International Journal of
HIGH PERFORMANCE
COMPUTING APPLICATIONS

Research Paper

The International Journal of High
Performance Computing Applications
2017, Vol. 31(6) 469–484

DOI: 10.1177/1094342017738610
journals.sagepub.com/home/hpc
SAGE

A look back on 30 years of the Gordon Bell Prize

Gordon Bell[1], David H Bailey[2], Jack Dongarra[4,5], Alan H Karp[6] and Kevin Walsh[3]

Abstract

The Gordon Bell Prize is awarded each year by the Association for Computing Machinery to recognize outstanding achievement in high-performance computing (HPC). The purpose of the award is to track the progress of parallel computing with particular emphasis on rewarding innovation in applying HPC to applications in science, engineering, and large-scale data analytics. Prizes may be awarded for peak performance or special achievements in scalability and time-to-solution on important science and engineering problems. Financial support for the US$10,000 award is provided through an endowment by Gordon Bell, a pioneer in high-performance and parallel computing. This article examines the evolution of the Gordon Bell Prize and the impact it has had on the field.

Keywords

High Performance Computing, HPC Cost-Performance, HPC Progress, HPC Recognition, HPPC Award. HPC Prize, Gordon Bell Prize, Computational Science, Technical Computing, Benchmarks, HPC special hardware

1. Introduction

The Association for Computing Machinery's (ACM) Gordon Bell Prize (ACM Gordon Bell Prize Recognizes Top Accomplishments in Running Science Apps on HPC, 2016; BELL Award Winners, n.d.). The prize chronicles the important innovations and transitions of high-performance computing (HPC), beginning in 1987 when the prize was first awarded for demonstrating that Amdahl's Law (1967) was not insurmountable. Almost every gain in parallelism since then has been recognized—from widely distributed workstations to machines with 10 million processor cores. In particular, the Bell Prize highlights the remarkably rapid transition from the traditional Seymour Cray-inspired, shared-memory, multivector supercomputer paradigm to today's massively parallel distributed-memory systems, which often use graphics processing units (GPUs) and other accelerators. Some, including Gordon Bell himself, one of the coauthors, hold that the prize has actually helped stimulate the rapid adoption of new paradigms like GPUs and possibly field-programmable gate arrays as well. The prize has also recognized the value of specialized hardware for certain targeted applications. And, most importantly, the prize has served to recognize the tremendous human effort required to efficiently exploit the performance potential of new and emerging HPC systems. This trend is most obvious, for instance, by observing the dramatic increase in the number of authors associated with the prize-winning papers in recent years. The prize has also had a positive effect on the careers of many of the participants.

2. What is the prize?

The Gordon Bell Prize is awarded each year to recognize outstanding achievements in HPC. The purpose of the award is to document scientific progress at the frontier of both supercomputing architectures and important computational science applications. It aims to track the progress of parallel computing over time, with particular emphasis on rewarding innovation in applying HPC to applications in science, engineering, and large-scale data analytics. Prizes

---

[1] Microsoft Researcher (retired)
[2] Lawrence Berkeley National Laboratory (retired) and University of California, Davis, CA, USA
[3] University of California, San Diego and the San Diego Supercomputer Center
[4] University of Tennessee and Oak Ridge National Laboratory, Oak Ridge, TN, USA
[5] University of Manchester UK
[6] Earth Computing Corporation, Palo Alto, CA, USA

Corresponding author:
Jack Dongarra, University of Tennessee and Oak Ridge National Laboratory, Oak Ridge, TN 37830, USA.
Email: dongarra@icl.utk.edu

may be awarded for peak performance or special achievements in scalability and time-to-solution on important science and engineering problems.

3. The evolution of the prize

The Gordon Bell Prize was rooted in a debate in the field of HPC in the early 1980s. Some members of the HPC community argued that Amdahl's Law fundamentally limited the speedup achievable on important applications, because according to Amdahl's Law, if even 1% of the computation cannot be run in parallel, the speedup cannot possibly exceed 100. Others argued that Amdahl's Law was not a fundamental barrier and proposed constructing multiple instruction, multiple data systems equipped with thousands of processors.

In 1985, Alan Karp, then a staff scientist at IBM (and one of the co-authors of the present article), challenged these claims in a letter to the editor of *Communications of the ACM* (1986). Having earlier attended the Society for Industrial and Applied Mathematics (SIAM) 1985 Conference on Parallel Processing, Karp observed that, while there was plenty of talk at the meeting about building 1,000-processor, 10,000-processor, and even 1,000,000-processor systems, to that date no one had shown that reasonable speedups could be obtained even on much smaller systems. Growing weary of such talk and seeing no real demonstrations, in 1986 Karp offered US$100 to the first person or group to demonstrate a speedup of at least 200 times on a real problem running on a general purpose parallel processor. The offer was to last for 10 years (http://www.netlib.org/benchmark/karp-challenge). One reporter asked Karp, "Why the prize?" Karp responded, "Because I don't think anyone can do it." "Why only $100?" "Because somebody might."

Gordon Bell thought Karp's challenge was a good idea, although he did not think anyone would win it. In 1986, it was clear that the future of computing would inevitably be based on distributed memory, scalability, and parallelism. Gordon Bell was then the founding Assistant Director of NSF's Directorate for Computer and Information Science and Engineering (CISE). To keep things interesting for the duration of the challenge, Gordon Bell offered his own prize—originally set at US$1,000 for the best speedup of a real application running on a real machine. The first year's prize would be awarded to the entrant that demonstrated the highest speedup. Subsequent winners would have to double the previous winner's speedup until either the speedup hit 200 times of the sequential application or 10 years had passed—whichever came first. The award was initially loosely specified to be: for speedup on a general purpose computer, and maximum scalability or parallelism including single instruction, multiple data (SIMD) systems. Prizes were given for price performance through 2002. The prize was initially funded for 10 years, and an additional category was added for special-purpose machines. Karp, who had formed a committee to judge his prize because of criticism that the rules specified in his challenge were too easy to manipulate (Communications of the ACM, 1987), agreed to have the committee judge Gordon Bell's challenge as well.

Bell had three goals in mind for his prize:

1. Reward practical use of parallel processors;
2. Encourage improvements in hardware and software; and
3. Demonstrate the usefulness of parallel processors for real problems.

Seven entries were submitted that first year, many of them confirming the wisdom of the original challenge and the Gordon Bell Award, as it was then called. These entries reported speedups of 2.5–16 times over the corresponding sequential implementations. A group from Yale University, Bell Labs, and Caltech showed that Karp's money was at risk by achieving speedups of 40 times and 100 times on two applications. There was, however, a clear winner.

A group from Sandia National Laboratories (SNL), consisting of Robert Benner, John Gustafson, and Gary Montry, took home both the inaugural Bell and Karp awards with demonstrated speedups of 400–600 times on three applications running on a 1024-node nCUBE machine. In addition, they pointed out that real users run larger problems on faster machines, and so they calculated potential speedups of close to 1,000 times if the problem size scaled with the number of processors.

So, how did the team from SNL break Amdahl's Law? As explained later (Gustafson, 1988), a machine with more processors lets one run larger problems, owing more to the larger amount of available memory than to the required execution time. The key point is that the parallelizable part of the computation grows faster than the sequential part. So, as the problem size grows, the parallel part becomes a larger fraction of the total computation time and the sequential part shrinks.

Bell attended the first panel meeting for the judges, which included Alan Karp (then at IBM), Jack Dongarra (then at Argonne National Laboratory), and Ken Kennedy (then at Rice University, now deceased). Since there was no entrant in the special purpose category, Bell gave out several special awards, because, as he said, "It's my money, and I can do whatever I want with it." One of these special awards was for performance, work that was never submitted for the prize, and another was for computations done on an SIMD machine. Details can be found in the first publication (Alan H Karp, 1988) of the series of articles on the Bell Award, first published in *IEEE Software* and later in *IEEE Computer*. The first 4 years of the Gordon Bell Prize were presented at the Institute of Electrical and Electronics Engineers (IEEE) COMPCON conference; subsequent awards were presented at the IEEE/ACM International Conference for High Performance Computing, Networking, Storage, and Analysis (SC) series.

The Gordon Bell Prize competition, as originally specified, was only supposed to run for 10 years. And even though the first winners exceeded Bell's mark for what was intended to be a decade-long competition, Bell decided to continue offering the prize. Since the original speedup criterion had been exceeded, the rules were changed to include categories for performance, price/performance, and compiler parallelization. These categories were flexible, in keeping with the spirit of the first year, and were changed from time to time.

The rules were modified following the first year but remained essentially unchanged until ACM assumed management of the prize in 2006. Each year the judges divided the total prize money (initially US$2,000) among two winners and any honorable mentions. The recipients are selected from entries in three categories: performance, price/performance, and compiler-generated speedup. These categories are explained in further detail, below.

- For the performance prize, entrants must convince the judges that they have run their application faster than anyone else has done previously. Computing performance is typically measured in millions, billions, trillions, and (at the present time) quadrillions of floating point operations per second (flop/s) (abbreviated Mflop/s, Gflop/s, Tflop/s, and Pflop/s, respectively).
- The price/performance prize encourages the development of cost-effective supercomputing, but the rules are setup to prevent unrealistic or "stunt" machines from winning the prize. For example, a parallel job running at 1 Kflop/s on two used Z-80 processors costing US$1 is not eligible in spite of its high-performance-per-dollar ratio.
- The compiler-generated speedup prize was intended to spur the development of compilers that can automatically parallelize sequential programs.

