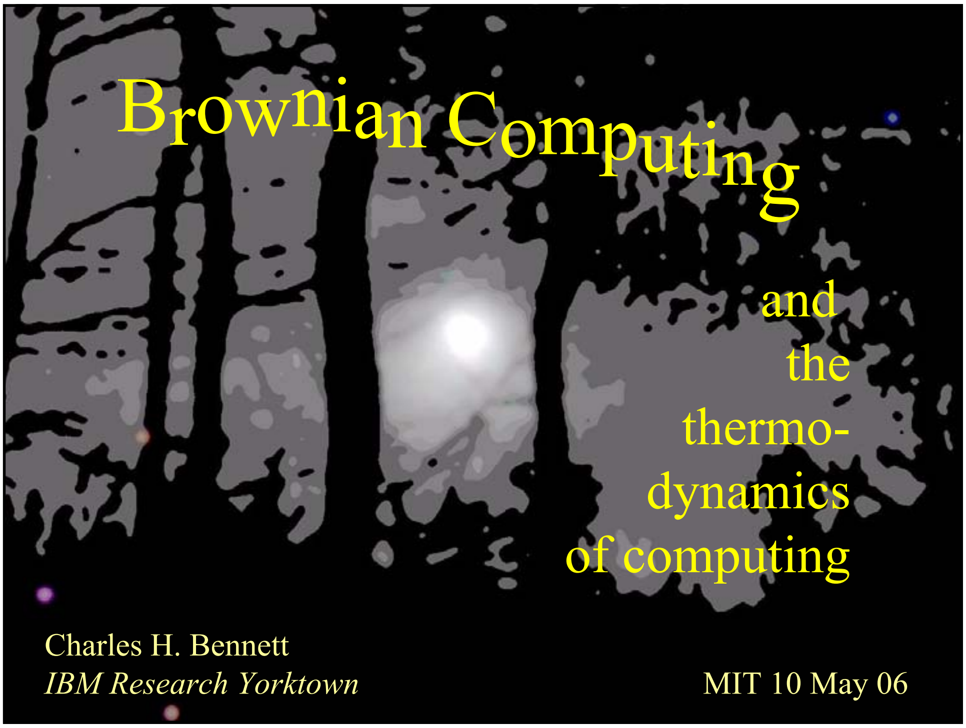


## 1981 Conference on Physics of Information (MIT Endicott House)

1 Freeman Dyson
2 Gregory Chaitin
3 James Crutchfield
4 Norman Packard
5 Panos Ligomenides
6 Jerome Rothstein
7 _ Hewitt?
8 Norman Hardy
9 Edward Fredkin
10 Tom Toffoli
11 Rolf Landuaer

12 J. Wallmark
13 Frederick Kantor
14 David Leinweber
15 Konrad Zuse
16 Bernard Zeigler
17 Carl Adam Petri
18 Anatol Holt
19 Roland Vollmar
20 Hans Bremerman
21 Donald Greenspan
22 Markus Buettiker
23 Otto Floberth

24 Robert Lewis
25 Robert Suaya
26 _ Kugell
27 Bill Gosper
28 Lutz Priese
39 Madhu Gupta
30 Paul Benioff
31 Hans Moravec
32 Ian Richards
33 Marian Pour-El
34 Danny Hillis
35 Arthur Burks

36 John Cocke
37 George Michael
38 Richard Feynman
39 Laurie Lingham
40 _ Thiagarajan
41 ?
42 Gerard Vichniac
43 Leonid Levin
44 Lev Levitin
45 Peter Gacs
46 Dan Greenberger

