Quantum Computing

Università della Svizzera italiana

Fabian Bosshard

July 1, 2026

This page is intentionally left blank.

Contents

Preface

This document contains unofficial student-made notes for the course Quantum Computing taught by Stefan Wolf with the assistance of Lorenzo Laneve in Spring 2026 (academic year 2025–2026) at the Università della Svizzera italiana.

These notes are mainly based on the course materials, especially [1], exercise sessions, and handwritten notes taken during the lectures.

Occasionally, additional references were consulted. Section 1.3 uses [2] for interaction-free measurements. Section 3 and Section 5 use [3, 4] for additional background on introductory quantum physics and standard quantum-mechanical postulates. Digression 5.1 uses [5] for the rigged-Hilbert-space interpretation of position kets. The digression on operators in Section 4 closely follows the corresponding sections in the textbook [6]. Section 6.1 uses [7] for the Bloch-sphere mixed-state discussion. Section 9.3 draws on [8, 9, 10, 11]. Section 11 and Section 11.6 use [12, 13, 14, 15, 16, 17]. Section 12 uses [18, 19].

Some figures are based on TikZ code from external sources. Section 1.1.4 uses [20] for the Malus-law drawing. Section 2.2 uses [21] for the gas microstate and macrostate drawings. Section 3.1 uses [22] for the cavity and [23] for the black-body radiation figures.

If you spot an error, please report it to `fabianlucasbosshard@gmail.com`. The LaTeX source is available at `https://github.com/fabianbosshard/usi-informatics-course-summaries`.

References

[1] Ulla Aeschbacher, Arne Hansen, and Stefan Wolf. Invitation to Quantum Informatics. vdf Hochschulverlag AG an der ETH Zürich, 2019. ISBN: 978-3-7281-3988-7. URL: https://vdf.ch/invitation-to-quantum-informatics-e-book.html.

[2] Avshalom C. Elitzur and Lev Vaidman. "Quantum Mechanical Interaction-Free Measurements". In: Foundations of Physics 23.7 (1993), pp. 987–997. URL: https://doi.org/10.1007/BF00736012.

[3] Lídia del Rio, Lorenzo Laneve, and Giulia Carocari. Quantum Physics for Non-Physicists. Lecture notes. 2020. URL: https://www.llaneve.net/files/qpnp-notes.pdf.

[4] Nouredine Zettili. Quantum Mechanics: Concepts and Applications. 2nd ed. Wiley, 2009. URL: https://openlibrary.org/isbn/9780470026793.

[5] Rafael de la Madrid. "The Role of the Rigged Hilbert Space in Quantum Mechanics". In: European Journal of Physics 26.2 (2005), pp. 287–312. URL: https://arxiv.org/abs/quant-ph/0502053.

[6] Walter Rudin. Functional Analysis. 2nd ed. McGraw-Hill, 1991. ISBN: 0-07-054236-8. URL: https://openlibrary.org/isbn/9780070542365.

[7] Bloch Sphere. Wikipedia. URL: https://en.wikipedia.org/wiki/Bloch_sphere (visited on 06/30/2026).

[8] Ernst Specker. "Die Logik nicht gleichzeitig entscheidbarer Aussagen". In: Dialectica 14.2–3 (1960), pp. 239–246. URL: https://doi.org/10.1111/j.1746-8361.1960.tb00422.x.

[9] Simon Kochen and Ernst P. Specker. "The Problem of Hidden Variables in Quantum Mechanics". In: Journal of Mathematics and Mechanics 17.1 (1967), pp. 59–87. URL: https://doi.org/10.1512/iumj.1968.17.17004.

[10] Renato Renner and Stefan Wolf. "Quantum Pseudo-Telepathy and the Kochen-Specker Theorem". In: (June 2004), p. 322. URL: https://crypto.ethz.ch/publications/files/RenWol04d.pdf.

[11] Renato Renner and Stefan Wolf. "Ernst Specker and the Hidden Variables". In: Elemente der Mathematik 67.3 (2012), pp. 122–133. URL: https://doi.org/10.4171/EM/201.

[12] Quantum Fourier Transform – Circuit Implementation. Wikipedia. URL: https://en.wikipedia.org/wiki/Quantum_Fourier_transform#Circuit_implementation (visited on 06/30/2026).

[13] Lorenzo Laneve. Quantum Algorithms. Lecture notes. 2026. URL: https://www.llaneve.net/qalg-notes.html.

[14] Peter W. Shor. "Polynomial-Time Algorithms for Prime Factorization and Discrete Logarithms on a Quantum Computer". In: SIAM Journal on Computing 26.5 (1997), pp. 1484–1509. URL: https://arxiv.org/abs/quant-ph/9508027.

[15] Victor Shoup. "Lower Bounds for Discrete Logarithms and Related Problems". In: Advances in Cryptology – EUROCRYPT '97. Vol. 1233. Lecture Notes in Computer Science. Springer, 1997, pp. 256–266. URL: https://doi.org/10.1007/3-540-69053-0_18.

[16] Minki Hhan, Takashi Yamakawa, and Aaram Yun. Quantum Complexity for Discrete Logarithms and Related Problems. 2023. URL: https://arxiv.org/abs/2307.03065.

[17] Euler's Totient Function – Growth Rate. Wikipedia. URL: https://en.wikipedia.org/wiki/Euler%27s_totient_function#Growth_rate (visited on 06/30/2026).

[18] Michael A. Nielsen and Isaac L. Chuang. Quantum Computation and Quantum Information. 10th anniversary. Cambridge University Press, 2010. URL: https://doi.org/10.1017/CBO9780511976667.

[19] Sevag Gharibian. Quantum Complexity Theory. Lecture notes, Paderborn University. 2024. URL: https://groups.uni-paderborn.de/fg-qi/data/QCT_Masterfile.pdf.

[20] Malus' Law. TikZ.net. URL: https://tikz.net/malus-law/ (visited on 06/30/2026).

[21] Izaak Neutelings. Gas Particles and Macrostates in TikZ. TikZ.net. URL: https://tikz.net/gas/ (visited on 06/30/2026).

[22] cfr and Stack Exchange Community. Drawing an Annotated Cuboid in TikZ. TeX Stack Exchange. URL: https://tex.stackexchange.com/a/288101 (visited on 03/25/2026).

[23] Izaak Neutelings. Black-Body Radiation Colors in TikZ. TikZ.net. URL: https://tikz.net/blackbody_plots/ (visited on 06/30/2026).

