

THE QCD TERAFLOPS PROJECT¹

STEVEN GOTTLIEB

*Physics Department, Brookhaven National Laboratory
Upton, Long Island, NY 11973, USA*

For the QCD Teraflops Collaboration

RECEIVED BY
JL ROOST
08 11 92
GOTT

ABSTRACT

Increased computer power is essential for future progress in lattice gauge theory and for other Grand Challenge applications. We address the physics that can be done with a computer capable of sustaining 1 Teraflops for QCD and the technology that will make it possible to construct such a computer within the next three years. Our collaboration has proposed to build a computer based on the Thinking Machines CM5 communication network, but with nodes 10 times faster.

1. Physics Goals

Lattice gauge theory was invented quite some time ago. It aspires to be a general purpose theoretical tool for dealing with nonperturbative quantum field theory. Quantum Chromodynamics (QCD) provides a testing ground for lattice techniques because many properties of the theory such as the spectrum of low lying hadrons are well determined by experiment. However, QCD is also an important application because there are a number of open questions such as the size of various weak matrix elements and the nature of QCD at high temperature whose answers should be predicted before the experiments are done. Beyond QCD, there are other important applications for nonperturbative field theory, particularly regarding symmetry breaking in the electroweak sector. Examples include strong higgs interactions, higgs mass bounds from triviality considerations and technicolor.

The current successes of QCD include the demonstration of confinement and chiral symmetry breaking from a first-principles calculation, characterization of the transition to quark matter at high temperature and qualitative results for the QCD spectrum and weak matrix elements. Nonetheless, it is clear that much greater computing power is needed in order to obtain quantitative control of all sources of systematic error in these calculations. We have, therefore, proposed to push the limits and natural evolution of computer development by building a machine capable of sustaining over a teraflops on QCD calculations.¹

Let us consider some of the projects that could be done with a Teraflops speed computer. The current state of the art in the quenched approximation (neglecting fermion loops) is a lattice size about $32^3 \times 64$ with gauge coupling $6/g^2$ between 5.7 and 6.5. For QCD with dynamical fermions, there are calculations on 20^4 or 32^4 lattices with a coupling of 5.7. With a teraflops speed computer, quenched calculations could easily be extended to a lattice size of $96^3 \times 192$. If a coupling of 6.5 is used with the quark mass set so that $m_\pi/m_\rho = 0.18$, then 100 independent quenched configurations can be generated in 15 days. An analysis of the staggered

¹Presented at DPF '92, Fermilab, November, 1992, to appear in the proceedings.

DISTRIBUTION OF THIS DOCUMENT IS UNLIMITED

MASTER

quark spectrum would take 150 hours, and the Wilson quark analysis could be done in 100 days.

Dynamical fermions are much more challenging and a calculation on a size of $32^3 \times 64$ with $m_\pi/m_\rho = 0.25$ would take about 800 hours for staggered quarks and 4000 hours with Wilson quarks. Such a calculation would provide a test of the effects of virtual quark loops. A test of asymptotic scaling with a coupling of 6.0 and $m_\pi/m_\rho = 0.5$ would be feasible.

With the above calculations we can study the glueball spectrum as well as the hadrons made from quarks. In addition, we can study weak matrix elements such as the dimensionless B parameters B_K and B_B , pseudoscalar decay constants f_B and f_{B_s} , and semileptonic form factors for the B and D mesons. Some calculations of weak matrix elements require a large volume such as $128^3 \times 256$ with $6/g^2 = 6.5$. Such a study would require about 3 months.

2. Hardware

The physicists in our collaboration are joined by computer scientists and engineers from the MIT Laboratory of Computer Science and Lincoln Laboratory. As detailed below, we propose to increase the speed of the CM5 nodes by a factor of 10 over the current design. This is done through the use of a redesigned floating point chip, new memory technology and new packaging methods.

The current CM5 processor is based upon a SPARC microprocessor that controls four vector units each of which is capable of 32 Megaflops in double precision (64 bit operands). Each vector unit is attached to an 8 Mbyte memory bank. The vector units and memory are actually on a separate board that is connected to the SPARC processor and network interface via an Mbus. The SPARC processor issues vector instructions to one or more vector units. To achieve peak speed, it is necessary to issue the same instruction to each vector unit. Thus, each processor has 32 Mbytes of memory and is capable of a peak speed of 132 Mflops.

