



Massive Ray Tracing in Fusion Plasmas on EGEE

J.L. Vázquez-Poletti (1), E. Huedo (2), R.S. Montero (1) and I.M. Llorente (1,2)

(1) Universidad Complutense de Madrid (Spain)


(2) Centro de Astrobiología (Spain)

What are we going to see?

MA-RA-TRA: a computational view
Using the LCG-2 Infrastructure
What if... GridWay?
Comparison
Conclusions
Our two cents

32444245304354

MA-RA-TRA: a computational view



MA-RA-TRA: “Massive Ray Tracing” in Fusion Plasmas

Application profile:

- Sizes

 - Executable (Truba) – 1.8 MB

 - Input files – 70 KB

 - Output files – about 459 KB

- Execution Time – about 26 minutes

 - Pentium 4 (3.2 GHz)

- 1 execution = 1 ray traced

Using the LCG-2 Infrastructure (1)



lcg2.1.9 User Interface C++ API

1 job = 1 ray

Procedure:

- Launcher script

 - Generates JDL files

- MRT framework

 - Launches them simultaneously

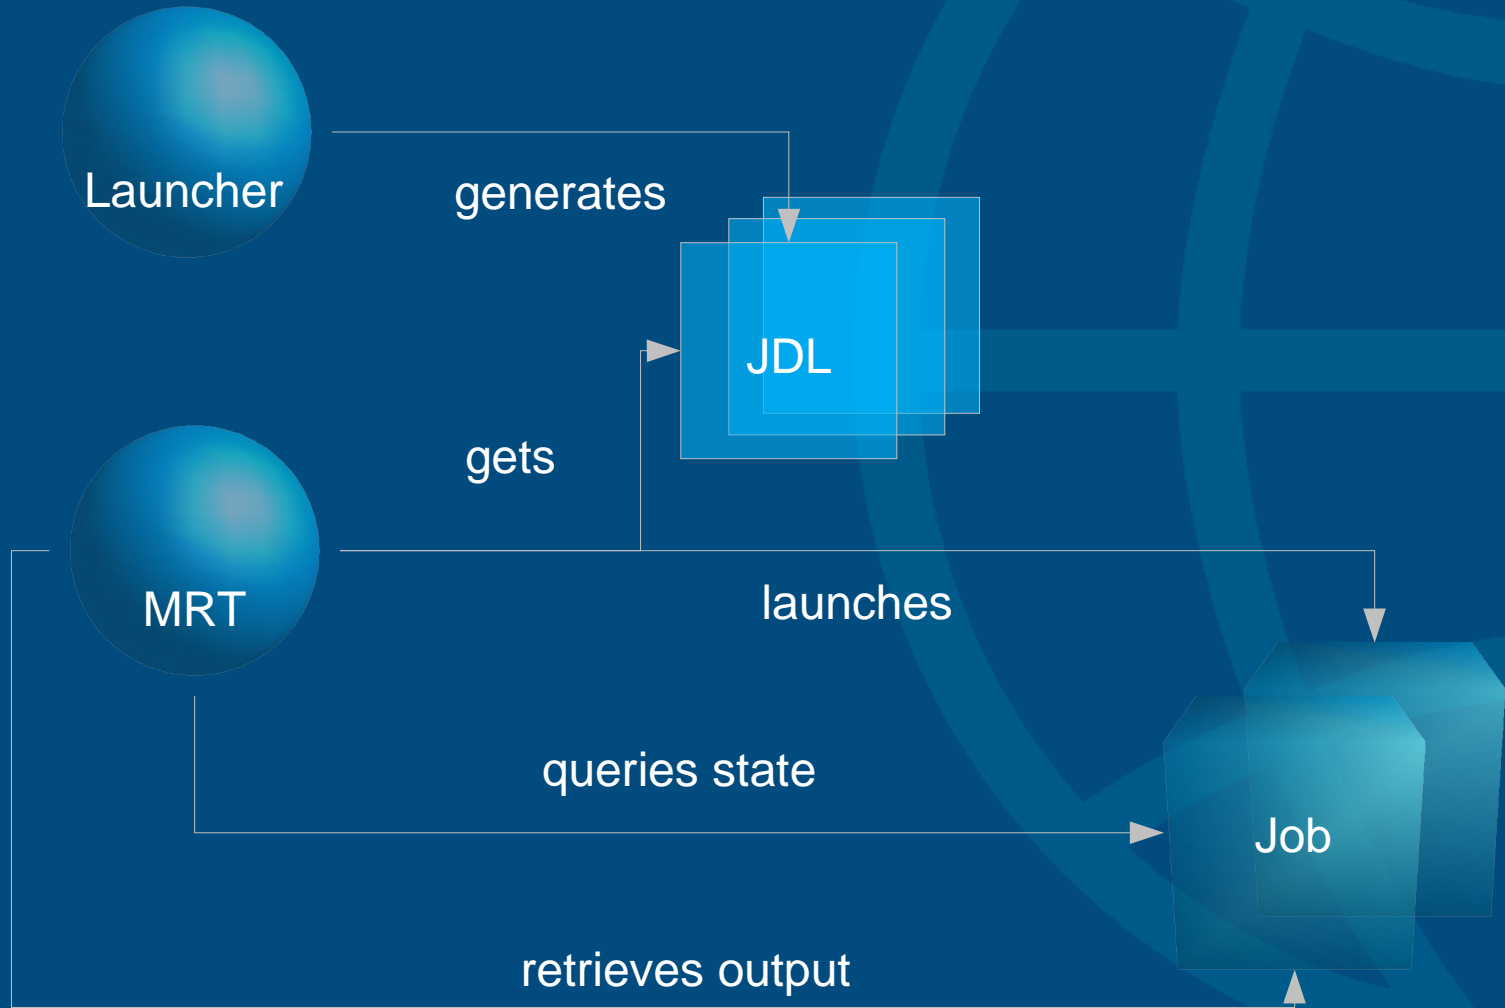
 - Queries each job's state periodically