1 What is Quantum Informatics?

Information Theory (Shannon) ⟷ Quantum Theory, Thermodynamics, Relativity

Landauer (1961): "Information is physical" →

← Wheeler (1983): "It from bit"

1.1 Stern/Gerlach Experiment

Measure *magnetic dipole moment* of silver atoms by sending a stream of them through an *inhomogeneous* magnetic field. Each atom is deflected from the path proportionally to its dipole in the direction of the field gradient.

1.1.1 Classical expectation

If we imagine that the moments of the atoms point in random directions, then the classical expectation is that the deflection pattern reaches a maximum in the middle (no deflection) and then symmetrically, monotonically and continuously decreases on the sides.

Classical mechanics. Angular momentum (pseudovector):

$$\boldsymbol{L} = \boldsymbol{r} \times \boldsymbol{p} = \boldsymbol{r} \times (m\,\boldsymbol{v})$$

Torque:

$$\boldsymbol{\tau} = \boldsymbol{r} \times \boldsymbol{F}$$

Relationship:

$$\boldsymbol{\tau} = \frac{\mathrm{d}}{\mathrm{d}t}\,\boldsymbol{L} \tag{1.1}$$

Classical electrodynamics. A charge q moving in a circle of radius R creates a current:

$$I = \frac{q}{T} \quad \text{where} \quad T = \frac{2\pi R}{v}$$

Magnetic dipole moment of a current loop:

$$\boldsymbol{\mu} = I\,A\,\hat{n} \quad \text{where} \quad A = \pi R^2 \quad \text{and} \quad \hat{n} \text{ is normal to the area.}$$

Thus

$$\boldsymbol{\mu} = \frac{q}{2\pi R/v}\,(\pi R^2)\,\hat{n} = \frac{q\,v\,R}{2}\,\hat{n}$$

Angular momentum:

$$\boldsymbol{L} = m\,\boldsymbol{r} \times \boldsymbol{v} = m\,R\,v\,\hat{n}$$

Therefore (the bridge):

$$\boxed{\boldsymbol{\mu} = \gamma\,\boldsymbol{L}} \tag{1.2}$$

with

$$\gamma = \frac{q}{2m}$$

Magnetic dipole dynamics. Torque:

$$\boldsymbol{\tau} = \boldsymbol{\mu} \times \boldsymbol{B} \tag{1.3}$$

Combining (1.1), (1.2) and (1.3) gives the precession equation:

$$\frac{\mathrm{d}}{\mathrm{d}t}\,\boldsymbol{L} = \gamma\big(\boldsymbol{L} \times \boldsymbol{B}\big) \quad \text{or equivalently} \quad \frac{\mathrm{d}}{\mathrm{d}t}\,\boldsymbol{\mu} = \gamma\big(\boldsymbol{\mu} \times \boldsymbol{B}\big)$$

Hence $\frac{\mathrm{d}\boldsymbol{\mu}}{\mathrm{d}t} \perp \boldsymbol{\mu}$, and therefore $|\boldsymbol{\mu}|$ is constant (only the direction changes).

Energy:

$$U = -\,\boldsymbol{\mu} \cdot \boldsymbol{B}$$

Force in general is

$$\boldsymbol{F} = \nabla\big(\boldsymbol{\mu} \cdot \boldsymbol{B}\big)$$

and if $\boldsymbol{\mu}$ is constant and $\boldsymbol{B} \approx [0, 0, B_z(z)]$, then

$$F_z \approx \mu_z \,\frac{\partial B_z}{\partial z}$$

So, as mentioned above, according to classical physics we would expect to measure

$$\mu_z \in [-|\boldsymbol{\mu}|,\, +|\boldsymbol{\mu}|]$$

i.e. a continuous range and thus a continuous vertical smear.

1.1.2 Observation

That is, however, not what was observed: there is no detection in (not even close to) the middle, but rather two sharp peaks at equal distances from the middle.

In quantum mechanics, $\boldsymbol{L}$ becomes an operator with three components $\hat{L}_x, \hat{L}_y, \hat{L}_z$ and they do not commute, e.g.

$$[\hat{L}_x, \hat{L}_y] := \hat{L}_x\hat{L}_y - \hat{L}_y\hat{L}_x = \imath\hbar\hat{L}_z$$

This means we cannot simultaneously know all three components exactly. What is possible is to simultaneously know one component and the total angular momentum operator:

$$\hat{L}^2 = \hat{L}_x^2 + \hat{L}_y^2 + \hat{L}_z^2$$

$\hat{L}^2$ has eigenvalues $\ell(\ell+1)\hbar^2$:

$$\hat{L}^2\,|\ell, m_\ell\rangle = \ell(\ell+1)\,\hbar^2\,|\ell, m_\ell\rangle$$

Projection quantization:

$$\hat{L}_z\,|\ell, m_\ell\rangle = m_\ell\,\hbar\,|\ell, m_\ell\rangle$$

with $m_\ell = -\ell, -\ell+1, \ldots, \ell-1, \ell$.

Even when orbital $\ell = 0$, particles can still have angular momentum (spin) $\boldsymbol{S}$. It also satisfies $\hat{S}^2 = \hat{S}_x^2 + \hat{S}_y^2 + \hat{S}_z^2$ and

$$\hat{S}^2\,|s, m_s\rangle = s(s+1)\,\hbar^2\,|s, m_s\rangle, \quad \hat{S}_z\,|s, m_s\rangle = m_s\,\hbar\,|s, m_s\rangle$$

For the electron, for example, $m_s = \pm\frac{1}{2}$ so $\hat{S}_z$ has eigenvalues $\pm\,\hbar/2$.

Magnetic moment in quantum mechanics:

$$\boldsymbol{\mu} = \boldsymbol{\mu}_L + \boldsymbol{\mu}_S$$

where $\boldsymbol{\mu}_L = \frac{q}{2m}\,\boldsymbol{L}$ and $\boldsymbol{\mu}_S = g\,\frac{q}{2m}\,\boldsymbol{S}$ (for the electron: $g \approx 2$).

For silver, the valence electron is in an s-orbital ($\ell = 0$), so $\boldsymbol{\mu}_L \approx \boldsymbol{0}$. The remaining contribution is spin: $s = \frac{1}{2}$ and $m_s = \pm\frac{1}{2}$, hence μ_z can take only two values. This is the reason for the two peaks.

1.1.3 Repeated measurements

If we split along z, take the z_+ beam, and split again along z, all atoms are deflected into the positive z direction. In this sense the Z-spin looks classical; it is *stable* with respect to repeated measurements.

