

High Performance Computing Services in a WAN Environment

*Distributed Object Group Metacomputing
Architecture (DOGMA)*

Brigham Young University

Mark Clement, Bryan Morse, Quinn Snell

Trends

- Telecommunication providers are expanding to offer additional features
 - Internet connectivity
 - Web page hosting
 - Messaging
 - Pager information (stock quotes/sports)
- Compute services can be offered by an administrator capable of harvesting idle cycles

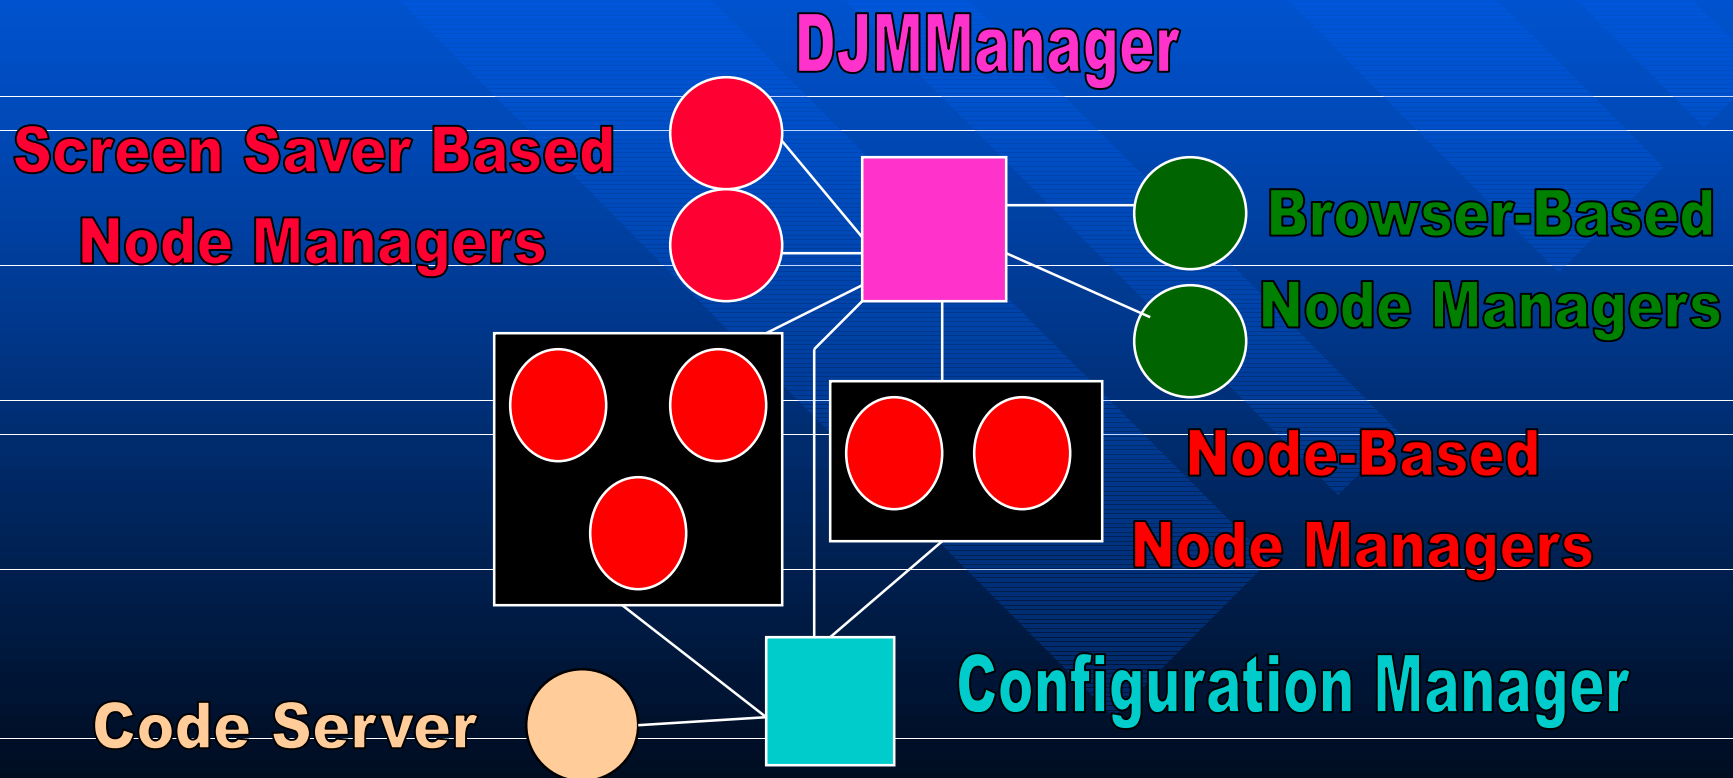
Overview

