Lawrence Berkeley National Laboratory

Recent Work

Title

The TOP500 List and Progress in High-Performance Computing

Permalink

<https://escholarship.org/uc/item/28b779q1>

Journal

Computer, 48(11)

ISSN

0018-9162

Authors

Strohmaier, E
Meuer, HW
Dongarra, J
[et al.](#)

Publication Date

2015-11-01

DOI

10.1109/MC.2015.338

Peer reviewed

The TOP500 List and Progress in High- Performance Computing

Erich Strohmaier, Lawrence Berkeley National Laboratory

Hans W. Meuer, University of Mannheim

Jack Dongarra, University of Tennessee

Horst D. Simon, Lawrence Berkeley National Laboratory

For more than two decades, the TOP500 list has enjoyed incredible success as a metric for supercomputing performance and as a source of data for identifying technological trends. The project's editors reflect on its usefulness and limitations for guiding large-scale scientific computing into the exascale era.

The TOP500 list (www.top500.org) has served as the defining yardstick for supercomputing performance since 1993. Published twice a year, it compiles the world's 500 largest installations and some of their main characteristics. Systems are ranked according to their performance of the Linpack benchmark,¹ which solves a dense system of linear equations. Over time, the data collected for the list has enabled the early identification and quantification of many important technological and architectural trends related to high-performance computing (HPC).²⁻⁴

Here, we briefly describe the project's origins, the principles guiding data collection, and what has made the list so successful during the two-decades-long transition from giga- to tera- to petascale computing. We also examine the list's limitations. The TOP500's simplicity has invited many criticisms, and we consider several complementary or competing projects that have tried to address these concerns. Finally, we explore several emerging trends and reflect on the list's potential usefulness for guiding large-scale HPC into the exascale era.

TOP500 ORIGINS

In the mid-1980s, coauthor Hans Meuer started a small and focused annual conference that has since evolved into the prestigious International Supercomputing Conference (www.isc-hpc.com). During the conference's opening session, Meuer presented statistics about the numbers, locations, and manufacturers of supercomputers worldwide collected from vendors and colleagues in academia and industry.

Initially, it was obvious that the supercomputer label should be reserved for vector processing systems from companies such as Cray, CDC, Fujitsu, NEC, and Hitachi that each claimed to have the fastest system for scientific computation by some selective measure. By the end of the decade, however, the situation became increasingly complicated as smaller vector systems became available from some of these vendors as well as new competitors (Convex, IBM) and as massively parallel systems with SIMD architectures (Thinking Machines, MasPar) and MIMD systems based on scalar processors (Intel, nCube, and others) entered the market. Simply counting the

installation base for systems of such vastly different scales did not produce any meaningful data about the market. New criteria for which systems constituted supercomputers were needed.

After two years experimenting with various metrics and approaches, Meuer and coauthor Erich Strohmaier concluded that the best way to provide a consistent, long-term picture of the supercomputer market was to maintain a list of systems up to a predetermined cutoff number, ranked according to their actual performance. On the basis of previous studies they determined that at least 500 qualified systems could be assembled, and so the TOP500 list was born.

RANKING SUPERCOMPUTER PERFORMANCE

The simplest and most universal ranking metric for scientific computing is floating-point operations per second (flops). More specialized metrics such as time to solution or time per iteration and time per gridpoint can be more meaningful in particular application domains and allow more detailed comparisons—for example, between alternative algorithms with different complexities—but are harder to define properly, more restricted in their use, and, due to their specialization, not applicable to the overall scientific computing market.

In addition to limiting performance measurement to flops, we decided to use actual measured values to avoid contaminating collected data with unsubstantiated and often outlandish performance “estimates” for systems that did not reliably function or even exist. In principle, measured results from different benchmarks or applications could be used to rank different systems,

FIGURE 1. Supercomputer performance over time as tracked by the TOP500. The red and orange lines show performance of the first and last systems, and the blue line average performance of all systems. Dashed lines are fitted exponential growth curves before and after 2008 for the orange line and before and after 2013 for the blue line.

but this would lead to inconsistent values and make comparisons difficult. To address this problem, we opted to select and mandate use of a single benchmark for all TOP500 editions.