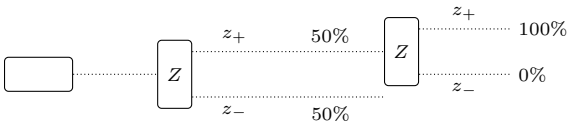

When the second magnet is rotated into the x-direction, again a 50–50 distribution occurs, which is not very surprising, since it just means that the two properties, Z-spin and X-spin, look *independent*.

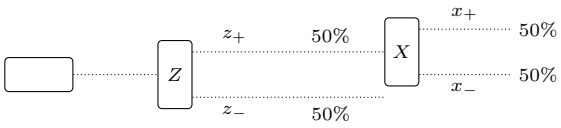

A surprising outcome occurs when the measurements are cascaded as follows:

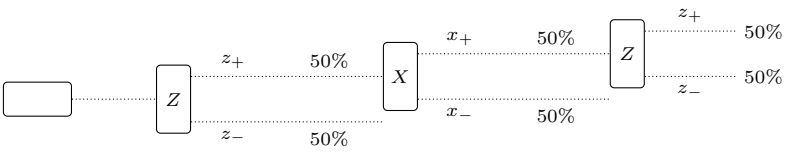

This questions both the *stability* and the *independence* of the spin property.

1.1.4 Superposition

The statistics found were surprising for particles, but they would not have been surprising for waves. With a beam of light and a polarizing filter, a similar thing happens. The angles are a bit different though, because classical polarization intensities obey Malus' law,

$$\frac{I}{I_0} = \cos^2(\varphi)$$

where φ is the angle between the intial polarization direction and the polarizer axis:

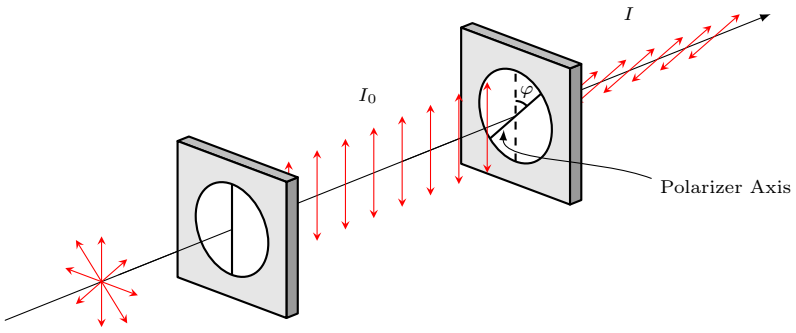

For a spin-$\frac{1}{2}$ particle the transition probability between "up" along two axes separated by an angle θ is

$$P = \cos^2 \frac{\theta}{2}$$

The extra factor of $\frac{1}{2}$ in the angle appears because spin states are *spinors* (an SU(2) representation): a physical rotation by θ acts on the state with half the angle in the two-dimensional Hilbert space.

The essential property of waves is that they can be linearly combined. Quantum mechanical states have the same property: they are elements in a vector space and we can derive probability distributions from them by squaring.

If, after a Z measurement, we perform an X measurement, then we want to know wether the silver atom is in a state $|x_+\rangle$ or $|x_-\rangle$. Both are superpositions:

$$|x_+\rangle = \frac{1}{\sqrt{2}}|z_+\rangle + \frac{1}{\sqrt{2}}|z_-\rangle \qquad |x_-\rangle = \frac{1}{\sqrt{2}}|z_+\rangle - \frac{1}{\sqrt{2}}|z_-\rangle$$

No matter the outcome of the first measurement, the X measurement yields one of both results with equal probability. Also in the other direction, if we first measure X and then Z, we get the same statistics, independent of any measurements before the X measurement.

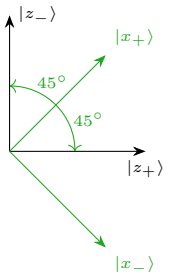

1.2 Quantum Key Distribution

If a quantum system is measured twice in the same basis, the same result is obtained with certainty. However, an intermediate measurement in a different basis disturbs the state and can change the outcome. So the interactions of a system with its environment become traceable. This can be used to detect an eavesdropper in a cryptographic key agreement protocol.

In 1984, Gilles Brassard and Charles Bennett developed the first application of quantum mechanics for cryptographic purposes with such a key agreement protocol (BB84).

Assume Alice and Bob want to agree on a secret key; e.g. for *symmetric encryption* (same key for encryption and decryption). They can proceed as follows.

1. Alice repeatedly measures a random bit in a random basis (e.g., Z or X) and sends the quantum system to Bob.
2. Bob measures the received system in a randomly chosen basis (Z or X).

For the bits where they used the same basis, the measurements must be the same, unless there has been an eavesdropper, Eve, measuring the system in a different basis during its transmission from Alice to Bob. Since Eve does not know the true basis, the best thing to do for her is to guess it, so her probability of guessing right is 50%. If, she guesses wrong, the bit Bob recieves will be incorrect with a chance of 50%, so in total, 25% of the bits will be wrong.

3. Alice and Bob communicate over an *public authenticated classical channel* which basis they used, without sharing the measurement results.
4. They randomly select some of the shared bits, reveal them publicly and check if they match.
5. If the error rate for those selected bits is high (around 25%), Eve is present, and they discard the key.
6. If the error rate is very low: No eavesdropping detected
7. They keep the remaining unrevealed bits as the secret key.

1.3 Mach Zehnder Interferometer

can be considered a variant of double-slit experiment.

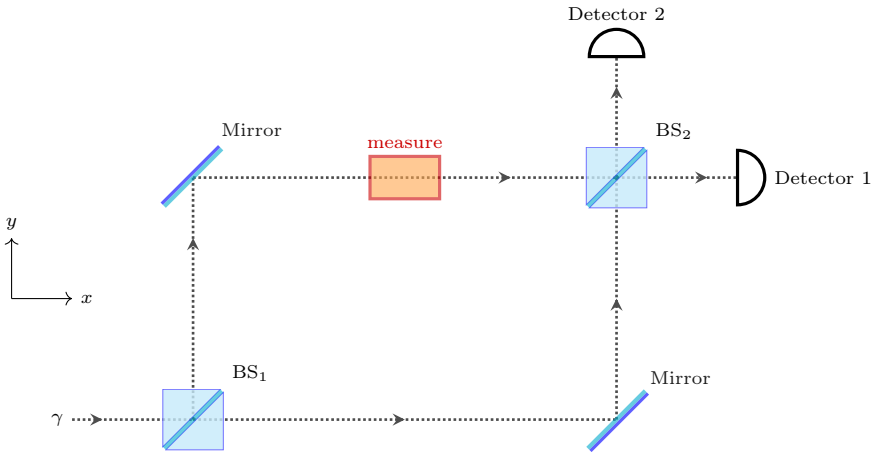

A symmetric 50/50 beam splitter acts on the state of the photon with

$$|x\rangle \mapsto \tfrac{1}{\sqrt{2}}|x\rangle + \imath\tfrac{1}{\sqrt{2}}|y\rangle$$
$$|y\rangle \mapsto \tfrac{1}{\sqrt{2}}|y\rangle + \imath\tfrac{1}{\sqrt{2}}|x\rangle$$

where the $\imath$ accounts for the 90° phase shift ($\frac{\pi}{2}$) a reflected photon picks up compare to the transmitted one. The effect of the mirrors is

$$|x\rangle \mapsto e^{\imath\phi_m}|y\rangle$$
$$|y\rangle \mapsto e^{\imath\phi_m}|x\rangle$$

where ϕ_m is the phase shift the photon picks up at the mirror (e.g. 0, $\frac{\pi}{2}$, π, depending on the physical implementation).

