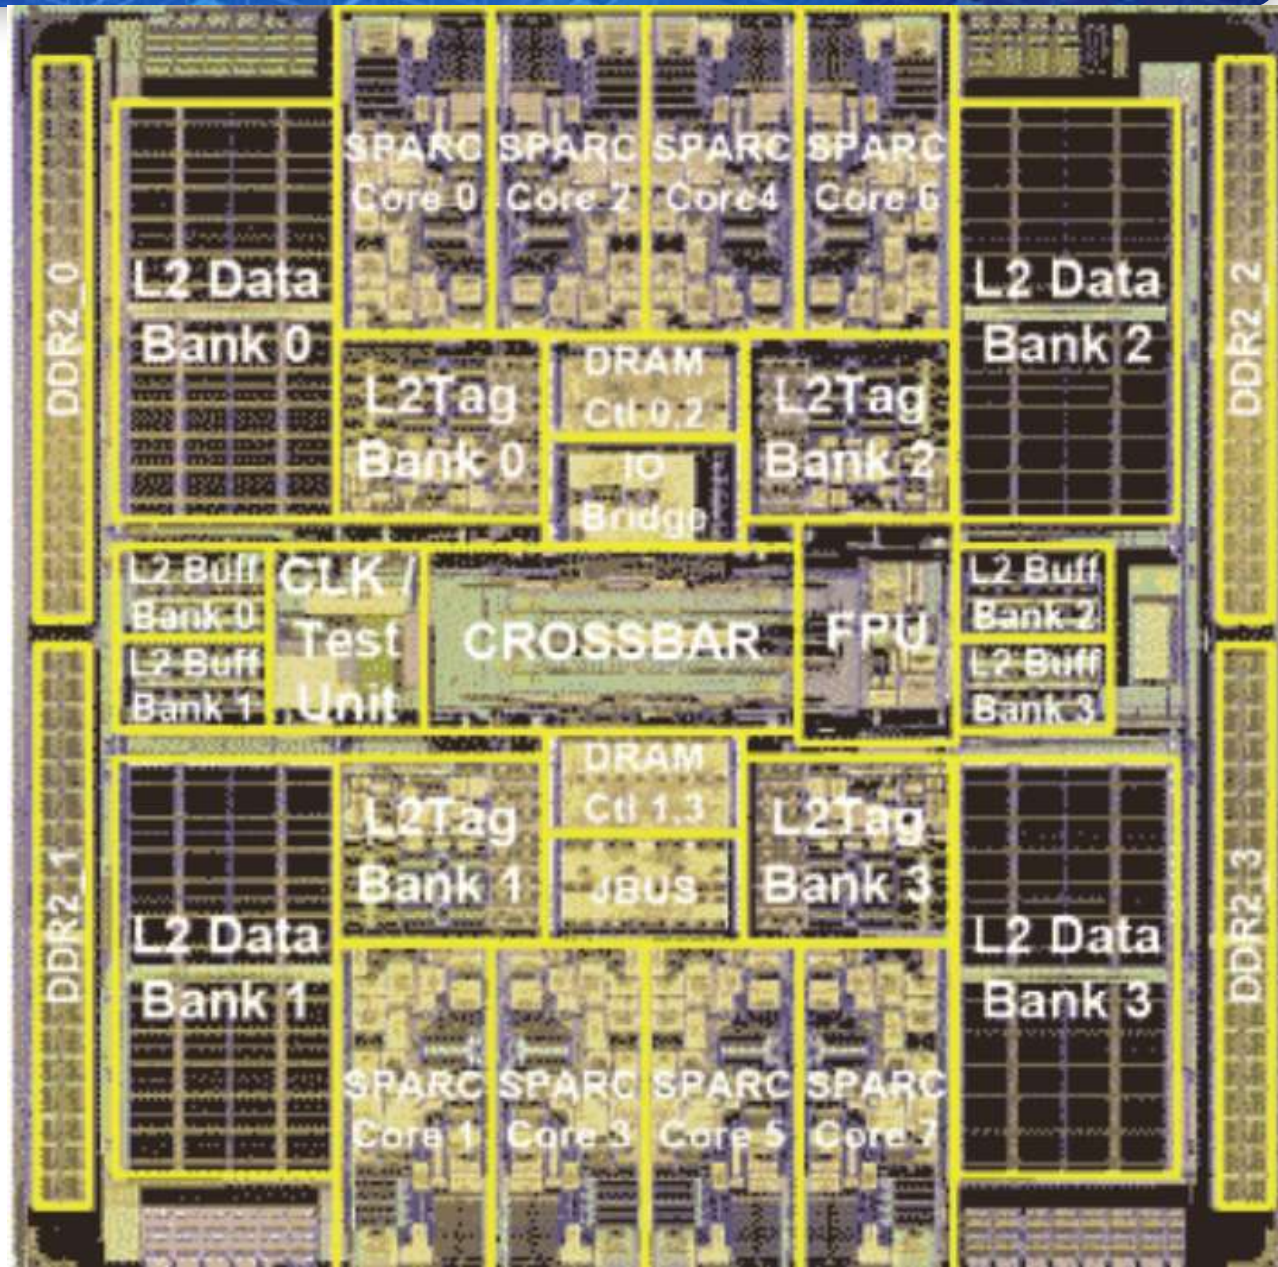


- ❑ Co-developing, Support for [ProActive Parallel Suite](#)
- ❑ Worldwide Customers: Fr, UK, Boston USA

□ Multi-Cores

Symmetrical Multi-Core: 8-ways Niagara II

- ❑ 8 cores
- ❑ 4 Native threads per core
- ❑ Linux see 32 cores!

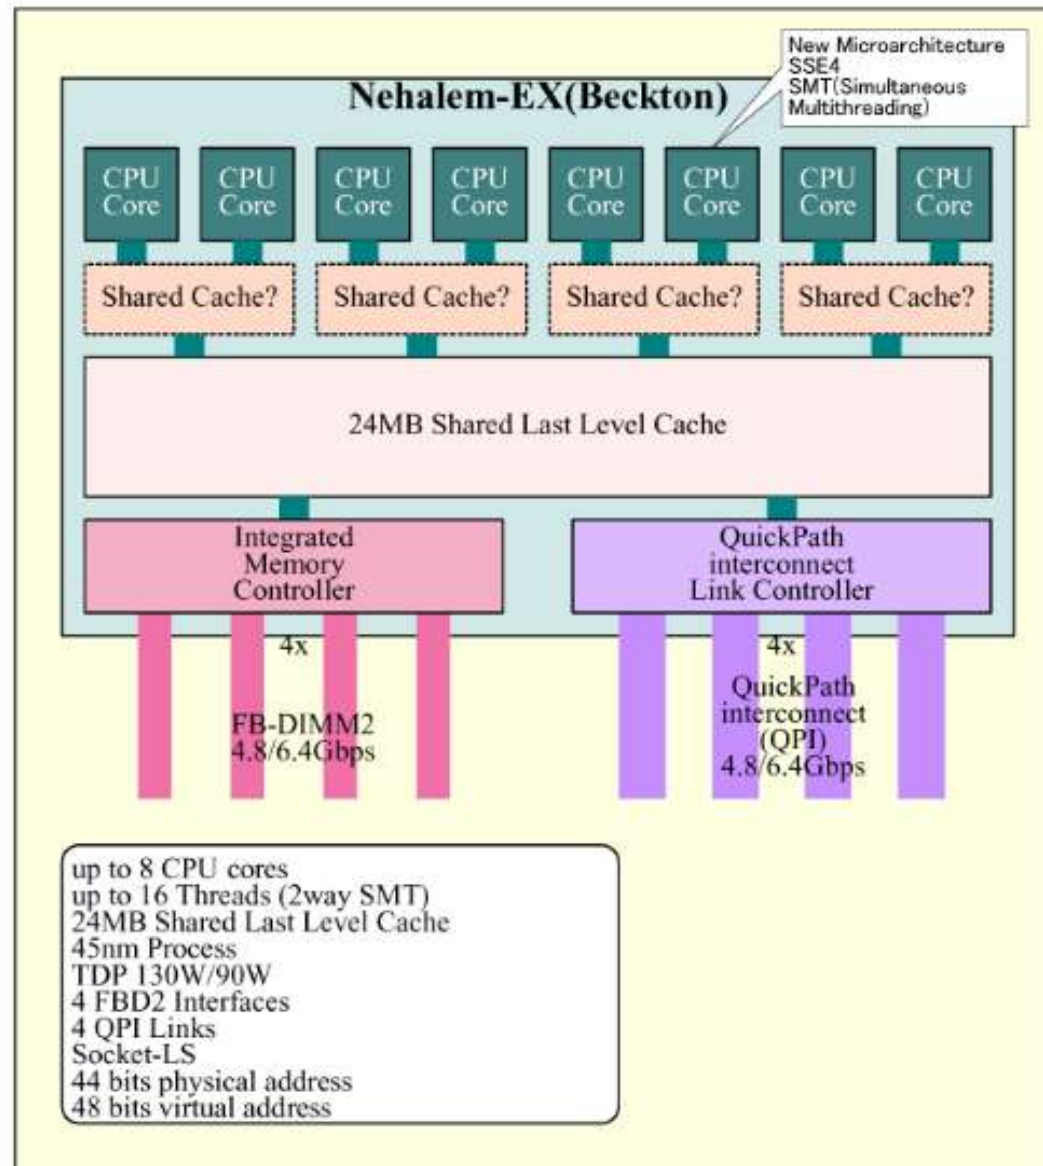


Sun 16-core Rock: Fall 2009

- ❑ 16 cores
- ❑ 4 native threads per core
- ❑
- ❑ → 64 “Cores” or “Native Threads” at OS level

Intel 8-cores, 16-thread Nehalem-based Xeon processor confirmed (Feb. 2009)

- ❑ Highly
- ❑ NUMA
- ❑ Not an SMP:
- ❑ L1,
- ❑ L2, then
- ❑ L3 attached
- ❑ to a given
- ❑ core



Multi-Cores

A Few Key Points

- ❑ Moore's Law rephrased:
 - ❑ Nb. of Cores double every 18 to 24 months
- ❑ Key expected Milestones: Cores per Chips (OTS)
 - 2010: 32 to 64
 - 2012: 64 to 128
 - 2014: 128 to 256
- ❑ 1 Million Cores Parallel Machines in 2012
- ❑ 100 M cores coming in 2020
- ❑ Multi-Cores are NUMA, and turning Heterogeneous (GPU)
- ❑ They are turning into SoC with NoC: NOT SMP!



Parallel Acceleration Toolkit in Java:

Parallelism:

Multi-Core+Distributed



Open Source Used in production by industry

OW2: Object Web + Orient Ware



Leading Open Source Middleware



ProActive Parallel Suite



PROGRAMMING

Java Parallel Frameworks

for HPC, Multi-Cores, Distribution, Enterprise Grids and Clouds.

Featuring: Async. comms, Master-Worker, Monte-Carlo, SPMD, components and legacy code wrapping.

OPTIMIZING

Eclipse GUI (IC2D)

for Developing, Debugging, Optimizing your parallel applications.

Featuring: graphical monitoring and benchmarking with report generation.

SCHEDULING

Multi-Language Scheduler

for Workflows made of C, C++, Java, Scripts, Matlab, Scilab tasks.

Featuring: graphical user interface, resource acquisition and virtualization.

ProActive Contributors

Abhijeet Gaikwad
(Option Pricing)
Abhishek-Rajeev Gupta
Antonio Cansado
Baptiste De Stefano
Bastien Sauvan
Brian Amedro (SPMD)
Cédric Dalmaso (Component)
Clement Mathieu (Core,
GCM Deployment)
Elaine Isnard
Elton Mathias
Eric Madelaine
Etienne Vallette-De-Osia
Fabien Viale (Matlab, Scilab)
Fabrice Huet (Mobility, P2P)
Florin Bratu
Franca Perrina
Francoise Baude
Germain Sigety (Scheduling)
Guillaume Laurent
Guilherme Perretti Pezzi
Imen Filiali
Jean-Luc Scheefer (Scheduling)

Jean-Michael Guillamume
Johann Fradj (Scheduling)
Jonathan Martin
Julian Krzeminski
Kamran Qadir
Khan Muhammad
Laurent Vanni
Ludovic Henrio
Marcela Rivera
Mario Leyton (Skeleton)
Maxime Menant
Nicolas Dodelin
Olivier Helin
Paul Naoumenko
Regis Gascon
Tomasz Dobek
Vasile Jureschi (Technical Writer)
Viet Dong Doan
Vincent Cave (Legacy Wrapping)
Virginie Contes (OSGi, WS)
Yu Feng
Yulai Yuan
Zihui Dai

Arnaud Contes (Core, Security, Debugging)
Christian Delbé (FT, Scheduling)
Emil Salageanu (IC2D, Agent)
Vladimir Bodnartchouk (IC2D, TimIt)

Alexandre di
Costanzo (P2P, B&B)
Boutheina Bennour
Guillaume Chazarain (DGC)
Julien Vayssiere
(MOP, Active Objects)

Lionel Mestre
Laurent Baduel (Group
Communications)
Matthieu Morel (Initial
Component Work)
Nadia Rinaldo
(Core, Deployment)
Romain Quilici

2. Distributed and Parallel Active Objects

ProActive Parallel Suite



PROGRAMMING

Java Parallel Frameworks

for HPC, Multi-Cores, Distribution, Enterprise Grids and Clouds.

Featuring: Async. comms, Master-Worker, Monte-Carlo, SPMD, components and legacy code wrapping.

OPTIMIZING

Eclipse GUI (IC2D)

for Developing, Debugging, Optimizing your parallel applications.

