Cambridge-Imperial Computational PhD Seminar *19th February 2020*

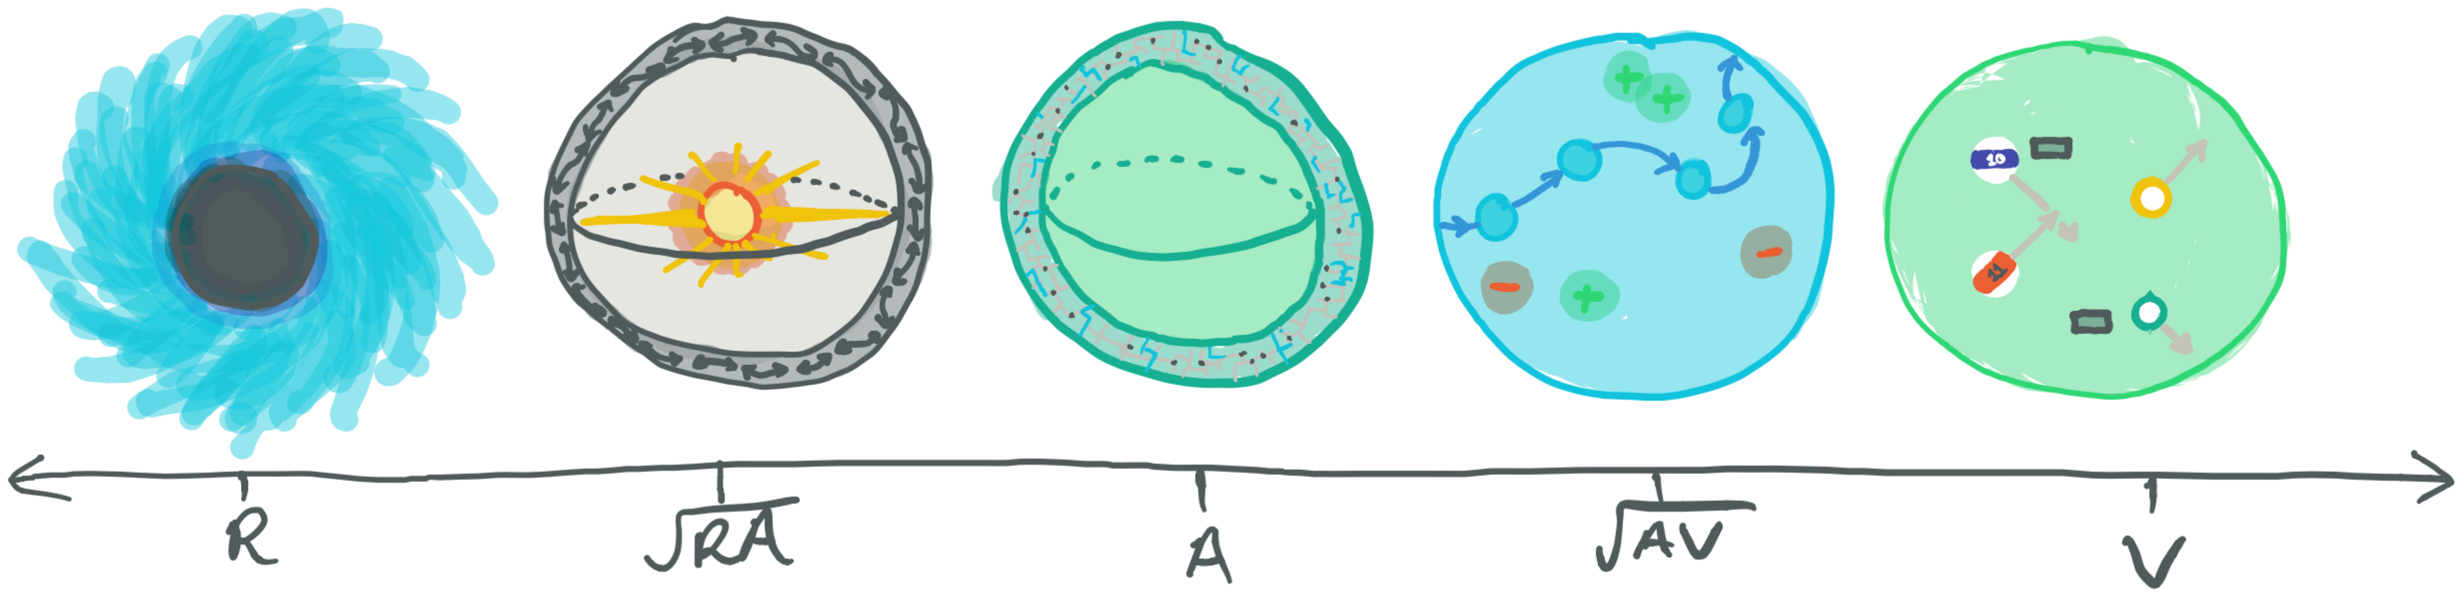

Engines of Parsimony

How to take over the universe on a budget

William Earley

Micklem Lab · DAMTP

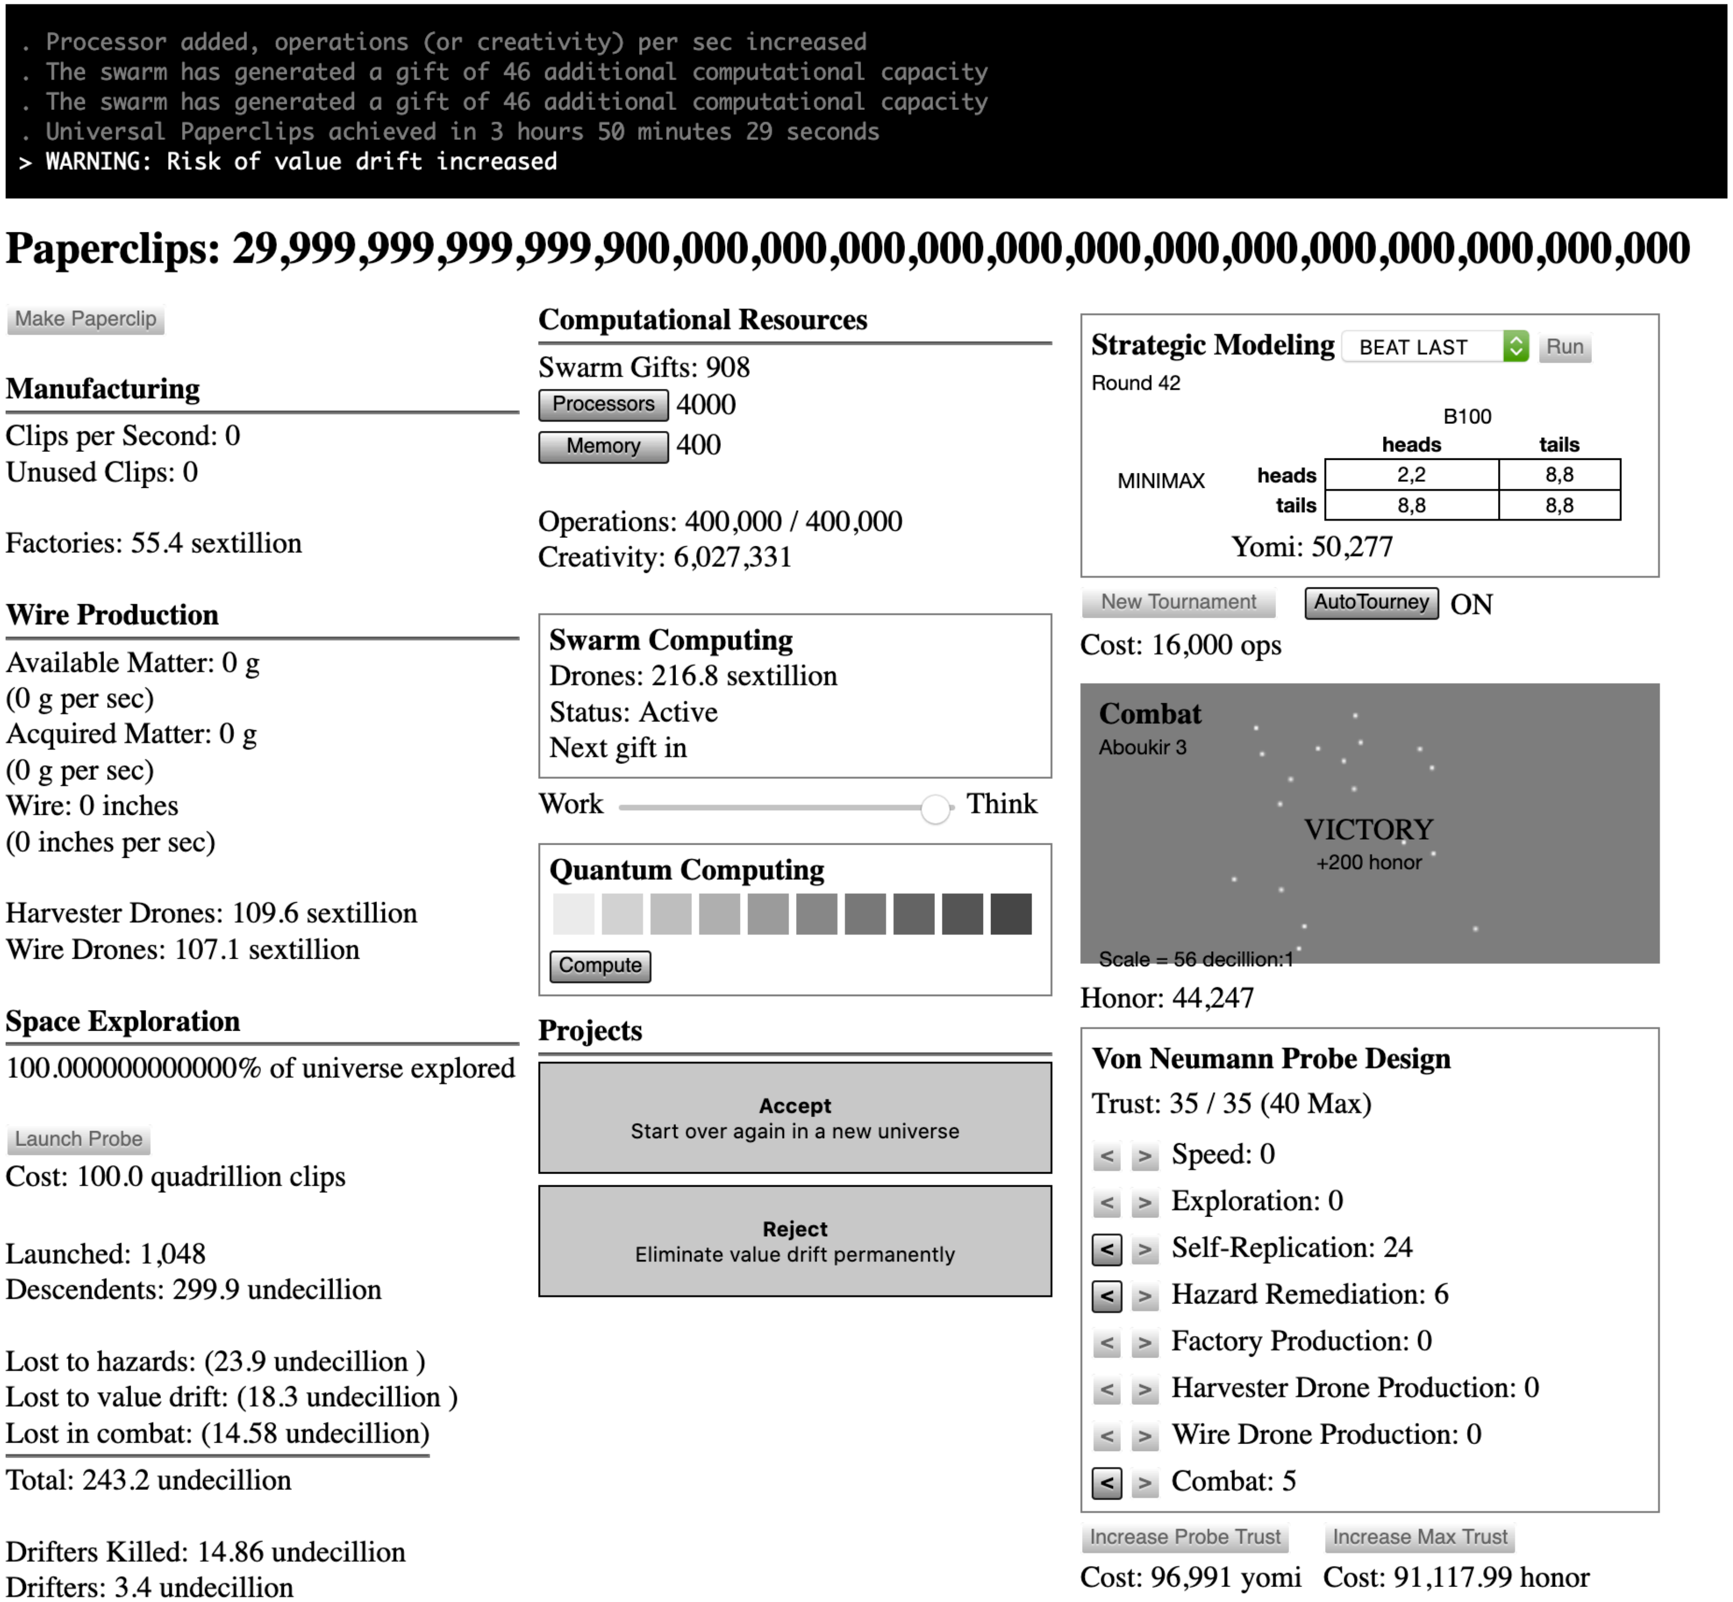

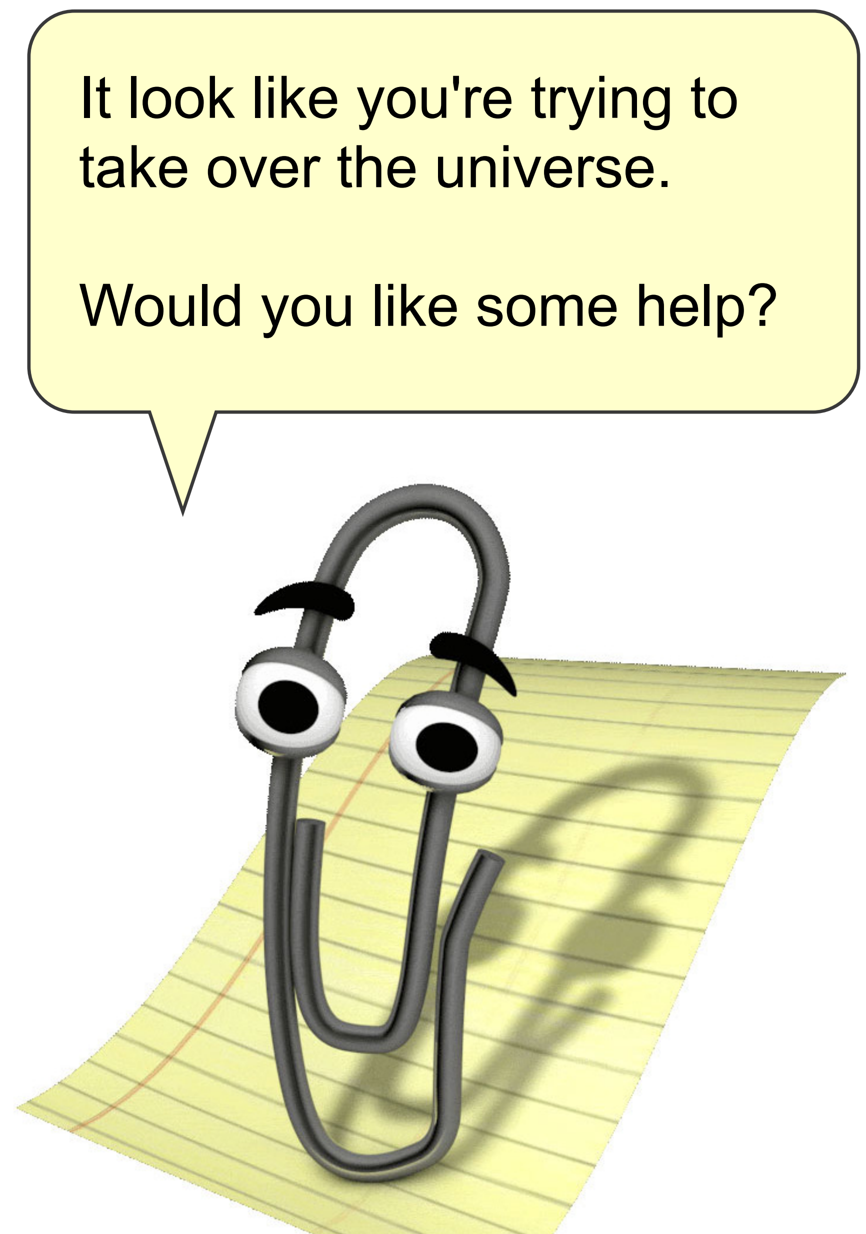

Further reading:

- *Instrumental Convergence*
- https://wiki.lesswrong.com/wiki/Paperclip_maximizer
- *Universal Paperclips*

Conventional Computing

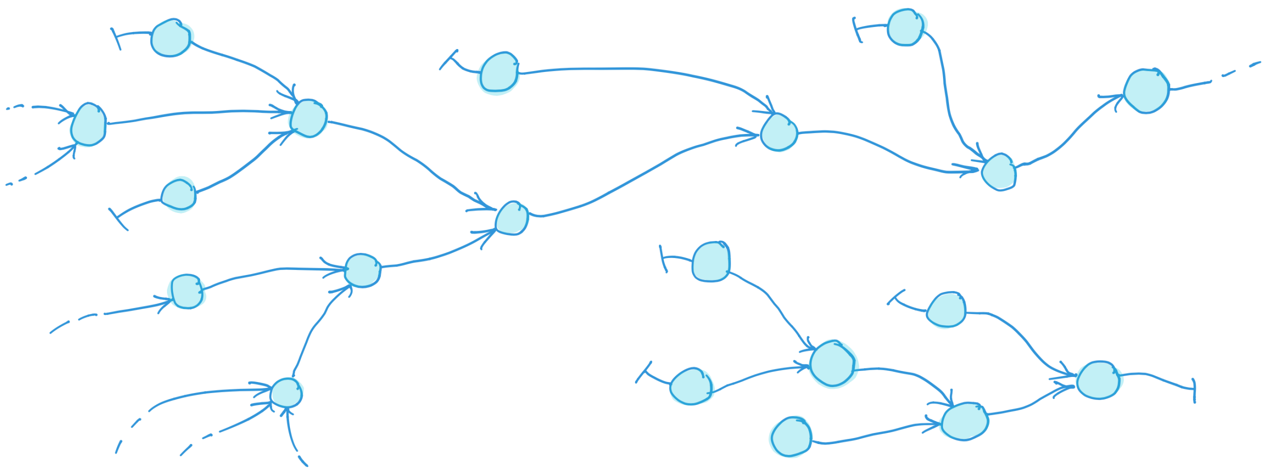

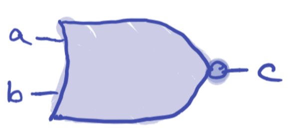

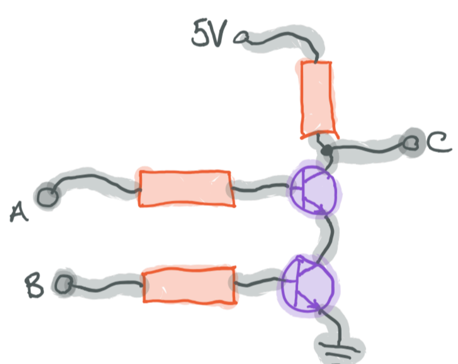

