CBPF
Brazilian Center for
Research in Physics

SANTA FE INSTITUTE

NExtComp
Molecular Dynamics Application for Long-Range Interacting Systems on a Computational Grid Environment

Marcelo Portes de Albuquerque

Alexandre Maia de Almeida
Leonardo Haas Peçanha Lessa
Márcio Portes de Albuquerque
Constantino Tsallis
Luis Gregorio Moyano
Nilton Alves

IV WORKSHOP ON COMPUTATIONAL GRIDS AND APPLICATIONS
Curitiba – June 2006

Road Map

1. Introduction
2. The Physical Problem
3. NExtComp Parallelization Strategy
4. Performance Analysis
5. Conclusion and Future Works

2

1. Introduction

CBPF

Scientific research in theoretical and experimental physics

Many physics research groups use intense and complex computational methods for numerical simulations or data analysis

Several scientific collaboration over the world
- C. Tsallis in Santa Fe Institute, New Mexico, USA.

Physics applications complexity is increasing
- With more FLOPS, need better algorithms
- Better algorithms lead to complex structure
- Need to be adaptive and optimizations
- Ambitious projects lead to dynamic behavior and multiple components

Typical applications needs
- Enormous processing power, Fast networks, Huge amounts of data storage

Create infra-structure for scientific computing
- SSolar Project – Linux Cluster
- Grid Project - Team qualification for operation and application development
- PoP of 2 important Academic Network
  - Rio Metropolitan Network
  - National Research and Education Network (POP-RJ/LNCC)

Projeto SSolar

SSolar Project

Modular infra-structure for Scientific Computing at CBPF
http://mesonpi.cat.cbpf.br/ssolar

Hardware Outline

CBPF Linux Cluster	40 AMD AthlonMP 1800+ 2 GBytes RAM / Processor 1 Gbps Ethernet
Statistical Physics Linux Cluster	10 AMD AthlonMP 2800+ 2 GBytes RAM / Processor 100 Mbps Ethernet 10 Xeon 3.2 GHz 2 GBytes RAM / Processor 100 Mbps Ethernet
CBPF 64 bits Linux Cluster	08 Opteron 64bits 3.2 GHz 2 GBytes RAM / Processor 1 Gbps Ethernet
INTEGRIDADE Grid Project	04 Pentium 4 3.2 GHz 2 GBytes RAM / Processor 100 Mbps Ethernet

New projects in 2006
- Cosmology Linux Cluster
- High Energy Physics Linux Cluster

CBPF Grid Project - Time Line

2003

Configure two clusters in a Grid environment [WCGA 2003]
- OpenPBS, Globus and MPICH-G2 were configured in the CBPF Cluster
- Configuration of firewalls rules and TCP Ports
- Keeps the CBPF cluster policy

Scientific application tests using CBPF Cluster in a Grid environment
- Scientific tests using a MPI numerical integration program in C

2004

Associate to a local Grid initiative and exchange experiences
- Connect to GridRio: UFF, LNCC and PUC-Rio

Search a physical problem to develop an application for the Grid
- Group 1: Magnetic Materials [WCGA 2004]
- Group 2: Statistical Physics
- Group 3: Magnetism and Image Processing
- Group 4: High Energy Physics

Start the Molecular Dynamics Project using MPI [WCGA 2005]
- Collaboration of the CBPF and the Centre for Computational Science of the University College London

CBPF Grid Project - Time Line

CBPF join 2 “RNP Giga Projects” in Grid Development
- **Grid Sinergia:** Computational environment to run existing scientific applications
  - UFF, PUC-Rio, UNICAMP, LNCC, NCSA
- **INTEGRIDADE:** development of a physic applications to the Grid
  - LNCC, NCSA, PUC-Rio, UFES, UFF, UNICAMP, UFRGS

2005

Start the NExtComp Project
- Understanding of the physical and computational problem
- Start the Molecular Dynamics Project using Charm++
- Tests in SSolar – Statistical Physics AMD Linux Cluster

2006

NCSA proposal to the NExtComp Project was accepted
- “Molecular Dynamics for Long-Range Interacting Systems and its Possible Connection with Non-Extensive Mechanics Theory”. NCSA Proposal Number: PHY060015.

