

# A scalable PC-based parallel computer for lattice QCD

Z. Fodor<sup>a</sup>, S.D. Katz<sup>b\*</sup> and G. Papp<sup>a</sup>

<sup>a</sup>Institute for Theoretical Physics, Eötvös University, Pázmány P. 1/A, H-1117 Budapest, Hungary

<sup>b</sup>Deutsches Elektronen-Synchrotron DESY, Notkestr. 85, D-22607, Hamburg, Germany

A PC-based parallel computer for medium/large scale lattice QCD simulations is suggested. The Eötvös Univ., Inst. Theor. Phys. cluster consists of 137 Intel P4-1.7GHz nodes. Gigabit Ethernet cards are used for nearest neighbor communication in a two-dimensional mesh. The sustained performance for dynamical staggered(wilson) quarks on large lattices is around 70(110) GFlops. The exceptional price/performance ratio is below \$1/Mflop.

## 1. Introduction

The most powerful computers for lattice gauge theory are industrial supercomputers or special purpose parallel computers (see e.g. [1,2,3]). It is more and more accepted that e.g. off-the-shelf PC systems can be used to build parallel computers for lattice QCD simulations [4,5,6].

Single PC hardware has excellent price/performance ratios. For recent review papers see [7,8].

We present our experiences and benchmark results on a scalable system, which uses nearest-neighbor communication through Gigabit Ethernet (GigE) cards. The communication is fast enough (consuming 40% of the total time in typical applications) and cheap enough (30% of the total price is spent on the communication). The system can sustain  $\approx 100$  Gflops on today's medium/large lattices. Its price/performance ratio is below \$1/Mflops for 32-bit applications (twice as much for 64-bit applications).

Already in 1999 a report was presented [5] on the PC-based parallel computer project at the Eötvös University, Budapest, Hungary. A machine was constructed with 32 PCs arranged in a three-dimensional  $2 \times 4 \times 4$  mesh. Each node had two special, purpose designed communication cards providing communication to the six neighbors. We used the "multimedia extension" instruction set of the processors to increase the performance. The communication bandwidth (16

Mbit/s) at that time was only enough for bosonic simulations. Here we report on a system which needs no hardware development and has two orders of magnitude larger communication bandwidth through GigE cards.

## 2. Hardware

Each node consists of an Intel KD850GB motherboard, Intel-P4-1.7GHz processor, 512 MB RDRAM, 100 Mbit Ethernet card, 20.4 GB IDE HDD and four SMC9452 GigE cards for the two dimensional inter-node communication. The 100 Mbit Ethernet cards with switches are used for job management only.

We have 137 PC nodes. A single machine is used for controlling. Two smaller clusters with 4 nodes are used for development. The remaining 128 machines can be used as one cluster with 128 nodes (or two clusters with 64 nodes or four clusters with 32 nodes) for mass production.

The best number of connected nodes for a given lattice can be determined by optimizing the surface to volume ratio of the local sub-lattice. Changing the node topology needs reconnecting cables, which can be done easily in a few minutes.

In April, 2002 the price of one node including the 100 Mbit Ethernet switches is \$380 (see [9]). The four GigE cards with cables cost additional \$160 for each node. The power consumption of the nodes (140W each) requires a cooling system which costs around \$13 per node. Thus, the total node (PC+communication+cooling) price is

---

\*on leave from Institute for Theoretical Physics, Eötvös University, Pázmány P. 1/A, H-1117 Budapest, Hungary

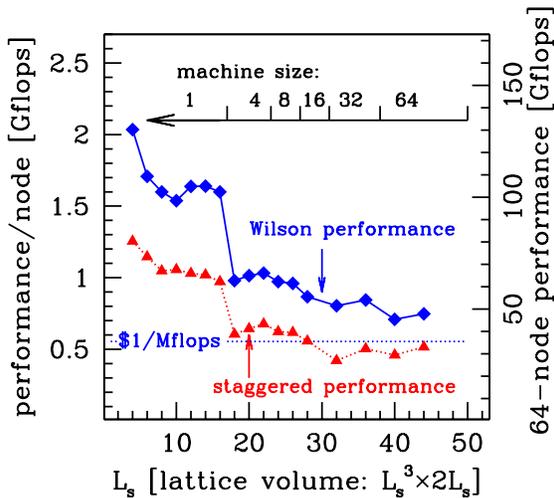


Figure 1. Full QCD performance results on a 64-node cluster. The node and the total performance as a function of the spatial extension  $L_s$ . The lattice sizes are chosen to be  $L_s^3 \times 2L_s$ . The dotted line represents the \$1/Mflops value. Wilson and staggered quarks are both indicated. The machine size (the number of communicating nodes) for a given lattice volume is shown by an inserted scale.

about \$553.

The key element of a cost effective design is an appropriate balance between communication and the performance of the nodes. We spent more than twice as much on the bare PC (\$393 including cooling) than for the Gigabit communication (\$160). These numbers are in strong contrast with Myrinet based PC systems, for which the high price of the Myrinet card exceeds the price of such a PC by a factor of two.

### 3. Software

The main operating system of the cluster is SuSE Linux 7.1, being installed on each node. Job management is done by the “main” computer through the 100 Mbit Ethernet network. We developed a simple job-management utility to distribute jobs on the nodes and collect the results.

