The Sciparc Project

Orsay, April 18th, 2005

R. Tripiccione

Physics Dept., University of Ferrara and INFN

Talk Overview

1. Sciparc: the official wisdom
2. Sciparc: the background
3. Sciparc scientific goals
4. Sciparc political objectives
5. Sciparc structure and workpackages
6. Suggestions for hardware related activities

Sciparc: the official view (I)

Proposal submitted to the EU

IST call 4 within the 6^{th} Framework program

Advanced Computer Architectures (ACA)

Project goals---->

Next-to-next generation computer architecture for

high-performance scientific & engineering applications

in the 10+ years time-frame

Sciparc: the official view (II)

Project participants

- DESY/NIC
- INFN
- CNRS
- Eurotech
- University of Padua
- Julich Computer Center
- University of Regensburg
- EPCC/Edinburgh
- INRIA
- IBM - Zurich

Sciparc Background (I)

The “European” community of Lattice QCD machine developers and users -> INFN, DESY, CNRS, Eurotech, INRIA, Edinburgh

+

Theoretical physics with large computational needs

---> INFN, Julich, Regensburg, INRIA, IBM- Zurich

+

Computer scientists with interest in “special” architectures

---> Padova

Lattice QCD background

Lattice QCD number crunchers have been developed in the last 20 years on both sides of the Atlantic Ocean, providing a large fraction of all CPU-cycles used in this field

Current generation (machines being installed right now):