 - Retrieves each job's output (Sandbox)

Using the LCG-2 Infrastructure (2)



32444245304354



Using the LCG-2 Infrastructure (3)



SWETEST VO

Sites involved in the experiments:

- CESGA (46 x Intel Pentium III 1133 MHz)
- IFIC (127 x AMD Athlon 1.2 GHz)
- PIC (172 x Intel Pentium 4 2.8 GHz)
- IFAE (11 x Intel Pentium 4 2.8 GHz)
- LIP (25 x Intel Xeon 2.8 Ghz)

Spanish sites conected by RedIRIS

- 622 Mbps to 2.5 Gbps

Using the LCG-2 Infrastructure (4)



Total Time: 220 minutes (3.67 hours)

Execution Time:

- Average: 30.33 minutes
- Std. Deviation: 11.38 minutes

Transfer Time:

- Average: 0.42 minutes
- Std. Deviation: 0.06 minutes


Avg. Productivity: 13.36 Jobs/hour

Avg. Overhead: 90.99 minutes/job

Using the LCG-2 Infrastructure (5)



Jobs Allocated (LCG-2)



3
2
4
4
4
2
4
5
3
0
4
3
5
4

Using the LCG-2 Infrastructure (6)



As in the real world, some jobs failed

- Jobs affected: 31
- Max resubmissions/job: 1

Problems encountered:

- LCG-2 Infrastructure:
 - Lack of opportunistic migration
 - Lack of fault tolerance mechanisms
- The API itself:
 - Submitting more than 80 jobs in a Collection

What if... GridWay? (1)



Light-weight framework

Works on top of Globus services

Performs:

- Job execution management
- Resource brokering

Allows unattended, reliable and efficient execution of:

- single jobs, array jobs, complex jobs
- on heterogeneous, dynamic and loosely-coupled grids

What if... GridWay? (2)



Works transparently to the end user

Adapts job execution to changing Grid conditions

- Fault recovery
- Dynamic scheduling
- Migration on-request

Scheduling using Information System (GLUE schema) from LCG-2

Stands on the client side

What if... GridWay? (3)



