

Artificial Intelligence (AI)

FPGA community forums and blogs have moved to the Altera Community. Existing Intel Community members can sign in with their current credentials.


Intel Community / Blogs / Tech Innovation / Artificial Intelligence (AI)

804 Discussions

Intel Senior Fellow Receives ACM Fellowship for Parallel Processing of Data-Intensive Applications

Subscribe

Article Options



Scott Bair 02-05-2024 Employee

6,096

Scott Bair is a key voice at Intel Labs, sharing insights into innovative research for inventing tomorrow's technology.

Highlights

- Intel Senior Fellow Pradeep K. Dubey was named a 2023 ACM Fellow for his lifelong technical contributions to emerging compute- and data-intensive applications and parallel processing computer architectures.
- At Intel Labs, Dubey has led a team of top researchers focused on state-of-the-art research to improve the performance and scalability of applications that require high compute power and handle large amounts of data.
- With this designation, Dubey joins two other Intel Senior Fellows who have been recognized as ACM Fellows.

Intel Senior Fellow Pradeep K. Dubey was named a 2023 Association for Computing Machinery (ACM) Fellow for his lifelong technical contributions to emerging compute- and data-intensive applications and parallel processing computer architectures. As director of the Parallel Computing Lab, a part of the Intel Labs organization at Intel Corporation, Dubey has led a team of top researchers focused on state-of-the-art research to improve the performance and scalability of applications that require high compute power and handle large amounts of data. The ACM Fellowship recognizes the top 1% of association members for their outstanding accomplishments in computing and information technology and/or exceptional service to ACM and the larger computing community. These contributions have advanced technologies used every day to improve the lives of people around the world.

At Intel Labs, Dubey and his team are responsible for defining computer architectures that can efficiently handle compute-intensive applications in emerging machine learning/artificial intelligence, and traditional high performance computing (HPC) applications for data-centric computing environments.

He holds more than 30 patents and has published more than 100 peer-reviewed technical papers. In 2014, Dubey was honored with an Outstanding Electrical and Computer Engineer Award from Purdue University, in addition to receiving an Intel Achievement Award for breakthroughs in parallel computing research in 2012. He was named a Fellow of the Institute of Electrical and Electronics Engineers (IEEE) in 2001 for his contributions to computer architecture supporting multimedia processing.

With this designation, Dubey joins two other Intel Senior Fellows who have been recognized as ACM Fellows. Geoff Lowney was honored in 2009 for his contributions to compiler technology and performance enhancement tools while Nick McKeown was acknowledged for his contributions to network switching and queueing in 2006.

Dubey shared with Intel Labs the innovative solutions he has contributed to the parallel processing of compute- and data-intensive applications.

You were recognized for your contributions to emerging compute- and data-intensive applications and parallel processing computer architectures. What are your major accomplishments in these areas?

I have had a singular passion the last two to three decades of my career: How do we transform the capabilities of high-volume, commodity compute infrastructure like PCs or the cloud such that we can deliver to the hardware-software system needs of compute- and data-hungry, real-world applications of mass relevance. In other words, how can we accelerate the benefits of Moore's Law to the masses? In the present day context, artificial intelligence (AI) offers a great means to this end, as it teaches us a feasible algorithmic method to leverage computing for learning globally to best serve locally.

My specific accomplishments are:

- Significant contributions to the design, architecture, and application performance of various microprocessors, including IBM® PowerPC®, Intel® i386TM, i486TM, Pentium®, and Xeon® line of processors.
- Significant contributions to novel algorithms for cyclic redundancy check (CRC) generation, Advanced Encryption Standard (AES) encryption, and double-blind internet transaction (while at IBM).
- Single instruction, multiple data (SIMD) instruction set extension that formed the basis of high-volume PowerPC-based gaming platforms in 1990s (while at IBM).
- Recognition-mining-synthesis (RMS) or tera-era applications enabled by manycore architecture research, leading to Xeon HPC roadmap in early 2010s.
- Big data-led rebirth of machine learning/artificial intelligence through novel support for irregular, low-precision, and matrix processing capabilities to x86 software-hardware platform for Intel's current and upcoming compute platforms.
- Application algorithm architecture co-design innovations delivering 10+ industry-wide performance records.
- Software innovations lowering the ninja barrier of high-performance programming.

What challenges did you and your team have to overcome?

