BNL-52678
Formal Report

Proceedings of RIKEN BNL Research Center Workshop Volume 46

Large-Scale Computations in Nuclear Physics using the QCDOC

September 26 - 28, 2002

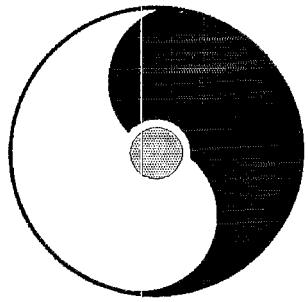

Organizers:

Y. Aoki, A. Baltz, M. Creutz, M. Gyulassy, and S. Ohta

RIKEN BNL Research Center

Building 510A, Brookhaven National Laboratory, Upton, NY 11973-5000, USA

DISCLAIMER

Available electronically at-

`http://www.doe.gov/bridge`

Available to U.S. Department of Energy and its contractors in paper from-

U.S. Department of Energy
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831
(423) 576-8401

Available to the public from-

U.S. Department of Commerce
National Technical Information Service
5285 Port Royal Road
Springfield, VA 22131
(703) 487-4650

Printed on recycled paper

Preface to the Series

The RIKEN BNL Research Center (RBRC) was established in April 1997 at Brookhaven National Laboratory. It is funded by the "Rikagaku Kenkyusho" (RIKEN, The Institute of Physical and Chemical Research) of Japan. The Center is dedicated to the study of strong interactions, including spin physics, lattice QCD, and RHIC physics through the nurturing of a new generation of young physicists.

During the first year, the Center had only a Theory Group. In the second year, an Experimental Group was also established at the Center. At present, there are seven Fellows and seven Research Associates in these two groups. During the third year, we started a new Tenure Track Strong Interaction Theory RHIC Physics Fellow Program, with six positions in the first academic year, 1999-2000. This program had increased to include ten theorists and one experimentalist in academic year, 2001-2002. With recent graduations, the program presently has eight theorists and two experimentalists. Beginning last year a new RIKEN Spin Program (RSP) category was implemented at RBRC, presently comprising four RSP Researchers and five RSP Research Associates. In addition, RBRC has four RBRC Young Researchers.

The Center also has an active workshop program on strong interaction physics with each workshop focused on a specific physics problem. Each workshop speaker is encouraged to select a few of the most important transparencies from his or her presentation, accompanied by a page of explanation. This material is collected at the end of the workshop by the organizer to form proceedings, which can therefore be available within a short time. To date there are forty-eight proceeding volumes available.

The construction of a 0.6 teraflops parallel processor, dedicated to lattice QCD, begun at the Center on February 19, 1998, was completed on August 28, 1998. A 10 teraflops QCDOC computer in under development and expected to be completed in JFY 2003.

T. D. Lee
November 22, 2002

***Work performed under the auspices of U.S.D.O.E. Contract No. DE-AC02-98CH10886.**

CONTENTS

WORKSHOP INTRODUCTION

Large-Scale Computations in Nuclear Physics Using the QCDOC

The massively parallel computer QCDOC (QCD On a Chip) of the RIKEN BNL Research Center (RBRC) will provide ten-teraflop peak performance for lattice gauge calculations. Lattice groups from both Columbia University and RBRC, along with assistance from IBM, jointly handled the design of the QCDOC. RIKEN has provided $5 million in funding to complete the machine in 2003. Some fraction of this computer (perhaps as much as 10%) might be made available for large-scale computations in areas of theoretical nuclear physics other than lattice gauge theory. The purpose of this workshop was to investigate the feasibility and possibility of using a supercomputer such as the QCDOC for lattice, general nuclear theory, and other calculations.

The lattice applications to nuclear physics that can be investigated with the QCDOC are varied: for example, the light hadron spectrum, finite temperature QCD, and kaon ($\Delta I = \frac{1}{2}$ and CP violation), and nucleon (the structure of the proton) matrix elements, to name a few. There are also other topics in theoretical nuclear physics that are currently limited by computer resources. Among these are *ab initio* calculations of nuclear structure for light nuclei (e.g. up to ~A = 8 nuclei), nuclear shell model calculations, nuclear hydrodynamics, heavy ion cascade and other transport calculations for RHIC, and nuclear astrophysics topics such as exploding supernovae.

The workshop ran from early Thursday morning to Saturday noon, September 26-28. Both physics and technical questions were discussed. On the technical side, Robert Mawhinney gave an introduction to the QCDOC on the first day and, along with Peter Boyle, was available for questions and feedback. Michael Creutz demonstrated the current QCDSP programming environment.

The physics topics were quite varied, ranging from simulations of stellar collapse by Douglas Swesty to detailed shell model calculations by David Dean, Takaharu Otsuka, and Noritaka Shimizu. Going outside traditional nuclear physics, James Davenport discussed molecular dynamics simulations and Shailesh Chandrasekharan presented a class of algorithms for simulating a wide variety of fermionic problems. Four speakers addressed various aspects of theory and computational modeling for relativistic heavy ion reactions at RHIC. Scott Pratt and Steffen Bass gave general overviews of how qualitatively different types of physical processes evolve temporally in heavy ion reactions. Denes Molnar concentrated on the application of hydrodynamics, and Alex Krasnitz on a classical Yang-Mills field theory for the initial phase. We were pleasantly surprised by the excellence of the talks and the substantial interest from all parties. The diversity of the audience forced the speakers to give their talks at an understandable level, which was highly appreciated. One particular bonus of the discussions could be the application of highly developed three-dimensional astrophysics hydrodynamics codes to heavy ion reactions.

It is too early to determine to what extent the QCDOC is a solution to problems other than lattice QCD; we all need some time to mull this over. It appears that for much of the simulation of hydrodynamical flow and cascades, conventional computer clusters are likely adequate, but that classical Yang-Mills theory might be a natural match to lattice techniques. The models of stellar collapse use techniques quite similar to those of lattice gauge theory, require substantial amounts of computer time, and again might be a natural match to the QCDOC. The shell model calculations appear to be outgrowing computer clusters. These calculations may be a candidate for future trial running on upcoming prototypes of the QCDOC.

Yasumichi Aoki
Anthony J. Baltz
Michael Creutz
Miklos Gyulassy
Shigemi Ohta

An Introduction to QCDOC

Robert D. Mawhinney
Columbia University

The QCDOC computer is being designed by a team composed of researchers from Columbia University, IBM, the RBRC and the UKQCD collaboration. QCDOC stands for Quantum Chromodynamics On a Chip and the machine's architecture will employ a single custom ASIC (Application Specific Integrated Circuit) manufactured by IBM. This QCDOC ASIC will contain the processor (a 440 Power PC with a 64-bit floating point unit), 4 MBytes of embedded memory, a 100 Mbit ethernet connection and 24, 500 MHz, bit-serial communication links for internode communication. Each processing node of QCDOC will consist of an ASIC plus external DDR SDRAM. It is expected that at least 128 MBytes of DDR will be available per node.

The most distinguishing architectural feature of QCDOC is the low-latency of the 24 internode communication links. These will be connected so that QCDOC is a six-dimensional hypercubic processor mesh, with simultaneous, bi-directional communication between nearest neighbors. The time interval between a user software call to start a data transfer and the arrival of the first data at a neighboring node is about 200 ns. This low latency allows efficient operation of QCDOC when many small data transfers are needed to keep the local floating point units running at high performance. Such a situation occurs when many thousands of processors are applied to a problem of a fixed size, in an effort to produce a result in a timely manner.

QCDOC also has special hardware features, called store and forward capability, to achieve fast global sums. These appear, for example, in taking the inner product of a vector over the entire machine. Having fast global sums is vital if thousands of nodes are being used on a problem.

Arbitrary communication can be done on QCDOC, but the hardware network supports nearest neighbor and global operations most efficiently. Other communication patterns are supported by multiple hops in the nearest neighbor hardware network. Many scientific applications involve extensive nearest neighbor communications and these should run well on QCDOC. Exploring such other applications is one of the goals of this workshop.

The members of the QCDOC design team are:

- **Columbia** Norman Christ, Saul Cohen, Calin Cristian, Zhihua Dong, Changhoan Kim, Xiaodong Liao, Hueywen Lin, Guofeng Liu, Robert Mawhinney, Azusa Yamaguchi
- **IBM** Dong Chen, Alan Gara
- **RBRC** Shigemi Ohta (KEK), Tilo Wettig (Yale)
- **UKQCD** Peter Boyle, Balint Joo
- **SciDAC (BNL/CU)** Chulwoo Jung, Kostya Petrov
- **BNL (SciDAC/DOE)** Dave Stampf, Rob Bennett

QCDOC Overview

- Uses IBM's System-On-a-Chip Technology
- Industry standard, 64-bit, 1 Gflop Power PC processor with L1 instruction and data caches.
- 4 MBytes of embedded DRAM on chip
- Up to 0.5 Gbytes of DDR SDRAM per node
- 100 Mbit Ethernet connection to each node
- High-bandwidth, low-latency 6 dimensional nearest neighbor communications network.

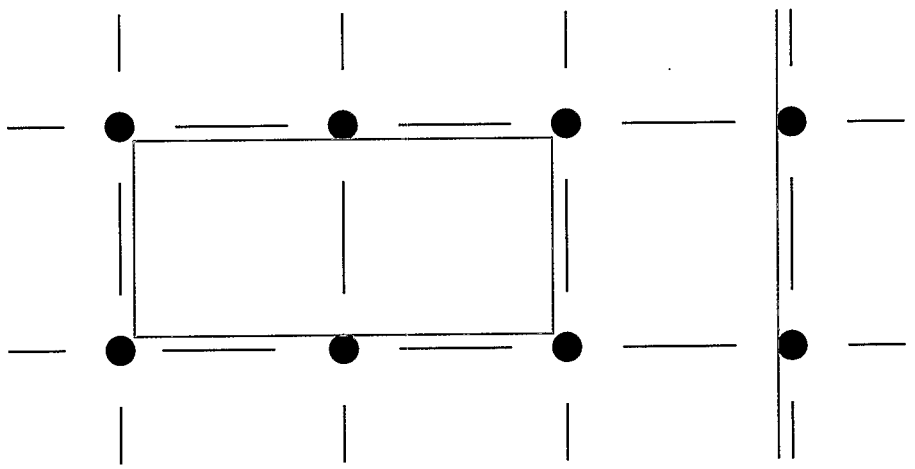

- Low-electrical power allows dense packing
- Evolution of successful QCDSP architecture.
- 20,000+ processor machines planned

QCDSP	QCDOC
• Faster processor	
TI C31	IBM PPC 440
0.05 Gflops	1.0 Gflops
32 bit precision	64 bit precision
0.7 Watts	$\leq$ 2 Watts
\$62 $\Rightarrow$ \$36	$\sim$ \$200
• Larger cache	
256 byte I-cache	32 KByte I-cache
0 word D-cache	32 KByte D-cache
• Faster communications	
50 MHz serial	500 Mhz serial
synchronous	phase locked loop
global sum, max	global sum, max
• Larger memory on chip	
2 Kwords SRAM	1 Mwords EDRAM
200 MBytes/sec bandwidth	8 GBytes/sec bandwidth
0 cycle latency	10 cycle latency

QCDSP	QCDOC
• Larger external memory	
2 Mbyte DRAM	$\geq$ 128 Mbyte DDR SDRAM
100 Mbytes/sec	2 Gbytes/sec
single bit EDC	single bit EDC
• I/O	
SCSI to some nodes	Ethernet to all nodes
Multi-layer host to node	Single-layer host to node
• Software tools	
C++/C/assembly	C++/C/assmbly
Tartan and TI	GNU, IBM XLC, ...
• Debuggers	
single DSP only	Possible on every node
	using RiscWatch
• Operating System	
Hardware debugging	Hardware debugging
Physical memory	Virtual memory
No protected modes	Protected modes

QCDSP	QCDOC
• Cost/performance (sust.)	
\$10/Mflops	\$1/Mflops
• Columbia machine	
8,192 processors	5,000 processors
0.4 Teraflops peak	5 Teraflops peak
• RBRC machine	
12,288 processors	10,000 processors
0.6 Teraflops peak	10 Teraflops peak
• UKQCD machine	
CRAY T3E, 216 processors	10,000 processors
0.2 Teraflops peak	10 Teraflops peak
• DOE Lattice QCD Topical Computing Center at BNL	
	20,000+ processors
	20+ Teraflops peak

QCDOC Performance and Scalability

From gate-level simulation of 500 MHz ASIC
$\rightarrow$ 1 Gflop peak (Peter Boyle, Lattice 2002)

- **Performance for lattice Dirac operators:**

Operation	Local Vol.	Perf/node (Mflops)
Wilson D_{eo}	2^4	470
Wilson D_{eo}	4^4	535
Clover D_{eo}	2^4	560
Clover D_{eo}	4^4	590
Staggered D_{eo}	2^4	370
Staggered D_{eo}	$2^2 \times 4^2$	430

- **Special hardware for global sums:**
  10μs for 4,000 nodes; 15μs for 32,000 nodes

- ***Estimate* for conjugate gradient with Wilson Dirac operator on a $32^3 \times 64$ lattice:**

Nodes	$M^\dagger M$+linalg	Global Sum	Sust. Tflops
4096	2620μs	10 μs	2.15
8192	1310μs	11.5μs	4.2
16384	680 μs	13 μs	8.1
32768	340 μs	15 μs	15.6

- **$\approx 50\%$ performance for fixed size problem while nodes change by $8\times$**

Potential CPU Needs for RHIC Modeling

Scott Pratt
Department of Physics and Astronomy
Michigan State University
pratts@pa.msu.edu

I. The hybrid nature of RHIC modeling

RHIC modeling requires a hybrid approach for several reasons:

1. The internal degrees for freedom switch from partonic to hadronic nearly half way through the collision.
2. The change of phase may involve non-trivial changes in vacuum energy which can make Boltzmann treatments problematic.
3. At the end of the collision, mean free paths become long, thus making even viscous hydrodynamic treatments irrelevant.
4. During much of the collision, Boltzmann treatments are questionable due to the necessity of including three and four body reactions.
5. During the first 0.5 fm/c particles are off-shell by an amount greater than the temperature due to the uncertainty principle.

As a crude illustration, consider the plot below which shows the regions of validity for several approaches.

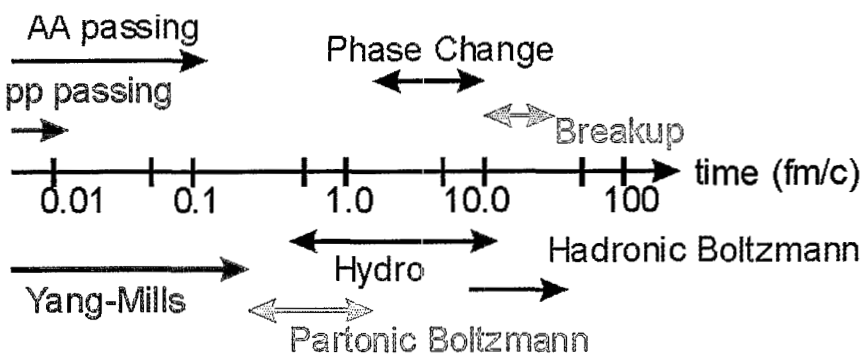

Given the variety of approaches it is difficult to accurately express CPU needs. This is especially true for approaches for the first one fm/c which are still in flux. However, we can make fairly reasonable estimates for hydrodynamic and Boltzmann approaches which cover the physics after the first 1.0 fm/c.

II. Status of Boltzmann models

Boltzmann approaches have provided the bulk of our detailed understanding of heavy ion collisions. Typically, these approaches involve simulating the collision with particles where one keeps track of the momenta and positions of the particles while invoking an oversampling the particles by a factor λ. As the cross section is scaled inversely with λ the simulations approach the Boltzmann limit of large λ.

It is only recently that these models have been checked to see whether the Boltzmann limit has indeed been achieved. (S. Cheng et al, PRC 034901, 2002, see talk of D. Molnar later in this workshop). To illustrate these studies, we show a plot of the transverse energy per particle for a 1-d expanding system. The result for λ =1 is shown in circles. The Boltzmann result should match the Navier-Stokes result, however due to the finite interaction range, $\sigma^{1/2}$, an effective viscosity ensues at high density. Repeating this calculation with λ ~20 indeed reproduces the Navier-Stokes result. Other observables were also sensitive to λ. For instance, the elliptic flow varied by a factor of 2 as λ was raised from one to ~ 20 where the result converged. If these calculations had been run from earlier times, even larger values of λ would have been needed for convergence.