From the beginning, the rules of the Gordon Bell Prize specifically stated that "toy problems" or "cooked-up" examples would not be allowed. The judges are looking for real problems running on real machines. Along this line, at one point some in the community rightly criticized the Bell Prize judges for what they saw as mechanically selecting the highest-performing application. In subsequent years, the judges responded by selecting submissions that were not necessarily the highest-performing submission, but which demonstrated a notable achievement for particularly challenging classes of applications.

4. ACM's management of the prize

In 2006, in cooperation with Gordon Bell, ACM assumed responsibility for managing and awarding the Bell Prize. The prize money is now awarded from an endowment granted to ACM by Gordon Bell, and the prize is now formally known as the "ACM Gordon Bell Prize." However, the overall goals of the award remain the same, 30 years after its inception. In 2012, the ACM policy changed, allowing only one prize per year. In 2017, the ACM states:

> The purpose of the award is to track the progress over time of parallel computing, with particular emphasis on rewarding innovation in applying high-performance computing to applications in science, engineering, and large-scale data analytics. Prizes may be awarded for peak performance or special achievements in scalability and time-to-solution on important science and engineering problems.

In recent years, the judges' panel has selected up to six finalists for the prize, and the finalists' papers have been included in the SC proceedings. Finalists are given the opportunity to provide updated results, usually in early August, prior to the final award decision in November. Before 2006, the panel members often made the final ruling on the winner(s) at SC itself, after listening to the individual talks presented by the finalists at the conference. Now, however, the decision is finalized before the SC meeting.

The prize is open to teams of up to 12 individuals each, and the submissions are evaluated based on rigorous criteria. Submissions are required to contain a number of components to be considered complete; please see the ACM website for full details (http://awards.acm.org/bell/nominations).

5. Charting the winners

As mentioned above, the Bell Prize has evolved through the years, and this evolution is clear when one looks at specific award designations, listed below, as they also changed over time:

- Scaling, HW (1987)
- Peak performance (1988)
- Price/performance (1988)
- Compiler parallelization (1989)
- Speedup (1992)
- Special purpose machine (1995)
- Special award for language (2002)
- Special achievement (lifetime, 2003)
- Algorithm innovation (2008)
- Sustained performance (2011)
- Scalability and time to solution (2011)
- Achievement in scalability (2015)
- 10-million core scalable atmospheric dynamics (2016)

Each year at SC, two key indicators of HPC performance are highlighted: (1) the TOP500 list (TOP 500) and (2) the ACM Gordon Bell Prize. It is interesting to compare the number one system on the TOP500 to the performance obtained for the Gordon Bell Prize over time (Figure 1). It is interesting to note that in 2001 the Gordon Bell Award achieved a performance greater than the number 1 Top500 system. In 2001, the Gordon Bell Award was given to the

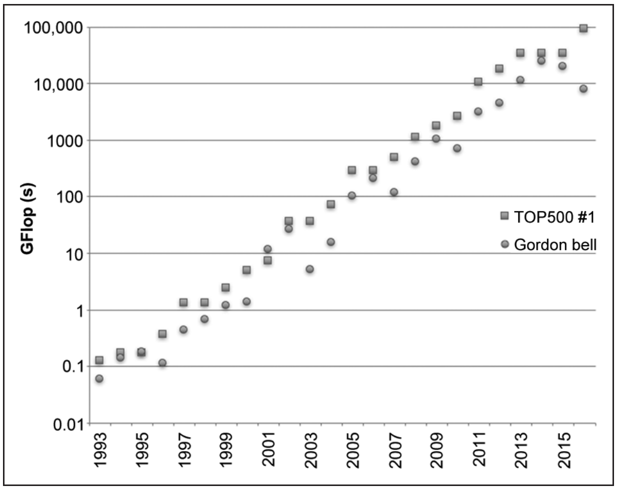

Figure 1. For the chart, we collected the application, system, and performance data from all of the ACM Gordon Bell Prize winners for best sustained performance since 1993. We also plotted the number one system from each November TOP500 list since 1993. We selected the November TOP500 editions because those are released at the same SC conferences where the Gordon Bell winners are announced. The chart shows a strong correlation between the ACM Gordon Bell Prize winners and the number one systems on the TOP500 list for each corresponding year. ACM: Association for Computing Machinery.

Grape-6 Computer. The Grape-6 (Gravity Pipe) computer is a special purpose computer which uses hardware acceleration to perform gravitational computations.

The two metrics track one another with a maximum delay of a couple of years between Linpack peak and the time it takes for prize winners to get the most performance. From the awarding of the first prize in 1987, the focus has been on parallelism. The first 4 years had substantial parallelism above conventional supercomputers, but the Top500 was not measured until 1993. Note that the Bell Prize is not judged on how effective an application is versus the peak machine. For example, the 2016 prize winning entry used all 10.5 million processors of Sunway TaihuLight to achieve a peak of 7.95 Pflop/s versus its Linpack peak of 93 Pflop/s.

6. Problem-specific hardware systems

In 1989, an honorable mention award was given for a PLA systolic array implementation that did pattern matching by comparing a new DNA sequence against a library of sequences for the closest match.

Beginning in 1995, Makino and Taiji at the University of Tokyo began building a series of Grape processors. The Grape-4 (for GRAvity piPE) system consisted of a host interface, a control board, and 40 processor boards. Each processor board held 48 custom Large Scale Integration (LSI) chips that calculated the gravitational force of astronomical bodies and its first derivative. Since each chip is capable of computing at over 0.6 Gflop/s, the total system of 1,920 processors has a peak performance of almost 1.3 Tflop/s (trillion floating-point operations per second). The calculation submitted used only 288 processors on six boards to simulate the motion of over 130,000 stars orbiting two black holes. The simulation ran at 112 Gflop/s, almost 60% of the peak rate. The Grapes evolved in 1996, 1999, 2002, 2004, and 2006. A 2,048 pipelined chip set called Grape 6, operated at a peak of 63.4 Tflop/s, to simulate 1.8 million planetesimals. Comparing the 2006 Grape, and add-in card with the 1999 system, the later system had a price performance of US$185/Gflop/s, whereas the earlier was US$7,000/Gflop/s.

The David E. Shaw Research group built two systems, Anton 1 and 2 in 2009 and 2014 for molecular dynamics (MD) simulation (Shaw, 2014). Like Grape that computes gravitational forces, MD computation calls for calculating the forces on each of the atom pairs in time steps. Anton 2 peak performance claimed:

> On a 512-node machine, Anton 2 achieves a simulation rate of 85 μs/day for dihydrofolate reductase (DHFR), a benchmark system containing 23,558 atoms that is near the limits of practical parallelism because it has less than one atom per processor core (a 512-node Anton 2 machine contains 33,792 processor cores). This rate represents a 4.5-fold improvement over Anton 1, and is 180 times faster than any implementation of MD on another platform.

7. Selected winning applications

The first three applications were submitted by the Sandia team in 1987 at Sandia. They achieved speedups of 400–600 times on an nCUBE 1024 node computer operating at 0.35 Mflop/s/node that was limited by the 128 Kbyte memory on each node. It was from these three demonstrations that Gustafson's Law was posited. Basically, the law says that a problem run on p processors can be sped up by a large fraction of p, by simply making the problem larger to cover the overhead of the serial or overhead part of the calculation.

Specialized hardware entries were pioneers for the N-body calculation. The Grapes at Tokyo University focused on gravitational interaction of stars. The DE Shaw Anton 1 and 2 computers were built for MD simulation that were used as servers for the biomolecular community. Anton 2 simulates multiple microseconds per day for systems of millions of atoms on a 1- to 5-femtosecond basis.

Programming the Sunway is a significant technical challenge to achieve the greatest scale in weather applications (Chao Yang, 2016). It simulates weather at ½ kilometer resolution with 770 billion unknowns, to deliver 0.07 simulated years per day or roughly a month per day. At the lowest level, computation is carried out by 260 core processor core boards of four, one plus an 8 × 8 array of processors with 8-TB memory. The system has 40,960 of these housed in 40 cabinets each with 4 × 256 processor boards or 40 cabinets × 4 × 256 core-processor boards × 4 × (1 + 8 × 8) core-processors/core-processor boards.

The simulation runs all 10.4 million processors at a rate 7.95 Pflop/s, versus 93 Pflop/s for Linpack.

8. The winning teams: Organization affiliation and size

Table 1 lists by year the Gordon Bell Prize winning applications with their performance. Table 2 gives the organizations associated with the prize or honorable mention when more prizes were awarded. All of the manufacturers that provided winning hardware or institutions hosting systems participated to some degree in winning prizes. For example, IBM researchers were involved in seven prizes. Nearly all of the institutions in the first column housed or manufactured HPC systems.

Unfortunately, no prizes have been given for the large cloud clusters at Amazon, Google, IBM, or Microsoft. This may be because that none of the applications use a large or noteworthy fraction of their cloud's facility, or that no one has bothered to note how much computing power is going into some of the applications.