This benchmark would not represent performance of an actual scientific application but coarsely embody scientific computing’s main architectural requirements. Because scientific computing is primarily driven by integrated large-scale calculations, we decided to avoid using simplistic benchmarks, such as embarrassingly parallel workloads, that could lead to very high rankings for systems otherwise unsuitable for scientific computing. Instead, we sought a benchmark that would showcase systems’ capabilities without being overly harsh or restrictive. Overall, the collected data should provide reasonable upper bounds for actual performance while penalizing systems unable to support a large fraction of scientific computing applications.

Obviously no single benchmark can ever hope to represent or approximate performance for most scientific computing applications, as the space of algorithms and implementations is too vast. The purpose of using a single benchmark in the TOP500 was never to claim such representativeness but to collect reproducible and comparable metrics.

Using a single benchmark that does not utilize all the system components necessary for most scientific applications or that maps better to particular computer architectures could lead to misleadingly high performance numbers for some systems, incorrectly indicating these systems’ suitability for scientific computing. To minimize such implicit bias, we decided that the benchmark should exercise all major system components and be based on a relatively simple algorithm that allows optimization for a wide range of architectures.

LINPACK

An evaluation of benchmarks suitable for supercomputing in the early 1990s found that Linpack¹ had the most documented results by a large margin and thus allowed immediate ranking of most of the systems of interest. The NAS Parallel Benchmarks (NAS PB)⁵ were also widely used, as most of them simulated actual application performance more closely, but none of them provided enough results to rank more than 40 percent of the systems.

Linpack solves a dense system of linear equations, which today is sometimes criticized as an overly simplistic problem. However, the benchmark is by no means embarrassingly parallel and it worked well with respect to reducing the rankings of loosely

FIGURE 2. Average number of processor sockets for new supercomputers in the TOP500, excluding systems with SIMD processors, vector processors, or accelerators. The exponential increase in the number of sockets up to 2008 accounts for the higher-than-expected growth rate in supercomputing performance during the same time period.

coupled architectures, which were of limited use to scientific computing in general. The High-Performance Linpack (HPL) implementation comes with a self-adjustable problem size, which allows it to be used seamlessly on systems of vastly different sizes as compared to discrete, fixed sizes for the NAS PB. Unlike many other benchmarks with variable problem sizes, HPL achieves its best performance on large-scale problems that use all of a system's available memory and not on small problems that fit into the cache. This greatly reduces the need for elaborate run rules and procedures to enforce full computer-system usage, which is similar to what many applications do. HPL also encodes an important set of optimization parameters and enables substantial performance optimization through their adjustments. All of these features made Linpack the obvious choice for our TOP500 ranking.

The selection of Linpack as the sole benchmark implies several other limitations. In Linpack, the number of operations is not measured but calculated with a simple formula based on the problem size and the original algorithm's computational complexity. In this case, changes and optimizations of the algorithm have to be limited so

that they do not reduce the number of floating-point operations performed. The TOP500 therefore cannot provide any basis for research into algorithmic improvements over time. Linpack and HPL could certainly be used to compare algorithmic improvements, but not in the context of the TOP500 ranking.

TOP500 TRENDS

Although we started the TOP500 to provide statistics about the HPC market at specific dates, it became immediately clear that the inherent ability to systematically track the evolution of supercomputers over time was even more valuable. Any TOP500 edition includes a mix of new and older systems and technologies. In the six months between successive editions of the list, the turnover rate has until recently (2012) averaged about 190 systems out of 500. The average age of systems since installation was 1.26 years. This continuous high replacement rate makes it possible to observe many trends by simply looking at the entire list. Using only the subsample of new systems on each list provides an overview of emerging technologies.