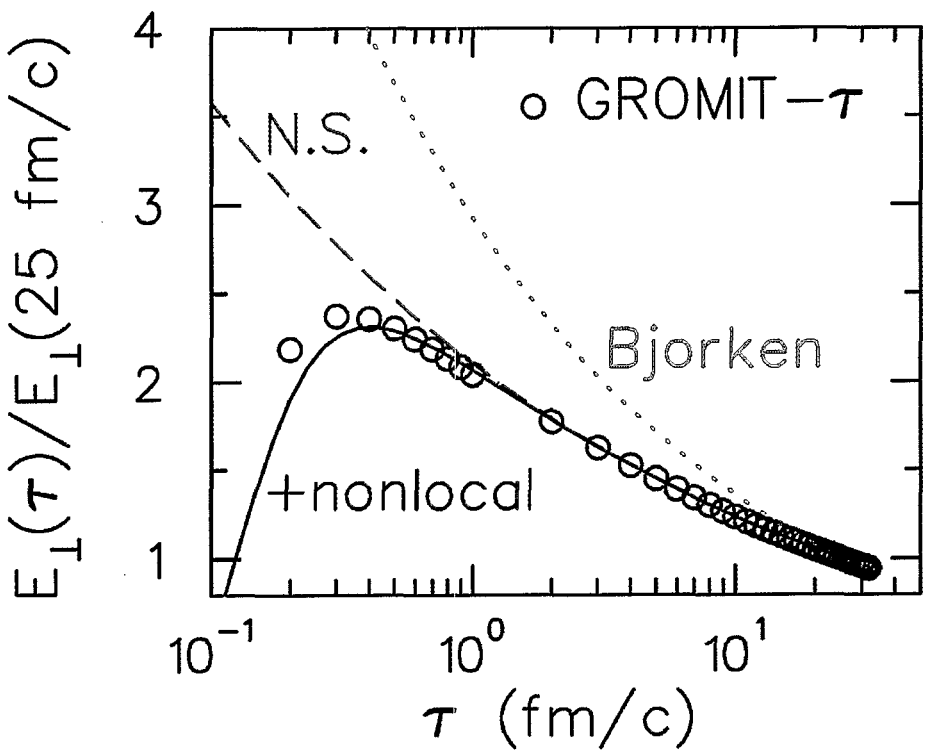

Since the calculations involve fairly large sampling factors, Boltzmann codes must be able to effectively sample small cross sections with a very large number of particles, ~ 100,000. In the last two years, several codes have been written with collision meshes so that this sampling can be performed efficiently. One such model, GROMIT, employs an expanding mesh so that it remains efficient over a large fraction of the collision. GROMIT is able to handle such a collision with a sampling rate of 20 in approximately one half hour on a PC. This is even faster than the $\lambda = 1$ codes like RQMD or URQMD which had required nearly a half hour to sample one event with no oversampling. Thus, GROMIT is effectively 20 times faster than these other approaches despite the additional challenge of oversampling.

Boltzmann codes have been run with both partonic and hadronic degrees of freedom. GROMIT was written to accommodate any choice of particles, i.e. the modular framework is independent of the physics. To read more about GROMIT, visit http://www.phy.duke.edu/rttc.

Boltzmann analyses are able to provide sufficient statistics for both spectra and correlations with approximately 10,000 events per impact parameter. Thus, a powerful PC can provide sufficient statistics with a few days running.

III. Hydrodynamic Calculations

Hydrodynamics provides an especially insightful means to model RHIC physics since the inherent vocabulary is directly related to the most pertinent questions addressed at RHIC, the equation of state and novel changes of phase. However, the mesoscopic nature of the collision volume demands that we one includes viscosity in the calculations. This only just beginning in the RHIC field. The Muller-Israel-Stewart equations are especially promising since they allow one to consider large asymmetries in the stress-energy tensor, which may be of the order of two at RHIC. (A. Muronga, PRL **88** 062302, 2002)

Our limited experience with these approaches suggests that a mesh of a few hundred cubed should be sufficient, which should add up to less than a GB. It is not expected that running such codes should require more than half a day of CPU on a PC. One of my motivations for attending this meeting is to

gain a better understanding of how our community can scavenge the astrophysics community for hydrodynamic modules.

IV. Modeling the first fm/c

The first one fm/c is especially challenging due to the highly un-thermalized nature and the inherently off-shellness of any particles. Due to the high phase space density at the beginning, it seems reasonable to solve classical Yang-Mills equations. However, it remains very open as to exactly how reasonable this approach is. Please see the talk of Alex Krasnitz.

Even if methods for the first fm/c never quantitatively attack the problem, one can still make progress by treating the conditions at 1.0 fm/c as an unknown and model the remaining physics with the aforementioned approaches. However, if these problems can be better understood, even by qualitatively understanding the trends which affect the conditions at 1.0 fm/c, it would very greatly affect our ability to interpret RHIC experiment.

V. Interpreting RHIC DATA

Unfortunately, nearly all RHIC observables are sensitive to a variety of unknown parameters and choice of scenarios which govern RHIC models. Therefore, it should be strongly stressed that running transport codes with one parameter set will certainly be insufficient. It is only by systematically varying the inputs to the model with minimum human bias that one can make firm quantitative conclusions the novel matter produced at RHIC.

A systematic search-and-destroy mission that would find the both the best parameters and the best model assumptions would represent both a tremendous numerical challenge as well as a daunting intellectual challenge. As soon as the community produces models which come close to the main features of the data, we will have to confront these issues.

At that point, it may prove that machines like the QCDOC or large BEOWULF clusters are required to make significant progress. If one is modeling observables for 10,000 parameter sets, it is not clear whether it is more cost effective to run the parameter sets consecutively on a machine like the QCDOC which would exploit the local-in-space-time nature of transport theory and divide a single run over many processors. Instead, one might be able to model the collision 10,000 times in parallel. Currently, it is difficult

to make definitive statements because several major puzzles have surfaced in the RHIC data, i.e. the HBT puzzle, the nearly hydrodynamic nature of elliptic flow and the strong jet quenching. In a few years, after the models have matured, it should be possible to determine an optimum strategy for confronting these problems.

Computational Radiation-Hydrodynamics for Core Collapse Supernova Simulations

F. Douglas Swesty
Dept. of Physics & Astronomy
SUNY Stony Brook

The explosion mechanism of core collapse supernovae remains elusive despite a nearly three decade long modeling effort. However, until recently we have lacked sufficient computational resources to carry out full multidimensional radiation-hydrodynamic models of supernovae. However, the recent expansion in the capabilities of MPP architectures, combined with algorithmic advances in the solution of both sparse linear and sparse non-linear systems, has rendered fully multidimensional radiation-hydrodynamics tractable for the first time.

Radiation-hydrodynamic models of supernovae rely on implementations of two major types of algorithms: explicit compressible fluid dynamics algorithms and implicit radiation transport algorithms. Both of these classes of algorithms are amenable to implementation on parallel architectures by means of spatial domain decomposition.

The implementation of explicit algorithms in a message-passing parallelism paradigm is straightforward. Under spatial domain decomposition the problem is broken up into a set of subdomain tiles (for 2-D problems) or blocks (for 3-D problems) that cover the global spatial domain. Each of these subdomains is instantiated on a separate processor using a Cartesian process topology. With this decomposition the only two communication patterns required to implement the algorithm. Most communication involves the exchange of ghost-zone information between nearest neighbors in the process topology. This communication must take place at several points throughout the timestep. The calculation of the timestep requires a global reduction operation to find the minimum CFL time across all processes.

The implicit radiation transport algorithms dominate the computational cost of supernova simulations. These models rely entirely on the use of Krylov subspace algorithms to solve the sparse linear systems or the use of Newton-Krylov algorithms to solve sparse non-linear systems. These algorithms require only three operations: matrix-vector multiplies (requires only nearest-neighbor communication), vector-additions (embarrassingly parallel), and vector inner inner products (requires a global reduction operation). The implementation of each of these operations is straightforward on any parallel architecture.

Because of the extensive use of Krylov subspace algorithms, computational radiation hydrodynamics simulations have much in common with lattice QCD simulations. Furthermore, both of these problems continue to require computing resources beyond the Tera scale in order to push the state of the field forward. Given the algorithmic commonalties between the two problems it seems likely that the QCDOC architecture could prove to be highly usable for any nuclear astrophysical simulations that rely on computational radiation-hydrodynamics.

TeraScale Supernova Initiative

Newtonian Hydrodynamics

$$\frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \mathbf{v}) = 0$$

Continuity Equation

$$\frac{\partial E}{\partial t} + \nabla \cdot (E\mathbf{v}) + P\nabla \cdot \mathbf{v} = \mathrm{S}$$

Gas Energy Conservation

$$\frac{\partial \rho v_i}{\partial t} + \nabla \cdot (\rho v_i \mathbf{v}) + (\nabla P)_i + \rho (\nabla \Phi)_i = \mathrm{A}_i$$

Gas Momentum Conservation

-Discretize variables on a 1-D, 2-D, or 3-D spatial mesh

-Solution via explicit finite difference or finite volume techniques

-Courant-Friedrichs-Lewy stability criterion on timestep size: $\Delta t < \frac{1}{2}\frac{\Delta x}{c_s}$

F. D. Swesty

STONY BROOK

QCDOC Meeting BNL 9/26/02

http://www.astro.sunysb.edu/dswesty
dswesty@mail.astro.sunysb.edu

TeraScale Supernova Initiative

Parallel Implementation

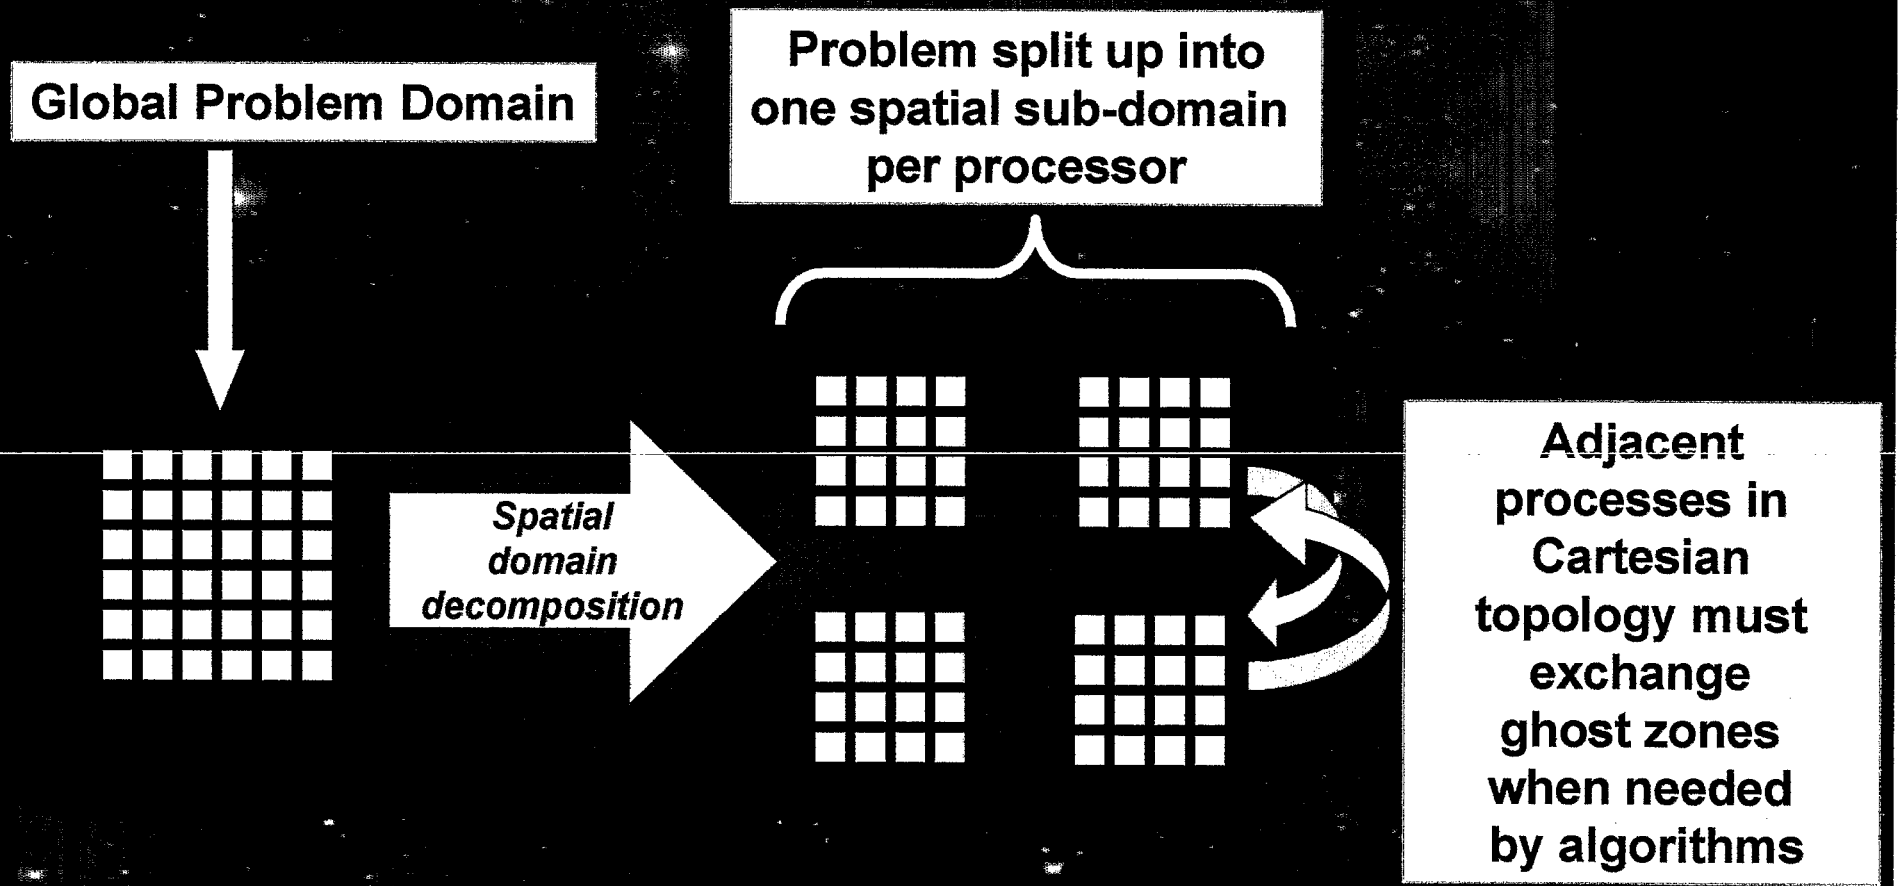

F. D. Swesty

STONY BROOK

QCDOC Meeting BNL 9/26/02

http://www.astro.sunysb.edu/dswesty
dswesty@mail.astro.sunysb.edu

TeraScale Supernova Initiative

Linear System Solution

- *Krylov Subspace Algorithms*
- *Idea: rather than solve* $Ax = b$
- *Minimize* $\Phi(x) = \frac{1}{2}x^T Ax - bx$
- *Use minimization algorithms to reduce the residual* $r = Ax - b$
- *Iteratively update the solution vector*
- *Requires only:*
  - ***Matrix-vector multiplies (matvecs)***
  - ***Vector additions (embarrassingly parallel)***
  - ***Vector inner products (global reduction operation)***
  - ***Matrix in operator form***

F. D. Swesty

STONY BROOK

http://www.astro.sunysb.edu/dswesty
dswesty@mail.astro.sunysb.edu

TeraScale Supernova Initiative

Neutrino Transport

We would like to solve the Boltzmann equation:

$$\frac{\partial f}{\partial t} + \frac{\partial f}{\partial x^\mu}\frac{\partial x^\mu}{\partial t} + \frac{\partial f}{\partial p^\mu}\frac{\partial p^\mu}{\partial t} = \left(\frac{\partial f}{\partial t}\right)_{coll}$$

Or in terms of the radiation intensity: $I = I(\epsilon, \mathbf{x}, \mathbf{\Omega}, t)$

$$\frac{1}{c}\frac{\partial I}{\partial t} + \mathbf{\Omega}\cdot\nabla I + a_\mu \frac{\partial I}{\partial p^\mu} = \left(\frac{\partial f}{\partial t}\right)_{coll}$$

- Equations are have 3 spatial + 2 angular + 1 spectral dimension
- Must solve 1 pair of equations for each neutrino flavor
- Too difficult for 3-D long timescale problems at present!
- For 2-D we can do this at the TeraScale level

$$E_\epsilon = \frac{1}{c}\int I\, d\Omega$$

In 3-D: Integrate over direction angle to reduce dimensionality

F. D. Swesty

STONY BROOK

http://www.astro.sunysb.edu/dswesty
dswesty@mail.astro.sunysb.edu

TeraScale Supernova Initiative

Relativistic Hydrodynamics

$$\frac{1}{\sqrt{-g}}\frac{\partial}{\partial x^0}(D\sqrt{-g}) + \frac{1}{\sqrt{-g}}\frac{\partial}{\partial x^i}(DV^i\sqrt{-g}) = 0$$

Baryon Conservation

$$\frac{1}{\sqrt{-g}}\frac{\partial}{\partial x^0}(E\sqrt{-g}) + \frac{1}{\sqrt{-g}}\frac{\partial}{\partial x^i}(EV^i\sqrt{-g}) + p\frac{\partial}{\partial x^\mu}(u^0V^\mu\sqrt{-g}) = 0$$

Gas Energy Cons.

$$\frac{1}{\sqrt{-g}}\frac{\partial}{\partial x^0}(S_\mu\sqrt{-g}) + \frac{1}{\sqrt{-g}}\frac{\partial}{\partial x^i}(S_\mu V^i\sqrt{-g}) + \frac{\partial p}{\partial x^\mu} + \frac{1}{2}\frac{\partial g^{\alpha\beta}}{\partial x^\mu}\frac{S_\alpha S_\beta}{S^0} = 0$$

Gas Mom. Cons.

$$u^\mu u_\mu = -1$$

4-Momentum Normalization

-Relativistic Hydro equations can be solved with explicit methods similar to those applied to the Newtonian equations

$$D = \rho u^0$$

$$E = \rho\epsilon u^0$$

$$S_\mu = \rho(1 + \epsilon + P/\rho)u_\mu u^0$$

$$V_\mu = u_\mu/u^0$$

F. D. Swesty

STONY BROOK

http://www.astro.sunysb.edu/dswesty
dswesty@mail.astro.sunysb.edu

Alexander Krasnitz
CENTRA and FCT Universidade do Algarve
Campus de Gambelas, P-8000 Faro, Portugal

krasnitz@ualg.pt

It was realized some time ago that, in a heavy nucleus in the infinite-momentum frame, gluons at low Bjorken x form a high-density state, dubbed a color glass condensate, that admits a description in terms of an effective classical field theory [1]. This effective description has been a focus of much interest lately [2]. When applied to gluon fields in the central-rapidity region of a heavy-ion collision, this approach, combined with real-time numerical methods [3], yielded a nonperturbative determination of the transverse energy density and of the momentum space gluon distribution [4].A study of the chiral charge produced in the central region has also been completed [5], as well as an analysis of differential and integral elliptic flows [6].