So as the photon moves through the device, its state evolves as follows

$$|x\rangle \xrightarrow{\text{BS1}} \frac{1}{\sqrt{2}}(|x\rangle + \imath|y\rangle) \xrightarrow{\text{Mirrors}} \frac{1}{\sqrt{2}}\left(e^{\imath\phi_m}|y\rangle + \imath e^{\imath\phi_m}|x\rangle\right) \xrightarrow{\text{BS2}} e^{\imath\left(\phi_m+\frac{\pi}{2}\right)}|x\rangle$$

and therefore it will be measured at Detector 1 with certainty.

If, however, one *measures* wether the photon has gone through the upper or lower arm of the interferometer, then the state before the second beam splitter is either $|x\rangle$ or $|y\rangle$. In the second beam splitter the state of the photon is then mapped to either $\frac{1}{\sqrt{2}}(|x\rangle + \imath|y\rangle)$ or $\frac{1}{\sqrt{2}}(|y\rangle + \imath|x\rangle)$. In both cases, the photon is detected in either of the detectors with equal probability.

Application 1.1 (Elitzur-Vaidman bomb tester [2])**.** A bomb in a box is very sensitive to light: it will explode as soon as it is hit by a photon. We would like to figure out whether the bomb is in the box without exploding. Towards this end, we repeat the process without measuring, so that the photon will pass through n beam splitters, and through the box with the bomb $n-1$ times. Moreover, we change the opacity of the beam splitters such that they change the state as follows:

$$U_\theta|B\rangle = \cos(\theta)|B\rangle + \imath\sin(\theta)|T\rangle$$
$$U_\theta|T\rangle = \cos(\theta)|T\rangle + \imath\sin(\theta)|B\rangle$$

where $|B\rangle$ and $|T\rangle$ denote the photon in the bottom and top paths. The ordinary 50/50 splitter from above corresponds to $\theta = \frac{\pi}{4}$, while the improved test uses n beam splitters with $\theta = \frac{\pi}{2n}$, arranged as follows:

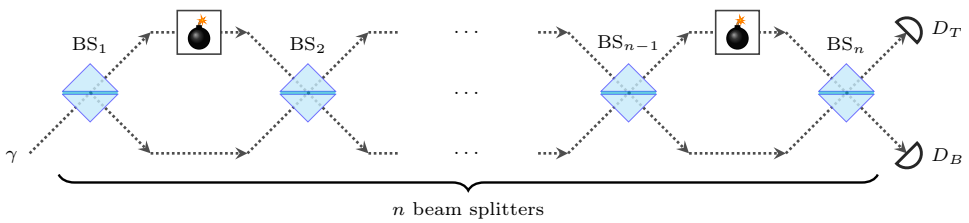

There are n beam splitters and $n-1$ passages through the box. Assume first that the bomb is present. Let S_{k-1} be the event that no explosion occurred before the k-th passage through the box, and let E_k be the event that the bomb explodes at the k-th passage, for $k = 1, \dots, n-1$. Conditioned on S_{k-1}, the photon has just been projected back to $|B\rangle$; after the next beam splitter,

$$|B\rangle \mapsto \cos\Big(\frac{\pi}{2n}\Big)\,|B\rangle + \imath \sin\Big(\frac{\pi}{2n}\Big)\,|T\rangle$$

Therefore

$$\mathbb{P}(E_k \mid S_{k-1}) = \sin^2\Big(\frac{\pi}{2n}\Big) \le \Big(\frac{\pi}{2n}\Big)^2 = \frac{\pi^2}{4n^2}$$

where we used $\sin x \le x$ for $x \ge 0$. This gives a bound on the unconditional probability as well, because $E_k \subseteq S_{k-1}$, i.e. the bomb can explode at the k-th passage only if there was no explosion before that passage. Thus

$$\mathbb{P}(E_k) = \mathbb{P}(S_{k-1})\,\mathbb{P}(E_k \mid S_{k-1}) \le \mathbb{P}(E_k \mid S_{k-1}) \le \frac{\pi^2}{4n^2}$$

Hence, by the union bound,

$$\begin{aligned}\mathbb{P}(\text{explode}) &= \mathbb{P}\left(\bigcup_{k=1}^{n-1} E_k\right)\\ &\le \sum_{k=1}^{n-1} \mathbb{P}(E_k)\\ &\le \sum_{k=1}^{n-1} \frac{\pi^2}{4n^2} = \frac{(n-1)\pi^2}{4n^2} \le \frac{\pi^2}{4n} = \mathcal{O}\Big(\frac{1}{n}\Big)\end{aligned}$$

Conditioned on no explosion, each passage through the box projected the state back to $|B\rangle$. Thus before the final beam splitter the state is $|B\rangle$, and after the final beam splitter it is

$$\cos\Big(\frac{\pi}{2n}\Big)\,|B\rangle + \imath \sin\Big(\frac{\pi}{2n}\Big)\,|T\rangle$$

Consequently

$$\mathbb{P}(D_B \mid \text{bomb present, no explosion}) = \cos^2\Big(\frac{\pi}{2n}\Big) = 1 - \sin^2\Big(\frac{\pi}{2n}\Big) \ge 1 - \frac{\pi^2}{4n^2}$$

If the bomb is not present, no intermediate measurement occurs and the n small rotations add coherently:

$$U_\theta^n\,|B\rangle = \cos(n\theta)\,|B\rangle + \imath \sin(n\theta)\,|T\rangle = \cos\Big(\frac{\pi}{2}\Big)\,|B\rangle + \imath \sin\Big(\frac{\pi}{2}\Big)\,|T\rangle = \imath\,|T\rangle$$

Hence $\mathbb{P}(D_B \mid \text{no bomb}) = 0$. ◀

1.4 Quantum Bit

To transfer a bit $b \in \{0, 1\}$ into the quantum world, we associate 0 and 1 with two orthogonal vectors, usually with the standard basis vectors,

$$|0\rangle = \begin{bmatrix}1\\0\end{bmatrix} \qquad |1\rangle = \begin{bmatrix}0\\1\end{bmatrix}$$

A general quantum state can then be written as a superpositions

$$|\psi\rangle = \alpha\,|0\rangle + \beta\,|1\rangle$$

with $\alpha, \beta \in \mathbb{C}$ and $|\alpha|^2 + |\beta|^2 = 1$. Measuring $|\psi\rangle$ in the standard basis yields 0 with probability $|\alpha|^2$ and 1 with probability $|\beta|^2$.

