from *SIAM News*, Volume 34, Number 9

Biannual Top-500 Computer Lists Track Changing Environments For Scientific Computing

By Jack Dongarra, Hans Meuer, Horst Simon, and Erich Strohmaier

In the last 50 years, the field of scientific computing has seen rapid, sweeping changes—in vendors, in architectures, in technologies, and in the users and uses of high-performance computer systems. The evolution of performance over the same 50-year period, however, seems to have been a very steady and continuous process. Moore's law and peak performance of various "supercomputers" over time.ic computing has seen rapid, sweeping changes—in vendors, in architectures, in technologies, and in the users and uses of high-performance computer systems. The evolution of performance over the same 50-year period, however, seems to have been a very steady and continuous process. Moore's law is often cited in this context, and, in fact, a plot of the peak performance of the various computers that could be considered the "supercomputers" of their times clearly shows that this law has held for almost the entire lifespan of modern computing (see Figure 1). On average, performance has increased every decade by about two orders of magnitude.

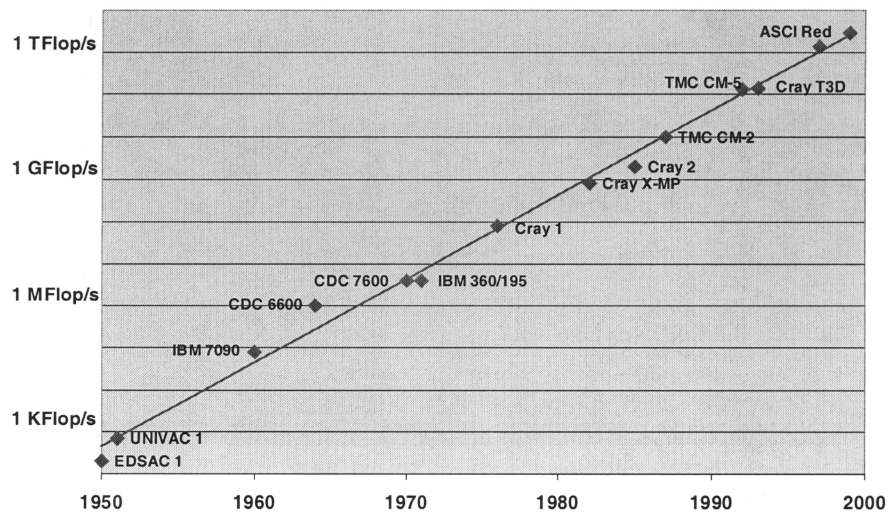

Figure 1. *Moore's law and peak performance of various "supercomputers" over time.*

From the First Vector Machines to Today's Cluster-based Systems

The modern supercomputing era dates back to the second half of the 1970s, with the introduction of vector computer systems from Cray Research and, later, from Control Data Corp. These systems offered performance advantages of at least one order of magnitude over conventional systems of the time. Raw performance was the main—if not the only—selling argument.

A supercomputer is typically used for scientific and engineering applications that require great amounts of computation; government research centers are the traditional use of these computers. With the 1980s came a push to integrate the vector systems into conventional computing environments. The manufacturers provided standard programming environments, operating systems, and tools that would allow effective use of these machines. As the environment became more user-friendly, the manufacturers were successful in attracting industrial, academic, and commercial customers—and in building an industry. The main routes to increased performance were improved chip technologies and the production of shared-memory multiprocessor systems.

By the late 1980s there had been a shift to highly parallel computing on scalable distributed-memory systems. This shift was fostered by several government programs, with the goal of moving to higher levels of performance. Improving performance and overcoming the hardware-scalability limitations of shared-memory systems were the main objectives of these new systems. The increased performance of standard microprocessors after the RISC revolution, together with the cost advantages of large-scale production, formed the basis for the "attack of the killer micro." Consequences were a transition from ECL to CMOS chip technology and the use of off-the-shelf microprocessors instead of custom-designed processors for these highly parallel systems, sometimes referred to as massively parallel processors, or MPPs.

With the acceptance of MPP systems not only for engineering applications but also for new commercial applications (for database applications in particular) came different criteria for market success: system stability, continuity of the manufacturer, and good price/performance ratios. Success in commercial environments became an important requirement for a successful supercomputer business. Given these factors, and the shrinking number of vendors in the market, hierarchical systems built with components designed for the broader commercial market began to replace homogeneous systems at the very high end of the performance spectrum. More recently, clusters built with commodity off-the-shelf (COTS) components have also gained more and more attention. Cluster-based computing has quickly spread and is replacing many of the traditional supercomputer machines. Driving this trend is the desire for low-cost computing. The cluster concept is an empowering force. It wrests high-level computing away from the privileged few and makes low-cost parallel-processing systems available to those with modest resources. Research groups, high schools, colleges, and small businesses can all build or buy their own clusters.

