B Thermodynamics of computation

B.1 Von Neumann-Landaur Principle

B.1.a Information and entropy

¶1. **Entropy:** A quick introduction/review of the entropy concept. We will look at it in more detail soon (Sec. B.2).

¶2. **Information content:** The information content of a signal (message) measures our "surprise," i.e., how unlikely it is.
$I(s) = -\log_b \mathcal{P}\{s\}$, where $\mathcal{P}\{s\}$ is the probability of s.
We take logs so that the information content of independent signals is additive.
We can use any base, with corresponding units *bits*, *nats*, and *dits* (also, hartleys, bans) for $b = 2, e, 10$.

¶3. **1 bit:** Therefore, if a signal has a 50% probability, then it conveys one bit of information.

¶4. **Entropy of information:** The *entropy of a distribution* of signals is their average information content:

$$H(S) = \mathcal{E}\{I(s) \mid s \in S\} = \sum_{s \in S} \mathcal{P}\{s\} I(s) = -\sum_{s \in S} \mathcal{P}\{s\} \log \mathcal{P}\{s\}.$$

Or more briefly, $H = -\sum_k p_k \log p_k$.

¶5. **Shannon's entropy:** According to a well-known story, Shannon was trying to decide what to call this quantity and had considered both "information" and "uncertainty." Because it has the same mathematical form as statistical entropy in physics, von Neumann suggested he call it "entropy," because "nobody knows what entropy really is, so in a debate you will always have the advantage."[9]
(This is one version of the quote.)

¶6. **Uniform distribution:** If there are N signals that are all equally likely, then $H = \log N$.
Therefore, if we have eight equally likely possibilities, the entropy is

[9] https://en.wikipedia.org/wiki/History_of_entropy (accessed 2012-08-24).

$H = \lg 8 = 3$ bits.
A uniform distribution maximizes the entropy (and minimizes the ability to guess).

¶7. **Describing state:** In computing, we are often concerned with the *state* of the computation, which is realized by the state of a physical system.
Consider a physical system with three *degrees of freedom* (*DoF*), each with 1024 possible values.
There are $N = 1024^3 = 2^{30}$ possible states, each describable by three 10-bit integers.

¶8. If we don't care about the distance between states (i.e., distance on each axis), then states can be specified equally well by six 5-bit numbers or 30 bits, etc. (or $30 \log_{10} 2 \approx 9.03$ digits).
Any scheme that allows us to identify all 2^{30} states will do.
There are 30 binary degrees of freedom.

¶9. In computing we often have to deal with things that grow exponentially or are exponentially large (due to combinatorial explosion), such as solution spaces.
(For example, NP problems are characterized by the apparent necessity to search a space that grows exponentially with problem size.)

¶10. In such cases, we are often most concerned with the *exponents* and how they relate.
Therefore it is convenient to deal with their logarithms (i.e., with logarithmic quantities).

¶11. The logarithm represents, in a scale-independent way, the degrees of freedom generating the space.

¶12. **Indefinite logarithm:**[10] Different logarithm bases amount to different units of measurement for logarithmic quantities (such as information and entropy).
As with other quantities, we can leave the units unspecified, so long as

[10]Frank (2005a) provides a formal definition for the indefinite logarithm. I am using the idea less formally, an "unspecified logarithm," whose base is not mentioned. This is a compromise between Frank's concept and familiar notation; we'll see how well it works!

we do so consistently.
I will use the notation "$\log x$" for an *indefinite logarithm*, that is, a logarithm with an unspecified base.
When I mean a specific base, I will write $\ln x$, $\lg x$, $\log_{10} x$, etc.

¶13. $\lg x = \log x / \log 2$, $\ln x = \log x / \log e$, etc.
(The units can be defined bit $= \log 2$, nat $= \log e$, digit $= \log 10$, etc.)

B.1.b The von Neumann-Landauer bound

¶1. **Thermodynamic entropy:** Thermodynamic entropy is unknown information residing in the physical state.
Macroscopic thermodynamic entropy S is related to microscopic information entropy H by Boltzmann's constant, which expresses the entropy in thermodynamical units (energy over temperature).
If H is measured in nats, then $S = k_{\mathrm{B}} H = k_{\mathrm{B}} \ln N$.
When using indefinite logarithms, I will drop the "B" subscript: $S = kH = k \log N$.
The physical dimensions of entropy are usually expressed as energy over temperature (e.g., joules per kelvin), but the dimension of temperature is energy per DoF (measured logarithmically), so the fundamental dimension of entropy is degrees of freedom, as we would expect.
(There are technical details that I am skipping.)

