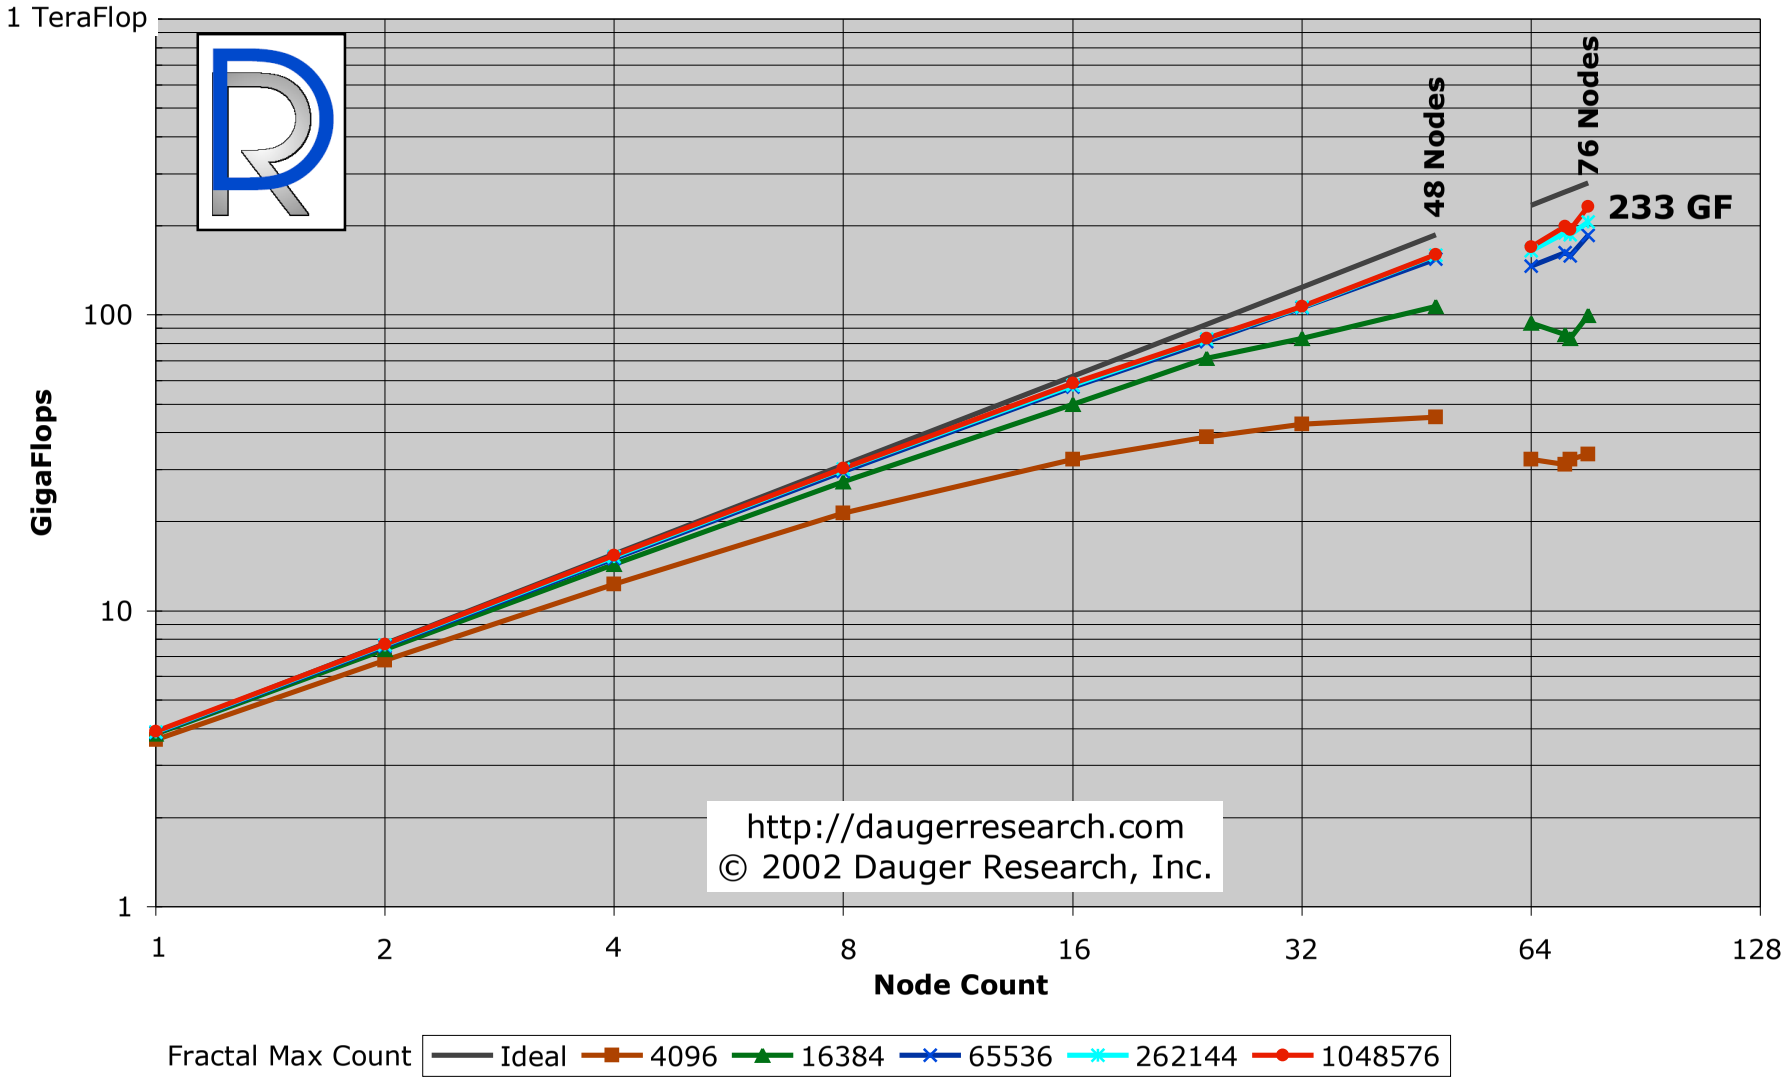
USC Macintosh Cluster Running the AltiVec Fractal Benchmark
1 TeraFlop
100
10
1
GigaFlops
48 Nodes
76 Nodes
233 GF
http://daugerresearch.com
© 2002 Dauger Research, Inc.
1
2
4
8
16
32
64
128
Node Count
Fractal Max Count
Ideal
4096
16384
65536
262144
1048576