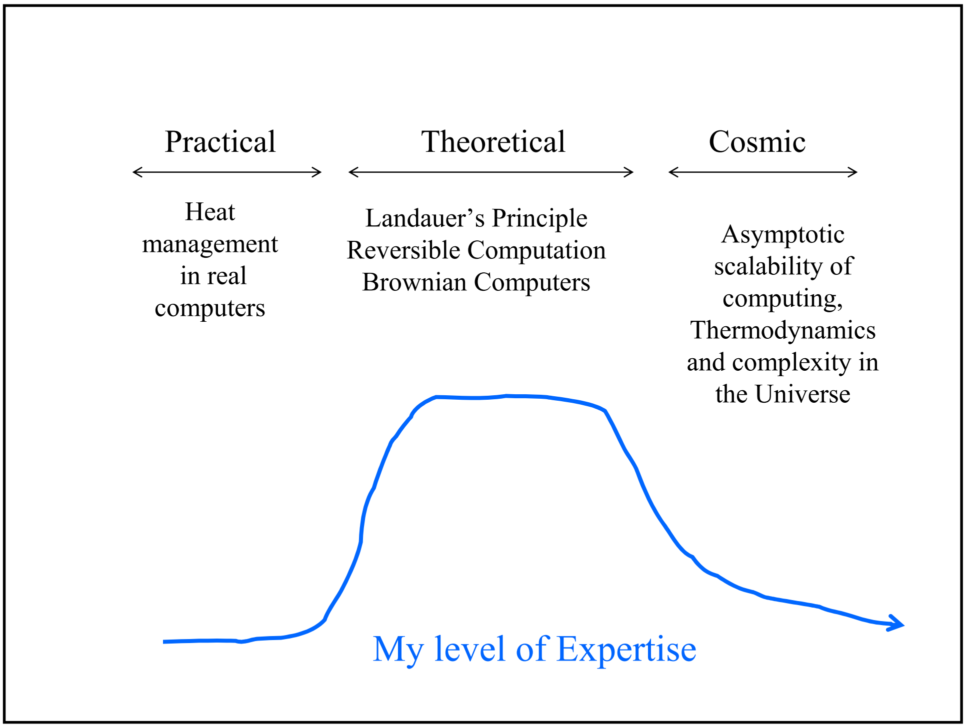


Heat generation is a serious problem in today's computers, limiting packing density and therefore performance. Combat it by:

- Making gates less dissipative, even if slower, can sometimes increase performance FLOPS/watt, while reduced clock speed is offset by increased parallelism (e.g. BlueGene/L)

- Dynamic Power Management—switching off clock where not needed or to let a hot region cool down.

- Resonant clock to reduce ½ $CV^2$ losses from non-adiabatic switching (see Michael Frank's talk)

- Thicker gates to reduce gate leakage current

- More conductive materials to reduce $I^2R$ resistive losses

# Dynamic Power Management in IBM Power5 (GR) microprocessor,

- A random test generator and exerciser program is run in Single Thread (ST) and Simultaneous Multiple Thread (SMT) mode with and without Dynamic Power Management (DPM).

- IR images clearly indicate significant reduction in power consumption and temperature rise due to DPM.