Note that the first prizes were done by a team of three and a single researcher. The largest team, 46, built Anton 2. In 2016, the winning team had 12 members from seven different institutions including computer science, computational mathematics, climate modeling, and geo computing that covered algorithms, applications, and architecture. The 2016 scalable atmospheric framework included four coupled models of the land, atmosphere, ocean, and ice. Each of these required detailed modeling of land biology and hydrological processes; space weather and atmospheric chemistry; marine biology; and dynamic ice, respectively.

9. Impact

There are a number of prestigious prizes that recognize outstanding accomplishments: the Nobel Prizes, the Fields Medal, and the ACM Turing Award just to name a few. However, it would be hard to argue that anyone starts a project in the realistic hope of winning one of them. In that sense, these other prizes do not have much of an impact on motivating the research itself. That's not the case for the ACM Gordon Bell Prize. Many, if not most, of the entrants adapted their research to push the limits of computation in order to have a better chance of winning the award.

Sometimes the contribution is more direct. For instance, the Sandia team who won the Karp Challenge and the first Gordon Bell Prize were about to be defunded but received permission to continue their work long enough to produce an entry. Not only did they win, but that group went on to be one of the premier contributors to HPC.

As part of preparing this article, the present authors reached out to many of the prize-winning participants of the past decade. We asked several brief questions, including whether or not the efforts to prepare a Bell Prize submission resulted in significant technical advances that would otherwise likely not have been achieved, and also whether the resulting prize advanced the careers of the people involved, and whether management was supportive.

Some of these responses were quite illuminating. Examples:

1. Erik Draeger of Lawrence Livermore National Laboratory, a co-winner in 2006, wrote:

> In attempting to prepare a strong submission, we certainly took more heroic measures to squeeze out the best performance [for our QBox code] than we likely would have simply for our local customers. Specifically, some of the final ScaLAPACK communication tuning and task mapping that pushed us from ∼195 TFlop/s to ∼207 TFlop/s. It originated in our doing more detailed profiling to make sure we understood communication performance at a deep enough level to prepare a strong submission that ultimately led to our identifying additional optimizations. As is often the case, it's not until you prepare something for public release that you find the gaps in your understanding.

2. Paul Kent of Oak Ridge National Laboratory, a co-winner in 2008, wrote:

> The Gordon Bell prize effort went hand in hand with the science applications. In this case the improved performance of DCA++ enabled much larger problems to be solved, which enabled far more conclusive statements about the Hubbard model and superconductivity to be made. However, we were driven to increase the performance by exactly this goal. So I wouldn't separate out these motivations.

3. Lin-Wang Wang of Lawrence Berkeley National Lab, a co-winner of a special award for "algorithm innovation" in 2008, wrote, "Entering the Gordon Bell prize competition certainly motivated us to improve the performance of the code. We did a lot of computer science/improvement work due to the Gordon Bell Prize competition."
4. Jeff Larkin, then of Cray, a co-winner in 2008 and 2009, wrote:

> Working at a vendor, we definitely put additional effort into making sure that the teams I worked with had the best possible libraries and other performance optimizations to try to ensure their success. In one case, we provided a specially tuned version of some BLAS routines that was well-tuned for the application's particular matrix sizes. This really pushed us afterward to develop an auto-tuning framework to try to provide this level of optimization to even more applications.

5. Peter W.J. Staar of IBM Zurich, a co-winner in 2015, wrote, "Without the goal of a Gordon Bell submission, I would not have optimized the DCA++/geophysics code to scale to the entire supercomputers, nor would I have optimized these codes to that level."

Virtually all who responded to our queries said that receiving the Bell Prize gave a significant boost to their careers and has been a positive force in the field. Thomas

Table 1. List, by year since 1993, of Gordon Bell Prize–winning applications and their performance together with the number one system in each corresponding TOP500 list.[a]

Year	TOP500 #1[b]	GB Perf	Ratio	System TOP500 (Top500 began in 1993)	System GBP	Pos	GBP application
	GFLOPS	GFLOPS	GBP/T1				
1987	X	.45	X	X	nCUBE 1K	X	Beam stress analysis, surface wave simulation, unstable fluid flow model
1988	X	1	X	X	Cray Y-MP	X	Static structures
1989	X	6	X	X	CM-2	X	Seismic data processing
1990	X	14	X	X	CM-2G	X	Seismic data processing
1991	X	No GB Prize this year[c]	X	X	X	X	X
1992	X	5	X	X	Intel Delta	X	Gravity simulation of evolution of the universe
1993	0.0597	0.06	100.5%	CM5	CM5	2	Modeling of a shock front using the Boltzmann equation
1994	0.1434	0.14	97.6%	Intel Paragon	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	0.1704	0.111	65.1%	Numerical Wind Tunnel	Numerical Wind Tunnel	2	Fluid dynamics problem
1997	1.338	0.43	32.1%	ASCI Red	ASCI Red	1	Motion of 322,000,000 self-gravitating particles
1998	0.891	0.657	73.7%	Cray T3E	Cray T3E	2	Modeling of metallic magnet atoms
1999	1.608	1.18	73.4%	ASCI Blue Pacific	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	35.86	15.2	42.4%	Earth-Simulator	Earth-Simulator	3	Simulation of geodynamo
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 Kelvin–Helmholtz instability in molten metals
2008	1059	409	38.6%	Cray XT4	Cray XT4	2	Simulations of disorder effects in high-Tc
2009	1759	1030	58.6%	Jaguar (Cray XT5)	Jaguar (Cray XT5)	1	Ab initio computation of free energies in nanoscale systems
2010	1759	700	39.8%	Jaguar (Cray XT5)	Jaguar (Cray XT5)	2	Direct numerical simulation of blood flow
2011	10510	3080	29.3%	K Computer	K Computer	1	First-principles calculations of electron states of a silicon nanowire
2012	10510	4450	42.3%	K Computer	K Computer	3	Astrophysical N-body simulation
2013	17173	11000	64.1%	Sequoia (BlueGene/Q)	Sequoia (BlueGene/Q)	3	Cloud cavitation collapse
2014	17590	24770	140.8%	Titan (Cray XK7)	DE Shaw Anton 2	2	Gravitational tree-code to simulate the milky way galaxy
2015	33863	20100	59.4%	Tianhe-2	BlueGene/Q	3	Implicit solver for PDEs: flow in Earth's mantle
2016	93000	7950	8.55%	TaihuLight	TaihuLight	1	10M core implicit solver nonhydrostatic atmospheric dynamics

[a] The "Pos" column denotes the position within the TOP500 list that each Gordon Bell Prize–winning system had. The entry for 2014 shows the finalist with the highest reported FLOPS. The actual winner ran a code on Anton 2, was selected based on a speedup metric, and no FLOPS number was available.

[b] Top500 list was started in 1993.

[c] No prize was awarded in 1991 because of a decision to move the announcement from the IEEE COMPCON conference to the IEEE/ACM SC meeting.

Table 2. Organizational affiliation of prize or honorable mention winners (number of prizes from an organization proceeds the name of the organization).

7 Caltech	2 CMU	Abuques	Ohio State
7 IBM	2 Cornell	Ansoft	Old Dominion U
6 LANL	2 D E Shaw	Bejing Normal U	Penn State
6 LLNL	2 FSU	Brown	Purdue
6 ORNL	2 Fujitsu	BTL	Rutgers
6 Sandia	2 HIT	Center of Earth System	Sandia
5 Earth Sim. Ctr.	2 JAMSTEC	Columbia	Tel Aviv U
	2 Max Planck Inst	Emory	THC
3 Japan AEC	2 Mobil	Fermilab	Traco
3 NEC	2 Nagoya U	Found MST	Tsingua U
3 Cray	2 Pittsburg SC	GA Tech	U Chicago
3 ETH	2 THC	Hiroshima U	U de Louviain
3 Japan Marine	2 TMC	HNC	U IL
3 NAL	2 Tokyo Inst of Tech	IDA	U Messina
3 NYU	2 Tokyo U	Inst fir Frontier Res	U MI
3 U Tokyo	2 Tskuba U	Japan NAL	U Milano
3 Yale	2 U Colorado	Keio U	U NM
	2 U MN	LBNL	U of Bristol
4 Argonne	2 U TX	MIT	U of Chinese Acad
4 Intel	2 UC/Berkeley	Munich Tech U	U of Electro-communication
4 Riken	(118)	Nagoya U	U of IL
		NAO Japan	U of TN
		NASA Ames	U Penn
		NASA Goddard	U Sydney
		NASA Langley	UC/Davis
		Nat Space Dev	United Tech
		NCAR	Vienna U of Tech
		Next Gen SC	Wills Phy Lab
		NRL	Yamagata U
		NSC	(62)

Sterling, now at Indiana University, wrote, "Receiving the Gordon Bell Prize (with my collaborators) . . . was a notable event for the emerging concept of commodity clusters and the practicality of their use and application—it made them seem more real." Haohuan Fu of Tsinghua University, a co-winner in 2016, wrote:

> I see a lot of positive impact of the award. . . . With such a prestigious award that recognizes the best application results on the best computer systems, it is a perfect accomplishment to be recognized by both computer and application communities.

David E. Shaw, leader of the Anton project, which received Bell prizes in 2009 and 2014, wrote:

> One of my career goals has been to generate interest within the computer architecture community in the design of special-purpose supercomputers for various applications. To the extent that our two Gordon Bell Awards have helped highlight the potential of such machines (at least in one application area), I would regard this as a meaningful positive contribution to the maximization of my personal objective function.