- *Distributed Object Group Metacomputing Architecture (DOGMA) provides connectivity and programming environment*
- *DOGMA is a parallel computing environment which integrates the power of heterogeneous clusters of workstations with that of idle nodes attached to the Internet.*
- *Java provides common executable.*
- *DOGMA supports two programming API's: MPI and Distributed Object Groups.*

Architecture

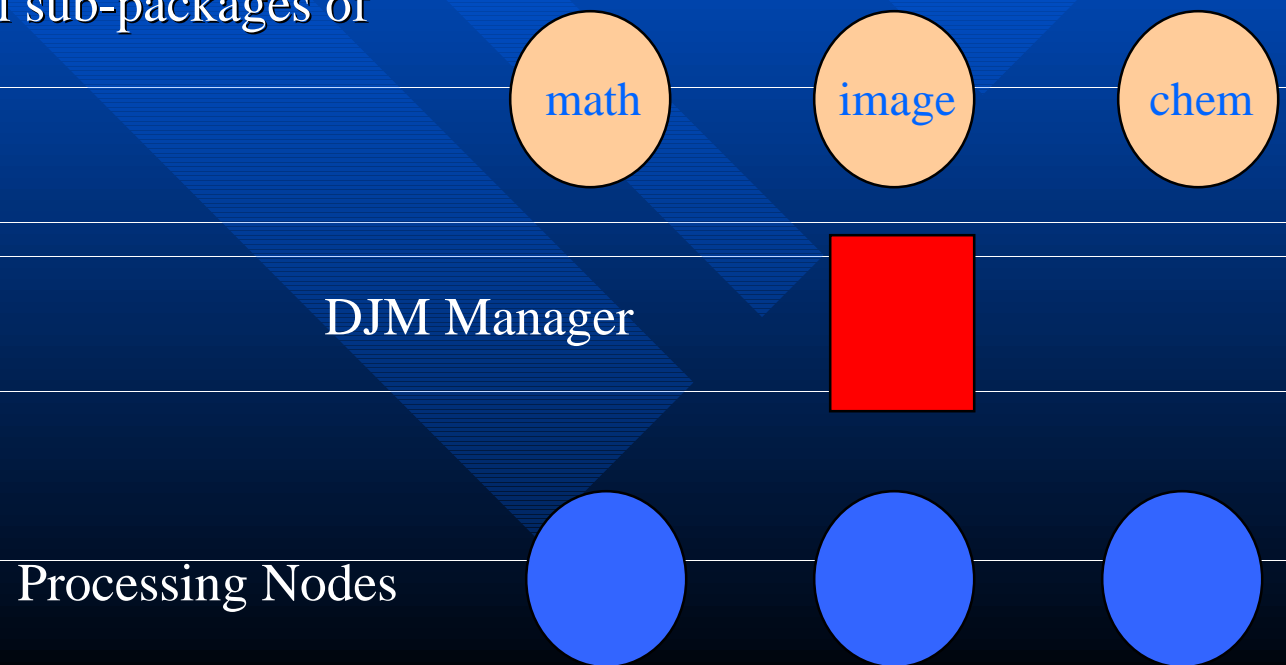
Distributed Object Groups	MPI
Distributed Java Machine	
Remote Method Invocation	Sockets