An analysis of supercomputing performance over time revealed that it grew noticeably faster than a direct interpretation of Moore's law would

predict. Figure 1 shows performance values for the first and last systems as well as average performance of all systems in the TOP500. Until 2008, these curves grew exponentially at a rate of 1.91 per year (multiplicative factor). Compared to the exponential growth rate of Moore's law at 1.59 per year, TOP500 system performance had an excess exponential growth rate of 1.20 per year. We suspected that this additional growth was driven by an increasing number of processor sockets in our system sample. (We use the term "processor sockets" to clearly differentiate processors from processor cores.)

To better understand this and other technological trends contained in the TOP500 data, we obtained a clean and uniform subsample of systems from each edition of the list by extracting the new systems and those systems that did not use special processors with vastly different characteristics including SIMD processors, vector processors, or accelerators (such as Nvidia GPUs and Intel Phi coprocessors). The average number of sockets for this subsample of systems is shown in Figure 2. The exponential increase in the number of sockets per system up to 2008 is 1.29 per year, which easily explains the observed growth rate above Moore's law.

Since 2008, however, the exponential growth rate of the last system on the TOP500 has been only 1.55 per year, slightly below that of Moore's law. The curve for average performance since 2013 indicates a similar slowdown. To find out why we examined per-socket performance, which, as Figure 3 shows, exhibits a constant exponential growth rate of 1.45 per year—just below what Moore's law would predict. Clearly, the increasing number of cores per socket has compensated for stagnant core

performance in the latter half of the past decade.

The changes in the overall growth rate since 2008 can be attributed mostly to a decline in the growth rate of the number of sockets and hence components in large-scale HPC systems, which has been a very modest 1.07 per year.

The impact of this observed slowdown is quite profound: prior to 2008, overall TOP500 system performance increased by a factor of about 1,000 over an 11-year time period; after 2008, it increased by only a factor of about 100 in the same time period (extrapolated to 11 years). Our data attributes this trend mostly to reduced growth in system size as measured by the number of processors and not to reduced performance growth for the processors themselves. This slowdown will only compound the projected decline in performance growth as we approach the end of the decade and Moore's law.

TOP500 FEEDBACK

TOP500 feedback falls mostly into two categories: requests for additional data and for different benchmarks.

Additional data

Requests for additional data usually center on system characteristics—memory, power consumption, disk space, costs, and so on. We agree that more information would be beneficial, but this first requires a consistent and precise definition of what to collect, which often is difficult to formulate. Also, such efforts are nontrivial in terms of time and manpower. In some cases we have added new data to the TOP500, while in others we have left the task to outside individuals and organizations. Some researchers have created

FIGURE 3. Maximal achieved performance (Rmax) of TOP500 systems on the Linpack benchmark. The increasing number of cores per socket has compensated for stagnant core performance in the latter half of the past decade.

geographically restricted top supercomputer lists—including for Russia (<http://top50.supercomputers.ru>), China (www.top500.org/blog/chinas-2014-hpc-top100-list-in-english), India (<http://topsupercomputers-india.iisc.ernet.in>), and Ireland (www.irishsupercomputerlist.org)—with substantially lower cutoff values, but otherwise largely follow our guidelines. Others augment our data with their own or publish a resorted list such as the Green500 (www.green500.org), which incorporates power consumption data. Still other researchers analyze TOP500 data for their own purposes—Peter Kogge and Tim Dysart's projections of architectural trends is one notable example.⁴

New benchmarks

Some requests for different benchmarks relate to applications in which flops performance is not the main metric of interest. The most prominent example is data-intensive applications. Since 2010, the Graph500 (www.graph500.org) has provided a complementary ranking of computing systems for such applications. As this project illustrates, however, developing a suitable benchmark (and metric) for a new application domain is nontrivial. Using a breadth-first search of an undirected graph as the base problem, the Graph500 team

has put substantial effort into specifying and redefining this benchmark but, as of July 2015, listed fewer than 200 systems for which measurements had been collected. Reported systems range from the largest supercomputers to single-node computers, including an iPad 3—which we certainly do not consider a supercomputer.