*see demo this afternoon by Maurice McGlashan-Powell*

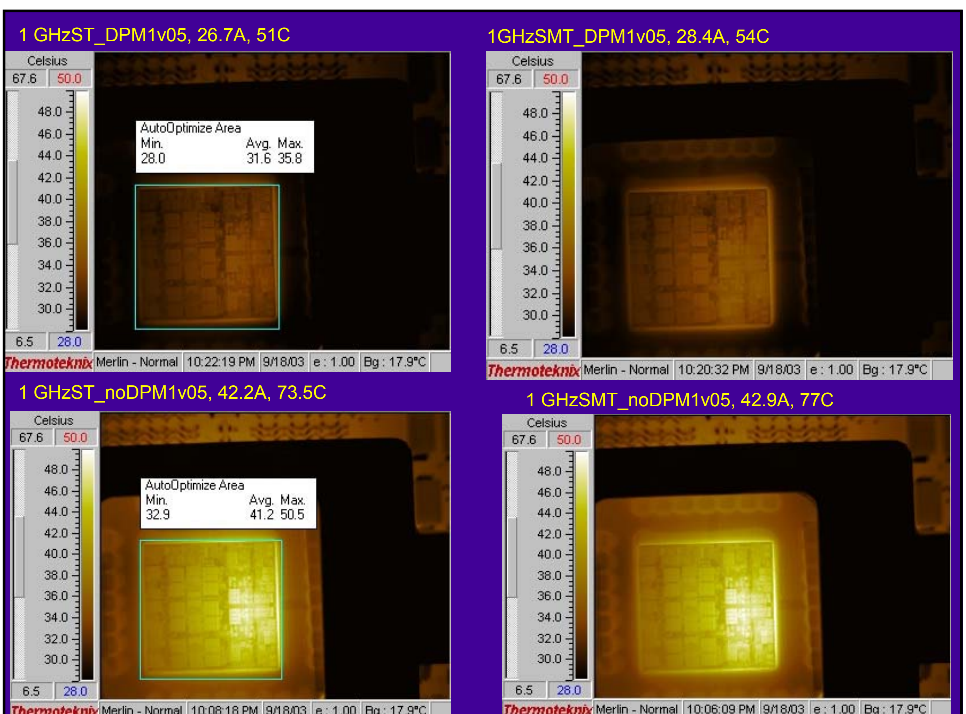

**"Information is Physical"** Rolf Landauer

**"It from bit"** John Archibald Wheeler

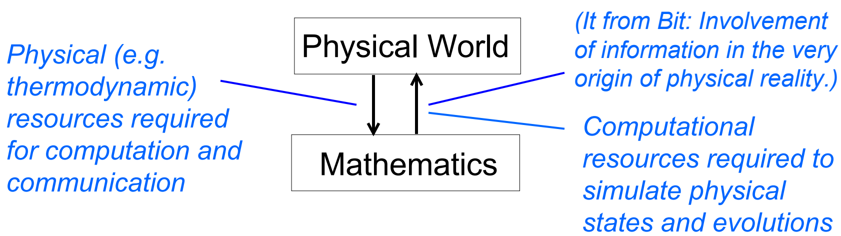


When Turing, Shannon, von Neumann and their contemporaries formalized the notions of information and computation, they left out notions of reversibility and quantum superposition

reversibility => thermodynamics of computation

superposition => quantum information/computation theory.

Conventional computer logic uses irreversible gates, eg NAND, but these can be simulated by reversible gates. Toffoli gate is universal.

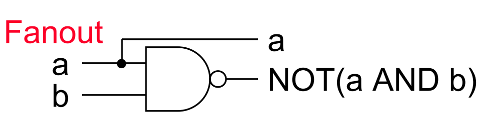


NAND gate

*no inverse*

*Reversible logic was used to show that computation is thermodynamically reversible in principle. Later needed for Quantum Computation*

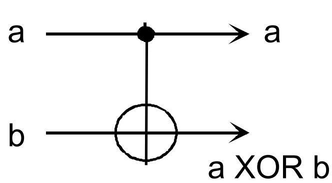


XOR gate

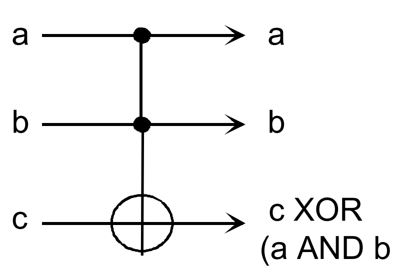


Toffoli gate

*self-inverse*

## Thermodynamics of Computation

• Landauer's Principle: each erasure of a bit, or other logical 2:1 mapping of the state of a physical computer, increases the entropy of its environment by k log 2.

• Reversible computers, which by their hardware and programming avoid these logically irreversible operations, can in principle operate with arbitrarily little energy dissipation per step.

## *Avatars of the Second Law of Thermodynamics*

No physical process has as its sole result is the conversion of heat into work.

It is impossible to extract work from a gas at constant volume if all parts are initially at the same temperature and pressure.

It is impossible to see anything inside a uniformly hot furnace by the light of its own glow.

No physical process has as its sole result the erasure of information.

Looking inside a pottery kiln

by its own glow

by external light

Ordinary irreversible computation can be viewed as an approximation or idealization, often quite justified, in which one considers only the evolution of the computational degrees of freedom and thus neglects the cost of exporting entropy to the environment.

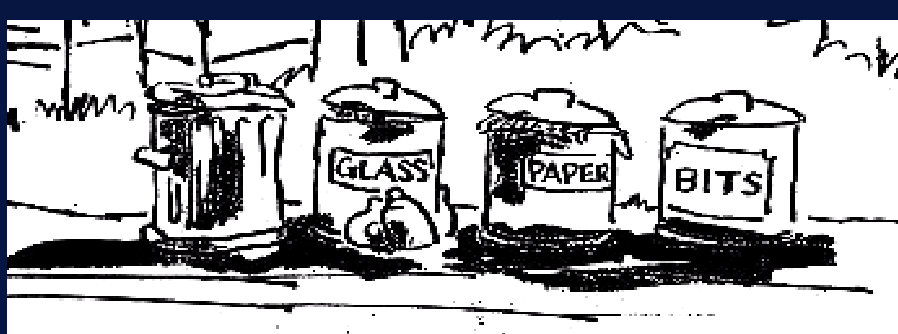

*Why study reversible classical computing, when Landauer erasure cost is negligible compared to other sources of dissipation in today's computers?*