Featuring: graphical monitoring and benchmarking with report generation.

SCHEDULING

Multi-Language Scheduler

for Workflows made of C, C++, Java, Scripts, Matlab, Scilab tasks.

Featuring: graphical user interface, resource acquisition and virtualization.

ProActive Parallel Suite



PROGRAMMING

Java Parallel Frameworks

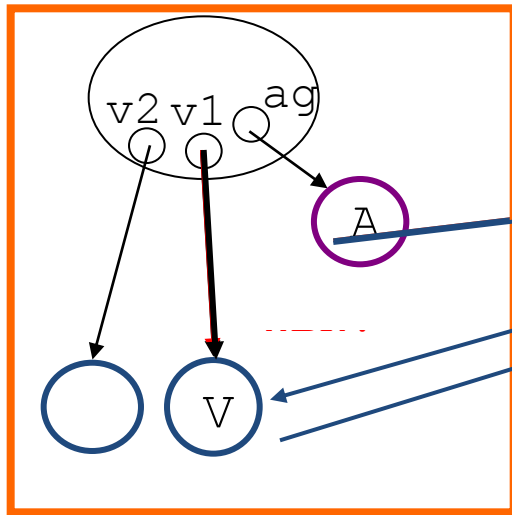
for HPC, Multi-Cores, Distribution, Enterprise Grids and Clouds.

Featuring: Async. comms, Master-Worker, Monte-Carlo, SPMD, components and legacy code wrapping.

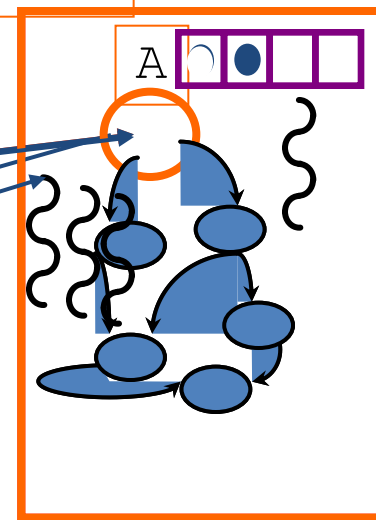
ProActive : Active objects

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

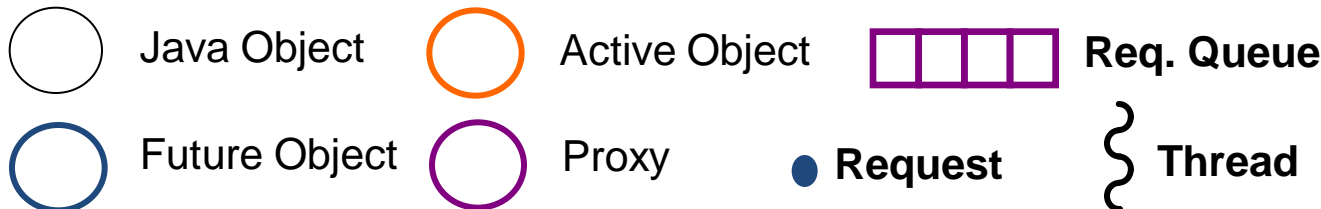
JVM



JVM



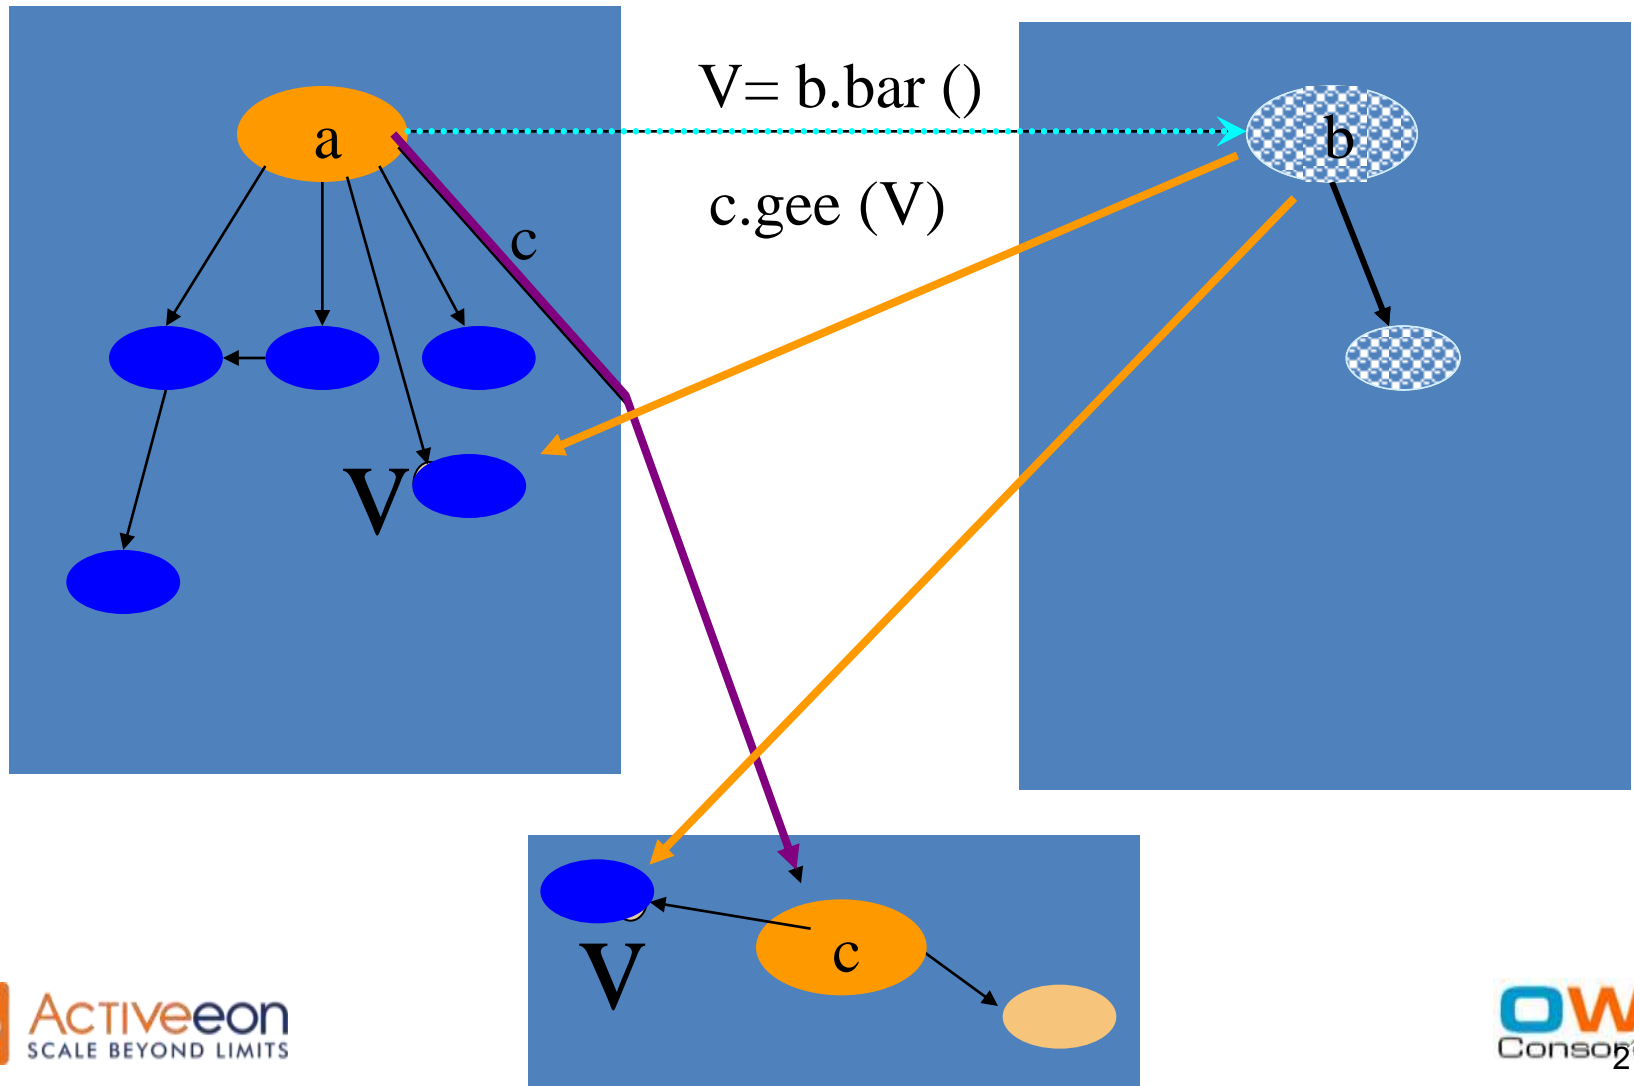
Wait-By-Necessity
is a
Dataflow
Synchronization



First-Class Futures Update

Wait-By-Necessity: First Class Futures

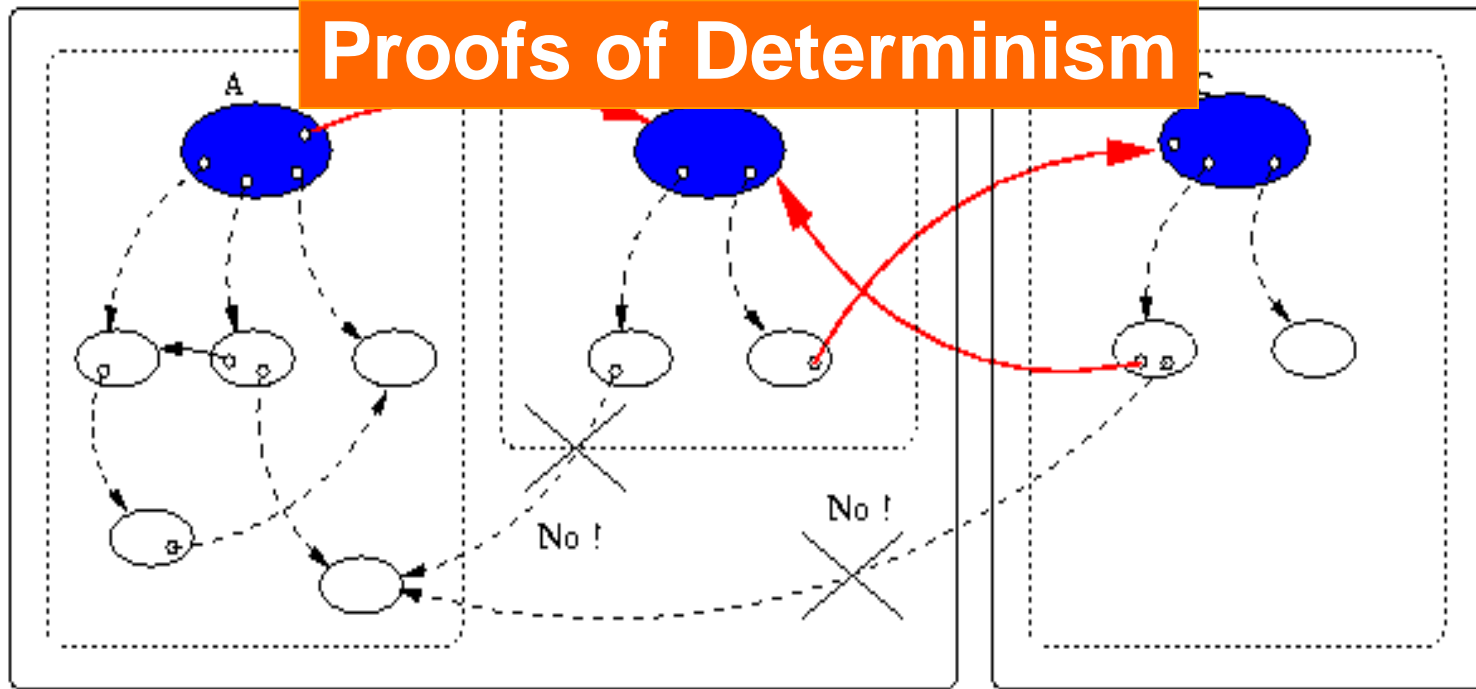
Futures are Global Single-Assignment Variables



