Unconventional Computing

A guide to programming molecules and turning back (computational) time

William Earley[1]

Department of Applied Maths and Theoretical Physics,
University of Cambridge

8th February 2019

[1]wje25@cam.ac.uk

Who Am I?

- I am a PhD student at Trinity Hall, based at DAMTP
- I also did my undergraduate at Trinity Hall, in Natural Sciences
- I seem to routinely find myself stretched between disciplines!
  - Throughout undergrad, I was half and half between biology and physics
  - My current research touches computer science, maths, physics and biochemistry
  - I guess this makes me a jack of all trades, though I hope to become a master of something by the end of my PhD!

Unconventional Computing

There are many types of computing:

- Von Neumann-style
- Analogue Computing
- Quantum Computing
- Artificial Neural Networks
- **Reversible Computing**
- **Molecular Computing**
- etc.

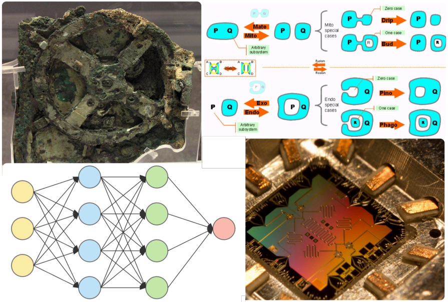

Here, I will discuss reversible and molecular computing, and show some of my work at the intersection of these two fields.

Outline

Outline

Irreversible computing

- The laws of the universe are fundamentally reversible, even quantum mechanics.
- This means that no information can ever be lost.
- Conventional computers, however, completely fail to be reversible.
- The transistors and logic gates that make up all modern computers actively discard information:

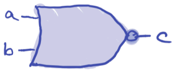

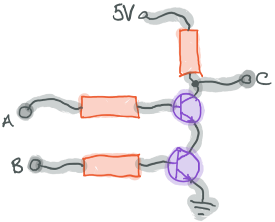

a	b	c
0	0	1
0	1	1
1	0	1
1	1	0

- Even the foundations of Computer Science rely on irreversibility...

Foundations of Computing: The Turing Machine

- Remarkably simple yet powerful.
- Consists of a tape of unlimited size, inscribed with symbols.
- A head scans these symbols and, depending on an internal state, may overwrite them and move the tape.
- By overwriting symbols and its internal state, the machine irreversibly forgets the past.
- It is famous for being one of the first models to be *computationally universal*.

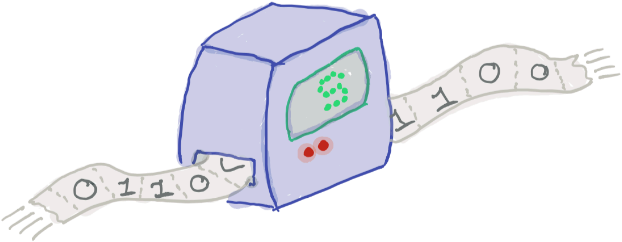

Reversible computing

- If the universe is reversible, how can computers be irreversible?
- As always, 'irreversibility' emerges from the laws of thermodynamics.
- The second law states that in any process, the entropy of the universe never decreases.

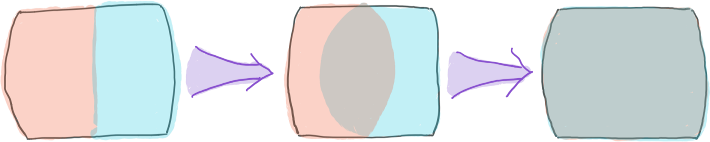

- There is more to this story, however.
- Time to see a dæmon about a box!

Maxwell's Dæmon

- Ever since the origins of thermodynamics, many have tried to find ways to circumvent it, or at least understand why certain things are disallowed.
- One famous thought experiment is Maxwell's dæmon (so named by Lord Kelvin!)
- A microscopic dæmon sits between two sides of a box, watching the particles closely.
- It then carefully opens the door in order to let fast particles through to the right side, and slow particles to the left side.
- Over time then, shouldn't we find a temperature difference, and so a decrease in entropy?

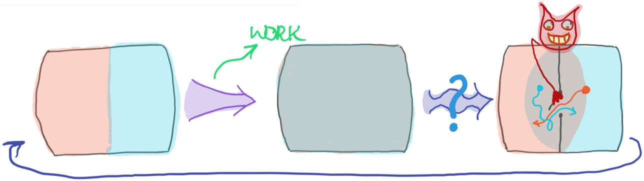

Banishing the Dæmon

- Many have tried to solve Maxwell's Dæmon, either in the positive or negative.
- Landauer's solution[2], however, was the first to satisfy most.
- He considers the simpler system to the right.
- Can any 'dæmon' X perform such a task without doing any work?
- Reversibility shows that this is impossible!
- We also see the entropy would decrease by $k_B \log 2$.