Almost all responses we received noted that their managers were supportive of their efforts. This is a very positive sign indeed, because it means that the incredible efforts required to prepare a submission for the Bell Prize, typically involving a team of 10 or more individuals (not to mention considerable efforts by computer center staff), are widely recognized as well worth the cost, even in an era of scarce research funds and equally scarce computer resources.

9.1. Oddities and anecdotes

The first meeting of the Bell Prize judges' panel made it clear that this was not going to be a typical prize committee. For instance, the performance winner never submitted an entry, and there was no performance category that year. At some point, Gordon had heard of "great work by some guy in Colorado" and decided to add a performance category. "It's my money," was all the justification that Gordon provided or was needed.

Entries that first year were completely unmanageable. The panel had never specified what constituted a valid entry, so most entrants assumed more was better. We got hundreds, even thousands, of hard copy pages from some people. Given the deadline the committee had given itself for picking the winners, we had no choice but to ignore most of the material. We learned our lesson, and future submissions were limited to four pages with time allocated for follow up when needed.

The judging procedure worked without a problem for several years. Only a modest fraction of the entries needed any follow up, and the entrants replied promptly. This was the case until 1990, a time when email access was not something one had on a mobile phone, when the judging process took place during the first Gulf War. One of the entries considered for a prize was submitted by an entrant from Tel Aviv University in Israel. The panel had many questions, but the submitter could not always answer promptly because he had to take shelter from Saddam's Scud missile attacks. He would wait with his family in a closet in his apartment that served as an air-raid shelter until the "all clear" and would then rush into the office to check his email and respond to the questions.

The panel was also once threatened with a lawsuit. One of the winners from a previous year submitted an entry claiming a speedup of 64,000 times on a Thinking Machines computer configured with that many one-bit processors. That claim would not have been unusual, except that the application was floating-point intensive, and the machine only had 2,000 floating-point coprocessors. The judges unanimously agreed that the claim was inflated and chose other entries for winners and honorable mentions. The submitter disagreed quite vehemently, but the panel eventually convinced him to drop his threatened lawsuit.

9.2. Summary

The Gordon Bell Prizes provide insight into how scientific computing capabilities have changed over the years. The prizes for peak performance usually go to researchers using the world's fastest supercomputer at the time. For example, the 1998 winning team got access to a 1024-processor Cray T3E that was still on the factory floor in Chippewa Falls and was able to use this machine for their submission.

With the debut of the Earth Simulator in 2002, teams running applications on that system took home the Gordon Bell Prize for peak performance in 2002, 2003, and 2004. A similar pattern ensued when BlueGene/L ended the Earth Simulator's dominance in late 2004. At SC16, several of the Gordon Bell Prize entries were based on results from the Chinese Sunway TaihuLight machine that earned the top spot on the June 2016 TOP500 list.

National pride resonates each year in the national origins of the machine used by the winning team. Japan and the United States have largely trade places over the years with the US winning the first seven years, with Japan winning 8 years in a leapfrog fashion throughout the 1990s and 2000s, culminating with the Fujitsu K in 2012. China has quickly emerged as the Linpack leader beginning in 2009, lapping the United States in 2010, only to be overtaken by Japan 6 months later, in 2011. The United States was back in the lead six months later in 2012, but swapped positions with Japan one more time. China is now on a run, outpacing all other countries since June of 2013.

The LINPACK Benchmark is the measure of a machine. The Gordon Bell Prize is a measure of what a team of committed HPC practitioners can do with the capabilities of the machine. Federal funding in China, the EU, Japan, and the United States for HPC, from terascale to exascale, has followed the Top 500 list and the Gordon Bell Prize.

What perhaps began as a US$100-dollar lark by Alan Karp in 1987, the Gordon Bell Prize competition has evolved to become a legacy in its own time. The prize represents a bar that moves each year in tandem with changes in computational technology and the continuous growth and maturity of the HPC community. Judging criteria has been codified to best address the pace of advances in HPC, and the ACM, a sponsoring organization of the SC conference now manages the prize as a conference program element. Like its namesake, the prize has become an architectural milestone to measure parallel computing performance. Moreover, the prize has become a measure of collective achievement built up over time.

Future prizes most likely will be won by larger teams that work on coupled models such as the 2016 prize, or entries that are a combination of simulation, data analysis, and visualization. Surprisingly, there have been no brain inspired massively parallel specialized computers. Data analytic entries that use the very large, commercial cloud have been noticeably absent. No doubt, a team will find a problem can be solved on a quantum computer within the next few years as researchers build and plan to use such systems for all kinds of applications, such as biomolecular simulation.

Yes, there is a lot of love in HPC. We do not do this for the money. When we receive institutional support for what we choose to do as our life's work, it is priceless. The true value of the Gordon Bell Prize lies within what it enables entrants to do—where the primary rewards are achievement and recognition by the HPC community. Yes, for many there have been tangible rewards, like tenure or advancement to full professor or project leadership—with funding! For the HPC community as a whole, though, the prize represents what is possible in our time in history, and in algorithms to "bring the future closer." (West, 1986) A complete list of Gordon Bell Prize Winners can be found in Appendix 1.

Appendix 1: Gordon Bell Prize Winners (1987–2016)

1987 (Dongarra et al., 1988)

General purpose computer

First place: Robert Benner, John Gustafson, Gary Montry, Sandia National Laboratories; "Beam Stress Analysis, Surface Wave Simulation, Unstable Fluid Flow Model," 400–600 speedup on a 1024 node nCUBE.

Honorable mention: Robert Chervin, NCAR; "Global Ocean Model," 450 Mflop/s on a Cray X/MP48.

Honorable mention: Marina Chen, Yale University; Erik Benedictus, Bell Labs; Geoffrey Fox, Caltech; Jingke Li, Yale University; David Walker, Caltech; "QCD and Circuit Simulation," Speedups ranging from 39–458 on three applications run on CM hypercubes.

Honorable mention: Stavros Zenios, University of Pennsylvania; "Nonlinear network optimization," 1.5 s. Execution time on a connection machine.

1988 (Kennedy et al., 1989)

Peak performance

First place: Phong Vu, Cray Research; Horst Simon, NASA Ames; Cleve Ashcraft, Yale University; Roger Grimes and John Lewis, Boeing Computer Services; Barry Peyton, Oak Ridge National Laboratory; "Static finite element analysis," 1 Gflop/s on 8-proc. Cray Y-MP, Running time reduced from 15 min to 30 s.

Price performance

Honorable mention: Richard Pelz, Rutgers University; "Fluid flow problem using the spectral method," 800 speedup on a 1024 node nCUBE compiler parallelization.

Honorable mention: Marina Chen, Young-il Choo, Jungke Li and Janet Wu, Yale University; Eric De Benedictus, Ansoft Corp.; "Automatic parallelization of a financial application," 350 times speedup on a 1024 nCUBE and 50 times speedup on a 64 node Intel iPSC-2.

1989 (David Kuck et al., 1990)

Peak performance

First place: Mark Bromley, Harold Hubschman, Alan Edelman, Bob Lordi, Jacek Myczkowski and Alex Vasilevsky, Thinking Machines; Doug McCowan and Irshad Mufti, Mobil Research; "Seismic data processing," 6 Gflop/s on a CM-2 (also, 500 Mflop/s/US$1 M).

Honorable mention: Sunil Arvindam, University of Texas, Austin; Vipin Kumar, University of Minnesota; V. Nageshwara Rao, University of Texas, Austin; "Parallel search for VLSI design," 1100 speedup on a 1024 processor CM.

Price performance

First place: Philip Emeagwali, University of Michigan; "Oil reservoir modeling," 400 Mflop/s/US$1 M on a CM-2.

Honorable mention: Daniel Lopresti, Brown University; William Holmes, IDA Supercomputer Research Center; "DNA sequence matching," 77 k MIPs/US$1 M.

1990 (J. Dongarra et al. 1991)

Peak performance

Honorable mention: Mark Bromley, Steve Heller, Cliff Lasser, Bob Lordi, Tim McNerney, Jacek Myczkowski, Irshad Mufti, Guy Steele, Jr. and Alex Vasilevsky, Thinking Machines; Doug McCowan, Mobil Research; "Seismic data processing," 14 Gflop/s on a CM-2.

Price performance

First place: Al Geist and G. Malcom Stocks, Oak Ridge National Laboratory; Beniamino Ginatempo, University of Messina, Italy; William Shelton, US Naval Research Laboratory; "Electronic structure of a high-temperature superconductor," 800 Mflop/s/US$1 M on a 128-node Intel iPSC/860.

Compiler parallelization

Second place: Gary Sabot, Lisa Tennies and Alex Vasilevsky, Thinking Machines; Richard Shapiro, United Technologies; "Grid generation program used to solve partial differential equations," 1900 speedup on a 2048 node CM-2 (2.3 Gflop/s).

Honorable mention: Eran Gabber, Amir Averbuch and Amiram Yihudai, Tel Aviv University; "Parallelizing Pascal Compiler," 25× on a 25 node sequent symmetry.

1991

No prize awarded

Up to this point he Bell Prize was awarded at the COMPCON meeting and the venue was changed in 1991 to the SC conference series. As a result no prize was awarded in 1991.

1992 (Miura et al., 1993)

Peak performance

First place: Michael Warren, Los Alamos National Laboratory; John K. Salmon, Caltech; "Simulation of 9 million gravitating stars by parallelizing a tree code," 5 Gflop/s on an Intel Touchstone Delta.