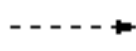
Standard system at Runtime: No Sharing

NoC: Network On Chip

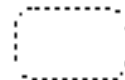
Proofs of Determinism



Active Object



Synchronous Call



Sub System



Passive Object



Asynchronous Call



Address Space

Key Point: Software Evolution

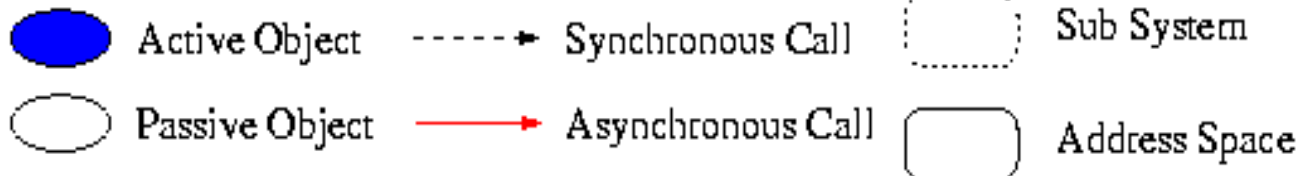
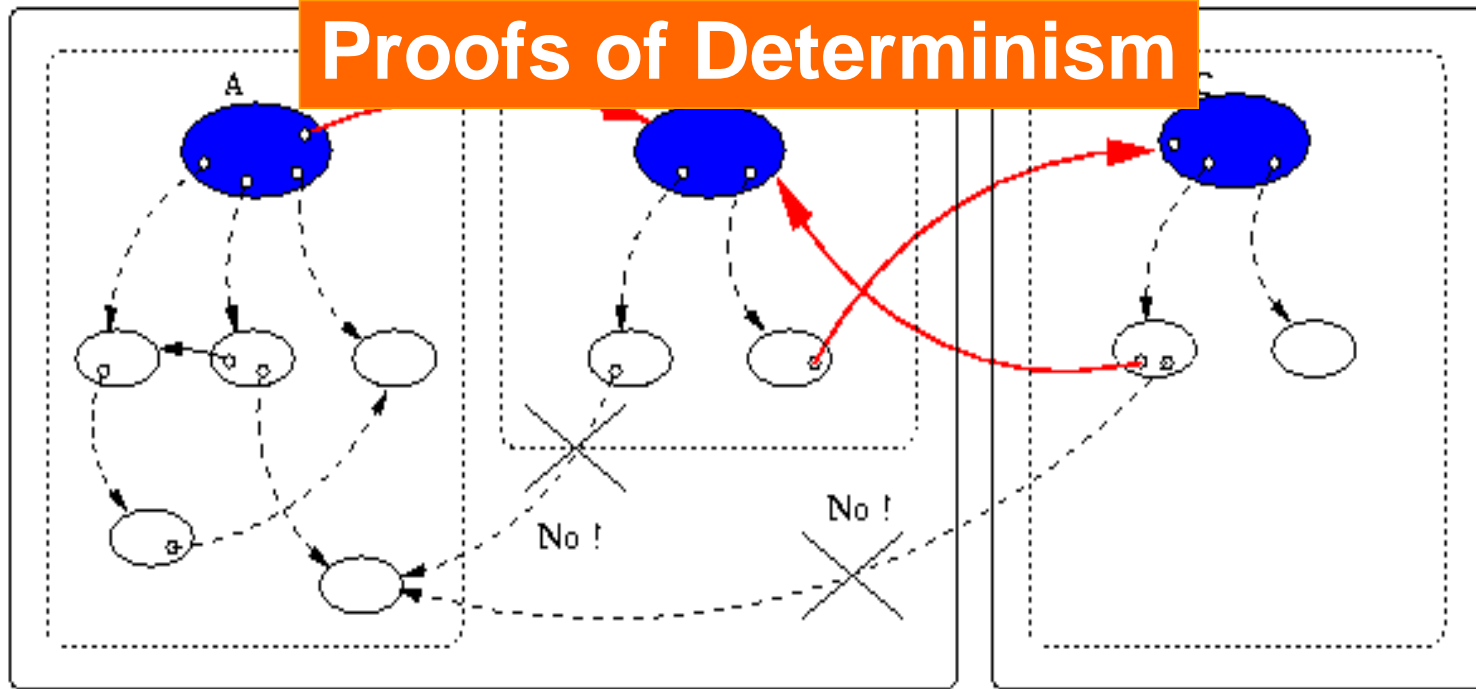
- ❑ **Distributed To Multicores**
- ❑ **Multi-Cores: 32 (2010) to 64 to 128 to 256 (2014)**
- ❑ **Shift the execution from several multi-cores executing**
- ❑ **the same application simultaneously to a single, larger**
- ❑ **multi-core chip. An application requiring 128 cores to**
- ❑ **correctly execute, can be executed in 2012 on four 32**
- ❑ **cores,**
- ❑ **and seamlessly executed in 2016 on a single 128-core**
- ❑ **chips**

- ❑ **→ Smooth evolutivity of applications**
- ❑ **Distributed and Multi-core Platforms**
- ❑ **→ Also for Data Scalability**

Standard system at Runtime: No Sharing

NoC: Network On Chip

Proofs of Determinism



(2) ASP: Asynchronous Sequential Processes

$$\frac{(a, \sigma) \rightarrow_S (a', \sigma')}{\alpha[a; \sigma; \iota; F; R; f] \parallel P \longrightarrow \alpha[a'; \sigma'; \iota; F; R; f] \parallel P} \text{ (LOCAL)}$$

Local

$$\frac{\begin{array}{l} \gamma \text{ fresh activity} \quad \iota' \notin \text{dom}(\sigma) \quad \sigma' = \{\iota' \mapsto AO(\gamma)\} :: \sigma \\ \sigma_\gamma = \text{copy}(\iota'', \sigma) \quad \text{Service} = (\text{if } m_j = \emptyset \text{ then } \text{FifoService} \text{ else } \iota''.m_j()) \end{array}}{\alpha[\mathcal{R}[\text{Active}(\iota'', m_j)]; \sigma; \iota; F; R; f] \parallel P \longrightarrow \alpha[\mathcal{R}[\iota']; \sigma'; \iota; F; R; f] \parallel \gamma[\text{Service}; \sigma_\gamma; \iota''; \emptyset; \emptyset; \emptyset] \parallel P} \text{ (NEWACT)}$$

Creating an Activity

$$\frac{\begin{array}{l} \sigma_\alpha(\iota) = AO(\beta) \quad \iota'' \notin \text{dom}(\sigma_\beta) \quad f_i^{\alpha \rightarrow \beta} \text{ new future} \quad \iota_f \notin \text{dom}(\sigma_\alpha) \\ \sigma'_\beta = \text{Copy\&Merge}(\sigma_\alpha, \iota' ; \sigma_\beta, \iota'') \quad \sigma'_\alpha = \{\iota_f \mapsto \text{fut}(f_i^{\alpha \rightarrow \beta})\} :: \sigma_\alpha \end{array}}{\alpha[\mathcal{R}[\iota.m_j(\iota')]; \sigma_\alpha; \iota_\alpha; F_\alpha; R_\alpha; f_\alpha] \parallel \beta[a_\beta; \sigma_\beta; \iota_\beta; F_\beta; R_\beta; f_\beta] \parallel P \longrightarrow \alpha[\mathcal{R}[\iota_f]; \sigma'_\alpha; \iota_\alpha; F_\alpha; R_\alpha; f_\alpha] \parallel \beta[a_\beta; \sigma'_\beta; \iota_\beta; F_\beta; R_\beta :: [m_j; \iota''; f_i^{\alpha \rightarrow \beta}]; f_\beta] \parallel P} \text{ (REQUEST)}$$

Sending a Request

$$\frac{R = R' :: [m_j; \iota_r; f'] :: R'' \quad m_j \in M \quad \forall m \in M, m \notin R'}{\alpha[\mathcal{R}[\text{Serve}(M)]; \sigma; \iota; F; R; f] \parallel P \longrightarrow \alpha[\iota.m_j(\iota_r) \uparrow f, \mathcal{R}[\Box]; \sigma; \iota; F; R' :: R''; f'] \parallel P} \text{ (SERVE)}$$

Service

$$\frac{\iota' \notin \text{dom}(\sigma) \quad F' = F :: \{f \mapsto \iota'\} \quad \sigma' = \text{Copy\&Merge}(\sigma, \iota ; \sigma, \iota')}{\alpha[\iota \uparrow (f', a); \sigma; \iota; F; R; f] \parallel P \longrightarrow \alpha[a; \sigma'; \iota; F'; R; f'] \parallel P} \text{ (ENDSERVICE)}$$

$$\frac{\sigma_\alpha(\iota) = \text{fut}(f_i^{\gamma \rightarrow \beta}) \quad F_\beta(f_i^{\gamma \rightarrow \beta}) = \iota_f \quad \sigma'_\alpha = \text{Copy\&Merge}(\sigma_\beta, \iota_f ; \sigma_\alpha, \iota)}{\alpha[a_\alpha; \sigma_\alpha; \iota_\alpha; F_\alpha; R_\alpha; f_\alpha] \parallel \beta[a_\beta; \sigma_\beta; \iota_\beta; F_\beta; R_\beta; f_\beta] \parallel P \longrightarrow \alpha[a_\alpha; \sigma'_\alpha; \iota_\alpha; F_\alpha; R_\alpha; f_\alpha] \parallel \beta[a_\beta; \sigma_\beta; \iota_\beta; F_\beta; R_\beta; f_\beta] \parallel P} \text{ (REPLY)}$$

Sending a Reply

TYPED ASYNCHRONOUS GROUPS

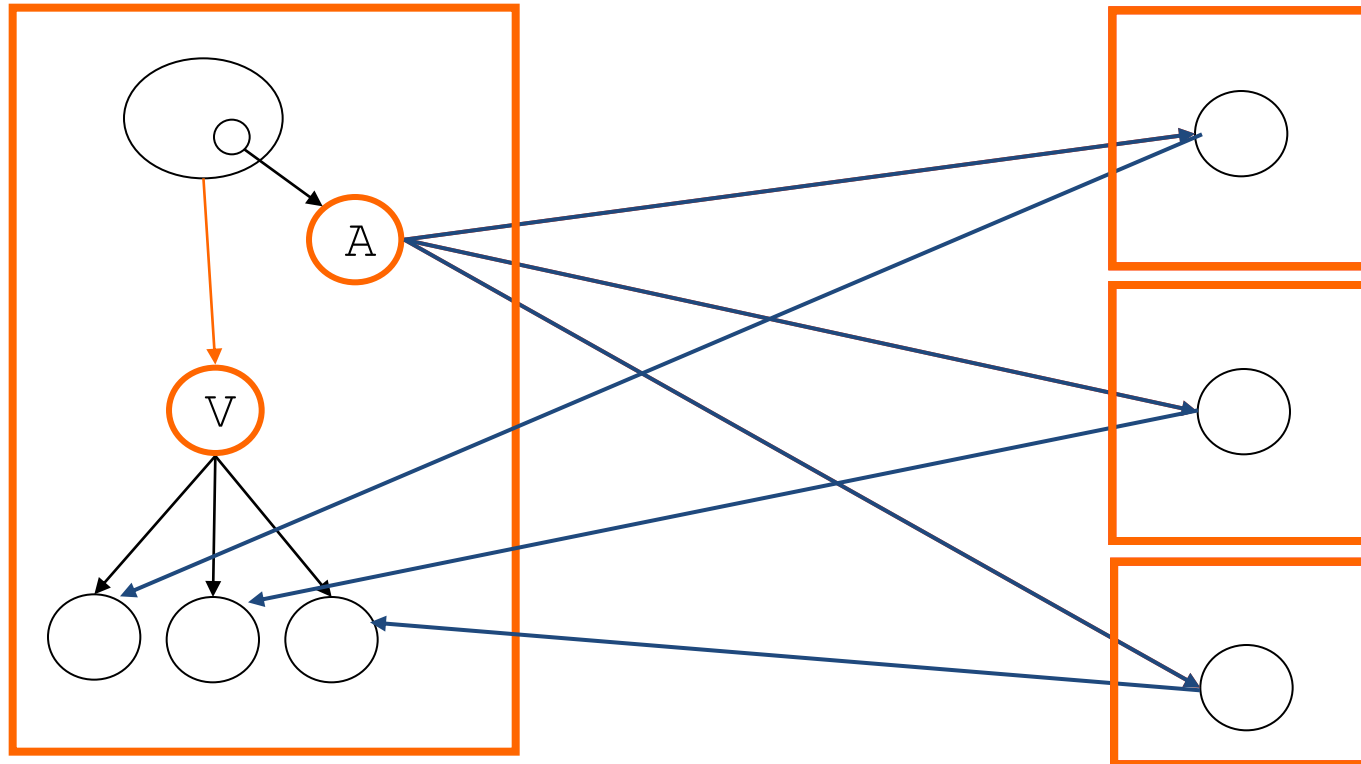
Creating AO and Groups

● A ag = `newActiveGroup` ("A", [...], VirtualNode)

● V v = ag.foo(param);