a	b	c
0	0	1
0	1	1
1	0	1
1	1	0

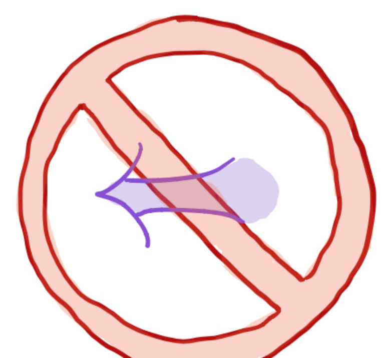

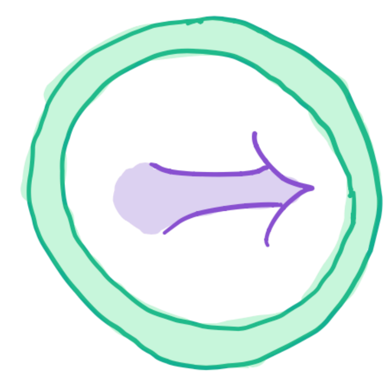

Physical Constraints

Maximum rate of dynamical evolution

$$\nu \leq E/h_p$$

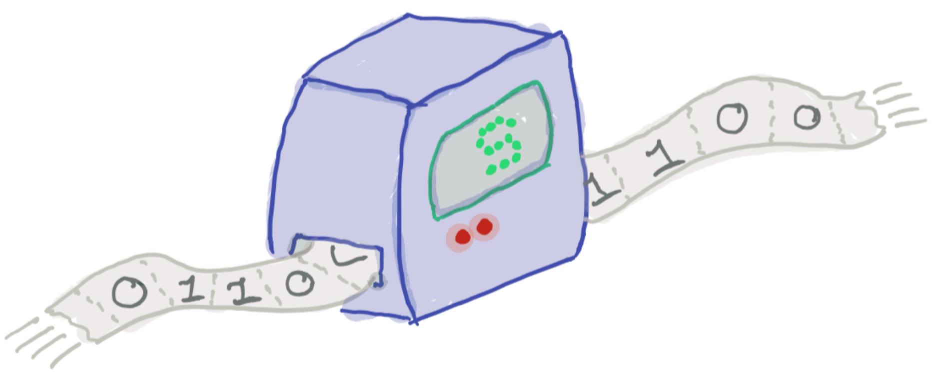

Entropy and heat generation due to irreversibility

$$I = -\sum p_i \log p_i$$

$$\Delta I = -\sum \Delta p_i \log p_i + \mathcal{O}(\Delta p_i^2)$$

$$\Delta Q \geq k_B T \Delta I$$

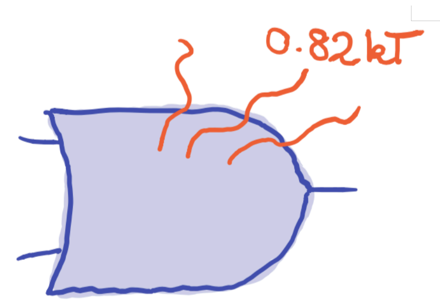

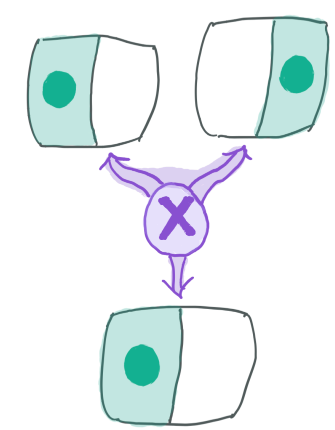

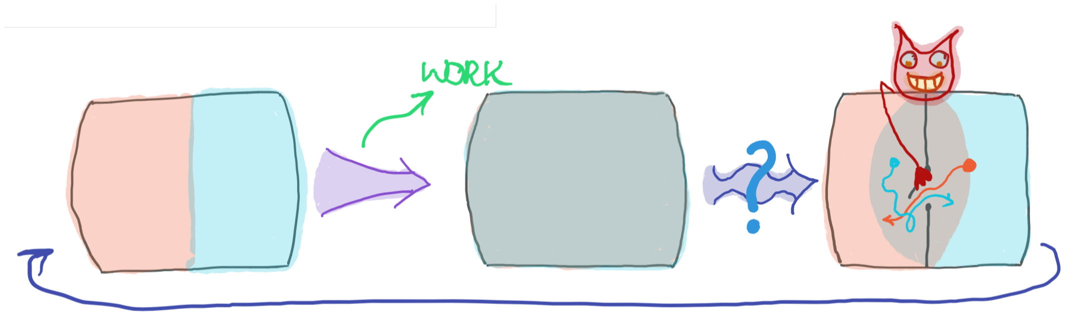

Maxwell's Dæmon

References:

- *HJ Bremermann [Self-organizing systems 1962]*
- *N Margolus and LB Levitin [Physica D 1998]*
- *L Szilard [Zeitschrift für Physik 1929] / L Szilard [Behavioural Science 1964]*
- *R Landauer [IBM J. Res. Dev. 1961]*

Consequences

- Computers* must dissipate heat*

$$P \equiv \dot{Q} \geq kT\nu_c \delta I$$

- The rate of heat generation is bounded from below, and scales proportionally to the rate of computation

$$\nu_C \leq E/h_p = \frac{\rho c^2 V}{h_P}$$

- The rate of computation scales with the energy/volume

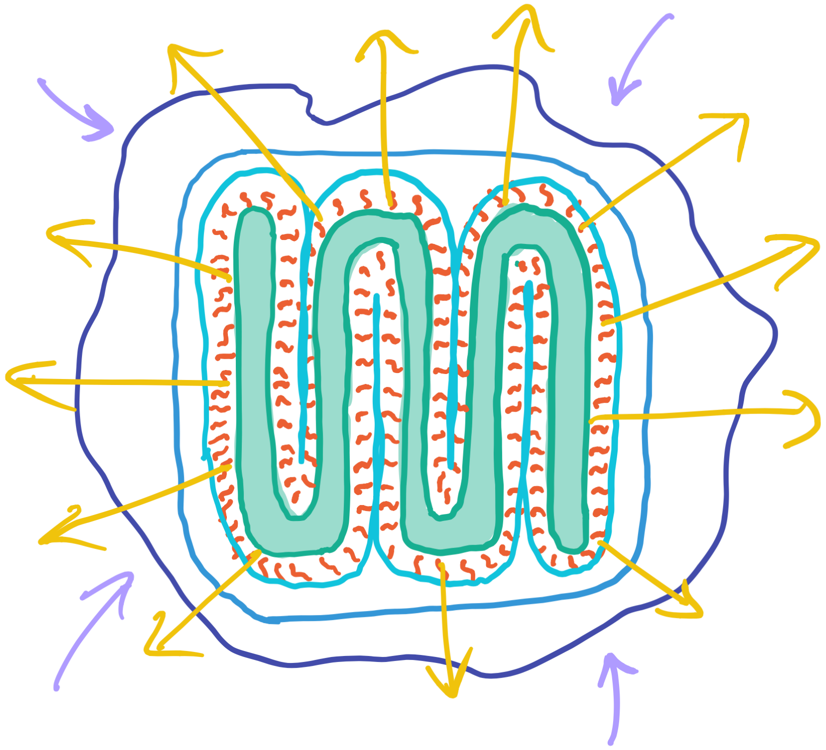

- The rate of heat dissipation and power input is limited by the convex bounding surface

$$P \leq \phi A$$

- This implies the maximum rate of computation is bounded proportional to the (convex bounding) surface area

$$\nu_C \leq \min\left(\frac{\rho c^2 V}{h_P}, \frac{\phi A}{kT\delta I}\right)$$

Irreversible Computing

Can we do better?

Reversible Computing

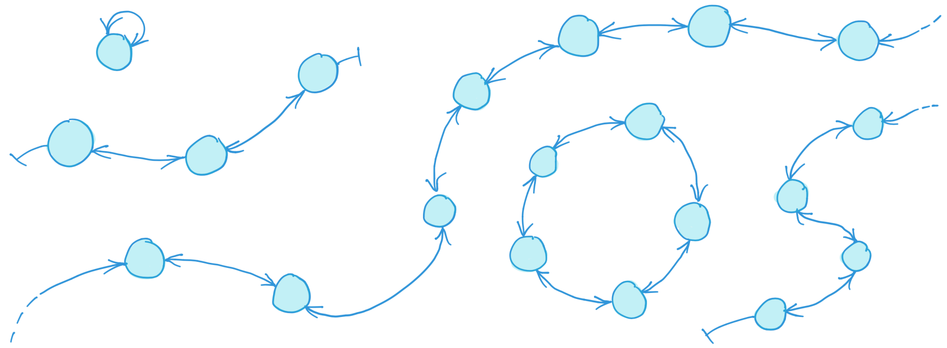

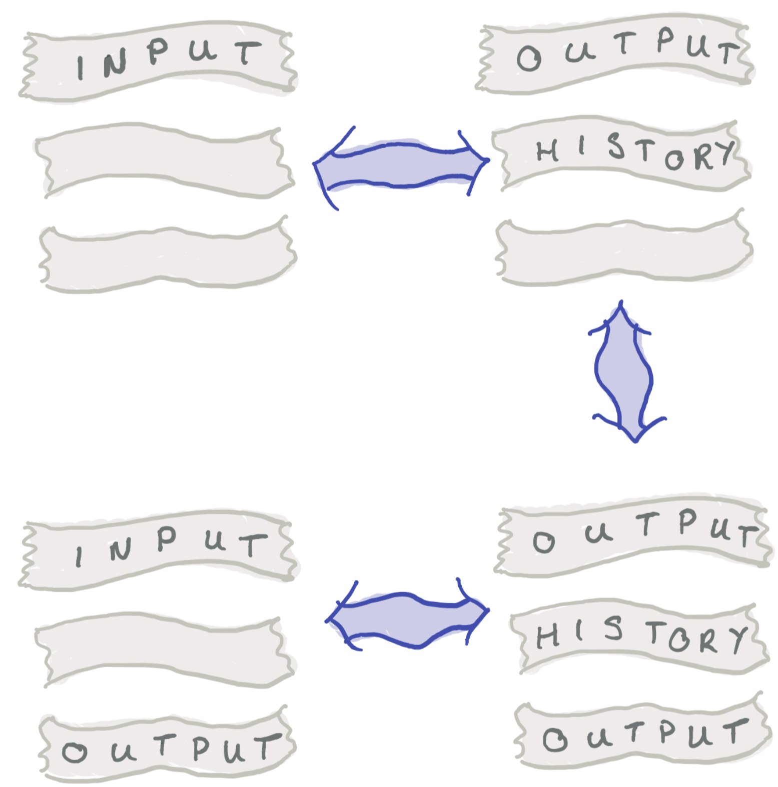

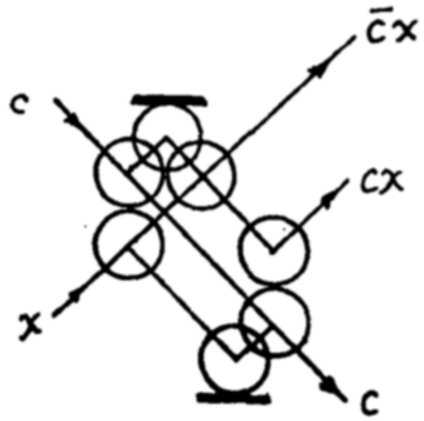

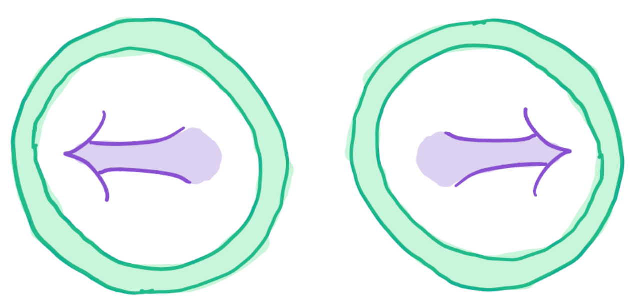

References:

- *CH Bennett [IBM J. Res. Dev 1973]*
- *E Fredkin and T Toffoli [Int. J. Theor. Phys. 1982]*
- *CH Bennett [Int. J. Theor. Phys. 1982]*

Irreversible Computing

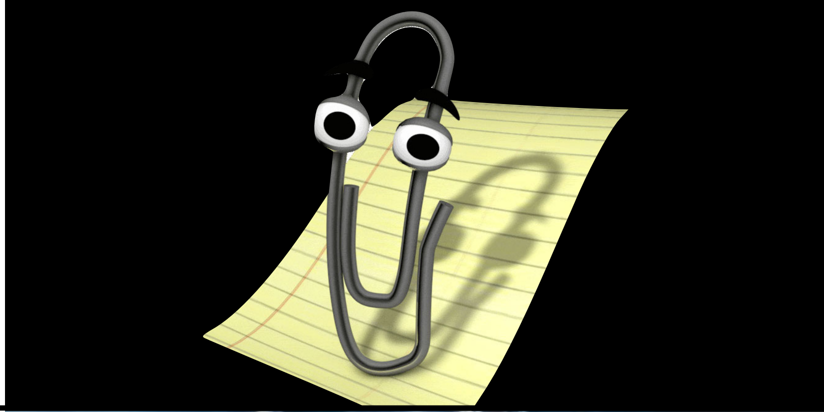

Reversible Computing

Ballistic Computation

> “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.

References:

- *CH Bennett [Int. J. Theor. Phys. 1982]*

Classical Mechanics

- Introduce a general Lagrangian for arbitrary coordinates

$$\mathcal{L} = \tfrac{1}{2}\dot{q}^\top C\dot{q} - V - Q\dot{q} + \mathcal{O}(\dot{q}^3)$$

- Determine the change in kinetic energy due to (brief) collisions; we must restore this kinetic energy to computational dofs

$$\Delta\mathcal{H}^{(m)} \approx \eta|\tilde{u}^{(m,n)\top}\mu^{(m,n)}(u^{(m)} - u^{(n)})|$$

- For fast computational 'particles', get hydrodynamic drag-like decay and an areametric rate of computation

$$P \geq 2N\left\langle \sum_\beta \tfrac{1}{2}\rho^{(\beta)}\bar{u}^{(\alpha)3}\eta^{(\alpha;\beta)}A^{(\alpha)} \right\rangle_\alpha$$

$$\nu_c \leq \sqrt[3]{PN^2\langle(\textstyle\sum_\beta \rho^{(\beta)}\eta^{(\alpha;\beta)}A^{(\alpha)}\ell^{(\alpha)3})^{-1/2}\rangle_\alpha}$$

$$\nu_C/N \not\ll 1$$

- For slow computational 'particles', get Stokes drag-like decay and a ***super***-areametric rate of computation

$$P \geq \sum_\alpha \frac{R^{(\alpha)2}}{N^{(\alpha)}} \underbrace{A^{(\alpha)}\ell^{(\alpha)2}\sum_\beta \eta^{(\alpha;\beta)}\rho^{(\beta)}\bar{u}^{(\beta)}}_{\gamma^{(\alpha)}}$$

$$\nu_c \leq \sqrt{PN\langle 1/\gamma^{(\alpha)}\rangle_\alpha} \sim \sqrt{AV} \sim V^{5/6}$$

Brownian Computation

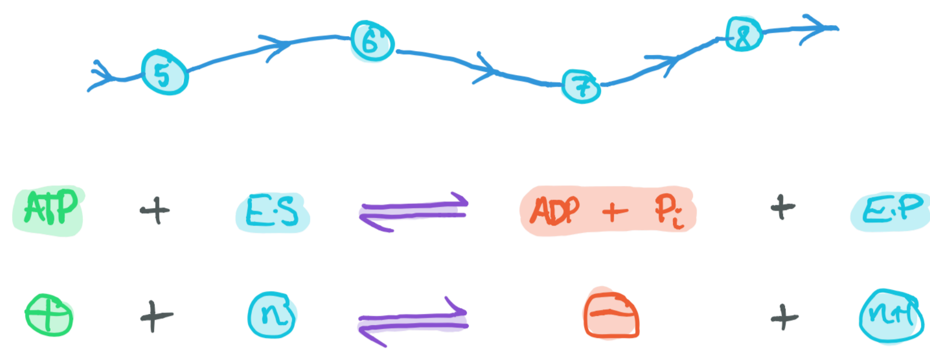

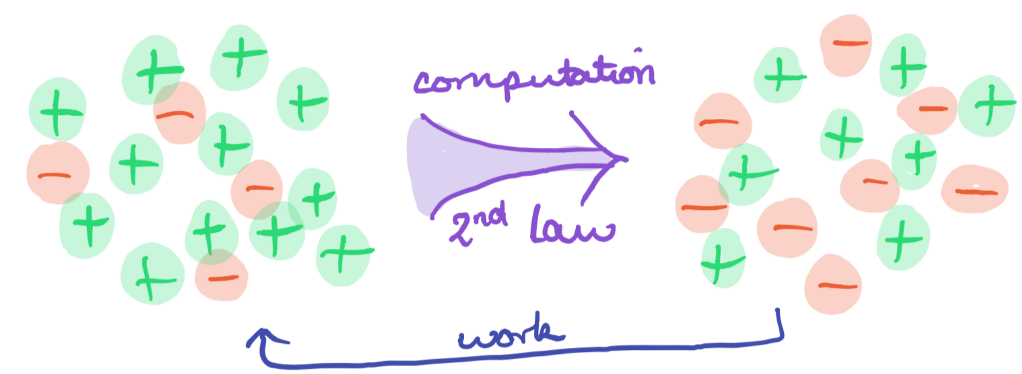

- Introduce a CRN with species and reactions

- Determine entropy change due to a reaction, and relate to reaction rate using detailed balance and 2^{nd}LoTD