In the early 1990s, as parallel vector systems were reaching their widest distribution, a new generation of MPP systems came on the market, claiming to be able to substitute for, or even surpass, parallel vector processors. It was during this time, with the goal of providing a better basis for statistics on high-performance computers, that the Top500 list (see www.top500.org) was born.

Who Uses the World's Most Powerful Computers?

The report, which has been updated twice a year since its June 1993 debut, lists the sites at which the 500 most powerful computer systems in the world are installed at the time of the report. Performance on a LINPACK benchmark (Dongarra, 1989) is a measure used to rank the computers. The benchmark attempts to measure the best performance of a computer in solving a system of linear equations. The problem size and software can be chosen to produce the best performance. The first Top500 list already included 156 highly parallel systems (31% of the systems on the list). Because it can be hard to appreciate the effect of exponen-tial growth, such as that predicted by Moore's law, we offer the following example: Many modern laptop computers would have been on the Top500 list back in 1993!

The Top500 lists for 1995 revealed remarkable changes in the distribution of the systems among customers of various types (academic sites, research labs, industrial/commercial users, vendor installations, and classified sites). Until June 1995, the major trend seen in the Top500 data had been a steady decrease in the number of industrial customers, matched by increasing numbers of government-funded research sites. This trend reflects the influence of the different government high-performance computing programs that enabled research sites to buy parallel systems, especially distributed-memory systems. Given that such systems were often far from mature or stable, industry was understandably reluctant to follow the researchers' lead. Hence, industrial customers stayed with their older vector systems, which gradually dropped off the Top500 list because of lower performance.

Beginning in 1994, however, several companies, among them SGI, Digital, HP/Convex, and Sun, had begun to sell symmetric multiprocessor (SMP) models of their major workstation families. From the very beginning, these systems were popular with industrial customers because of the maturity of the architectures and their superior price/performance ratios. At the same time, IBM SP2 systems, highly parallel systems with distributed memory, appeared at a number of industrial sites. While the SP was sold initially for numerically intensive applications, successful expansion to a larger market, including database applications, began in the second half of 1995.

Subsequently, the number of industrial customers listed in the Top500 increased from 85, or 17%, to the current level of 236, or 47%. These machines have been placed in a variety of settings, such as the aerospace and automotive industries, and used for financial analysis, oil exploration, and drug design, among many other applications.

The increasing presence of Top500 systems in industrial settings can be attributed to several factors:

- MPP systems replaced the vector systems formerly in place at a substantial number of industrial sites. The change in architecture reflects the fact that parallel systems were ready for commercial use and environments.

- The most successful companies (Sun, IBM, and SGI) had many industrial customers. Their success was built on the use of standard workstation technologies for their MPP nodes. This approach provides a smooth migration path for applications, from workstations up to parallel machines.
- The maturity of these advanced systems and the availability of software tools for key applications made the systems appealing to commercial customers. Especially important are database applications, which can make effective use of highly parallel systems (more than 128 processors).

2010—The First Petaflops Machine?

Despite the dramatic changes in many aspects of the high-performance computing market over time, performance, as mentioned earlier, seems to have evolved in quite good agreement with empirical laws, such as Moore's law. The Top500 provides an ideal data base for verifying such an observation. Looking at the computing power of the individual machines in the Top500 and the evolution of the total installed performance, we plotted the performance of the systems at positions 1 and 500 in the list, as well as the total accumulated performance. In Figure 2, the curve for position 500 shows an average yearly increase of a factor of two; the curves for all other positions show growth rates of 1.8 ± 0.07 per year.

To compare these growth rates with those predicted by Moore's law, we needed to separate the influence of two factors that affect total performance: (1) increased processor performance and (2) increased numbers of processors per system. To get meaningful numbers, we excluded SIMD systems (which tend to have extremely large numbers of processors and extremely low processor performance) from our analysis. We have found that these two factors contribute almost equally to the total annual