● ● ● ●

JVM ● v.bar(); //Wait-by-necessity



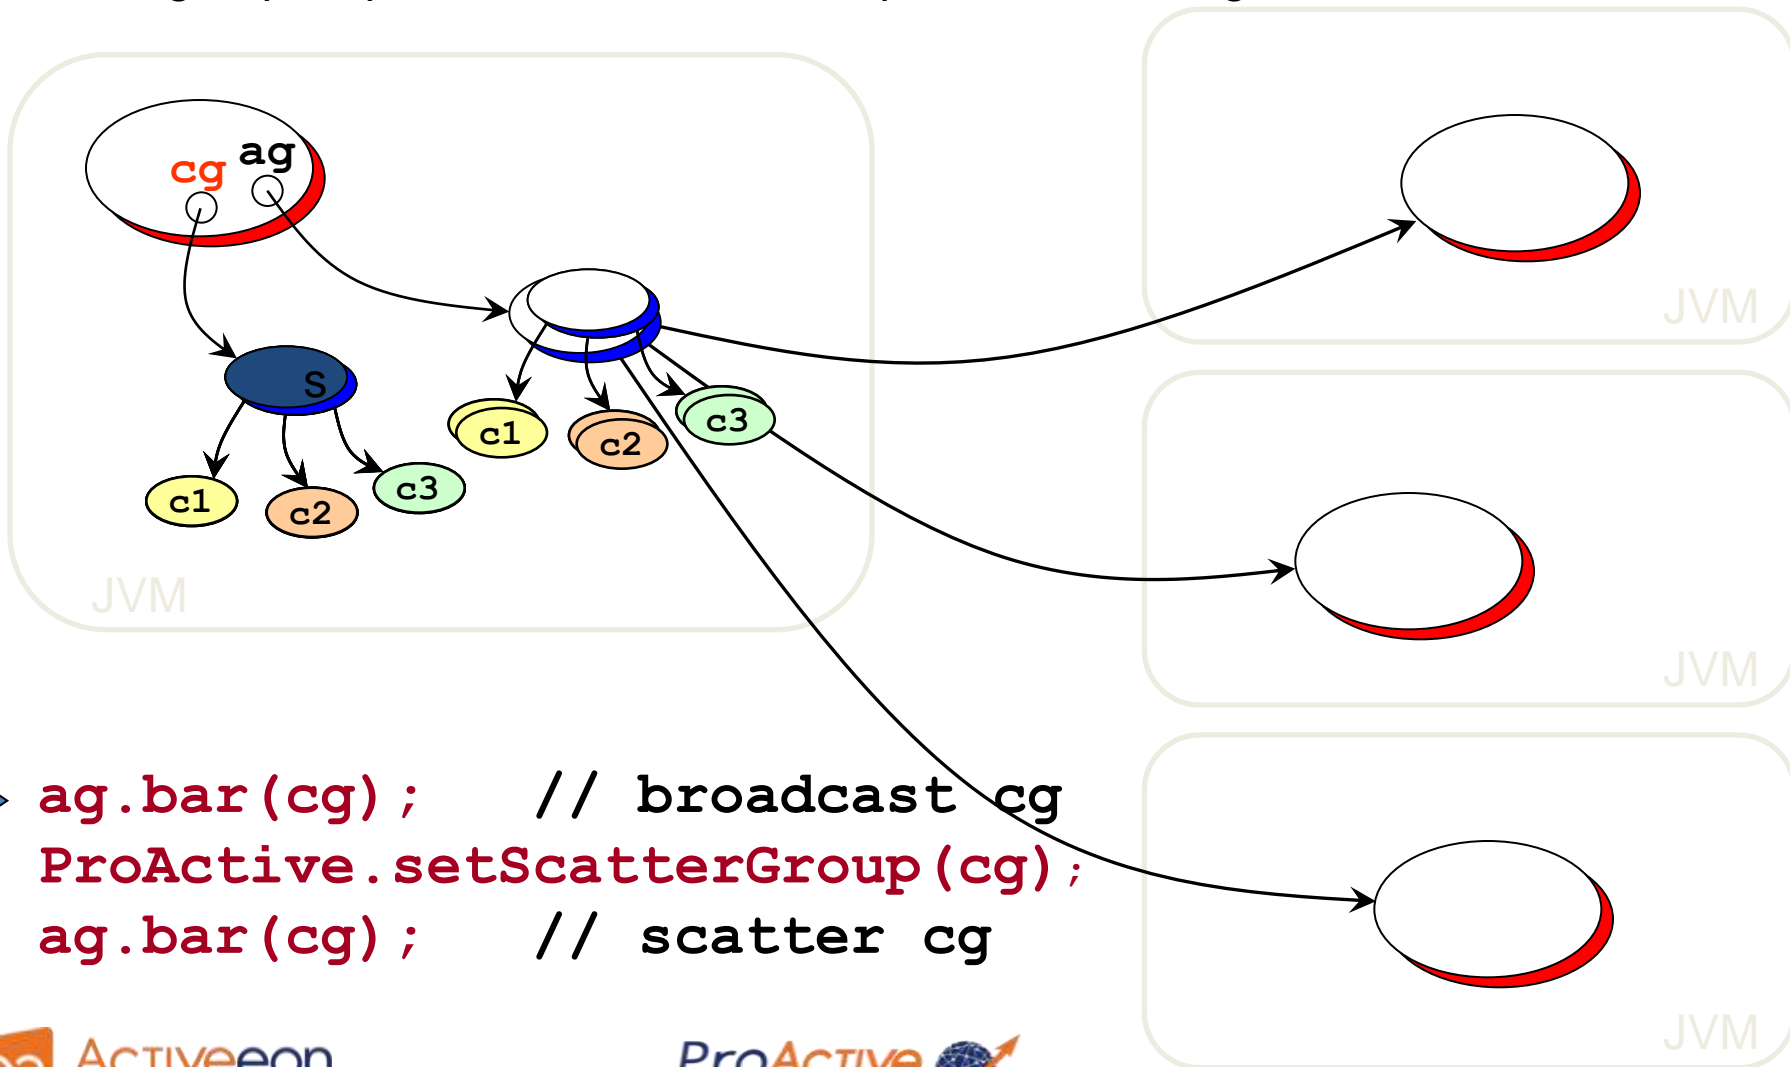
○ Typed Group ○ Java or Active Object

Group, Type, and Asynchrony
are crucial for Composition

Broadcast and Scatter

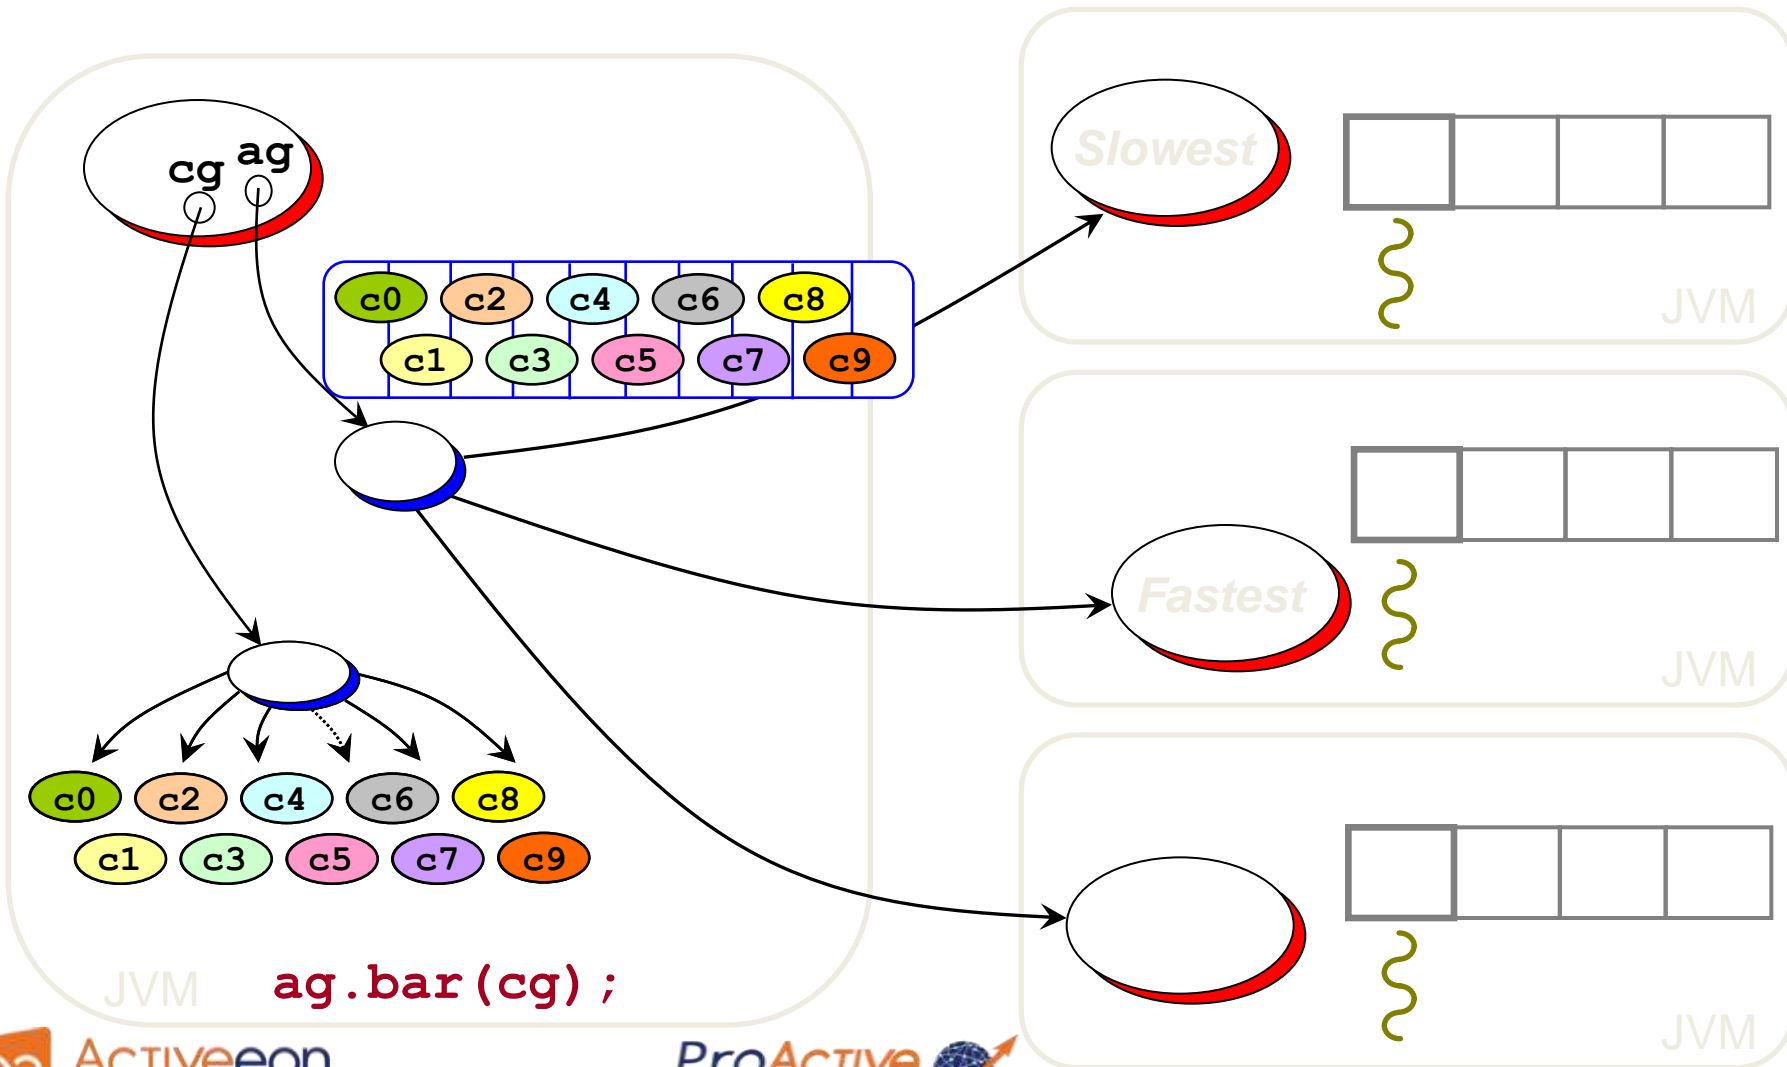
Broadcast is the default behavior

Use a group as parameter, Scattered depends on rankings



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

Dynamic Dispatch Group



JVM `ag.bar (cg) ;`

Abstractions for Parallelism

The right Tool to do the Task right

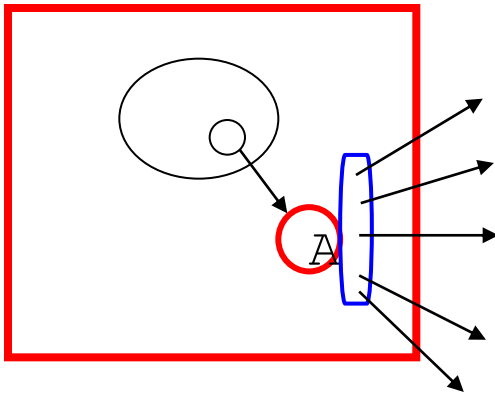
Object-Oriented SPMD

Key Point

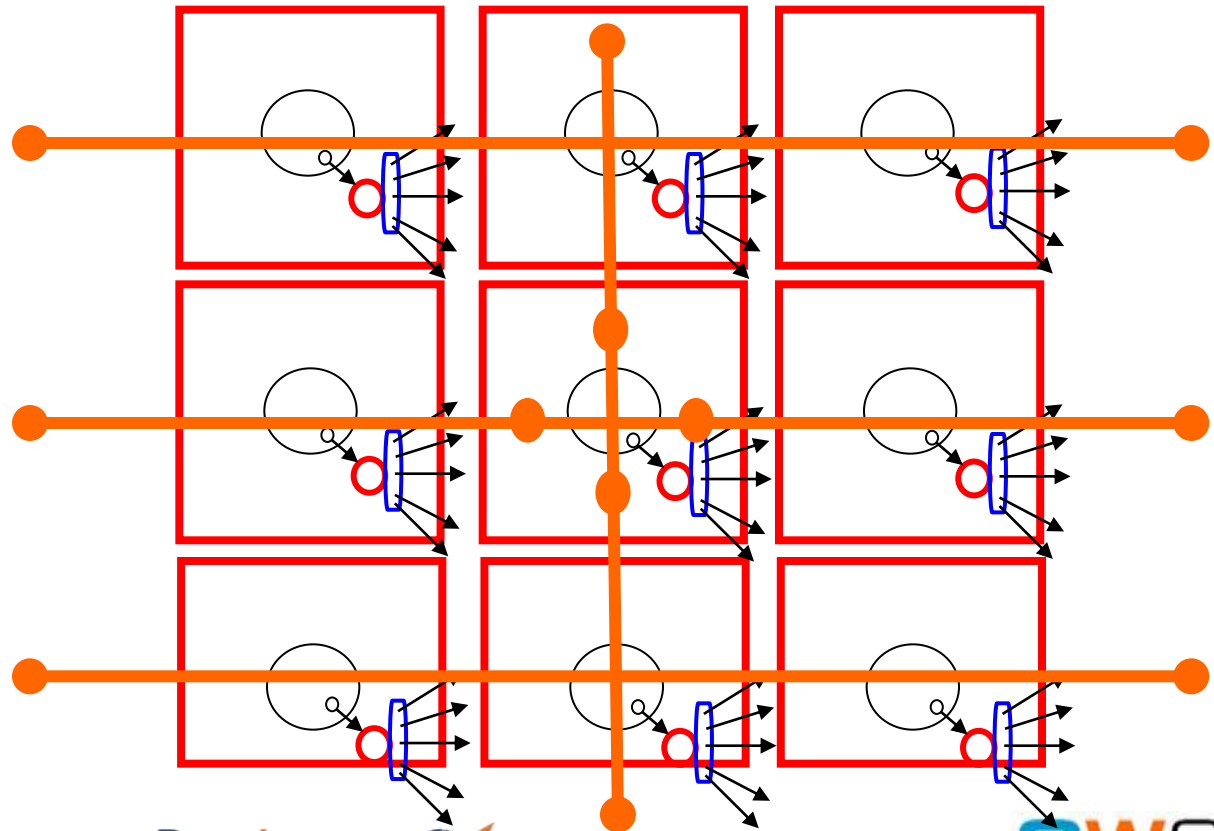
- ❑ **“MPI and programming languages from the 60’s will not make it”**
- ❑ **Jack Dongarra, 2/13/2009,**
- ❑ **Wake Forest University talk**