¶2. **Macrostates and microstates:** Consider a macroscopic system composed of many microscopic parts (e.g., a fluid composed of many molecules).
In general a very large number of *microstates* (or *microconfigurations*) — such as positions and momentums of molecules — will correspond to a given *macrostate* (or *macroconfiguration*) — such as a combination of pressure and termperature.
For example, with $m = 10^{20}$ particles we have $6m$ degrees of freedom, and a $6m$-dimensional phase space.

¶3. **Microstates representing a bit:** Suppose we partition the microstates of a system into two macroscopically distinguishable macrostates, one representing 0 and the other representing 1.
For example, whether the electrons are on one plate of a capacitor or the other.

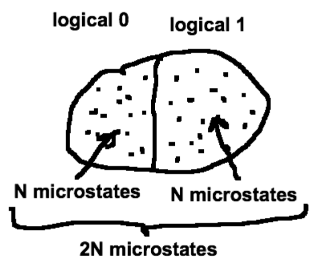

Figure II.3: Physical microstates representing logical states. Setting the bit decreases the entropy: $\Delta H = \lg N - \lg(2N) = -1$ bit. That is, we have one bit of information about its microstate.

¶4. Suppose N microconfigurations correspond to each macroconfiguration (Fig. II.3).
This could be all the positions, velocities, and spins of the many electrons, which we don't care about and cannot control individually.

¶5. If we confine the system to one half of its microstate space, to represent a 0 or a 1, then the entropy (average uncertainty in identifying the microstate) will decrease by one bit.
We don't know the exact microstate, but at least we know which half of the statespace it is in.

¶6. **IBDF and NIBDF:** In general, in physically realizing a computation we distinguish *information-bearing degrees of freedom* (IBDF), which we control, and *non-information-bearing degrees of freedom*, which we do not control (Bennett, 2003).

¶7. **Erasing a bit:** Consider the erasing or clearing a bit (i.e., setting it to 0, no matter what its previous state).

¶8. We are losing one bit of physical information. The physical information still exists, but we have lost track of it.

Suppose we have N physical microstates per logical macrostate (logical 0 or logical 1).
Before the bit is erased it can be in one of $2N$ possible microstates.
There are only N microstates representing its final state.

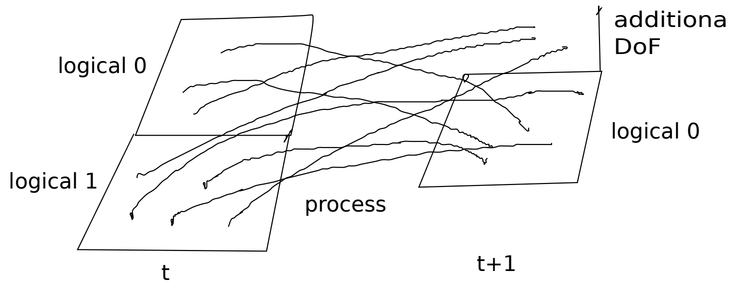

Figure II.4: Thermodynamics of erasing a bit. On the left is the initial state (time t), which may be logical 0 or logical 1; on the right (time $t+1$) the bit has been set to logical 0. In each case there are N microstates representing each prior state, so a total of $2N$ microstates. However, at time $t+1$ the system must be on one of N posterior microstates. Therefore N of the microstate trajectories must exit the defined region of phase space by expanding into additional, uncontrolled degrees of freedom. Therefore entropy of the environment must increase by at least $\Delta S = k\log(2N) - k\log N = k\log 2$. We lose track of this information because it passes into uncontrolled degrees of freedom.

The laws of physics are reversible,[11] so they cannot lose any information.
Physical information can't be destroyed, so it must go into NIBDF (e.g., the environment) (Fig. II.4).
The trajectories have to expand into other DoF (NIBDF) to maintain the phase space volume.

¶9. The information lost, or dissipated into NIBDF (typically as heat), is $\Delta S = k\log(2N) - k\log N = k\log 2$.
($\Delta S = k_{\mathrm{B}} \ln 2 \approx 10$ yJ/K.)
Therefore the increase of energy in the device's environment is $\Delta Q = \Delta S \times T_{\mathrm{env}} = k_{\mathrm{B}} T_{\mathrm{env}} \ln 2 \approx 0.7 k T_{\mathrm{env}}$.
At 300K, $k_{\mathrm{B}} T_{\mathrm{env}} \ln 2 \approx 18$ meV $\approx 3 \times 10^{-9}$ pJ $= 3$ zJ
Recall, for reliable operation we need $40 k_{\mathrm{B}} T_{\mathrm{env}}$ to $100 k_{\mathrm{B}} T_{\mathrm{env}}$, vs. the vNL limit of $0.7 k_{\mathrm{B}} T_{\mathrm{env}}$.