Architecture



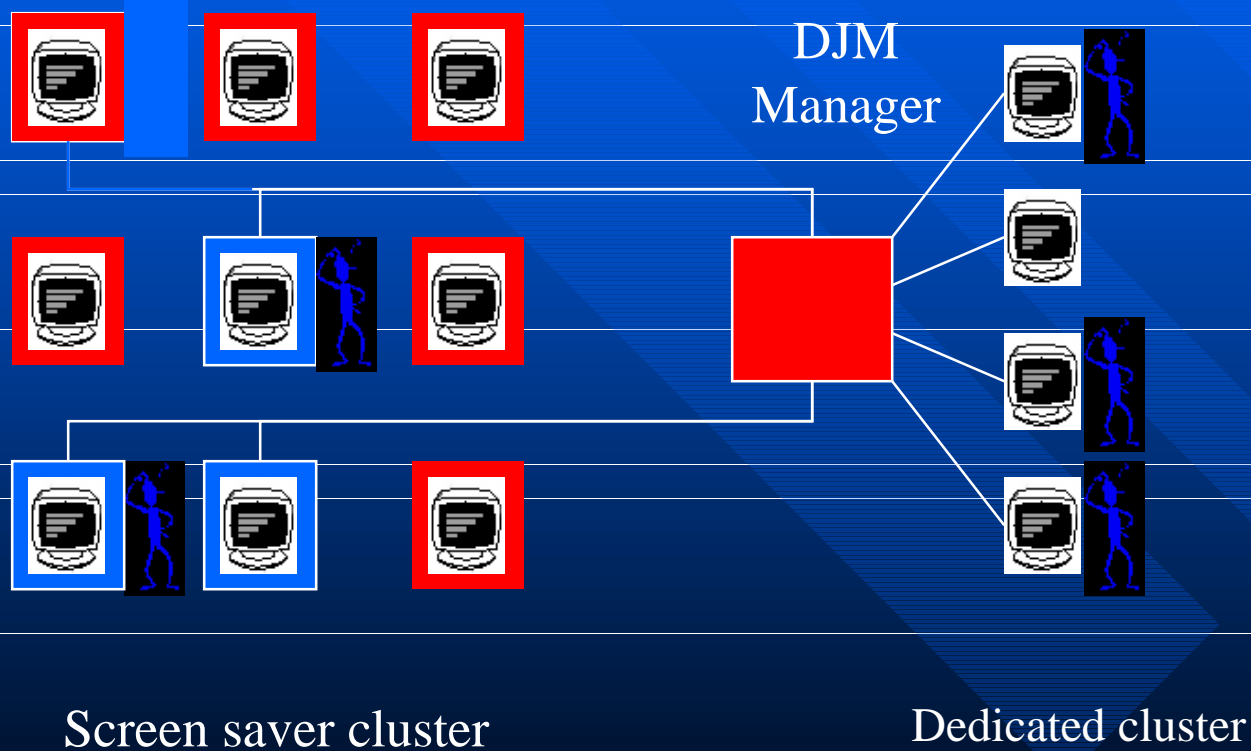
Code Servers

- Code servers allow application binary code to reside in multiple locations eliminating the need for a shared file system.
- Each code server serves up specific packages (and all sub-packages of that package).



Screen Saver

As “transitory” nodes activate (when the machine moves out of idle mode), “migratable” objects move to available nodes. Semi-dedicated nodes allow all objects to complete execution before stopping DOGMA.



Screen Saver

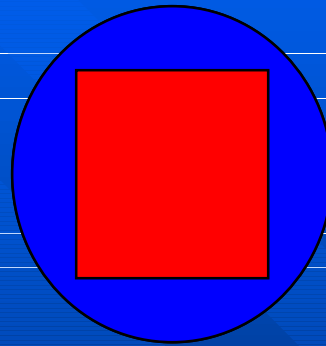
- Many computers spend a large amount of time idle.
- The DOGMA Screen provides a simple means of putting idle CPU cycles to work.
- Two modes of operation for screen saver termination:
- Transitory - Objects must migrate off of the node in a fixed amount of time after which DOGMA shuts down.
- Semi-Dedicated - Screen saver shuts down, but DOGMA runs in the background until all objects on the node have completed.

Distributed Object Groups


- Simplify parallel programming.
- Provide asynchronous group method invocation.
- Allow arguments to be automatically partitioned.
- Data is partitioned according to the power of nodes on which the group resides.
- Return values can be automatically assembled.
- Alternatively an array of results may be returned.
- Support for multiple partitioning schemes.

