JPL XServe Cluster Running the AltiVec Fractal Benchmark

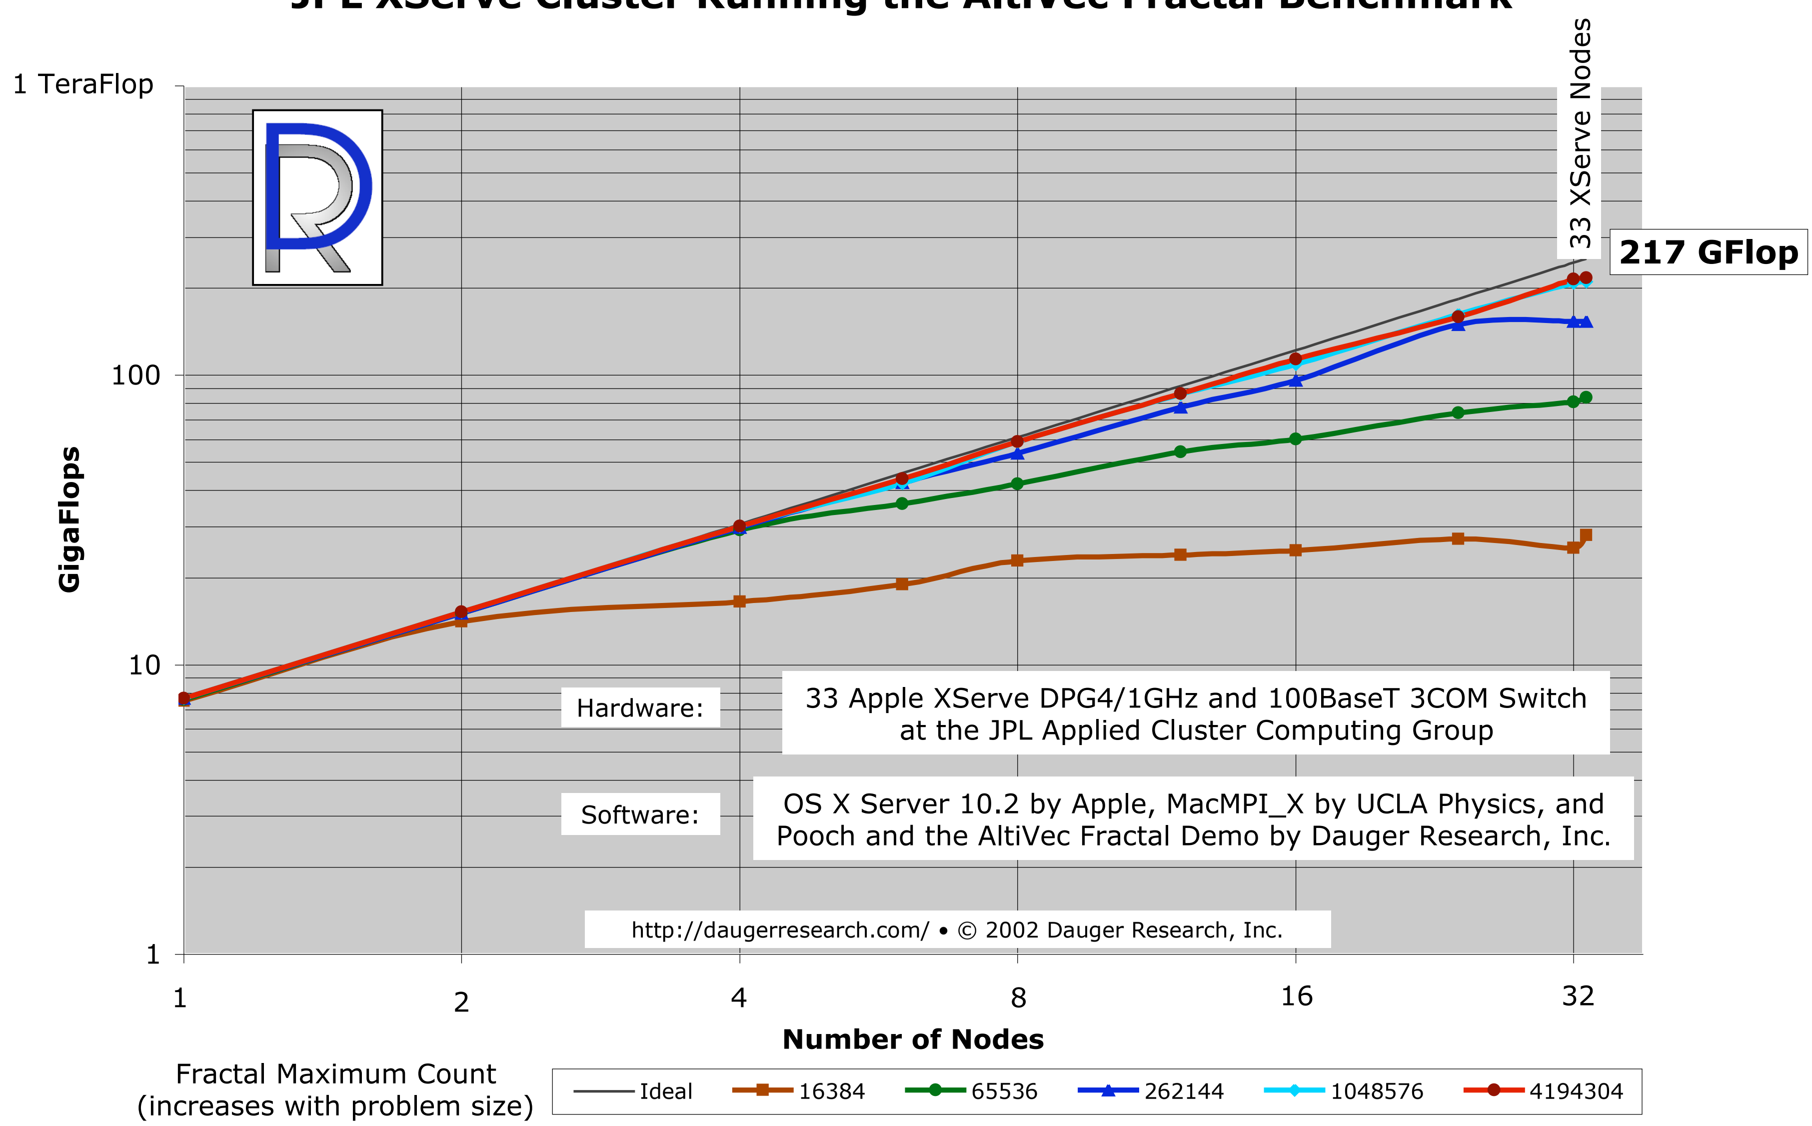

JPL XServe Cluster Running the AltiVec Fractal Benchmark

XServe cluster achieves over 1/5 TeraFlop using 66 1-GHz G4's and demonstrates excellent potential scalability

Hardware:

Nodes:	**33 Dual-Processor XServe G4/1GHz's**
Network:	**100BaseT 3COM 48-port Switch**

Software:

Application:	**AltiVec Fractal Carbon demo**
Communications/Message-Passing Library:	**MacMPI_X.c**
Cluster Support and Management:	**Pooch Application**
Operating System:	**Mac OS X Server 10.2**

Location and Host:

Jet Propulsion Laboratory's Applied Cluster Computing Group

Date:

November 5, 2002

The above figure illustrates the potential performance and scalability of clusters using Apple's new XServe. The Applied Cluster Computing Group (formerly known as the High-Performance Computing Group) at NASA's Jet Propulsion Laboratory (JPL) recently acquired and assembled a cluster using 33 XServe's. Using Pooch, the group ran the AltiVec Fractal Carbon demo and achieved over 217 billion floating-point operations per second on this XServe cluster, the largest result yet accomplished using an XServe cluster.

Members of the Applied Cluster Computing Group, building on their experience using other parallel computers and cluster types, spearheaded the planning, purchase, and construction of this XServe cluster. They plan to make this XServe cluster available for other parts of JPL in addition to using it for their own parallel codes using both MacMPI_X and mpich. Viktor Decyk of the Plasma Physics Group at University of California, Los Angeles, (UCLA) is also a member of this group and contributed to the planning and successful operation of this cluster. In addition, the group invited Dean Dauger of Dauger Research, Inc., & UCLA to help with the software assembly and configuration of these XServe's. As of this writing, this is the largest XServe cluster known to exist.

About the Benchmark

The different colored lines indicate the fractal benchmark code operating on different problem sizes. As expected on any parallel computer running a particular problem type, larger problems scale better. The AltiVec Fractal Carbon demo uses fractal computations that are iterative in nature. For a portion of the fractal image, these iterations may continue *ad infinitum*; therefore, a maximum iteration count is imposed. In the AltiVec Fractal Carbon demo, this limit is specified using the Maximum Count setting. Increasing the Maximum Count setting to 65536, then 262144, and so on, increases the problem size. It was clear that, given sufficient problem size, the XServe cluster was able to acheive over 1/5 TeraFlop (1 TF = 1000 GF = one trillion floating-point calculations per second).