Performance Analysis of the NExtComp Program
- Tests in SSolar – Statistical Physics Xeon Linux Cluster
- Tests in NCSA Xeon Linux Cluster

2. The Physical Problem

- **Classical physics, and particularly statistical mechanics, studies systems formed by elements that interact through forces**
- **Usually, these forces have a dependency with the distance between any two elements**
  - Strong when the inter-particle distance is small
  - Weak when the elements are far apart
- **Depending on the intensity of these forces the interaction may be classified as short or long range interaction**
- **Examples of systems with long-range interactions**
  - Gravitational Systems, Coulombian Systems, Magnetic Systems, Fractures, etc.
- **Many properties of these systems still remain to be explained**
- **The main challenge regarding these systems**
  - Construction of a thermodynamics that may describe them correctly
  - Explain the similarities and differences with their short-range counterparts

Nonextensive Statistical Mechanics

This is one of the main points of interest in Nonextensive Statistical Mechanics

- Long Range Interacting Systems

Nonextensive Statistical Mechanics is a formalism formulated by Professor Tsallis in 1988, that generalizes the usual Boltzmann-Gibbs (BG) statistical mechanics

This formalism is based in a generalization of the conventional entropy

$$S = k \sum p \ln p$$

that includes a parameter *q*

$$S_q = k\,(1 - \sum_i p_i^q)\,(q-1)^{-1}$$

$$S_q \to S \quad \text{when} \quad q \to 1$$

Many publications are available in this area

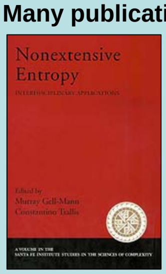

Nonextensive Entropy Interdisciplinary Applications

Edited by Murray Gell-Mann and Constantino Tsallis

Pub. Date: July 2004

Publisher: Oxford University Press

http://www.cbpf.br/GrupPesq/StatisticalPhys/TEMUCO.pdf

The Long-Range System

Hamiltonian Mean Field (HMF)

$$H = \sum_{i=1}^{N} \frac{v^2{}_i}{2} + \frac{1}{N}\sum_{i,j=1}^{N}[1-\cos(\theta_i-\theta_j)]$$

System formed by N planar classical rotators

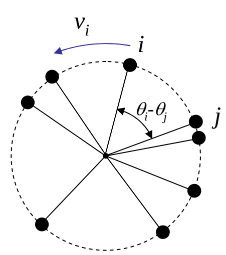

- The interaction force between rotators *i* and *j* is proportional to the angle difference
- The force in each rotator is influenced by every other rotator

$$F_i = \frac{1}{2N}\sum_{j=1}^{N}\sin(\theta_i-\theta_j)$$

- The sum includes every rotator (infinite interaction)

This simple model reflects many realistic characteristics of systems with long-range interactions

The Anomalous Behaviors

Given a specific energy, the value of any mean macroscopic observable such, as the temperature, may be predicted for equilibrium

But it is known that, for certain values of initial conditions, the system may be trapped in states where the mean microscopic quantities stay approximately constant for long periods of time with different values than those predicted by the BG theory

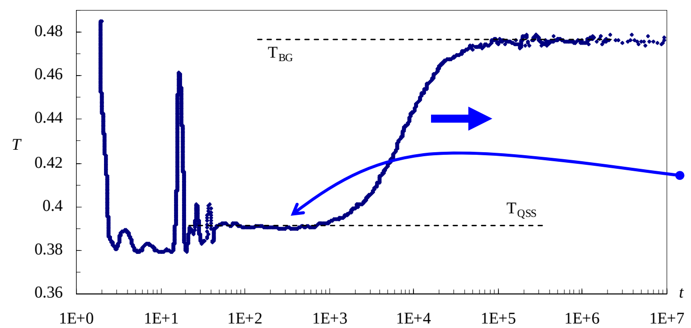

$$T = \frac{\langle 2K\rangle}{N} = \frac{1}{N}\sum_{j=1}^{N} v(i)^2$$

The duration of the quasistationary state grows with system size N

Simulate the evolution of this system for large values of *N* to verify the applicability of nonextensive statistical mechanics to this model

Numeric Simulation

Numeric simulation of the Hamiltonian equation

Differential equations defining the rotators movements

$$\frac{d}{dt}\theta_i = v_i \qquad 1 \le i \le N$$

$$\frac{d}{dt}v_i = m_y \cos(\theta_i) - m_x \sin(\theta_i) \quad \text{where} \quad \bar{m}_{x,y} = \frac{1}{N}\sum_{j=1}^{N}[\cos(\theta_j), \sin(\theta_j)]$$