Distributed Object Groups

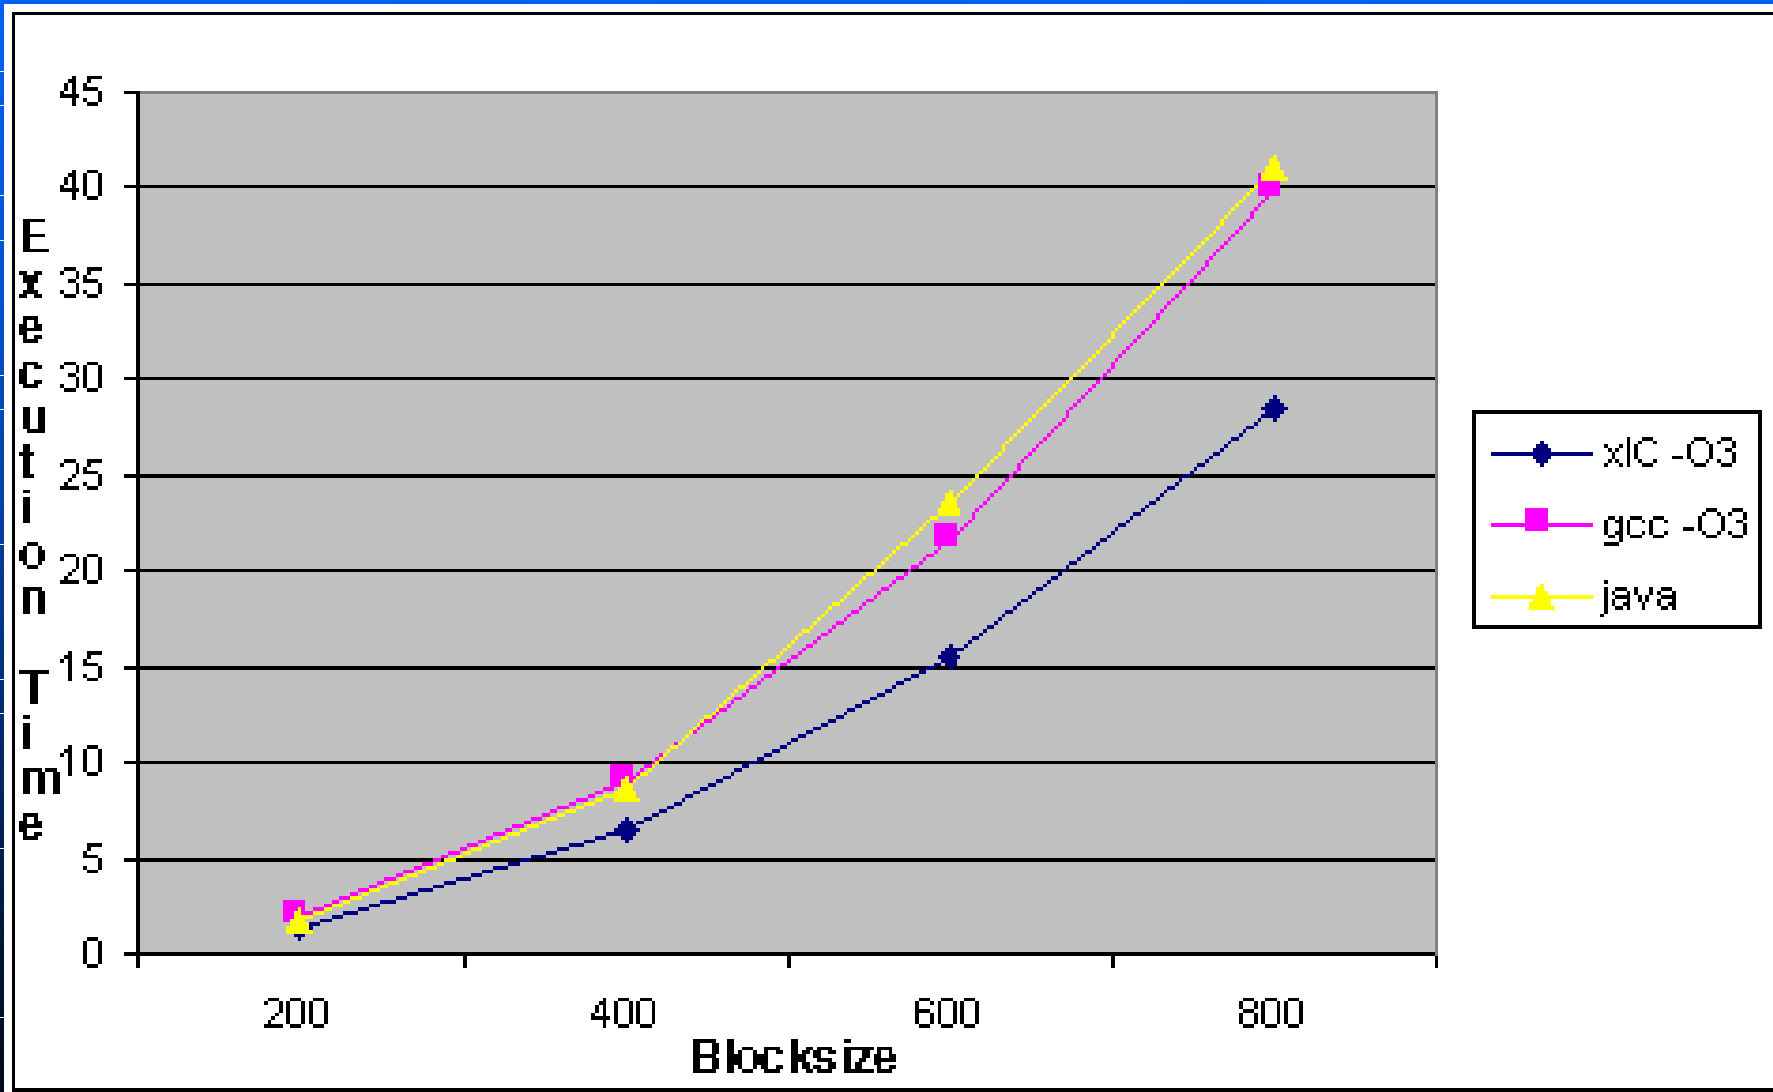
DOG Master



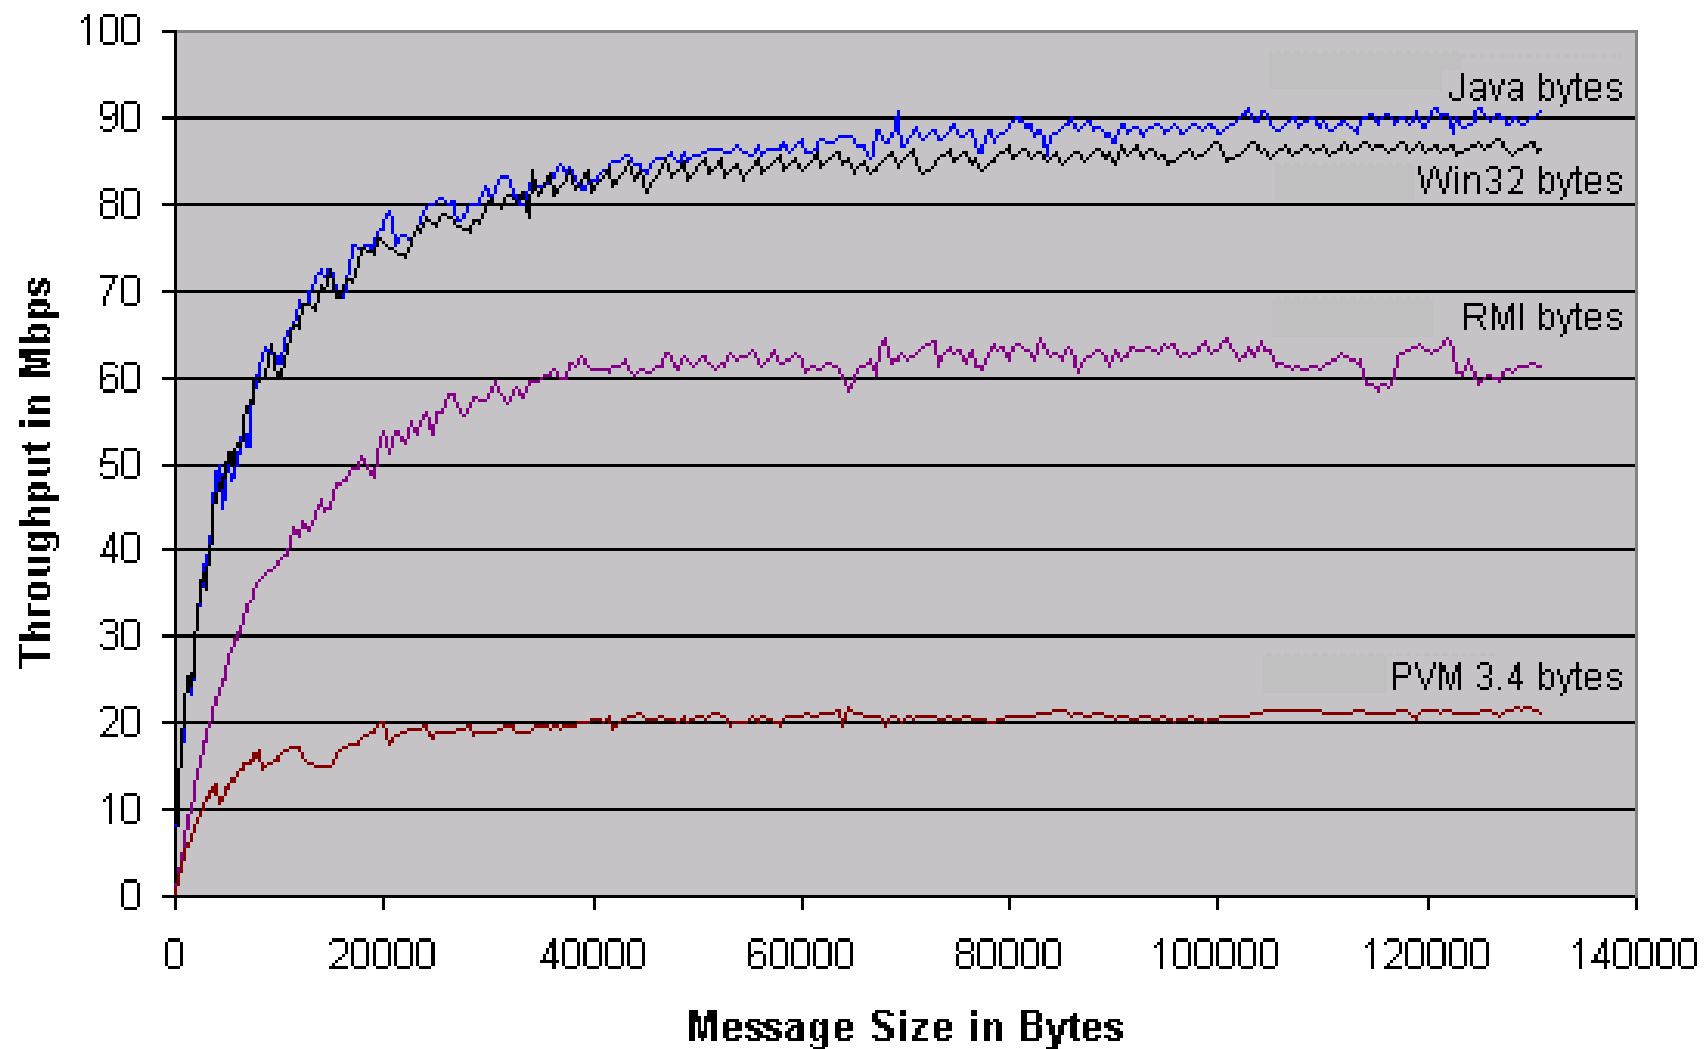
DOG
Elements




Computation




Communication



Jacobian



Gaussian Elimination



Applications

- Web page indexing
- Image Processing
- Complex speech recognition
- Wall Street predictions
- Virtual Reality
- Scientific applications

Cost Model

- Credit users who make idle cycles available
- Charge for telecommunications network time
- Optimize communication routes for network utilization

Conclusions

- DOGMA provides functionality for integrating heterogeneous clusters with anonymous idle nodes.
- Performance is very competitive with traditional parallel programming environments.
- Compute services can be offered through Java based systems
- <http://zodiac.cs.byu.edu/DOGMA>