[11] This is true in both classical and quantum physics. In the latter case, we cannot have $2N$ quantum states mapping reversibly into only N quantum states.

¶10. **von Neumann – Landauer bound:** We will see that this is the minimum energy dissipation for any irreversible operation (such as erasing a bit).
It's called the *von Neumann – Landauer (VNL) bound* (or sometimes simply the *Landauer bound*).
VN suggested the idea in 1949, but it was published first by Rolf Landauer (IBM) in 1961.[12]

¶11. "From a technological perspective, energy dissipation per logic operation in present-day silicon-based digital circuits is about a factor of 1,000 greater than the ultimate Landauer limit, but is predicted to quickly attain it within the next couple of decades." (Berut et al., 2012)
That is, current circuits are about 18 eV.

¶12. **Experimental confirmation:** In research reported in March 2012 Berut et al. (2012) confirmed experimentally the Landauer Principle and showed that it is the erasure that dissipates energy.

¶13. They trapped a 2μ silica ball in either of two laser traps, representing logical 0 and logical 1.
For storage, the potential barrier was greater than $8k_{\mathrm{B}}T$.
For erasure, the barrier was lowered to $2.2k_{\mathrm{B}}T$ by decreasing the power of the lasers and tilting the device to put it into the logical 0 state.
See Fig. II.5.

¶14. At these small sizes, heat is a stochastic property, so the dissipated heat was computed by averaging over multiple trials the trajectory of the particle:

$$\langle Q \rangle = \left\langle - \int_0^\tau \dot{x}(t) \frac{\partial U(x,t)}{\partial x} \mathrm{d}t \right\rangle_x .$$

¶15. "incomplete erasure leads to less dissipated heat. For a success rate of r, the Landauer bound can be generalized to"

$$\langle Q \rangle^r_{\text{Landauer}} = kT[\log 2 + r \log r + (1-r)\log(1-r)] = kT[\log 2 - H(r, 1-r)].$$

"Thus, no heat needs to be produced for $r = 0.5$" (Berut et al., 2012).

[12]See Landauer (1961), reprinted in Leff & Rex (1990) and Leff & Rex (2003), which include a number of other papers analyzing the VNL principle.

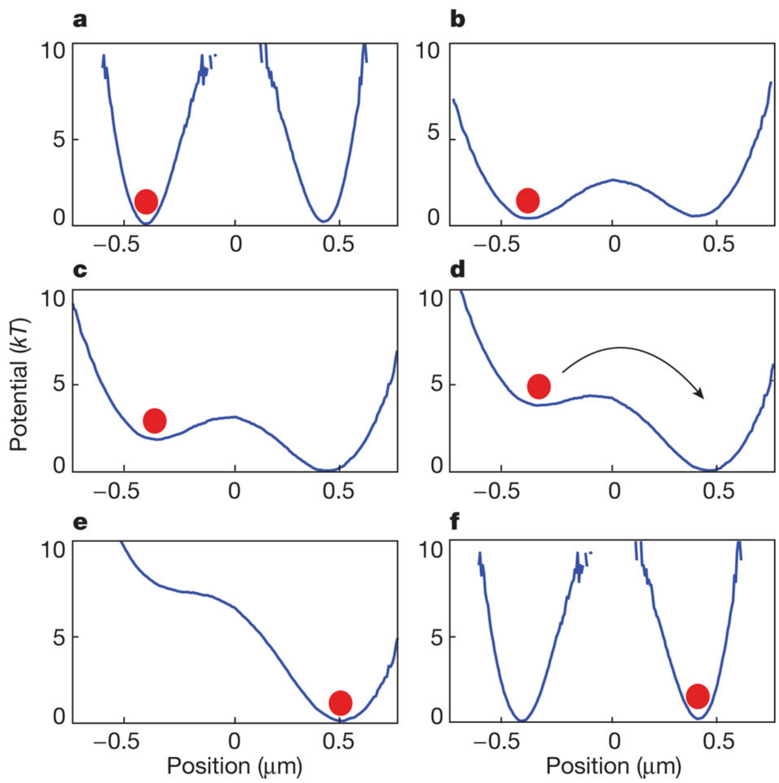

Figure II.5: Erasing a bit by changing potential barrier. (Figure from Berut et al. (2012).)