Quantum circuits are composed of quantum gates, i.e, *unitary maps*.

An important one is the Hadamard gate,

$$\underline{\boldsymbol{H}} = \frac{1}{\sqrt{2}}\begin{bmatrix}1 & 1\\ 1 & -1\end{bmatrix}$$

which maps the basis states to superpositions

$$H\,|0\rangle = \frac{1}{\sqrt{2}}\,(|0\rangle + |1\rangle) \qquad H\,|1\rangle = \frac{1}{\sqrt{2}}\,(|0\rangle - |1\rangle)$$

and when applied twice, yields the basis states again, i.e., $H^2\,|0\rangle = |0\rangle$ and $H^2\,|1\rangle = |1\rangle$.

Another interesting gate F has coordinate matrix

$$\underline{\boldsymbol{F}} = \frac{1}{\sqrt{2\imath}} \begin{bmatrix} 1 & \imath \\ \imath & 1 \end{bmatrix}$$

because applying it twice yields the NOT-gate,

$$\underline{\boldsymbol{F}}^2 = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$$

which is why F is called "square root of NOT", $F = \sqrt{\mathrm{NOT}}$. There is no classical gate that yields the NOT-gate in this way.

2-Qbit system:

$$|\psi\rangle_{AB} = |a\rangle \otimes |b\rangle = |ab\rangle$$

1.5 The Aspect/Gisin/Zeilinger Experiments

A pair of Qbits (e.g. polarized photons) can exhibit correlations that cannot be explained as merely two independent systems plus shared classical randomness. This happens when the two Qbits are prepared in an *entangled* state.

A first example is $\frac{1}{\sqrt{2}}(|0\rangle \otimes |0\rangle + |1\rangle \otimes |1\rangle)$. If Alice and Bob both measure in the standard basis, they always obtain the same bit. This looks similar to a classical correlation, but quantum theory says that the individual outcomes are still random and only become fixed upon measurement.

Claim 1.1 (Einstein, Podolsky, Rosen, 1935)**.** Quantum theory is incomplete and must be augmented by "hidden parameters" completely determining the measurement outcomes of all alternative measurements. ◁

Almost 30 years later, Bell came up with the idea that Claim 1.1 can be tested by allowing Alice and Bob to choose between different measurement bases.

Claim 1.2 (Bell, 1964)**.** Claim 1.1 is in doubt: There exist quantum correlations that go beyond the explanatory power of shared classical information. ◁

For that purpose, consider the *singlet* state

$$|\Psi^-\rangle := \frac{|0\rangle \otimes |1\rangle - |1\rangle \otimes |0\rangle}{\sqrt{2}}$$

We rewrite the singlet in rotated measurement bases of Alice and Bob.

$$\begin{bmatrix} |0\rangle \\ |1\rangle \end{bmatrix} = \begin{bmatrix} \cos\alpha & -\sin\alpha \\ \sin\alpha & \cos\alpha \end{bmatrix} \begin{bmatrix} |0'\rangle \\ |1'\rangle \end{bmatrix}$$

Let Alice rotate her basis by an angle α_A and Bob by α_B. Substituting this into the singlet state yields

$$\begin{aligned} |\Psi^-\rangle &= \frac{1}{\sqrt{2}}\Big(|0\rangle_A \otimes |1\rangle_B - |1\rangle_A \otimes |0\rangle_B\Big) \\ &= \frac{1}{\sqrt{2}}\Big((\cos\alpha_A\,|0'_A\rangle - \sin\alpha_A\,|1'_A\rangle) \otimes (\sin\alpha_B\,|0'_B\rangle + \cos\alpha_B\,|1'_B\rangle) \\ &\qquad - (\sin\alpha_A\,|0'_A\rangle + \cos\alpha_A\,|1'_A\rangle) \otimes (\cos\alpha_B\,|0'_B\rangle - \sin\alpha_B\,|1'_B\rangle)\Big) \\ &= \frac{1}{\sqrt{2}}\Big(\sin(\alpha_A - \alpha_B)\,|0'_A\rangle \otimes |0'_B\rangle + \cos(\alpha_A - \alpha_B)\,|0'_A\rangle \otimes |1'_B\rangle \\ &\qquad - \cos(\alpha_A - \alpha_B)\,|1'_A\rangle \otimes |0'_B\rangle + \sin(\alpha_A - \alpha_B)\,|1'_A\rangle \otimes |1'_B\rangle\Big) \end{aligned}$$

Thus the probabilities depend only on the difference of the angles, $\Delta := \alpha_A - \alpha_B$:

$$\begin{aligned} \mathbb{P}(\text{same}) &= \mathbb{P}(0'_A 0'_B) + \mathbb{P}(1'_A 1'_B) = \sin^2(\Delta) \\ \mathbb{P}(\text{opposite}) &= \mathbb{P}(0'_A 1'_B) + \mathbb{P}(1'_A 0'_B) = \cos^2(\Delta) \end{aligned}$$

In particular, if both parties measure in the same basis $\Delta = 0$ then $\mathbb{P}(\text{opposite}) = 1$, so their outcomes are perfectly anti-correlated.

For Bell's experiment, Alice chooses between the standard basis and a basis rotated by $30°$, while Bob chooses between the standard basis and a basis rotated by $-30°$.

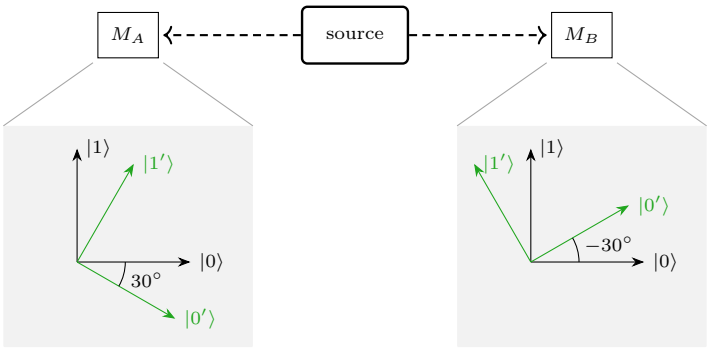

Thus the angle between the bases for the pairs (A_1, B_2) and (A_2, B_1) is $30°$, while the angle between (A_2, B_2) is $60°$.

