A HINT of Things to Come

In today's high-tech society, performance is everything. From the car that goes from 0 to 60 mph in five seconds to the computer that solves complex equations in the blink of an eye, performance sets the tone for how consumers measure a product's efficiency and usefulness.

Perhaps nowhere is the quest for performance more critical than in the highly competitive, consumer-driven computer industry. Whether for simple word processing or complex scientific calculations, finding a way to effectively measure computer performance could spark revolutionary changes in the ways engineers design computers and the ways consumers compare and buy them.

At the forefront of the quest to effectively measure computer performance is a new benchmarking tool called HINT, Hierarchical INTegration. Developed at the Ames Laboratory by John Gustafson, a Lab computational scientist and an Iowa State University adjunct professor of computer science, HINT measures a computer's full range of performance. "When you buy a car there are so many things involved, such as quality, comfort and styling. Just looking at the highest number printed on the speedometer would be a really dumb way to buy a car," says Gustafson, facetiously. "But up to now that's how people have selected computers. They've looked at the highest speed printed on the 'speedometer,' in this case the megahertz rating, and hoped it had something to do with how fast the computer operates. There's a lot more than that to buying a computer, just as there is to buying a car."

Ames Lab scientists John Gustafson (right) and Quinn Snell say HINT is possibly the most portable computer benchmark ever written. With HINT, scientists can measure computer performance on sophisticated supercomputers as easily as they can on simple pocket calculators.

The HINT Difference

Most computer benchmarks measure a computer's performance by fixing the size of the problem to be run, regardless of whether the computer can perform a thousand or a billion operations per second. "That makes it impossible to compare computers on a fixed problem," says Gustafson, who in 1990 developed a benchmark called SLALOM, which changed the method for measuring computer performance by fixing the calculation time at one minute and varying the job size. "But one minute is kind of an arbitrary number, and maybe it's more important to go for an hour," explains Gustafson.

HINT builds on the SLALOM benchmark but is faster to execute. Based on a fundamentally different concept for measuring speed, HINT fixes neither the problem size nor the calculation time. Instead it uses a work measure called QUIPS (Quality Improvement Per Second), also developed by Gustafson, to measure performance based on the amount of work a computer can perform over time. "What it suddenly helps you understand is that most computers start fast, just like you do when you start sprinting, and they slow down when they get out of fast memory and start to use main memory. Then they slow further when they have to go to disk," says Gustafson, who emphasizes that all those different ranges of speed are visible on the HINT benchmark.

"John's findings are important, particularly for high performance computing where performance depends enormously on how memory is accessed," says Norris Smith, editor at large of HPC wire, an international on-line publication on high-performance computing. "Judging from John's charts, we're seeing some very interesting, new approaches to measuring computer performance. I think HINT may shake up the whole benchmark industry."

The HINT benchmark also opens the door to the "personality" of the computer. "It allows you to get a feel for example, of whether a certain computer is a fast starter but runs out of gas easily," says Quinn Snell, an Ames Lab research assistant and co-author of the HINT program.

Snell remembers spending countless hours working with Gustafson to hash out the code for HINT. "When it worked I was quite excited and shocked because I didn't know what to expect out of it," he says. "The theory was John's, and he had an idea of what he wanted out of it. When we hit the first smooth graph, it was like the light came on. We could now tell somebody how fast a computer was for a problem of a given size. And you couldn't do that with other benchmarks."

Computer Industry Reacts

Although the jury is still out on whether HINT is the ultimate computer benchmark, one certainty is that HINT has the computer industry's attention. Client/Server Today (CST) Labs, which is devoted exclusively to hands-on product testing, has adopted HINT as its primary computer system benchmark. "When we showed HINT to CST Labs, they were impressed," says Snell, who demonstrated HINT on four different computers in two hours. "They'd never seen that before. They'd always taken longer to apply a benchmark."

"HINT is a giant leap in thought in benchmarking techniques," says Tim Sylvester, labs director for Client/Server Today magazine, who adds that CST Labs has already tested HINT on eight to ten systems and plans to test it on a number of other systems soon. "We are impressed with its ability to give a more accurate picture of the performance of a system, and we will definitely continue to use it."

The HINT benchmark is simple to use. The measured program is only two pages long. It's also scalable and easily ported to a variety of computer architectures. "Quinn deserves credit for writing a program that is possibly the most portable ever. You can run this on a pocket calculator or a supercomputer," says Gustafson. "No one has ever achieved that kind or portability before."

HINT was recently put to use in the acquisition of a new parallel computer for Ames Lab's Scalable Computing Laboratory (SCL). Scientists at the SCL asked vendors to adhere to specific product requirements, including providing the highest HINT benchmark rating possible. After careful consideration, the SCL purchased a Silicon Graphics POWER Onyx graphics super-computer. Gustafson has also benchmarked the IBM SP2 and Cray T3D computers, which are among the fastest computers in the world.

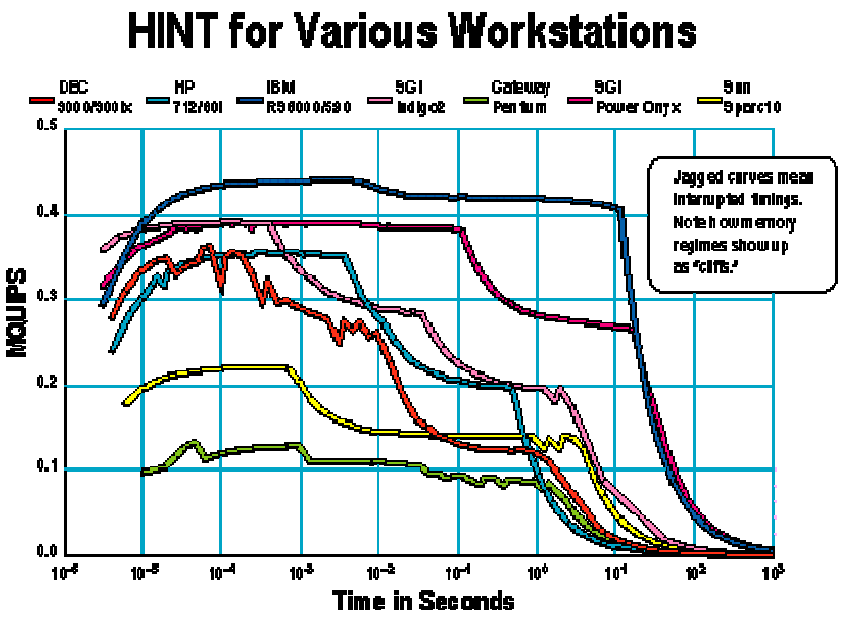

By plotting QUIPS on a variety of computer workstations, HINT shows how performance decreases as a problem moves through the various stages of computer memory.

A Bright Future

"I hope HINT will affect everyone who uses computers," emphasized Gustafson, who wants to publish an annual computer almanac, similar to a Consumer Reports guide to cars, that will include HINT benchmarking numbers for all commercially available computers. "Or maybe you'll be able to walk into a store and see in a bubble pack the HINT benchmark program for your computer," he says. "You buy it, pop it in and find out what your computer's HINT rating is and why. Or you can pop it into another machine and find out why that one may be twice as fast. Not only will you be able to use the benchmark to tune your machine, but you will be able to determine right there on the spot what type of computer you want to buy."

Published by Inquiry, winter 1995
Contact: john.gustafson@sun.com
The URL for this page is: http://www.scl.ameslab.gov/Publications/Inquiry,winter1995.html
Pages prepared by Maria E. Blanco.