B.1.c Irreversible operations

¶1. **Macroscopic and microscopic entropy:** Suppose the phase space is divided into M macrostates of size $N_1, N_2, \ldots, N_M$, where $N = N_1 + N_2 + \cdots + N_M$.

¶2. Let p_{ij} be the probability the device is in microstate i of macrostate j. The total entropy is

$$S = -k \sum_{ij} p_{ij} \log p_{ij}. \tag{II.2}$$

¶3. We can separate this into the *macroscopic entropy* associated with the macrostates (IBDF) and the *microscopic entropy* associated with the microstates (NIBDF).

¶4. Let $P_j = \sum_{i=1}^{N_j} p_{ij}$ be the probability of being in macrostate j. Then Eq. II.2 can be rearranged (exercise):

$$S = -k \sum_j P_j \log P_j - k \sum_j P_j \sum_{i=1}^{N_j} \frac{p_{ij}}{P_j} \log \frac{p_{ij}}{P_j} = S_\mathrm{i} + S_\mathrm{h}.$$

The first term is the macrostate entropy (IBDF):

$$S_\mathrm{i} = -k \sum_j P_j \log P_j.$$

The second is the microstate entropy (NIBDF):

$$S_\mathrm{h} = -k \sum_j P_j \sum_{i=1}^{N_j} \frac{p_{ij}}{P_j} \log \frac{p_{ij}}{P_j}.$$

¶5. When we erase a bit, we go from a maximum S_i of 1 bit (if 0 and 1 are equally likely), to 0 bits (since there is no uncertainty).
Thus we lose one bit, and the macrostate entropy decreases $\Delta S_\mathrm{i} = -k \log 2$.
Since according to the Second Law of Thermodynamics $\Delta S \geq 0$, we have a minimum increase in microstate entropy, $\Delta S_\mathrm{h} \geq k \log 2$.
Typically this is dissipated as heat, $Q \geq kT \log 2$.
The information becomes inaccessible, unusable.

¶6. The actual entropy decrease can be less than 1 bit if the 0 and 1 are not equally likely initial states.

¶7. **Standard logic gates:** The standard logic gates (AND, OR, XOR, NAND, etc.) have two input bits and one output bit. Therefore the output will have lower entropy than the input, and so these gates must dissipate at least 1 bit of entropy, $kT \log 2$ energy.

¶8. Consider AND. If the four inputs 00, 01, 10, 11, are equally likely, then the input entropy is $H_{\mathrm{i}} = 2$ bits.
However the output entropy will be $H_{\mathrm{o}} = -(1/4)\lg(1/4) - (3/4)\lg(3/4) = 0.811$, so the entropy lost is 1.189 bits.
Hence $\Delta Q \geq T(kH_{\mathrm{o}} \ln 2) \approx 0.83kT$.
For each gate, we can express H_{o} in terms of the probabilities of the inputs and compute the decrease from H_{i} (exercise).

¶9. If the inputs are not equally likely, then the input entropy will be less than 2 bits, but we will still have $H_{\mathrm{i}} > H_{\mathrm{o}}$ and energy dissipated. (Except in a trivial, uninteresting case. What is it?)

¶10. **Irreversible operations:** More generally, any irreversible operation (non-invertible function) will lose information, which has to be dissipated into the environment.
If the function is not one-to-one (injective), then at least two inputs map to the same output, and so information about the inputs is lost.

¶11. **Assignment operation:** Changing a bit, that is, overwriting a bit with another bit, is a fundamental irreversible operation, subject to the VNL limit.

¶12. **Convergent flow of control:** When two control paths join, we forget where we came from, and so again we must dissipate at least a bit's worth of entropy (Bennett, 2003).

¶13. **Reversible operations:** The foregoing suggests that reversible operations might not be subject to the VNL limit, and this is in fact the case, as we will see.

¶14. **Maxwell's Demon:** The preceding observations have important connections with the problem of Maxwell's Demon and its resolution.

Briefly, the demon has to reset its mechanism after each measurement in preparation for the next measurement, and this dissipates at least $kT \log 2$ energy into the heat bath for each decision that it makes.
The demon must "pay" for any information that it acquires.
Therefore, the demon cannot do useful work.
Further discussion is outside the scope of this class, so if you are interested, please see Leff & Rex (2003) and Leff & Rex (1990) (which have a large intersection), in which many of the papers on the topic are collected.