OO SPMD: Object-Oriented SPMD

```
A ag = newSPMDGroup ("A", [...], VirtualNode)
// In each member
myGroup.barrier ("2D"); // Global Barrier
myGroup.barrier ("vertical"); // Any Barrier
myGroup.barrier ("north", "south", "east", "west");
```



Still,
not based on raw
messages, but
Typed Method Calls
==> Components



OO SPMD: Object-Oriented SPMD

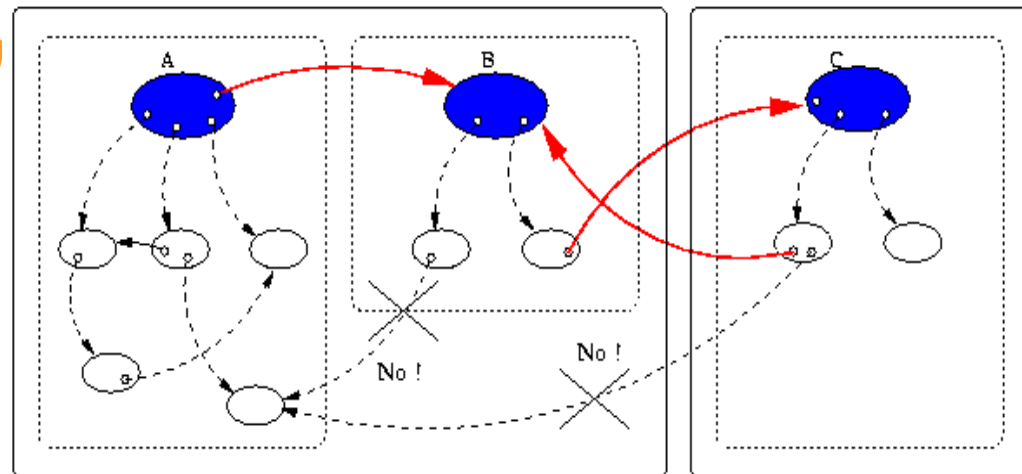
- ❑ Motivation
- ❑ Use Enterprise technology (Java, Eclipse) for Numerical Parallel Computing
- ❑ Able to express in Java MPI's Collective Communications:
 - **broadcast**
 - **scatter**
 - **gather**
 - **reduce**
 - **allscatter**
 - **allgather**
- ❑ Together with
- ❑ **Barriers, Topologies.**

Application Semantics rather than Low-Level Architecture-Based Optimization

- ❑ MPI: MPI_Send MPI_Recv MPI_Ssend MPI_Irecv
- ❑ MPI_Bsend MPI_Rsend MPI_Isend MPI_Ibsend
- ❑ What we propose:
 - High-level Information from Application Programmer
 - Tower Self-Adapting parallel applications

- ❑ Examples:
 - ro.foo (*ForgetOnSend* (params));
 - *ActiveObject.exchan*

- ➔ Optimizations for Both
- ➔ Distributed &
- ➔ Multi-Core



Key Point:

Infrastructure Independence

Application Abstractions

- I give you this data and I no longer need it

Not Infrastructure Abstractions

- I asynchronously send you this and I do not lock the buffer

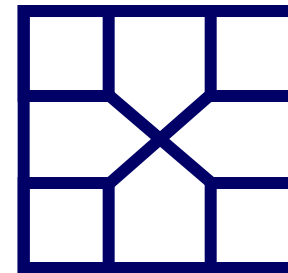
NAS Parallel Benchmarks

- Experimented on 3D ElectroMagnetism, and Nasa Benchmarks
- **Designed by NASA to evaluate benefits of high performance systems**
- **Strongly based on CFD**
- **5 benchmarks (kernels) to test different aspects of a system**
- **2 categories or focus variations:**
 - **communication intensive and computation intensive**

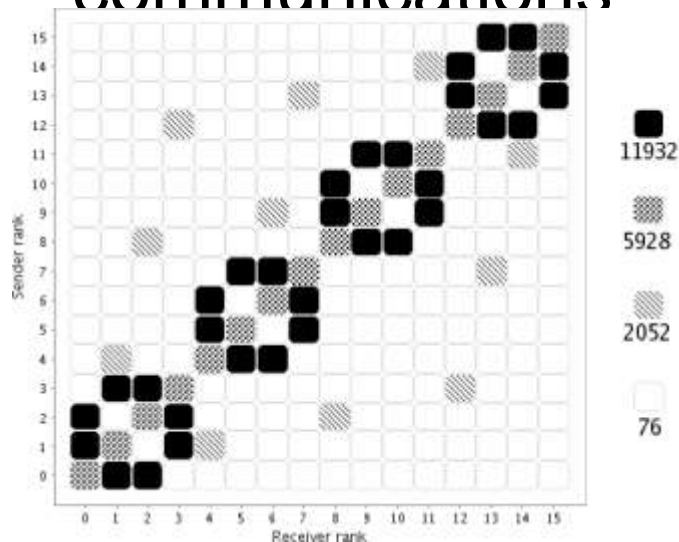


Communication Intensive CG Kernel (Conjugate Gradient)

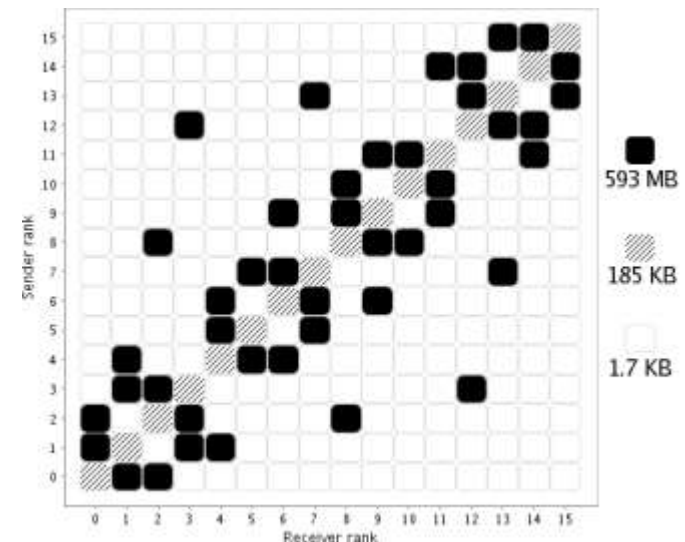
- ❑ Floating point operations
- ❑ Eigen value computation
- ❑ High number of unstructured communications



- 12000 calls/node
- 570 MB sent/node
- 1 min 32
- 65 % comms/WT

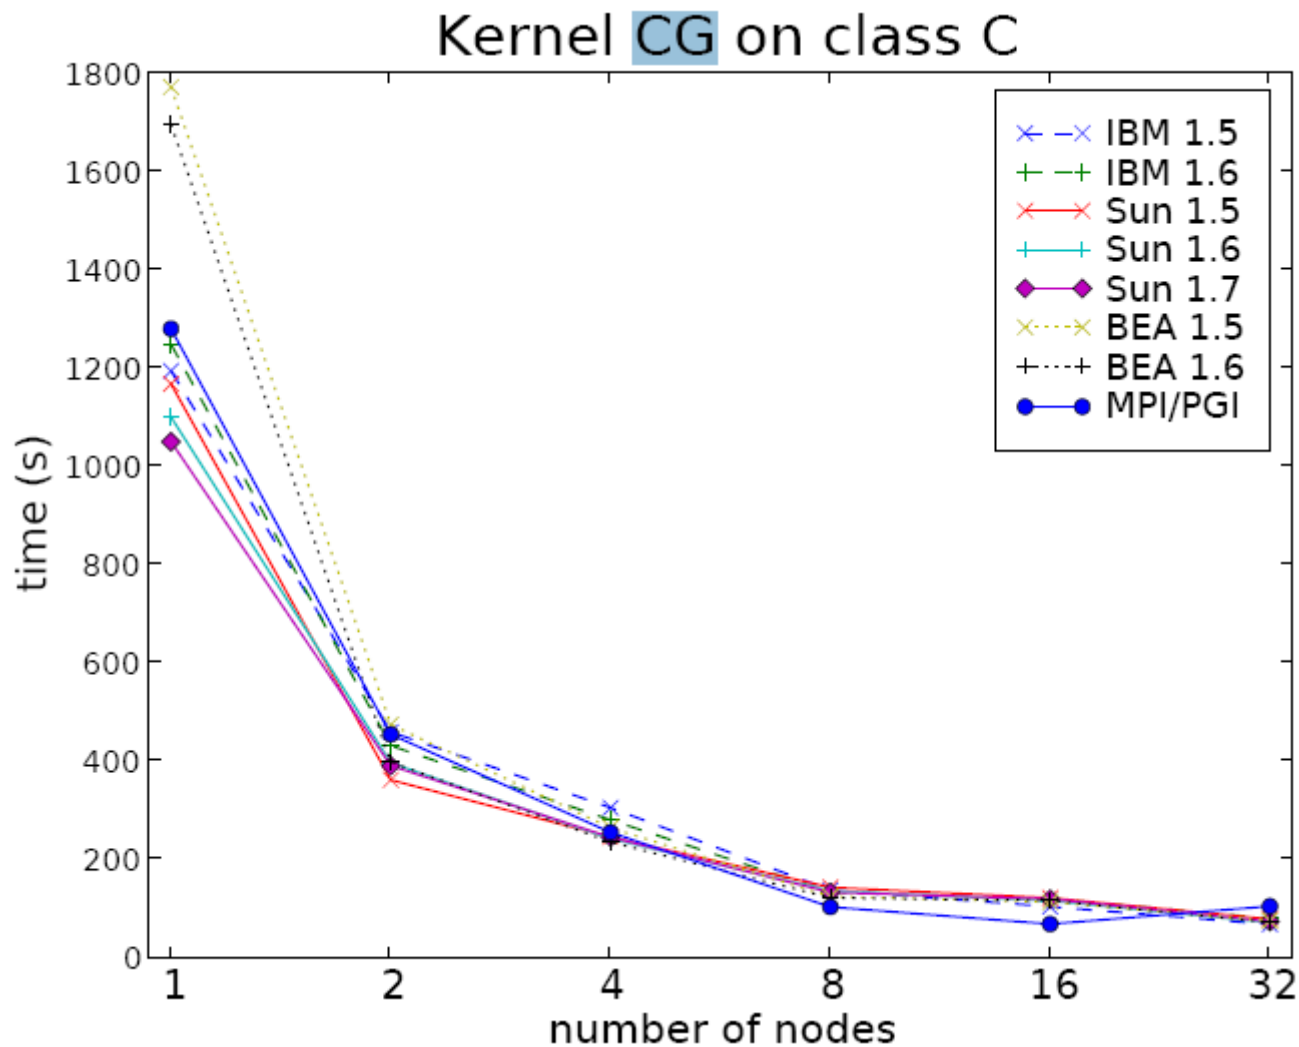


Message density distribution



Data density distribution

Communication Intensive CG Kernel (Conjugate Gradient)



← Comparable Performances

Key Point: Locality will more than ever be Fundamental

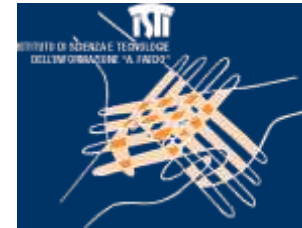
- Let the programmer control it
- No global shared memory
- PGAS like**
- Partitioned Global Address Space**
- But with more
- Flexibility, Dynamicity and Control