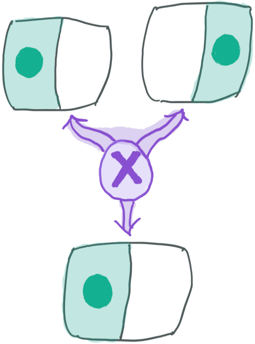

[2] Landauer 1961.

Landauer's Principle

- What is the intuitive reason for Maxwell's dæmon not existing?
- Landauer showed that it had to do with the connection between information and entropy
- If a quantity of information I is 'erased', then the entropy (volume of phase space - $S = k_B \log W$) has decreased by $k_B I$, violating Liouville's theorem!
- We can therefore only move information. If we want to forget it, we need to dump it somewhere.
- The environment is always a good dumping ground...
- Landauer's principle states that forgetting information I requires dissipating at least $k_B T I$ in heat.

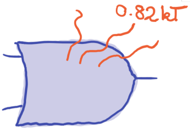

$$I = -\sum p_i \log p_i$$

$$\Delta q_h \geq k_B T \Delta I$$

Does Logic require Irreversibility?

- Landauer argued yes, on the basis that a reversible computer would get cramped.
- If so, then the efficiency of computers has a limit (though we're currently 8 orders of magnitude above Landauer's limit!)
- Charles Bennett[3], often regarded as the founder of reversible computing, showed that reversible computing was both possible and practical.
- To do so, he came up with a *reversible* Turing Machine.
- He also showed how to use it to simulate any irreversible program efficiently.

[3] Bennett 1973.

Bennett's Algorithms

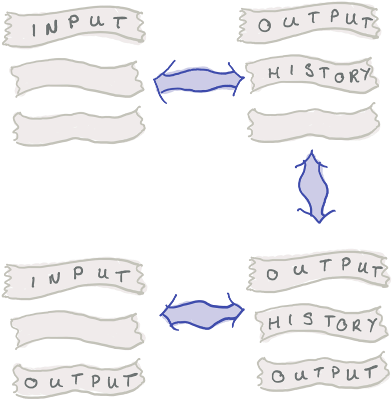

TABLE 2
Reversible simulation in time $O(T^{\log 3/\log 2})$ and space $O(S \cdot \log T)$.

Stage	Action	Checkpoints in storage (0 = initial ID, checkpoint $j = (jm)$th step ID)								
0	Start	0								
1	Do segment 1	0	1							
2	Do segment 2	0	1	2						
3	Undo segment 1	0		2						
4	Do segment 3	0		2	3					
5	Do segment 4	0		2	3	4				
6	Undo segment 3	0		2		4				
7	Do segment 1	0	1	2		4				
8	Undo segment 2	0	1			4				
9	Undo segment 1	0				4				
10	Do segment 5	0				4	5			
11	Do segment 6	0				4	5	6		
12	Undo segment 5	0				4		6		
13	Do segment 7	0				4		6	7	
14	Do segment 8	0				4		6	7	8
15	Undo segment 7	0				4		6		8
16	Do segment 5	0				4	5	6		8
17	Undo segment 6	0				4	5			8
18	Undo segment 5	0				4				8
19	Do segment 1	0	1			4				8
20	Do segment 2	0	1	2		4				8
21	Undo segment 1	0		2		4				8
22	Do segment 3	0		2	3	4				8
23	Undo segment 4	0		2	3					8
24	Undo segment 3	0		2						8
25	Do segment 1	0	1	2						8
26	Undo segment 2	0	1							8
27	Undo segment 1	0								8

[3] Bennett 1989.

How can we build one?

Fredkin and Toffoli[5] gave some of the first examples

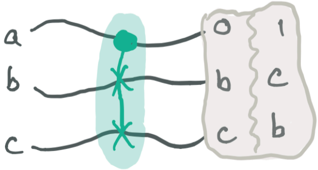

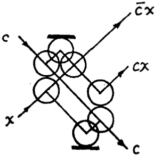

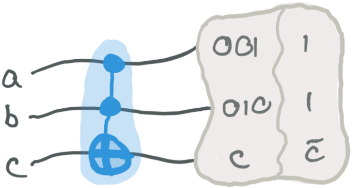

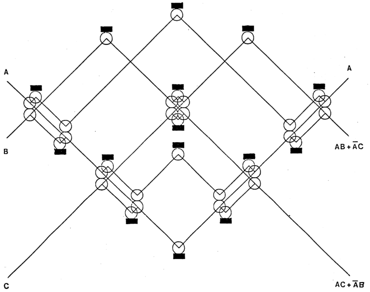

[5]Fredkin and Toffoli 1981.

Can we build one?

- The billiard ball only requires classical mechanics.
- An ambitious master's project by Ressler[6] even managed to design a full fledged CPU, complete with arithmetic unit and memory stores using the formalism!
- In principle, such a computer could compute without any dissipation. In practice, though....[7]

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.

[7] Ressler 1981.

[7] Bennett 1982.

Quantum Computing: Reversible?

