Disrupting Algorithms and Software: Dispatch from IPDPS18

*In May 2018, another ACM SIGHPC in- cooperation meeting, IEEE IPDPS, took place in Vancouver, Canada. We caught up with **Bora Uçar** (CNRS and ENS Lyon), who served as this year's general chair, for highlights from the meeting.*

Though we didn't plan for it per se, at this year's IPDPS there emerged a general theme on the potentially disruptive effects of emerging and future hardware.

It began with three invited talks that aligned in perfect harmony. **Michael Bender** (Stony Brook University) discussed how to optimize the write operations in algorithms and software, to save both time and energy. Bender's talks are always entertaining and this one was no exception according to social <https://twitter.com/hpcgarage/status/998965513953755136> **media** <https://twitter.com/search?q=%23ipdps%20bender&src=typd>. **Keren Bergman** (Columbia University) explained how energy is spent in current architectures and the role photonic interconnects may play in rebalancing architectures. She concluded that these changes will affect how we program. That proved to be a prescient segue to the next invited talk: **Bruce Hendrickson** (Lawrence Livermore National Laboratory) gave his views on the challenges to increasing performance and speculated on the viable paths forward. He credited five orders of magnitude in improvements to Moore and Dennard scaling, compared to just one order by building bigger machines and another one by architectural improvements (see this slide <http://bit.ly/BAH-DennardScaling>). He argued that while we may achieve exascale performance in the near future with familiar technologies and programming models, to go beyond that, we will need radically new computer architectures that will disrupt to our approach to algorithms and software.

Consistent with this theme, **Janice McMahon** (Emu Technology) gave a tutorial (slides are here <http://bit.ly/JMM-EmuTutorial>) on her company's new architecture, which promises an order of magnitude improvement for data intensive problems, such as computations involving graphs, sparse matrices, and sparse tensors. This system tries to "move code to data" rather than the other way around.

Many of the 28 technical paper sessions, consisting of 113 papers, reinforced the theme. These papers underwent a two-round evaluation and revision, resulting in a high-quality program. There were the usual workshops on Monday and Friday,

which covered a wide spectrum of topics. Refer to the **online program** <http://bit.ly/IPDPS2018Program> for details.

Whether you attended or not, we hope you'll give us your feedback. Use <http://bit.ly/IPDPS18> to do so online. On behalf of the IPDPS steering committee, I look forward to hearing from you!