The total energy of this system needs to be conserved

- **The differential equations must be discretized, and this may have as consequence a poor total energy conservation → incorrect dynamics**

Need a special algorithm to solve the differential equations conserving the total energy

- **Symplectic Yoshida Integrator**

  H. Yoshida (1990), *Phys. Lett. A* **150**, 262;

Averages from several realizations of the same simulation

- **Reduce the statistical fluctuations of the macroscopic observables**

3. NExtComp Parallelization Strategy

Sequential NExtComp MD

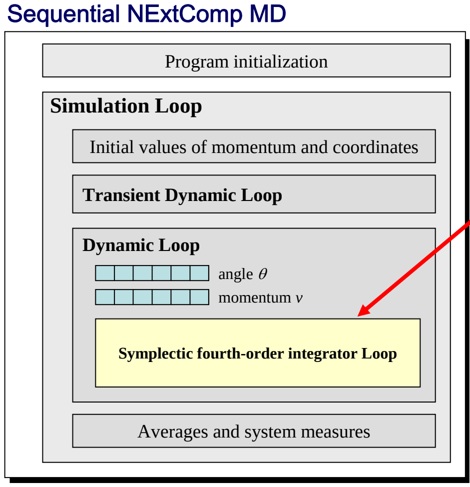

Figure: Simplified diagram

1. $\theta(t+1)$
2. Calculate Magnetization → $m_{x,y}(t+1)$
3. Calculate interactions Forces → $F(t+1)$
4. $v(t+1)$

To calculate Magnetization we need to access every rotator

Computational times might be unreachable

A possible solution for this model is parallel programming

Execution in high-performance machines

We implemented NExtComp using a parallel language

[1] C. Tsallis, "***Entropy, nonlinear dynamics, complexity and all that***", 17th Symposium on Comp. Arch. and High Perf. Computing (SBAC-PAD 2005)

Suitable Features of Charm++

Parallel programming language based on C++
- Objects communicate with each other via messages

Is target to tightly coupled and high-performance parallel machines

Is portable to a wide variety of parallel machines

Allow the scalability up to thousands of processors
- NAMD project for Biomolecular Simulations §1

Use of object array (chare arrays) distributed over all processors
- Optimized communication for collective operations → reductions

When an object is waiting for some incoming data other ready objects are free to execute

Performance prediction on large machines
- Performance tuning without continuous access to a large machine
- Develop a parallel application for a non-existent machine

Design the NExtComp-MD using object oriented techniques in Charm++

[1] Kale L. V. et al. , "***NAMD: Biomolecular Simulation on Thousands of Processors***", Proceedings of Supercomputing (2002).

NExtComp-MD-π Overview

Rotator class instantiated as a chare array

Yoshida Symplectic Integrator in a parallel algorithm
- Integration in 2 Stages: 1) computes the angle θ and makes a reduction in $m_{x,y}$;
  2) compute Forces and Momentums

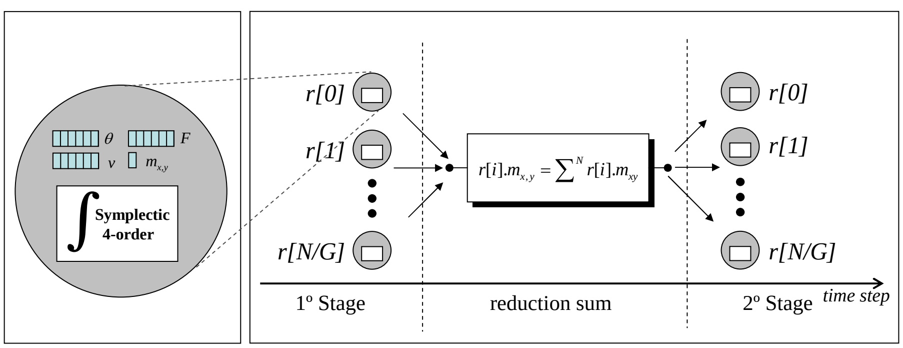

Synchronization processes in order to compute $m_{x,y}$

NExtComp-MD-π is a tightly coupled parallel program
- All rotators objects needs to exchange data at regular intervals

4. Performance Analysis

The NExtComp MD-π should

- Solve the physical problem faster than its sequential version
- Deal with a larger physical system that could not be attained before

We conducted 3 sets of experiments:

1. Measurement of the total execution time and thus the speedup of the system
2. Analysis of the object execution time and the performance distribution for all object entry points
3. Monitor the processors activities inspecting the program time line
   - Details of CPU utilizations using Charm++'s *Projections* visualization tool

Measures were carried out

- CBPF SSolar Statistical Physics Linux Cluster (#P = 10 Processors)
- NCSA Xeon Linux Cluster (#P = 20 Processors)
- With the size of the system $\rightarrow$ N=[10^2,10^3, 10^4, 10^5, 10^6 and 10^7] rotators
- With fewer time-steps

Goal

- Validates the NExtComp-MD-π version
- Optimize the program for long time execution

15

First Experiment

Runs NExtComp-π and measure the execution time $\rightarrow$ Speedup

The amount of computation per object = N/#P

Speedup				
N	#P			
	1	5	10	20
10^2	-	0,08	0,06	0,07
10^3	-	0,46	0,35	0,36
10^4	-	3,05	3,07	3,54
10^5	-	**4,53**	**8,51**	13,84
10^6	-	**4,88**	**9,57**	**17,85**
10^7	-	**4,96**	**9,85**	**19,35**

Table: Speedup as function of system size N and #P

The speedup is better for large systems (N $\gg$1)

N=10^8 $\rightarrow$ needs to change de initialization procedures

Linear speedup: expected a downward trends when the:

- System size N and #P increases
- Communication among processors became a significant factor

16

Second Experiment

Analysis of objects entry points (EP's) execution time

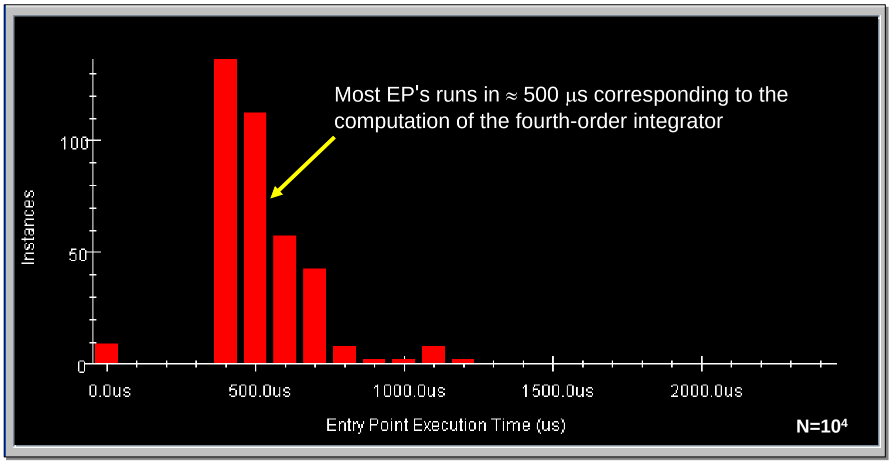

Figure: Entry Points Execution time histogram for NExtComp-MD-π

Charm ++ Projections - performance and visualization analysis tool

Third Experiment

Charm++ Timeline Tool → detailed view of the application overall processors

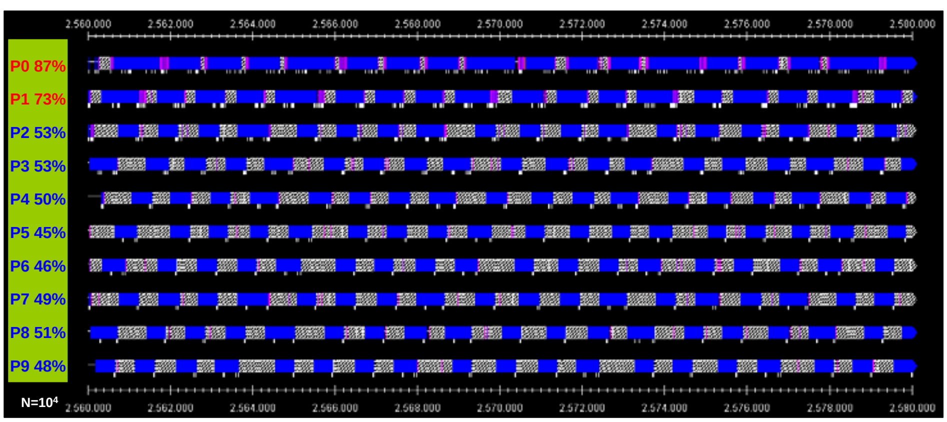

Figure: Tracedata window of Δt=20ms for NExtComp-MD-π running on NCSA Xeon Linux Cluster