In a boost-invariant approximation the numerical problem is that of a 2+1-dimensional lattice gauge theory, whose solution involves familiar algorithms and techniques of lattice QCD. There are a number of dimensional parameters in the problem, forming a hierarchy of scales. As a result, a realistic calculation within the model requires computational resources beyond those of a single-processor machine. A massively parallel computation will be necessary in order to relax the assumption of the strict boost invariance, or to couple the classical fields to hard particles, or to study quark and hard gluon production in the classical background. As usually is the case in lattice gauge theory, the problem lends itself easily to parallelization. It therefore is ideally suited for QCDOC and is likely to derive maximal benefit from the QCD-specific features of the machine.

References

[1] L. McLerran, R. Venugopalan, Phys. Rev. D49 (1994) 2233; ibid, 3352.

[2] For a recent review, see L. McLerran, Nucl. Phys. A702 (2002) 49.

[3] A. Krasnitz, R. Venugopalan, Nucl. Phys. B557 (1999) 237.

[4] A. Krasnitz, R. Venugopalan, Phys. Rev. Lett. 84 (2000) 4309, Phys. Rev. Lett. 86 (2001) 1717; A. Krasnitz, Y. Nara, R. Venugopalan, Phys.Rev.Lett.87 (2001) 192302, hep-ph/0209269.

[5] D. Kharzeev, A. Krasnitz, R. Venugopalan, hep-ph/0109253, to appear in Phys. Lett. B.

[6] A. Krasnitz, Y. Nara, R. Venugopalan, hep-ph/0204361.

Classical Numerical Gluodynamics of Nuclear Collisions

- Formulation of the problem.
- Physics so far: multiplicity, elliptic flow, chiral charge.
- Numerics so far: algorithms, problem size, benchmarks.
- Future physics: rapidity fluctuations, particle d.o.f., pair production.
- Future numerics: parallelization.

A.K., R. Venugopalan, NPB **557** (1999) 237; PRL **84** (2000) 4309; PRL **86** (2001) 1717.

A.K., Y. Nara, R. Venugopalan, PRL **87** (2001) 192302; hep-ph/0204361; hep-ph/0209269.

D. Kharzeev, A.K., R. Venugopalan, hep-ph/0109253 (PLB).

Formulation of the problem

Classical physics at small x (a reminder of MV model)

High-x partons are considered recoilless sources of color charge. For a large, Lorentz-contracted nucleus this color charge distribution is Gaussian in the transverse plane.

$$P\left([\rho]\right) \propto \exp\left[-\frac{1}{2\Lambda_s^2}\int \mathrm{d}^2 r_t \rho^2(r_t)\right].$$

This color charge is static in the transverse plane.

A nucleus is considered infinitely thin in the longitudinal direction. This assumption can be relaxed if necessary.

Λ_s is a dimensional parameter which determines the color charge density and satisfies a self-consistency relation

$$\Lambda_s^2 = \frac{\alpha_s \pi^2}{2\sigma C_F} x G(x, \Lambda_s^2).$$

YM equations for low-x fields in the random color-charge background:

$$D_\mu F_{\mu\nu} = J_\nu,$$

Discretize on a 2d lattice.

Dimensional quantities in the classical lattice theory:

- Λ_s
- R, the nuclear radius
- l, the color neutrality scale (a recent development!)
- a, the lattice cutoff

Hierarchy of scales (ideal): $1/a \gg \Lambda_s \gg 1/l \gg 1/R$

In the units of a, in the continuum limit $\Lambda_s \to 0$, $R \to \infty$, but $\Lambda_s R$ is constant.

For any well-defined P of dimension d

$$P = (\Lambda_s)^d f_P(\Lambda_s R),$$

where $f_P(\Lambda_s R)$ contains all the non-trivial physical information.

- RHIC – $\Lambda_s \approx 1.4$ GeV
- LHC – $\Lambda_s \approx 2.2$ GeV

Physics so far

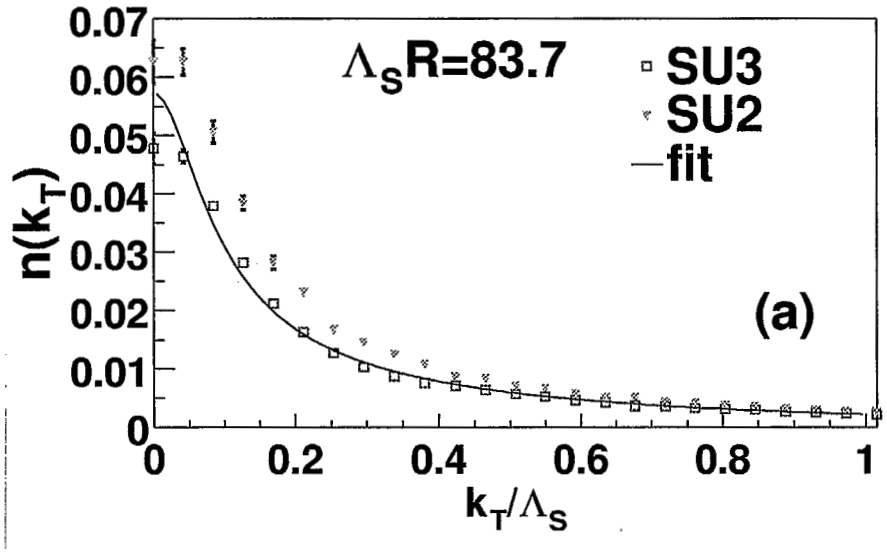

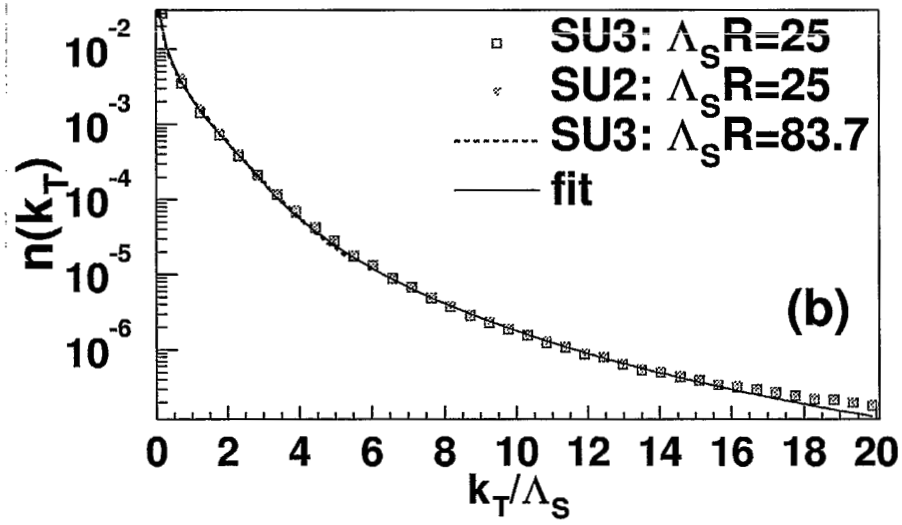

Gluon multiplicity $n\left(k_T\right) = \tilde{f}_n/(N_c^2 - 1)$

The SU(3) gluon momentum distribution can be fitted by the following function,

$$\frac{1}{\pi R^2}\frac{dN}{d\eta d^2 k_T} = \frac{1}{g^2}\tilde{f}_n(k_T/\Lambda_s)\,, \qquad (1)$$

where $\tilde{f}_n(k_T/\Lambda_s)$ is

$$\tilde{f}_n = \begin{cases} a_1\left[\exp\left(\sqrt{k_T^2+m^2}/T_{\rm eff}\right)-1\right]^{-1} \\ (k_T/\Lambda_s \le 3) \\ \\ a_2\,\Lambda_s^4 \log(4\pi k_T/\Lambda_s) k_T^{-4} \\ (k_T/\Lambda_s > 3) \end{cases} \qquad (2)$$

with $a_1 = 0.0295$, $m = 0.067\Lambda_s$, $T_{\rm eff} = 0.93\Lambda_s$, and $a_2 = 0.0343$.

Numerics so far

Numerical effort

Layout: adopted a MILC-like strategy $\longrightarrow$ lattice is an array of sites.

Memory: The lattice site structure is 190 bytes for the SU(2) gauge group, 600 bytes for the SU(3) gauge group.

A "reasonable" simulation at RHIC or LHC regimes requires a 1024×1024 lattice $\longrightarrow \approx 200$ Mb for SU(2), ≈ 600 Mb for SU(3).

To do even slightly better would require parallelization!

Flop count: Proper-time evolution (leapfrog algorithm) requires $\approx$ 500 Flop per step per site (SU(2)).

$1024 \times 1024 \times 600$ steps $\times$ 500 configurations $\times$ 2 (measurements) $\approx 5\times10^{14}$ Flop

For SU(3), an identical simulation would require $\approx 2 \times 10^{15}$ Flop.

Future physics

Relax the exact boost invariance $\longrightarrow$ study rapidity dependence of observables $\longrightarrow$ data, Flop $\times$ 10, extra dimension.

Include high-p_t d.o.f. in a "quenched approximation" $\longrightarrow$ determine pair production rates for quarks, gluons in the classical background $\longrightarrow$ expect effort to grow faster than L^2.[1]

Include high-p_t d.o.f. dynamically, as classical particles or as a distribution function $\longrightarrow$ likely data, Flop $\times$ 10.[2]

Potentially, a 100 Gb (per configuration), 10^{17} Flop problem!

Our future physics goals require parallelization! (Effort, data size far beyond capacities of a single processor).

The problem is LGT $\longrightarrow$ lends itself easily to parallelization, requires mostly nearest-neighbor communications.

Our current code has been written with future parallelization in mind.

Would benefit from QCDOC!

[1] Discussions with P. Levai
[2] Discussions with L. Mclerran

Classical Molecular Dynamics on QCDOC

J. W. Davenport[a], Y. Deng[b], J. Glimm[a,b], E. Santos[c]

[a]Center for Data Intensive Computing, Brookhaven National Laboratory, Upton, NY 11973, [b]Department of Applied Mathematics and Statistics, University at Stony Brook, Stony Brook, NY 11794, [c]Department of Computer Science, Virginia Tech, Blacksburg, VA 24061

Molecular dynamics simulations [1,2] are ideal for massively parallel computing because the time spent on communication can be a small fraction of the time spent on computation. However, most molecular dynamics (MD) simulations have been performed on machines where this potential has not been realized because of slow communication and/or small number of processors.

We explore here the use of a new machine designed for quantum chromodynamics (QCD) by a team of elementary particle physicists [3] mainly at Columbia University, The Riken BNL Research Center and IBM. This machine consists of approximately 10,000 IBM PowerPC processors with extremely fast nearest neighbor communication and 4 megabytes of memory located on the same chip. For this reason it has been named QCDOC for quantum chromodynamics on a chip. It is expected to achieve of order 5 teraflops (or 5×10^{15} floating point operations per second) at a cost of roughly $500 per processor.

For a class of MD simulations described below we find that calculations on 10^6 atoms with simulation times up to 10 microseconds are feasible with 5 - 6 weeks of computing time.

Such simulations are important in areas such as nanoscience and biology where the complexity of the atomic arrangements, the difficulty in performing experiments in non crystalline environments, and the need to understand crystal growth and other dynamic processes are crucial.

References

1. M. P. Allen and D. J. Tildesley, Computer Simulation of Liquids, Oxford, NY, 1987.

2. D. Frenkel and B. Smit, Understanding Molecular Simulation, Academic, San Diego, 1996.

3. D. Chen, N. H. Christ, C. Cristian, Z. Dong, A. Gara, K. Garg, B. Joo, C. Kim, L. Levkova, X. Liao, R. D. Mawhinney, S. Ohta, and T. Wettig, hep-lat/0011004.

On the Use of QCDOC in Condensed Matter Physics

RIKEN-BNL Workshop, Brookhaven National Lab

September 26-28 2001

J. W. Davenport[1], Y. Deng[2]
J. Glimm[2], E. Santos[3]

1.CDIC, Brookhaven, 2. Applied Math and Statistics, Stony Brook, 3. Computer Science Virginia Polytechnic Institute

Why QCDOC?

- 10 teraflops
  - ~ 10,000 processors
  - ~ 1 gigaflops each
  - 4 megabytes memory on chip
- fast communication
  - 500 megabits/sec simultaneously to each neighbor in a 6d mesh
  - Low latency – 0.200 μsec
- but…nearest neighbor only

Molecular Dynamics

- Interatomic Potential
  - Embedded Atom
    - Pair term
      - Lennard-Jones, Morse
    - Embedding term
      - Density dependent
    - Cut off after several neighbors
  - Long range terms (e.g 1/r)
    - Ewald techniques
- Force = grad (Potential)

Newton's 2nd Law

- $F(t) = M\, a(t)$
- $x(t + d) = x(t) + v(t)\, d + \frac{1}{2}\, a(t)\, d^2$
- $x(t - d) = x(t) - v(t)\, d + \frac{1}{2}\, a(t)\, d^2$
- $x(t + d) = 2\, x(t) - x(t - d) + a(t)\, d^2$
- Start
  - random positions
  - Velocities drawn from a Gaussian
  - $3/2\, k_B T = \langle P^2/2M \rangle$

How to Parallelize?

- Domain decomposition
  - 100 particles on each of 10,000 processors
  - Easily fits on 4 MB
  - Computation @ each time step
    - ~ 10^3 operations / particle
    - ~ 10^5 ops
    - = $0.2\ 10^{-3}$ sec @ 500 megaflops
  - Communication
    - 3 x 100 particles x 64 bits / 500Mb / sec
    - x 3 corner cells
    - = 0.12×10^{-3} sec

Long time simulations

- For d = 10^{-15} sec
- 10^{10} time steps
  - 800 hours = 5 weeks
- 10 μ sec simulated time!
- ~3 teraflops
- Much longer than most current simulations

Long Range Forces

- Long range treated as Fourier sum
  - Small number of terms
  - Need to communicate across all cells
  - Multi hops (NlogN)
  - Other dimensions
  - Fewer evaluations (Berne et al)
- Increased communication, decreased computation
- Comparable times

Conclusion

- With QCDOC can do classical MD on large systems for up to 10 μsec!
- Other applications should be investigated
- Maybe its not such a special purpose machine after all!