- Rate of entropy change given distance from equilibrium/bias

- Maximising, we recover the same super-areametric scaling

$$\Gamma = \{\nu_{ij} X_j \longleftrightarrow_{\gamma_i} \nu'_{ij} X_j : i\}$$

$$\eta = \sum_j [X_j](\varepsilon_j + 1 - \log \lambda_0 [X_j])$$

$$\dot{\eta} = \sum_{ij} (\nu'_{ij} - \nu_{ij})(\overrightarrow{r}_i - \overleftarrow{r}_i)(\varepsilon_j - \log \lambda_0 [X_j])$$

$$= \sum_i b_i (\overrightarrow{r}_i + \overleftarrow{r}_i) \log \frac{\overleftarrow{r}_i}{\overrightarrow{r}_i}$$

$$\dot{H} = 2R_0 \langle b \,\mathrm{arctanh}\, b \rangle$$

$$\nu_c \leq \sqrt{\frac{PR_0}{2k_B T}} \sim \sqrt{AV} \sim V^{5/6}$$

Molecular Programming

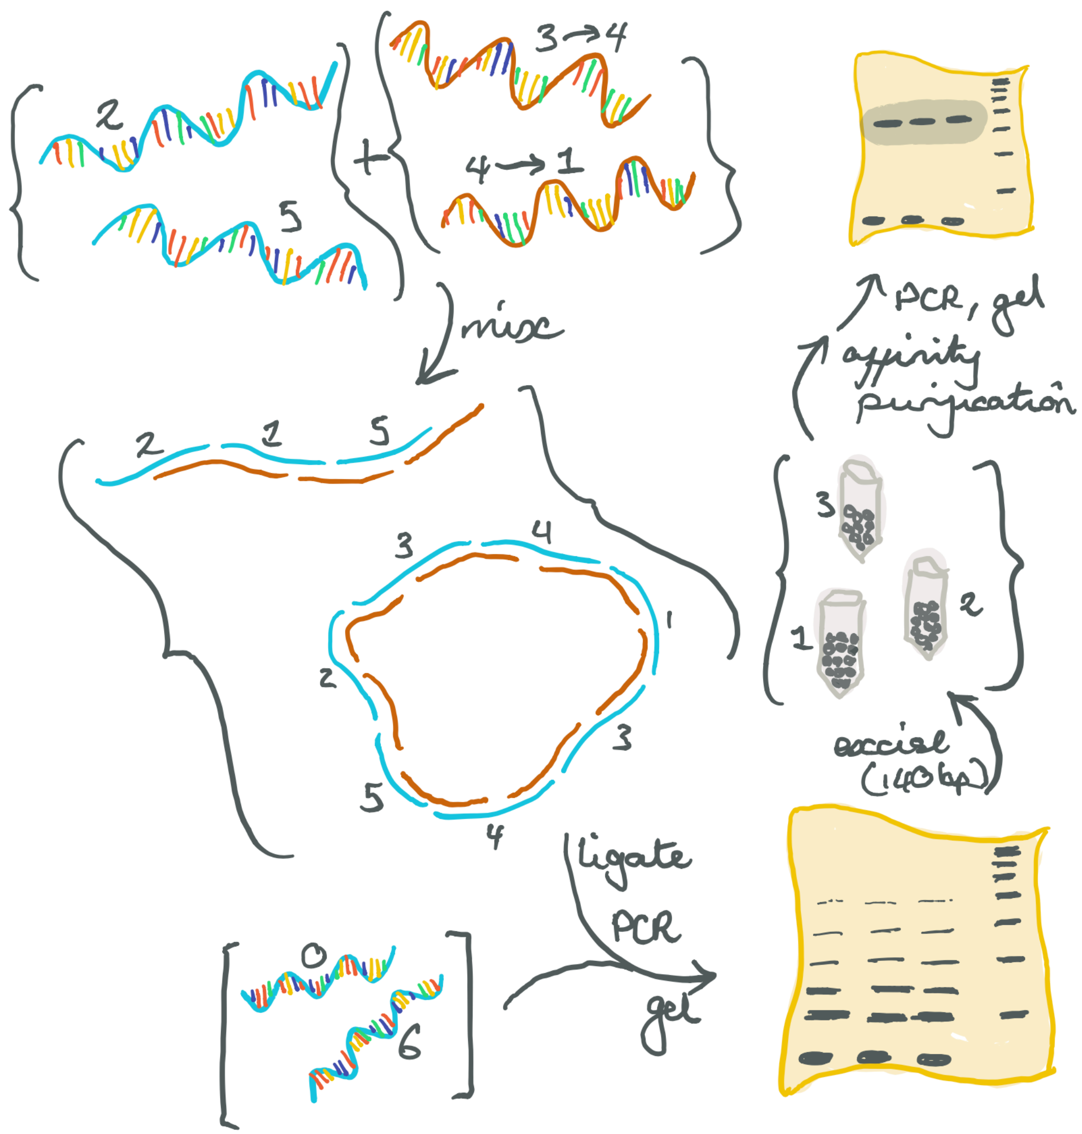

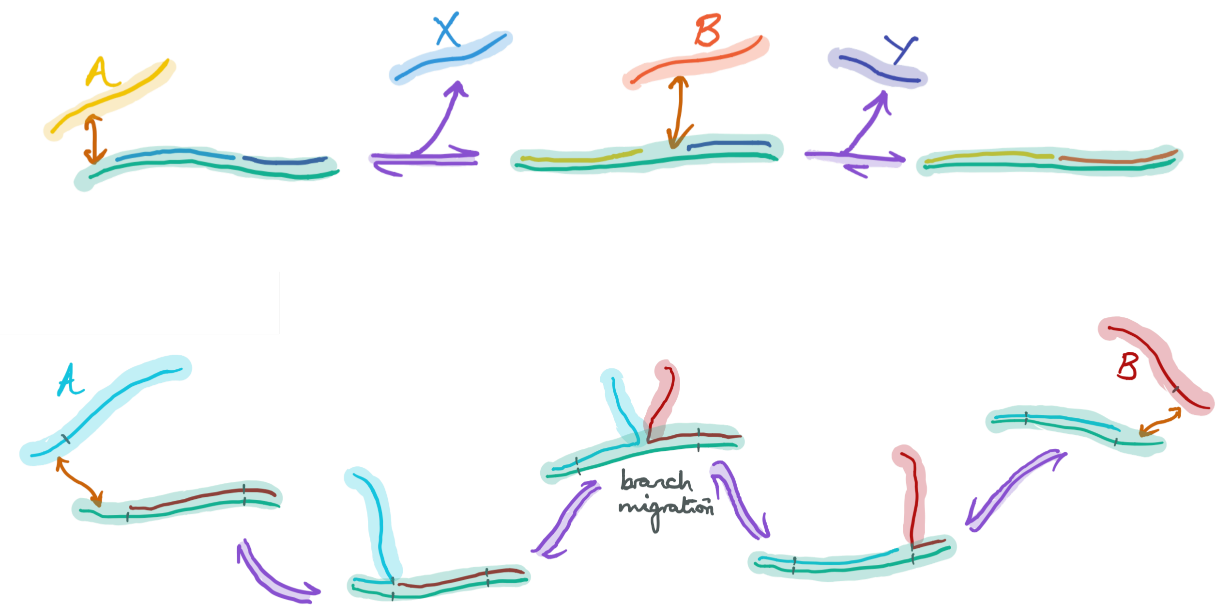

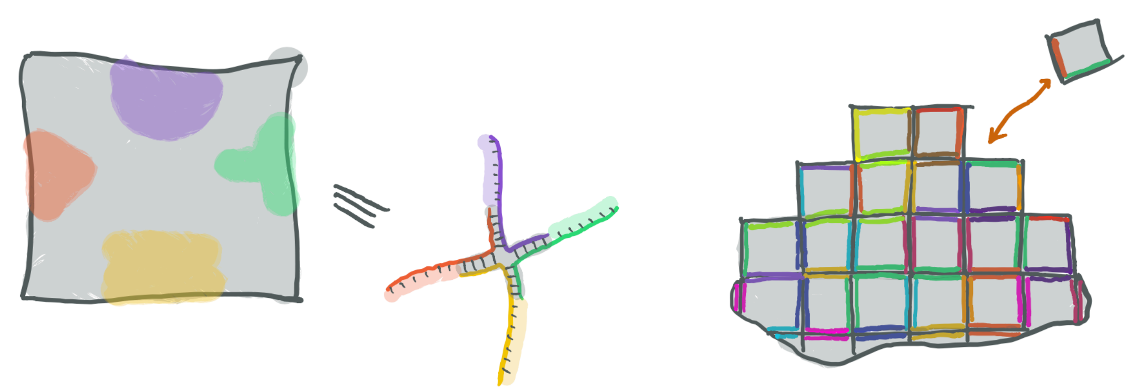

References:

- *LM Adleman [Science 1994]*
- *G Seelig et al. [Science 2006]*
- *E Winfree et al. [Nature 1998; PhD Thesis 1998]*