Price performance

First place: Hisao Nakanishi and Vernon Rego, Purdue University; Vaidy Sunderam, Emory University; "Simulation of polymer chains parallelized over a heterogeneous collection of distributed machines," 1 Gflop/s/US$1 M.

Speedup

First place: Mark T. Jones and Paul Plassmann, Argonne National Laboratory; "Large, sparse linear system solver that enabled the solution of vortex configurations in superconductors and the modeling of the vibration of piezoelectric crystals," 4 Gflop/s on an Intel Touchstone Delta. Speedups between 350 and 500.

1993 (Heller et al., 1994)

Peak performance

First place: Lyle N. Long and Matt Kamon, Penn. State University; Denny Dahl, Mark Bromley, Robert Lordi, Jack Myczkowski and Richard Shapiro, Thinking Machines; "Modeling of a shock front using the Boltzmann Equation," 60 Gflop/s on a 1024 processor CM-5.

Honorable mention: Peter S. Lomdahl, Pablo Tamayo, Niels Gronbech-Jensen and David M. Beazley, Los Alamos National Laboratory; "Simulating the micro-structure of grain boundaries in solids," 50 Gflop/s on a 1024 processor CM-5.

Price/performance

First place: Robert W. Means and Bret Wallach, HNC Inc.; Robert C. Lengel Jr., Tracor Applied Sciences; "Image analysis using the bispectrum analysis algorithm," 6.5 Gflop/s/US$1 M on a custom-built machine called SNAP.

1994 (Heath et al., 1995)

Peak performance

First place: David Womble, David Greenberg, Stephen Wheat and Robert Benner, Sandia National Laboratories; Marc Ingber, University of New Mexico; Greg Henry and Satya Gupta, Intel; "Structural mechanics modeling using the boundary element method," 140 Gflop/s on a 1904 node Intel Paragon.

Price/performance

First place: Stefan Goedecker, Cornell University; Luciano Colombo, Università di Milano; "Quantum mechanical interactions among 216 silicon atoms," 3 Gflop/s/US$1 M on a cluster of eight HP workstations.

Honorable mention: H. Miyoshi, Foundation for Promotion of Material Science and Technology of Japan, M. Fukuda, T. Nakamura, M. Tuchiya, M. Yoshida, K. Yamamoto, Y. Yamamoto, S. Ogawa, Y. Matsuo and T. Yamane National Aerospace Laboratory; M. Takamura, M. Ikeda, S. Okada, Y. Sakamoto, T. Kitamura and H. Hatama, Fujitsu Limited; M. Kishimoto, Fujitsu Laboratories Limited; "Isotropic Turbulence and other CFD codes," 120 Gflop/s on a 140 processor Numerical Wind Tunnel.

1995 (Geist et al., 1996)

Peak performance

First place: Masahiro Yoshida, Masahiro Fukuda and Takashi Nakamura, National Aerospace Laboratory (Japan); Atushi Nakamura, Yamagata University; Shini Hoiki, Hiroshima University; "Quantum chromodynamics simulation," 179 Gflop/s on 128 processors of the Numerical Wind Tunnel.

Price/performance

First place: Panayotis Skordos, MIT; "Modeling of air flow in flue pipes," 3.6 Gflop/s/US$1 M on a cluster of 20 HP workstations.

Special-purpose machines

First place: Junichiro Makino and Makoto Taiji, University of Tokyo; "Simulation of the motion of 100,000 stars," 112 Gflop/s using the Grape-4 machine with 288 processors.

1996 (Bailey et al., 1997)

Peak performance

First place: Toshiyuki Iwamiya, Masahiro Yoshida, Yuichi Matsuo, Masahiro Fukuda and Takashi Nakamura, National Aerospace Laboratory (Japan); " Fluid dynamics problem," 111 Gflop/s on 166-processor Numerical Wind Tunnel.

Honorable mention: Toshiyuki Fukushige and Junichiro Makino, University of Tokyo; "Simulation of the motion of 780,000 stars," 333 Gflop/s using the Grape-4 machine w/1269 processors.

Price/performance

First place: Adolfy Hoisie, Cornell University; Stefan Goedecker and Jurg Hutter, Max Planck Institute; "Electronic structures calculations," 6.3 Gflop/s/US$1 M on an SGI Power Challenge with 6 MIPS R8000 processors.

1997 (Karp et al., 1998)

Peak performance

First prize-part 1: Michael S. Warren, Los Alamos, National Laboratory; John K. Salmon, Caltech; "Simulating the motion of 322,000,000 self-gravitating particles," 430 Gflop/s on ASCI Red using 4096 processors.

Price/performance

First prize: Nhan Phan-Thien and Ka Yan Lee, University of Sidney; David Tullock, Los Alamos National Laboratory; "Modeling suspensions," 10.8 Gflop/s/US$1 M on 28 DEC Alpha machines.

First prize-part 2: Michael S. Warren, Los Alamos, National Laboratory; John K. Salmon, Caltech; Donald J. Becker, NASA Goddard; M. Patrick Goda, Los Alamos National Laboratory; Thomas Sterling, Caltech; Gregoire S. Winckelmans, Universite Catholique de Louvain (Belgium); "Two problems: vortex fluid flow modeled with 360,000 particles; galaxy formation following 10,000,000 self-gravitating particles," 18 Gflop/s/US$1 M on a cluster of 16 Intel Pentium Pros (200 MHz).

1998

Peak performance

First prize: Balazs Ujfalussy, Xindong Wang, Xiaoguang Zhang, Donald M. C. Nicholson, William A. Shelton and G. Malcolm Stocks, Oak Ridge National Laboratory; Andrew Canning, Lawrence Berkeley National Laboratory; Yang Wang, Pittsburgh Supercomputing Center; Balazs L. Gyorffy, H. H. Wills Physics Laboratory, UK; "First principles calculation, of a unit cell (512 atoms) model of non-collinear magnetic arrangements for metallic magnets using a variation of the locally self-consistent multiple scattering method," 657 Gflop/s on a 1024-PE Cray T3E system (600 MHz).

Second prize: Mark P. Sears, Sandia National Laboratories; Ken Stanley, University of California, Berkeley; Greg Henry, Intel; "Electronic structures: a silicon bulk periodic unit cell of 3072 atoms, and an aluminum oxide surface unit cell of 2160 atoms, using a complete dense generalized Hermitian eigenvalue-eigenvector calculation," 605 Gflop/s on the ASCI Red machine with 9200 processors (200 MHz).

Price/performance

First prize: Dong Chen, MIT; Ping Chen, Norman H. Christ, George Fleming, Chulwoo Jung, Adrian Kahler, Stephen Kasow, Yubing Luo, Catalin Malureanu and Cheng Zhong Sui, Columbia University; Robert G.

Edwards and Anthony D. Kennedy, Florida State University; Alan Gara, Robert D. Mawhinney, John Parsons, Pavlos Vranas and Yuri Zhestkov, Columbia University; Sten Hansen, Fermi National Accelerator Laboratory; Greg Kilcup, Ohio State University; "3 lattice quantum chromodynamics computations," 79.7 Gflop/s/US$1 M on a custom system with 2048 PE's using a Texas Instruments chip (32-bit floating point ops.).

Second prize: Michael S. Warren, Timothy C. Germann, Peter S. Lomdahl and David M. Beazley, Los Alamos National Laboratory; John K. Salmon, Caltech; "Simulation of a shock wave propagating through a structure of 61 million atoms," 64.9 Gflop/s/US$1 M using a 70 PE system of DEC Alpha's (533 MHz).

1999

Peak performance

First prize: A. A. Mirin, R. H. Cohen, B. C. Curtis, W. P. Dannevik, A. M. Dimits, M. A. Duchaineau, D. E. Eliason and D. R. Schikore, Lawrence Livermore National Laboratory; S. E. Anderson, D. H. Porter and R. Woodward, University of Minnesota; L. J. Shieh and S. W. White, IBM; "Very high resolution simulation of fluid turbulence in compressible flows," 1.18 Tflop/s on short run on 5832 CPU's on ASCI Blue Pacific, 1.04 Tflop/s sustained on 1-h run, 600 Gflop/s on 1-week run on 3840 CPUs.

Price/performance

First prize: Atsuchi Kawai, Toshiyuki Fushushige and Junichiro Makino, University of Tokyo; "Astrophysical n-body simulation," 144 Gflop/s/US$1 M on custom-built GRAPE-5 32-processor system.

Special

First prize, shared: W. K. Anderson, NASA Langley Research Center; W. D. Gropp, D, K. Kaushik, B.F. Smith, Argonne National Laboratory; D. E. Keyes, Old Dominion University, Lawrence Livermore National Laboratory, and ICASE, NASA Langley Research Center; "Unstructured tetrahedral mesh fluid dynamics using PETSc library," 156 Gflop/s on 2048 nodes of ASCI Red, using one CPU per node for computation.

First prize, shared: H. M. Tufo, University of Chicago; P. F. Fischer, Argonne National Laboratory; "Spectral element calculation using a sparse system solver," 319 Gflop/s on 2048 nodes of ASCI Red, using two CPU's per node for computation.

2000

Peak performance

Competitors for this year's prize for best performance tied, each achieving 1.34 Tflop/s.

