Performance Trade-offs in Reversible Amorphous Computers

William Earley · w.earley@damtp.cam.ac.uk · Department of Applied Mathematics and Theoretical Physics · University of Cambridge

Speed Limits

An important metric of a computer's performance is its raw speed in operations per unit time. It has been shown [1] that the maximum rate of quantum state transitions per unit mass-energy is $\nu \leq 2c^2/h = 2.71 \times 10^{50}\,\mathrm{s}^{-1}\,\mathrm{kg}^{-1}$, thus placing an upper bound on the rate of computation.

This suggests that to maximise performance, one must build ever larger computers. We investigated the physical limits of computational speed for asymptotically large computers. A key consideration was the generation and removal of heat. In order to maintain system integrity, we constrain our computers to be isothermal.

Reversible computers outperform irreversible computers for a given size

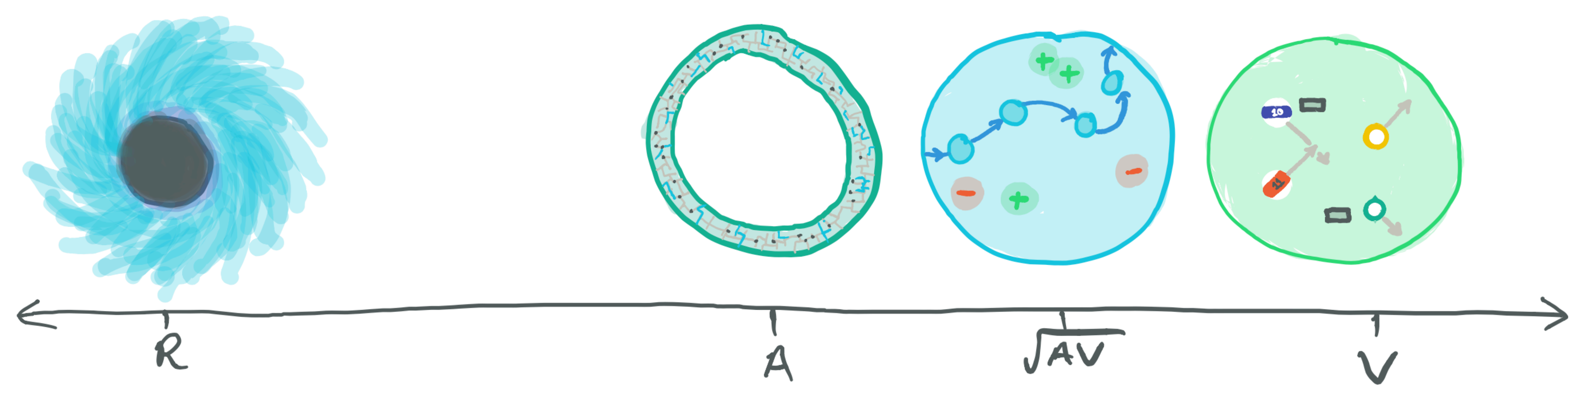

It can be shown that the rate of heat removal for a system scales with its convex-bounding surface area, A.

By consideration of Landauer's principle, we can show that an irreversible computer of volume V can only perform useful computation in a thin shell, restricting its total computational rate to $\nu \sim A$.

In contrast, a reversible computer can in principle generate no heat. In practice, however…

> *Even if classical balls could be shot with perfect accuracy into a perfect apparatus, fluctuating tidal forces from turbulence in the atmospheres of nearby stars would be enough to randomise their motion within a few hundred collisions. Needless to say, the trajectory would be spoiled much sooner if stronger nearby noise sources (e.g., thermal radiation and conduction) were not eliminated.*
>
> – Bennett, 1982 [2]

We proceed to surmise that reversible computers must operate near thermal equilibrium, with each individual computational entity operating at a net rate $b\lambda$, where λ is the gross transition rate and b is a bias factor. We show that we can maximise the rate of the whole system by letting b scale as $V^{-1/6}$, giving the total rate as $\nu \sim V^{5/6}$. It can therefore be seen that reversible computers of a given size can outperform irreversible ones.