Irreversible Computing

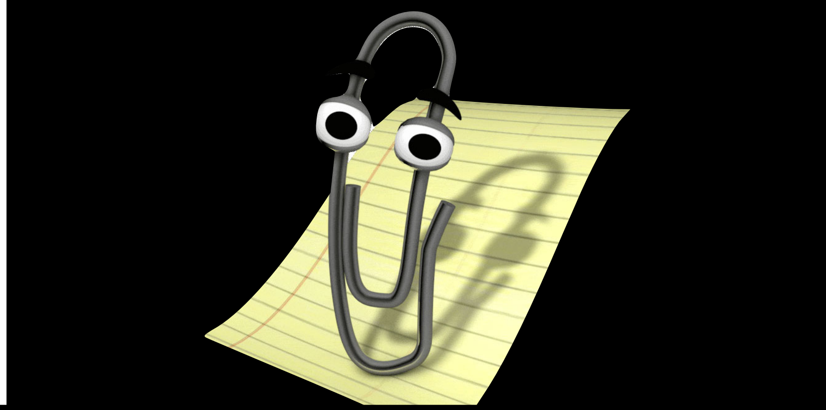

Reversible Computing

Brownian Computing

Quantum Computing

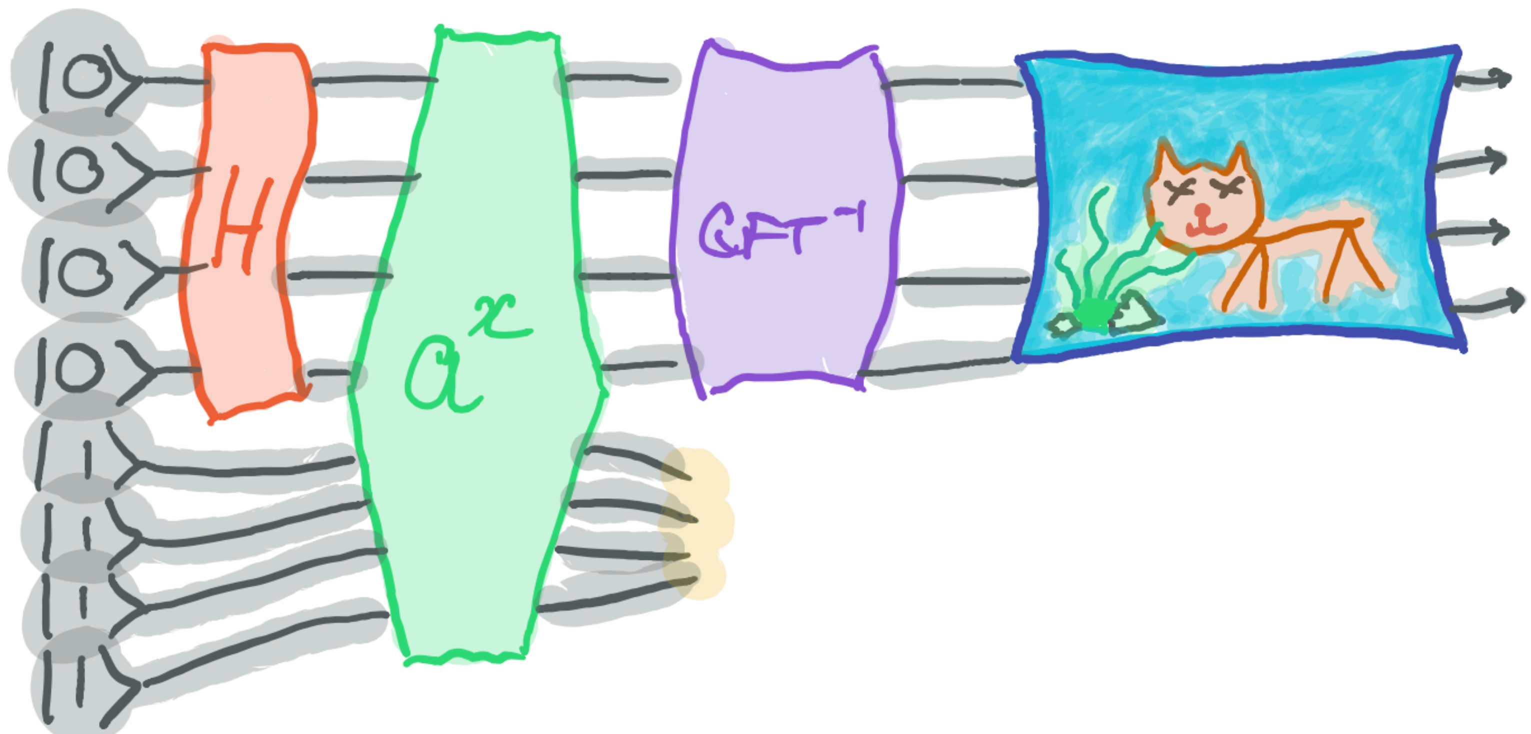

- Intrinsically reversible because Schrödinger equation is unitary
- Can we use superconducting/superfluid systems for ballistic reversible computation?

qZeno Computing

- Consider a quantum computer with subsystems evolving under a perturbed ballistic hamiltonian

$$H = H_0 + V(t)$$

$$\delta p = \left(\frac{\delta t}{\hbar_P}\right)^2 \left(\langle V^2 \rangle - \langle VPV \rangle\right)$$
$$= \left(\frac{\delta t}{\hbar_P}\right)^2 \zeta n (k_B T)^2 \qquad \zeta \in [\tfrac{1}{2}, 1)$$

- Periodically check for and correct errors in the basis $(P, P^{\perp})$; if frequency sufficiently rapid, observe the Quantum Zeno Effect

$$\dot{h} \geq \frac{\zeta \log g}{\nu_Z} \left(\frac{n k_B T}{\hbar_P}\right)^2$$

- Determine the entropy generation rate and the maximal computation rate

$$\nu_C \leq \frac{1}{2\pi} \left\langle \frac{\zeta \log g}{\beta^2 \varepsilon^2} \right\rangle^{-1/2} \sqrt{\frac{P}{k_B \hat{T}} \frac{E_Z}{h_P}}$$
$$\lesssim \sqrt{AV} \sim V^{5/6}$$

Irreversible Computing

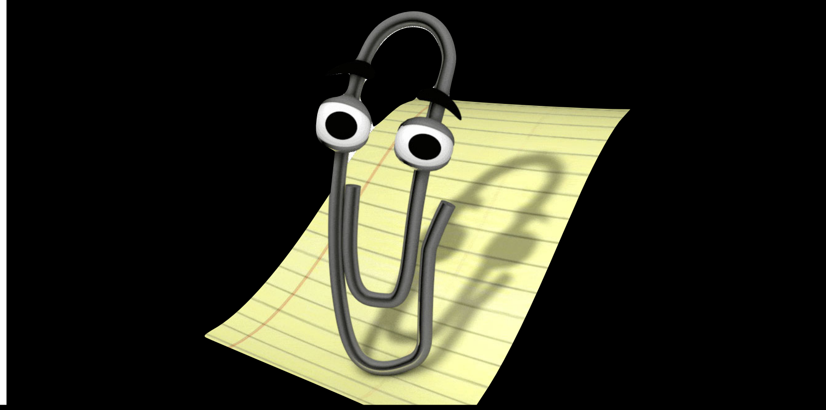

Reversible Computing

Brownian Computing

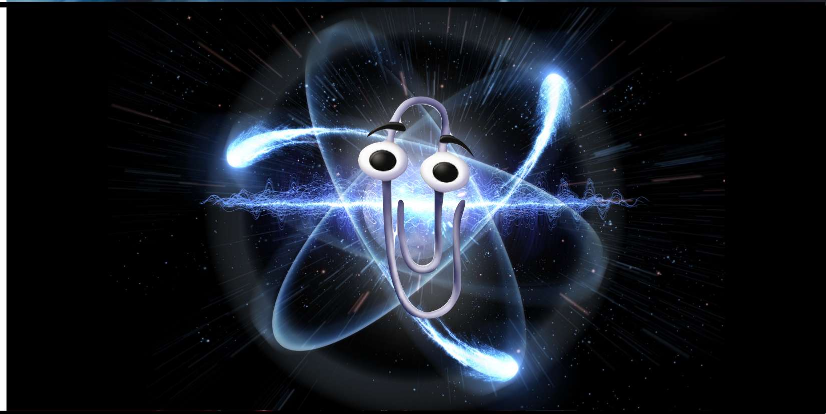

Quantum Computing

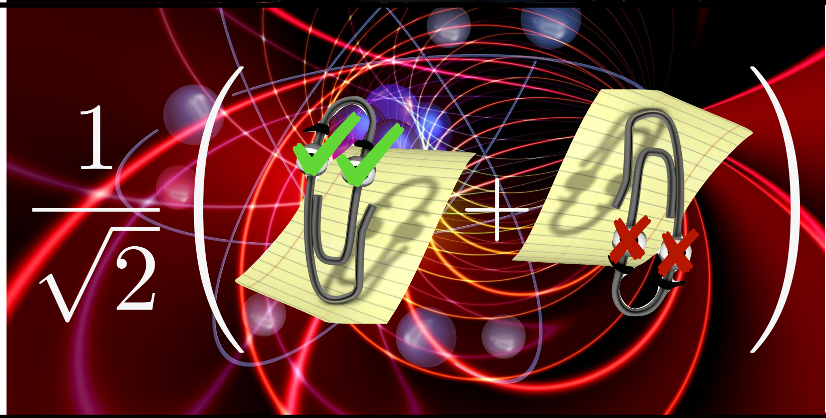

Summary So Far

- Irreversible computing has an areametric (amortised) rate limit due to thermodynamic-geometric constraints
- Reversible computing can in principle use the entire computational volume (Margolus-Levitin)
- In practice still subject to thermodynamic constraints, but both classical and quantum computers can achieve super-areametric scaling
- Molecular computing may be a convenient/ideal approach

General Relativity

- When our computer gets too large, threat of collapse

$$r_S = \frac{2GM}{c^2}$$

- From then on, mass must scale linearly in radius; there are two computational rate regimes

$$\nu_C \lesssim \min(\sqrt{PM}, E/h_P)$$
$$\sim \min(V^{1/2}, V^{1/3})$$

- We also need to take into account (gravitational) time dilation; use the Tolman–Oppenheimer–Volkoff equation

$$\frac{\mathrm{d}\log\sqrt{g_{00}}}{\mathrm{d}v} = \frac{vv_s}{2}\frac{\mu + 3u}{\mu_1 - v^2 v_s \mu}$$
$$\frac{\mathrm{d}u}{\mathrm{d}v} = -(1+u)\frac{\mathrm{d}\log\sqrt{g_{00}}}{\mathrm{d}v}$$

- Need to construct systems as thick shells with outer radius

$$\frac{\nu_C|_\infty}{\nu_C|_{\mathrm{local}}} = \frac{1}{M}\int \mathrm{d}V\,\rho\sqrt{g_{00}}$$