First place: Tetsu Narumi, Ryutaro Susukita, Takahiro Koishi, Kenji Yasuoka, Hideaki Furusawa, Atsushi Kawai and Thoshikazu Ebisuzaki; "Molecular Dynamic Simulation for NaCl for a Special Purpose Computer: MDM," 1.34 Tflop/s.

First place: Junichiro Makino, Toshiyuki Fukushige and Masaki Koga; "Simulation of Black Holes in a Galactic Center on GRAPE-6," 1.349 Tflop/s.

Price/performance

First place: Douglas Aberdeen, Jonathan Baxter and Robert Edwards; "92 cents/Mflops Ultra-Large Scale Neural Network Training on a PIII Cluster."

Honorable Mention: Thomas Hauser, Timothy I. Mattox, Raymond P. LeBeau, Henry G. Dietz and P. George Huang, University of Kentucky; "High-Cost CFD on a Low-Cost Cluster."

Special

Alan Calder, B.C. Curtis, Jonathan Dursi, Bruce Fryxell, G. Henry, P. MacNeice, Kevin Olson, Paul Ricker, Robert Rosner, Frank Timmes, Henry Tufo, James Truran and Michael Zingale; "High-Performance Reactive Fluid Flow Simulations Using Adaptive Mesh Refinement on Thousands of Processors."

2001

Peak performance

Toshiyuki Fukushige and Junichiro Makino; "Simulation of black holes in a galactic center," 11.55 Tflop/s.

Price/performance

Joon Hwang, Seung Kim and Chang Lee, "Study of impact locating on aircraft structure," by low-cost cluster cost 24.6 cents $\mathrm{Mflop}^{-1}\ \mathrm{s}^{-1}$, or less than 25 cents per 1-million floating operations per second.

Special

Gabrielle Allen, Thomas Dramlitsch, Ian Foster, Nick Karonis, Matei Ripeanu, Edward Seidel and Brian Toonen for supporting efficient execution in the heterogeneous distributed computing environments with Cactus and Globus.

2002

Peak performance

Satoru Shingu, Yoshinori Tsuda, Wataru Ohfuchi, Kiyoshi Otsuka, Earth Simulator Center, Japan Marine Science and Technology Center; Hiroshi Takahara, Takashi Hagiwara, Shin-ichi Habata, NEC Corporation; Hiromitsu Fuchigami, Masayuki Yamada, Yuji Sasaki, Kazuo Kobayashi, NEC Informatec Systems; Mitsuo Yokokawa, National Institute of Advanced Industrial Science and Technology; Hiroyuki Itoh, National Space Development Agency of Japan. "A

26.58 Tflops Global Atmospheric Simulation with the Spectral Transform Method on the Earth Simulator," 26.58 Tflop/s simulation of a complex climate system using an atmospheric circulation model called AFES.

Special award for language

Hitoshi Sakagami, Himeji Institute of Technology; Hitoshi Murai, Earth Simulator Center, Japan Marine Science and Technology Center; Yoshiki Seo, NEC Corporation; Mitsuo Yokokawa, Japan Atomic Energy Research Institute; "14.9 Tflop/s Three-dimensional Fluid Simulation for Fusion Science with HPF on the Earth Simulator," 14.9 Tflop/s run of a parallelized version of IMPACT-3D, an application written in High Performance Fortran that simulates the instability in an imploding system, such as the ignition of a nuclear device.

Special

Mitsuo Yokokawa, Japan Atomic Energy Research Institute; Ken'ichi Itakura, Atsuya Uno, Earth Simulator Center, Japan Marine Science and Technology Center; Takashi Ishihara, Yukio Kaneda, Nagoya University; "16.4-Tflops Direct Numerical Simulation of Turbulence by a Fourier Spectral Method on the Earth Simulator." New methods for handling the extremely data-intensive calculation of a three-dimensional Fast Fourier Transform on the Earth Simulator have allowed researchers to overcome a major hurdle for high-performance simulations of turbulence.

Manoj Bhardwaj, Kendall Pierson, Garth Reese, Tim Walsh, David Day, Ken Alvin, James Peery, Sandia National Laboratories; Charbel Farhat, Michel Lesoinne, University of Colorado at Boulder; "Salinas: A Scalable Software for High Performance Structural and Solid Mechanics Simulation." The structural mechanics community has embraced Salinas, engineering software over 100,000 lines long that has run on a number of advanced systems, including a sustained 1.16 Tflop/s performance on 3,375 ASCI White processors.

James C. Phillips, Gengbin Zheng, Sameer Kumar, Laxmikant V. Kale, University of Illinois at Urbana-Champaign; "NAMD: Biomolecular Simulation on Thousands of Processors." Researchers achieved unprecedented scaling of NAMD, a code that renders an atom-by-atom blueprint of large biomolecules and biomolecular systems.

2003

Peak performance

Dimitri Komatitsch, Chen Ji, and Jeroen Tromp, California Institute of Technology; and Seiji Tsuboi, Institute for Frontier Research on Earth Evolution, JAMSTEC; "A 14.6 Billion Degrees of Freedom, 5 Tflop/s, 2.5 Terabyte Earthquake Simulation on the Earth Simulator." The researchers used 1,944 processors of the Earth Simulator to model seismic wave propagation resulting from large earthquakes.

Special achievement

Volkan Akcelik, Jacobo Bielak, Ioannis Epanomeritakis, Antonio Fernandez, Omar Ghattas, Eui Joong Kim, Julio Lopez, David O'Hallaron and Tiankai Tu, Carnegie Mellon University; George Biros, Courant Institute, New York University; and John Urbanic, Pittsburgh Supercomputing Center; "High Resolution Forward and Inverse Earthquake Modeling on Terascale Computers." The researchers developed earthquake simulation algorithms and tools and used them to carry out simulations of the 1994 Northridge earthquake in the Los Angeles Basin using 100 million grid points.

Special achievement ("lifetime")

Junichiro Makino and Hiroshi Daisaka, University of Tokyo; Eiichiro Kokubo, National Astronomical Observatory of Japan; and Toshiyuki Fukushige, University of Tokyo; "Performance Evaluation and Tuning of GRAPE-6—Towards 40 'Real' Tflop/s." The researchers benchmarked GRAPE-6, a sixth-generation special-purpose computer for gravitational many-body problems, and presented the measured performance for a few real applications with a top speed of 35.3 Tflop/s.

2004

Peak performance

Akira Kageyama, Masanori Kameyama, Satoru Fujihara, Masaki Yoshida, Mamoru Hyodo, and Yoshinori Tsuda, JAMSTEC; "A 15.2 Tflops Simulation of Geodynamo on the Earth Simulator," 15.2 TFlop s^{-1} on 4096 processors of the Earth Simulator.

Special

Mark. F. Adams, Sandia National Laboratories; Harun H. Bayraktar, Abuqus Corp.; Tony M. Keaveny and Panayiotis Papadopoulos, University of California, Berkeley; "Ultrascalable implicit finite element analyses in solid mechanics with over half a billion degrees of freedom."

2005

Peak performance

Frederick H. Streitz, James N. Glosli, Mehul V. Patel, Bor Chan, Robert K. Yates, Bronis R. de Supinski, Lawrence Livermore National Laboratory; James Sexton and John A. Gunnels, IBM; "100+ TFlop Solidification Simulations on BlueGene/L." The team achieved up to 107 Tflop/s (trillion operations per second) with a sustained rate of 101.7 Tflop/s over a 7-h run on the IBM BlueGene/L's 131,072 processors.

2006 (ACM takes over the Gordon Bell Prize)

Price/performance

Francois Gygi University of California, Davis; Erik W. Draeger, Martin Schulz and Bronis R. de Supinski, Lawrence Livermore National Laboratory; John A. Gunnels, Vernon Austel and James C. Sexton, IBM Watson Research Center; Franz Franchetti, Carnegie Mellon University; Stefan Kral, Christoph W. Ueberhuber and Juergen Lorenz; Vienna University of Technology; “Large-scale electronic structure calculations of high-Z metals on the BlueGene/L platform.” A sustained peak performance of 207.3 Tflop/s was measured on 65,536 nodes, corresponding to 56.5% of the theoretical full machine peak using all 128 k CPUs.

https://dl.acm.org/citation.cfm? id=1188502

Honorable mention: Tetsu Narumi, Yousuke Ohno, Noriaki Okimoto, Takahiro Koishi, Atsushi Suenaga, Futatsugi, Ryoko Yanai, Ryutaro Himeno, Shigenori Fujikawa and Makoto Taiji, all of RIKEN; and Mitsuru Ikei, Intel Corp.; “A 185 Tflop/s Simulation of Amyloid-forming Peptides from Yeast Prion Sup35 with the Special-Purpose Computer System MD-GRAPE3.”