Assume now, in the spirit of Claim 1.1, that the outcomes of all possible measurements are already fixed in the preparation stage.

Then Alice has two potential output bits A_1, A_2 and Bob analogously B_1, B_2. Here index 1 denotes the standard basis, while index 2 denotes the rotated basis.

In any single run each party reveals only one of these bits, but the hidden-variable assumption says that all four bits already have predetermined values.

For convenience, we flip Bob's output bit, so that equal bases lead to correlation instead of anti-correlation. Therefore, $\mathbb{P}(A_1 = B_1) = 1$, $\mathbb{P}(A_1 = B_2) = \cos^2(30°) = \frac{3}{4}$, $\mathbb{P}(A_2 = B_1) = \cos^2(30°) = \frac{3}{4}$.

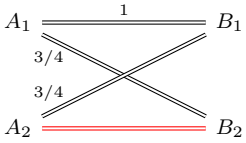

Under the hidden-variable assumption, the four bits A_1, A_2, B_1, B_2 exist simultaneously. If $A_2 \neq B_2$, then at least one of the equalities $A_1 = B_1$, $A_1 = B_2$, or $A_2 = B_1$ must fail. Hence we obtain (the simplest example of) a so-called Bell inequality

$$\mathbb{P}(A_2 \neq B_2) \leq \mathbb{P}(A_1 \neq B_1) + \mathbb{P}(A_1 \neq B_2) + \mathbb{P}(A_2 \neq B_1) = 0 + \frac{1}{4} + \frac{1}{4} = \frac{1}{2}$$

by the transitivity of equality and the union bound.

But quantum theory predicts

$$\mathbb{P}(A_2 = B_2) = \cos^2(60°) = \frac{1}{4}$$

and therefore

$$\mathbb{P}(A_2 \neq B_2) = \frac{3}{4}$$

which violates the Bell inequality.

Thus the observed correlations cannot be explained by pre-shared classical information.

Quantum correlations are stronger than any correlations obtainable from local hidden variables.

This phenomenon is called *Bell non-locality*. It shows that entanglement produces correlations that arise only upon measurement and cannot be understood as predetermined data carried by the particles.

This challenges the "Reichenbach principle":

> Any correlation between two observed events must arise either from a common cause in their shared past, or from a direct causal influence between the events.

Bell's result rules out the first explanation. The correlations cannot be explained by shared classical information created in the common past.

The second explanation would contradict the spirit of relativity theory, since the speed of that influence would have to be highly superluminal.

Experiments have repeatedly confirmed the predictions of quantum theory. In some experiments the measurements are spacelike separated, meaning that in each observer's reference frame their own measurement occurs first.

Remark 1.2 (No Signalling)**.** Whatever the observable that Alice (or Bob, respectively) measures on her (his) part of the singlet: The probability for both results is 1/2. In particular, Alice's measurement result does not depend on measurements on Bob's side and vice versa. Thus, Alice and Bob cannot transmit messages this way, possibly faster than light. This saves us from problems with relativity. ◀

To study how strong the correlations can be, Popescu and Rohrlich proposed an idealized device called a *PR box*:

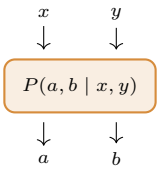

It has two inputs $x, y \in \{0, 1\}$ and two outputs $a, b \in \{0, 1\}$ with the constraint

$$a \oplus b = xy$$

This gives four equations

$$\begin{aligned} a_1 \oplus b_1 &= x_1 y_1 \\ a_1 \oplus b_2 &= x_1 y_2 \\ a_2 \oplus b_1 &= x_2 y_1 \\ a_2 \oplus b_2 &= x_2 y_2 \end{aligned}$$

which, when added modulo 2 yield

$$0 = (x_1 \oplus x_2)(y_1 \oplus y_2)$$

But the RHS equals 1, since $x_1 \neq x_2$ and $y_1 \neq y_2$. Therefore the system cannot be satisfied simultaneously.

At most three out of the four equations can hold. Thus any classical strategy can succeed with probability at most $\frac{3}{4}$.

2 Information is Physical

Quantum physics describes fundamental particles and their interactions, which ultimately involve the exchange and transformation of energy. Even mass itself is a form of energy. However, energy alone does not fully describe a physical system. We also need information: a specification of the system's possible states. Thus, physical reality is characterized by the interplay between energy and information.

To understand Law 2.2, we need a precise notion of entropy. It is often described as a "measure of disorder", but this interpretation is subjective and can be misleading. A more precise viewpoint is that entropy quantifies the amount of information required to describe a system.

In the 19$^{\text{th}}$ century, Ludwig Boltzmann connected thermodynamics with the microscopic behavior of particles. He showed that entropy is related to the number of possible microscopic configurations (microstates) corresponding to a macroscopic state (macrostate). If Ω denotes this number, then

$$\boxed{S = k_{\text{B}} \ln(\Omega)}$$

which is the famous equation that is engraved on his tombstone. Thus, systems with many possible configurations have high entropy, since specifying their exact state requires more information.

Law 2.1 (First law of thermodynamics: Energy)**.** In a closed system, the total energy is constant. ◁

Law 2.2 (Second law of thermodynamics: Entropy)**.** In a closed system, the total entropy ("disorder") does not decrease. ◁

A remarkable feature of Law 2.2 is that it is *asymmetric* in time, whereas, e.g., the laws of classical mechanics are time-symmetric.