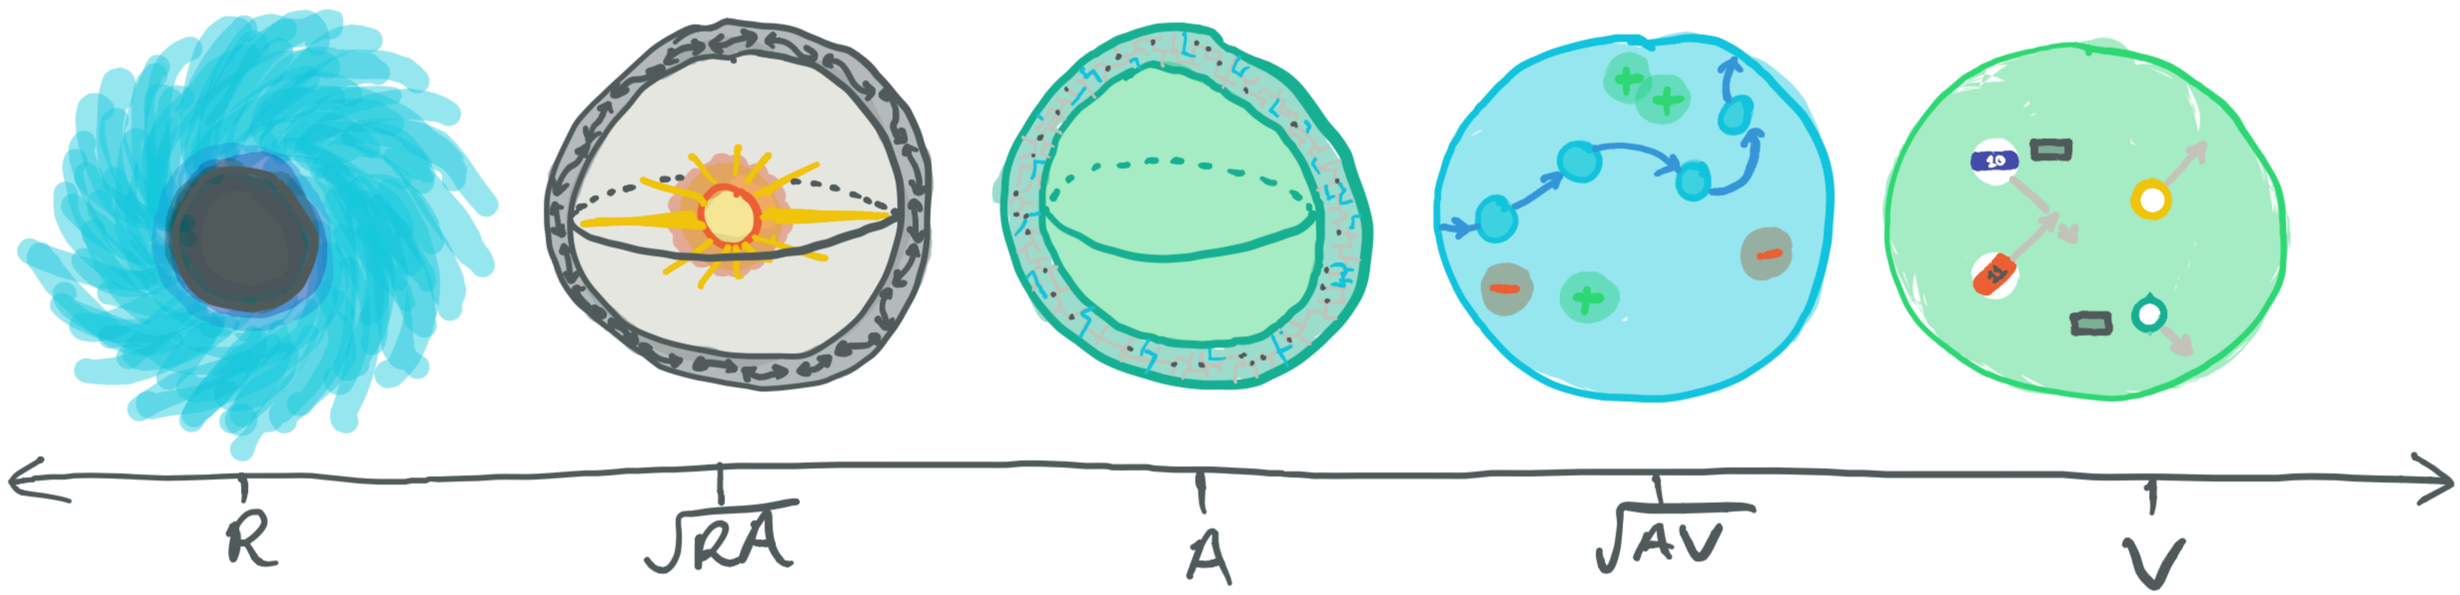
R
√RA
A
√AV
V

Scaling

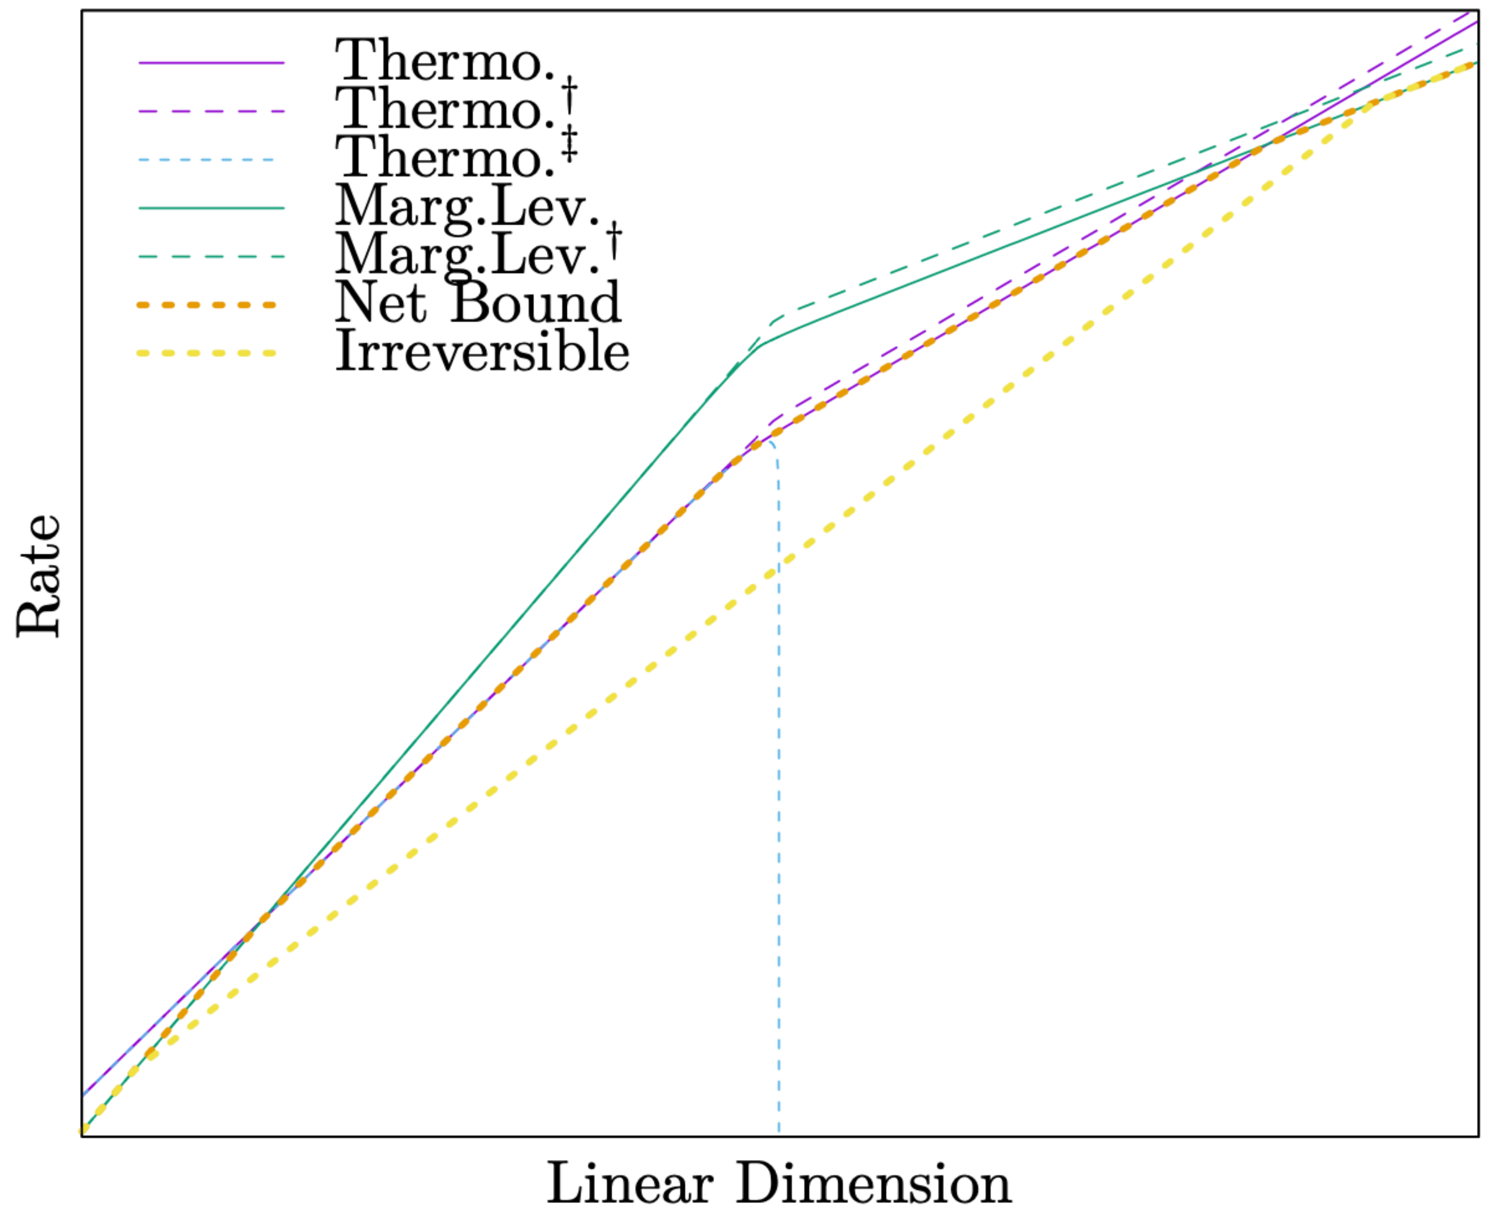
Thermo.
Thermo.†
Thermo.‡
Marg.Lev.
Marg.Lev.†
Net Bound
Irreversible
Rate
Linear Dimension

Irreversible Computing

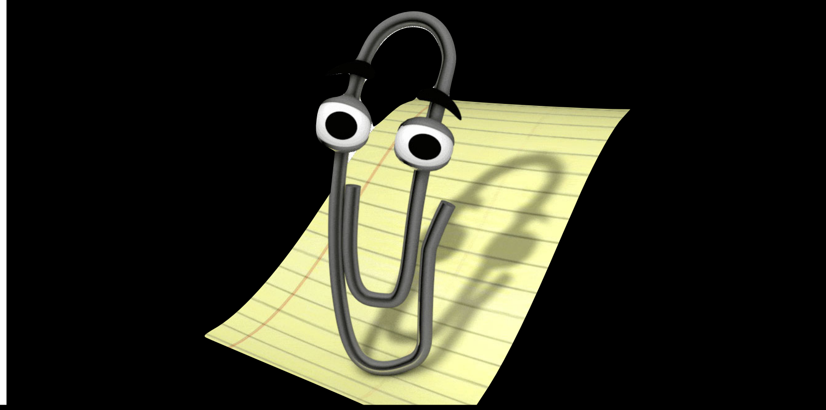

Reversible Computing

Brownian Computing

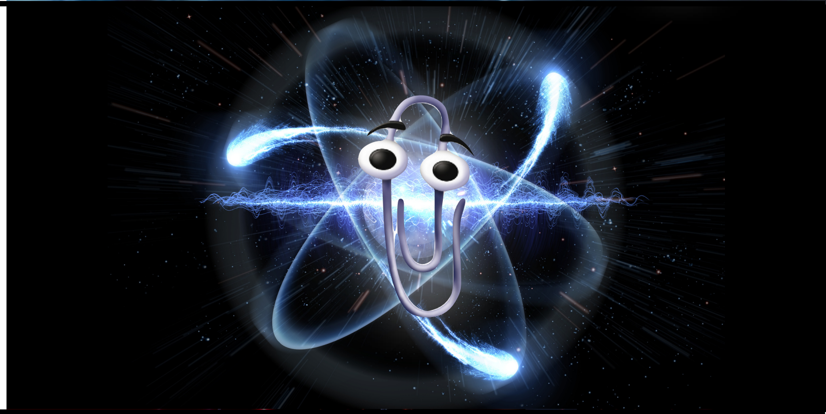

Quantum Computing

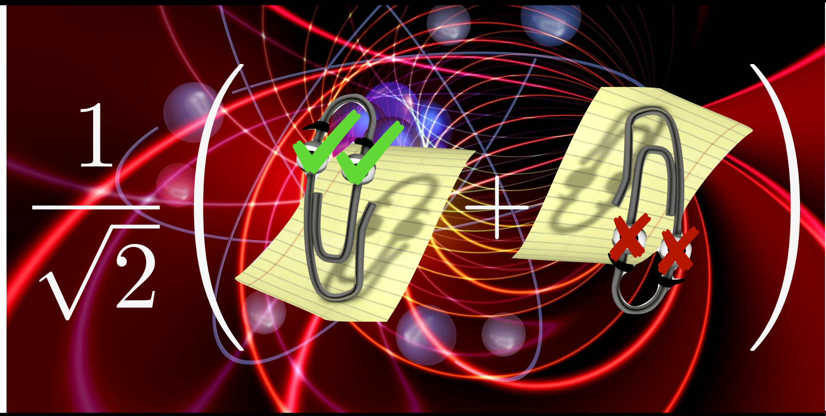

Peri-Black Hole Computing

Credit: Microsoft; BBC; UNM Newsroom; CC0/phys.org; Unknown

Further Work

- The results here concern independent computation

- Things are complicated when the individual computational loci interact such as by communication and resource distribution

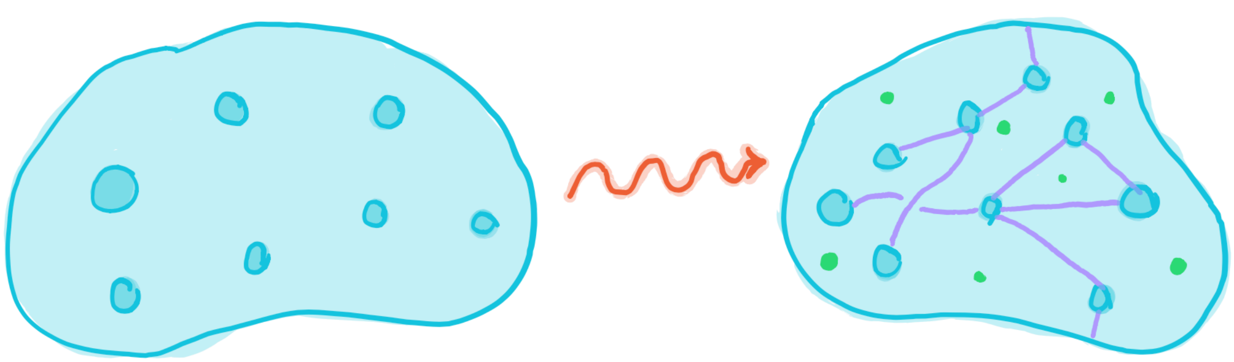

- In fact, these interaction events can only occur at an areametric rate, imposing significant constraints on the design and programming of such systems

Thank you!

