



The SNOW Project

2001 -- 2004

UFSC/LISHA
UFRGS/GPPD
Altus Automation Ltda

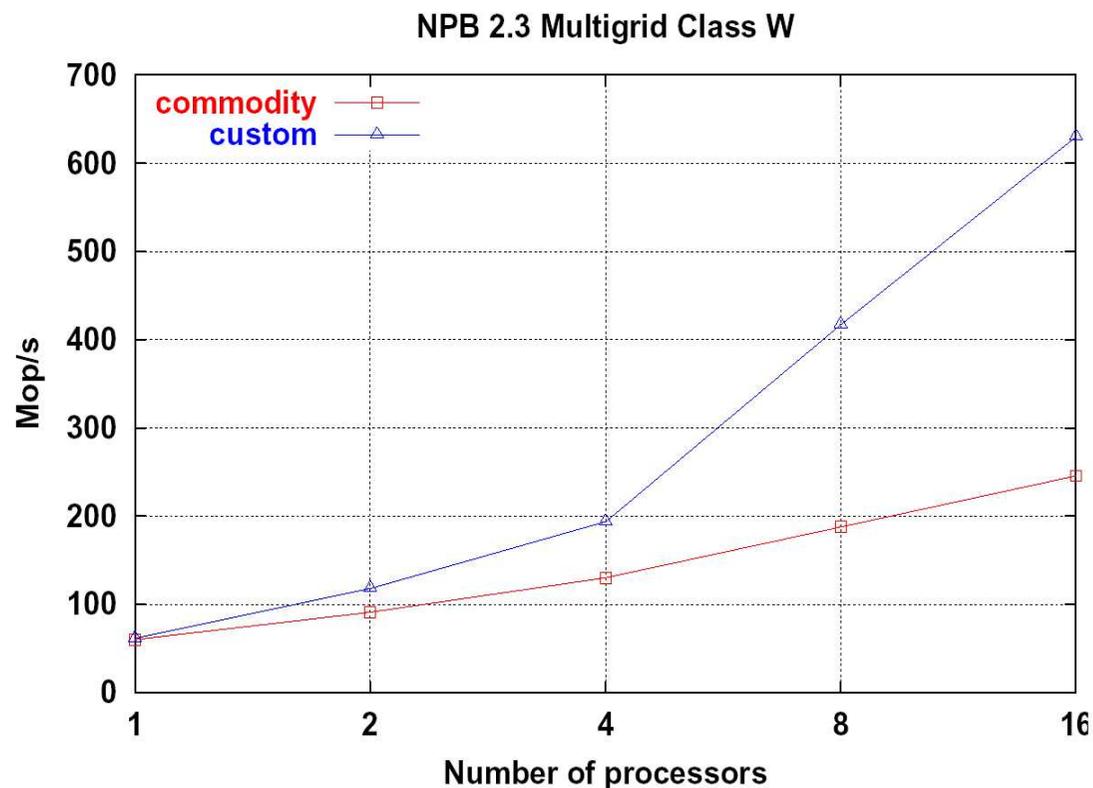
UEN/DSOS
Fraunhofer FIRST
Pure-Systems GmbH

snow@lisha.ufsc.br
<http://snow.lisha.ufsc.br/>



Motivation

- Parallel computing performance revisited in 2000
 - a cluster of commodity workstations
 - running on **commodity** and **custom** run-time systems





Commodity Hardware, Custom Run-time System

- Commodity X Custom Hardware
 - Convergence
- Commodity X Custom Run-time Systems
 - Commodity
 - multi-{user,tasks,...}, interactive, web-aware
 - more distributed than parallel
 - Custom
 - can deliver high performance and low latency
 - specially designed to support parallel computing
- Clusters do need **dedicated** RTS in order to be as efficient as traditional supercomputers