Paradox 2.1 (Maxwell's demon)**.** Imagine a creature sitting at a separation wall in a gas container with a frictionless door. The demon can open whenever a molecule moves towards the door. In linear time in the numbrer of particles the gas will be "sorted", i.e., compressed - revoltingly unfaithful to Law 2.2. ◀

The understanding of what is wrong with this argument is one of the first success stories of the synergy between physics and

2.1 Information Theory

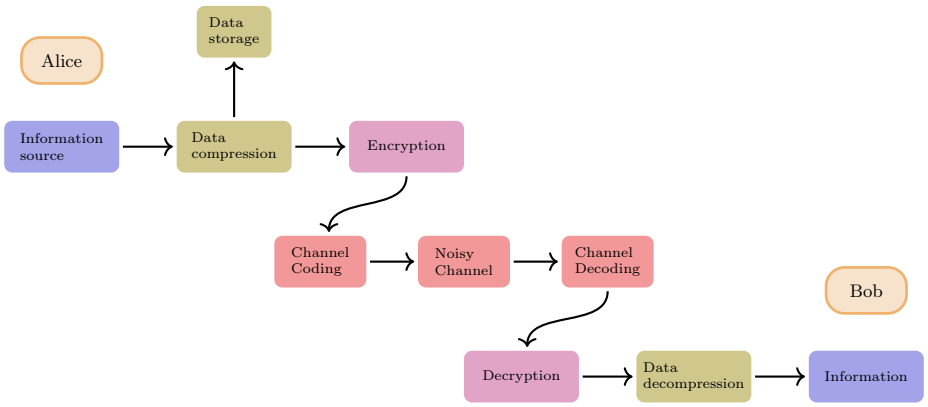

Definition 2.2 (Entropy)**.** The entropy (or uncertainty) of a random variable X over $\mathcal{X}$ with PMF $p_X(\cdot)$ is

$$\boxed{H(X) \coloneqq \mathbb{E}[-\log_2 p_X(X)] = -\sum_{x \in \mathcal{X}} p_X(x) \log_2 p_X(x)}$$

◀

Entropy describes the uncertainty of a random experiment with given PMF.

An important observation is that the entropy of X does not depend on the actual values of X, but only on the probabilities $p_X(x)$. Hence, Definition 2.2 can easily be extended to random vectors, by just replacing X with $\boldsymbol{X}$ and x with $\boldsymbol{x}$ in the equation.

Definition 2.3 (Mutual information)**.** The mutual information between two random variables X, Y is

$$I(X;Y) \coloneqq H(X) - H(X \mid Y) = H(Y) - H(Y \mid X)$$ ◀

Another nice form for the mutual information is

$$I(X;Y) = H(X) + H(Y) - H(X,Y)$$

because it shows its symmetry.

Definition 2.4 (Kullback-Leibler divergence)**.** Let $P(\cdot)$ and $Q(\cdot)$ be two PMFs over the same finite or countable alphabet $\mathcal{X}$. The *relative entropy* or *Kullback-Leibler divergence* between P and Q is

$$D_{\mathrm{KL}}(P \parallel Q) \coloneqq \mathbb{E}_{X\sim P}\left[\log_2 \frac{P(X)}{Q(X)}\right] = \sum_{x\in\mathcal{X}} P(x) \log_2 \frac{P(x)}{Q(x)}$$ ◀

Unlike Mutual information, Kullback-Leibler divergence is not symmetric. It should not be thought of as a distance, but rather as an "energy" (it actually describes the inefficiency of assuming that the PMF is $Q(\cdot)$ when the true distribution is $P(\cdot)$.).

Mutual information and Kullback-Leibler divergence are related to one another via

$$I(X;Y) = D_{\mathrm{KL}}(p_{X,Y} \parallel p_X \cdot p_Y)$$

Mutual information can therefore be interpreted as deviance of the actual joint distribution from if they were treated independently (while keeping the same marginals).

Example 2.1. In case of a uniform distribution, i.e. $p_X(x) = 1/|\mathcal{X}|$ for all $x \in \mathcal{X}$, the entropy of X is the logarithm of the size of its range, i.e., $H(X) = \log_2 |\mathcal{X}|$. ◀

2.2 Connection to Thermodynamics

Example 2.1 can be carried over to thermodynamics: "The entropy of a macrostate is the logarithm of the number of microstates corresponding to it." The *microstate* of a physical system, e.g., a gas, specifies the position and momentum of each molecule. The *macrostate* is simply a set of microstates, typically with a short description (specifying, e.g., the gas type, volume, temperature and pressure).

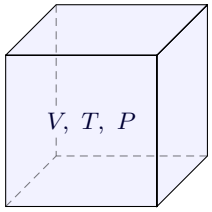

Macrostate: coarse description of the system by macroscopic quantities such as volume, temperature, and pressure.

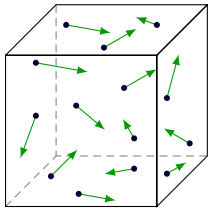

Microstate: complete specification of the positions and momenta of all individual molecules.

If all microstates in a macrostates "look roughly the same" (this condition is the translation of the uniformity condition from Example 2.1), we have

$$H(\text{macrostate}) \approx \log_2(\#\text{ corresponding microstates})$$

The "physical entropy", as opposed to the information-theoretic one, is usually denoted by S and has an additional factor of $k_{\mathrm{B}} \ln(2)$, where $k_{\mathrm{B}} \approx 1.38 \cdot 10^{-23}$Joule/Kelvin is the Boltzmann constant and the $\ln(2)$ is a common view at the border between nature and abstraction since 2 is a "logical" constant and e is a "natural" constant. If we denote $\Omega \coloneqq \#$ microstates, we obtain the famous formula

$$S = k_{\mathrm{B}} \ln(2) H = k_{\mathrm{B}} \ln(\Omega)$$

derived by Boltzmann in the 1870s and engraved on his tombstone.

Example 2.2 (1-Molecule Gas)**.** Consider a "one-molecule gas" in a container of volume V. The first macrostate corresponds to the particle being anywhere in the container, where for the second, smaller macrostate, the particle is confined to the left half of the container.

Remark: We cannot simply *count* the number of possible positions and momenta of the particle, since it is (uncountably) infinite for both macrostates. Indeed, there also exist definitions of entropy for the continuous case, but we do not need that here as long as we are only interested in entropy *differences*: Assume that we distinguish different microstates only up to some fixed finite precision, i.e., we put a grid of certain "density" (e.g. cubes with side length δ) into the respective volumes and the particles position is described by the coordinates of the cube it is in.

Then, no matter the actual value of δ, there are exactly twice as many microstates corresponding to the first macrostate than to the second. If we denote the number of microstates corresponding to the first macrostate by Ω_1 and the number of microstates corresponding to the second macrostate by $\Omega_2 = \Omega_1/2$, we have, the entropy difference thus is

$$\Delta H = H(\text{macrostate 2}) - H(\text{macrostate 1}) = \log_2(\Omega_2) - \log_2(\Omega_1) = \log_2\left(\frac{\Omega_2}{\Omega_1}\right) = -1$$

or, in terms of physical entropy, $\Delta S_{\text{gas}} = k_{\text{B}} \ln(2) \Delta H = -k_{\text{B}} \ln(2)$. ◀

Example 2.3 (N-Molecule Gas)**.** Consider a gas of N molecules in a container of volume V. We regard the gas as the *system* and the surroundings as a heat reservoir (the *environment*) at temperature T.