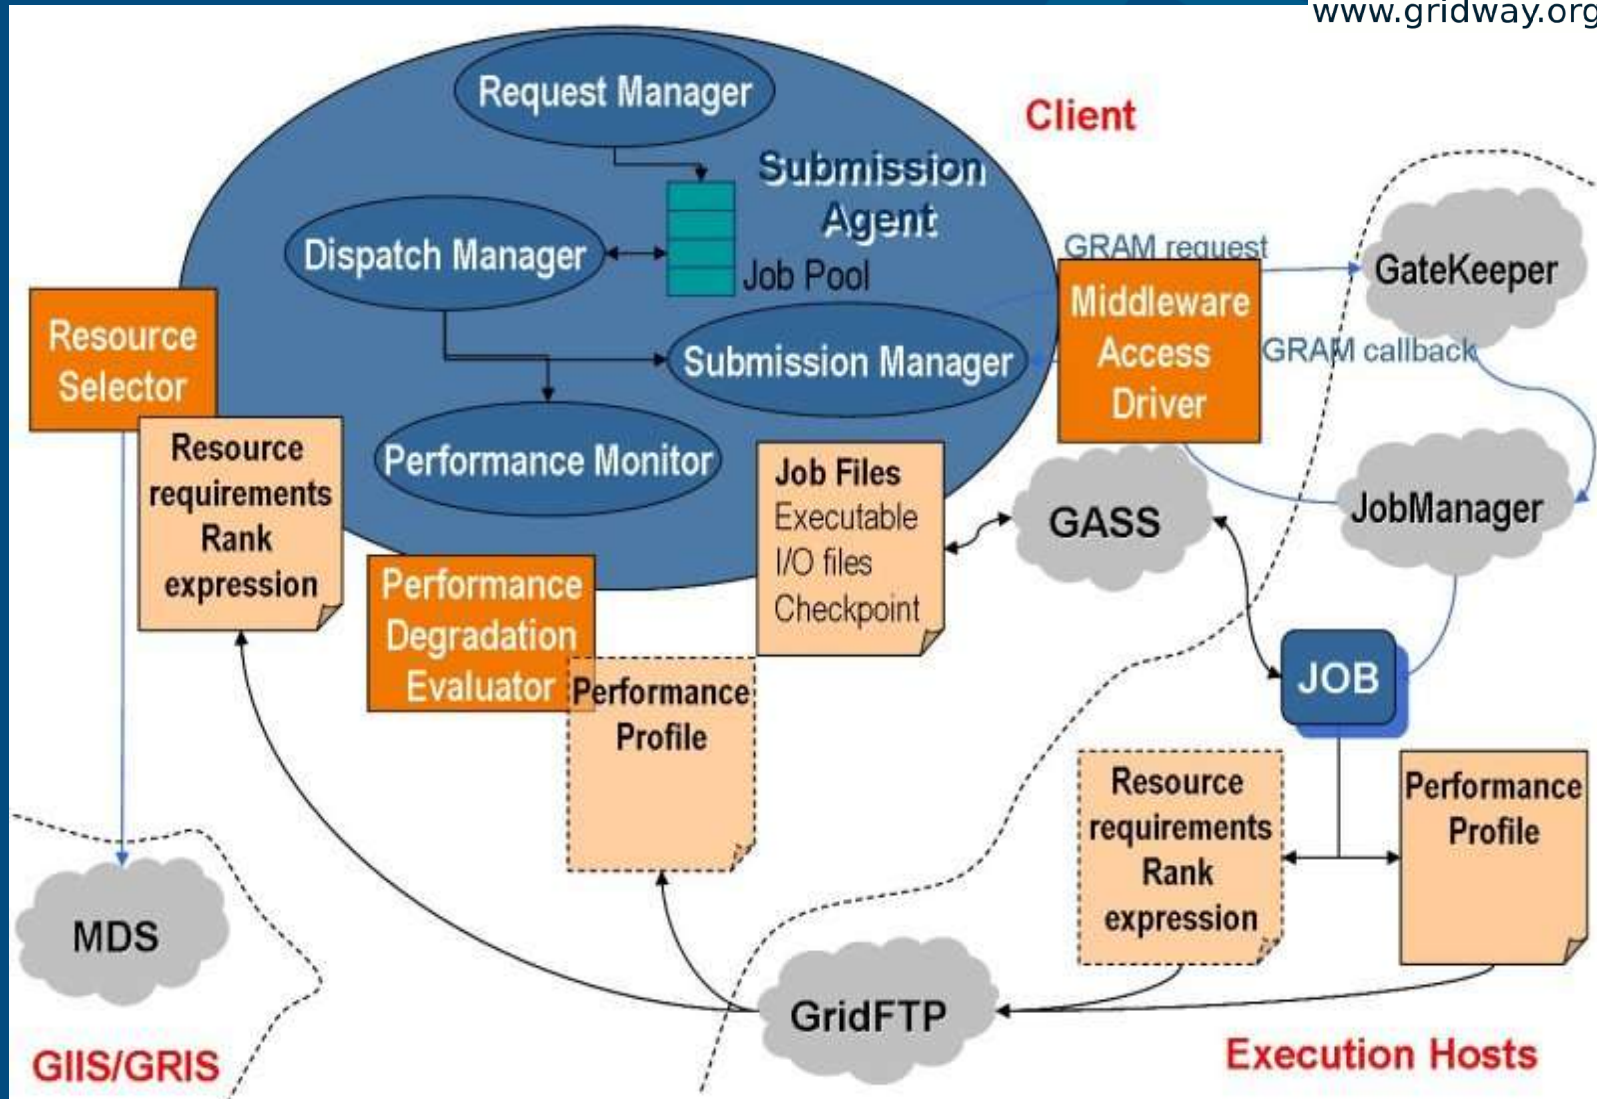
Execution as seen by GridWay:

- Prolog: prepares remote system
 - Creates directory
 - Transfers input files and executable
- Wrapper: executes job and gets exit code
- Epilog: finalizes remote system
 - Transfers output files
 - Cleans up directory

What if... GridWay? (4)



www.gridway.org



32444245304354

What if... GridWay? (5)



SWETEST VO

Sites involved in the experiments:

- CESGA (46 x Intel Pentium III 1133 MHz)
- IFIC (127 x AMD Athlon 1.2 GHz)
- INTA-CAB (4 x Intel Pentium 4 2.8 GHz)
- IFAE (11 x Intel Pentium 4 2.8 Ghz)
- USC (100 x Intel Pentium III 1133 Mhz)

Spanish sites conected by RedIRIS

- 622 Mbps to 2.5 Gbps

What if... GridWay? (6)



Total Time: 123.43 minutes (2.06 hours)

Execution Time:

- Average: 36.8 minutes
- Std. Deviation: 16.23 minutes

Transfer Time:

- Average: 0.87 minutes
- Std. Deviation: 0.51 minutes


Avg. Productivity: 23.82 Jobs/hour

Avg. Overhead: 25.99 minutes/job

What if... GridWay? (7)



Jobs Allocated (GridWay)



32444245304354


What if... GridWay? (8)



Also with GridWay, some jobs failed

- Jobs affected: 9
- Total reschedules: 22
- Max. reschedules/job: 4

GridWay Resched Reasons

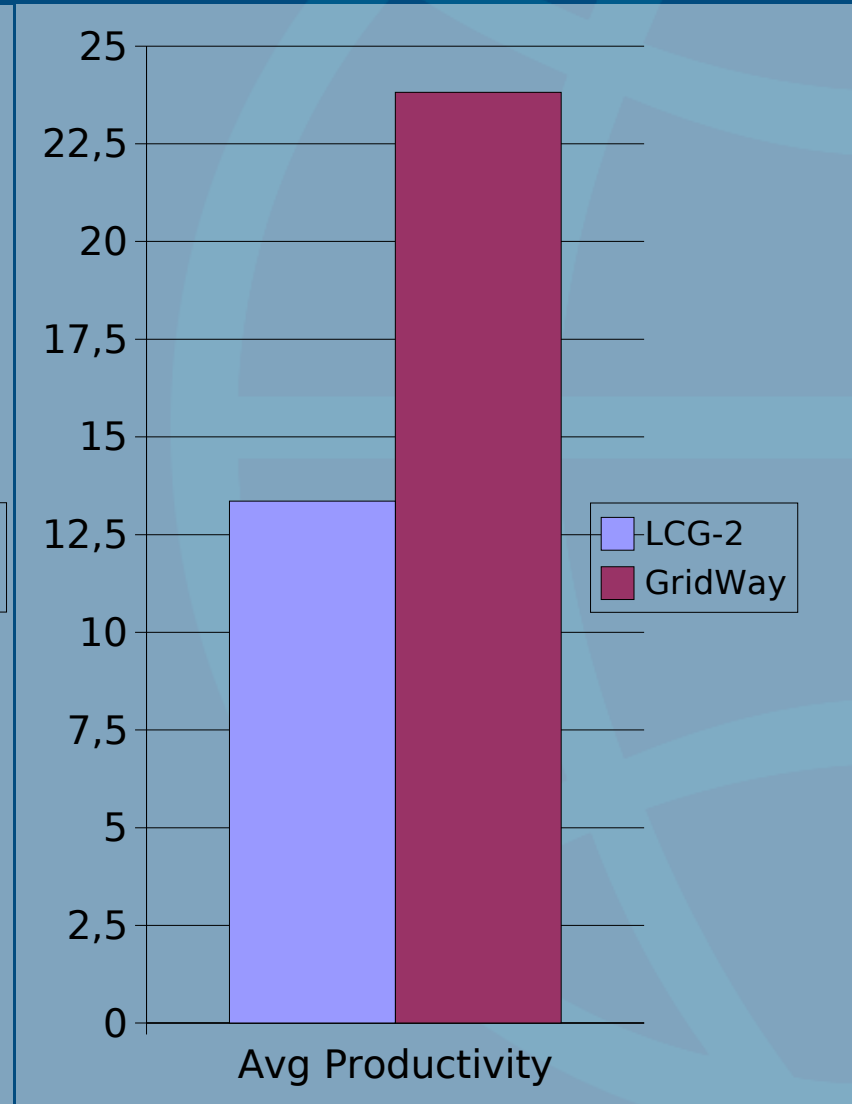



32444245304354

Comparison




32444245304354




Comparison