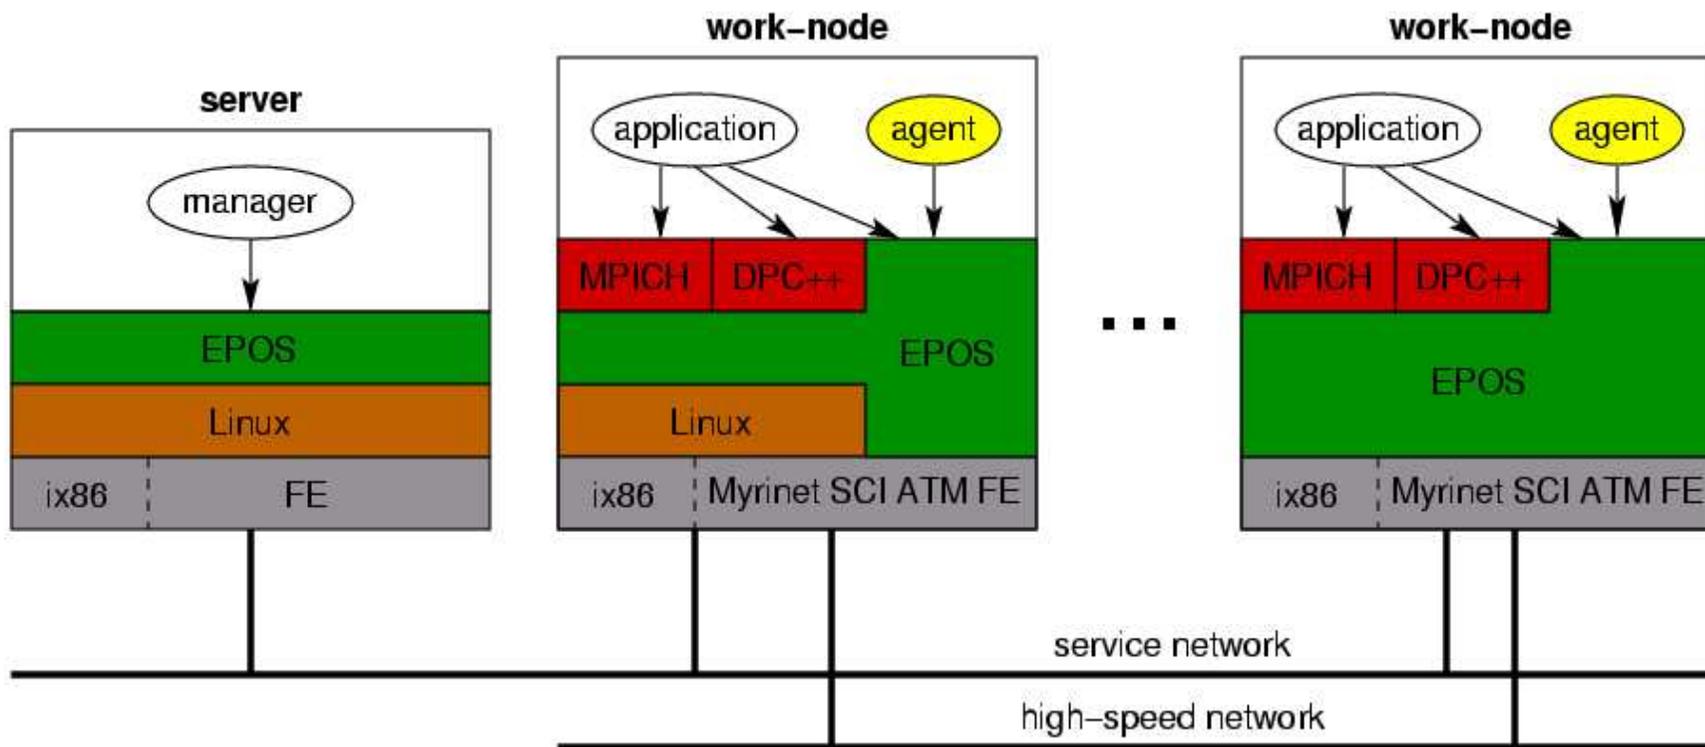


Project Goals

- Developing an application-oriented parallel programming environment for clusters of workstations
 - run-time support system
 - programming language
 - management tools
- Bringing cluster performance closer to traditional supercomputers
- Validated by selected parallel applications
 - computational biology
 - control of complex industrial processes