Intel has been a bottoms-up company, driven by Moore's Law innovations. Whereas my team has been singularly focused on complementing this focus with top-down applications-driven algorithm architecture co-design innovations. This has been a daunting task, especially in our early days, to get our architects and design teams to realize that they were designing processors for three to four years in the future, focusing on benchmarks that were five to 10 years lagging in terms of the applications they represented. In other words, there was a 10+ year gap in what machines were designed for versus how they were actually used. The second challenge we face is an ongoing one: Bringing technical transparency and fairness in claimed speedups and performance comparisons across two platforms.

What innovative solutions did your team introduce to solve these problems?

We led the industry by creating a more forward-looking set of benchmarks (RMS Bench, which later became PARSEC, managed by Princeton). We introduced novel features in our x86 architecture and middleware such as the Intel® Math Kernel Library (MKL) and deep neural networks (DNN) libraries. We also worked with academia and the ecosystem to mainstream parallel and low-precision computing. We addressed the challenges posed by traditional HPC computing, which was pulling in the opposite direction toward higher precision versus lower precision computing.

What are the applications of these solutions and how do they affect our daily lives?

The applications are primarily aimed at making people more productive in their daily lives by delivering machines that don't just do fast and accurate computations, but also can make better and faster decisions. Transforming real-time decision making of complex problems will have very broad impact on what computing can do for us, from discovering novel materials, to safer driving, to personalized healthcare.

What limitations or opportunities are you and your team tackling today that will affect the future of this field of work?

Our research focus looking ahead is growing the energy efficiency of AI computations by more than one order of magnitude. Most of the energy spent today is not on the computing itself, but rather on bringing the data to the compute unit and communicating the result to the next unit. Hence, our specific focus has shifted to gaining energy efficiency for the infrastructure devoted to moving and communicating the data through memory and network.

The vast majority of programmers today are not engineers or computer scientists, rather they are data scientists or other non-computing domain experts. This further affects how machines will get programmed tomorrow at an abstraction level much higher than today. At the same time, these new programmers are even more demanding in terms of their compute performance needs than a niche HPC programmer. Bridging this performance-productivity gap is our second focus, leveraging increased use of machine learning for automated code generation, verification, and debug.

How did key mentors, colleagues, and internal support systems at Intel assist you in your technical journey?

My first and last academic mentor was Prof. Mike Flynn at Stanford University. I was already a productive member of an x86 architecture team working happily at Intel. I left my job to pursue a doctorate degree and then a research career at the IBM Thomas J. Watson Research Center, followed by Intel Labs, primarily guided by Prof. Flynn in every step of my decision. I wish all aspiring researchers had a mentor like him. Besides my schoolteachers, my Intel seniors, managers, peers, juniors, plus my friends and family members have had tremendous influence on my career and achievements as well.

What key takeaways or advice would you give to the next generation in this field of work?

We are living through exciting times as technologists. Machine learning and big data are changing not just what computing can do for us, but how computing is done. AI in essence is all about the art and science of learning globally to best serve locally.

My advice for the next generation: Find your passion first! Then do some homework to find a real-world, high-impact problem aligned with your passion. The rest will just follow. You will devote all your time and energy to gain the expertise needed such that you cannot be replaced by any human or machine!

Translate

Tags: AI Intel_Labs

0 Kudos

About the Author



Scott Bair is a Senior Technical Creative Director for Intel Labs, chartered with growing awareness for Intel's leading-edge research activities, like AI, Neuromorphic Computing and Quantum Computing. Scott is responsible for driving marketing strategy, messaging, and asset creation for Intel Labs and its joint-research activities. In addition to his work at Intel, he has a passion for audio technology and is an active father of 5 children. Scott has over 23 years of experience in the computing industry bringing new products and technology to market. During his 15 years at Intel, he has worked in a variety of roles from R&D, architecture, strategic planning, product marketing, and technology evangelism. Scott has an undergraduate degree in Electrical and Computer Engineering and a Masters of Business Administration from Brigham Young University.

You must be a registered user to add a comment. If you've already registered, sign in. Otherwise, register and sign in.

Comment

Community support is provided Monday to Friday. Other contact methods are available here.

Intel does not verify all solutions, including but not limited to any file transfers that may appear in this community. Accordingly, Intel disclaims all express and implied warranties, including without limitation, the implied warranties of merchantability, fitness for a particular purpose, and non-infringement, as well as any warranty arising from course of performance, course of dealing, or usage in trade.

For more complete information about compiler optimizations, see our Optimization Notice.