Execution Time

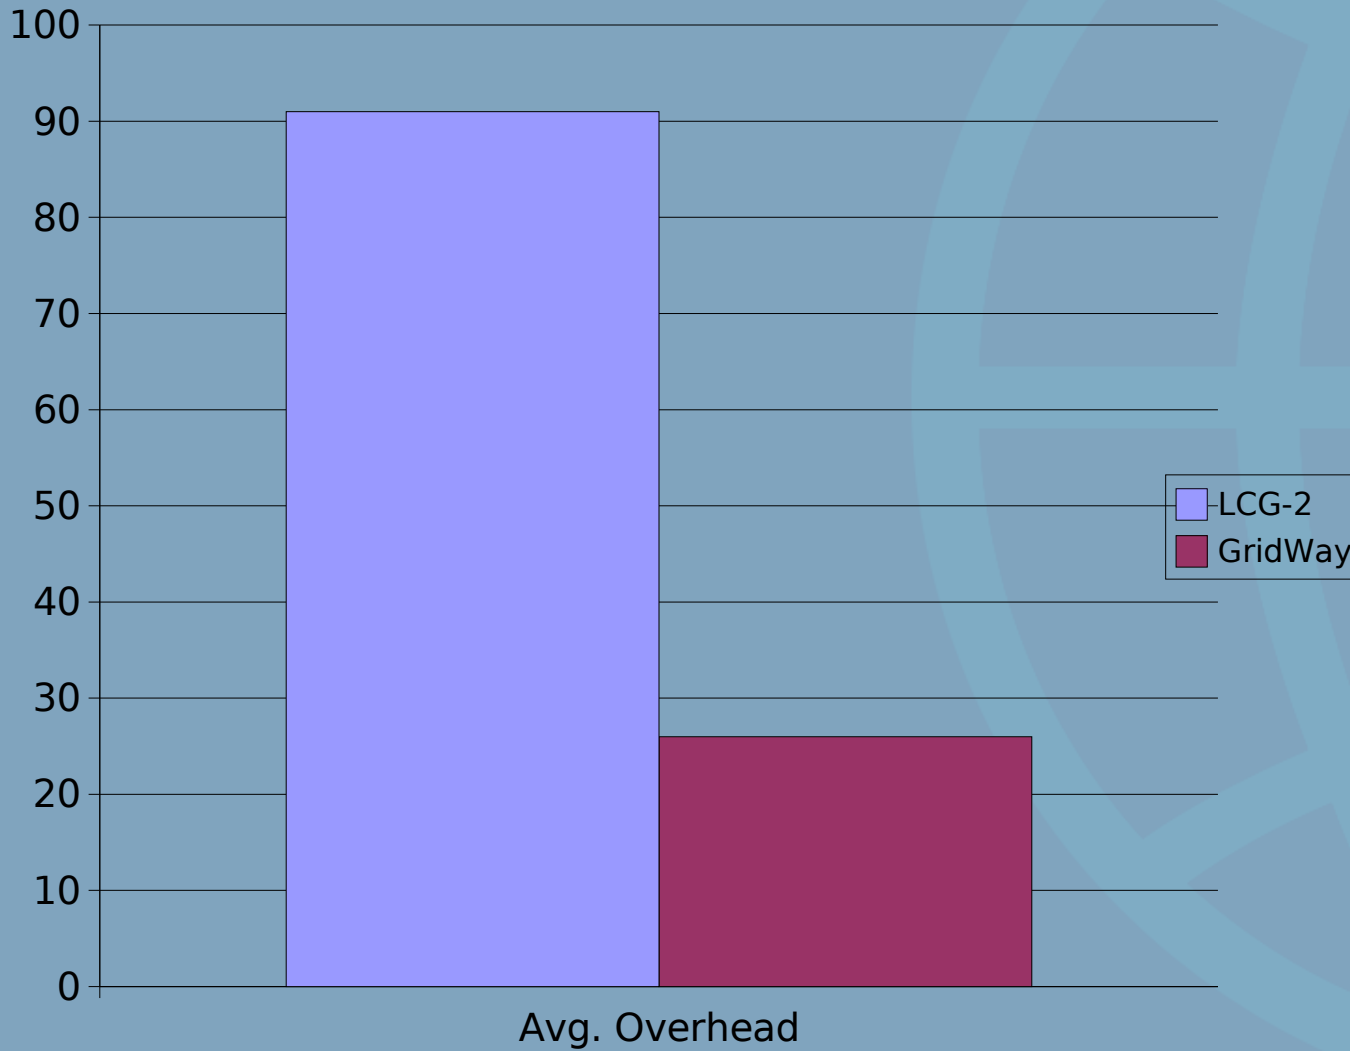


Transfer Time



32444245304354

Comparison

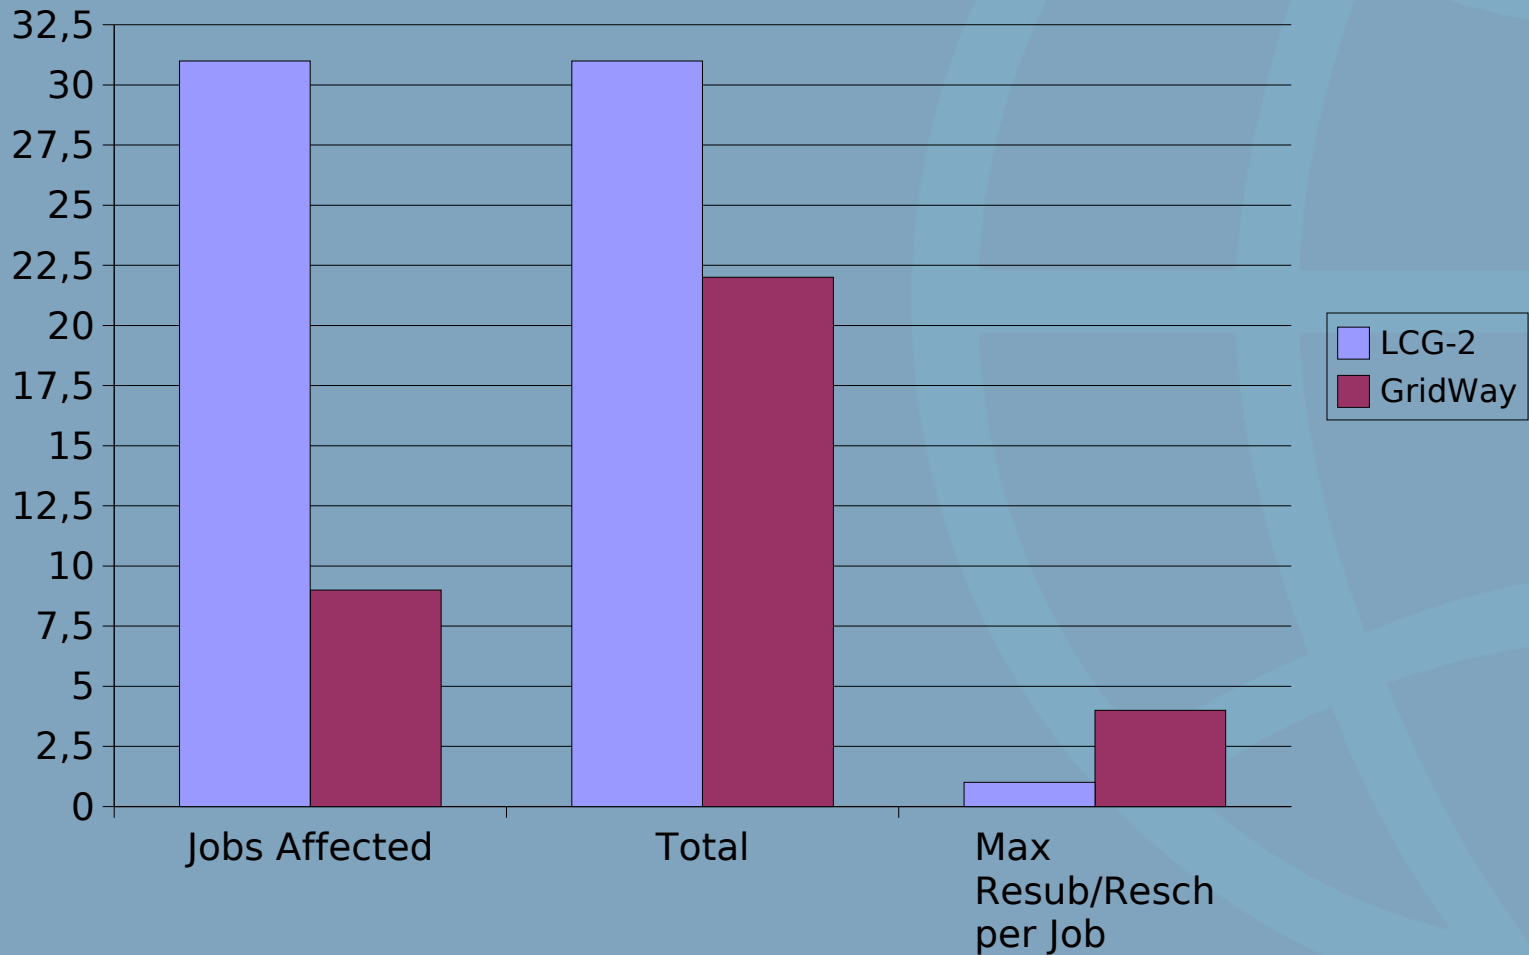


32444245304354

Comparison



Resubmissions/Reschedules

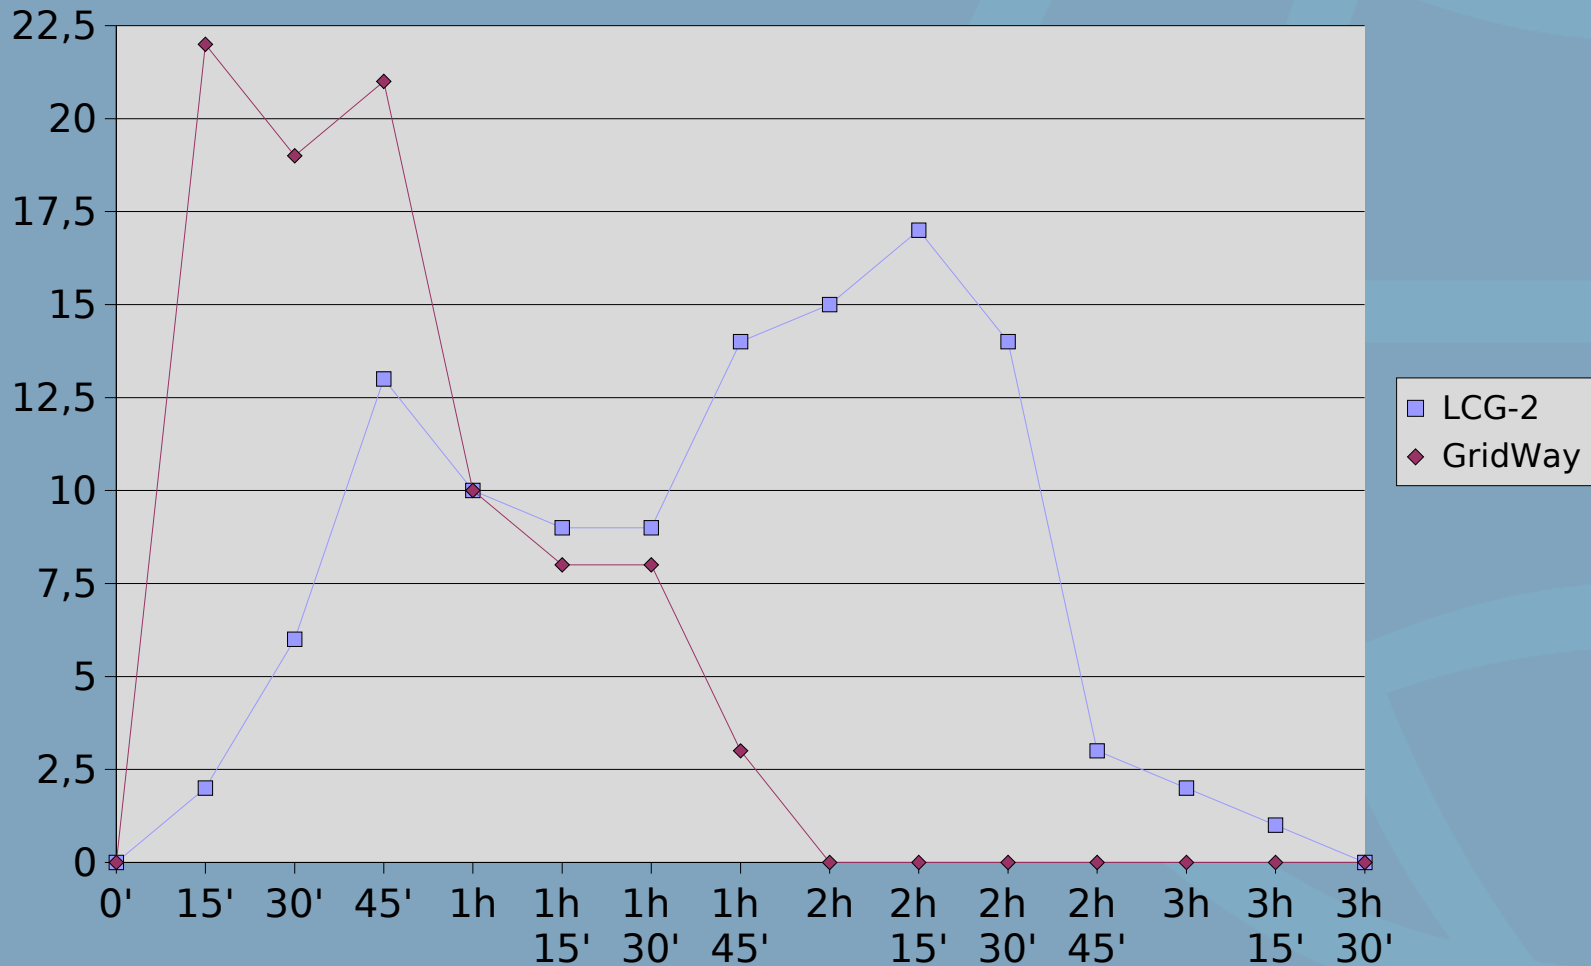


32444245304354

Comparison



Jobs Allocated (Every 15')




32444245304354

Comparison



Productivity (Every 15')



32444245304354

Conclusions



GridWay obtains higher productivity

- Reduces number of nodes and stages
- Mechanisms

Opportunistic migration

Fault tolerance

API's

- LCG-2: Relays on specific middleware
- DRMAA implementation: doesn't
GGF standard
Job sync, termination and suspension

Our two cents

Data from Information System should:

- be updated more frequently
- represent the real situation



32444245304354

Thank you for your attention

32444245304354