Overview of a SNOW Cluster





Original Project Partners

- Germany
 - Fraunhofer FIRST (SWT)
 - University of Magdeburg (DSOS)
 - Genias Software GmbH (industry)
- Brazil
 - Federal University of Santa Catarina (LISHA)
 - Federal University of Rio Grande do Sul (GPPD)
 - University of São Paulo (LSI)
 - Altus Automation Ltda (industry)



Project Partners as of Today

- Germany
 - University of Erlangen-Nürnberg (DSOS)
 - Fraunhofer FIRST (SWT)
 - Pure-Systems GmbH (industry)
- Brazil
 - Federal University of Santa Catarina (LISHA and BioInfo)
 - Federal University of Rio Grande do Sul (GPPD)
 - Altus Automation Ltda (industry)



University of Erlangen-Nürnberg - DSOS

- Tasks
 - Run-time support system
 - Reuse of PURE components
- Team
 - Prof. Dr. Wolfgang Schröder-Preikschat
 - Olaf Spinczyk
- Expertise
 - Family-based operating systems (PURE)
 - Software components



Fraunhofer FIRST

- Tasks
 - Run-time support system
 - Software refactoring tools
- Team
 - Friedrich Schön
- Expertise
 - Parallel operating systems (PEACE)
 - Software tools



Pure-Systems GmbH

- Tasks
 - Parallel embedded applications
 - Performance analysis and validation
- Team
 - Dr. Danilo Beuche
 - Holger Papajewski
- Expertise
 - Deeply embedded systems
 - Software families



UFSC-LISHA

■ Tasks

- Run-time support system
- Reuse of EPOS components
- MPI adaptation layer

■ Team

- Prof. Dr. Antônio Augusto Fröhlich
- André Gobbi Sanches
- Fernando Secco
- Fauze Valério Polpeta
- Tiago Stein D'Agostini

■ Expertise

- Application-orinetated operating systems (EPOS)
- Myrinet clusters



UFSC-BioInfo

- Tasks
 - Parallel computational biology applications
 - Web front-end
- Team
 - Prof. Dr. Edmundo Carlos Grisard
 - Charles Ivan Wust
 - Ingrid Thais Beltrame Botelho
- Expertise
 - Molecular biology (protozoology and genomics)
 - Bioinformatics



UFRGS-GPPD

- Tasks
 - DPC++ port and adaptation
 - Parallel run-time library
- Team
 - Prof. Dr. Philippe O. A. Navaux
 - Rafael Bohere Ávila
- Expertise
 - Parallel programming languages
 - Myrinet and SCI clusters