- Quantum mechanics is time symmetric as well.
- What about wavefunction collapse?
- Quantum computers cannot be allowed to mix with their environment at all (tricky!).
- This means that all quantum computers must be reversible!

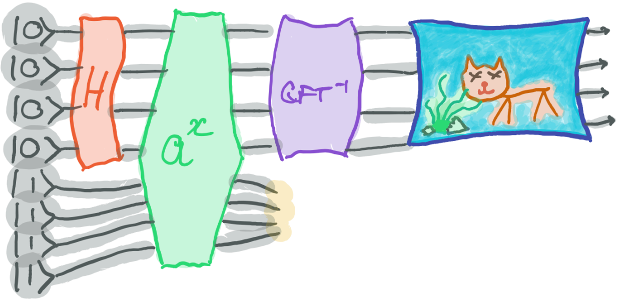

Outline

Early days: Adleman[8]

- One of the earliest
- Solve for Hamiltonian paths with DNA!
- Experimentally verified!

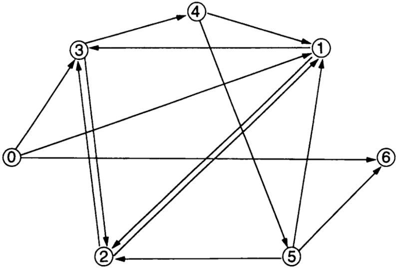

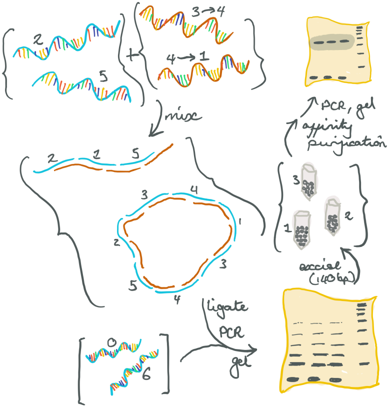

[7] Adleman 1994.

Early days: Rothemund's DNA Turing Machine[9]

- DNA and restriction enzyme system
- Quite complicated!

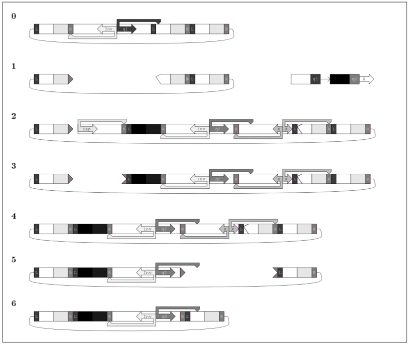

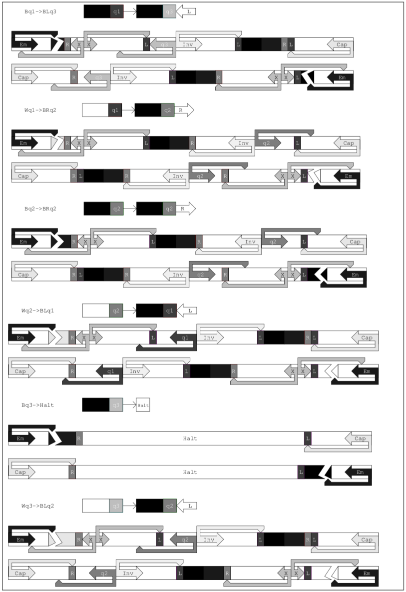

[8]Rothemund 1995.

Modern Approaches

- The previous two examples demonstrate that molecular computing is possible.
- Unfortunately they're not very practical!
- Luckily, autonomous molecular computation is possible.
- Over the last 20 years or so, much work has been done on dynamic DNA nanotechnology.
- We will look at the two most popular systems for molecular computation that have emerged:
  - DNA Strand Displacement (DSD),
  - The Tile Assemble Model (TAM).

DNA Strand Displacement

- DSD[10] has emerged as a near standard after over a decade of work by many pioneers.
- It is built from the primitive operation shown below, of ‘toehold’-mediated strand exchange.

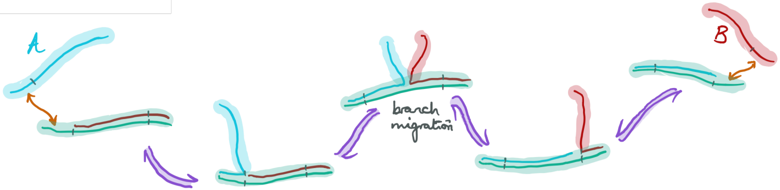

[10] Seelig et al. 2006.

Chemical Reaction Networks

- Is DSD expressive enough to compute?
- It turns out it can implement any CRN.
- A CRN is an abstraction of chemical reactions.
- It is defined by a set of species, and a set of reactions between those species.
- E.g. $\mathsf{A} + 2\,\mathsf{B} + \mathsf{C} \longrightarrow 3\,\mathsf{D} + \mathsf{E}$, $\mathsf{A} \longrightarrow 2\,\mathsf{A}$, ...
- Why are CRNs useful? Well, Soloveichik[11] showed that any Register Machine can be simulated by a CRN...

