Contact: Jim Ormond
                212-626-0505
                ormond@hq.acm.org

ACM TURING AWARD HONORS JACK J. DONGARRA FOR PIONEERING CONCEPTS AND METHODS WHICH HAVE RESULTED IN WORLD-CHANGING COMPUTATIONS

Dongarra's Algorithms and Software Fueled the Growth of High-Performance Computing and Had Significant Impacts in Many Areas of Computational Science from AI to Computer Graphics

New York, NY, March 30, 2022 – ACM, the Association for Computing Machinery, today named Jack J. Dongarra recipient of the 2021 ACM A.M. Turing Award for pioneering contributions to numerical algorithms and libraries that enabled high performance computational software to keep pace with exponential hardware improvements for over four decades. Dongarra is a University Distinguished Professor of Computer Science in the Electrical Engineering and Computer Science Department at the University of Tennessee. He also holds appointments with Oak Ridge National Laboratory and the University of Manchester.

The ACM A.M. Turing Award, often referred to as the "Nobel Prize of Computing," carries a $1 million prize, with financial support provided by Google, Inc. It is named for Alan M. Turing, the British mathematician who articulated the mathematical foundation and limits of computing.

Dongarra has led the world of high-performance computing through his contributions to efficient numerical algorithms for linear algebra operations, parallel computing programming mechanisms, and performance evaluation tools. For nearly forty years, Moore's Law produced exponential growth in hardware performance. During that same time, while most software failed to keep pace with these hardware advances, high performance numerical software did—in large part due to Dongarra's algorithms, optimization techniques, and production-quality software implementations.

These contributions laid a framework from which scientists and engineers made important discoveries and game-changing innovations in areas including big data analytics, healthcare, renewable energy, weather prediction, genomics, and economics, to name a few. Dongarra's work also helped facilitate leapfrog advances in computer architecture and supported revolutions in computer graphics and deep learning.

Dongarra's major contribution was in creating open-source software libraries and standards which employ *linear algebra* as an intermediate language that can be used by a wide variety of applications. These libraries have been written for single processors, parallel computers, multicore nodes, and

multiple GPUs per node. Dongarra's libraries also introduced many important innovations including autotuning, mixed precision arithmetic, and batch computations.

As a leading ambassador of high-performance computing, Dongarra led the field in persuading hardware vendors to optimize these methods, and software developers to target his open-source libraries in their work. Ultimately, these efforts resulted in linear algebra-based software libraries achieving nearly universal adoption for high performance scientific and engineering computation on machines ranging from laptops to the world's fastest supercomputers. These libraries were essential in the growth of the field—allowing progressively more powerful computers to solve computationally challenging problems.

"Today's fastest supercomputers draw headlines in the media and excite public interest by performing mind-boggling feats of a quadrillion calculations in a second," explains ACM President Gabriele Kotsis. "But beyond the understandable interest in new records being broken, high performance computing has been a major instrument of scientific discovery. HPC innovations have also spilled over into many different areas of computing and moved our entire field forward. Jack Dongarra played a central part in directing the successful trajectory of this field. His trailblazing work stretches back to 1979, and he remains one of the foremost and actively engaged leaders in the HPC community. His career certainly exemplifies the Turing Award's recognition of 'major contributions of lasting importance.'"

"Jack Dongarra's work has fundamentally changed and advanced scientific computing," said Jeff Dean, Google Senior Fellow and SVP of Google Research and Google Health. "His deep and important work at the core of the world's most heavily used numerical libraries underlie every area of scientific computing, helping advance everything from drug discovery to weather forecasting, aerospace engineering and dozens more fields, and his deep focus on characterizing the performance of a wide range of computers has led to major advances in computer architectures that are well suited for numeric computations."

Dongarra will be formally presented with the ACM A.M. Turing Award at the annual ACM Awards Banquet, which will be held this year on Saturday, June 11 at the Palace Hotel in San Francisco.