Special achievement

Pavlos Vranas, Gyan Bhanot, Matthias Blumrich, Dong Chen, Alan Gara, Philip Heidelberger, Valentina Salapura and James C. Sexton, all of IBM Watson Research Center; “The BlueGene/L supercomputer and quantum ChromoDynamics,” QCD simulation that achieved 12.2 Tflop/s sustained performance with perfect speedup to 32 k CPU cores.

https://dl.acm.org/citation.cfm? id=1188507

2007

Peak performance

James N. Glosli, David F. Richards, Kyle J. Caspersen, Robert E. Rudd and Frederick H. Streitz, all of Lawrence Livermore National Laboratory; and John Gunnels of IBM Watson Research Center; “Extending Stability Beyond CPU Millennium: A Micron-Scale Simulation of Kelvin-Helmholtz Instability.” The team that won the 2005 Gordon Bell Prize for a simulation investigating the solidification in tantalum and uranium at extreme temperatures and pressure, with simulations ranging in size from 64,000 atoms to 524 million atoms, used an expanded machine to conduct simulations of up to 62.5 billion atoms. The optimized ddcMD code is benchmarked at 115.1 Tflop/s in their scaling study and 103.9 Tflop/s in a sustained science run.

https://dl.acm.org/citation.cfm? id=1362700

2008

Peak performance

Gonzalo Alvarez, Michael S. Summers, Don E. Maxwell, Markus Eisenbach, Jeremy S. Meredith, Thomas A. Maier, Paul R. Kent, Eduardo D’Azevedo and Thomas C. Schulthess, all of Oak Ridge National Laboratory; and Jeffrey M. Larkin and John M. Levesque, both of Cray, Inc.; “New Algorithm to Enable 400+ Tflop/s Sustained Performance in Simulations of Disorder Effects in High-Tc.”

https://dl.acm.org/citation.cfm? id=1413433

Algorithm innovation

Lin-Wang Wang, Byounghak Lee, Hongzhang Shan, Zhengji Zhao, Juan Meza, Erich Strohmaier, and David H. Bailey, Lawrence Berkeley National Laboratory; “Linear Scaling Divide-and-Conquer Electronic Structure Calculations for Thousand Atom Nanostructures,” for special achievement in HPC for their research into the energy harnessing potential of nanostructures. Their method, which was used to predict the efficiency of a new solar cell material, achieved impressive performance and scalability.

https://dl.acm.org/citation.cfm? id=1413437

2009

Peak performance

Markus Eisenbach and Donald M. Nicholson, Oak Ridge National Laboratory; Cheng-gang Zhou, J.P. Morgan Chase; Gregory Brown, Florida State University; Jeffrey Larkin, Cray Inc.; and Thomas Schulthess, ETH Zurich; “A scalable method for ab initio computation of free energies in nanoscale systems,” on the Cray XT5 system at ORNL, sustaining 1.03 Pflop/s in double precision on 147,464 cores.

https://dl.acm.org/citation.cfm? id=1654125

Price/performance

Tsuyoshi Hamada, Nagasaki University; Tetsu Narumi, University of Electro-Communications, Tokyo; Rio Yokota, University of Bristol; Kenji Yasuoka, Keio University, Yokohama; Keigo Nitadori and Makoto Taiji, RIKEN Advanced Science Institute; “42 Tflop/s hierarchical N-body simulations on GPUs with applications in both astrophysics and turbulence.” The maximum corrected performance is 28.1 Tflop/s for the gravitational simulation, which results in a cost performance of 124 Mflop/s/US$1 M.

https://dl.acm.org/citation.cfm? id=1654123

Special category

David E. Shaw, Ron O. Dror, John K. Salmon, J. P. Grossman, Kenneth M. Mackenzie, Joseph A. Bank, Cliff Young, Martin M. Deneroff, Brannon Batson, Kevin J. Bowers, Edmond Chow, Michael P. Eastwood, Douglas J. Ierardi, John L. Klepeis, Jeffrey S. Kuskin, Richard H. Larson, Kresten Lindorff-Larsen, Paul Maragakis, Mark A. Moraes, Stefano Piana, Yibing Shan and Brian Towles, all of D.F. Shaw Research; “Millisecond-scale molecular dynamics simulations on Anton.”

https://dl.acm.org/citation.cfm? id=1654099

2010

Peak performance

Abtin Rahimian and Ilya Lashuk, Georgia Tech; Shravan Veerapaneni, NYU; Aparna Chandramowlishwaran, Dhairya Malhotra, Logan Moon and Aashay Shringarpure, Georgia Tech; Rahul Sampath and Jeffrey Vetter, Oak Ridge National Laboratory; Richard Vuduc and George Biros, Georgia Tech; Denis Zorin, NYU; "Petascale Direct Numerical Simulation of Blood Flow on 200 k Cores and Heterogeneous Architectures," achieved 0.7 Pflop/s of sustained performance on Jaguar.

https://dl.acm.org/citation.cfm? id=1884648

Honorable mention (first): Anton Kozhevnikov, Institut for Theoretical Physics, ETH Zurich; Adolfo G. Eguiluz, The University of Tennessee, Knoxville; and Thomas C. Schulthess, Swiss National Supercomputer Center and Oak Ridge National Laboratory; "Toward First Principles Electronic Structure Simulations of Excited States and Strong Correlations in Nano- and Materials Science."

Honorable mention (second): Tsuyoshi Hamada, Nagasaki University; and Keigo Nitadori, RIKEN Advanced Science Institute; "190 Tflops Astrophysical N-body Simulation on a Cluster of GPUs."

2011

Sustained performance

Yukihiro Hasegawa, Next-Generation Supercomputer R&D Center, Riken; Jun-Ichi Iwata, Miwako Tsuji and Daisuke Takahashi, University of Tskuba; Atsushi Oshiyama, University of Tokyo; Kazuo Minami, Taisuke Boku, University of Tskuba; Fumiyoshi Shoji, Atsuya Uno and Motoyoshi Kurokawa, Next-Generation Supercomputer R&D Center, Riken; Hikaru Inoue and Ikuo Miyoshi, Fujitsu Ltd.; and Mitsuo Yokokawa, Next-Generation Supercomputer R&D Center, Riken; "First-principles calculations of electron states of a silicon nanowire with 100,000 atoms on the K computer." A 3.08 Pflop/s sustained performance was measured for one iteration of the SCF calculation in a 107,292-atom Si nanowire calculation using 442,368 cores, which is 43.63% of the peak performance of 7.07 Pflop/s.

https://dl.acm.org/citation.cfm? id=2063386

Scalability and time to solution

Takashi Shimokawabe, Takayuki Aoki, Tomohiro Takaki, Toshio Endo, Akinori Yamanaka, Naoya Maruyama, Akira Nukada, and Satoshi Matsuoka, all of Tokyo Institute of Technology; "Petascale phase-field simulation for dendritic solidification on the TSUBAME 2.0 supercomputer," simulations on the GPU-rich TSUBAME 2.0 supercomputer at the Tokyo Institute of Technology have demonstrated good weak scaling and achieved 1.017 Pflop/s in single precision for our largest configuration, using 4000 GPUs along with 16,000 CPU cores.

https://dl.acm.org/citation.cfm? id=2063388

Because of the unusually high quality of all of the ACM Gordon Bell Prize finalists, the committee took the unusual step of awarding Honorable Mentions to the remaining three finalists papers:

"Atomistic nanoelectronics device engineering with sustained performances up to 1.44 Pflop/s" by Mathieu Luisier et al., "Petaflop biofluidics simulations on a two million-core system," by Simone Melchionna et al., and "A new computational paradigm in multiscale simulations: Application to brain blood flow," by Leopold Grinberg et al.

2012

Scalability and time to solution

Tomoaki Ishiyama, Keigo Nitadori, University of Tskuba; and Junichiro Makino, Tokyo Institute of Technology; "4.45 Pflops astrophysical N-body simulation on K computer: the gravitational trillion-body problem," The average performance on 24,576 and 82,944 nodes of K computer are 1.53 and 4.45 Pflop/s, which correspond to 49% and 42% of the peak speed.

https://dl.acm.org/citation.cfm? id=2388996.2389003

2013

Best performance of a high-performance application

Diego Rossinelli, Babak Hejazialhosseini, Panagiotis Hadjidoukas and Petros Koumoutsakos, all of ETH Zurich; Costas Bekas and Alessandro Curioni of IBM Zurich Research Laboratory; and Steffen Schmidt and Nikolaus Adams of Technical University Munich; "11 Pflop/s simulations of cloud cavitation collapse," high throughput simulations of cloud cavitation collapse on 1.6 million cores of Sequoia reaching 55% of its nominal peak performance, corresponding to 11 Pflop/s.

https://dl.acm.org/citation.cfm? id=2503210.2504565

2014

Best performance of a high-performance application

David E. Shaw, J.P. Grossman, Joseph A. Bank, Brannon Batson, J. Adam Butts, Jack C. Chao, Martin M. Deneroff, Ron O. Dror, Amos Even, Christopher H. Fenton, Anthony Forte, Joseph Gagliardo, Gennette Gill, Brian Greskamp, C. Richard Ho, Douglas J. Ierardi, Lev Iserovich, Jeffrey S. Kuskin, Richard H. Larson, Timothy Layman, Li-Siang Lee, Adam K. Lerer, Chester Li, Daniel Killebrew, Kenneth M. Mackenzie, Shark Yeuk-Hai Mok, Mark A. Moraes, Rolf Mueller, Lawrence J. Nociolo, Jon L. Peticolas, Terry Quan, Daniel Ramot, John K. Salmon, Daniele P.