[11] Soloveichik, Seelig and Winfree 2010.

Simulating CRNs with DSD

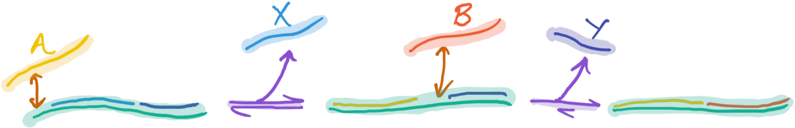

- By cascading variants of these, we can implement any CRN reaction $\sum_i \alpha_i X_i \longrightarrow \sum_i \beta_i Y_i$, perhaps with some additional fuel and waste strands.
- This is not the only construct that can be used in DSD. VisualDSD[12] is a tool to compile any CRN into a DSD scheme using any of the various approaches.

Reaction	Inputs	Outputs	Fuel	Waste
AND: A + B $\longrightarrow$ Y	A, B	Y		X
FANOUT: A $\longrightarrow$ X + Y	A	X, Y	B	

[12] Lakin et al. 2011.

Calculating Square Roots with DSD![13]

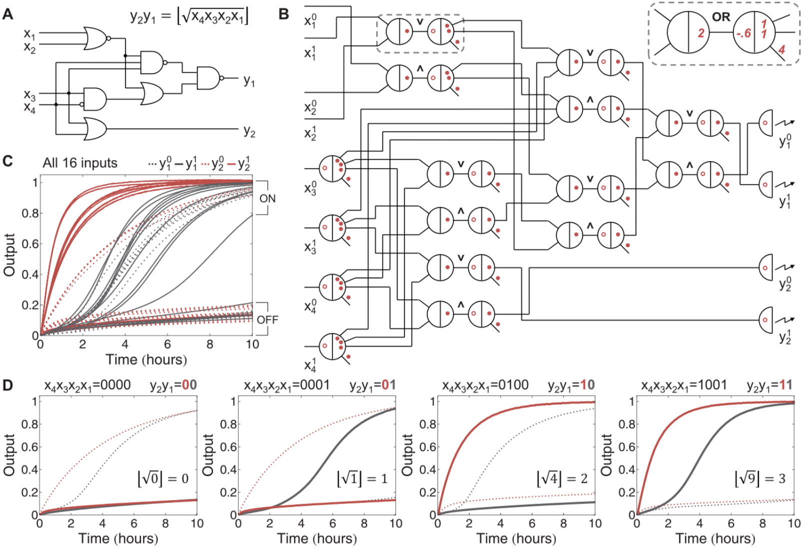

[12]Qian and Winfree 2011.

Neural Networks with DSD![14]

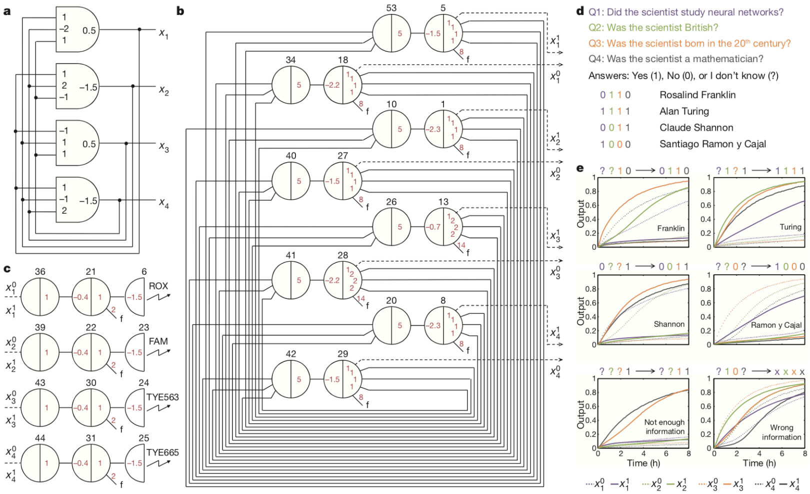

---

[13] Qian, Winfree and Bruck 2011.

The Tile Assembly Model[15]

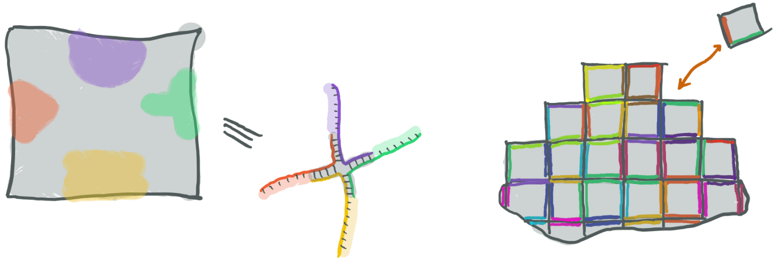

- TAM was developed by Erik Winfree for his PhD Thesis.
- Abstractly, it consists of a set of square tiles with 'coloured' edges. These are implemented in DNA as above.
- Like coloured edges can associate via their sticky ends.