SELECT TECHNICAL CONTRIBUTIONS
For over four decades, Dongarra has been the primary implementor or principal investigator for many libraries such as LINPACK, BLAS, LAPACK, ScaLAPACK, PLASMA, MAGMA, and SLATE. These libraries have been written for single processors, parallel computers, multicore nodes, and multiple GPUs per node. His software libraries are used, practically universally, for high performance scientific and engineering computation on machines ranging from laptops to the world's fastest supercomputers.

These libraries embody many deep technical innovations such as:

Autotuning: through his 2016 Supercomputing Conference Test of Time award-winning ATLAS project, Dongarra pioneered methods for automatically finding algorithmic parameters that produce linear algebra kernels of near-optimal efficiency, often outperforming vendor-supplied codes.

Mixed precision arithmetic: In his 2006 Supercomputing Conference paper, "Exploiting the Performance of 32 bit Floating Point Arithmetic in Obtaining 64 bit Accuracy," Dongarra pioneered harnessing multiple precisions of floating-point arithmetic to deliver accurate solutions more quickly. This work has become instrumental in machine learning applications, as showcased recently in the [HPL-AI benchmark](), which achieved unprecedented levels of performance on the world's top supercomputers.

Batch computations: Dongarra pioneered the paradigm of dividing computations of large dense matrices, which are commonly used in simulations, modeling, and data analysis, into many computations of smaller tasks over blocks that can be calculated independently and concurrently. Based on his 2016 paper, "Performance, design, and autotuning of batched GEMM for GPUs," Dongarra led the development of the [Batched BLAS Standard]() for such computations, and they also appear in the software libraries MAGMA and SLATE.

Dongarra has collaborated internationally with many people on the efforts above, always in the role of the driving force for innovation by continually developing new techniques to maximize performance and portability while maintaining numerically reliable results using state of the art techniques. Other examples of his leadership include the Message Passing Interface (MPI) the de-facto standard for portable message-passing on parallel computing architectures, and the Performance API (PAPI), which provides an interface that allows collection and synthesis of performance from components of a heterogeneous system. The standards he helped create, such as MPI, the LINPACK Benchmark, and the Top500 list of supercomputers, underpin computational tasks ranging from weather prediction to climate change to analyzing data from large scale physics experiments.

Biographical Background

Jack J. Dongarra has been a University Distinguished Professor at the University of Tennessee and a Distinguished Research Staff Member at the Oak Ridge National Laboratory since 1989. He has also served as a Turing Fellow at the University of Manchester (UK) since 2007. Dongarra earned a B.S. in Mathematics from Chicago State University, an M.S. in Computer Science from the Illinois Institute of Technology, and a Ph.D. in Applied Mathematics from the University of New Mexico.

Dongarra's honors include the IEEE Computer Pioneer Award, the SIAM/ACM Prize in Computational Science and Engineering, and the ACM/IEEE Ken Kennedy Award. He is a Fellow of ACM, the Institute of Electrical and Electronics Engineers (IEEE), the Society of Industrial and Applied Mathematics (SIAM), the American Association for the Advancement of Science (AAAS), the International Supercomputing Conference (ISC), and the International Engineering and Technology Institute (IETI). He is a member of the National Academy of Engineering and a foreign member of the British Royal Society.

About the ACM A.M. Turing Award

[The A.M. Turing Award]() is named for Alan M. Turing, the British mathematician who articulated the mathematical foundation and limits of computing, and who was a key contributor to the Allied cryptanalysis of the Enigma cipher during World War II. Since its inception in 1966, the Turing Award has honored the computer scientists and engineers who created the systems and underlying theoretical foundations that have propelled the information technology industry.

About ACM

ACM, the Association for Computing Machinery, is the world's largest educational and scientific computing society, uniting computing educators, researchers, and professionals to inspire dialogue, share resources, and address the field's challenges. ACM strengthens the computing profession's collective voice through strong leadership, promotion of the highest standards, and recognition of technical excellence. ACM supports the professional growth of its members by providing opportunities for life-long learning, career development, and professional networking.

###