Most new benchmark requests are rooted in the argument that Linpack is a poor proxy for application performance. This is not surprising, as Linpack was never meant to be an application proxy; rather, we chose it for the TOP500 because it is a well-defined, easily optimizable, nontrivial benchmark requiring a well-integrated computer system for execution. Nowadays, it is clear that Linpack primarily stresses per-core flops performance and, secondarily, network performance. Due to its optimized implementations, it is very insensitive to memory performance (bandwidth and latency alike). A benchmark that more equally balances all three characteristics might represent overall application performance slightly better, and we are certainly open to and involved in efforts to develop such a benchmark—in particular, HPCG (High-Performance Conjugate Gradients)⁶ and HPGMG (High-Performance Geometric Multi-grid)⁷—but constructing a reasonably simple benchmark based on a relevant

GRAND CHALLENGES IN SCIENTIFIC COMPUTING

FIGURE 4. Performance of Gordon Bell Prize (GBP) winners and top-ranked systems from the TOP500 since 1993. The correlation between these two series of performance values is near perfect, suggesting that the Linpack benchmark still serves as a reasonably close upper bound for actual application performance.

computational algorithm that also displays most of Linpack’s additional benefits has proved to be quite difficult.

Any attempts to define new benchmarks for the TOP500 must also consider their potential impact on the rankings and derived trends. A benchmark that only results in a small, local, and hence trivial reordering might not yield enough substantive additional data to justify the effort. Radically different benchmarks might raise the same question Linpack has with respect to representation of the HPC community at large. Assuming that an alternative benchmark would include as large a variety of computing architectures as Linpack, we believe that the benchmark’s details should have very little influence on the major observed trends.

BEST APPLICATION PERFORMANCE

The TOP500 is often criticized because the published performance numbers for Linpack are far lower than what is achievable for actual applications. Unfortunately, no other benchmark with a consistent collection of data covering a reasonable subset of

supercomputers spanning two and a half decades exists. However, we can construct a series of “best application performance” data by looking at recipients of the ACM Gordon Bell Prize (GBP; <http://awards.acm.org/bell>), awarded annually since 1987 at the SC conference. This award is not based on a single benchmark—rather, each year the award committee selects a different application with the best performance. However, “best” does not necessarily imply highest performance, as the award considers all aspects of the application and system used. The GBP, however, clearly focuses on real-world scientific applications and can therefore serve as an indicator of how far Linpack has departed from actual peak application-performance levels and how the growth rate observed in the TOP500 would be affected.

Table 1 shows the system and performance data of each year’s GBP-winning application in the Sustained Performance category, together with the top-ranked system from the November TOP500 list, since 1993. We selected this edition of the list as it is released at the same SC conference where the GBP winner is

announced. Because GBP submissions can take 6–12 months to prepare and must occur several months prior to announcement of the winner, the June TOP500 edition is arguably more appropriate, and indeed the table contains many GBP winners whose performance was measured on the previous June TOP500 list’s top-ranked system. We matched performance values of GBP winners to those of systems on the June TOP500 list and obtained a very similar outcome. For brevity, we present here only the unmatched November TOP500 data.

Figure 4 plots the performance values of GBP winners together with those of number-one systems in the TOP500 over time. The correlation rate between these two series of performance values (derived from log10-transformed data to eliminate the raw data’s overwhelming scale effects) is near perfect at 0.99. Fitting exponential growth rates to both data-sets yields annual growth rates of 1.89 for TOP500 systems and 1.85 for GBP winners, which is reasonably close considering that we made no effort to clean up the raw GBP data.