[14] Winfree 1998.

Wang Tiles[17]

- Is TAM sufficiently expressive to compute?
- Tiling models have been studied for millennia.
- TAM turns out to be isomorphic to Wang tilings.
- Wang asked whether all tilesets would give a periodic pattern.
- It turns out that the answer to this is no, because a tileset can be constructed to simulate a Turing Machine![16]

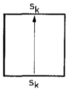

Fig. 12. Alphabet tile

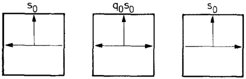

Fig. 15. Starting tiles for blank tape

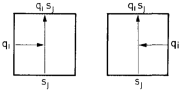

Fig. 13. Merging tiles

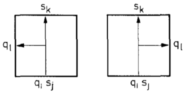

Fig. 14. Action tiles

---

[15] Wang 1961.

[16] Robinson 1971.

Comparison

DSD

- \+ Impressive feats
- − Not very composable
- − Doesn't parallelise well
- − Very error prone

TAM

- \+ Compact algorithms
- \+ Localised $\implies$ parallelism
- \+ Basic 'error correction'
- − Keeps computation history
- − Not very dynamic

- What does the future hold?
- I am seeking a new model to combine the strengths of DSD and TAM, but finding systems as robust as them is tricky!

Outline

How far can Moore's law go?

- Historically, processor speeds have tended to double around every 2 years.
- Though they have stagnated recently, is there any upper limit?
- In 1962, Bremermann[18] used the uncertainty principle to give the first estimate, $\nu \lesssim E/h$.
- Margolus and Levitin[19] then refined this, $\nu \leq 2E/h$.
- Numerically, this gives $\nu \leq 2.71 \times 10^{50}\,\mathrm{kg}^{-1}\,\mathrm{s}^{-1}$.
- Lloyd[20] took this to its logical extreme, analysing the properties of the 'ultimate laptop' - a kilogram of matter operating at this limit, compressed into a black hole, and performing 10^{32} operations on 10^{16} bits in $10^{-19}\,\mathrm{s}$ at an apparent temperature of $10^{9}\,\mathrm{K}$!

[18] Bremermann 1962.
[19] Margolus and Levitin 1998.
[20] Lloyd 2000.

How big can we go?

- If we can't make our computers faster, can we make them bigger?
- Authors such as Sandberg[21] describe concepts of 'Jupiter brains' - immense spheres filled with computational matter.
- Suppose this matter is the Intel Xeon E5-2699 v4 (catchy)...
  - Est. stats: $\nu = 1.6 \times 10^{14}\,\mathrm{bit\,s^{-1}}$, $T = 300\,\mathrm{K}$, $P = 145\,\mathrm{W}$, $m = 0.1\,\mathrm{kg}$, $V = 5\,\mathrm{cm}^3$.
  - Note - Landauer predicts $P = 0.6\,\mu\mathrm{W}$, a 2×10^8 difference!
- Our 'Jupiter' would contain 3×10^{29} of them, and radiate $4 \times 10^{31}\,\mathrm{W}$, or $7 \times 10^{14}\,\mathrm{W\,m^{-2}}$.
- The sun only outputs $4 \times 10^{26}\,\mathrm{W}$ or $7 \times 10^{7}\,\mathrm{W\,m^{-2}}$...
- Our Jupiter would then have a surface temperature of $3 \times 10^{5}\,\mathrm{K}$, let alone its core temperature!
- The Landauer limit is not much better, $T \sim 3000\,\mathrm{K}$.

---

[21] Sandberg 1999.

Geometry of computing

- Volumetric irreversible computing is unsustainable, unless extreme temperatures can be tolerated.
- Heat can only be removed from its surface, which only scales with r^2.
  - This is part of the reason why CPUs aren't stacked.
  - A large irreversible computer must be shell-like...
- What about reversible computers?
  - Reversible computers can in principle compute without dissipation.
  - In practice though, some energy is needed to keep things running smoothly.
  - Doesn't this imply the same scaling?

Can we do better?

- Can't run a reversible computer without dissipation, is there a lower limit?
- Bennett[22] was perhaps the first to point out that a reversible computer could be run close to thermodynamic equilibrium.
- This would not work for ballistic computers, but is appropriate for (bio)chemical computers.

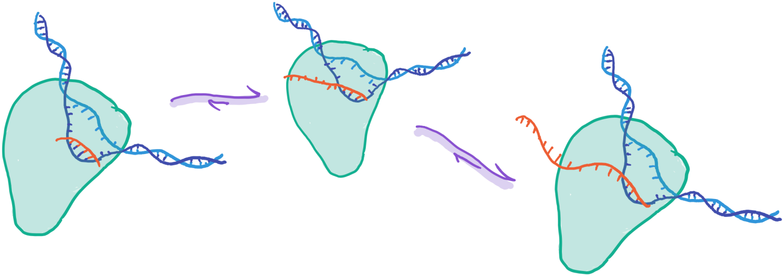

[22] Bennett 1973.

