23.08.2021

Philippe Faist

Thermodynamics of Quantum Information Processing

Goals:

* Understand the fundamental energy cost of computation.
* To what extent is thermodynamics and the second law universal?

  [Thermodynamics is usually presented as a theory of macroscopic gases — why can we also apply it to solid state physics, chemistry, electromagnetic radiation ("gas of photons"), and even black holes?]

  → Role of the observer in thermodynamics

  → Resolve the paradox of Maxwell's demon
* Develop robust tools to analyze the work cost of arbitrary quantum processes.

References:

* Feynman lectures on computation, Chap. 5
* C. Bennett, Notes on Landauer's principle, reversible computation and Maxwell's demon (2003) & The thermodynamics of computation — a review (1982)
* R. Landauer, Irreversibility and Heat Generation in the Computing Process. (1961)
* L. del Rio et al., The thermodynamic meaning of negative entropy, Nature (2011)
* Ph.F. et al., The minimal work cost of information processing, (2015)

Lecture I: Landauer's Principle.

1. The Szilárd engine. (Szilárd, 1929)

Consider a single-particle gas enclosed in a box:

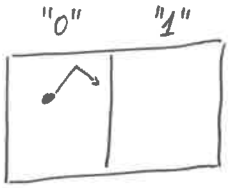

the particle can be on the left ("0") or on the right ("1") side of the box, representing one bit of information.

[It is ok to consider the single particle as a gas here, because we will be considering very slow operations, and we will be averaging over long periods of time.]

* Extracting work from knowledge: Say we know that the particle is on the left:

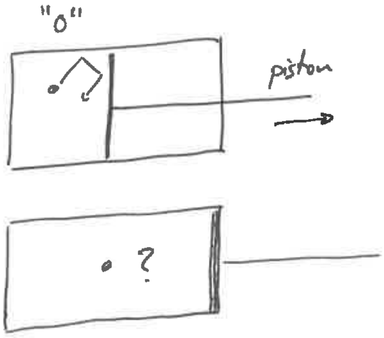

1. Attach a piston to the separator
2. Let the one-particle gas expand isothermally, extracting work (keep system in contact with an environment at temperature T)

We can calculate the amount of work extracted using the corresponding formula for the isothermal expansion of an ideal gas:

$$\underline{W_{extr.}} = k T \cdot \ln\left(\frac{V_{final}}{V_{initial}}\right) = \underline{k T \cdot \ln(2)}$$

Boltzmann's constant

But the particle can now be anywhere: We lost the knowledge of the bit's value.

* Resetting to Zero. If we don't know where the particle is, we can reset it to the left side ("0") by compressing the gas.

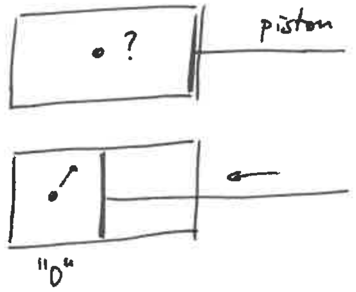

We need to perform work:

$$\underline{W_{reset}} = kT \ln\left(\frac{V_{final}}{V_{initial}}\right) = \underline{-kT\ln(2)}$$

[the reverse of the work extraction process.]

→ We can trade the knowledge of the bit's value ("0" or "1") for $kT\ln(2)$ work.

* Flip "0" to "1": We can transform "0" to "1":

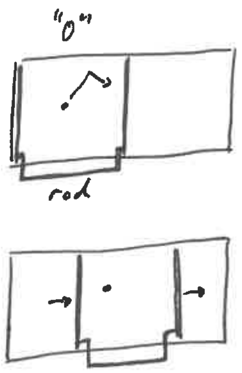

1. Insert two separators, in the middle and on the left edge. Connect them with a rigid rod.
2. Give a nudge to the right. The two separators will set in motion and arrive to the other side, carrying the particle to the "1" state.

→ Switching "0" to "1" costs no work!

