Message from the Workshop Chair

Welcome to the 10th Workshop on Fault-Tolerance for HPC at eXtreme Scale (FTXS 2020), organized in conjunction with SC20. Addressing failures in extreme-scale systems remains a significant challenge to reaching exascale. Current projections suggest that at the scale necessary to sustain exaflops of computation, systems could experience failures as frequently as once per hour. As a result, robust and efficient fault tolerance techniques are critical to obtaining acceptable application performance. Additionally, it is imperative that we develop an understanding of trends in hardware devices that may affect the reliability of future systems. The emergence of high-bandwidth memory devices, the continued deployment of burst buffers, and the development of near-threshold devices to address power concerns will all impact fault tolerance on new systems. These design trends coupled with increases in the number, variety, and complexity of components required to compose next-generation extreme-scale systems mean that these systems will experience significant increases in aggregate fault rate, fault diversity, and the complexity of isolating root causes.

Further complications arise from the fact that many of these hardware trends also mean that the likelihood of undetected errors, e.g., silent data corruption, is growing. Power limitations mean that future systems may not have room in their power budgets to deploy powerful hardware correction mechanisms (e.g., chipkill). Less protective (and more power-efficient) mechanisms have been shown to be more susceptible to undetected errors. As a result, application developers are increasingly less confident that they can rely on hardware devices to produce "correct" results.

Based on these trends in extreme-scale systems, FTXS attracts work from scientists and engineers studying fault tolerance around the world. We would like to thank all the authors who submitted their work to FTXS 2020. The quality of the submissions that we received is a testament to the excellent work that is currently being done on fault tolerance for extreme-scale systems. We would also like to acknowledge the hard work of the members of our Program Committee in reviewing these submissions.

Our program consists of six papers that examine several important aspects of fault tolerance. The topics addressed by these papers include: resilience for asynchronous programming models, understanding design patterns for resilience, leveraging compiler analysis for collecting and writing checkpoint data, and developing algorithms that can withstand faults.

Although FTXS 2020 will be held virtually this year due to the pandemic, we hope that it will still provide an opportunity for experts and novices alike to engage in thought-provoking discussions of new and innovative ideas on fault tolerance for next-generation extreme-scale systems.

Scott Levy
Sandia National Laboratories
FTXS 2020

978-1-6654-2289-5/20/$31.00 ©2020 IEEE