Using the QCDSP

(outside the Columbia physics environment)

Michael Creutz

Brookhaven National Laboratory

Resources

- Users Guide, Mawhinney and Joo:
  http://phys.columbia.edu/~cqft/qcdsp/qcdsp.htm
- on hosts: /qcdsp/sfw/<os version>/example/ hello_world
- under qshell: "qhelp"
- sysfunc.h
- memory map:
  http://thy.phy.bnl.gov/~creutz/qcdsp/memory.ps

Write a C/C++ program

- that's the compiler we have (Tartan, from TI)
- QCDOC will have full power-pc software environment

Compile on workstation qmemcpy()

- "#if ONQCDSP" for special parts
- dummy system functions for workstation
- debug: your favorite tools

Move to qcd host machine

- compile with Tartan compiler
- debug: compiler diagnostics

Obtain machine

- start qshell
- run program
- debug: printf(), qshell tools

Interprocessor communication

4d mesh, 8 wires per node, $(\pm X, \pm Y, \pm Z, \pm T)$

I. System functions

- described in sysfunc.h

II. Write to serial communication unit (SCU) registers
- SCU_BASE=0x813000
- Registers hold block size, stride, number of blocks

```
for (i = 0; i < 8; i++) {
    *((int *) (SCU_BASE + 0x50 + i)) =
        (1 << 22) | (sizeof(message) << 12) | 1;
    *((int *) (SCU_BASE + 0x48 + i)) = 0;
}
```

Address registers for location of data
- start reception at `*dest` with:

```
*((float**) (SCU_BASE + wire)) = (float*) dest;
```

- start sending from `*source` with

```
*((float**) (SCU_BASE + 0x8 + wire)) = (float*) source;
```

Read communication status at (SCU_BASE + 0x40)
- bits map onto wires
- do other things while communication proceeds

QCDOC differences
- standard commodity power-pc processor
- more compiler choices: gnu, ibm
- machine partitioning in os: user requests size
- 6-d basic machine
- two wires for each neighbor
- simultaneous sends and/or receives
- interupts handled differently

```
//------------------------------------------------------------------------
//  test_printf.C   7/15/97  RDM
//
//  Testing the print buffers on QCDSP for qos.5 and the qcsh
//------------------------------------------------------------------------


#include <math.h>
#include <stdio.h>
#include <stdlib.h>
#include <sysfunc.h>


//------------------------------------------------------------------------
//  Print commands
//------------------------------------------------------------------------


main()
{
  int i, j;
  float x, y;
  char s[50];

  printf("\nHello world\n");
  printf("Motherboard and daughterboard numbers for this node:  %d %d\n"
    , MbNum(), DbNum() );
  printf("Unique ID for this node:  %d\n", UniqueID() );
  printf("Physics coordinates for this node:  %d, %d, %d, %d\n",
    CoorT(), CoorX(), CoorY(), CoorZ() );

  for ( i = 0; i < 10; i++ )
    printf("The square of %d is %d\n", i, i * i );

  sprintf( s, "End of hello world from node %d\n", UniqueID() );

  j = ( UniqueID() % 10 ) * 100000 + 100000;

  y = 0;

  for ( i = 0; i < j; i++ ) {

    x = i * i +  3.764 * i - (7.81 * i ) / 3.45;

    y += sin(x);
  }

//  StringExit( UniqueID() , s );

  InterruptExit( UniqueID() , s );

  exit( UniqueID() );
}
```

Cluster Algorithms for Strongly Coupled Lattice Gauge Theory

Shailesh Chandrasekharan
Duke University

Matter is made up of fermions at various scales starting from quarks and leptopns, all the way to nuclear, atomic and molecular physics. Phenomena like high T_c superconductivity and superfluidity in ^3He arise due to the fermionic nature of the underlying degrees of freedom. Some of the theories that describe these systems, like QCD, the Hubbard model, many-body quantum mechanics involving the coulomb potential remain unsolved due to lack of computational method for fermions.

The current state of Fermion algorithms is quite primitive. It begins by integrating the fermions out and works using non-local bosonic degrees of freedom. Recently, a new alternative has emerged in a limited class of models [1]. In this method one works directly with fermionic degrees of freedom. Since fermions obey the Pauli principle this leads to two problems. Firstly, one is forced to work with constraints since two identical fermions cannot be found at the same position. Secondly, one has to grapple with sign problems since a permutation of fermions leads to a negative sign. However, when one can solve these problems, the new methods that become available are far superior to conventional algorithms.

There are many examples where the new approach has been successful. Here we propose that lattice gauge theories in the strong coupling limit provide a unique opportunity for us to learn about this new approach to fermion dynamics in QCD. We show that one can constuct a very efficient algorithm in the chiral limit and compute quantities with unprecidented accuracy.

[1] See for example cond-mat/0201360 for a recent review of the method.

Performance of the Hybrid Monte Carlo in QCD

- Lattice 2001 result: Resources necessary in (Teraflops Years).

$$N_{Flops} = 2.8 \left(\frac{\#\text{Confs}}{1000}\right) \left(\frac{0.6}{M_\pi/M_\rho}\right)^6 \left(\frac{L}{3fm}\right)^5 \left(\frac{1/a}{(2GeV)}\right)^7$$

  Physical $M_\pi/M_\rho \sim 0.18$.

  The current algorithms breakdown close to the chiral limit.

- Current estimates are that we will need around 10-100Tlop years to do the first round of calculations at "light enough" quark masses!

  "Algorithmic breakthrough essential for dynamical fermions".

  Lüscher, Lattice 2001.

Cluster Approach to Path Integrals

- Rewrite the partition function using "cluster" variables:

$$Z = \sum_{[\phi]} W[\phi] = \sum_{[\mathcal{C}]} W'[\mathcal{C}]$$

ϕ is a local field,
$W[\phi]$ contains non-local interactions

$\mathcal{C}$ is a non-local cluster variable,
$W'[\mathcal{C}]$ contains local interactions

- Cluster Variables represent correlated points and are very useful
  - they help in non-local updates
  - observables are related to properties of clusters
  - often one can work directly at $T = T_c$ and/or $m = 0$ etc..
  - can sometimes solve "sign problems"
- Usefulness of cluster variables have been demonstrated in many models:

Monomer-Dimer Representation at strong couplings ($\beta = 0$)

Reference: P. Rossi, U. Wolff Nucl. Phys. **B248** (1984) 105.

- At strong couplings in $U(N)$ gauge theory it is possible to show

$$Z = \sum_{[n,b]} \prod_{x,\mu} \frac{(N - b_{x,\mu})!}{b_{x,\mu}! N!} \prod_x \frac{N!}{n_x!} m^{n_x}$$

site variable: $n_x = 0, 1.., N$

bond variable: $b_{x,\mu} = 0, 1, ..., N$

- For $N = 1$ and $d = 2$ configurations involve monomers and dimers

$$Z = \sum_{[n,b]} \prod_x m^{n_x}$$

Fermions lead to constraints in the configurations

$$n_x + b_{x,1} + b_{x,2} + b_{x-\hat{1},1} + b_{x-\hat{2},2} = 1$$

Chiral Susceptibility vs. Lattice size at $m = 0$ for $N = 1$

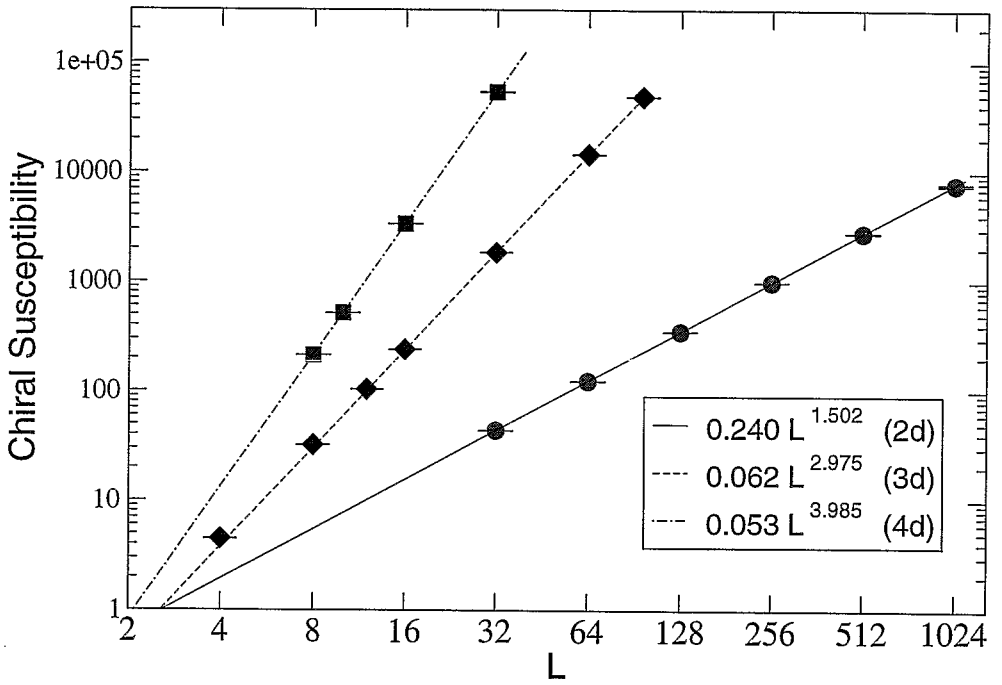

Chiral Symmetry breaking indeed occurs for $d = 3, 4$ but not in $d = 2$.

Strong Coupling QCD

- Must consider SU(3) gauge group instead of $U(3)$.
- Baryonic Loops are allowed and can have negative signs.
- Sign problem can be solved using cluster techniques

Reference: F. Karsch, K. Mütter Nucl. Phys. **B313** (1989) 541.

Reference: S.C, U.-J. Wiese, Phys. Rev. Lett. **83** (1999) 3116.

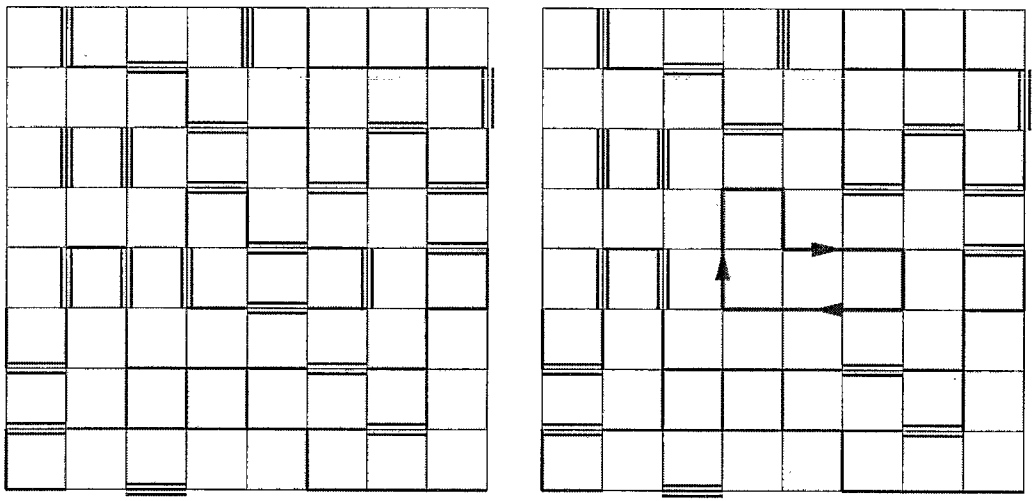

Solving nonlinear dynamical equations with QCDOC horsepower

Dénes Molnár

Ohio State University, Columbus, OH 43210

The interpretation of experimental data from relativistic heavy-ion collider experiments requires a Lorentz-covariant dynamical model. Even the simplest such models lead to nonlinear equations, which are rather challenging to solve. The numerical difficulty increases rapidly as one proceeds from the simplest equilibrium theory, Eulerian hydrodynamics, toward general nonequilibrium elastic and inelastic classical transport theory. I discuss the feasibility of solving these theories on the QCDOC computer for initial conditions expected in Au+Au at $\sqrt{s} \sim 100 - 200A$ GeV at RHIC based on recent elliptic flow and pion interferometry data.

References:

- PRC 62, 054907 ('00) [nucl-th/0005051]
- NPA 697, 495 ('02) [nucl-th/0104073]
- PRC 65, 024901 ('02) [nucl-th/0107001]
- nucl-th/0204062

See complete talk on WWW at
http://nt3.phys.columbia.edu/people/molnard

Lesson One: avoid black boxes

Must separate the model (physics) and the algorithm (code): (Pang)

model $\equiv$ mathematical equation

algorithm $\equiv$ numerical technique to solve the equation

Test: model against data $\Rightarrow$ meaningful physics conclusions

algorithm against analytic solutions $\Rightarrow$ ensures correct numerics

1st RBRC meeting, 1997: Open Standard Codes And Routines (OSCAR)

- **postulated the "bare minimum" requirements:**
  - **accessibility of code and documentation**
  - **conservation laws**
  - **common output format**
- **compliant codes available from OSCAR site at**

http://www-cunuke.phys.columbia.edu/OSCAR

Covariant classical transport theory

Pang, Zhang, Gyulassy, D.M., Vance, Csizmadia, Pratt, Cheng , ...

Simplest nonequilibrium dynamics $\rightarrow$ next step beyond hydrodynamics.

Physics:
- **energy-momentum and charge conservation**
- **finite rates & MFP $\Rightarrow$ thermal/chemical equil. not assumed**
- **on-shell physics clear $\leftrightarrow$ e.g., VNI cascade** Geiger, Mueller, Shrivastava, Bass
- **natural decoupling, MFP$(t \rightarrow \infty) \rightarrow \infty$**
- **no wave effects, no particle correlations, no phase transitions**
- **needs transition probabilities and initial conditions**

Equation: Lorentz-covariant Boltzmann eqn. for 1-particle distr., $f(x,\vec{p})$

$$p_1^\mu \partial_\mu f(x,\vec{p}_1) = S(x,\vec{p}_1) + C^{2\rightarrow 2}[f](x,\vec{p}_1) + C^{2\leftrightarrow 3}[f](x,\vec{p}_1) + ...$$

$$\text{elastic } 2 \rightarrow 2 \qquad \text{inelastic } 2 \leftrightarrow 3$$

$\Rightarrow$ partial integro-differential eqn. for 6+1D phase space evolution

Resource requirements for $2 \to 2$

Largely depend on initial conditions, transport parameters (opacity), and algorithmic optimizations.

- **Au+Au RHIC initial conditions, one MPC event w/ $2 \to 2$, 1-GHz CPU:**

$b = 12$ **fm**	$\chi \approx 3$, $\ell = 5500$ $(>\sim 40)$	~ 2 **hours**	$\sim 120 - 150$ **MB**
$b = 8$ **fm**	$\chi \approx 20$, $\ell = 450$ $(>\approx 250)$	~ 6 **hours**	$\sim 150 - 200$ **MB**
$b = 4$ **fm**	$\chi \approx 40$, $\ell = 150$ $(>\sim 300)$	~ 12 **hours**	$\sim 200 - 250$ **MB**
$b = 0$ **fm**	$\chi \approx 48$, $\ell = 100$ $(>\sim 500)$	~ 1 **day**	$\sim 300 - 400$ **MB**

Red: insufficient subdivision $\to$ numerical artifacts

[$N_{test} \sim$ a few hundred thousand]

- **Scaling with subdivision: CPU – $\sim \ell^{3/2} \to \sim$ a week**
  memory – between $\ell^{3/2}$ and $\ell^2 \to \sim$ 4-10 GB

$\Rightarrow$ single-CPU machines are not sufficient, might be barely

Resource estimate for $3 \leftrightarrow 2$ (RHIC)

$3 \to 2$ **is more challenging than** $2 \to X$**:**

- **extra loop when processing** $3 \to 2$ **(in 2nd step) – few times slower**
- **typically large 3-body range** $\Rightarrow$ **larger cells, much reduced efficiency,** $1/d^3$
- **slower** $\ell^{-1/5}$ **scaling toward Lorentz covariance**
- **larger subdivisions** $\sim 10^4 - 10^5$ $\Rightarrow$ **need** $10 - 100$ **more memory**

Σ : $\sim 2 - 4$ **orders of magnitude more demanding than** $2 \to 2$

$\Rightarrow$ **requires QCDOC horsepower**

Summary

- classical transport theory:
  - well-defined covariant nonequilibrium dynamical theory
  - addresses relevant questions for RHIC: equilibration, v_2 & HBT puzzles

- Molnar's Particle Cascade (MPC):
  - first covariant technique for elastic $2 \to 2$ and inelastic $3 \leftrightarrow 2$ inter.

- RHIC simulations on a single PC:
  - limited to elastic $2 \to 2$ and moderate $\chi \sim 20$ opacities
  - insufficient for central Au+Au, where recent data suggest $\chi \sim 45$

- Further progress requires teraflops of parallel computers $\to$ QCDOC
  - urgent for realistic parton equilibration studies $\leftrightarrow$ wishful arguments
  - needed anyway for future $\to$ Vlasov, Yang-Mills + Boltzmann, etc.