- Practice for quantum computing
- Improving the thermodynamic efficiency of computing at the practical ½ $CV^2$ level (rather than the kT level)
- Understanding ultimate limits and scaling of computation and, by extension, self-organization

Classification of Computers from thermodynamic viewpoint

A. Irreversible (eg. PC, Mac…)

B. Reversible

1. Ballistic (e.g. Billiard ball model) dynamical trajectory isomorphic to desired computation

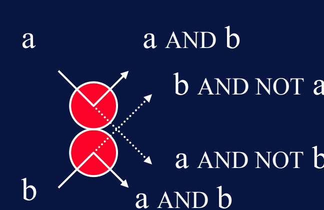


2. Brownian (e.g. RNA polymerase) random walk in a low-energy labyrinth in configuration space, isomorphic to desired computation

3. Intermediate, like walk on a 1d lattice with mean free path >1 (e.g. Feynman's quantum computer)

The chaotic world of Brownian motion, illustrated by a molecular dynamics movie of a synthetic lipid bilayer (middle) in water (left and right)

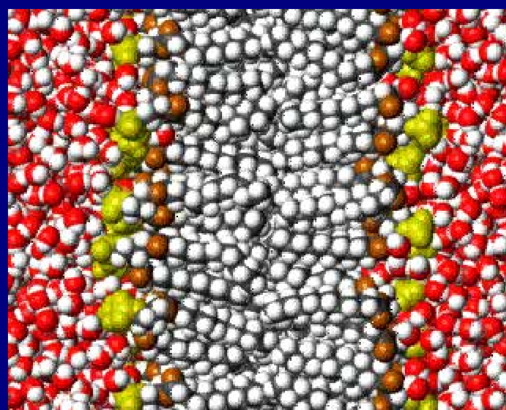

dilauryl phosphatidyl ethanolamine in water

http://www.pc.chemie.tu-darmstadt.de/research/molcad/movie.shtml

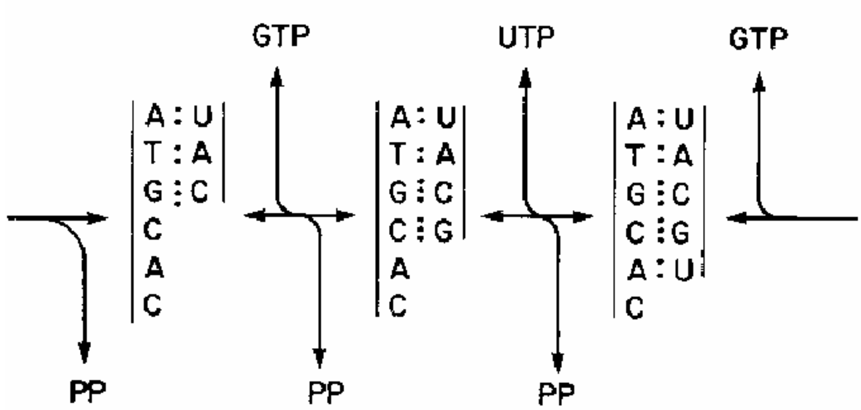


. RNA polymerase reaction viewed as a one-dimensional random walk.

