Qeios

Peer Review

Review of: "Evaluation of Bfloat16, Posit, and Takum Arithmetics in Sparse Linear Solvers"

Aditya Kashi[1]

1. National Center for Computational Sciences, Oak Ridge National Laboratory, Oak Ridge, United States

The article compares fixed-precision and tapered precision formats in various (mixed-precision) solver methods for sparse linear systems from the Suitesparse collection. It does a nice job of systematically developing the test setup and making largely fair comparisons. The conclusion that takums typically give better numerical behaviour than posits is a nice contribution, in addition to the more predictable fact that tapered precision formats do better than fixed-precision formats.

The biggest weakness of the work is that there is no comment on runtime performance or energy efficiency. While I understand that hardware for tapered-precision formats is currently unavailable, some discussion of this essential aspect (efforts such as CalligoTech could have been mentioned) would have been appreciated. I would even hope that some performance comparison is possible for memory bandwidth-limited workloads such as these. On a different track, is it possible to systematically predict the runtime performance on some realizable architecture?

Section II.C. The sparse linear solver does not use pivoting, but it is claimed to be “maintaining parity” with UMFPACK. While I understand that the permutations and scaling from UMFPACK are used, LU decomposition without pivoting is not a representative solver for direct solvers.

Section II.D. Similar issue as above. The SPQR solver algorithm is not used, but rather something that's claimed to “effectively emulate the behaviour” of SPQR, without justification or explanation.

Related to the above two points, I would ideally like to see UMFPACK and SPQR converted to a templated or parameterized (on the scalar format) code so that the actual same algorithms can be tested, though I understand this takes more work. As of now, it's hard to judge whether this is a fair comparison.

doi.org/10.32388/J1042N

Minor:

In the results figures, I like the idea of infinity_omega and infinity_sigma as y-axis points.

Fig 3a: it would be good to zoom into what's happening in the first ~20% on the x-axis and maybe clip the y-axis for clarity.

Declarations

Potential competing interests: No potential competing interests to declare.