Resource Distribution

Computers in the Turing-machine sense need access to unlimited memory (at least in principle). In practice, memory is limited, and this is even more pronounced for our amorphous computational entities.

Some amorphous computers use local interactions to solve this, but our results for communication show this to be impractical. Instead, our computational entities must exchange resources in order to function effectively.

Passive resource pools lead to trapping of computational states

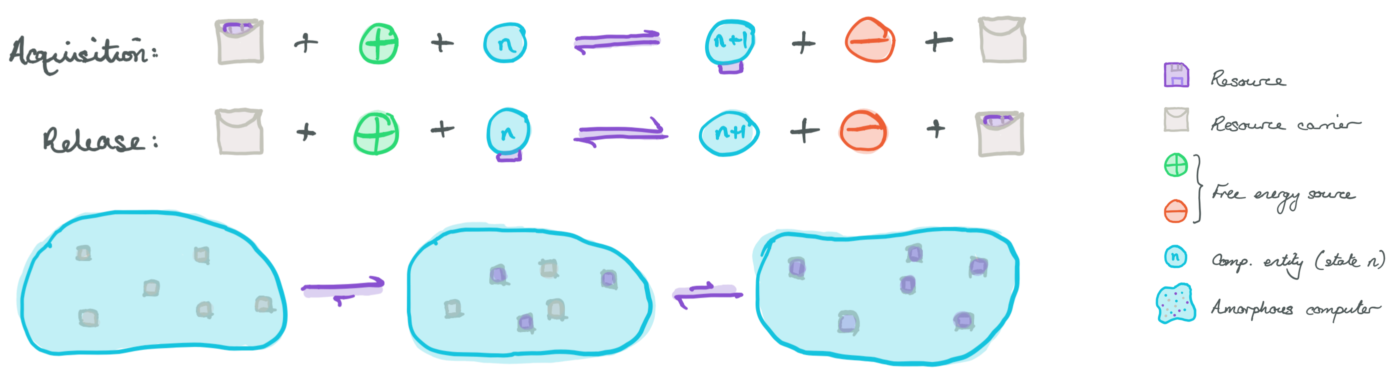

An example passive scheme analogous to the supply of amino acids (purple floppies) in cells by tRNA carriers (grey cases). Plus and minus tokens correspond to a source of free energy driving computation, shown here by the transition from state n to $n+1$. Meanwhile, the resource pool is entropically driven towards a state of half-full.

The above figure shows an example passive resource pool, and how it gives rise to a competing entropic driving force. The interaction of this with the computational driving force can lead to computational entities being 'attracted' towards states by their resource usage, rather than towards completion, thus undermining computation. We show that any passive resource pool results in this behaviour for sufficiently large systems.

A computational-entity mediated approach provides processive resource availability

We propose a server-client architecture whereby a special class of computational entities act as a distributed network of resource banks. In the forthcoming paper, an example implementation of this architecture written in `alethe` is presented and proven to provide efficient and processive resource availability.

Communication & Parallelism

As amorphous computers are composed of simple (and often slow) computational entities, cooperation between these entities is very important. This may manifest as entities meeting in physical space and or state space, often with an exchange of information. Due to the high level of reversibility, however, these meetings would normally be fleeting. As such, they may be regarded as synchronisation processes involving a constriction in phase space.