2. Landauer's Principle. (Landauer, 1961)

Consider a macroscopic device which stores information:

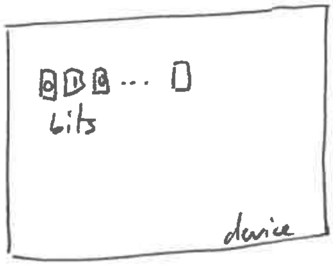

For instance, each bit is realized by a ferromagnet

$\boxed{0} = \boxed{\uparrow\uparrow\uparrow \cdots \; \cdots \uparrow\uparrow}$ $\qquad$ $\boxed{1} = \boxed{\downarrow\downarrow\downarrow \cdots \; \cdots \downarrow\downarrow\downarrow}$

As a macroscopic object, the device is a thermodynamic system. (E.g., it has a temperature.) There are many microscopic configurations which are irrelevant macroscopically: These form the device's thermodynamic entropy $S_{th} = k \cdot \ln$ (# of available microscopic configurations). They are <u>inaccessible degrees of freedom</u>. The bit values are <u>accessible degrees of freedom</u>.

Suppose we would like to reset an unknown bit to zero. Let's count the available microstates:

before: $\underbrace{2}_{\text{possible bit values}} \times \underbrace{N_{micro}}_{\text{\# of microstates in inaccessible d.o.f.'s (environment)}}$ $\longrightarrow$ after: $\underbrace{1}_{\text{only one possible bit value}} \times \underbrace{N'_{micro}}_{\text{(might have changed)}}$

But the laws of physics are reversible for a closed system, and the # of microscopic configurations cannot have decreased.

→ must have $N'_{micro} = 2 N_{micro}$

→ causes increase in thermodynamic entropy of the environment

$$\Delta S_{th} = k \ln (N'_{micro}) - k \ln (N_{micro}) = k \ln (2)$$

→ minimal heat dissipation $\underline{\Delta Q} \geq T \Delta S = \underline{k T \ln (2)}$ $\quad$ $\sim 3 \cdot 10^{-21}\ \mathrm{J} \sim 0.02\ \mathrm{eV}$ at room temperature

Landauer's Principle: In any logically irreversible process (= one where the input cannot be deduced from the output), entropy is transferred from the accessible degrees of freedom to the inaccessible degrees of freedom, resulting in the dissipation of heat.

This energy must be provided as work. E.g., resetting one bit of information ("erasure of information"), with an environment at temperature T, costs $kT\ln(2)$ work.

Note: No dissipation is required for reversible computation!

3. The work value of information.

* We've seen that
  - → 1 known bit can be traded for $kT\ln(2)$ work
  - → reversible operations require no work

  (See also Matteo's and Ralph's lectures for additional models and protocols.)

* How can we reset an arbitrary quantum state ρ to the fixed state $|0\rangle$? → compress the data first!

  1) data compression: ρ can be encoded onto $\approx H(\rho)$ qubits (given by the information entropy)

  2) reset only those qubits → total work cost $= kT\ln(2) \times H(\rho)$

information entropy

* Example: n qubits are either in the state $|00\ldots0\rangle$ or in the state $|11\ldots1\rangle$.
  1) apply sequence of C-NOT gates → qubit is in either states $|00\ldots0\rangle$ or $|100\ldots0\rangle$
  2) reset only the first bit → total work cost = $kT\ln(2)$

[→ Correlations between the bits can be exploited to reduce the work cost of erasure.]

4. Erasure with side information.

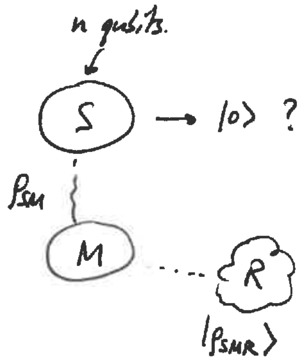

We would like to reset a quantum system S to the pure state $|0\rangle$. We are given a memory system M that is correlated with S.

We can act on M, but we must preserve its reduced state as well as any correlations with any other system R.

del Rio et al. (2011) → protocol with cost $\sim \underbrace{H(S|M)_\rho}_{\text{conditional entropy}} \cdot kT\ln(2)$

Construction:

1) Use the decoupling theorem, a powerful tool of quantum information theory:
   (i) Apply a random unitary on S and trace out k qubits.
   (ii) Thm: as long as $k \lesssim \frac{1}{2}(n - H(S|M))$, then

$$\rho_{K,R} \approx \frac{\mathbb{1}_k}{2^k} \otimes \rho_R$$

(actually $H^{\varepsilon}_{\max}(S|M)$. more on that later.)