QCDOP (Columbia + IBM + Edinburgh) --> 10 + 10 Tflops

apeNEXT (Infn + Desy + Orsay) --> 8.5 + 2.0 + 3.5 + ?? Tflops

Lattice QCD background

L

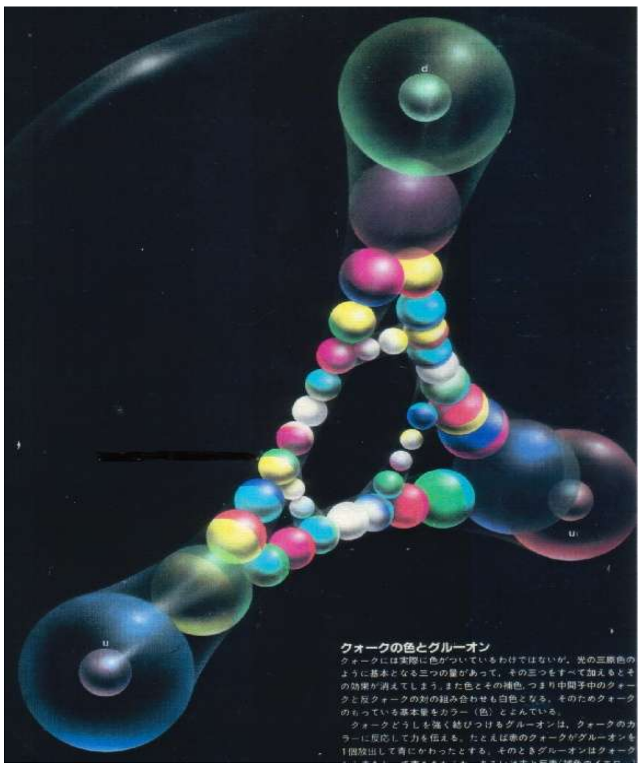

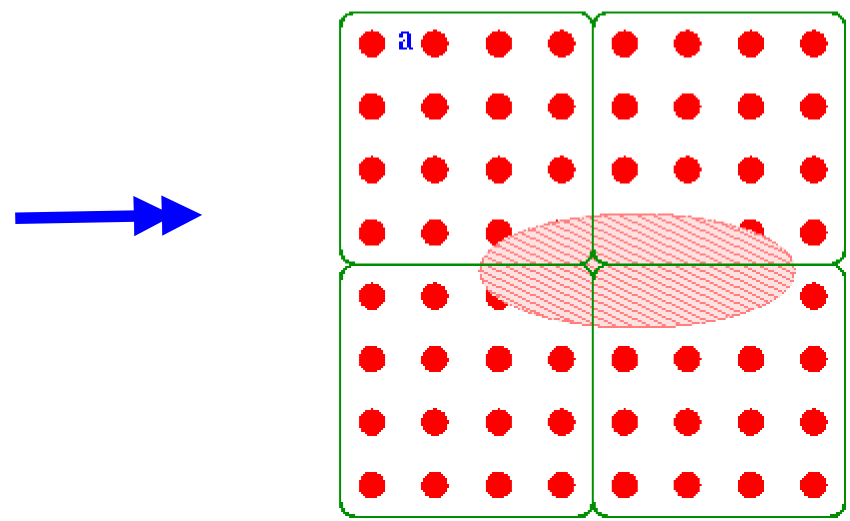

Lattice QCD background

Lattice-QCD is not a model and therefore physical quantities can be computes from first principles without arbitrary assumptions. It provides a method for predicting physical quantities (decay constants, form factors) within a unique coherent theoretical framework …….