Potential Calculations for QCDOC

David J. Dean
Physics Division
Oak Ridge National Laboratory

- **Some comments on nuclear structure**
- **GFMC (light nuclei)**
- **Coupled Cluster techniques (new stuff)**
- **Auxiliary Field Monte Carlo for nuclei**

Overall motivations and goals

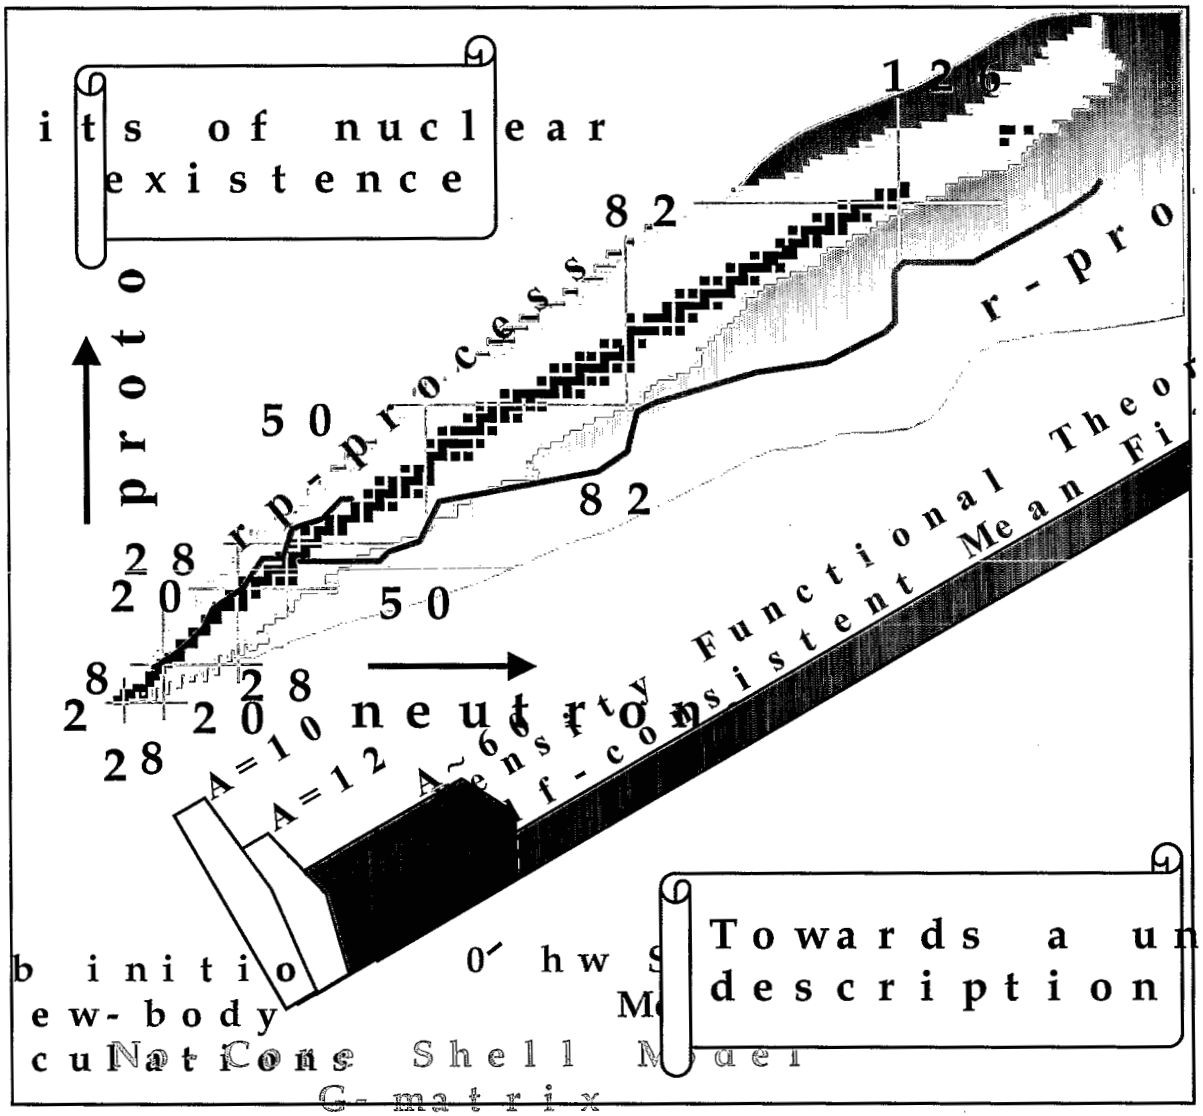

- Understand emergent phenomena in many-body systems starting from the nucleon-nucleon interaction
- How about Vnnn
- Connection to QCD?
- Predict limits of existence
- Understand nuclear effects in e.g. CKM matrix
- Describe nuclear effects in astrophysical scenarios:
  - nucleosynthesis
  - neutrino processes in stars
  - neutron matter
- Calculate density of states (for nuclear reactions)

Following three methods may be used to address these questions

Nuclear Coupled Cluster Method
Computational/Science considerations

- Currently runs at 200 Mflops/processor at NERSC (SP)
  - expect 3-4 speed up in matrix-matrix implementation.
- Utilize domain decomposition of the matrix elements to fit memory/proc.
- Load balancing is not an issue

Problem is mainly matrix-matrix or vector-vector multiplies. Modest size problems are ideal for the QCDOC cluster.

Science:

- **Ab initio calculations of A>16 nuclei**
- **Extension to open-shell systems will allow for studies of weakly bound nuclei.**
- **Application to reaction theory??**

Contacts: Dean (ORNL), Hjorth-Jensen (Oslo), Bishop (Manchester)

Nuclear AFMC Computational/Science considerations

fp-gds model space: Ca40 as a core
Nuclei range from A=40 to A=120

- single-particle space: 50x50 matrices
- One MC sample : 0.135 T-ops
- 5000 statistically independent and thermalized samples
  initial thermalization requires an extra 0.16 Tops;
- so for 5000 samples one needs 675 Tops.
- Small memory requirements (order $4N^3$)

Ideal for the QCDOC cluster.

Science:

- **Structure of medium mass nuclei (neutron-rich)**
- **Electron capture (and neutrino reactions) on these nuclei**
- **Thermal properties (density of states)**

Contacts: Dean (ORNL), Alhassid (Yale), Langanke (Aarhus), Ormand (LLNL)

Greens Function Monte Carlo
Computational/Science Considerations

- **Main computational load: extremely sparse matrix, H, with 8x8 and 4x4 block structure (not necessarily continuous) operating on a large vector, Ψ.**
- **MC sign problem (alleviated by constrained path techniques)**

$$H\left|\Psi\left(\Omega_1,\Omega_2,\ldots,\Omega_N\right)\right\rangle = \left|\Psi'\left(\Omega_1',\Omega_2',\ldots,\Omega_N'\right)\right\rangle$$

$$\Omega_i = \left\{\vec{r}_i,\vec{p}_i,\vec{\sigma}_i,\vec{\tau}_i\right\}$$

H includes 2 and 3 body interaction terms AV18+Illinois V3N

Science:

- **A=12 systems (50 Tflop-hours/state)**
- **12n + 1p (nuclear matter) (200 Tflop-hours)**
- **C-12 triple alpha burning (500 Tflop-hours)**
- **A=12 requires 2 Gbyte/processor**

Contacts: Pieper, Wiringa (ANL), Carlson (LANL)

Prospectus

Two established methods:

GFMC: Monte Carlo sampling in r,p,spin,isospin space

AFMC: Monte Carlo sampling in Slater Determinant space

One new method: Coupled-Cluster theory

Solution of coupled non-linear equations

Each method enables vast amounts of scientific discovery in various nuclear structure and nuclear astrophysics applications.

- **GFMC/CCM geared towards state-specific properties**
- **AFMC appropriate for thermal properties**

Peter Boyle

University of Edinburgh, Edinburgh, UK and Columbia University, New York, USA
e-mail: pab@phys.columbia.edu

0.1 Discussion Session Summary

Some slides were presented as a stimulus for discussion. The first two slides were intended to summarise the broad classes of problem that would suit the characteristics of QCDOC.

Several "anecdotal" slides of optimisation for QCDOC were presented, based on experiences tuning QCD code in the VHDL simulation of QCDOC. The dominant factors that should be considered are cache reuse and the memory access pattern.

Some code examples were presented of the O/S communications "SCU" calls, which enable internode communication and global summation from high level languages.

Several questions were asked.

- It was confirmed the QCDOC can carry out communication on multiple wires simultaneously, and concurrently with computation.
- It was queried why the DAXPY vector operation only obtains 190 Mflops. The reason for this is that there are two loads and one store for every "fmadd" instruction. This means that 33% usage of the FPU is the upper limit achievable on this instruction mix, and the subroutine is load/store bound. On top of this it was found to be beneficial to prefetch data into cache, and one prefetch must be issued for every four loads, and "eats" 2 cache access cycles. This leaves 2/8 load store slots "wasted", and hence the 190Mflops was obtained from a maximum *for this instruction mix* of 250Mflops, and relatively little effort has been spent on this kernel.
- A technical question on the memory protection hardware was asked. Using a simple run kernel will allow both memory protection against "bad" pointers, and simple DMA from user arrays. The 64 entry TLB (memory translation hardware) in the 440 CPU is sufficiently full featured that the entire memory space can be mapped without taking TLB misses.
- The speaker was asked to compare clusters using a mesh of ethernet links to QCDOC. It was his opinion that they are architecturally very similar and will be pretty reasonable propositions for systems of a few

hundred nodes. However, the combination of lower hardware and software communication latency, better bus bandwidth to the communications hardware, the lower power consumption and hardware global sum support allows QCDOC to be a practical proposition with 20-100 times more nodes.

Parallelisation Schemes

- Geometrical decomposition in some coordinate system
- Both fairly local and load balanced in that coordinate system
  $\Rightarrow$ Ideally, mesh should be continuously deformable into a N-dim grid
- Coordinate space physics involves local operators - ∂_μ, ∇^2

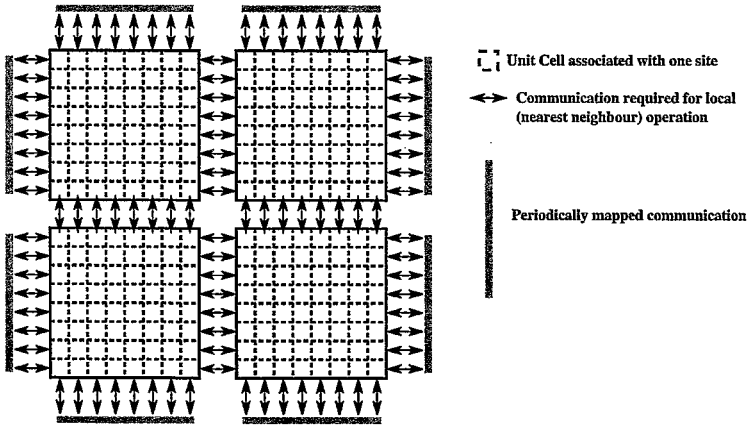

- Other schemes possible but comms may not map trivially to the hardware

Possible application areas

- Fluid Dynamics, Solid State, MHD/Plasma, Field theory
- FFT's & Long range forces (Galaxy dynamics, MD) open issues*
- Mesh depends on boundary conditions and source terms
  Arbitrary tesselation of the latest SUV may not map well
- Iterative sparse matrix solvers require
  Fast coordinate space matrix multiply + Gsum
- Other algorithms certainly possible - need either
  i) Locality
  Or
  ii) Good compute to communicate ratio

*Lesson: dont ignore gauge field dynamics!

QCDOC On-Chip memory system

- 4MB On-chip Edram: 16GB/s raw bandwidth
- Custom multi-ported prefetching Edram controller, IBM Yorktown
  Two read streams on each port
  Writes merged in 128byte buffers
  raw 8GB/s CPU Read + 8GB/s Write
  raw 2.6GB/s PLB Read + 2.6GB/s Write + DMA

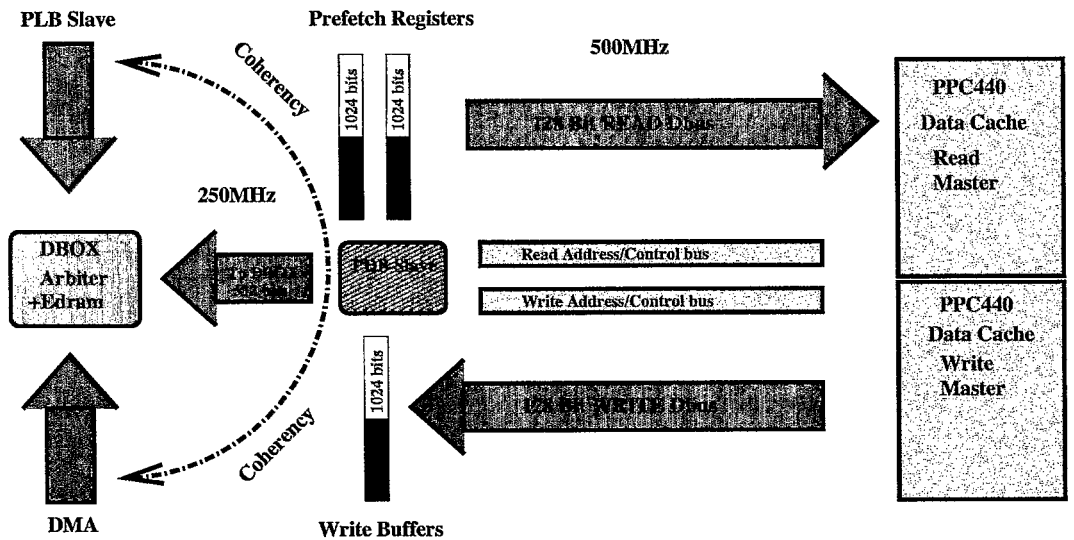

Optimisation Anecdotes

32 FP reg OoO pipelined super-scalar risc: 1 FMADD + 1 L/S
Flop scheduling pretty easy (FP Latency 5).
L/S slots tight on many kernels
Cache allows 3 outstanding fetches of 32 byte lines
Queue `dcbt` every 6 cycles *while* keeping FPU fed *without* stalling
Each `dcbt` eats 2 cache access cycles
Very high fractions of peak achievable on many kernels

- Use the Edram
- **Do** cache prefetch: xlc built-in `dcbt(void *)`
- **Don't** use `dcbz` - rely on store gathering
- Tune memory layout:
  128byte alignment for PEC + padding inner indices
- Tune memory access order for linear read access
  Scatter not Gather (also Pentium4, Athlon, Power3/4, T3d/e)
- Characteristics similar to Alpha's running from L2 cache

- Performance characteristics at small local volumes are critical to scaling.
- We show estimated performance per node as a function of local volume.
- Sparse matrix application takes as little as 20 μs, with eight communications in this time.

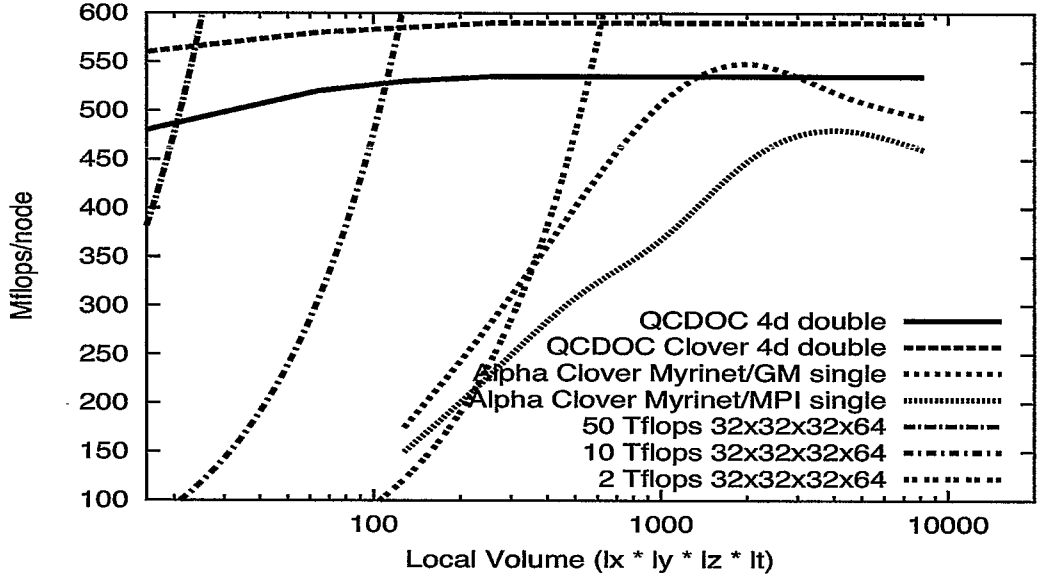

Contours of constant aggregate performance on a $32^3 \times 64$ lattice are superimposed.

```
double U[Npsite][Ncolor][Ncolor][Ncmplx];
double Chi[Npsite][N2spin][Ncolor][Ncmplx];
double Phi[Npsite][N2spin][Ncolor][Ncmplx];
int idx[Npsite];

Phi[site] = U[idx[site]] * Chi[idx[site]]
```


Optimize as

```
#define Npad 1

double U[Npsite][Ncolor][Ncolor][Ncmplx];
double Chi[Npsite][N2spin][Ncolor+Npad][Ncmplx] PEC_ALIGN;
double Phi[Npsite][N2spin][Ncolor+Npad][Ncmplx] PEC_ALIGN;
int idxprime[Npsite];

Phi[idxprime[site]] = U[site] * Chi[site]
```

Simulator performance estimates

"Nominal" 500MHz.
All figures include communication and cache flush overhead as appropriate.
Pin-pin wire delay 20 ns = 3 m @ speed of light in medium.

Operation	Local Volume	Performance/node (Mflops)
Wilson D_{eo}	2^4	470
Wilson D_{eo}	4^4	535
Clover D_{eo}	2^4	560
Clover D_{eo}	4^4	590
Staggered(Cristian) D_{eo}	2^4	370
Staggered D_{eo}	$2^2.4^2$	430
SU3-SU3	-	800
SU3-2spinor	-	780
DAXPY	-	190
ZAXPY	-	450
DAXPY-Norm	-	350
CloverTerm/asm	-	790
CloverTerm/gcc, no dcbt	-	150
CloverTerm/xlc, no dcbt	-	300

Scalability

Global sum
4k nodes: 10μs
16k nodes: 13μs
32k nodes: 15μs

Estimate for Wilson CG on $32^3 \times 64$ Lattice

Nodes	$M^\dagger M$+linalg	Global Sum	Sust. Tflops
4096	2620μs	10 μs	2.15
8192	1310μs	11.5μs	4.2
16536	680 μs	13 μs	8.1
32768	340 μs	15 μs	15.6

Both Clover and DWF *more* scalable!

