



Contribution ID: 128

Type: **Keynote Talk**

The Third Pillar of HPC

Wednesday, 4 December 2024 19:00 (45 minutes)

The advent of Exascale computing in 2022 marks a major milestone in HPC but also demonstrates its limitations for future progress. Historically, conventional MPP (and large commodity clusters) have achieved enhanced performance by a factor of 2 every two years. In addition to CPUs, GPUs have extended this through streaming SIMD computations for certain classes of application algorithms. But the end of Moore's Law as well as Dennard Scaling is severely constraining future progress, especially with respect to cost as Frontier approaches 8,000 square feet and only one other, Aurora, has been announced since then. The major class of supercomputer computation not adequately addressed is that of dynamic adaptive graph processing required for advanced forms of machine intelligence; hence, the third pillar of computation. Graphs exhibit neither much spatial locality nor temporal locality but suggest what may be called "logical locality" as the data structures explicitly define their own topologies. However, a new approach to computer architecture, a non-von Neumann family, that is both dynamic and adaptive can easily provide an order of magnitude performance to cost advantage over current methods. While this form of improvement is particularly advantageous for dynamic graph processing, it also can enhance more typical matrix processing. This closing Keynote address will introduce the foundational concepts of the Active Memory Architecture which is being pursued by the Texas Advanced Computing Center. Questions will be addressed from the participants throughout the presentation.

Student or Postdoc?

Email address

Co-Authors

CHPC User

CHPC Research Programme

Workshop Duration

Primary author: Prof. STERLING*, Thomas (TACC)

Presenter: Prof. STERLING*, Thomas (TACC)

Session Classification: Keynote

Track Classification: HPC Technology