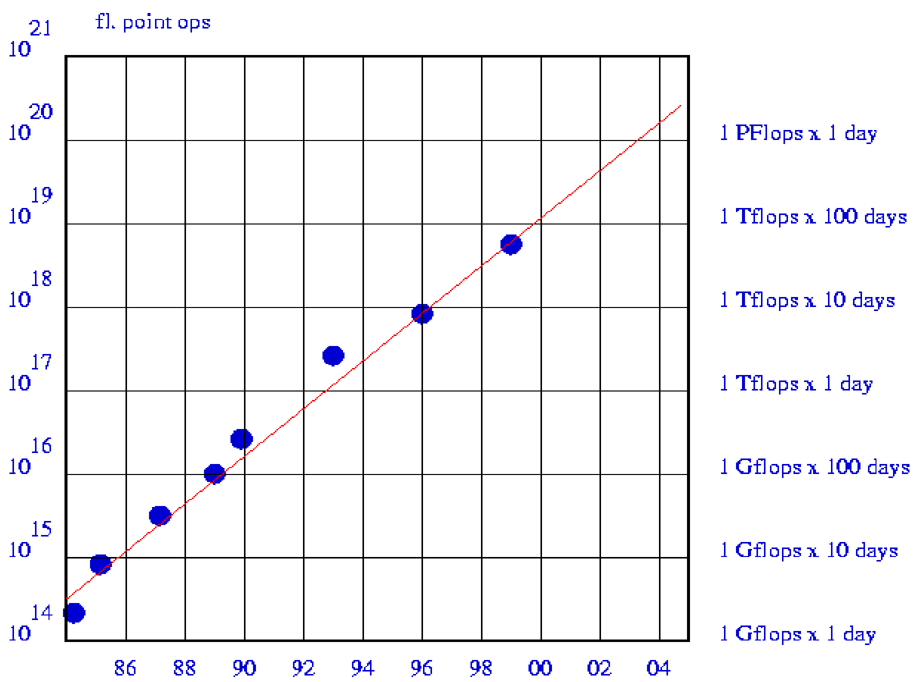

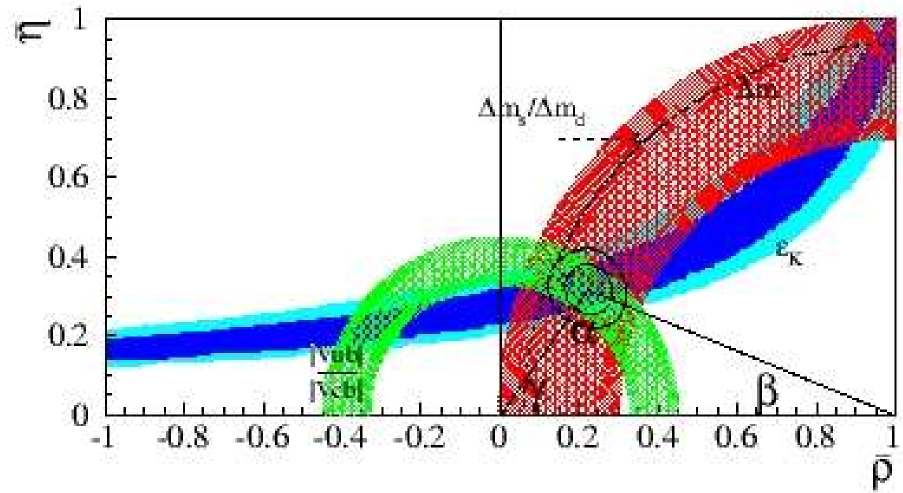

Lattice QCD machines

L

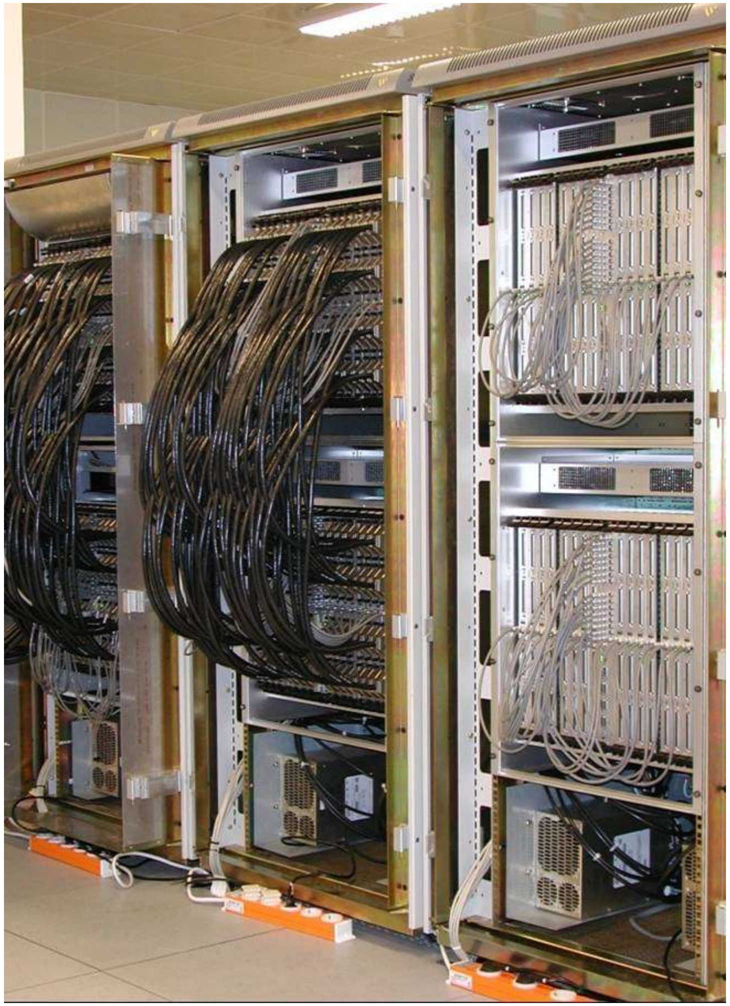

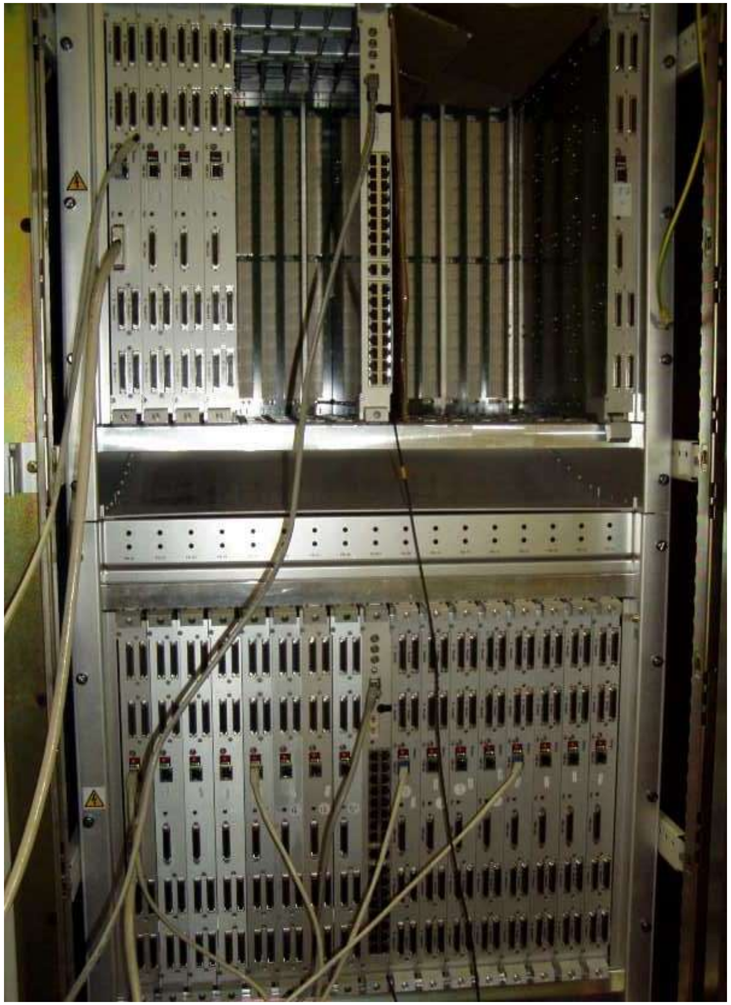

Sciparc scientific goals

Is it possible/useful/efficient to define a high-performance computer architecture for scientific computing?

Look at the problem from the application point of view:

- *Almost explicit parallelism*
- *Regular communication pattern*
- *Regular addressing structure*
- *Favourable ratio of computing/communication*
- *Long/regular kernel loops*

Sciparc scientific goals

Beautiful match with present or forecast technology development

Almost explicit parallelism	*Network-on-chip of processors/SIMD clusters*
Regular communication pattern	*Network-on-chip of processors*
Regular addressing structure	*prefetch queues*
Favourable ratio of computing/stirage	*embedded DRAM*
Long/regular kernel loops	*no speculative techniques. Save on static power*

Sciparc political wisdom

Is it possible to create an enlarged scientific community with serious computational problems, able to:

- define its ambitious computer requirements
- suggest architectural solutions
- define a viable project to:
- design/develop/test/commission/use top-of-the range computer systems giving O(10 – 100) better performance than available on the main road
- and large enough to be able to survive

Sciparc structure

Sciparc is structured in 3 Workpackages:

WP2 Advanced Hardware Architecture

WP3 System Software Architecture

WP4 Computing Requirements of Scientific Applications

Sciparc structure/WP4

WP4 Computing Requirements of Scientific Applications

Characterize, parametrize, simulate and benchmark the computing requirements of a “large enough” set of scientific applications:

LQCD

Quantitative Biology

Material sciences/Computational Chemistry

Fluid-dynamics

(Astrophysics)

WP4 coordinator: H. Simma

Sciparc structure/WP3

WP3 System Software Architecture

Develop the tools that make scientific computing efficient on large massively parallel architectures

- Expose parallelism through language directives
- Application libraries
- Assembly generators
- Compile-time optimization & scheduling
- Minimal structure of operating systems

WP3 coordinator: F. Bodin

Sciparc structure/WP2

WP2 Advanced hardware architecture

Design, parametrize, characterize and simulate a realistic architecture for massively parallel scientific computing

- Task 1: Global architectural design
- Task 2: Node structure
- Task 3: Interconnection structure
- Task 4: Better have a few sanity checks

Sciparc structure/WP2

In short, Tasks 1 thru 3 are an attempt to design a visionary but reasonable scientific architecture able to live on a 5+ years frame

Task 4 is a badly needed set of sanity checks:

- Are foreseen technology trends going in the right direction ?
- Are our architectural solutions viable ?
- Can we leverage on commercially available building blocks ?

WP2/Task 4

WP2/Task4 is the place for hardware developments:

Relevant questions:

- Are we able to VHDL-model/simulate our most critical blocks ?
- Do we need to validate blocks/structures/protocols with FPGA techniques ?
- Are "commercial" processors a viable solution (Clearspeed, IBM/Cell, Efficeon ....) ?
- Are we able to connect efficiently these processors ?

What next after apeNEXT? (I, theory)

Find the size of the optimal lattice size contained in each processor

Expected performance: 200 Gflops/processor

Expected cost: see plot (200 X better than today)