USC Macintosh Cluster Running the AltiVec Fractal Benchmark

Large Macintosh cluster achieves over 1/5 TeraFlop on 152 G4's and demonstrates excellent scalability.

Hardware:

Nodes:	**56 Power Mac Dual-Processor G4/533's + 20 Power Mac Dual-Processor G4/450's**
Network:	**100BaseT Switched**

Software:

Application:	**AltiVec Fractal Carbon demo**
Message-Passing Library:	**MacMPI_X.c**
Cluster Support and Management:	**Pooch Application**
Operating System:	**Mac OS 9.2.1 (Basic)**

Location:

USC Language Arts Center and other facilities

Dates:

December 17, 19, 20, & 27, 2001

The above figure illustrates the potential performance and scalability of Macintosh clusters. Over Christmas Break 2001, Kay Ferdinandsen, Curtis Safford, and Tom Katsouleas of the Unversity of Southern California (USC) invited Viktor Decyk of the University of California, Los Angeles, (UCLA) and Dean Dauger of Dauger Research, Inc., & UCLA to perform these and other benchmarks on the Macs residing in computer labs at USC. Except for Decyk's physics codes, this was the first time the software was run on that many nodes or processors.

Note that, at 48 nodes and below, we were able to use a homogenous cluster of DP G4/533's; however, beyond 56 nodes, we combined DP G4/450's with the 533's. As a result, that hetereogeneous cluster cannot be expected to perform as evenly as a homogeneous one of the same size. Nevertheless, we were able to acheive over 1/5 TeraFlop (1 TF = 1000 GF = one trillion floating-point calculations per second).

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, whose the default value is 4096 iterations. By increasing the Maximum Count setting to 16384, then 65536, and so on, we increased the problem size.

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. When running on over 50 nodes at the 4096 setting, the total time was less than a half second, so it was clear that communications time became similar to the computation time. By increasing the problem size significantly, the computation time was once again much greater than the communications time.

The grey "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 76-node Macintosh 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 cluster could be found, indicating that Macintosh clusters are capable of excellent scalability in performance.

Acknowledgements

The above could not be accomplished without the help of others. Many thanks to:

- Kay Ferdinandsen - ISD Program Director and HPCC Facilitator,
- Curtis Safford - ISD Macintosh System Adminstrator, and
- Thomas Katsouleas - Professor of Electrical Engineering

from the University of Southern California, and Tim Parker, Steve Cook, and Frank Callaham from Apple Computer, Inc.

Copyright © 2002 Dauger Research, Inc.

---

updated: January 31, 2002