- One can envision: Spatial view of multicore



□ Research for
□ High-Level Parallel Abstractions

GridCOMP Partners



THE UNIVERSITY OF
MELBOURNE

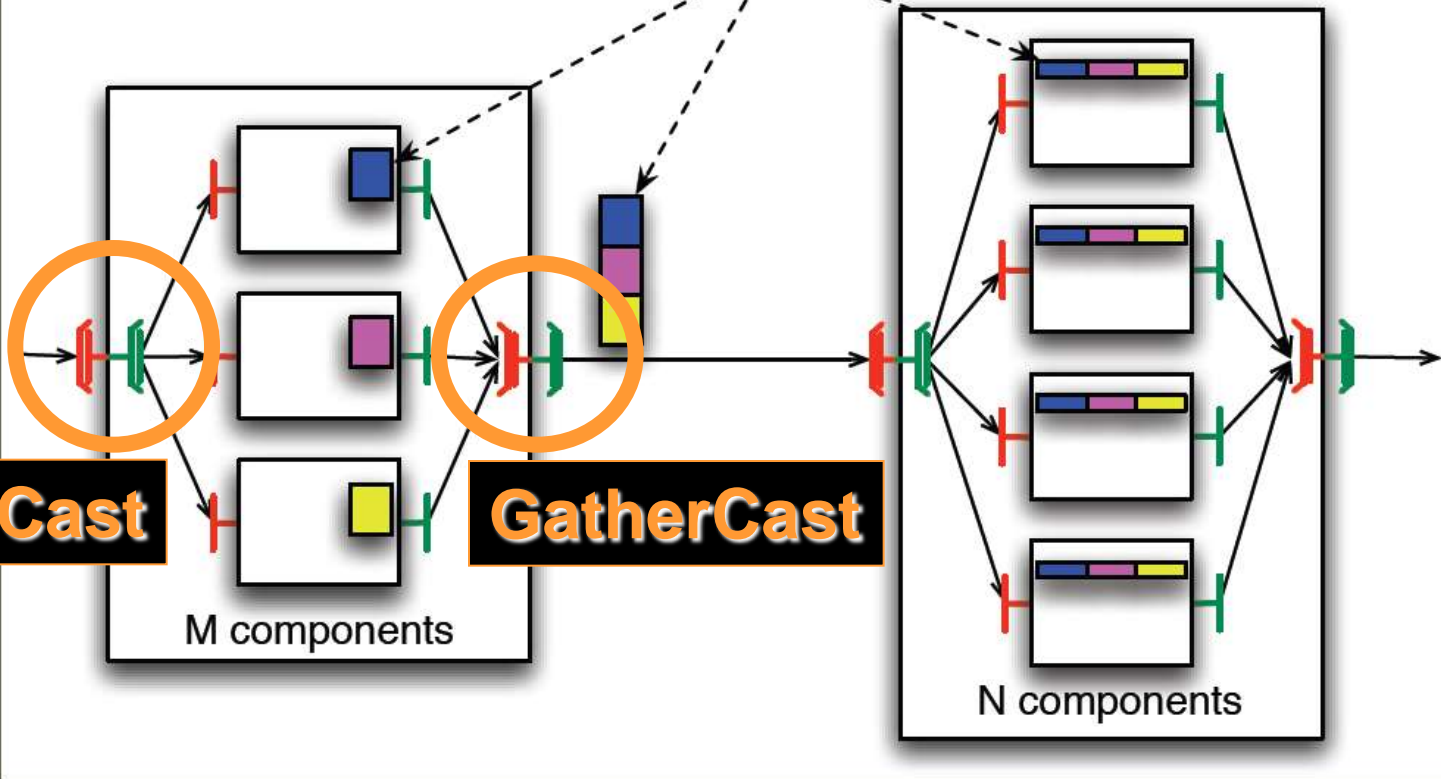


M

Jem3D

Steering and Visualisation GUI

invocation parameters

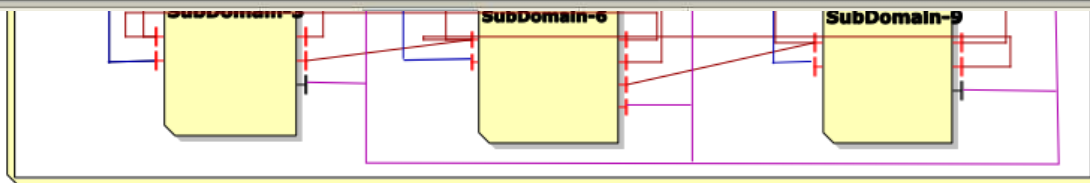


MultiCast

GatherCast

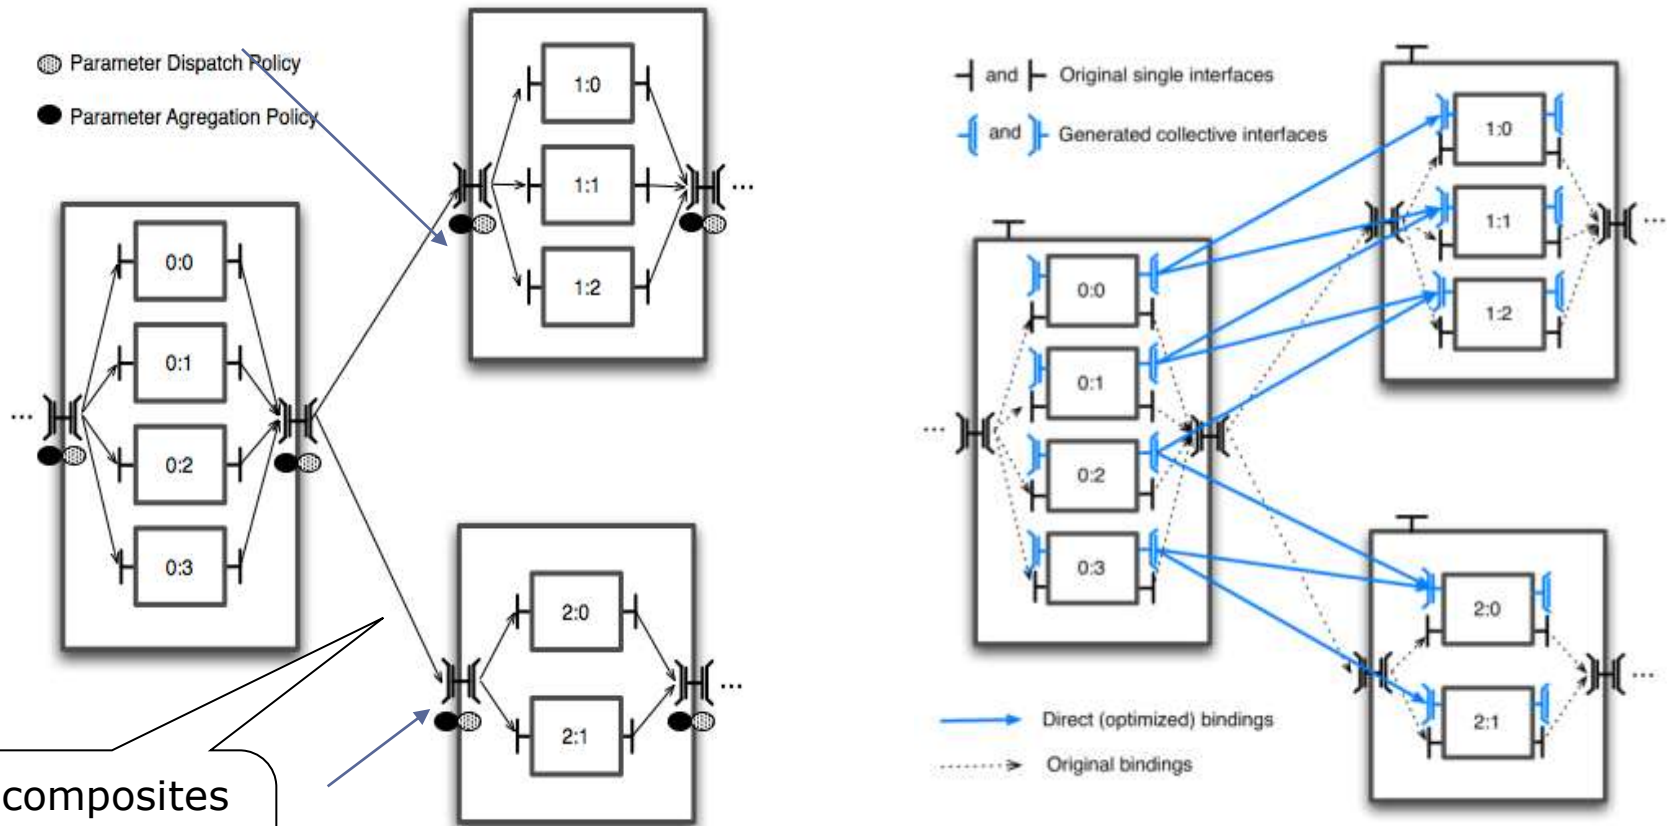
M components

N components



Sc
 Gr
 No
 Ab
 Co
 Mu

Optimizing MxN Operations



2+ composites can be involved in the Gather-multicast

Key Points about Parallel Components

- ❑ Parallelism is captured at the Module interface
- ❑ Identical to Typing for functional aspects
- ❑ Composition, in a parallel word, becomes possible
- ❑ Configuration of the Parallel *aspect*

Optimizing

ProActive Parallel Suite

PROGRAMMING

Java Parallel Frameworks

for HPC, Multi-Cores, Distribution, Enterprise Grids and Clouds.

Featuring: Async. comms, Master-Worker, Monte-Carlo, SPMD, components and legacy code wrapping.

OPTIMIZING

Eclipse GUI (IC2D)

for Developing, Debugging, Optimizing your parallel applications.

Featuring: graphical monitoring and benchmarking with report generation.

The screenshot shows the Eclipse IDE interface with the Monitoring View and Job Monitoring View. The Monitoring View displays a network diagram of virtual nodes and their connections. The Job Monitoring View shows a tree structure of job components.

Monitoring View:

- Virtual nodes: Renderer, InitialVM, Dispatcher, User
- Host: bebita.inria.fr:1099:OS u...
- Node PA_JVM1357457629_... (Node Node60562498...):
 - DinnerLayout#2
 - Table#3
 - Philosopher#4
 - Philosopher#5
 - Philosopher#6
 - Philosopher#7
 - Philosopher#8
- Node PA_JVM436155261_... (Node Renderer-127...): C3DRendering...
- Node PA_JVM1672076495_... (Node Dispatcher-5...): C3DDispatche...
- Node PA_JVM294719007_... (Node User16026446...): C3DUser#13
- Node PA_JVM1631909824_... (Node Renderer1307...): C3DRendering...

Job Monitoring View:

- DefaultVN (JOB-1357457629!)
- bebita.inria.fr:1099:OS un
 - PA_JVM1357457629_ (blue square)
 - Node Node60562498... (blue square)
 - DinnerLayout#2 (red circle)
 - Table#3 (JOB-13... (red circle)
 - Philosopher#4 (red circle)
 - Philosopher#5 (red circle)
 - Philosopher#6 (red circle)
 - Philosopher#7 (red circle)
 - Philosopher#8 (red circle)
 - sidonie.inria.fr:1099:OS u...
 - Dispatcher (JOB--1672076495... (blue square)
 - User (JOB--294719007) (blue square)
 - bebita.inria.fr:1099:OS un
 - PA_JVM294719007_... (blue square)
 - Node User1602644... (blue square)
 - C3DUser#13 (red circle)
 - Renderer (JOB--1672076495... (blue square)
 - bebita.inria.fr:1099:OS un
 - PA_JVM1631909824_... (blue square)

IC2D

The screenshot displays the IC2D monitoring application interface. The main window is titled "Monitoring#1" and shows a network topology of virtual nodes. The nodes are connected in a hierarchical structure, including PA_JVM1820960857, Node Node-632703901, Node matrixNode15..., OctTree#2, Domain#3, Domain#4, Domain#5, Domain#6, Maestro#7, and BigMaestro#8. The legend on the right lists active objects and pending requests, with colors corresponding to the nodes in the topology. The "Timer Tree View" on the right shows a detailed breakdown of request processing times for Domain#5 and Domain#4. The "Bar Chart" in the bottom left shows the distribution of request types for Domain#4. The "Time Line View" at the bottom shows the execution timeline for various components, including BigMaestro#8, Maestro#7, Domain#6, Domain#5, Domain#4, Domain#3, and OctTree#2.