Near-Equilibrium Computation

- A volumetric computer has to divide its energy ($\propto A$) throughout its volume.
- So larger computers run arbitrarily close to equilibrium.
- Is this useful!?
- Let's build a model...

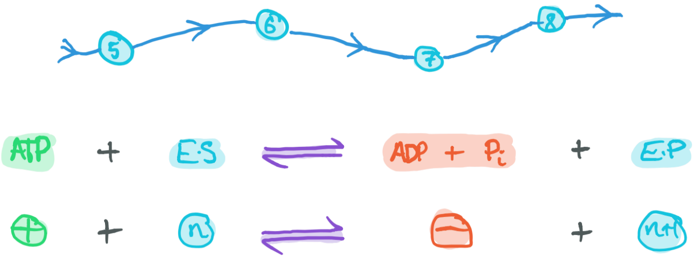

Maintaining a Bias

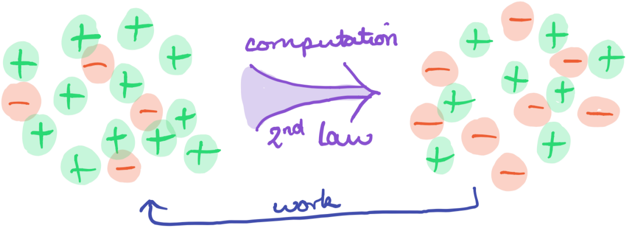

- We start off with a net bias, but this dissipates over time...
- We will need to do work to maintain the bias! How much work?

Maintaining a Bias

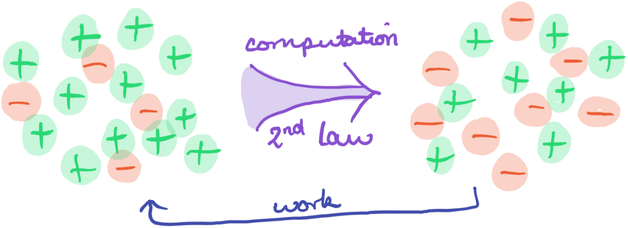

- Appealing to information theory, the information content/entropy of each token has increased...

$$I = -p \log p - q \log q \qquad \delta I = -\delta b \, \text{arctanh}\, b + \mathcal{O}\left(\delta b^2\right)$$

- We need to 'reset' each token back to its original state!

$$\delta E \geq k_B T \delta I \qquad P = \dot{E} \geq -k_B T \dot{b} \, \text{arctanh}\, b$$

Maintaining a Bias

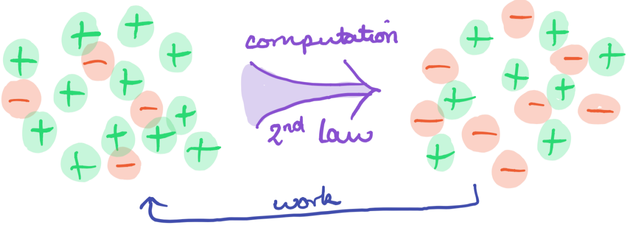

- What is $\dot{b}$? Time to do some IA chemistry!

$$\partial_t[+] = k[C][-] - k[C][+] = -kb[C][\pm] \implies \dot{b} = -2kb[C]$$

$$P \geq 2kN_C N_\pm k_B b \operatorname{arctanh} b \approx \alpha N_\pm b^2$$

- Now, $P \sim A$ but $N_\pm \sim V$, so $b \sim 1/\sqrt{\ell}$ or $R \sim V^{5/6}$

What on Jupiter-Brain does that mean?

- $b \sim 1/\sqrt{\ell}$ shows that each individual computer is getting slower...
- But $R \sim V^{5/6}$ shows that the total computation rate is getting faster, faster than expected even!
- An irreversible body would only have $R \sim V^{4/6}$.
- So we're halfway between irreversible and ballistic!

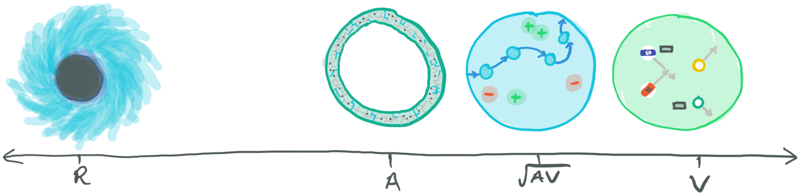

The numbers

- ‘Unambitious’ Biocomputer
  - Power dissipation, $500\,\mathrm{W\,m^{-2}}$
  - Raw speed, $1\,\mathrm{bit\,s^{-1}}$ per $5\,\mathrm{nm^3}$ unit
  - 1 metre^3 computer, $10^{25}\,\mathrm{bit\,s^{-1}}$
- ‘Typical’ ARM Chip
  - Power dissipation, $2 \times 10^5\,\mathrm{W\,m^{-2}}$ ($1\,\mathrm{W}$ per $5.2\,\mathrm{mm^2}$)
  - (Landauer overhead 10^8)
  - Speed, $2.2 \times 10^{12}\,\mathrm{bit\,s^{-1}}$
  - 1 metre^3 computer, $10^{18}\,\mathrm{bit\,s^{-1}}$
- Not directly comparable
  - Parallel vs Serial
  - In raw terms though, reversible wins!

