Pingpong MPI Benchmark - SEM vs. "grid"

The Supercomputing Engine for Mathematica (SEM), like modern supercomputers, supports an "all-to-all" communications network between *Mathematica* kernels. Practitioners in high-performance computing know the communications network ultimately limits the addressable problem size. SEM provides supercomputing infrastructure for Wolfram's *Mathematica*.

We tested the performance of SEM versus gridMathematica using the pingpong and swap communications benchmarks on a cluster with 4, 8, and 16 processors. While SEM's performance remained steady for increased cluster size, gridMathematica's performance decreased due to its "master" kernel becoming a bottleneck for all communications.

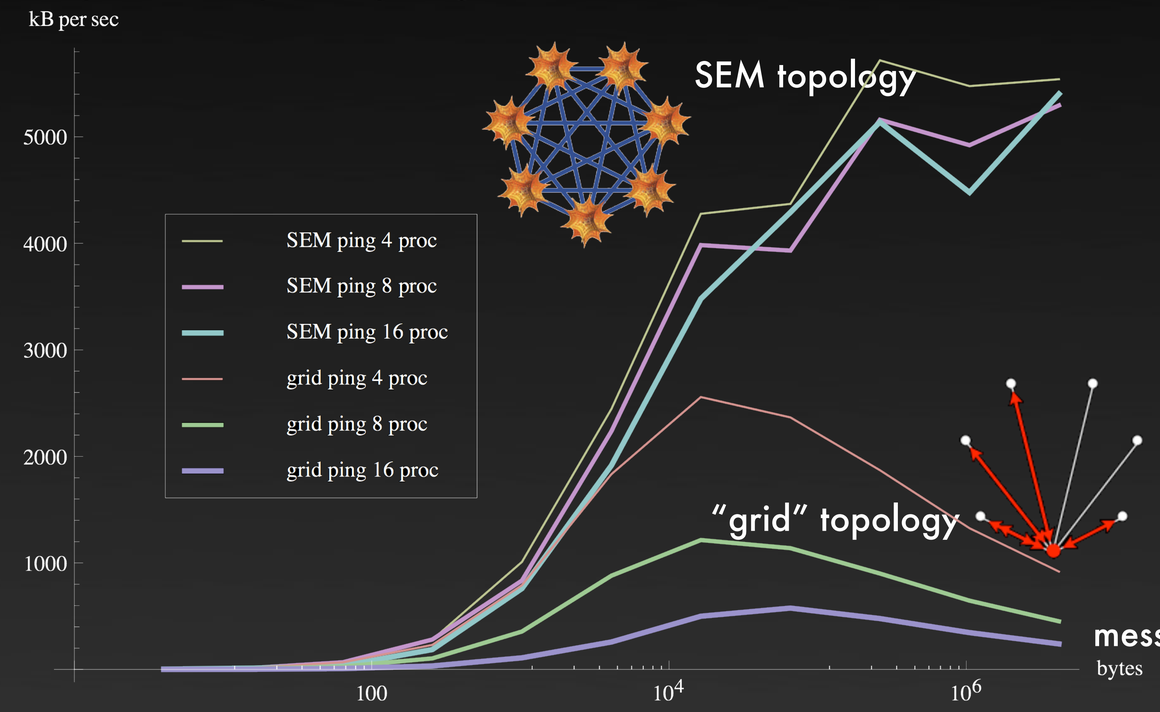

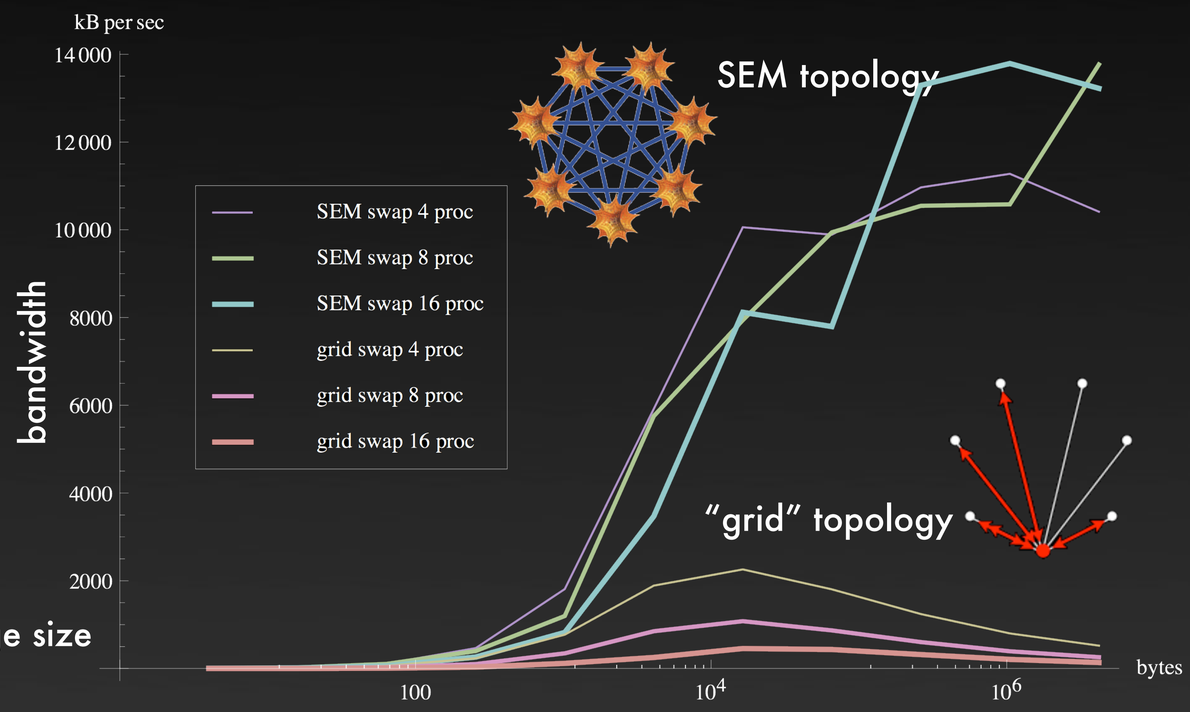

A high-performance computing benchmark called "pingpong", developed by Viktor K. Decyk of the UCLA Plasma Physics Group, stresses the network of a parallel computer in two stages after dividing the processing elements into the even and odd processors. The pingpong stage has each even one send a message to an odd one (0 to 1, 2 to 3, ...) and back again. The message is varied in size from a few bytes to many megabytes while each operation is timed and averaged. The second stage performs a "swap", meaning each pair of processes sends and receives messages from each other simultaneously, if possible. Again the message size is varied and the results are timed and averaged.

Patent-Pending

http://daugerresearch.com/pooch/mathematica/