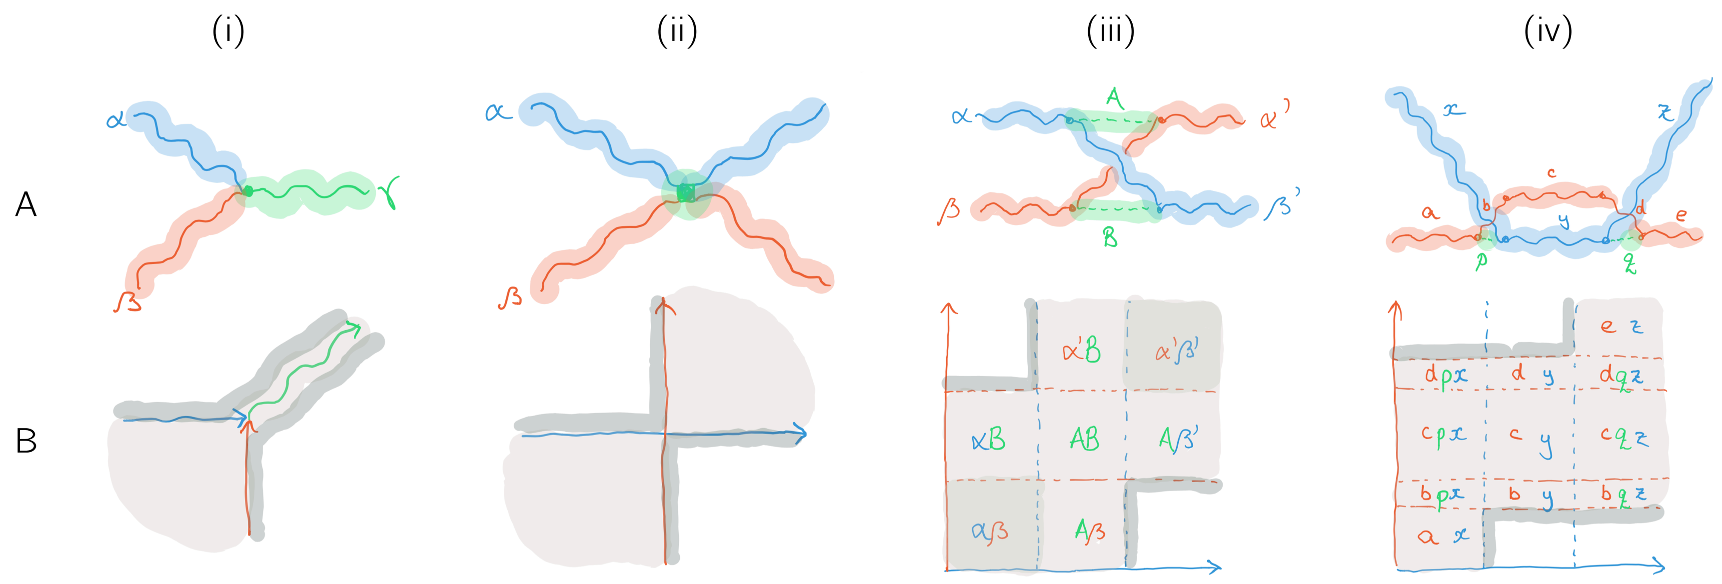

A range of increasingly sophisticated synchronisation processes. The top diagrams, A, schematically represent the worldlines of computational entities in a manner analogous to Feynman diagrams. The bottom diagrams, B, show the same scenarios in phase space.

(i) An entity fission process (its reverse being the corresponding fusion process).
(ii) A symmetric communication event.
(iii) A fly-by communication event; the dashed green lines indicate the deposit of a *dormant* entity, i.e. a non-computing entity carrying a data payload without the creation of an additional dimension in phase space. Notice the wider constriction.
(iv) An asymmetric double fly-by communication event. The double fly-by permits the two entities to converse on their shared state before departing.

In diagrams B, the drift through phase space is northeast. Passage through a constriction requires both entities to simultaneously enter a specific state or set of states. Whilst in an irreversible system this is not unreasonable, a little thought shows that this may be a rare event in reversible systems.

We conjectured that such passage events incur a time penalty on top of the time to pass through unconstrained phase space, being $x/b\lambda$ where x is the phase distance traversed[1]. We proceeded to investigate it for continuous and discrete phase space, proving it for the former and confirming it experimentally in the latter.

Continuous phase space

We used a Fokker-Planck approach to find the mean first passage time (MFPT) for any symmetric 2D geometry, expressed as an integral. For the symmetric fly-by case, we obtain a penalty of

$$\frac{2}{3\lambda}\frac{1}{b^2}\log\left(1+\frac{2x}{w}\right)$$

where the log term corresponds to the entropy loss of a phase particle squeezed through the constriction, with w being the constriction width.

Discrete phase space

Solving for discrete MFPTs is markedly harder than the continuous case. We have made progress towards an analytical solution for the $w=1$ case. For general w, we have employed a Markov chain Monte Carlo (MCMC) approach to experimentally measure the penalty. Results are congruent with that for continuous phase space for moderate x (as indicated by the trend lines), but for larger x the penalty appears to plateau. We hope that an analytical solution will explain this discrepancy.