To take advantage of the Gigabit communication from applications (e.g. C, C++ or Fortran), a simple C library was written using the

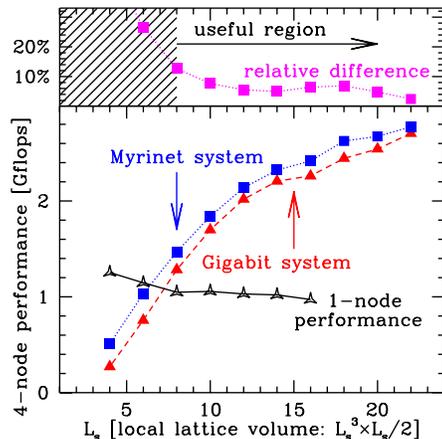


Figure 2. Sustained staggered performance of a 4-node ( $2 \times 2$ ) system as a function of the local lattice volume  $V=L_s^3 \times L_s/2$  (lower panel). Squares show the performance of a myrinet, triangles that of a gigabit system. The one-node performance is also shown by stars. Communication is of no use if the one-node performance is better than the 4-node one. Thus, the useful region is given by local lattices  $8^3 \times 4$  or larger. In this useful region the relative difference between the performances of the myrinet and gigabit systems (upper panel) is always less than 12%.

standard Linux network interface. Currently, we are using a standard socket based communication with Transmission Control Protocol (TCP) widely used in network applications. The typical bandwidth that we can reach in QCD applications is around 400 Mbit/s in contrast to the theoretical 1000 Mbit/s. Writing a lower level driver for the gigabit cards which would increase the bandwidth is in progress. The functions of the library can be used to open and close the communication channels and transfer data between the neighbors.

### 4. Performance

Figure 1 gives a summary of our benchmark runs. Single precision is used for local variables (gauge links and vectors, see [10]) and dot products are accumulated in double precision. Using double precision gauge links and vectors reduces

the performance by approximately 50%.

For benchmarking we started from the MILC code [10]. To increase the performance of the code we modified it by three different techniques.

First of all, we used the “multimedia extension” (*sse*) instructions of the Intel-P4 processors. As we pointed out in 1999 [5] this capability can accelerate the processor by a large factor. We rewrote almost the whole conjugate gradient part of the program to assembly (including also loops over the lattice and not only elementary matrix operations). This way we obtained a speedup factor of  $\approx 2$ . (see similar results of e.g. Ref. [8,11]).

Secondly, an important speedup was obtained by changing the data structure as suggested by S. Gottlieb [12]. In the original “site major” concept all the physical variables of a given site are stored in one structure and the lattice is an array of these structures. Instead of this concept one should use “field major” variables. The set of a given type of variable of the different sites are collected and stored sequentially. This increases the number of cache hits. Similarly to Ref. [12] this change leads to a speedup factor of  $\approx 2$ .

The third ingredient of our improvement was a better cache management by extensive use of the “prefetch” instruction.

One of the most obvious features of Figure 1 is the sharp drop of the performance when one turns on the communication. The most economic solution is to turn on the communication only if the memory is insufficient for the single-node mode. Smaller local lattices can be also studied by using the communication (for instance for thermalization or parameter tuning); however, in these cases the communication overhead increases somewhat. The performance as a function of the local lattice volume is shown on Fig. 2 (we also made a direct comparison between our architecture and a Myrinet system at DESY, Hamburg). Two orders of magnitude smaller local lattice means factor of  $\approx 2$  drop in the performance. Thus, for local lattices as small as  $8^3 \times 4$  the performance is still acceptable. This indicates clearly the scalability of this architecture. Note that the performance of a Myrinet based system is only 5-10% better than that of the GigE based system.

## 5. Conclusion

We report on the status of our PC-based parallel computer project for lattice QCD (for more details see ref. [13]). Nearest-neighbor communication is implemented in a two-dimensional mesh using Gigabit Ethernet cards. This architecture presents a good compromise between computation and communication. The saturated sustained performances on large lattices are around 0.5(0.8) GFlops/node for staggered (Wilson) fermions, which gives a price/performance ratio better than \$1.0(0.7)/Mflop.

**Acknowledgment.** We thank F. Csikor and Z. Horváth for their continuous help. This work was supported by Hungarian Science Foundation Grants under Contract Nos. OTKA-T37615/T34980/T29803/M37071/OMFB1548/OMMU-708.

## REFERENCES

1. R. Alfieri et al., hep-lat/0102011.
2. <http://www.rccp.tsukuba.ac.jp>.
3. <http://phys.columbia.edu/~cqft>.
4. N. Eicker et al., Nucl. Phys. (Proc. Suppl.) 83 (2000) 798 [hep-lat/9909146].
5. F. Csikor et al., Comp. Phys. Comm. 134 (2001) 139 [hep-lat/9912059].
6. X.Q. Luo et al., cs.DC/0109004.
7. N. Christ, Nucl. Phys. (Proc. Suppl.) 83 (2000) 111; S. Gottlieb, Nucl. Phys. (Proc. Suppl.) 94 (2001) 833.
8. M. Lüscher, Nucl. Phys. Proc. Suppl. 106 (2002) 21.
9. <http://www.pricewatch.com>
10. <http://physics.indiana.edu/~sg/milc.html>.
11. M. Di Pierro, hep-lat/0110116.
12. S. Gottlieb, hep-lat/0112026, hep-lat/0112038, hep-lat/0112039.
13. Z. Fodor, S. D. Katz and G. Papp, hep-lat/0202030.