Outline

Cooperative Chemical Computers

- The previous section covered a body of isolated computers.
- In practice, we will want them to interact...
  - Resource sharing
  - Communication
- How well do reversible computers fare?

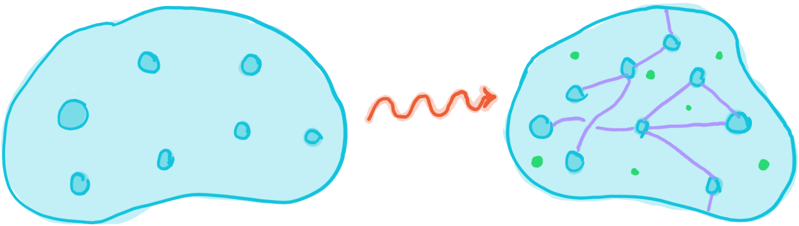

Outline

Resource Sharing

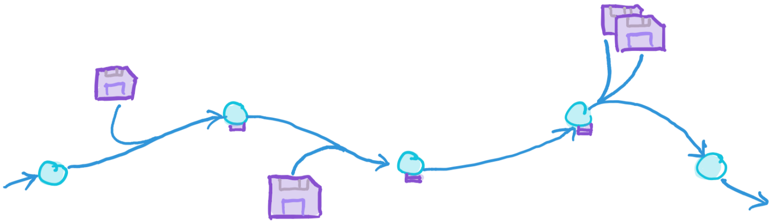

- Potentially unbounded memory is essential to computers
- Small chemical computers are clearly limited
- Can we engineer them to reversibly acquire and release additional memory when needed?

Resource Sharing: Attempt #1

- As a naïve attempt, let's scatter resources freely throughout the medium... Compare with tRNAs, nucleotides, etc.
- The bias is low, so each computer doesn't have much energy to spare. We should aim for no energy difference then between the two states.

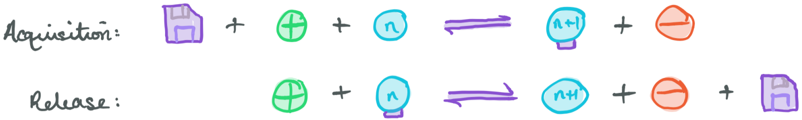

Resource Sharing: Attempt #1

- Even though there's no energy difference, there's still a free energy difference!
- This leads to an unavoidable entropic force.
- Computers will get stuck in local free energy minima.

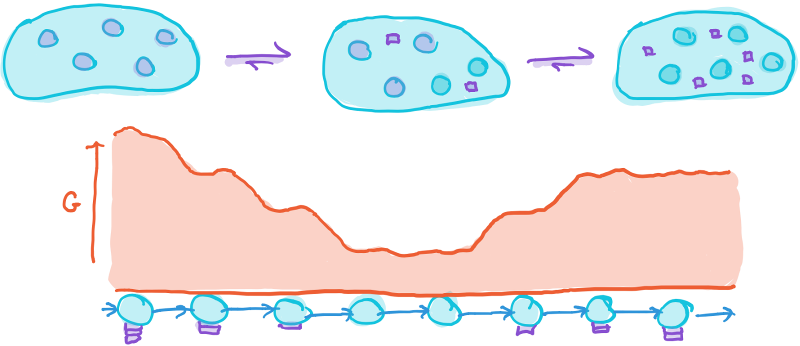

Resource Sharing: Attempt #2

- Ok, so attempt #1 was subject to a chemical potential...
- What if we kept particle number the same?

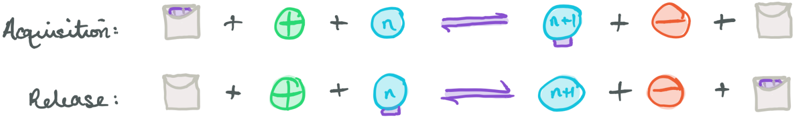

Resource Sharing: Attempt #2

- Unfortunately, there's still an entropic driving force!
- Intuitively because of the varying difficulty in acquiring and releasing resources...

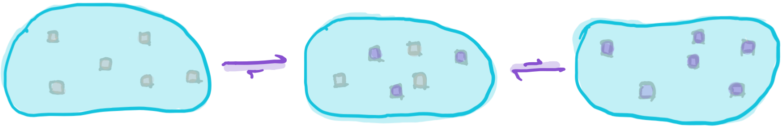

Resource Sharing: Attempt #3

- Digging our hole deeper, let's try stacking resources together...
- Now the difficulty for acquisition changes more slowly...
- To work out how slowly, we need to do some calculations...