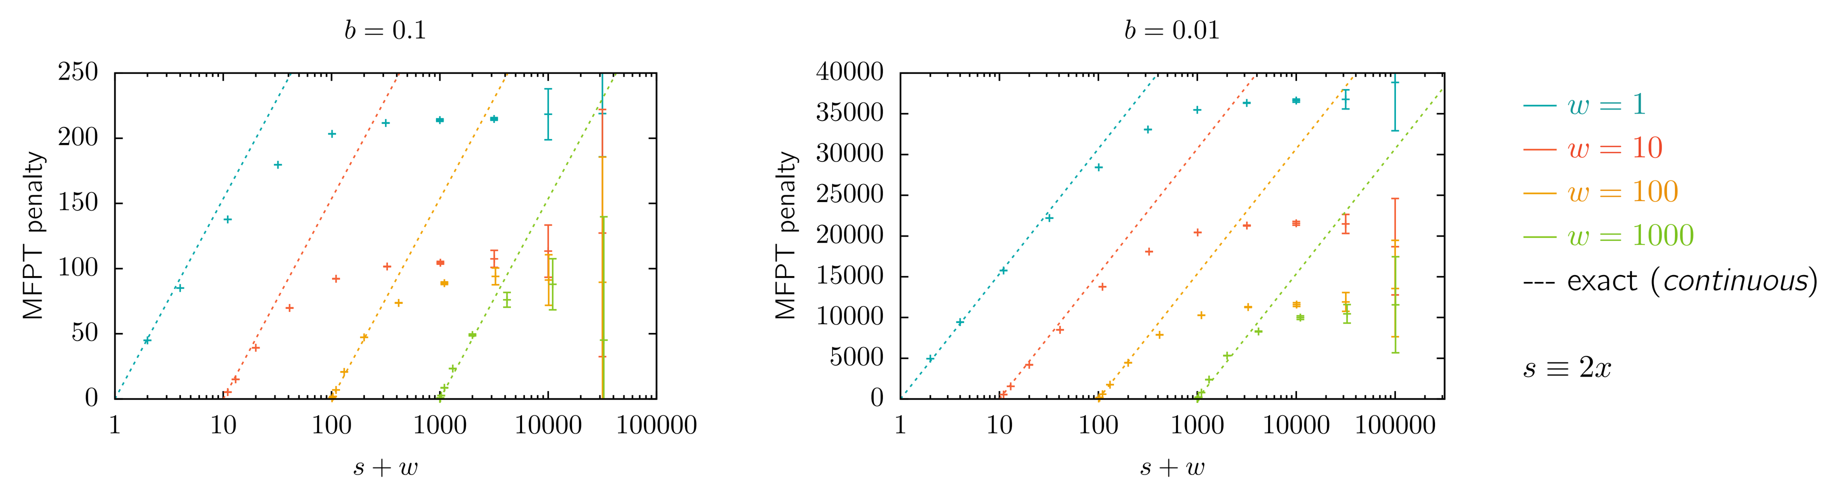

As $b \to 0$, communication 'freezes out'

For large systems, synchronisation events become prohibitively expensive and a scarce resource in comparison to concurrent serial computation. A range of management techniques will be discussed in the forthcoming paper.

[1] Assuming the traversal is northeasterly. For other directions, this may be reduced by up to 50%.

UNIVERSITY OF CAMBRIDGE

Department of Applied Mathematics and Theoretical Physics (DAMTP)

MICKLEM LAB

EPSRC

The author would like to acknowledge the support of the Engineering and Physical Sciences Research Council (EPSRC), project reference 1781682 for "*Modelling approaches to molecular computation*".

References

[1] Norman Margolus and Lev B Levitin. 'The maximum speed of dynamical evolution'. In: *Physica D* 120.1 (Sept. 1998), pp. 188–195.
[2] Charles H Bennett. 'The thermodynamics of computation—a review'. In: *Int. J. Theor. Phys.* 21.12 (Dec. 1982), pp. 905–940.