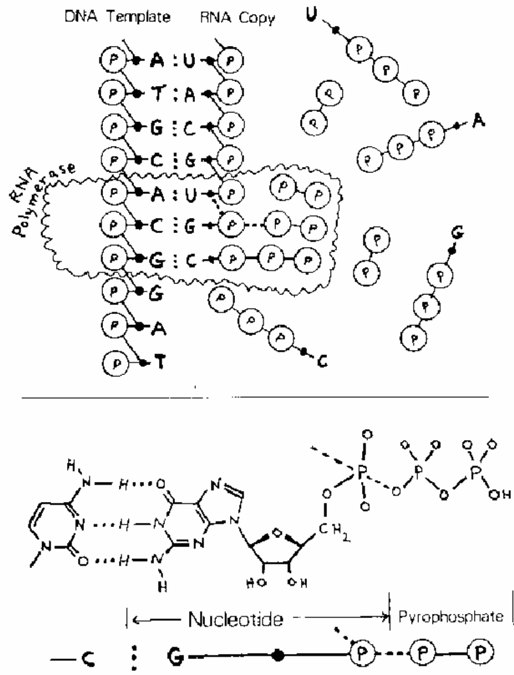


Kinds of computation graph for Brownian computers

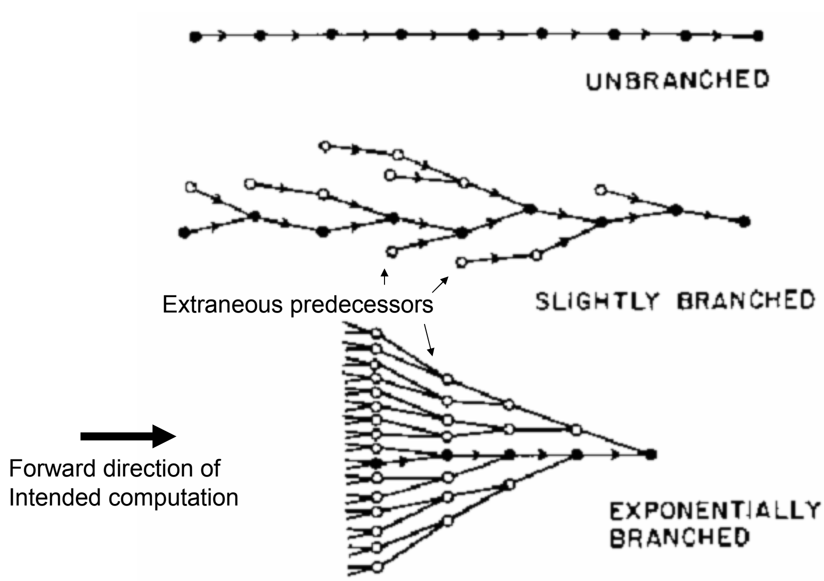

## Potential Energy Landscape for Brownian Computer

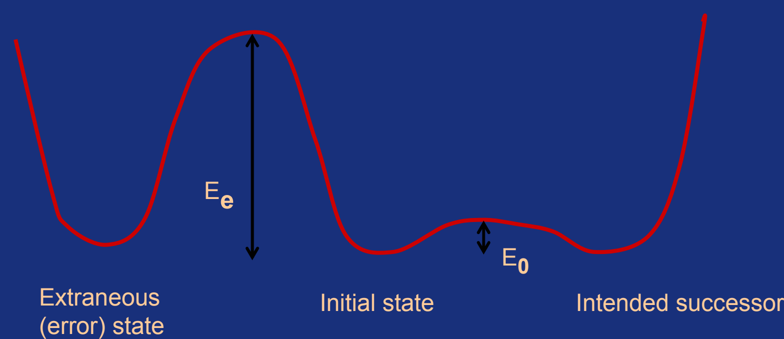


Error probability per step is approx. $\exp[(E_0 - E_e)/kT]$

Error correction is logically many-to-one, so it has a thermodynamic cost, by Landauer's principle.

Conversely, and less obviously, a system's "desire" to make errors is itself a thermodynamic driving force that can be partly harnessed to reduce the cost of correcting the errors.

## Proofreading in DNA Replication

Polymerase (1) tries to insert correct base, but occasionally (2) makes an error. Exonuclease (3) tries to remove errors, but occasionally (4) removes correct bases. When both reactions are driven hard forward the error rate is the product of their individual error rates.