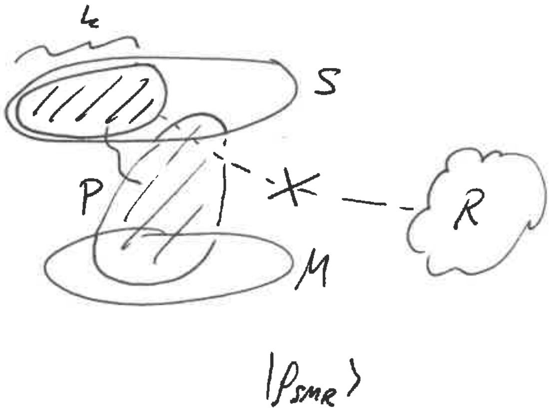

2) Because the k qubits are decoupled from R, there must exist a purification within $S \otimes M$

$$\rho_{SM} \approx |\phi\rangle\langle\phi|_{k,P} \otimes \rho_{rest}$$

3) Extract work by transforming

$$|\phi\rangle_{k,P} \longrightarrow \frac{\mathbb{1}_{k,P}}{2^{2k}} \qquad (2k \text{ bits})$$

$$\text{extracted work} = 2k \cdot k_B T \ln(2) \sim \left(n - H(S|M)\right) \cdot kT\ln(2)$$

Observe that ρ_{MR} remains unchanged because even before the work extraction, P's local reduced state was maximally mixed.

4) Reset S to $|0\rangle$ $\rightarrow$ costs work $= n \cdot kT\ln(2)$

$\rightarrow$ total work cost $\sim$ $\underline{\underline{H(S|M) \cdot kT\ln(2)}}$

(actually, $\sim H^{\varepsilon}_{max}(S|M)$)

<u>Examples:</u>

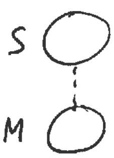

$\rho_{SM} = \frac{1}{2}|00\rangle\langle 00| + \frac{1}{2}|11\rangle\langle 11|$

$H(S|M) = 0$

no work cost to reset S to $|0\rangle$!

$|\rho_{SM}\rangle = \frac{1}{\sqrt{2}}\left(|00\rangle + |11\rangle\right)$

$H(S|M) = -1$

can <u>extract</u> work by resetting S to $|0\rangle$!

Lecture II: The Work Cost of Quantum Processes.

1. Framework – free operations

* To derive fundamental limits, we want a permissive choice of free operations.

* We only consider systems with a fully degenerate Hamiltonian $H=0$ (for now)

* Are allowed for free: any completely positive (c.p.), trace-nonincreasing (t.n.) map $\Phi_{S\to S'}$ that satisfies

$$\Phi_{S\to S'}(\mathbb{1}_S) \leq \mathbb{1}_{S'} . \qquad (*)$$

(Recall: $A \leq B$ means that $B-A$ is positive semidefinite.)

* Why? If we allowed any c.p., t.p. map Φ' with $\Phi'(\text{thermal state}) \neq (\text{thermal state})$, then one could extract arbitrary work using Φ'^{Gen} (cf. e.g. Matteo's lecture.)
* Ultimately, $(*)$ is an assumption that should be justified by some underlying (microscopic) physical model.
* Any trace-nonincreasing (and c.p.) map $\Phi_{S\to S'}$ that satisfies $\Phi(\mathbb{1}_S) \leq \mathbb{1}_{S'}$ can be dilated to a trace-preserving, c.p. map $\tilde{\Phi}_{SK\to S'L}$ (on a larger system) that satisfies $\tilde{\Phi}(\mathbb{1}_{SK}) = \mathbb{1}_{S'L}$

[→ we shouldn't worry about Φ maps being trace nonincreasing or not obeying $\Phi(\mathbb{1})=\mathbb{1}$, or even if $\dim(S) \neq \dim(S')$]

* *Examples.* Q a qubit:

→ $\Phi_{Q\to Q}(\sigma) = \mathrm{tr}(\sigma)\cdot|0\rangle\langle 0| \quad \forall\sigma$ (reset to $|0\rangle$)

$\Phi_{Q\to Q}(\mathbb{1}) = 2|0\rangle\langle 0| \not\leq \mathbb{1}$ → not a free operation

→ $\Phi_{Q\to Q}(\sigma) = U\sigma U^\dagger$ fixed unitary U

$\Phi_{Q\to Q}(\mathbb{1}) = U\mathbb{1}U^\dagger = \mathbb{1} \leq \mathbb{1}$ ✓ → free operation.

