ABSTRACT – 2022 ENERGY CONSEQUENCES OF INFORMATION WORKSHOP

The Reversible Computing Scaling Path: Challenges and Opportunities

Thursday, November 18, 2021

Michael P. Frank
Center for Computing Research
Sandia National Laboratories
mpfrank@sandia.gov

Hannah Earley
Dept. of Appl. Math.
& Theoretical Physics
Cambridge University
wje25@cam.ac.uk

Karpur Shukla
Lab. for Emerging Technologies
Brown University
karpur_shukla@brown.edu

Abstract

We know, as a direct logical consequence of fundamental physical and information-theoretic considerations, that general digital computing technology can only potentially avoid having its energy efficiency (*i.e.*, computational performance per unit power consumption) be limited as a consequence of certain fundamental thermodynamic constraints, such as Landauer's principle, if that technology is rearchitected on the basis of the principles of *reversible computing*. Conventional *non*-reversible digital technology appears likely to approach its practical efficiency limits within the next decade or so, and several exemplar design methodologies for reversible digital hardware have been prototyped for both semiconducting and superconducting technology platforms. So, the time seems ripe to begin developing reversible computing technology in earnest.

It is important for this enterprise to realistically understand the broad, high-level picture of the long-term technology scaling path that reversible computing offers. Careful analysis shows that indeed, reversible technology can potentially offer a roadmap towards improving the energy efficiency (and cost efficiency!) of digital computing technology indefinitely, with no known fundamental limits, at least for machines below an astronomically large size. This remains true despite the existence of various overheads inherent to the reversible paradigm, including adiabatic trade-offs between energy dissipation and speed, as well as the algorithmic overheads that come into play when transforming (worst-case) non-reversible algorithms into equivalent reversible ones.

Despite the overall optimistic long-term picture that this scaling analysis paints, a number of concrete, and quite significant, related scientific and engineering challenges remain.

One fundamental scientific question asks whether we can derive technology-independent limits on dissipation that apply to *all possible* reversible digital machines, as a function of various relevant timescales of interest, directly from fundamental theory (such as, *e.g.*, the general formal apparatus of non-equilibrium quantum thermodynamics in open systems). A related question is whether exotic quantum-mechanical phenomena can be usefully harnessed in the design of new fundamental physical mechanisms of operation that could be utilized in novel device technologies that aim to perform as well as possible within the aforementioned technology-independent limits.

Meanwhile, on the engineering side, a key challenge is to describe a systematic methodology by which the *effective quality factor* of energy-recovering, ballistic driving mechanisms (*e.g.*

resonators or flywheels) for synchronous machines can be increased to indefinitely high levels. Although we don't yet know any fundamental limits on these quality factors, we also don't yet have a clear recipe for lifting them to unboundedly high levels. An alternative paradigm of *asynchronous* ballistic reversible computing remains to be fully explored, but fundamental considerations may ultimately turn out to prevent it from scaling as well as the synchronous approach.

Another, less uncertain, but still significant challenge is simply to do the significant amount of reworking of electronic design automation (EDA) tools and methodologies that will be required in order to facilitate the development of reversible machines of substantial complexity. Coupled to this is the significant effort that will need to be put into workforce development to train or retrain engineers to work within the new paradigm.

Artificial intelligence may very well end up playing a key role in facilitating design optimization, workforce education, and eventually even design discovery within this new field, thus helping to drive the reversible computing revolution forwards. In turn, reversible technology will feed back to help boost the efficiency of AI technology to previously unthinkable levels, in a virtuous cycle that stands to become one of the key engines that will drive the growth of the digital economy throughout the 21st century, and perhaps even far beyond.

In this talk, we survey and discuss the above issues, and conclude with remarks about the next steps that will be required in order for the reversible computing future to blossom.

Funding Statement

This research was funded in part by the Laboratory Directed Research and Development (LDRD) and Advanced Simulation and Computing (ASC) programs at Sandia National Laboratories, a multimission laboratory managed and operation by National Technology and Engineering Solutions of Sandia, LLC., a wholly owned subsidiary of Honeywell International, Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA-0003525. It was also supported in part by the U.S. Army Research Office (ARO) under cooperative agreement W911NF-14-2-0075 and BAA W911NF-19-S-0007, in part by the U.S. Air Force Office of Scientific Research (AFOSR) under grant number FA9550-19-1-0355, and in part by the U.K. Engineering and Physical Sciences Research Council (EPSRC) under project reference 1781682. This document describes objective technical results and analysis. Any subjective views or opinions that might be expressed in this document do not necessarily represent the views of the U.S. Department of Energy or the United States Government. Approved for public release, SAND2021-14732 A.