The performance is determined by the total number of floating-point calculations performed that contribute to the answer and the time it takes to construct the answer. This time includes not only the time it takes to complete the computation, but also the time it takes to communicate the results to the screen on node 0 for the user to see. Also note that we quote the actual achieved performance, a practical measure of true performance while solving a problem, rather than the theoretical peak performance.

The time it takes to compute most of these fractals is roughly proportional to the Maximum Count setting, yet, since the number of pixels is the same, the communications time remains constant. For the smallest problem sizes on a large number of nodes, it was clear that communications time became greater than the computation time. By increasing the problem size significantly, the computation time was once again much greater than the communications time.

The dark "Ideal" line is an extrapolation multiplying the node count by the performance of one node alone. As shown in the graph, the cluster's performance while solving the larger problems closely approach that "Ideal" extrapolation. That observation tells us we can find no evidence of an intrinsic limit to the size of a Mac cluster.

Conclusion

After running a series of numerically-intensive trials on a 33-node XServe cluster, we were able to achieve over 1/5 TeraFlop on certain problems. These results were very repeatable. No evidence of an intrinsic limit to the size of a Macintosh-based cluster could be found. Building on a previous result using 76 Power Macs at USC, this finding is further evidence that Macintosh-based clusters are capable of excellent scalability in performance.

Acknowledgements

The above could not be accomplished without involvement of many people. Many thanks goes to NASA's Jet Propulsion Laboratory and its Applied Cluster Computing Group. Also, Ron Ustach and Monika Mohler from Apple Computer, Inc., faciliated the purchase of and assisted with the cluster.

---

updated: November 9, 2002