→ $\Phi_{Q\to Q}(\sigma) = \langle 0|\sigma|0\rangle\cdot\rho$ for fixed ρ

(prepare state ρ as long as input is $|0\rangle\langle 0|$.)

$\Phi_{Q\to Q}(\mathbb{1}) = \rho \leq \mathbb{1}$ ✓ → free operation.

2. Framework — work and processes

* Work is stored in a battery (cf. Matteo's lecture)

$$\underbrace{|0|0| \cdots |0|0|0|}_{\text{pure qubits in the } |0\rangle \text{ state}}\underbrace{?|?|?|?| \cdots |?|}_{\text{maximally mixed qubits}}$$ ← register of qubits.

The battery is required to always be in a state of the form

$$\tau(\lambda) = 2^{-\lambda}\mathbb{1}_{2^\lambda} = \text{uniform state of rank } 2^\lambda$$

($\hat{=}$ λ maximally mixed qubits, the rest is pure.)

λ is the "depletion state" of the battery.

(Many other battery models are possible, and many are equivalent.)

* The most general process in this framework has the form:

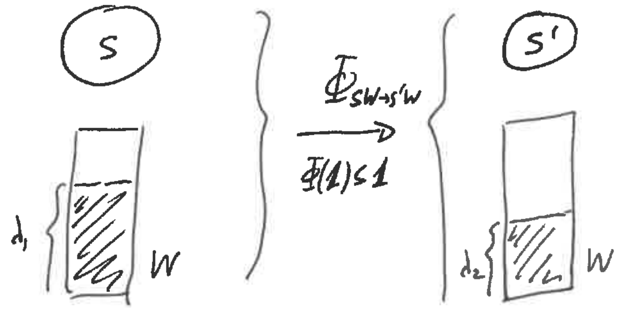

* The resulting induced process on $S \to S'$ is

$$\mathcal{T}^{[\Phi]}_{S\to S'}(\sigma_S) = \mathrm{tr}_W\left\{ \Phi\left(\sigma_S \otimes 2^{-d_1}\mathbb{1}_{2^{d_1}}\right) \cdot \mathbb{1}_{2^{d_2}} \right\}$$

"asserts"/"checks" that W doesn't have a depletion state exceeding d_2

* <u>The task</u> we consider is that of implementing a given cp. t.n map $\mathcal{E}_{S\to S'}$ on an input state σ_S (preserving correlations with a reference system).

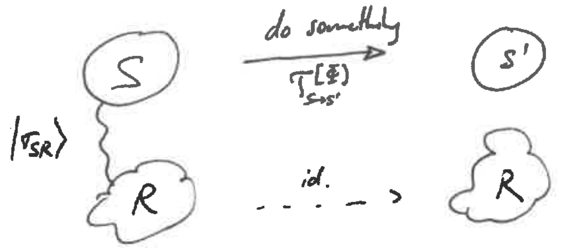

$$\mathcal{T}^{[\Phi]}_{S\to S'}(\sigma_{SR}) \overset{!}{=} \mathcal{E}_{S\to S'}(\sigma_{SR})$$

requirement

What is the minimal work cost over implementations (Φ, W, d_1, d_2) that satisfy our requirement?

* <u>Thm</u>: The minimal work cost is $\boxed{W_{cost}(\mathcal{E},\sigma) = \log \| \mathcal{E}(\Pi^{\sigma}_S) \|_\infty}$

$\|A\|_\infty$ = max. singular value of A.