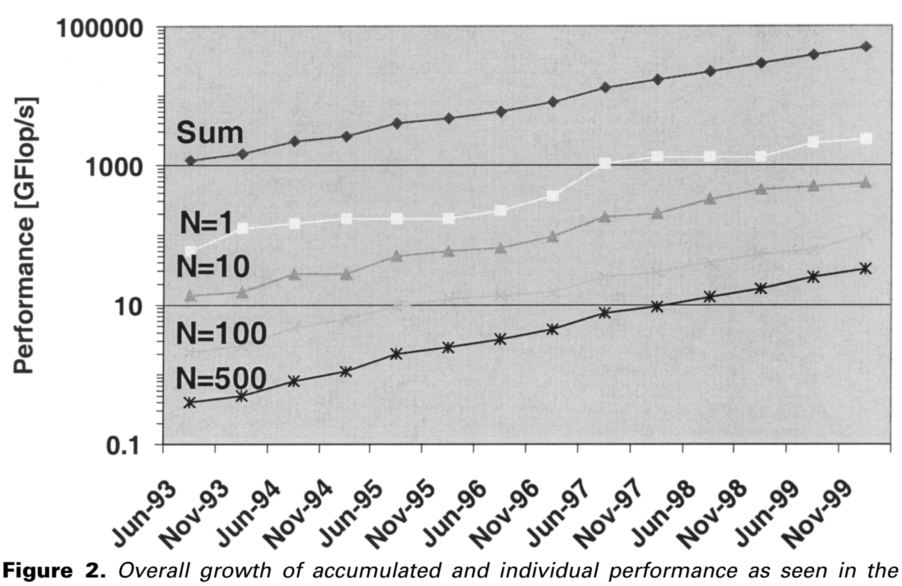

Figure 2. *Overall growth of accumulated and individual performance as seen in the Top500.*

performance growth factor of 1.82. The number of processors grows, on average, by a factor of 1.30 each year, and processor performance increases by 1.40 per year, as compared with the 1.58 predicted by Moore's law.

Based on Top500 data covering the last six years, and the assumption that current performance developments will continue for some time to come, we can now extrapolate from observed performance levels and compare these values with the goals of the government programs mentioned earlier. For Figure 3, using linear regression on the logarithmic scale, we extrapolated from the observed performance values—that is, we fit exponential growth to all levels of performance in the Top500. This simple fitting of the data produced surprisingly consistent results. Based on extrapolation from these fits, we can expect to have the first 100-TFlop/s system by 2005 (one or two years later than called for in the Department of Energy's ASCI "path forward" plans (which, however, use peak performance rather than best Linpack performance). Also by 2005, we can predict that no system smaller then 1 TFlop/s will make the Top500.

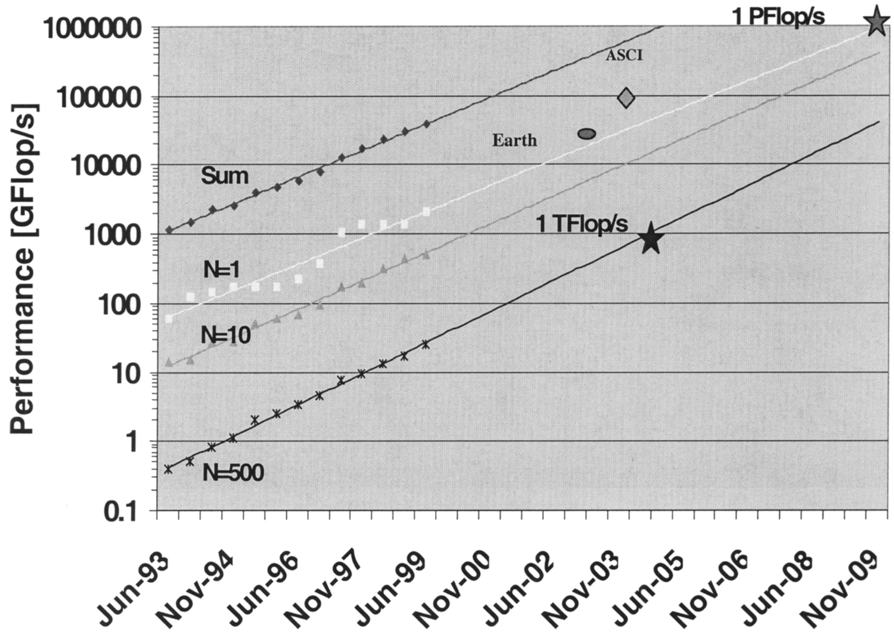

Figure 3. *Extrapolation from recent performance growth rates seen in the Top500.*

Looking even further into the future, we speculate that, based on the current progress and rate of increase of processors, the first PETAFlop/s (10^{15} floating-point operations per second) system should be available in about 2010. Given the rapid changes in the technologies used in HPC systems, however, no reasonable projection can be made at this point for the architecture of such a system. There seems to be no end in sight for the rapid cycles of redefinition that have changed the face of the high-performance computing market so substantially since the introduction of the Cray 1 two and a half decades ago.

Jack Dongarra is a professor of computer science at the University of Tennessee, Knoxville. Hans Meuer is a member of the Fakultät für Mathematik und Informatik, Universität Mannheim, Germany. Horst Simon and Erich Strohmaier are researchers at the National Energy Research Supercomputing Center, Lawrence Berkeley National Laboratory.