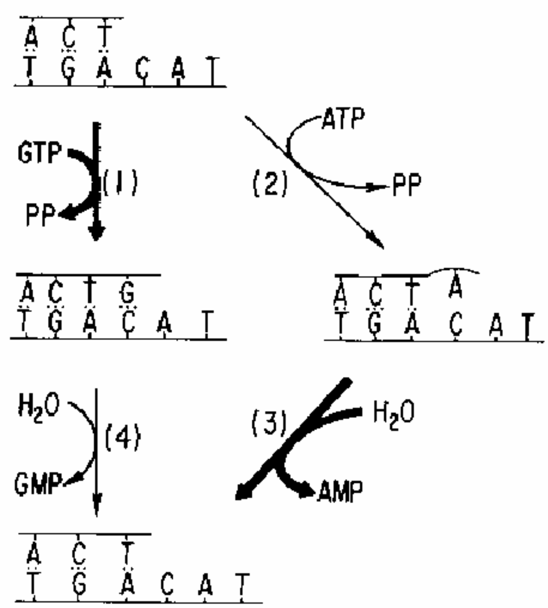

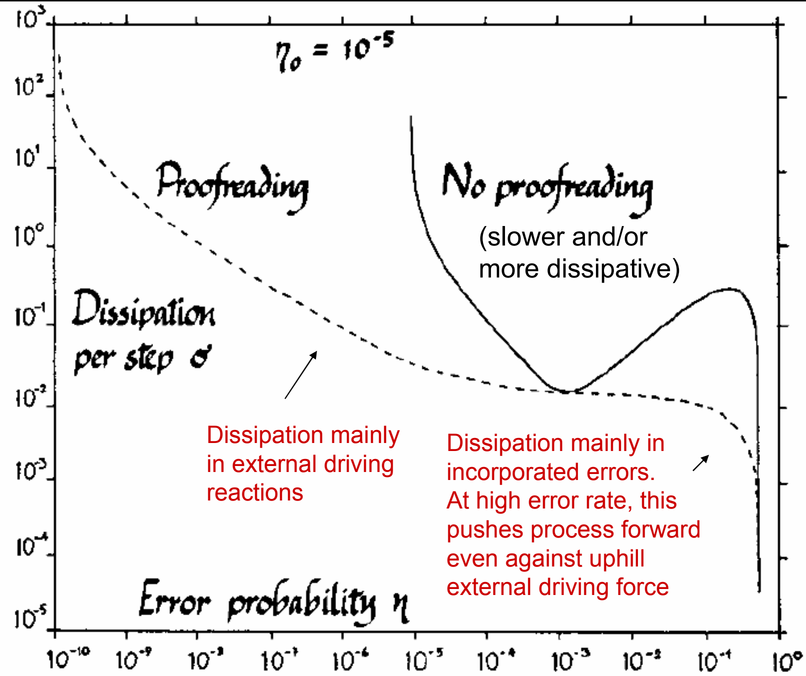


**Fig. 3. Dissipation per digit copied, with and without proofreading.**

For any given hardware environment, e.g. CMOS, DNA polymerase, there will be some tradeoff among dissipation, error, and computation rate. More complicated hardware might reduce the error, and/or increase the amount of computation done per unit energy dissipated.

This tradeoff is largely unexplored, except by engineers.

# *The Thermodynamics of Distillation*

Fractional still

Simple still

Old Seagram's warehouse, Waterloo, Ontario

Simple Distillation

Input

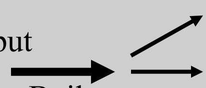

Boil

Vapor (richer in alcohol, condensed to make whiskey

Remaining liquid (richer in water)

Fractional distillation

Input

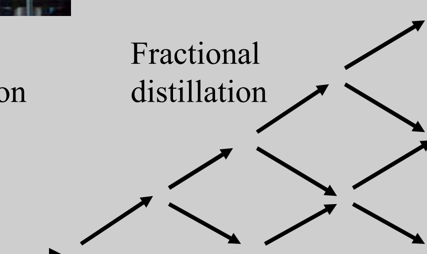

Can approach ideal efficiency in the limit of zero speed:
Reversible: mixture separation cost = −free energy of mixing.
Real stills operate less efficiency but at finite speed.

## Practical Fractional Stills

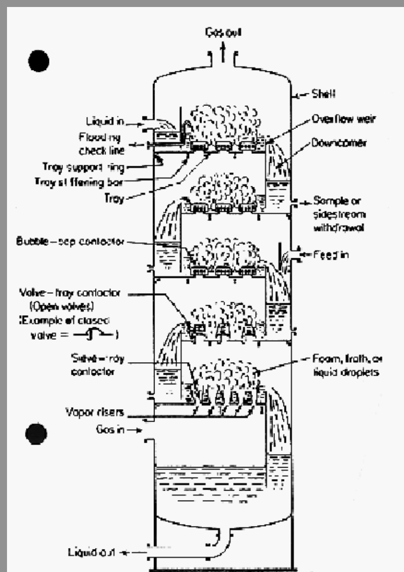


*Ultimate scaling of computation.*

Obviously a 3 dimensional computer that produces heat uniformly throughout its volume is not scalable.