Timer Tree View Data:

Name	Time [ms]	Total [%]	Invocations	Parent [%]
Domain#5				
Total	142212.28	100.00	1	0.00
WaitForRequest	21627.76	15.21	2056	15.21
Serve	120543.91	84.76	5352	84.76
SendReply	0.00	0.00	0	0.00
WaitByNecessity	17050.55	11.99	5340	14.14
SendRequest	101773.58	71.56	16054	84.43
Domain#4				
Total	142228.27	100.00	1	0.00
WaitForRequest	21249.88	14.94	2114	14.94
Serve	120936.36	85.03	5353	85.03
SendReply	0.00	0.00	0	0.00
GroupOneWayCall	0.00	0.00	0	0.00
GroupAsyncCall	0.00	0.00	0	0.00
WaitByNecessity	16765.29	11.79	5348	13.86
SendRequest	102320.24	71.94	16057	84.61
Serialization	1101.89	0.77	5352	1.08
LocalCopy	2471.16	1.74	10705	2.42
BeforeSerializati	20631.26	14.51	5352	20.16

Bar Chart Data (Domain#4):

Request Type	Time
GroupAsyncCall	1.28ms
GroupOneWayCall	1.28ms
AfterSerialization	1.19ms
Serialization	1.19ms
BeforeSerialization	20.63ms
LocalCopy	2.47ms
WaitForRequest	21.24ms
WaitByNecessity	16.76ms
SendReply	1.79ms
SendRequest	2.81ms
Serve	2.97ms
Total	2.97ms

Time Line View Data:

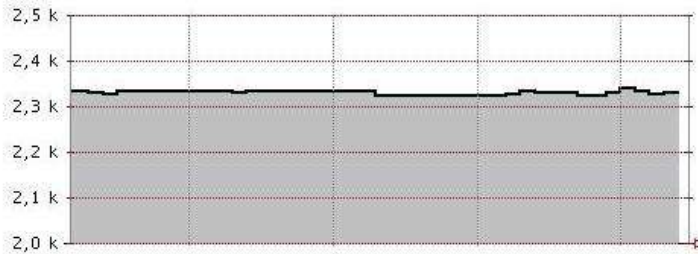
Component	Start Time	End Time
BigMaestro#8	0ms	746.4ms
Maestro#7	0ms	746.4ms
Domain#6	0ms	746.4ms
Domain#5	0ms	746.4ms
Domain#4	0ms	746.4ms
Domain#3	0ms	746.4ms
OctTree#2	0ms	746.4ms

ChartIt

PA_JVM251111462

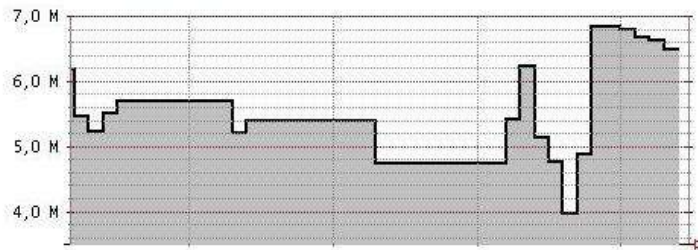
Charts

Chart#1 [LoadedClassCount]



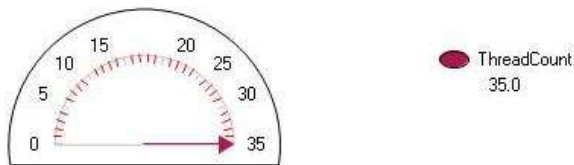
LoadedClassCount

Chart#2 [UsedHeapMemory]

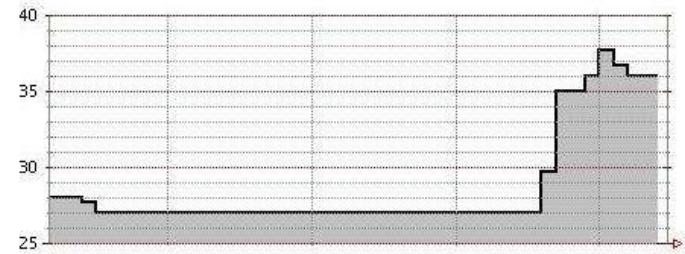


UsedHeapMemory

Chart#5 [ThreadCount] Meter

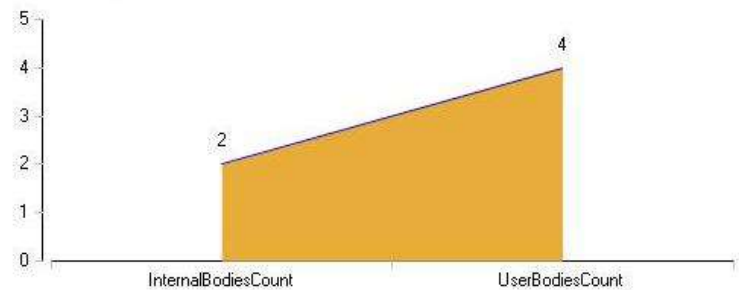


Chart#3 [ThreadCount]

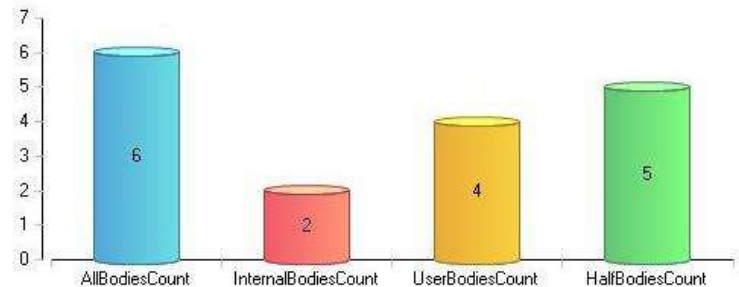


ThreadCount

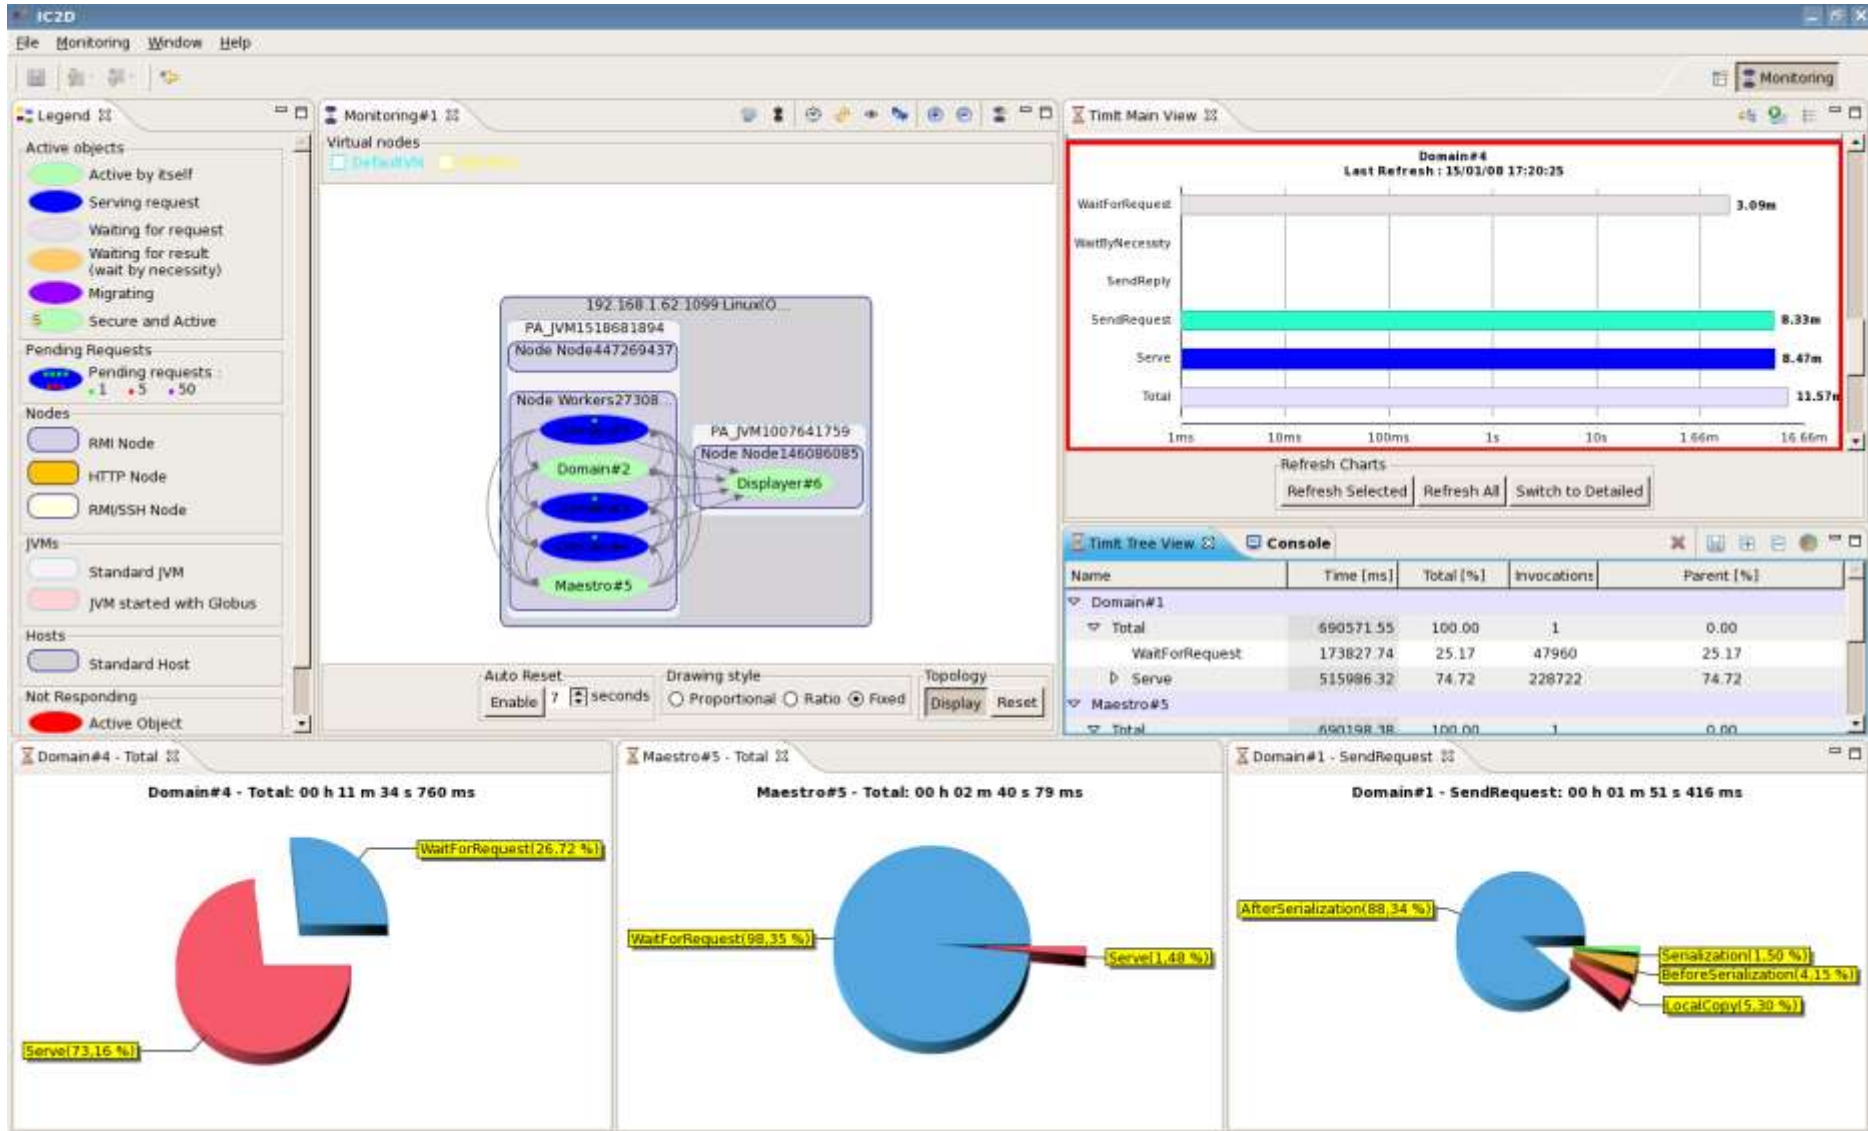
Chart#4 [InternalBodiesCount, UserBodiesCount]



Chart#6 [AllBodiesCount, InternalBodiesCount, UserBodiesCount, HalfBodiesCount]



Pies for Analysis and Optimization



Video 1: IC2D Optimizing Monitoring, Debugging, Optimizing



3. Scheduling

ProActive Parallel Suite

PROGRAMMING

Java Parallel Frameworks

for HPC, Multi-Cores, Distribution, Enterprise Grids and Clouds.

Featuring: Async. comms, Master-Worker, Monte-Carlo, SPMD, components and legacy code wrapping.

OPTIMIZING

Eclipse GUI (IC2D)

for Developing, Debugging, Optimizing your parallel applications.

Featuring: graphical monitoring and benchmarking with report generation.

SCHEDULING

Multi-Language Scheduler

for Workflows made of C, C++, Java, Scripts, Matlab, Scilab tasks.

Featuring: graphical user interface, resource acquisition and virtualization.

Scheduler: User Interface

ProActive Scheduler

File Window Help

Scheduler

Jobs

Pending (8)

Id	State	User	Priority	Name
172	Pending	user1	Low	job_2_tasks
173	Pending	user1	Low	job_2_tasks
174	Pending	user1	Low	job_2_tasks
176	Pending	user1	Low	job_2_tasks
177	Pending	user1	Low	job_2_tasks
178	Pending	user1	Low	job_2_tasks
179	Pending	user1	Low	job_2_tasks
180	Pending	user1	Low	job_2_tasks

Running (13)

Id	State	Progress	# Finishes	User	Priority	Name
54	Running	<div style="width: 50%;"></div>	1/2	user1	Low	job_2_t
55	Running	<div style="width: 50%;"></div>	0/2	user1	Low	job_2_t
56	Running	<div style="width: 50%;"></div>	1/2	user1	Low	job_2_t
160	Running	<div style="width: 50%;"></div>	1/2	user1	Low	job_2_t
161	Running	<div style="width: 50%;"></div>	1/2	user1	Low	job_2_t
162	Running	<div style="width: 50%;"></div>	1/2	user1	Low	job_2_t
163	Running	<div style="width: 50%;"></div>	1/2	user1	Low	job_2_t
164	Running	<div style="width: 50%;"></div>	1/2	user1	Low	job_2_t
165	Running	<div style="width: 50%;"></div>	1/2	user1	Low	job_2_t
166	Running	<div style="width: 50%;"></div>	1/2	user1	Low	job_2_t
168	Running	<div style="width: 50%;"></div>	0/2	user1	Low	job_2_t
169	Running	<div style="width: 50%;"></div>	0/2	user1	Low	job_2_t
170	Running	<div style="width: 50%;"></div>	0/2	user1	Low	job_2_t

Finished (11)

Id	State	User	Priority	Name
152	Finished	user1	Low	job_2_tasks
167	Finished	user1	Normal	job_2_tasks
171	Finished	user1	Normal	job_2_tasks
153	Finished	user1	Low	job_2_tasks
175	Finished	user1	Normal	job_2_tasks
154	Finished	user1	Low	job_2_tasks
155	Finished	user1	Low	job_2_tasks
156	Finished	user1	Low	job_2_tasks
157	Finished	user1	Low	job_2_tasks
158	Finished	user1	Low	job_2_tasks
159	Finished	user1	Low	job_2_tasks

STARTED

Console Tasks Users

job 55 has 2 tasks

Id	State	Name	Host name	Start time	Finished time	Re-run	Description
55000	Running	task1	eon8.inria.fr	16:09:28 08/27/08	Not yet	0/3	task WaitAndPrint - will sleep for 3s
55000	Running	task2	eon8.inria.fr	16:09:28 08/27/08	Not yet	0/1	task WaitAndPrint - will sleep for 20s

Job Info Result Preview

Property	Value
Id	55
State	Running
Name	job_2_tasks
Priority	Low
Pending tasks number	0
Running tasks number	2
Finished tasks number	0
Total tasks number	2
Submitted time	16:09:28 08/27/08

Video 2: Scheduler, Resource Manager



4. Enterprise Grids, Clouds: Standards & Amazon EC2

GCM Standardization

Grid Component Model



Overall, the standardization is supported by industrials:

BT, FT-Orange, Nokia-Siemens, NEC, Telefonica, Alcatel-Lucent, Huawei ...