Resource Sharing: Attempt #3

- This is a non-equilibrium system, but we can assume the resource carriers are in a quasi-steady state. At any point in time, there will be an average $\bar{n}$ resources per carrier (depending on the current demand for resources).
- We can then use the principle of maximum entropy to find the distribution of resources, getting a similar distribution to that of energy in an ideal gas...

$$\text{Pr}\{n\} = Ae^{-\beta n} = \left(\frac{1}{1+\bar{n}}\right)\left(\frac{\bar{n}}{1+\bar{n}}\right)^n$$

$$\mathbb{P}_0 = \frac{1}{1+\bar{n}} \qquad \mathbb{P}_+ = \frac{\bar{n}}{1+\bar{n}}$$

Resource Sharing: Attempt #3

- We can then find the rate equation for the resource reactions to find the *effective bias* of our resourceful reversible computers.
- The release reaction always has positive bias,
  $$b_{\mathrm{rel.}} = b\mathbb{P}_+ + \tfrac{1}{2}(1+b)\mathbb{P}_0$$
  The acquisition reaction is not so lucky,
  $$b_{\mathrm{acq.}} = b\mathbb{P}_+ - \tfrac{1}{2}(1-b)\mathbb{P}_0$$
- We find that the acquisition reactions are only processive for $\bar{n} \gtrsim \sqrt{\ell/\ell_0}$.
- So we need to stuff our system with at least $(V/V_0)^{7/6}$ resources!

Resource Sharing: Attempt #4

- Is all lost? The previous attempts show how futile fighting entropy can be!
- It turns out there *is* a way to evade entropy here though...
- How? We implement resource sharing on top of the computers!

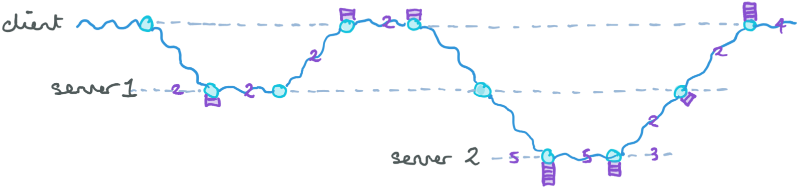

Outline

Communication

- Communication is perhaps more important than resource sharing.
- The first realisation is that free floating computers cannot communicate effectively unless actively propelled.
- Chemotaxis shows one way such problems could be solved, but it is actively dissipative!
- Therefore we must introduce some fixed lattice, not too unprecedented!

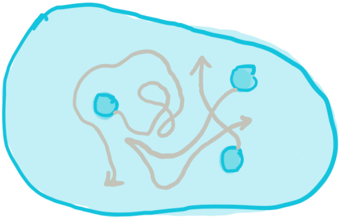

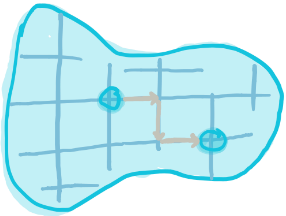

Lattice Communication

- Does using a lattice solve the problem? Unfortunately no...
- I am still actively researching this, but the crux is that most communication scenarios intrinsically rely on a decrease in entropy.
- The resource server example is a rare isentropic exception.
- It looks like the cost to communication is an unavoidable time penalty...

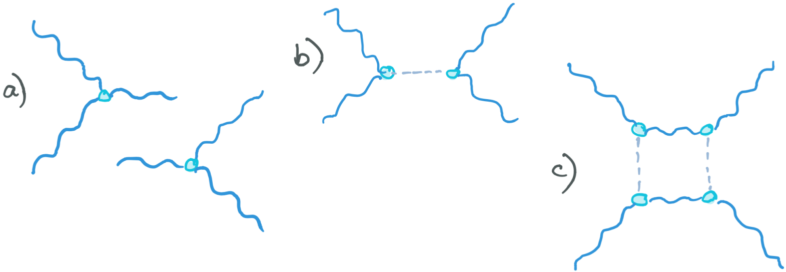

Outline

Programming a Reversible Computer

- As this is a NatSci presentation, I won't go too much into programming!
- Programming reversibly not *too* much different from normal programming.
- Need to be more careful about manipulation of information and merging control flow...
- Time for a live demo?

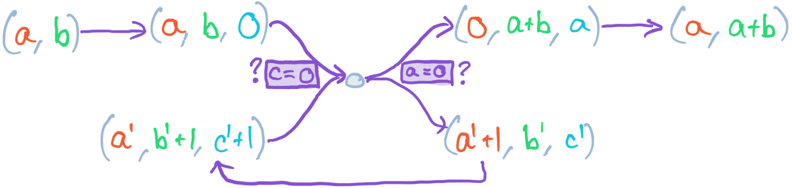

Thanks!

UNIVERSITY OF CAMBRIDGE

Department of Applied Mathematics and Theoretical Physics (DAMTP)

MICKLEM LAB