A 1- or 2- dimensional computer can dispose of heat by radiation, if it is warmer than 3K.

Conduction won't work unless a cold reservoir is nearby. Convection is more complicated, involving gravity, hydrodynamics, and equation of state of the coolant fluid.

Fortunately 1 and 2- dimensional fault tolerant universal computers exist:

i.e. cellular automata that correct errors by a self-organized hierarchy of majority voting in larger and larger blocks, even though all local transition probabilities are positive.
(P. Gacs math.PR/0003117)

(For quantum computations, two dimensions appear sufficient for fault tolerance

*Dissipation without Computation*

50 C  Simple system: water heated from above

*Temperature gradient is in the wrong direction for convection. Thus we get static dissipation without any sort of computation, other than an analog solution of the Laplace equation.*

10 C

Dissipation-error Tradeoff for Computation

50 C But if the water has impurities

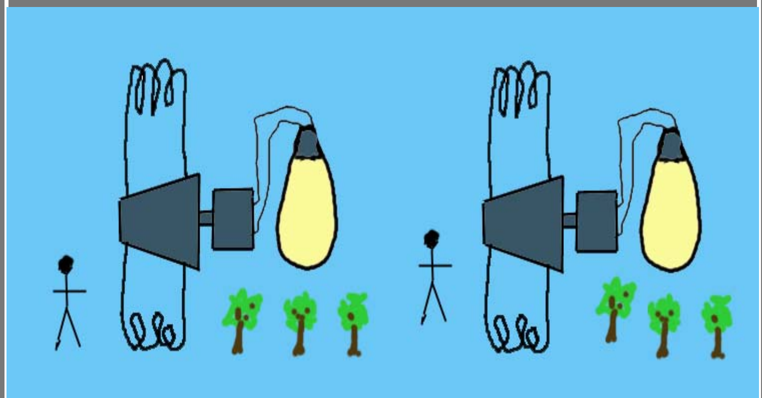

10 C Turbine civilization can maintain and repair itself, do universal computation.

Applying this reversible dynamics to an initial condition (left edge) that is periodic except for a small defect (green) creates a complex deterministic wake in spacetime.

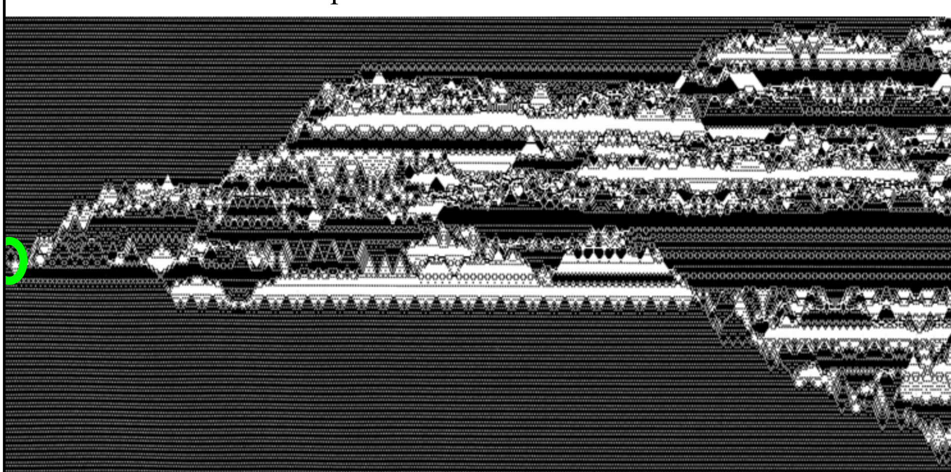

A snapshot of the wake at a later time is *logically deep*, in the sense of containing internal evidence of a nontrivial dynamical history, requiring many computational steps to simulate.

Heat death: a world at thermal equilibrium is no fun.
Our world is only fun because it's (still) out of equilibrium.

For a fully equilibrated system, a single snapshot is typically random, but a pair of snapshots far apart in time, when taken together (as a single 2n bit string) can contain evidence of a nontrivial intervening history.