Scarpazza, U. Ben Schafer, Naseer Siddique, Christopher W. Snyder, Jochen Spengler, Ping Tak Peter Tang, Michael Theobald, Horia Toma, Brian Towles, Benjamin Vitale, Stanley C. Wang and Cliff Young: all of D.E. Shaw Research; "Anton 2: raising the bar for performance and programmability in a special-purpose molecular dynamics supercomputer." Anton 2 is the first platform to achieve simulation rates of multiple microseconds of physical time per day for systems with millions of atoms. Demonstrating strong scaling, the machine simulates a standard 23,558-atom benchmark system at a rate of 85 μs day^{-1}—180 times faster than any commodity hardware platform or general-purpose supercomputer.

https://dl.acm.org/citation.cfm? id=2683599

2015

Outstanding achievement in HPC scalability

Johann Rudi and Tobin Isaac, Omar Ghattas University of Texas at Austin; A. Cristiano I. Malossi, Peter W. J. Staar, Yves Ineichen, Costas Bekas, Alessandro Curioni, IBM Research, Zurich; Georg Stadler, New York University; and Michael Gurnis, Caltech; "An extreme-scale implicit solver for complex PDEs: highly heterogeneous flow in earth's mantle," scaled to 1.5 million cores for severely nonlinear, ill-conditioned, heterogeneous, and anisotropic PDEs.

https://dl.acm.org/citation.cfm? id=2807675

2016

Outstanding achievement in HPC scalability

Chao Yang Chinese Academy of Sciences, China and University of Chinese Academy of Sciences, China; Wei Xue Tsinghua University, China and National Supercomputing Center in Wuxi, China; Haohuan Fu Tsinghua University, China and National Supercomputing Center in Wuxi, China; Hongtao You National Research Center of Parallel Computer Engineering and Technology, China; Xinliang Wang Tsinghua University, China; Yulong Ao Chinese Academy of Sciences, China and University of Chinese Academy of Sciences, China; Fangfang Liu Chinese Academy of Sciences, China and University of Chinese Academy of Sciences, China; Lin Gan Tsinghua University, China and National Supercomputing Center in Wuxi, China; Ping Xu Tsinghua University, China; Lanning Wang Beijing Normal University, China; Guangwen Yang Tsinghua University, China and National Supercomputing Center in Wuxi, China; Weimin Zheng Tsinghua University, China; "10M-core scalable fully-implicit solver for nonhydrostatic atmospheric dynamics," The fully-implicit solver successfully scales to the entire system of the Sunway TaihuLight supercomputer with over 10.5 M heterogeneous cores, sustaining an aggregate performance of 7.95 Pflop/s in double-precision, and enables fast and accurate atmospheric simulations at the 488-m horizontal resolution (over 770 billion unknowns) with 0.07 simulated-years-per-day.

https://dl.acm.org/citation.cfm? id=3014904.3014912

Declaration of Conflicting Interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The author(s) received no financial support for the research, authorship, and/or publication of this article.

References

ACM Gordon Bell Prize Recognizes Top Accomplishments in Running Science Apps on HPC (2016) From SC16 supercomputing conference. Available at: http://sc16.supercomputing.org/2016/08/25/acm-gordon-bell-prize-recognizes-top-accomplishments-running-science-apps-hpc/ (accessed November 2016).

Amdahl GM (1967) Validity of the single-processor approach to achieving large scale. In: *AFIPS Conference Proceedings,* Vol. 30, Atlantic City, NJ, 18–20 April 1967, pp. 483–485, AFIPS Press, Reston, Virginia.

Bailey D, Karp AH and Geist A (1997) 1996 Gordon Bell Prize Winners. *Computer* 30:. 80–85. DOI:10.1109/MC.1997.562930.

Browne J, Dongarra J, Karp A, et al. (1989) 1988 Gordon Bell Prize. *IEEE Software* 6(3): 78–85.

BELL Award Winners (2016) Available at: http://awards.acm.org/ bell/award-winners

Chao Yang WX (2016) SC '16. In: *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis,* Salt Lake, UT, 13–18 November 2016. DOI: 10.1109/SC.2016.5.

Communications of the ACM (1986) The Karp challenge. *Communications of the ACM Forum* 29: 87.

Communications of the ACM (1987) Reply to Chanderjit Bajaj and John Rice's comments on "Contest Rules". *Communications of the ACM* 30(1): 7–9.

Dongarra J, Karp AH and Kennedy K (1988) Soft News. *IEEE Software* 5: 102–112, May/June 1988, DOI:10.1109/MS.1988.10021 The GB part is 108–113.

David Kuck AH (1990) Special report: 1989 Gordon Bell Prize. *IEEE Software* 7: 100–104, 110

Dongarra J, Simon H, Miura K, et al. (1991) Special Report: 1990 Gordon Bell Prize Winners. *IEEE Software* 8: 92–97. DOI:10.1109/MS.1991.10026.

Geist A, Heath M and Karp AH (1996) 1995 Gordon Bell Prize Winners. *Computer* 29: 79–85, DOI:10.1109/MC.1996.481494.

Heller D, Karp AH and Simon H (1994) 1993 Gordon Bell Prize Winners. *Computer* 27: 69–75. DOI:10.1109/MC.1994.10000.

Heath M, Karp AH, Heller D, et al. (1995) 1994 Gordon Bell Prize Winners. *Computer* 28: 68–74. DOI:10.1109/2.362627.

John LG, Montry GR and Benner RE (1988) Development of Parallel Methods For a 1024-Processor Hypercube. *SIAM Journal on Scientific and Statistical Computing* 9(4): 609–638. ustafson

Karp AH, Bailey DH and Lusk E (1988) 1997 Gordon Bell Prize Winners. *Computer* 31: 86–92. DOI:10.1109/MC.1998.641981.

Kennedy K, Karp A, Kuck D, et al. (1988) Gordon Bell Prize. *IEEE Software* 6: 78–85. DOI:10.1109/52.28127.

Kuck D, Karp AH, Dongarra J, et al. (1990) Special Report: 1989 Gordon Bell Prize. *IEEE Software* 7: 100–104. DOI:10.1109/MS.1990.10021.

Miura K, Karp AH and Simon H (1993) 1992 Gordon Bell Prize Winners. *Computer* 26: 77–82. DOI:10.1109/MC.1993.10006.

Paul W (1986) Supercomputer., *October3*, 1986. Page 23.

Shaw DE (2014) Anton 2: Raising the Bar for Performance and Programmability in a Special-Purpose Molecular Dynamics, Supercomputer. In: *Proceeding SC ‘14 Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, New Orleans, LA, 16–21 November 2014, pp. 41–53. IEEE Press Piscataway, NJ, USA.

TOP 500 (n.d.) Available at: www.top500.org, June 2017.

Author biographies

Gordon Bell is a Microsoft researcher emeritus, and former Digital Vice President of R&D, where he led the development of the first mini- and time-sharing computers. As NSF’s founding Director for Computing (CISE), he led then plan for NREN (Internet). Professional interests: computer architecture, high-tech startup companies, and lifelogging. He is a member of the American Academy of Arts and Sciences, the National Academy of Engineering, National Academy and Science, and received The 1991 National Medal of Technology. He is a founding trustee of the Computer History Museum, Mountain View, CA.

David H Bailey recently retired from the Lawrence Berkeley National Laboratory, and a Research Associate with the University of California, Davis, has published over 200 papers in high-performance computing, numerical algorithms, and computational mathematics. He has received the Sidney Fernbach Award from the IEEE Computer Society, the Gordon Bell Prize from the Association for Computing Machinery, the Chauvenet Prize and Merten Hesse Prizes from the Mathematical Association of America, and the Levi L. Conant Prize from the American Mathematical Society.

Jack Dongarra holds an appointment at the University of Tennessee, Oak Ridge National Laboratory, and the University of Manchester. He specializes in numerical algorithms in linear algebra, parallel computing, use of advanced-computer architectures, programming methodology, and tools for parallel computers. He was awarded the IEEE Sid Fernbach Award in 2004; in 2008 he was the recipient of the first IEEE Medal of Excellence in Scalable Computing; in 2010 he was the first recipient of the SIAM Special Interest Group on Supercomputing’s award for Career Achievement; in 2011 he was the recipient of the IEEE Charles Babbage Award; and in 2013 he received the ACM/IEEE Ken Kennedy Award. He is a Fellow of the AAAS, ACM, IEEE, and SIAM and a foreign member of the Russian Academy of Science and a member of the US National Academy of Engineering.

Alan H Karp is a principal architect at Earth Computing, where he is designing a novel network fabric for the datacenter. He was a principal scientist in the Office of the CTO of HP’s Enterprise Services and the technical architect of its service delivery platform. Before that, he worked in HP Labs on a variety of subjects, including usable security, automated negotiation, and processor architecture. Dr. Karp was the chief scientist at Hewlett–Packard’s E-speak Operation and one of the architects of the chips in Intel’s Itanium processor. Dr. Karp received his PhD in Astronomy from the University of Maryland, spent 2 years at IBM Research, and 1 year as an assistant professor of physics at Dartmouth College before joining IBM’s Palo Alto Scientific Center where he was awarded two outstanding innovation awards. He has published over 100 papers and conference proceedings and holds more than 70 patents.

Kevin Walsh is a student of the history of high-performance computing. He is the supercomputing history project lead for the 30th Anniversary of the IEEE/ACM Supercomputing Conference in 2018. Previously a systems engineer at the San Diego Supercomputer Center, he is currently at the Institute of Geophysics and Planetary Physics at the Scripps Institution of Oceanography. He is a member of the ACM, the IEEE Computer Society, and the Society of History of Technology.