Nearest Neighbour code example

```
SCUDirArgIR DirArg_send_f[4];
SCUDirArgIR DirArg_recv_f[4];
SCUDirArgIR *DA_f_p[8];
SCUDirArgMulti wfm_multi_f;

static const SCUDir  plus_dirs[] = { SCU_XP, SCU_YP, SCU_ZP, SCU_TP } ;
static const SCUDir minus_dirs[] = { SCU_XM, SCU_YM, SCU_ZM, SCU_TM } ;
...
/*Forwards xfer send buffer*/
DirArg_send_f[mu].Init((void *)wp->send_f[mu], plus_dirs[mu], SCU_SEND,
                       (HALF_SPINOR_SIZE*sizeof(Float)), wp->nbound[mu],
                       ((PAD_HALF_SPINOR_SIZE-HALF_SPINOR_SIZE )*sizeof(Float))
                       );

wfm_multi_f.Init(DA_f_p,8);
...
void wfm_comm_forward_start()
{
  wfm_multi_f.StartTrans();
}
void wfm_comm_forward_complete()
{
  wfm_multi_f.TransComplete();
}
```

Global sum code examples

```
#include <gsum64.h>

  double SpaceSum, SpaceTimeSum, SixDSum;
  double MyVal;

/*3d*/
  SCUAxis reduce_xyz [] = { SCU_X,SCU_Y,SCU_Z };
  Gsum64  Gsum_xyz (reduce_xyz ,3);
  SpaceSum =  Gsum_xyz.Sum(MyVal);

/*4d*/
  SCUAxis reduce_xyzt[] = { SCU_X,SCU_Z,SCU_Y,SCU_T };
  Gsum64  Gsum_xyzt(reduce_xyzt,4);
  SpaceTimeSum = Gsum_xyzt.Sum(MyVal);

/*6d*/
  Gsum64  Gsum_xyztsw;
  SixDSum = Gsum_xyztsw.Sum(MyVal);
```

Large-Scale Shell-Model Calculations and Alphleet Project

Noritaka Shimizu RIKEN

Takaharu Otsuka Univ. of Tokyo/RIKEN

Takahiro Mizusaki Senshu Univ.
Michio Honma Aizu Univ.
Yutaka Utsuno JAERI

Susumu Shimoura CNS
Yasushi Watanabe RIKEN
Eiji Ikezawa RIKEN
Shigeo Numata RIKEN
Masatake Henmi RIKEN
Syuuji Kase RIKEN

Alphleet-2 since 2002

Compaq ES-45(4cpu) x 13
Myrinet network

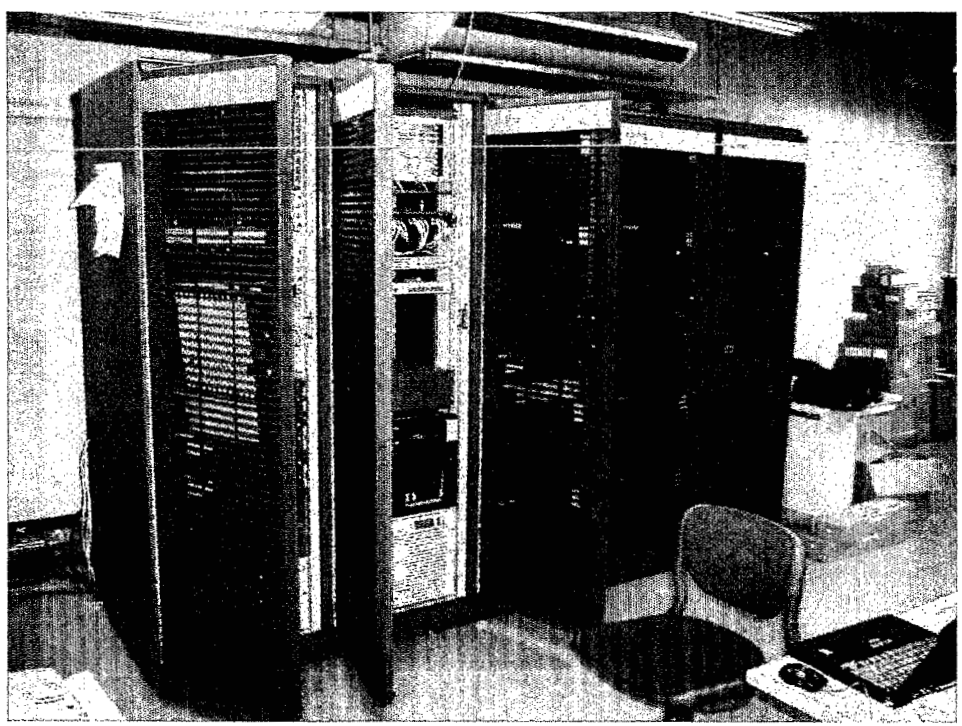

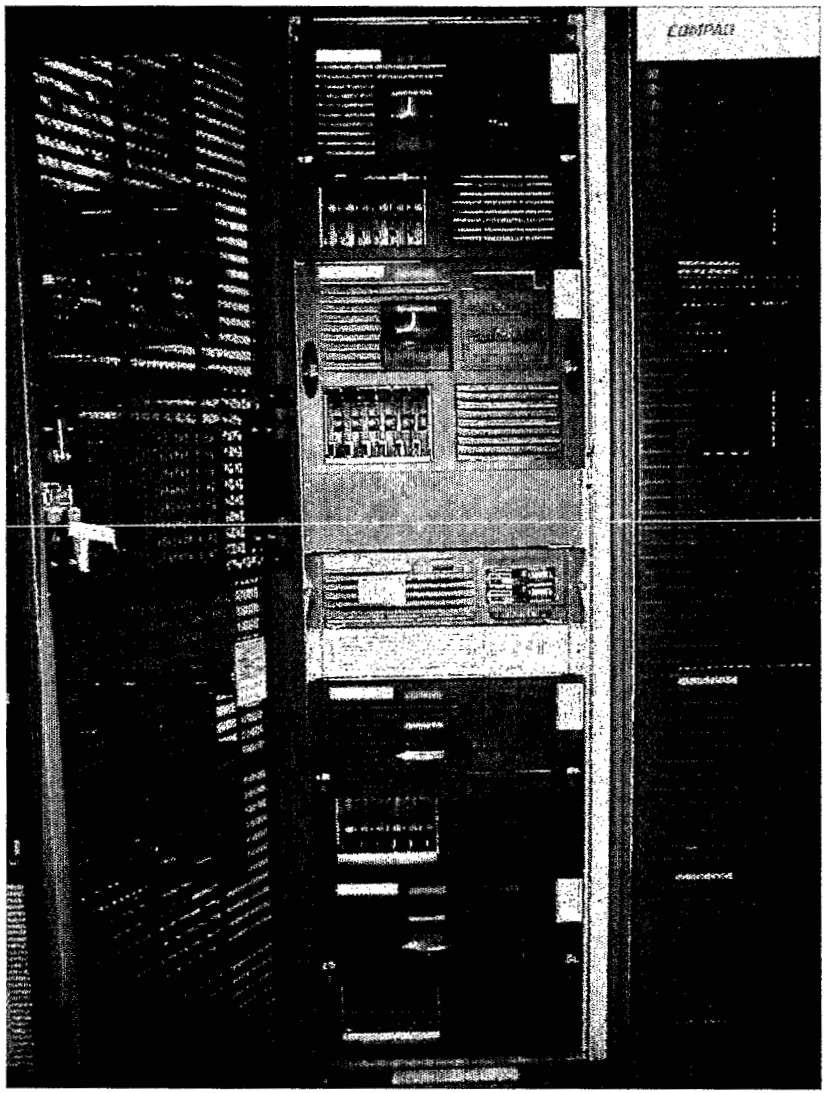

Comparison between conventional and Monte-Carlo shell-model calculations

- Conventional method

all components (Slater determinants)

$H =$ 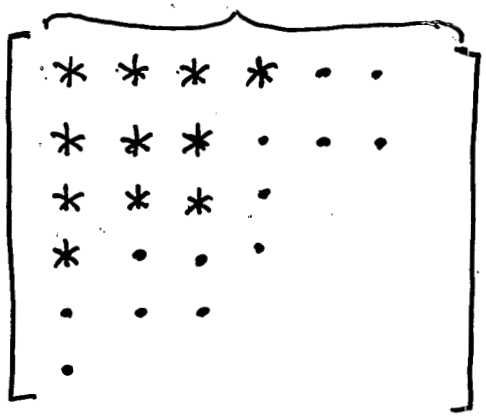 diagonal-ize 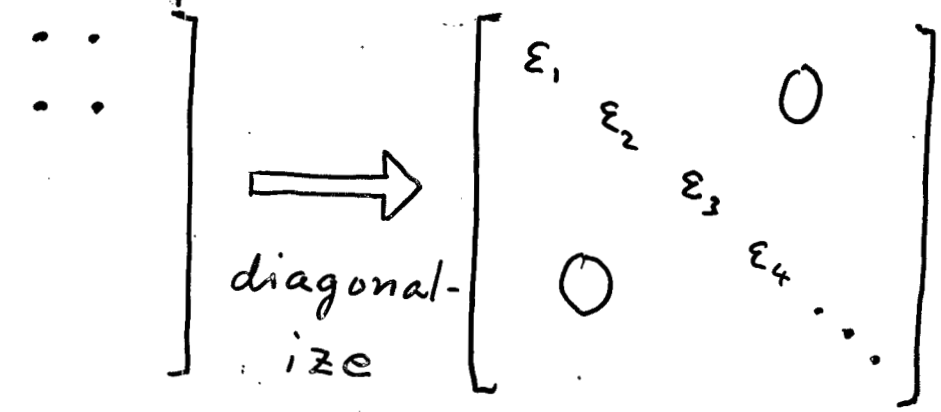

dimension 10^4 (^{24}Mg)

$\sim 10^9$ (^{32}Mg) $\sim \cdots$

- MCSM : importance truncation

$\psi \rightarrow e^{-\beta h(\sigma)} \psi$

Slater det. Slater det.

$H =$ 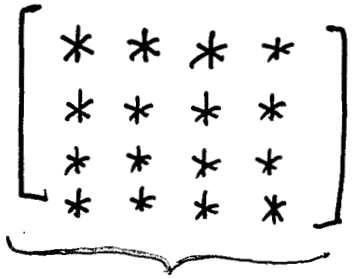 diagonalize 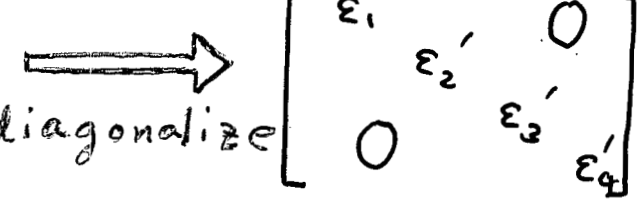