projector onto the support of σ_S

* Proof of $W_{cost}(\mathcal{E},\sigma) \geq \log \| \mathcal{E}(\Pi_S^\sigma) \|_\infty$: For any implementation $(\Phi, W, \lambda_1, \lambda_2)$:

$$\| \mathcal{E}(\Pi_S^\sigma) \|_\infty \underset{\substack{\uparrow \\ \text{our requirement}}}{=} \| \mathcal{T}^{[\Phi]}(\Pi_S^\sigma) \|_\infty$$

$$= \| \mathrm{tr}_W \{ \Phi(\underbrace{\Pi_S^\sigma}_{\leq \mathbb{1}_S} \otimes 2^{-\lambda_1} \underbrace{\mathbb{1}_{2^{\lambda_1}}}_{\leq \mathbb{1}_W}) \cdot \mathbb{1}_{2^{\lambda_2}} \} \|_\infty$$

$$\leq 2^{-\lambda_1} \| \mathrm{tr}_W \{ \underbrace{\Phi(\mathbb{1}_{SW})}_{\leq \mathbb{1}_{SW}} \cdot \mathbb{1}_{2^{\lambda_2}} \} \|_\infty$$

$$\leq 2^{-\lambda_1} \underbrace{\mathrm{tr}(\mathbb{1}_{2^{\lambda_2}})}_{= 2^{\lambda_2}} \cdot \underbrace{\| \mathbb{1}_{S'} \|_\infty}_{=1} = 2^{-(\lambda_1 - \lambda_2)}$$

$$\Longrightarrow \lambda_2 - \lambda_1 \geq \log \| \mathcal{E}(\Pi_S^\sigma) \|_\infty \qquad \text{true for all implementations}$$

$$\Longrightarrow W_{cost}(\mathcal{E},\sigma) \geq \log \| \mathcal{E}(\Pi_S^\sigma) \|_\infty .$$

Proof that $\exists\, \Phi, W, \lambda_1, \lambda_2$ s.t. $W_{cost}(\mathcal{E},\sigma) = \lambda_2 - \lambda_1 = \log \| \mathcal{E}(\Pi_S^\sigma) \|_\infty$ (or at least, arbitrarily close) $\longrightarrow$ exercise.

3. Entropic picture.

* A c.p. t.n. map $\mathcal{E}_{S\to S'}$ has a Stinespring dilation :

$$\mathcal{E}_{S\to S'}(\cdot) = \mathrm{tr}_E \{ V_{S\to ES'} (\cdot) V^\dagger \}$$

if $\mathcal{E}$ is t.p. $\to$ V is an isometry
if $\mathcal{E}$ is t.n. $\to$ $V^\dagger V \leq \mathbb{1}$

* Given an input state $|\sigma_{XR}\rangle$ (including reference system R), the output including the environment and the reference is

$$|\rho_{S'ER}\rangle = V_{S\to ES'} |\sigma_{SR}\rangle$$

* Alternative expression for the work cost:

$$\underline{W_{cost}(\mathcal{E},\sigma)} = \log \|\mathcal{E}(\Pi_S^\sigma)\|_\infty = \log \| \mathrm{tr}_E\{V_{S\to S'E}\, \Pi_S^\sigma V^\dagger\}\|_\infty$$

$$= \log \| \mathrm{tr}_E\{\Pi_{S'E}^\rho\}\|_\infty = \underline{H_{max,0}(E|S')_\rho}$$

projector onto the support of $\rho_{S'E}$

max-entropy (Rényi-0) Renner '05

→ Work cost = information discarded to the environment.

* Robustness to ε-approximations. If the implementation is ε-close to the required $\mathcal{E}_{S\to S'}(\sigma_{SR})$, it should still be acceptable. Maybe we can save some work this way!

$$\underline{W_{cost}^{(\varepsilon)}(\mathcal{E},\sigma)} = \min_{\tilde\rho_{S'R} \approx_\varepsilon \rho_{S'R}} H_{max,0}(E|S')_\rho = \underline{\underline{H^\varepsilon_{max,0}(E|S')_\rho}}$$

suitable metric, e.g. trace distance

smooth max-entropy (Rényi-zero) Renner '05

cf. also Nelly's lecture.

4. Remarks.

* Is the bound achievable in a less permissive framework?

Yes, using conditional erasure!