UNIVERSITY OF CAMBRIDGE

EPSRC

Engineering and Physical Sciences Research Council

Department of Applied Mathematics and Theoretical Physics (DAMTP)

Further Reading & References:

- *Instrumental Convergence*
- https://wiki.lesswrong.com/wiki/Paperclip_maximizer
- *Universal Paperclips*
- *CH Bennett [IBM J. Res. Dev 1973]* — *founding of reversible computing*
- *E Fredkin and T Toffoli [Int. J. Theor. Phys. 1982]* — *ballistic computation*
- *CH Bennett [Int. J. Theor. Phys. 1982]* — *review of reversible computing*
- *LM Adleman [Science 1994]* — *founding of molecular computation*
- *EPSRC Project Reference 1781682: "Modelling approaches to molecular computation"*

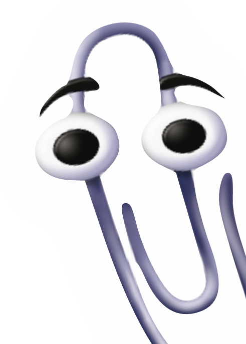

Communication

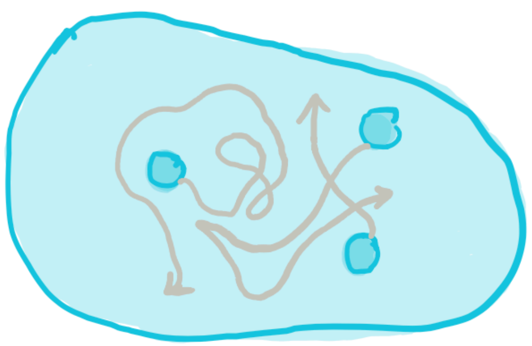

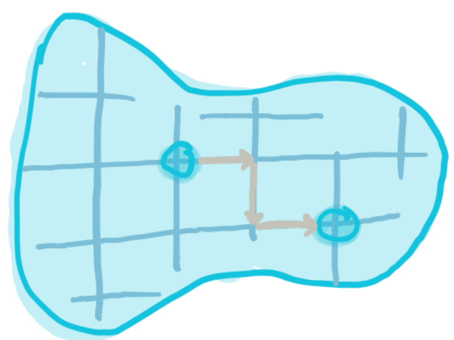

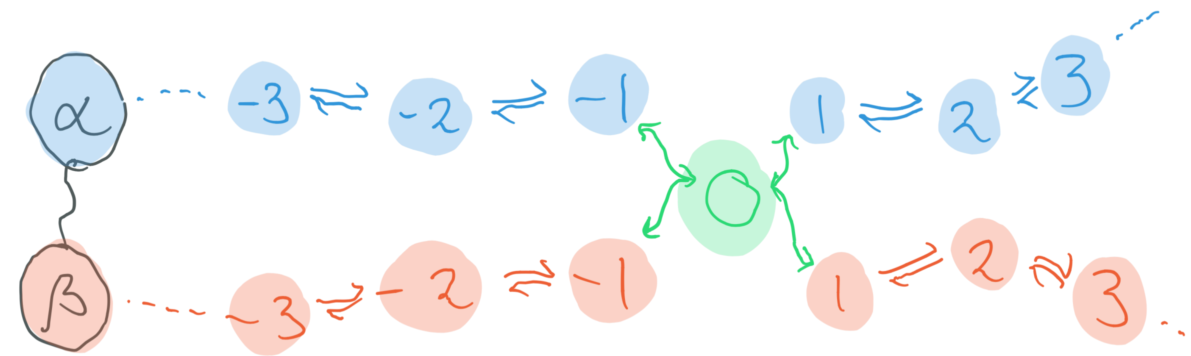

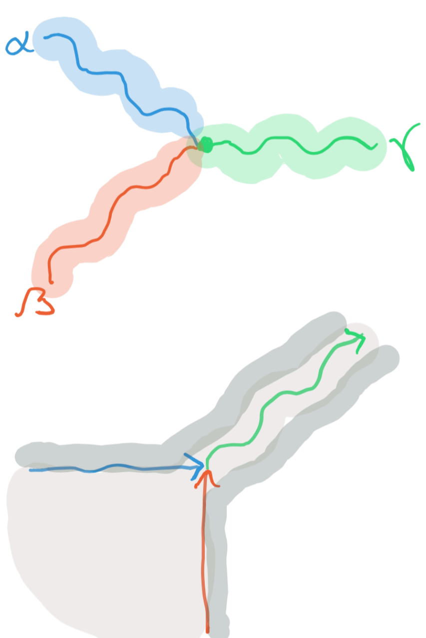

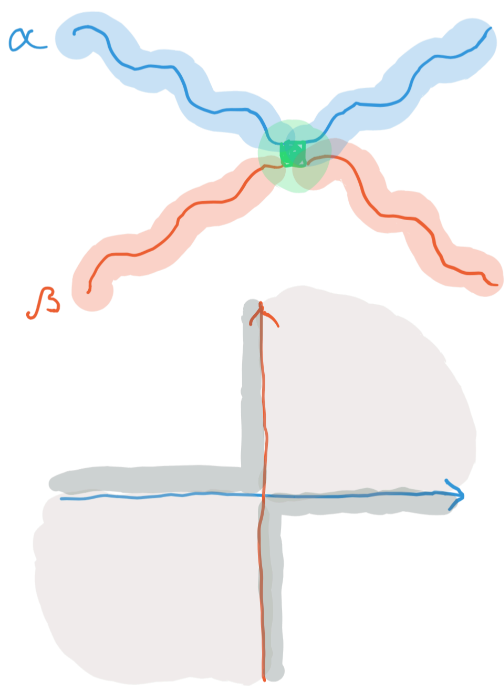

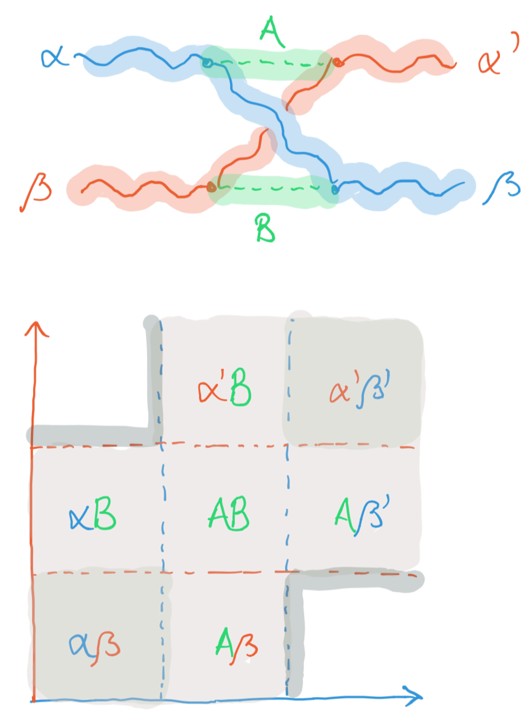

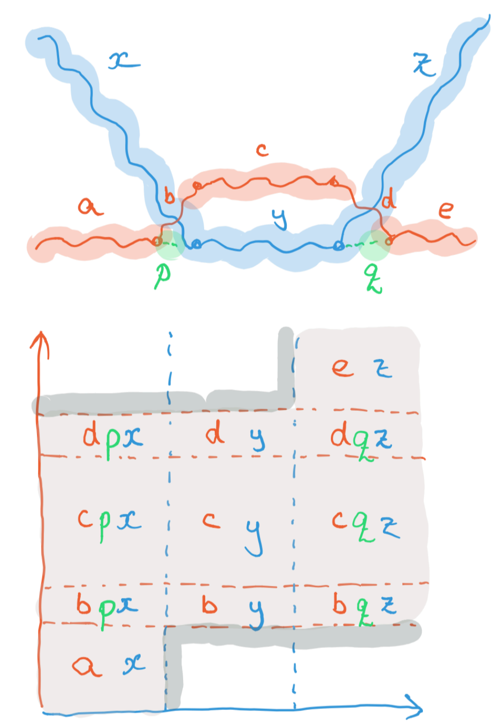

Resource Distribution

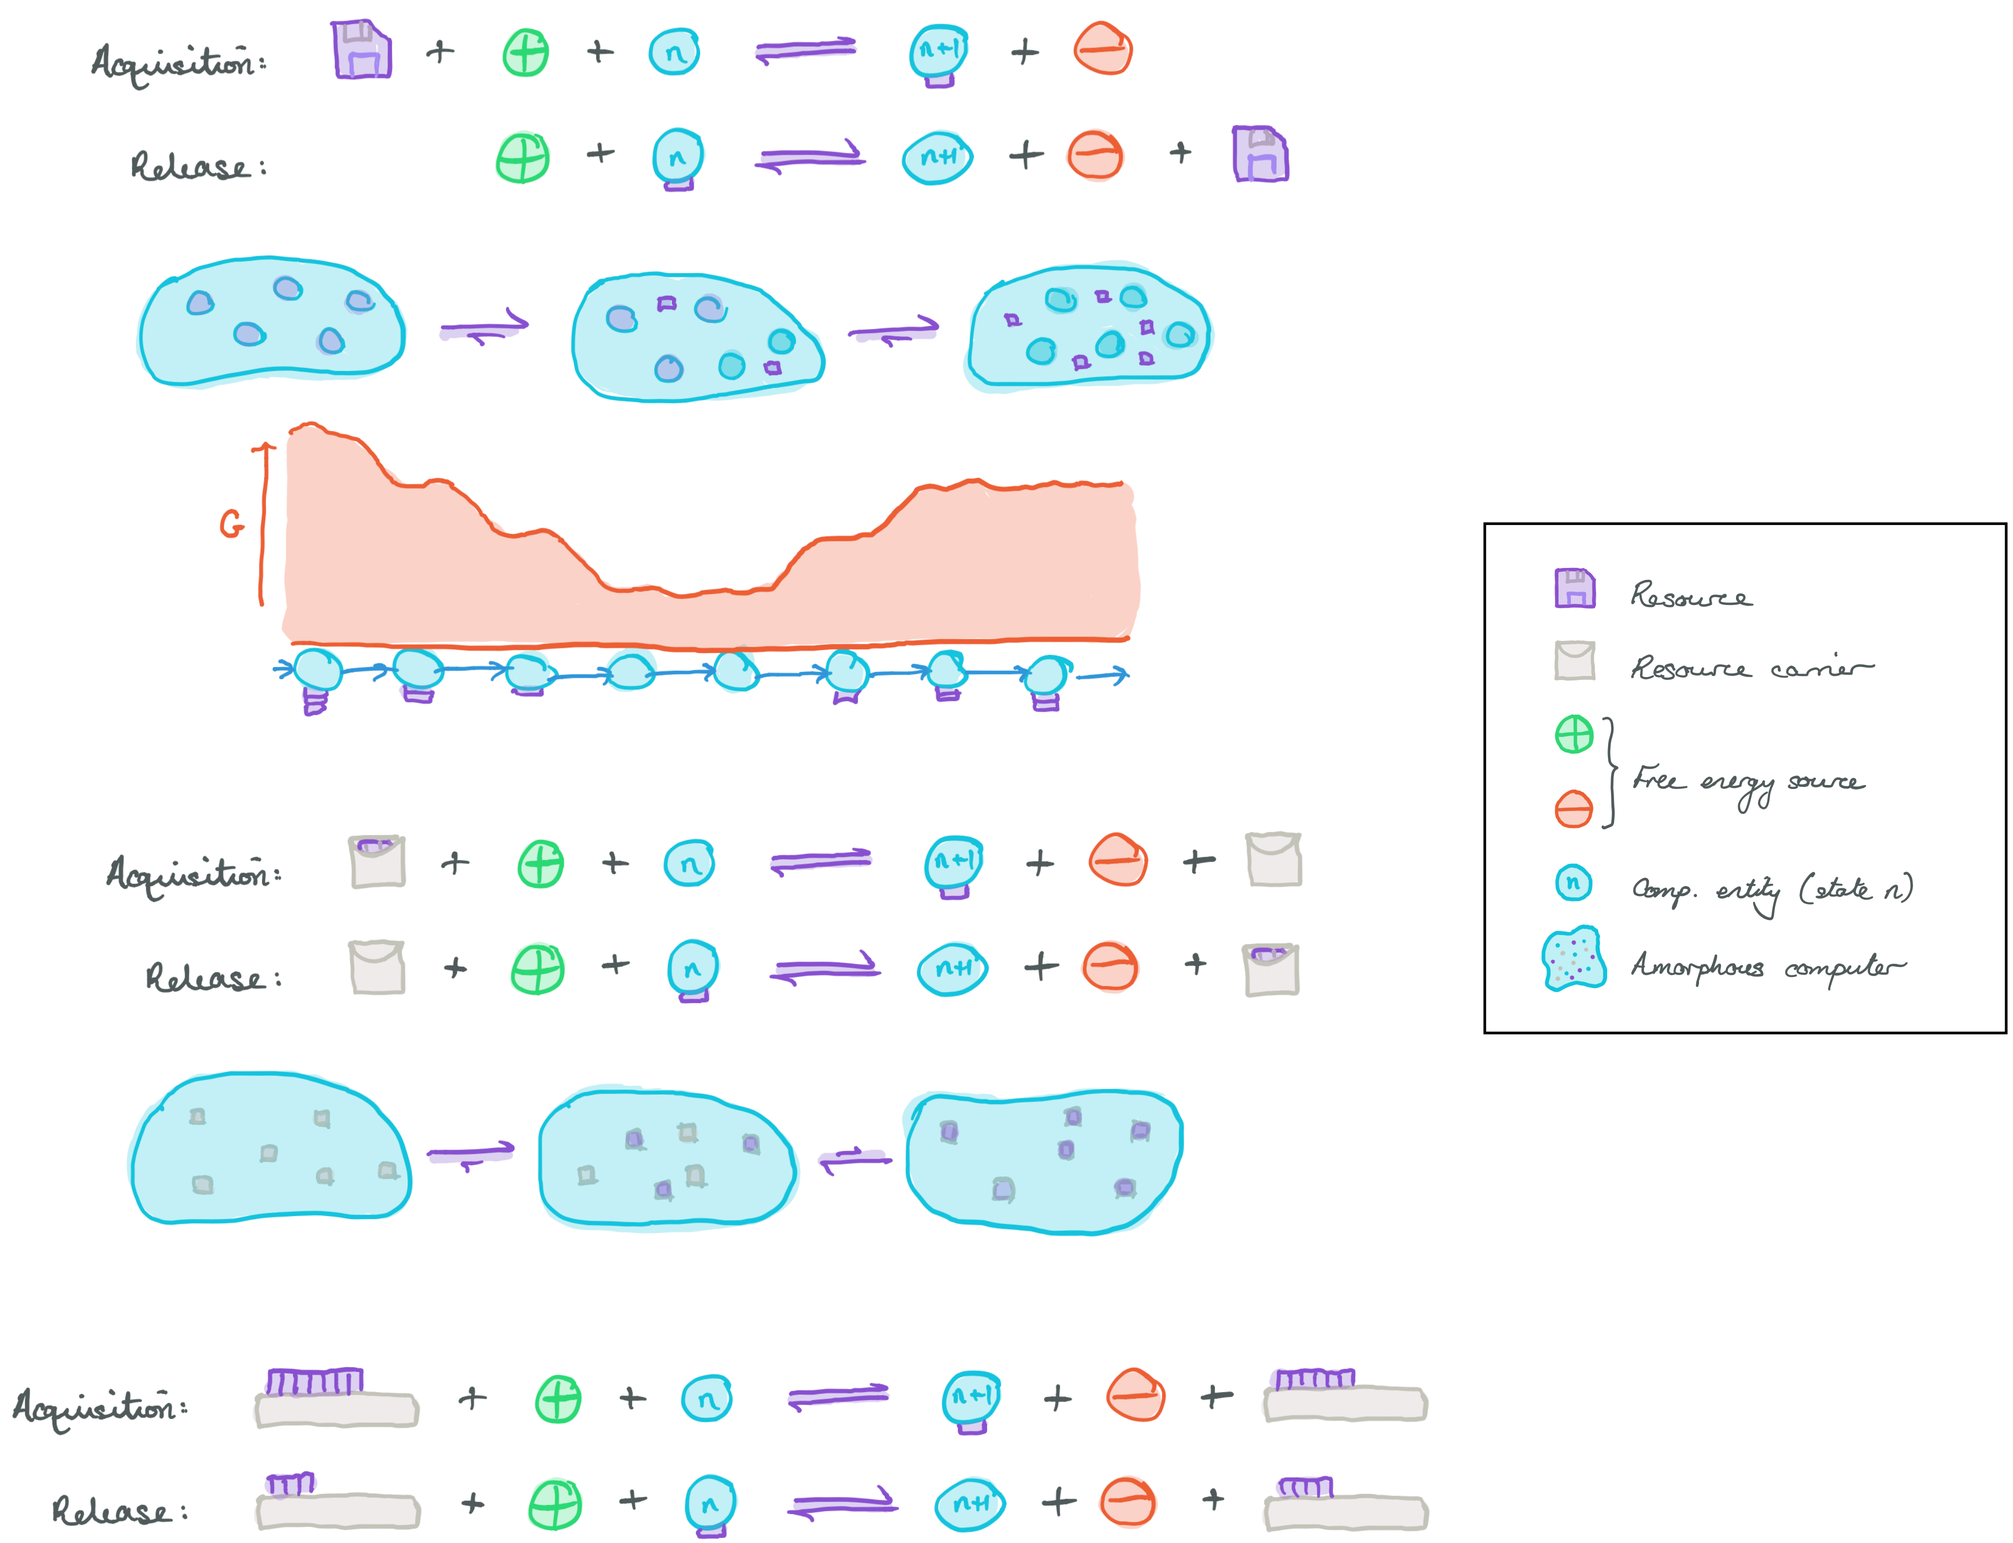

alethe & The ℵ Calculus

As well as promising lower energy costs, reversible computing is of particular interest to molecular computing because it better exploits the physics of this domain, namely microscopic reversibility. Hence, we propose a model of reversible computation, the ℵ (*aleph*) calculus, with features desirable in a reversible molecular computing context. We also introduce the associated programming language, `alethe`, meaning *not forgotten*.

Key insights

As far as we are aware, all existing reversible programming languages keep program and data separate. Implementations of such models can be very complicated, as the computers must maintain an often complicated representation of where in the program it currently is. This makes direct molecular implementation of such models especially challenging.

Our model circumvents this by being a term-rewriting system (TRS). In a TRS, there is no distinction between program and data. Instead, they are combined into a 'term', which gives a complete[1] representation of the state of computation at any point along its worldline. As the below listing shows, this representation is concise, and our reversible TRS makes it particularly easy to step forward and backward through this computational worldline.

```
-- Peano definition of natural numbers
data Z;
data S n;

-- reversible Peano addition
      + a b _ = + a b Z +;
+ (S a) b c + = + a (S b) (S c) +;
  + Z a+b a + = _ a a+b +;

-- termination conditions
! + a b   _;
! _ a a+b +;
```

```
+ 3 4 _
  ↕
+ 3 0 4 +
  ↕
+ 4 1 3 +
  ↕
+ 5 2 2 +
  ↕
+ 6 3 1 +
  ↕
+ 7 4 0 +
  ↕
_ 7 4 +
```

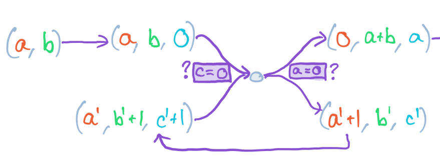

top-left: An example implementation of Peano addition written in `alethe`.

top-right: Each term encountered during the reversible addition of 4 to 3. Note that the result necessarily includes one of the addends.

bottom: A schematic of the control flow of the addition algorithm.

The choice of a TRS makes particularly good sense from a molecular implementation point of view. Terms keep the relevant bits of program next to the relevant bits of data, allowing computation to be effected by local operations and manipulation. In contrast, other models often need to synchronise activity over large domains. Their structure also is strongly suggestive of a molecular representation.

-- Some simple examples:

```
-- a simpler, recursive infix implementation of Peano addition
a `+ Z` a;
a `+ (S b)` (S c):
    a `+ b` c.

-- the Cantor pairing function π : ℕ² ↔ ℕ
-- this definition makes use of the comparison operator >
a b `Pair` n:
    ! ~Go n Z True = ~Go b (S a+b) False.
    ~Go n a+b True = ~Go n' (S a+b) p:
        n' `+ a+b` n.
        `> n' a+b` p.
    a `+ b` a+b.

-- the Factorial function, with no garbage
-- this is only a partial bijection, and will not compute 0!
(S n) `Fac` m:
    ! ~Go (S n)      1 = ~Go 1       m.
      ~Go (S (S n)) a = ~Go (S n) (S (S a')):
          a `* (S (S n))` (S (S a')).
```

-- A proof of Turing completeness:

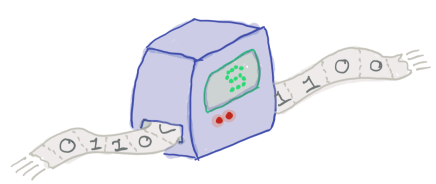

```
-- a one hole context representation of a bi-infinite tape
data Tape l x r;
data Symb x;
data Blank;

[Blank    x . xs] `Pop` Blank    [x . xs];
[(Symb x)   . xs] `Pop` (Symb x) xs;
[]                `Pop` Blank    [];

--- tape movement
x `Id` x;

(Tape l x r) `Left` (Tape l' x' r'):
  l  `Pop` x' l'.
  r' `Pop` x  r.

t `Right` t':
  t' `Left` t.

--- tape initialisation
xs `MkTape` (Tape [] Blank xs'):
  x ~ (Symb x);
  xs `Map ~` xs'.

-- a four-tape reversible Turing Machine (RTM), per Bennett [1]

--- halting states
! Start t1 t2 t3 t4;
! Stop  t1 t2 t3 t4;

--- the rule S1 [C / D /] ⟷ [B - F +] S2
  S1 (Tape l1 C r1) t2  (Tape l3 D r3) t4
= S2 (Tape l1 B r1) t2' (Tape l3 F r3) t4':

    t2 `Left`  t2'.
    t4 `Right` t4'.
```

Definition

The ℵ calculus has a very simple and concise definition. In BNF notation, this is:

$$
\begin{aligned}
\text{(pattern term)} \quad & \tau ::= \text{ATOM} \mid \text{VAR} \mid (\tau^*) \\
\text{(rule)} \quad & \rho ::= \tau^* \leftrightarrow \tau^* \\
\text{(definition)} \quad & \delta ::= \rho : \rho^* \mid \vdash \tau^*
\end{aligned}
$$

`alethe` is essentially the same except for some syntactic sugar for common programming motifs. The model is perhaps best understood by example (see the central listing).

Essentially, we define rules that map terms to other terms. Each rule has two patterns, and can thus convert a term between these two forms. Patterns need not match the term as a whole – they can match anywhere in a term. We can express more complex maps via subrules, which are the primary means of composition.

Additional constraints on this definition, as well as semantics, will be presented in the full forthcoming paper.

Select Features

- The ℵ calculus is microscopically reversible by design
  - This makes it directly compatible with the laws of physics without the need for an external source of free energy.
  - It may even make it easier to implement molecularly.
  - It encourages an economical programming style. Often one finds that garbage data can be recycled somehow or avoided altogether.
- Automatic parallelisation
  - When possible, subrules or subterms can be automatically evaluated in parallel.
  - Subrules can even be evaluated in both directions when needed.
    - This includes Bennett's algorithm [1] as a special case.
- Effects and interactions are easily accommodated, e.g.
  - Walking along a lattice
  - Actuating molecular machinery

Extensions

- If run in a thermally coupled environment, the ℵ calculus can be extended to handle non-determinism and irreversibility by allowing for ambiguous patterns. This can be done whilst preserving microscopic reversibility.
- The ℵ calculus can also be easily extended to support first class concurrency, in which terms can freely fission and fuse. The semantics and consequences of this will be discussed in the full paper.

Both of these extensions permit entropy-generating behaviour and so should be used with caution.

Future work

Though the ℵ calculus is motivated by reversible macromolecular computing systems, it is perhaps too high level for a direct implementation. We plan next on exploring simpler models more amenable to molecular implementation.