Oliver Wong
CS352 - Homework #1

Benchmark Critique: Apple Power Mac G5

Apple's recent announcement of the G5 was accompanied by a benchmark figures obtained through Veritest. I believe that Veritest did an excellent job documenting their benchmarking technique, which allows me to get a good idea of how they benchmarked and to provide critiques on their methods.

The good thing they did was that they applied similar optimizations (such as compiler flags) to both the Dells and the Apples they tested. Their choice of the SPEC benchmarks was also good because it is well-known and (somewhat?) an industry standard. Apple also made a good move in outsourcing the benchmarking to another company as that supposedly will give more credibility to the scores since the hired company *should* be totally objective... although I doubt that is the actual case.

I believe there were various things wrong with the way Veritest benchmarked the systems however, giving an advantage to the G5's. The SPEC tests are meant to test the CPU, memory architecture, and compiler (SPEC, www.spec.org/cpu2000/docs/readme1st.html#Q5, 2003). Veritest says that "to directly compare the performance of the hardware systems, the same compiler - GCC, with similar settings were used on both platforms." I do not think that using the "same" compiler for both platforms allows for a direct comparison of the hardware. Even though GCC is used in both, the way GCC handles x86 in no way relates to the way it handles PPC. GCC's performance for x86 and PPC is can not be assumed to be equal just because it is the same compiler. It might even be the case that different people work on x86 GCC than PPC GCC. I believe a better way to eliminate the compiler from the equation would be to use the best available compiler for both platforms. That way, people would be able to get an idea of what would be the best feasible performance they could get from each system.

Another thing I think Veritest did wrong was that they used a special implementation of the malloc library for the G5's that "is geared for speed rather than memory efficiency and is single-threaded which makes it unsuitable for many uses" (Veritest, www.veritest.com/clients/reports/apple/apple_performance.pdf, 2003). This is neither the default malloc library nor suggested for regular use. The SPEC Run and Reporting Rules state that the vendor should encourage the optimization for general use and must provide a suitable environment for running typical C, C++ programs. Using this special malloc library violates those two conditions and thus would make the figures obtained with it invalid.

Overall, Veritest's tests show that the G5 is a solid performer, however, they most likely gave an intentional advantage to Apple's hardware over Dell/x86 in their testing setup.