![Association for Computing Machinery — Advancing Computing as a Science & Profession]

NEWS RELEASE

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

KAZEM CHESHMI, MADHURIMA VARDHAN AND KEREN ZHOU NAMED
RECIPIENTS OF 2020 ACM-IEEE CS GEORGE MICHAEL MEMORIAL HPC FELLOWSHIPS

New York, NY, **October 8, 2020** – ACM, the Association for Computing Machinery, announced today that Kazem Cheshmi of the University of Toronto, Madhurima Vardhan of Duke University, and Keren Zhou of Rice University are the recipients of the 2020 ACM-IEEE CS George Michael Memorial HPC Fellowships. Cheshmi is recognized for his work building a Sympiler that automatically generates efficient parallel code for sparse scientific applications on supercomputers. Vardhan is recognized for her work developing a memory-light massively parallel computational fluid dynamic algorithm using routine clinical data to enable high-fidelity simulations at ultrahigh resolutions. Zhou is recognized for his work developing performance tools for GPU-accelerated applications. The Fellowships are jointly presented by ACM and the IEEE Computer Society.

Kazem Cheshmi
In mathematics, a matrix is a grid (represented in a table of rows and columns) that is used to store, track and manipulate various kinds of data. In computer science, matrices are especially used in graphics, where an image is represented as a matrix in which each datapoint on the matrix table would directly correspond to the color and/or intensity of a given pixel. Matrix computations have a wide range of practical uses. For example, a 3D graphics programmer would hold all the datapoints related to an image as elements of the matrix and might make matrix computations to cause the image to rotate or scale. Matrix computations also play an essential role in computer vision, a branch of AI in which a computer learns to identify an image.

Historically, mathematicians would develop algorithms for matrix computations, and software engineers would write programs to make the algorithms run on powerful parallel computers. However, the emergence of massive datasets has meant that traditional approaches to matrix computation are often inadequate for the enormous matrices, requiring complex algorithms that are increasingly used today in areas such as data analytics, machine learning, and high performance computing.

To address this problem, Cheshmi has developed Sympiler, a domain-specific compiler (a program that translates the source code from a programming language to a code the computer can understand). Cheshmi's Sympiler generates high performance codes for sparse numerical methods and can process complex matrix computations derived from massive datasets. Sympiler is extended to nonlinear

optimization algorithms and performs faster than existing nonlinear optimization tools and is scalable to some of the most powerful high performance computers. Cheshmi's work was also accepted to SIGGRAPH 2020, where he demonstrated how he is using Sympiler in robotics and graphics applications.

Madhurima Vardhan
Despite recent advances, cardiovascular disease (CVD) remains the leading cause of deaths worldwide. In the field of high performance computing, some researchers develop algorithms that are processed on powerful supercomputers to create visual simulations of complex biological processes. These simulations can be useful tools to help researchers better understand how to treat disease. Currently, a form of simulation called a computational fluid dynamic (CFD) simulation is used in health clinics to provide noninvasive diagnosis of CVDs.

However, existing state-of-the-art CFD simulations do not provide high-fidelity real-time diagnosis of CVDs. These limitations stem from a variety of factors, including problems with model accuracy based on the patient images that comprise the datasets; the extensive memory requirements of these kinds of simulations; the long runtimes on high performance computers that are required for these kinds of simulations; and teaching physicians how to effectively use these simulations.

To address these problems, Vardhan is developing a new kind CFD algorithm using routine patient image datasets that can develop high-fidelity simulations at ultra-high resolutions. Her algorithm is memory-light (that is, using less memory than existing algorithms), and massively parallel (proven to scale on supercomputers). As part of her PhD work, she also completed a study to determine how physicians interact with simulation data, and how physician behavior might be modified in treatment planning.

Keren Zhou
In the last 10 years, graphics processing units (GPUs) have become a critical component in high performance computing systems. For example, five of the top 10 supercomputers in the world today use GPUs to accelerate the performance of applications in various domains. These systems must be designed to avoid common GPU performance problems, and identifying specific performance problems can be challenging.

Working with his advisor John Mellor-Crummey and others, Zhou has taken the lead in developing performance tools for GPU-accelerated supercomputing to help programmers detect program inefficiencies and provide optimization advice.

Their work has already been well received in academia and industry. Zhou and his colleagues have published three papers in top-tier conference proceedings. They are also collaborating with GPU vendors, including AMD, Intel, and NVIDIA; they have submitted a collection of bug reports and offered advice about how to improve their GPU hardware and software measurement interfaces.

About the ACM-IEEE CS George Michael Memorial HPC Fellowship
The ACM-IEEE CS George Michael Memorial HPC Fellowship is endowed in memory of George Michael,

one of the founders of the SC Conference series. The fellowship honors exceptional PhD students throughout the world whose research focus is on high performance computing applications, networking, storage, or large-scale data analytics using the most powerful computers that are currently available. The Fellowship includes a $5,000 honorarium and travel expenses to attend the SC conference, where the Fellowships are formally presented.

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.

About SC

SC, the International Conference for High Performance Computing, sponsored by ACM and IEEE-CS offers a complete technical education program and exhibition to showcase the many ways high performance computing, networking, storage and analysis lead to advances in scientific discovery, research, education and commerce. This premier international conference includes a globally attended technical program, workshops, tutorials, a world class exhibit area, demonstrations, and opportunities for hands-on learning.

###