When does an unconventional substrate compute?

D. Horsman * Susan Stepney † Viv Kendon ‡

Many diverse substrates are proposed for unconventional computation, from relativistic and quantum systems, to chemical reactions and slime moulds. But what does it mean to say that such substrates are specifically computing, as opposed to merely performing the physical processes of that substrate?

Our theoretical framework [1] encompasses and distinguishes the acts of performing scientific experiments to test a theory, engineering physical systems in the presence of a theory, and performing computation with physical systems. In our framework we have a space of physical entities P, behaving according to (known or unknown) physical laws $\mathcal{H} : P \to P$, and a space of abstract mathematical or computational entities A, behaving as prescribed by the relevant (theory-dependent) mathematical or computational dynamics $\mathcal{C} : A \to A$. We have a (theory-dependent) modelling representation relation $M : P \to A$, that models physical entities as abstract mathematical objects. We can define a (theory-dependent and non-primitive) reverse instantiation relation $\widetilde{M} : A \to P$, that instantiates representational abstract mathematical objects as corresponding physical entities. The different possible relationships between P, A, M and $\widetilde{M}$ distinguish the processes of science, prediction, engineering, and computation, as described in [1].

We define *physical computing* to be the use of a (well-characterised and well-engineered) physical system's dynamics to *predict* the outcome of an abstract dynamics (desired computation) [1]. We represent our abstract problem in an abstract machine description $A = M(P)$. We encode this into P using $\widetilde{M}$, and the physical dynamics $\mathcal{H}(P)$ gives the final physical state P'. We decode this to $M(P')$, which, for a well-engineered physical system, is *sufficiently close* to $M'(P)$, the desired a result.

We can use our framework to analyse various physical devices to see if and how they are being used as computers. The use of a physical system as a computer is a use of *technology*: computers are highly engineered devices.

The physical theory C for classical digital computers is well developed. It is used to scale devices, and run programs that are new. In contrast, unconventional computing substrates generally have a theory that is much less well

*Department of Computer Science, University of Oxford, OX1 3QD, UK. Supported by the CHIST-ERA DIQIP project, and by the FQXi Large Grant "Time and the Structure of Quantum Theory"

†Department of Computer Science, University of York, YO10 5GH, UK. Partially supported by EU FP7 FET Coordination Activity TRUCE, project reference number 318235

‡School of Physics and Astronomy, University of Leeds, LS2 9JT, UK

developed. This covers both the scientific theory of the native substrate, and the theory of it in its engineered state, which can require significant modification. This leads to problems of confidence in scale and composition that can cast doubt on the use of a particular substrate as a computer.

A *phenomenological theory* is one derived from experimental observation, without much or any underlying explanatory model. A phenomenological theory can have predictive power (this input is taken to that output); however, without an underlying physical theory, it has to do a lot more work to convince that all relevant changes have been taken into account and that the computation can be relied on.

Well-characterised physical substrates tend to have deep theories. Any phenomenological theories involved in the engineering of these substrates is usually valid at the relevant production scales, too. However, less well-characterised chemical and biological substrates, and unconventional materials, tend to have phenomenological theories only, and valid only at smaller scales. This means it is difficult to take any results about computational ability or performance of these substrates at small laboratory sizes, and scale these up to large scale performance, with any confidence.

Some unconventional devices do have a physical model behind them, but this claimed physical model does not capture the actual physical behaviour, particularly at large scales. A theory that is 'sufficiently good' for its original purpose might nevertheless not be 'sufficiently good' for computational purposes, particularly at larger scales, or at the extremes necessary to gain the claimed computational properties. In this case we have a bug in the physical implementation, not because it has been incorrectly engineered according to theory, but because the underlying simplistic physical theory is wrong. Analysis of the operation of unconventional devices needs to be aware of the domain of applicability of the underlying theory, and when the operation of the device falls outside this domain.

Another reason that a theory that is 'sufficiently good' for its original purpose might nevertheless not be 'sufficiently good' for computational purposes is that computational use might require a dramatically better precision, particularly when multiple physical steps are composed, in order to minimise the propagation of errors.

Additionally, a highly engineered substrate may well include multiple kinds of components and connections between components, each with their underlying physical theories known to a greater or lesser extent. Composition of such theories introduces yet more problems, and can further reduce confidence in the resulting system. More research is needed in order to fully characterise the computation happening within and between multiple composed substrates.

References

[1] Horsman, D., Stepney, S., Wagner, R.C., Kendon, V.: When does a physical system compute? arXiv:1309.7979 [cs.ET] (2013)