The average ratio of GBP to Linpack performance values given in Table 1 is 50 percent, and even eliminating obvious outliers from the list of GBP winners (a system without Linpack in 1995, special-purpose systems in 2000 and 2001, and a high-performing 35.3-Tflops system in 2003) results in an average ratio of 45 percent. Ratios rarely fall below 25 percent, and little if any systematic change in values is apparent: average ratios for the first 10 years are only about 10–20 percent higher than those for the last 10 years. While GBP performance values arguably are no more representative than Linpack performance values, this high

TABLE 1. Gordon Bell Prize (GBP) winners from 1993 to 2014 and top-ranked systems from the corresponding TOP500 list.

Year	TOP500 no. 1 system perf. (Gflops)	GBP winnerperf. (Gflops)	Perf. ratio GBP winner/ TOP500 no. 1 system (%)	TOP500 no. 1 system	GBP system	GBP system pos. in TOP500	GBP application
1993	0.1245	0.06	48.2	Numerical Wind Tunnel	CMS	2	Modeling of a shock front using the Boltzmann equation
1994	0.1704	0.14	82.2	Numerical Wind Tunnel	Intel Paragon	2	Structural mechanics using the boundary element method
1995	0.1704	0.179	105.0	Numerical Wind Tunnel	Numerical Wind Tunnel	1	Quantum chromodynamics simulation
1996	0.3682	0.111	30.1	CP-PACS	Numerical Wind Tunnel	2	Fluid dynamics problem
1997	1.338	0.43	32.1	ASCI Red	ASCI Red	1	Motion of 322 million self-gravitating particles
1998	1.338	0.657	49.1	ASCI Red	Cray T3E	2	Modeling of metallic magnet atoms
1999	2.3796	1.18	49.6	ASCI Red	ASCI Blue Pacific	2	Fluid turbulence in compressible flows
2000	4.938	1.349	27.3	ASCI White, SP Power3	Grape-6	–	Simulation of black holes in a galactic center
2001	7.226	11.55	159.8	ASCI White, SP Power3	Grape-6	–	Simulation of black holes in a galactic center
2002	35.86	26.58	74.1	Earth Simulator	Earth Simulator	1	Global atmospheric simulation with the spectral transform method
2003	35.86	5	13.9	Earth Simulator	Earth Simulator	1	Earthquake simulation
2004	70.72	15.2	21.5	BlueGene/L beta	Earth Simulator	3	Geodynamo simulation
2005	280.6	101.7	36.2	BlueGene/L	BlueGene/L	1	Solidification simulations
2006	280.6	207.3	73.9	BlueGene/L	BlueGene/L	1	Large-scale electronic structure calculations of high-Z metals
2007	478.2	115.1	24.1	BlueGene/L	BlueGene/L	1	Simulation of Kevin–Helmholtz instability in molten metals
2008	1,105	409	37.0	Roadrunner	Cray XT4	2	Simulation of disorder effects in high-temperature superconductors
2009	1,759	1,030	58.6	Jaguar (Cray XT5)	Jaguar (Cray XT5)	1	Ab initio computation of free energies in nanoscale systems
2010	2,566	700	27.3	Tianhe-1A	Jaguar (Cray XT5)	2	Direct numerical simulation of blood flow
2011	10,510	3,080	29.3	K Computer	K Computer	1	First-principles calculations of silicon nanowire electron states
2012	17,590	4,450	25.3	Titan (Cray XK7)	K Computer	3	Astrophysical <i>N</i> -body simulation
2013	33,862.7	11,000	32.5	Tianhe-2 (MilkyWay-2)	Sequoia (BlueGene/Q)	3	Cloud cavitation collapse
2014*	33,862.7	24,770	73.1	Tianhe-2 (MilkyWay-2)	Titan (Cray XK7)	2	Gravitational tree code to simulate the Milky Way

*Finalist with the highest reported flops metric; no flops number was available for the actual winner, which ran code on Anton-2 and was selected based on a speedup metric.

IN MEMORIAM: HANS WERNER MEUER (1936–2014)

Hans W. Meuer started the TOP500 list together with Erich Strohmaier, Jack Dongarra, and Horst D. Simon in 1993. The project grew directly out of his work as cofounder and organizer in 1986 of the first Mannheim Supercomputer Seminar, an annual meeting that in 2001 became the International Supercomputing Conference (ISC).