selected components

$\varepsilon_i' \approx \varepsilon_i$

Developments of Shell Model calculations vs. Evolution of Computers

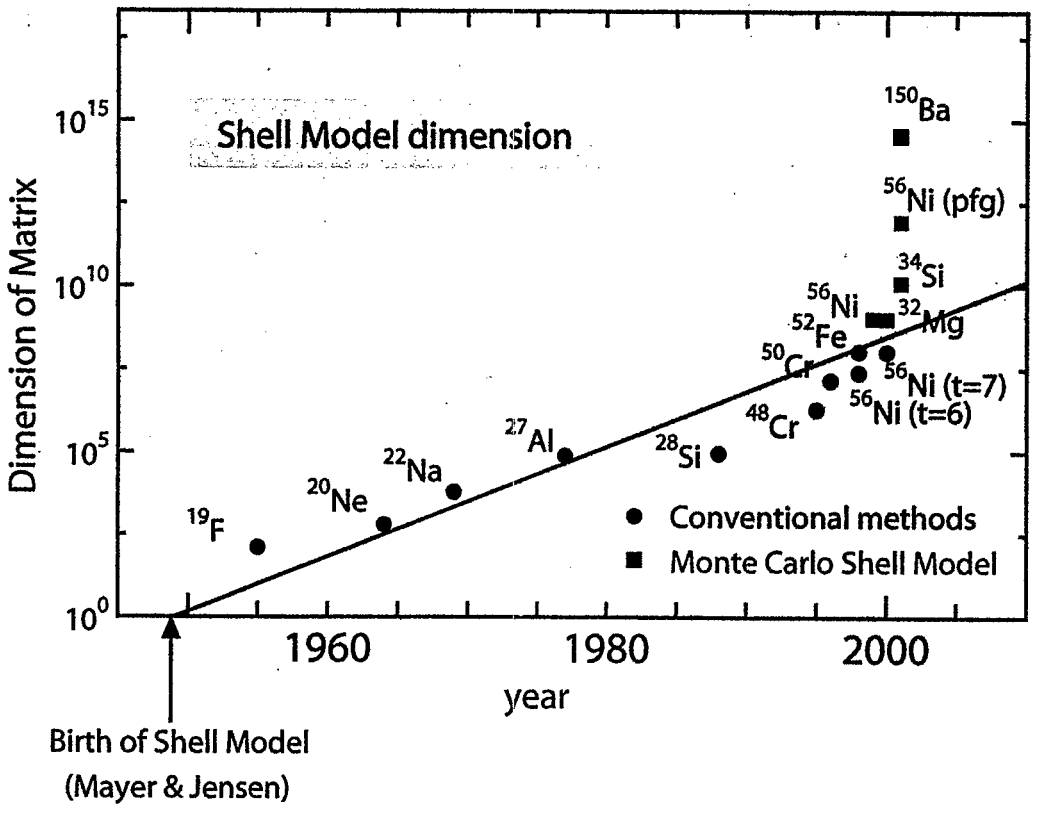

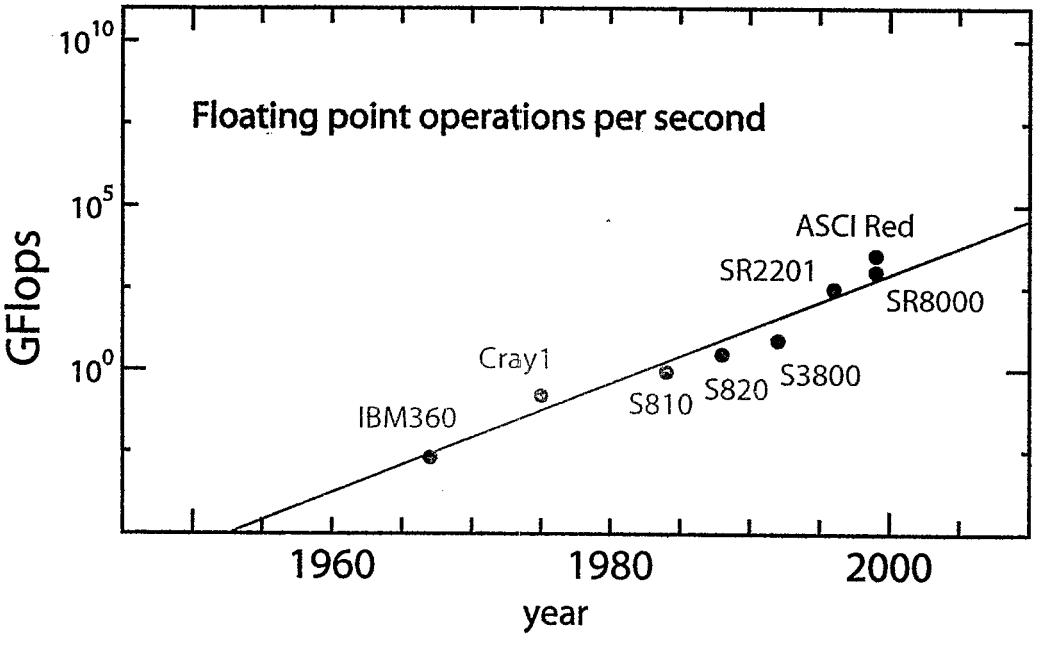

Lines of both figures indicate a growth of 10^{5}/30 yeas

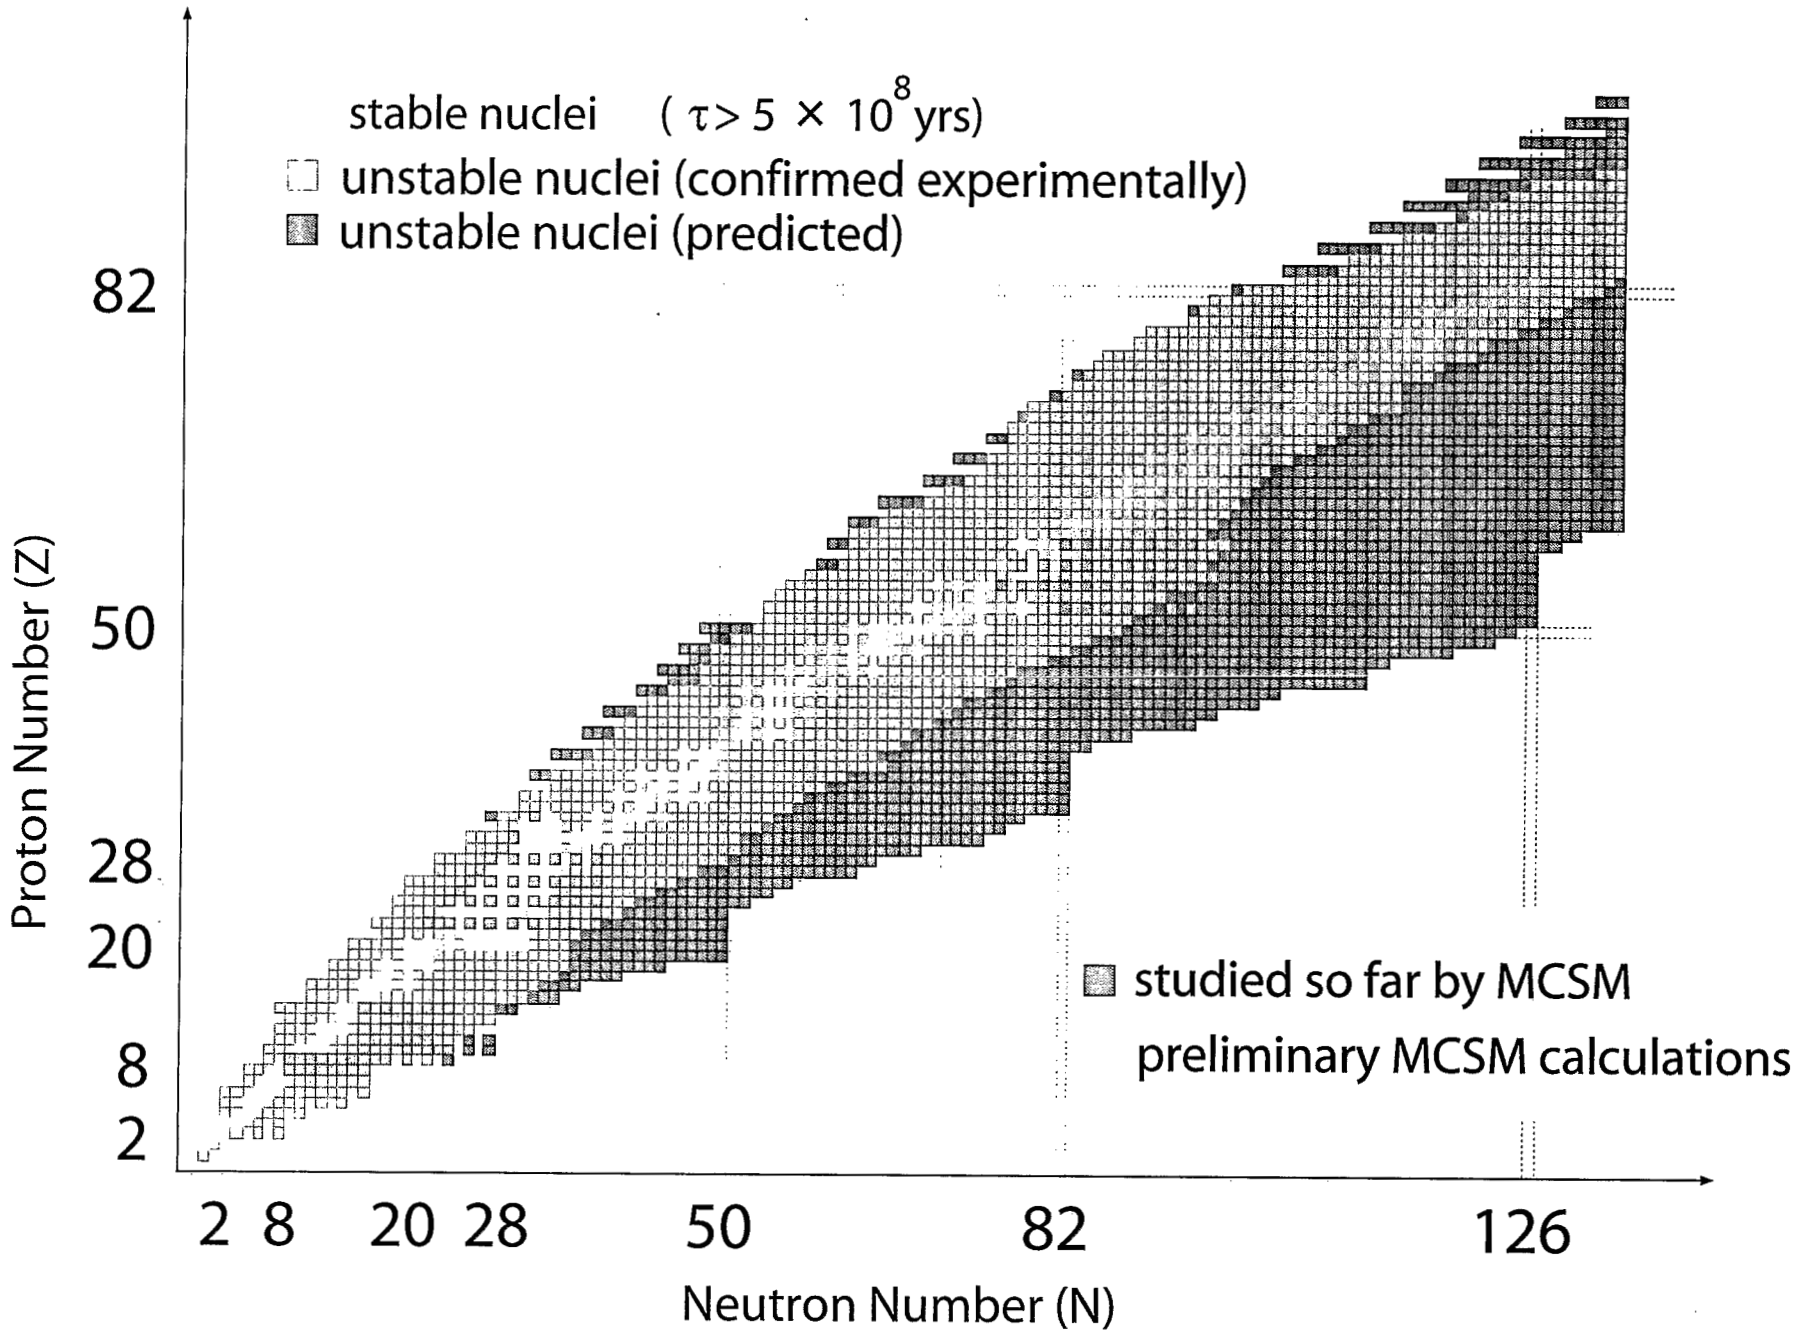
stable nuclei (τ> 5 × 10^8 yrs)
unstable nuclei (confirmed experimentally)
unstable nuclei (predicted)
studied so far by MCSM
preliminary MCSM calculations
Proton Number (Z)
Neutron Number (N)
2
8
20
28
50
82
126

Summary

MCSM (Monte Carlo Shell Model)

20~50 bases with projection onto J and M

against the dimension of the original Hilbert space : 10^4 ~ 10^9 ~ 10^{14} ~

^{24}Mg ^{32}Mg ^{56}Ni rare earth

diagonalization of Hamiltonian ⇒

levels, wave functions,
transition matrix elements

MCSM : *Importance truncation* scheme to the full calculation

- many valence orbits

  more than one major shell

- many valence particles

  so far up to 26

- **natural relation to mean field theories** (HF, GCM, HFB, *etc.*)

- no minus sign problem

High-Performance Computing in Modeling of Ultra-Relativistic Heavy-Ion Collisions

Steffen A. Bass

Duke University &
RIKEN BNL Research Center

- Overview: Transport-Theory at RHIC
- String/Hadron Transport Models
- The PCM: Fundamentals & Computational Requirements
- Nuclear Fluid Dynamics
- Summary and Outlook

Transport Theory at RHIC

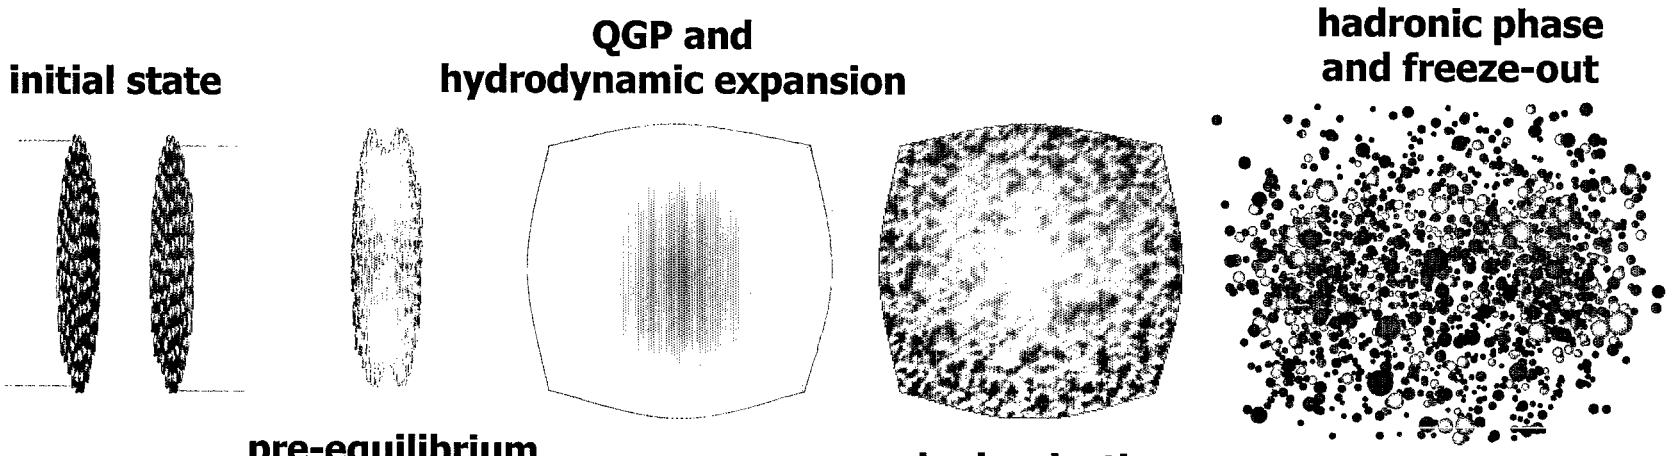

CYM & LGT

PCM & clust. hadronization

NFD

NFD & hadronic TM

string & hadronic TM

PCM & hadronic TM

String/Hadron Transport Models

- elementary degrees of freedom: hadrons, const. (di)quarks
- classical trajectories in phase-space (relativistic kinematics)
- initial high energy phase of the reaction is modeled via the excitation and fragmentation of strings
- 55 baryon- and 32 meson species, among those 25 N*, Δ* resonances and 29 hyperon/hyperon resonance species
- full baryon-antibaryon and isospin symmetry

➢ ideal for the description of excited hadronic matter

- main physics input and parameters:
  - **cross sections**: total and partial cross sections, angular distributions
  - **resonance parameters**: total and partial decay widths
  - **string fragmentation scheme**: fragmentation functions, formation time

- An interaction takes place if at the time of closes approach d_{min} of two hadrons the following condition is fulfilled:

$$d_{\min} = \sqrt{\frac{\sigma_{tot}}{\pi}} \quad \text{with} \quad \sigma_{tot} = \sigma_{tot}(\sqrt{s}, |h_1\rangle, |h_2\rangle$$

Computational Requirements in String/Hadron Models

Memory:

- Depends on collision system and beam energy
- Depends on particle information stored (beyond 6N+1 phase-space)
- RHIC: less than 50 MB

CPU:

- collision finder dominates CPU needs
- initial run scales with N^2, updates with N log(N)
- RHIC: 80 minutes per event (on AMD Athlon 1.5 GHz)

Storage:

- RHIC: 0.5 – 1.0 MB per event (final output, compressed ASCII)

Scalability:

- trivial parallelization – can utilize broad range of cpu's

Overview of PCM Implementations I

- tools for theoretical studies:

  ZPC:
  - simplified initial conditions
  - elastic gg scattering with screening
  - no initial/final state radiation
  - no hadronization

  MPC:
  - same features as ZPC
  - includes $2 \leftrightarrow 3$ processes

- HIJING:
  - initial conditions from PDFs
  - includes initial/final state radiation
  - no space-time evolution (not a PCM!)
  - ➢ predicts # of hard scattered partons utilizing Glauber geometry
  - ➢ can be used to provide *initial* conditions for PCMs

Overview of PCM Implementations II

- PCMs for modeling RHIC and LHC heavy ion physics:

AMPT:
- HIJING as initial condition to generate parton distribution
- PCM evolution utilizing ZPC
- hadronization via HIJING fragmentation scheme

GROMIT:
- new multi-purpose framework for RHIC transport calculations
- HIJING as initial condition to generate parton distribution

VNI:
- initial conditions from PDFs
- full space-time evolution of gluons, sea- and valence quarks
- pQCD matrix elements for gg, qg, qq scattering
- initial/final state radiation
- cluster-hadronization model
  (not maintained after passing away of K. Geiger)

VNI/BMS:
- based on VNI, developed by SAB, BM & DKS
- extensive rewrite of initialization, collision & shower generation
- provisions for saturation, screening & initial state scale corrections

Basic Principles of the PCM

- degrees of freedom: quarks and gluons
- classical trajectories in phase space (with relativistic kinematics)
- initial state constructed from experimentally measured nucleon structure functions and elastic form factors
- an interaction takes place if at the time of closest approach d_{min} of two partons

$$d_{\min} \leq \sqrt{\frac{\sigma_{tot}}{\pi}} \quad \text{with} \quad \sigma_{tot} = \sum_{p_3,p_4} \int \frac{d\sigma\left(\sqrt{\hat{s}}; p_1, p_2, p_3, p_4\right)}{d\hat{t}} d\hat{t}$$

- system evolves through a sequence of binary (2→2) elastic and inelastic scatterings of partons and initial and final state radiations within a leading-logarithmic approximation (2→N)
- binary cross sections are calculated in leading order pQCD with either a momentum cut-off or Debye screening to regularize IR behaviour
- guiding scales: initialization scale Q_0, p_T cut-off p_0 / Debye-mass μ_D, intrinsic k_T / saturation momentum Q_s, virtuality > μ_0

PCM Computational Requirements

Memory:

- Depends on collision system and beam energy
- Depends on particle information stored (beyond 6N+1 phase-space)
- RHIC: approx. 100-150 MB / LHC: approx. ?? MB

CPU:

- collision finder & cross section integrator dominates CPU needs
- initial run scales with N^2, updates with N log(N)
- RHIC: 80 - 240 minutes per event (on AMD Athlon 1.5 GHz)

Storage:

- RHIC: 0.5 – 1.0 MB per event (final output, compressed ASCII)

Scalability:

- trivial parallelization – can utilize broad range of cpu's

Nuclear Fluid Dynamics

- transport of macroscopic degrees of freedom
- based on conservation laws: $\partial_\mu T^{\mu\nu}=0$ $\partial_\mu j^\mu=0$

- for ideal fluid: $T^{\mu\nu}= (\varepsilon+p)\, u^\mu\, u^\nu - p\, g^{\mu\nu}$ and $j_i^\mu = \rho_i\, u^\mu$
- **Equation of State** needed to close system of PDE's: $p=p(T,\rho_i)$

- assume local thermal equilibrium
- initial conditions (i.e. thermalized QGP) required for calculation

- simple case: scaling hydrodynamics
  - assume longitudinal boost-invariance
  - cylindrically symmetric transverse expansion
  - no pressure between rapidity *slices*
  - conserved charge in each *slice*

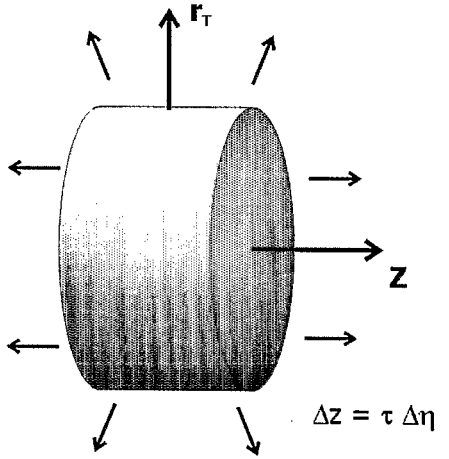

NFD Computational Requirements

Memory:

- momenta, densities and conserved quantities need to be stored in a large coordinate space grid
- Depends on utilized symmetries (1+1 boost-invariant vs. 3+1D)
- RHIC: at least 500 MB

CPU:

- integrator dominates CPU needs (i.e. time step width)
- modest needs: 2-4 hours at RHIC

Storage:

- RHIC: 1-2 GB (final output of hypersurface, compressed ASCII)

Scalability:

- difficult – usually not necessary in heavy-ion context

Summary of Computational Requirements

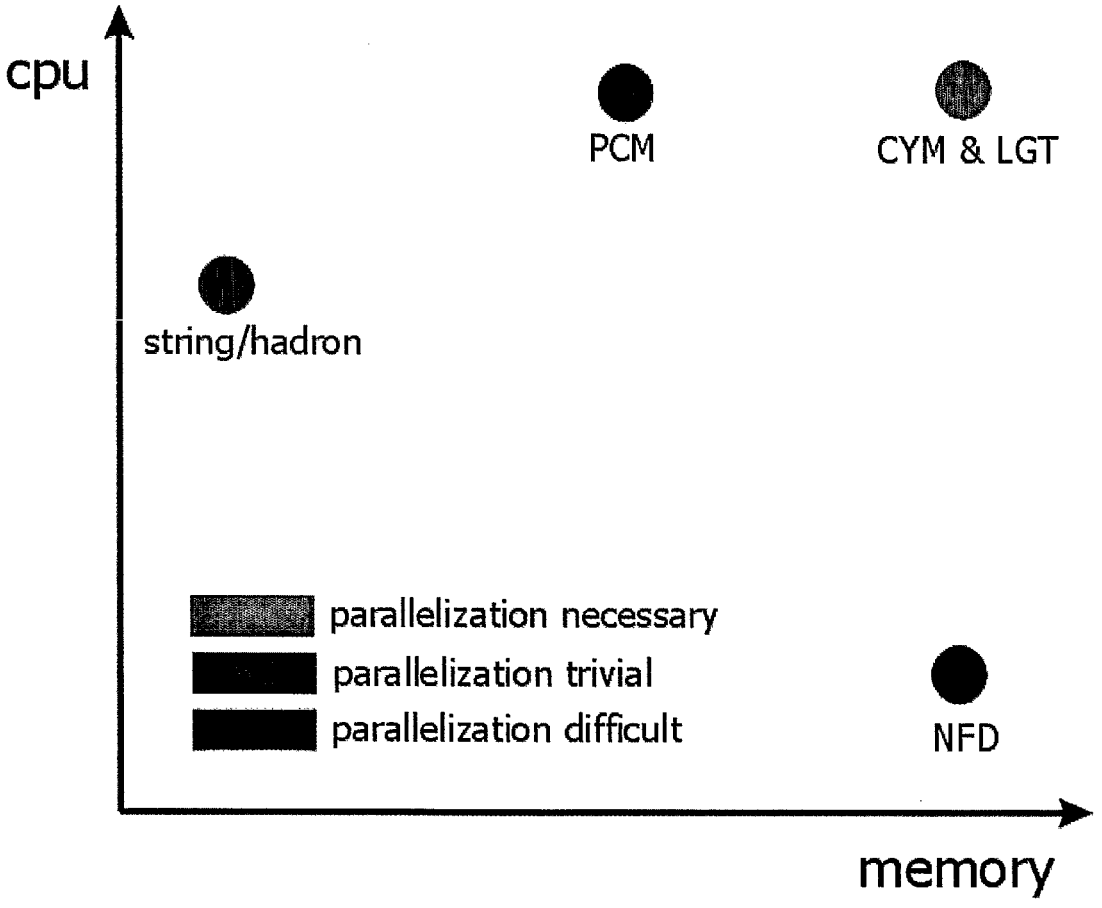

- most microscopic approaches already benefit from trivial parallelization, i.e. heterogeneous clusters
- novel transport schemes with field degrees of freedom require truly parallel processing

Status & Future Directions ...

Currently, Linux based cpu-farms provide a flexible and cost-effective computing environment for most transport applications

- Problem: not enough resources available (clusters/availability of access) for large scale systematic studies and detailed comparison to data

Future transport approaches will treat field degrees of freedom and microscopic particles simultaneously with equal sophistication and will require algorithms combining lattice/grid technology with microscopic features (collision finders, Monte-Carlo integration)

- need parallel processing capabilities!

Large-Scale Computations in Nuclear Physics using the QCDOC
RIKEN BNL Research Center Workshop
September 26-28, 2002

Organizers: Yasumichi Aoki, Anthony Baltz, Michael Creutz, Miklos Gyulassy & Shigemi Ohta

PARTICIPANTS

Name	Mailing Address	E-Mail Address
Yasumichi Aoki	Physics – 510A Brookhaven National Laboratory Upton, NY 11973	yaoki@bnl.gov
Anthony Baltz	Physics – 510A Brookhaven National Laboratory Upton, NY 11973	baltz@bnl.gov
Steffen Bass	Physics Department Duke University Durham, NC 27708-0305	bass@bnl.gov or bass@phy.duke.edu
Robert Bennett	ITD - 515 Brookhaven National Laboratory Upton, NY 11973	robertb@bnl.gov
Tom Blum	Physics – 510A Brookhaven National Laboratory Upton, NY 11973	tblum@bnl.gov
Peter Boyle	Department of Physics Columbia University New York, NY 10027	pab@phys.columbia.edu
Shailesh Chandrasekharan	Box 90305, Physics Department Duke University Durham, NC 27708-0305	sch@phy.duke.edu
Michael Creutz	Physics – 510A Brookhaven National Laboratory Upton, NY 11973	creutz@bnl.gov
James Davenport	ALD – 463B Brookhaven National Laboratory Upton, NY 11973	jdaven@bnl.gov
Chris Dawson	Physics – 510A Brookhaven National Laboratory Upton, NY 11973	cdawson@bnl.gov
David Dean	Physics, Box 2008, Bldg 6011, MS 6373 Oak Ridge National Laboratory Oak Ridge, TN 37831-6373	deandj@ornl.gov
Yuefan Deng	Applied Math Dept. SUNY, Stony Brook Stony Brook, NY 11794	deng@ams.sunysb.edu

Valeriya Gadiyak	Department of Physics Columbia University New York, NY 10027	gadiyak@theory1.phys.columbia.edu
Urs Heller	American Physical Society 4 Research Road Ridge, NY 11961	heller@aps.org
Taku Izubuchi	Physics – 510A Brookhaven National Laboratory Upton, NY 11973	izubuchi@bnl.gov
Chulwoo Jung	Physics – 510A Brookhaven National Laboratory Upton, NY 11973	chulwoo@phys.columbia.edu
Thomas Kirk	Physics – 510F Brookhaven National Laboratory Upton, NY 11973	tkirk@bnl.gov
Alexander Krasnitz	CENTRA and FCT Universidade do Algarve Campus de gambelas P-8000 Faro, Portugal	krasnitz@ualg.pt
James Lattimer	Dept. of Physics & Astronomy SUNY, Stony Brook Stony Brook, NY 11794	lattimer@astro.sunysb.edu
T.-D. Lee	Physics – 510A Brookhaven National Laboratory Upton, NY 11973	tdl@cuphyb.phys.columbia.edu
Robert Mawhinney	Department of Physics Columbia University New York, NY 10027	rdm@physics.columbia.edu
Ed McFadden	ITD - 515 Brookhaven National Laboratory Upton, NY 11973	emc@bnl.gov
Michael McGuigan	ITD - 515 Brookhaven National Laboratory Upton, NY 11973	mcguigan@bnl.gov
Larry McLerran	Physics – 510A Brookhaven National Laboratory Upton, NY 11973	mclerran@bnl.gov
Denes Molnar	Ohio State University 178 West 18th Ave Columbus, OH 43210	molnard@phys.columbia.edu
Eric Myra	Dept. Physics & Astronomy SUNY, Stony Brook Stony Brook, NY 11794	emyra@mail.astro.sunysb.edu

Jun-Ichi Noaki	Physics – 510A Brookhaven National Laboratory Upton, NY 11973	noaki@rccp.tsukuba.ac.jp
Shigemi Ohta	Physics – 510A Brookhaven National Laboratory Upton, NY 11973	ohta@bnl.gov
Kostas Orginos	Physics – 510A Brookhaven National Laboratory Upton, NY 11973	orginos@bnl.gov
Taka Otsuka	Department of Physics University of Tokyo Hongo, Bunkyo-ku Tokyo 113-0033, Japan	otsuka@phys.s.u-tokyo.ac.jp
Konstantin Petrov	Physics – 510A Brookhaven National Laboratory Upton, NY 11973	petrov@bnl.gov
Scott Pratt	4208A Biophysical Science Bldg. Michigan State University East Lansing, MI 48823	pratts@pa.msu.edu
Eunice Santos	Dept of Computer Science Virginia Polytechnic Institute Blacksburg, VA 24061	santos@cs.vt.edu
Thomas Schaefer	Physics – 510A Brookhaven National Laboratory Upton, NY 11973	tschaefer@bnl.gov
Edward Shuryak	Dept. Physics & Astronomy SUNY, Stony Brook Stony Brook, NY 11794	shuryak@tonic.physics.sunysb.edu
Noritaka Shimizu	RI-Beam Science Lab RIKEN 2-1, Hirosawa, Wako-shi Saitama, 351-0198, Japan	nshimizu@postman.riken.go.jp
Dave Stampf	ITD - 515 Brookhaven National Laboratory Upton, NY 11973	drs@bnl.gov
Douglas Swesty	Dept. Physics & Astronomy SUNY, Stony Brook Stony Brook, NY 11794	douglas.swesty@stonybrook.edu
Raju Venugopalan	Physics – 510A Brookhaven National Laboratory Upton, NY 11973	raju@bnl.gov
Tilo Wettig	Physics – 510A Brookhaven National Laboratory Upton, NY 11973	tilo.wettig@yale.edu

Large-Scale Computations in Nuclear Physics using the QCDOC

A RIKEN BNL Research Center Workshop
Physics Department, Bldg. 510, Room 2-160
September 26-28, 2002

Organizers: Yasumichi Aoki, Anthony Baltz, Michael Creutz, Miklos Gyulassy, and Shigemi Ohta

AGENDA

~ ~ ~

THURSDAY, SEPTEMBER 26, 2002

9:00 Registration & Coffee

Morning Session Chair – Anthony Baltz, BNL and RBRC

9:30 Thomas Kirk, BNL ..... *Welcome*

9:40 Robert Mawhinney, Columbia Univ ..... *An Introduction to QCDOC*

10:40 Scott Pratt, Michigan State Univ ..... *Potential CPU Needs for RHIC Modeling*

11:30 Douglas Swesty, SUNY SB ..... *Multi-D Hydrodynamics and Radiation Transport in Core Collapse Supernova Simulations*

12:20 Lunch

Afternoon Session Chair – Michael Creutz, BNL

1:30 Alexander Krasnitz, Algarve Univ ..... *Classical Numerical Gluodynamics of Nuclear Collisions*

2:20 James Davenport, BNL ..... *Classical Molecular Dynamics on QCDOC*

3:10 Coffee

3:30 Robert Mawhinney, Columbia Univ ..... *Introduction to QCDOC II: Discussion & Feedback*

6:00 Workshop Dinner

FRIDAY, SEPTEMBER 27, 2002

Morning Session Chair – Chulwoo Jung, BNL

9:30 Michael Creutz, BNL ..... *Experience Using the QCDSP for Other than Lattice Applications*

10:20 Shailesh Chandrasekharan, Duke Univ ..... *Cluster Algorithms for Strongly Coupled Lattice Gauge Theories*

11:10 Denes Molnar, Ohio State Univ ..... *Solving Nonlinear Dynamical Equations with QCDOC Horsepower*

12:00 Lunch

Afternoon Session Chair – Kostas Orginos, RBRC

1:30 Tour of QCDSP

2:00 David Dean, ORNL ..... *The Computational Many-Body Problem*

2:50 Thomas Schaefer, SUNY SB and RBRC ..... *Short Contribution*

Additional Short Contributions

3:20 Coffee

3:30 Peter Boyle, Univ Edinburgh & Columbia Univ ..... *Discussion & Feedback*

SATURDAY, SEPTEMBER 28, 2002

Session Chair - Shigemi Ohta, KEK and RBRC

8:30 Continental Breakfast

9:30 Takaharu Otsuka, Univ of Tokyo and Noritaka Shimizu, RIKEN, Japan ..... *Large-Scale Shell Model Calculations for Atomic Nuclei and Alphleet Project*

10:40 Steffen Bass, Duke Univ & RBRC ..... *Heavy Ion Transport*

11:15 Yasumichi Aoki, RBRC ..... *Discussion & Feedback*

12:00 Adjourn

Additional RIKEN BNL Research Center Proceedings:

Volume 45 – Summer Program: Current and Future Directions at RHIC – BNL-
Volume 44 – RHIC Spin Collaboration Meetings VIII, IX, X, XI – BNL-
Volume 43 – RIKEN Winter School – Quark-Gluon Structure of the Nucleon and QCD – BNL-52672
Volume 42 – Baryon Dynamics at RHIC – BNL-52669
Volume 41 – Hadron Structure from Lattice QCD – BNL-52672
Volume 40 – Theory Studies for RHIC-Spin – BNL-52662
Volume 39 – RHIC Spin Collaboration Meeting VII – BNL-52659
Volume 38 – RBRC Scientific Review Committee Meeting – BNL-52649
Volume 37 – RHIC Spin Collaboration Meeting VI (Part 2) – BNL-52660
Volume 36 – RHIC Spin Collaboration Meeting VI – BNL-52642
Volume 35 – RIKEN Winter School – Quarks, Hadrons and Nuclei – QCD Hard Processes and the Nucleon Spin – BNL-52643
Volume 34 – High Energy QCD: Beyond the Pomeron – BNL-52641
Volume 33 – Spin Physics at RHIC in Year-1 and Beyond – BNL-52635
Volume 32 – RHIC Spin Physics V – BNL-52628
Volume 31 – RHIC Spin Physics III & IV Polarized Partons at High Q^2 Region – BNL-52617
Volume 30 – RBRC Scientific Review Committee Meeting – BNL-52603
Volume 29 – Future Transversity Measurements – BNL-52612
Volume 28 – Equilibrium & Non-Equilibrium Aspects of Hot, Dense QCD – BNL-52613
Volume 27 – Predictions and Uncertainties for RHIC Spin Physics & Event Generator for RHIC Spin Physics III – Towards Precision Spin Physics at RHIC – BNL-52596
Volume 26 – Circum-Pan-Pacific RIKEN Symposium on High Energy Spin Physics – BNL-52588
Volume 25 – RHIC Spin – BNL-52581
Volume 24 – Physics Society of Japan Biannual Meeting Symposium on QCD Physics at RIKEN BNL Research Center – BNL-52578
Volume 23 – Coulomb and Pion-Asymmetry Polarimetry and Hadronic Spin Dependence at RHIC Energies – BNL-52589
Volume 22 – OSCAR II: Predictions for RHIC – BNL-52591
Volume 21 – RBRC Scientific Review Committee Meeting – BNL-52568
Volume 20 – Gauge-Invariant Variables in Gauge Theories – BNL-52590
Volume 19 – Numerical Algorithms at Non-Zero Chemical Potential – BNL-52573
Volume 18 – Event Generator for RHIC Spin Physics – BNL-52571
Volume 17 – Hard Parton Physics in High-Energy Nuclear Collisions – BNL-52574
Volume 16 – RIKEN Winter School - Structure of Hadrons - Introduction to QCD Hard Processes – BNL-52569
Volume 15 – QCD Phase Transitions – BNL-52561

Additional RIKEN BNL Research Center Proceedings:

Volume 14 – Quantum Fields In and Out of Equilibrium – BNL-52560
Volume 13 – Physics of the 1 Teraflop RIKEN-BNL-Columbia QCD Project First Anniversary Celebration – BNL-66299
Volume 12 – Quarkonium Production in Relativistic Nuclear Collisions – BNL-52559
Volume 11 – Event Generator for RHIC Spin Physics – BNL-66116
Volume 10 – Physics of Polarimetry at RHIC – BNL-65926
Volume 9 – High Density Matter in AGS, SPS and RHIC Collisions – BNL-65762
Volume 8 – Fermion Frontiers in Vector Lattice Gauge Theories – BNL-65634
Volume 7 – RHIC Spin Physics – BNL-65615
Volume 6 – Quarks and Gluons in the Nucleon – BNL-65234
Volume 5 – Color Superconductivity, Instantons and Parity (Non?)-Conservation at High Baryon Density – BNL-65105
Volume 4 – Inauguration Ceremony, September 22 and Non -Equilibrium Many Body Dynamics – BNL-64912
Volume 3 – Hadron Spin-Flip at RHIC Energies – BNL-64724
Volume 2 – Perturbative QCD as a Probe of Hadron Structure – BNL-64723
Volume 1 – Open Standards for Cascade Models for RHIC – BNL-64722

For information please contact:

Ms. Pamela Esposito
RIKEN BNL Research Center
Building 510A
Brookhaven National Laboratory
Upton, NY 11973-5000 USA

Phone: (631) 344-3097
Fax: (631) 344-4067
E-Mail: pesposit@bnl.gov

Ms. Tammy Heinz
RIKEN BNL Research Center
Building 510A
Brookhaven National Laboratory
Upton, NY 11973-5000 USA

(631) 344-5864
(631) 344-2562
theinz@bnl.gov

Homepage: http://quark.phy.bnl.gov/www/riken/index.html

RIKEN BNL RESEARCH CENTER

Large-Scale Computations in Nuclear Physics using the QCDOC

September 26 – 28, 2002

Li Keran

Nuclei as heavy as bulls
Through collision
Generate new states of matter.
T.D. Lee

Speakers:

S. Bass
J. Davenport
D. Molnar
D. Swesty
P. Boyle
D. Dean
T. Otsuka
S. Chandrasekharan
A. Krasnitz
S. Pratt
M. Creutz
R. Mawhinney
N. Shimizu

Organizers: Y. Aoki, A. Baltz, M. Creutz, M. Gyulassy and S. Ohta