GRIDS for Finance
& Telecommunications

V

GRID
20-24 Oct
INRIA – Sophia



Information & Registration: www.etsi.org/plugtests/GRID2008/GRID.htm

With the support of



European partners

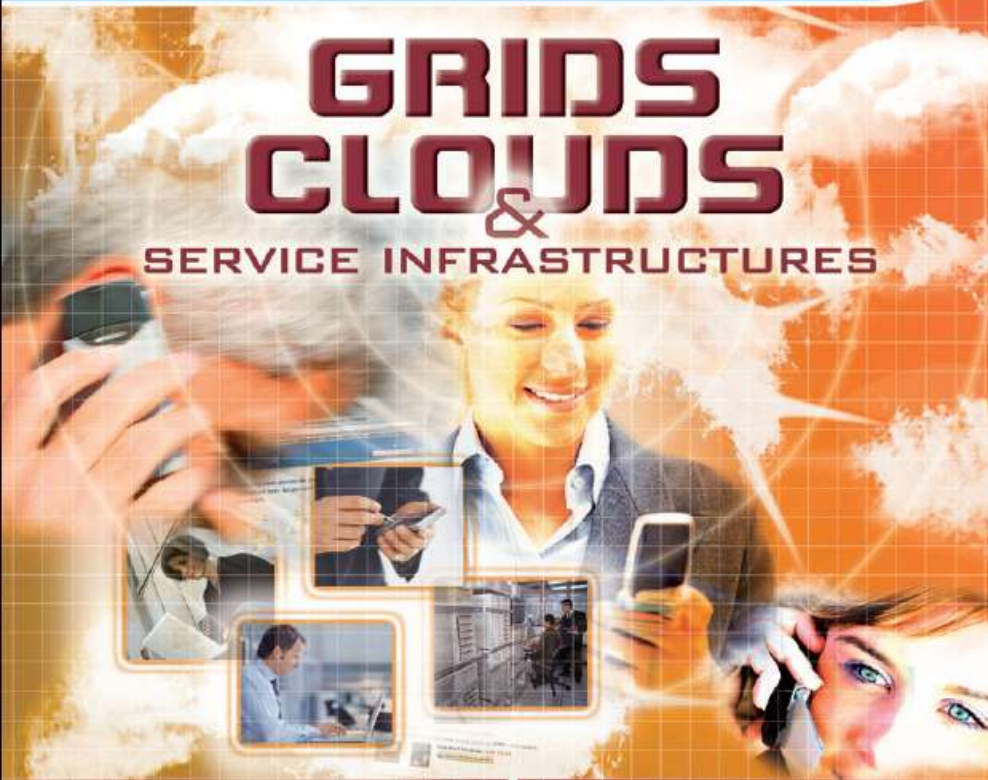


Industrial sponsors



GRIDS & CLOUDS

SERVICE INFRASTRUCTURES



PLUGTESTS
30 NOV. - 02 DEC.

WORKSHOP
02 - 03 DEC.

Information & registration at <http://www.etsi.org/plugtests/GRID09/GRID.htm>

30 NOV. - 03 DEC. 2009
ETSI - FRANCE



Summary

Summary of Key Points

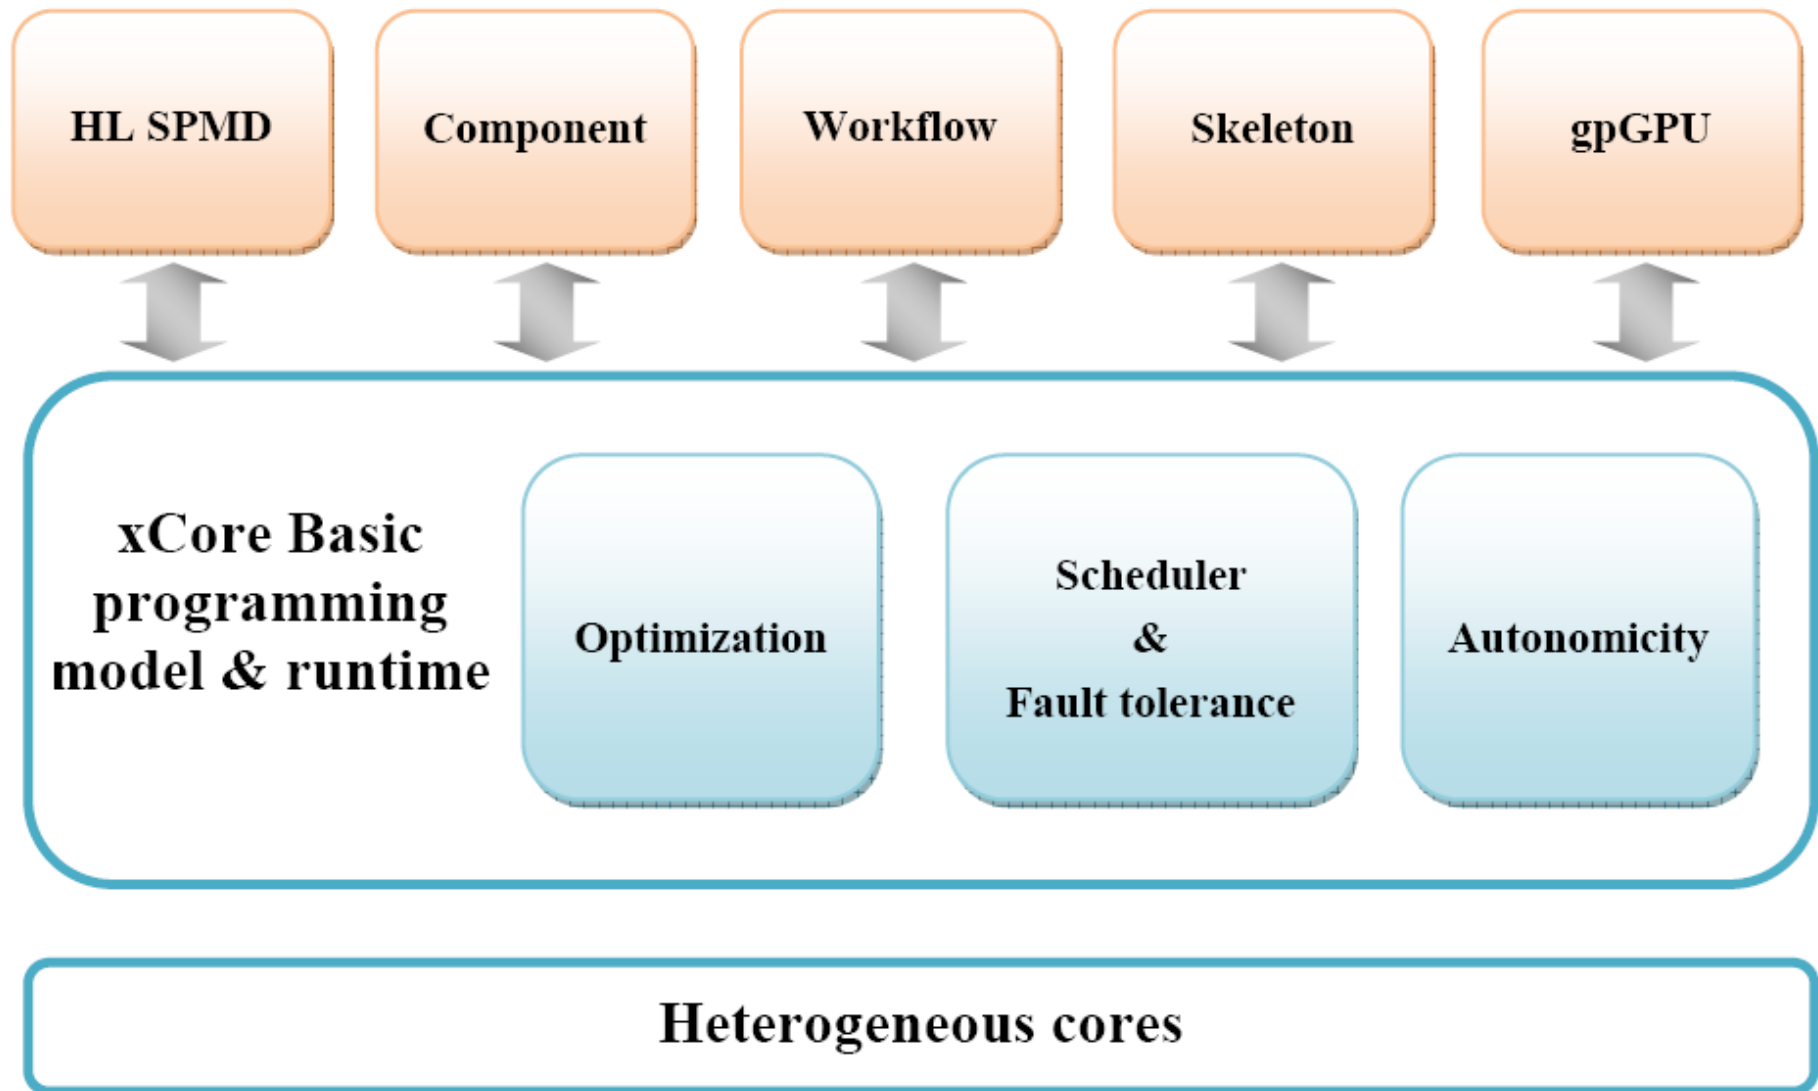
- ❑ **Multi-Cores are NUMA, and turning Heterogeneous (GPU)**
 - They are turning into SoC with NoC: NOT SMP!
 - Smooth evolution needed: Distributed to Multi-core

- ❑ **A need for a unified Parallel Abstraction:**
- ❑ **Multi-Core + Distributed**
- ❑ **Shall MPI and OpenMP RIP**
 - Application Abstractions Not Infrastructure Abstractions
 - Maintain strong Programmer control on Locality

Other Evolutions

- ❑ **Scheduling of Asynchronous Tasks, Workflows, Dynamic Data Driven Execution**
- ❑ **Fault-Tolerance + Need for QoS and SLA:**
 - ❑ **→ Self-Adapting Auto-Tuning systems**

Summary and Perspectives: On-going

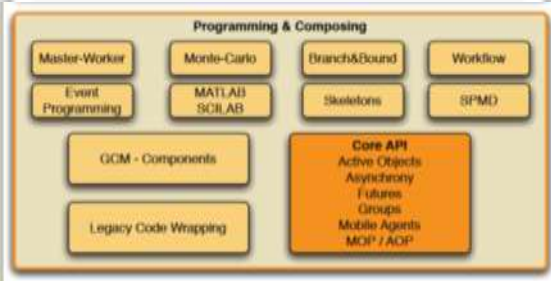


Conclusion: Currently Available



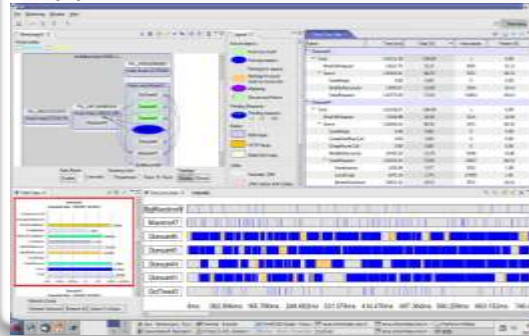
PROGRAMMING

Java Parallel Frameworks
for HPC, Multi-Cores,
Distribution, Enterprise
Grids and Clouds.



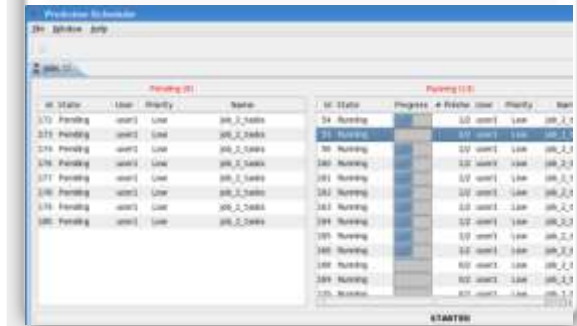
OPTIMIZING

Eclipse GUI (IC2D)
for Developing, Debugging,
Optimizing your parallel
applications.



SCHEDULING

Multi-Language Scheduler
for Workflows made of C,
C++, Java, Scripts, Matlab,
Scilab tasks.



**Further into the direction of:
Multi-Core + Distributed**

AGOS: Grid Architecture for SOA

Building a Platform for Agile SOA with Grid

□ AGOS Solutions



In Open Source with Professional Support