In addition to serving as general chair of ISC from its inception until his passing, Meuer was director of the Computing Center and a professor of mathematics and computer science at the University of Mannheim, Germany, from 1976 until his retirement in 2000. From 1998 until 2013, he also served as managing director of Prometheus GmbH, a company that runs a series of conferences in fields closely associated with high-performance computing.

Prior to joining the University of Mannheim, Meuer studied mathematics, physics, and politics at the universities of Marburg, Giessen, and Vienna. From 1962 until 1973, he was a research specialist, project leader, and group and department chief at the Jülich Research Center, during which he received his PhD in applied mathematics from RWTH Aachen University. From 1986 until 2004, he served as editor in chief of the professional IT journal *Pik (Praxis der Informationsverarbeitung und Kommunikation)*. Meuer was a member of ACM and the German Informatics Society. He published numerous articles in the areas of mathematics, data processing, and computer science.

Meuer passed away at the age of 77 at his home in Daisbach, southern Germany, on 20 Jan. 2014 after a brief battle with cancer.

correlation indicates that Linpack still serves as a reasonably close upper bound for actual application performance. A more balanced benchmark might close the existing gap by a factor of 2 but would probably be regarded as overly demanding if its performance levels fell too far below the GBP levels. Figure 4 indicates that correcting performance by such a small factor would not materially affect the observed long-term trends.

TOWARD EXASCALE COMPUTING

As we approach the exascale era, the rate of increase in peak and application performance of our largest systems has clearly slowed. Our analysis suggests that over the next decade we will likely fall short of historical trends for performance increases by almost an order of magnitude (100× instead of 1,000×). An even more substantial slowdown can be expected once any change or an end to Moore's law affects increases in per-socket performance.

Any such slowdown will eventually open up opportunities for companies

to explore competitive advantages through stronger architectural differentiation. For scientific computing, this most likely will increase the difficulty of judging the appropriateness of different architectures and their potential performance without detailed measurements of specialized benchmarks and applications. At first sight, this might indicate less relevance for Linpack or the TOP500, but neither of these was ever meant to guide procurement decisions or architectural design. For these critical tasks, we must examine the scientific applications we use every day.

Augmenting the TOP500 with a more balanced benchmark might correct rankings by a small factor and provide additional angles for analysis. But to be useful and manageable, any such benchmark must be simple and scalable, enable a broad variety of implementations and optimizations, and be useable on a large set of architectures.

The current approach for compiling the TOP500 clearly cannot address truly novel architectures such as

neuromorphic or quantum systems. Should a market for such systems develop, very domain-specific benchmarking and ranking techniques would need to be developed, a situation similar to that of data-intensive computing.

A project such as the TOP500 that collects data for analysis of the HPC market and long-term technological trends is best served by building upon a broadly applicable benchmark that is flexible enough to avoid handicapping otherwise well-designed systems but demanding enough to penalize architectures that do not adequately support the scientific computing community. We acknowledge that there are limitations to the TOP500 and actively work to improve it, but we also believe that the TOP500 continues to provide important insights about how to move forward into the exascale era that are based on actual data instead of marketing claims.

The TOP500 has enjoyed incredible success as a metric for supercomputing performance for

See www.computer.org/computer-multimedia for multimedia content related to this article.

more than two decades. The trends it exposes, the focused optimization efforts it inspires, and the publicity it brings to the HPC community are critical. As increasingly diverse architectures emerge, appropriate benchmarks that match application needs are more necessary than ever.

HPL encapsulates some aspects of real scientific applications—strong demands for system reliability and stability, flops performance, and to some extent network performance—but no longer tests memory performance adequately. Alternative benchmarks that complement HPL could help correct individual rankings and improve our understanding of systems, but they are not likely to significantly alter observed technological trends.