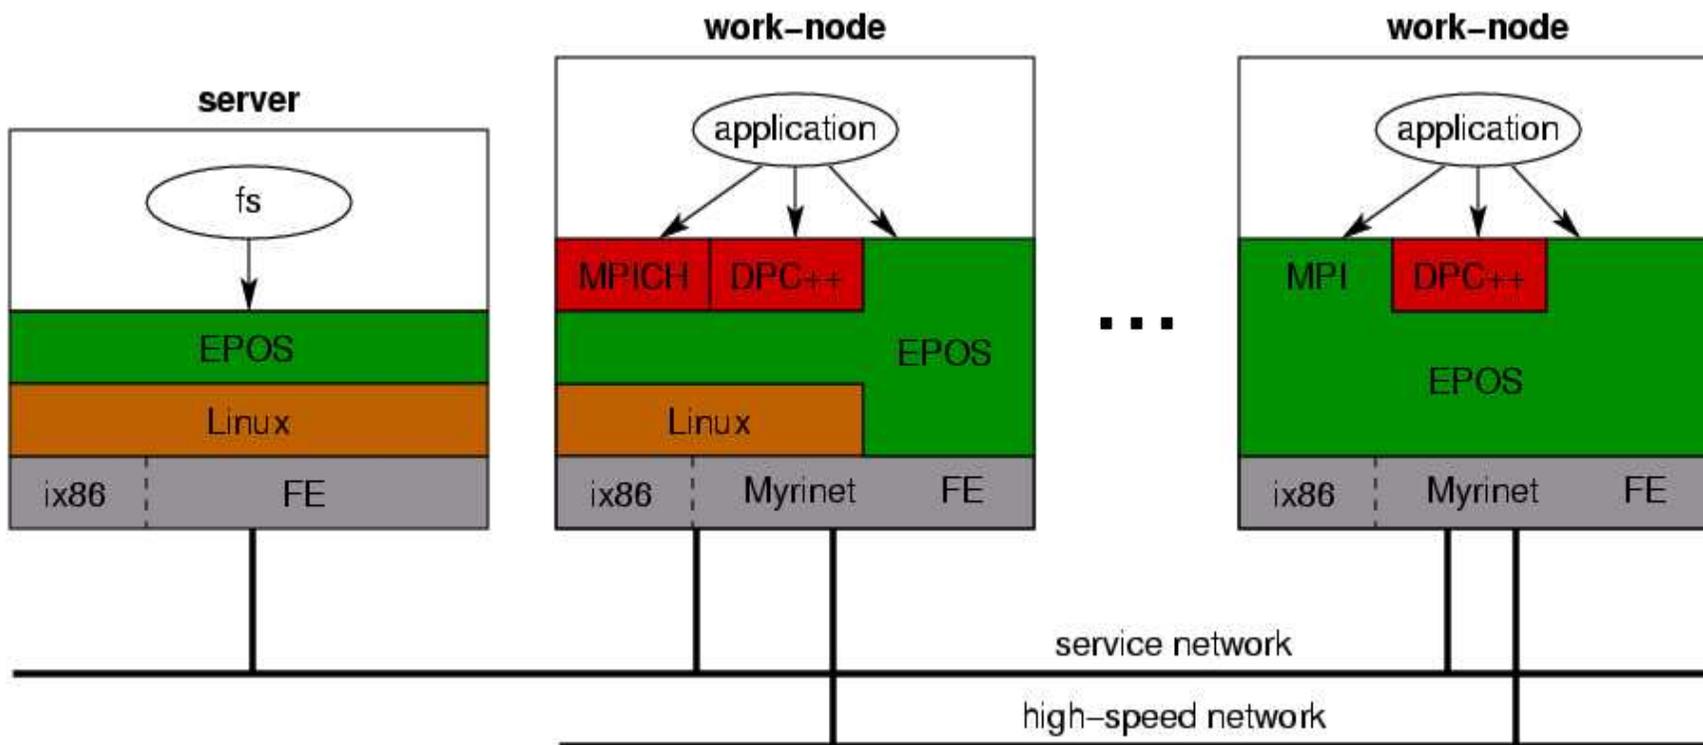


Altus Automation Ltda

- Tasks
 - Industrial control applications
 - Performance analysis and validation
- Team
 - Eng. Luiz Francisco Gerbase
- Expertise
 - Industry automation
 - Brazilian market leader



Overview of a SNOW Cluster (as of today)





Project Support

- International cooperation activities
 - German-Brazilian Cooperation Program on IT (DLR/CNPq)
 - Basically mobility
 - Fine at the German size (7 missions implemented)
 - Unstable at the Brazilian side (5 troublesome missions implemented)
- Local research activities
 - Independent research projects supported through different strategies
 - Very heterogeneous
 - Each partner devoted a different amount of resources to the project
 - Impact on results



Technical Remarks

- SNOW is operational as regards
 - Run-time support system (EPOS)
 - MPI on EPOS
 - DPC++
- SNOW is able to support high-performance applications
 - RST < 10 KB
 - Thread scheduling < 375 ns
 - Dynamic memory allocation < 100 ns
 - Round-trip latency < 5 us
- But a real application must confirm it!



Operational Remarks

- Partners with distinct expectancies about the project
 - Collection of individual efforts instead of a real cooperation
 - A real challenge that motivates all partners would be desirable
- German-Brazilian Cooperation Program on IT must run more smoothly
 - More investment and less bureaucracy from the Brazilian side
- Local project support options should be synchronized
- Do we want SNOW-II?



SNOW-II

- Two perspectives
 - Parallel computing
 - SNOW-I is a very good basis
 - Real applications in computational biology
 - New partners required
 - Embedded computing
 - SNOW-I dedicated RTS has much in common with a embedded RTS
 - No real application so far
 - Current partners could do it
- Why not both?
 - SNOW-II => GBCPIT
 - ESTK => DAAD