The communication network of the CM5 is a "fat tree." At the lowest level, within a group of four processors there is a peak capacity of 20 Mbytes/sec between any pair of processors. This is for one way communication and both pairs, for example, can be communicating in both directions at the same time. Communication outside of the group of 4 nodes, but within the group of 16 involves using an additional level of the tree. Here, the peak capacity of the link out of the group of 4 nodes is 40 Mbytes/sec in each direction. However, since there are four nodes sharing that capacity, the peak rate for each node is 10 Mbytes/sec. At the next level of the tree, the data bandwidth again doubles as the number of nodes in the cluster quadruples, so the peak rate is 5 Mbytes/sec for each node. At higher levels of the tree, the bandwidth also quadruples at each level.

Our proposed computer will have 2048 nodes and can be completed by the Spring of 1995.² We expect to achieve a speed of 1.5 teraflops on the quenched calculations described above and 0.5 teraflops on the dynamical fermion calculations.

We plan to design a new floating point chip that will run at 40 MHz. There will be eight vector units per node. Since the chip has both an adder a multiplier, it will be capable of 80 Mflops in double precision; however, we will design the chip so that there are instructions that allow it to do 2 single precision operations in

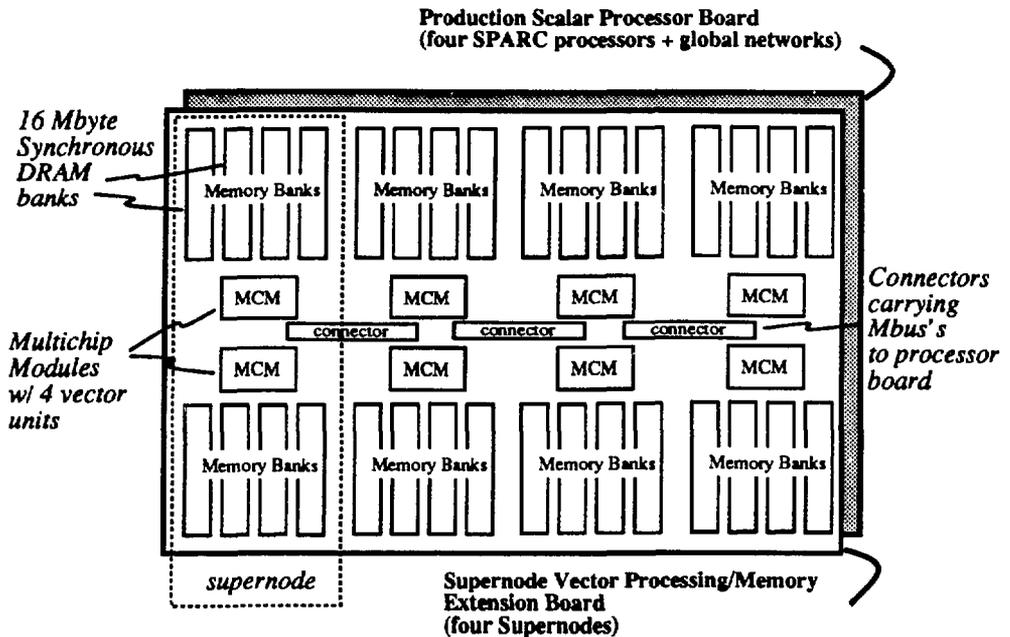


Figure 1: Layout of a supernode vector board containing 4 supernodes.

each unit rather than a double precision operation. Thus, with the same memory bandwidth, we get 160 Mflops speed for each vector unit or 1.28 Gflops for each node. There are two important technical problems that arise with these changes. First, we now require much more memory bandwidth. This will be done using a new technology called synchronous DRAM. At least eight vendors have plans to produce such chips with speeds up to 100 MHz within about a year. The second issue is how to pack the increased memory and floating point chips into the same space as the current CM5 vector boards. This will be done by packing four vector units together into a "multichip module." Space for the memory chips will be no problem since 16 Mbit chips will be used in place of 4 Mbit chips. Figure 1 shows the mechanical layout of the cluster of 4 supernodes that will replace the current production CM5 vector boards.

Brookhaven National Lab is supported under DOE Contract DE-AC02-76CH00016. The author's permanent address is Department of Physics, Indiana University, Bloomington, IN 47405 and he is supported under DOE grant FG02-91ER 40661.

References

1. S. Aoki *et al*, *Intl. J. Mod. Phys. C* 2 (1991) 829.
2. Subject to funding, of course.

DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.