Parallel tasks have sections of idle time

- Frequent periods of communication because of the fourth-order integrator

Decreasing the Idle Time

Using independent tasks

- Estimation of physical parameters needs several realizations of the simulation

Charm++ enhance task distribution and reduce the processors idle time

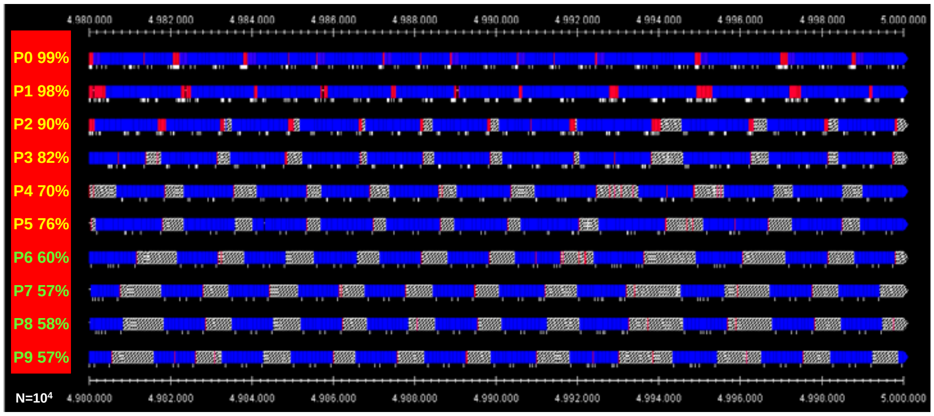

Figure: Tracedata window of Δt=20ms for 5 instances of the NExtComp-MD-π running on NCSA Xeon Linux Cluster

19

5. Conclusion and Future Works

Nonextensive (NExt) statistical mechanics (SM) is actually a field of intense activities in physics and are concerned with long-range interactions systems

We are modeling this kind of system to verify the applicability of NExt-SM

- Investigations need many identical elements and several time steps

Computational time can be unreachable

- This is a kind of problem for high performance parallel computing
- And for Grid deployment **(?)**

We created the NExtComp Project

- The development of a parallel algorithm for molecular dynamics simulation

NExtComp application needs scalability to machines with thousands of processors

- We decided to use Charm++ runtime system
- The explored techniques show good parallel scalability allowing simulations of very large physical systems ($N=10^7$).

20

Conclusion and Future Works

We implemented experiments for performance analysis

- **Achieved a linear speedup in the range of #P=1 to #P=20.**
  - This is a good indicative of the performance
  - We expect a downward trends as N and #P increases
- **Analysis of objects entry points execution time**
  - Compute the symplectic fourth-order integrator in parallel is the program hotspot
  - Re-design the algorithm taking advantages of symmetries of the physical problem
- **Time Analyze – examine CPU utilization**
  - CPU Idle-Time: diminish introducing independents tasks running simulations in parallel
  - We are studying Charm++ synchronization techniques and load balancing strategies

Further Experiments

- Verify communication latencies → in CBPF and NCSA Cluster

Future Work

- Test NExtComp-MD-π in a Grid environment - INTEGRIDADE and SINERGIA Project
- Run tightly-coupled applications on Grids need an algorithm-level modifications
- Grid topologies with near neighbor connections → avoiding wire-length delays
- Test the NextComp-MD-π using the Myrinet Network in the NCSA Cluster

21

Acknowledgments

Grid is a collaborative work and its important to thanks our partners…

SFI	Santa Fe Institute	USA
NCSA	National Center for Supercomputing Applications	USA
INTEGRIDADE Project		
LNCC	National Laboratory for Scientific Computing	Petrópolis
SINERGIA Project		
UFF	Fluminense Federal University	Niterói

22

CBPF
Brazilian Center for
Research in Physics

SANTA FE
INSTITUTE

NExtComp

Molecular Dynamics Application for Long-Range Interacting Systems on a Computational Grid Environment

Marcelo Portes de Albuquerque

Alexandre Maia de Almeida
Leonardo Haas Peçanha Lessa
Márcio Portes de Albuquerque
Constantino Tsallis
Luis Gregorio Moyano
Nilton Alves

IV WORKSHOP ON COMPUTATIONAL GRIDS AND APPLICATIONS
Curitiba – June 2006