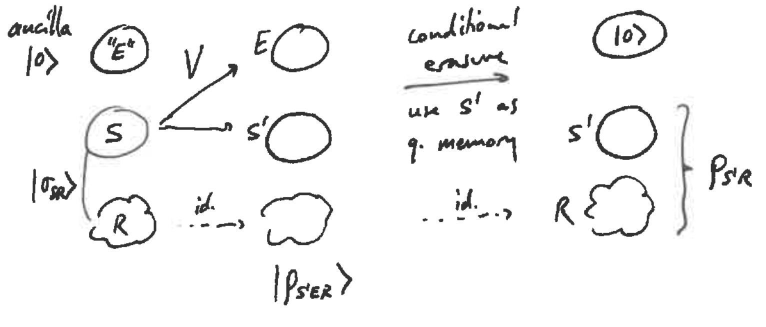

work cost (cf. Lecture I)

$$\approx \underline{H^\varepsilon_{max,0}(E|S')_\rho}$$

* I.i.d. limit. If we consider many independent copies of the process, $\mathcal{E} \to \mathcal{E}^{\otimes N}$, $\sigma \to \sigma^{\otimes N}$, we have

$$H^{\varepsilon}_{max,0}(E^n|S'^n)_{\rho^{\otimes n}} = n\, H(E|S')_\rho + O(\sqrt{n})$$

asymptotic equipartition property

So the work cost per copy is

$$\frac{1}{n} W^{\varepsilon}_{cost}(\mathcal{E}^{\otimes n}, \sigma^{\otimes n}) \xrightarrow{n\to\infty} H(E|S')_\rho = \underbrace{H(ES')_\rho}_{=H(S)_\sigma} - H(S')_\rho = H(S) - H(S')$$

= difference between input & output entropy

→ the work cost derives from a thermodynamic potential, the von Neumann entropy.
(It is independent of the details of $\mathcal{E}$, only depends on σ and $\mathcal{E}(\sigma)$.)

→ emergence of macroscopic thermodynamic behavior

* Nontrivial Hamiltonians. We can generalize this approach to $H \neq 0$. Replace:

free op. : $\Phi(\mathbb{1}) \leq \mathbb{1}$ ⟶ $\Phi(\Gamma_S) \leq \Gamma_{S'}$ "Gibbs sub-preserving maps" $\Gamma_S = e^{-\beta H_S}$ Gibbs weights

work cost : $W^{\varepsilon}_{cost}(\mathcal{E}, \sigma) = H^{\varepsilon}_{max,0}(E|S')$ ⟶ it gets messy.

i.i.d. cost : $H(S)_\sigma - H(S')_\rho$ ⟶ $D(\rho_{S'} \| \Gamma_{S'}) - D(\sigma_S \| \Gamma_S)$

In general, the achievability of this work cost in a more permissive framework for $H \neq 0$ is unclear.

* A consistent treatment of <u>observers</u> in thermodynamics:

Each observer can individually assess whether $\Phi(\mathbb{1}) \leq \mathbb{1}$ (or $\Phi(\Gamma) \leq \Gamma$) is a good assumption <u>at their scale</u> (eg. microscopic/macroscopic), and derive <u>their</u> version of "thermodynamics" applicable to their viewpoint.

<u>Example</u>: memory register

|0|0|1|1|0|?|?|?|?| $\longrightarrow$ say: bits are thermodynamic components (eg. ferromagnet)

$\boxed{0} = \begin{vmatrix} \uparrow\uparrow & - & - \\ - & - & \uparrow\uparrow \end{vmatrix}$ $\quad$ $\boxed{1} = \begin{vmatrix} \downarrow\downarrow & \cdots & \\ \cdot & \cdot & \downarrow\downarrow \end{vmatrix}$

microscopic picture $\longrightarrow$ say we have thermalizing dynamics (standard, no surprises)

then $\Phi(e^{-\beta H}) \leq e^{-\beta H}$ ✓

for the natural dynamics that can happen on the system

macroscopic picture $\longrightarrow$ we can only "see" the bit values (and manipulate them)

then $\Phi(\mathbb{1}_{\text{memory bits}}) \leq \mathbb{1}_{\text{memory bits}}$ is implied by the microscopic description.

$\rightsquigarrow$ <u>Landauer's Principle</u>, <u>work cost of computation</u>

but if we had hidden "Maxwell demons" at the microscopic level performing fine-tuned operations, then $\Phi(\mathbb{1}_{\text{mem.}}) \nleq \mathbb{1}_{\text{mem.}}$ and the "demons" could trick us into thinking we witnessed a violation of the second law (and Landauer's principle would appear violated).

$\rightarrow$ When does the second law hold? $\rightarrow$ (at least) when $\Phi(\mathbb{1}) \leq \mathbb{1}$ is a justified assumption.

cf. Gibbs' paradox (Jaynes 1992) $\rightsquigarrow$ exercise.