REFERENCES

1. J.J. Dongarra, P. Luszczek, and A. Petitet, "The LINPACK Benchmark: Past, Present and Future," *Concurrency and Computation: Practice and Experience*, vol. 15, no. 9, 2003, pp. 803–820.
2. E. Strohmaier et al., "The Marketplace of High-Performance Computing," *Parallel Computing*, vol. 25, 1999, pp. 1517–1544.
3. E. Strohmaier et al., "Recent Trends in the Marketplace of High-Performance Computing," *Parallel Computing*, vol. 31, 2005, pp. 261–273.
4. P.M. Kogge and T.J. Dysart, "Using the TOP500 to Trace and Project Technology and Architecture Trends," *Proc. Int'l Conf. High Performance Computing, Networking, Storage and Analysis (SC 11)*, 2011; doi:10.1145/2063384.2063421.
5. D.H. Bailey et al., "The NAS Parallel Benchmarks—Summary and Preliminary Results," *Proc. ACM/IEEE*

ABOUT THE AUTHORS

ERICH STROHMAIER, cofounder and coeditor of the TOP500, is a senior scientist and leads the Performance and Algorithms Research Group at Lawrence Berkeley National Laboratory (LBNL). His research focuses on performance characterization, evaluation, modeling, and prediction for high-performance computing (HPC) systems; analysis of advanced computer architectures and parallel programming paradigms; and classification of and programming patterns for scientific computational kernels. Strohmaier received a PhD in theoretical physics from the University of Heidelberg. He is a member of the American Association for the Advancement of Science (AAAS), ACM, IEEE, and the American Physical Society (APS). Contact him at estrohmaier@lbl.gov.

HANS W. MEUER was a professor emeritus of mathematics and computer science at the University of Mannheim at the time of his passing on 14 January 2014. A renowned figure in the HPC community, he cofounded the annual International Supercomputing Conference as well as the TOP500 project. For more about Meuer's rich life and legacy, see the sidebar "In Memoriam: Hans Werner Meuer (1936–2014)" and www.top500.org/blog/remembering-hans-meuer.

JACK DONGARRA, coeditor of the TOP500, is an American University Distinguished Professor of Computer Science in the Department of Electrical Engineering and Computer Science and directs both the Innovative Computing Laboratory and the Center for Information Technology Research at the University of Tennessee. He is also a Distinguished Research Staff member in the Computer Science and Mathematics Division at Oak Ridge National Laboratory and holds the Turing Fellowship in the schools of Computer Science and Mathematics at the University of Manchester. Dongarra specializes in numerical algorithms in linear algebra, parallel computing, use of advanced computer architectures, programming methodology, and tools for parallel computers. He received a PhD in applied mathematics from the University of New Mexico. Dongarra is a Fellow of the AAAS, ACM, IEEE, and the Society for Industrial and Applied Mathematics (SIAM) and a member of the National Academy of Engineering (NAE). Contact him at dongarra@eecs.utk.edu.

HORST D. SIMON, coeditor of the TOP500, is deputy director of LBNL. His research interests are in the development of sparse-matrix algorithms, algorithms for large-scale eigenvalue problems, and domain decomposition algorithms for unstructured domains for parallel processing. He was a member of the NASA team that developed the NAS Parallel Benchmarks. Simon received a PhD in mathematics from the University of California, Berkeley. Contact him at hdsimon@lbl.gov.

6. J. Dongarra and M.A. Heroux, *Toward a New Metric for Ranking High Performance Computing Systems*, tech. report SAND2013-4744, Sandia National Laboratories, June 2013; www.sandia.gov/~maherou/docs/HPCG-Benchmark.pdf.
7. M.F. Adams et al., *HPGMG 1.0: A Benchmark for Ranking High Performance Computing Systems*, tech. report LBNL-6630E, Lawrence Berkeley National Laboratory, May 2014; http://crd.lbl.gov/assets/pubs_presos/hpgmg.pdf.