Precision *vs.* Flop/s Performance

Tale of three Gordon Bell prize finalists in 2025: Effects of floating-point (FP) and brain-float (BF) precisions on Flop/s (floating-point operations per second) performance on the Aurora supercomputer

- "Advancing quantum many-body GW calculations on exascale supercomputing platforms," B. Zhang (USC) *et al.*

  0.71 Exaflop/s on 9,600 nodes (FP64)

- "Multiscale light-matter dynamics in quantum materials: from electrons to topological superlattices," T. Razakh (USC) *et al.*

  1.87 Exaflop/s on 10,000 nodes (FP64/FP32/BF16)

- **"**AERIS: Argonne earth systems model for reliable and skillful predictions," V. Hatanpaa (Argonne) *et al.*

  10.2 Exaflop/s on 10,080 nodes (FP32/BF16)

Hybrid-precision accumulation/ matrix-multiplication

$$\overbrace{\mathbf{M} \mathrel{+}= \sum_i}^{\text{FP32}} \overbrace{\mathbf{L}_i \times \mathbf{R}_i}^{\text{BF16}}$$

BF16 provides a wider dynamic range with lower precision