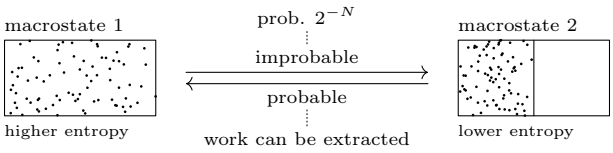

Compare the following two macrostates:

- macrostate 1: all molecules may be anywhere in the container (volume V)
- macrostate 2: all molecules are confined to the left half of the container (volume $V/2$)

As in the one-molecule case, we discretize position space into cells of volume δ^3. Then the number of position microstates for a single molecule is proportional to the accessible volume. Since the momentum space is the same in both macrostates, it cancels in the ratio, so it suffices to consider position space.

Assuming the molecules are non-interacting so that their positions are independent, the total number of microstates is $\Omega_1 \propto V^N$, $\Omega_2 \propto (V/2)^N$. Therefore, the entropy difference is

$$\Delta H = \log_2(\Omega_2) - \log_2(\Omega_1) = \log_2\left(\frac{\Omega_2}{\Omega_1}\right) = \log_2\left(2^{-N}\right) = -N$$

or, in physical units, $\Delta S_{\text{gas}} = k_{\text{B}} \ln(2)\, \Delta H = -N k_{\text{B}} \ln(2)$. This transition (compressing all molecules into half the volume) is highly improbable, with probability 2^{-N}, but it can be enforced by investing work, e.g., by using a piston to compress the gas. We now interpret this entropy decrease thermodynamically and relate it to the energy required to enforce the compression. The gas is then not an isolated system: work is performed on it and, after thermalization, energy is transferred as heat to the environment, increasing the entropy around the gas container, so there is no contradiction to Law 2.2.

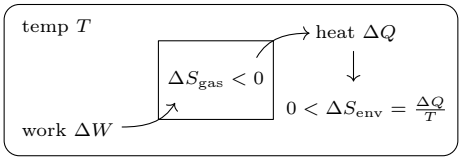

A certain amount of work ΔW is invested to compress the gas, which, after thermalization, is dissipated as heat ΔQ to the environment (Law 2.1). It is that heat transfer that is responsible for the increase of entropy in the environment, which is necessary to compensate for the decrease of entropy in the gas and thus to avoid a violation of Law 2.2.

For the environment, modeled as a heat reservoir at temperature T, the entropy change is

$$\Delta S_{\text{env}} = \frac{\Delta Q}{T}$$

where $\Delta Q > 0$ denotes the heat transferred to the environment. A heat reservoir is an idealized system that remains in internal equilibrium at fixed temperature T while exchanging heat with other systems.

Law 2.2 requires

$$\Delta S_{\text{tot}} = \Delta S_{\text{gas}} + \Delta S_{\text{env}} \geq 0$$

where equality holds only in the reversible case, i.e. without any additional entropy production (i.e. excluding irreversible processes such as heat flow across a finite temperature difference, friction, diffusion, turbulence, chemical reactions, ...). Thus, the minimal required work to enforce this entropy decrease is obtained in the reversible case:

$$\Delta W = \Delta Q = -T\Delta S_{\text{gas}} = Nk_{\text{B}}T\ln(2)$$

We obtain the same result by using the ideal gas law to compute the work needed to compress the gas from volume V to volume $V/2$:

$$pV = Nk_{\text{B}}T \quad \Longleftrightarrow \quad p(V) = \frac{Nk_{\text{B}}T}{V}$$

and thus

$$-\int_V^{V/2} p(V)\mathrm{d}V = -\int_V^{V/2} \frac{Nk_{\text{B}}T}{V}\mathrm{d}V = Nk_{\text{B}}T\ln(2)$$

◀

We can interpret the molecule in Example 2.2 as storing one bit of information, 0 if it is in the left halft and 1 if it is in the right half of the container. If we force the molecule to be in the left half, it corresponds to erasing that bit of information, in the sense that its state is 0 at the end, regardless of its initial state.

Claim 2.3 (Landauer's principle)**.** Erasing one bit - irrelevant by what physical system it is stored - requires at least

$$k_{\text{B}}T\ln(2) \quad (\approx 3\cdot 10^{-21}\text{Joule} \quad \text{at room temperature})$$

of free energy, which, in the process, must be dissipated as heat to the environment. ◁

2.3 Converse of Landauer's principle

The inverse process of erasure, let's call it *randomization*, allows for gaining free energy.

Example 2.4. The work value of the all-0-string of length N, however it is represented physically, is $Nk_{\text{B}}T\ln(2)$. This amount of environmental heat energy can be transformed into work. ◀

The Bennett work value of a general string s of length N is approximately

$$W \approx (N - K(s))\, k_{\text{B}}T\ln(2)$$

where $K(s)$ denotes the Kolmogorov complexity of the string s, i.e., the length of the shortest program that outputs s.

Claim 2.3 and its converse can be extended to b-ary digits: If, instead of binary digits, we work with b-ary digits, the work value of the all-0-string of length N is $Nk_{\text{B}}T\ln(b)$.

Example 2.5. The string π_N admits a short description: there exists a fixed program that computes the digits of π, together with an encoding of the number N specifying how many digits to output. Therefore,

$$K(\pi_N) = O(\log N)$$

since encoding the integer N requires about $\log N$ bits.

Therefore the Bennett work value of the first N decimal digits of π is approximately

$$W \approx \big(N\ln(10) - O(\log N)\big)\, k_{\text{B}}T$$

which, for large N, is essentially $Nk_{\text{B}}T\ln(10)$, i.e.,maximal. ◀

Example 2.5 shows that the work value is related to *reversible data compression*: Exaclty those strings have work value that can be compressed to a shorter length in a lossless fashion.

Example 2.6. A string of length N resulting from flips of an unfair coin with probabilities p and $(1-p)$ has, most likely, work value $(1-h(p))Nk_{\mathrm{B}}T\ln(2)$, where

$$h(p) \coloneqq -p\log_2 p - (1-p)\log_2(1-p)$$

is the binary entropy function. ◀

Example 2.7. r copies of the same, perhaps incompressible string of length N have work value at least $(r-1)Nk_{\mathrm{B}}T\ln(2)$. ◀

2.4 Benett's Solution to Maxwell's Demon

In the face of Landauer's principle, Paradox 2.1 disappears: The demon must erase the information it gathers during the sorting process. This erasure requires a minimal amount of work and leads to heat dissipation into the environment, which compensates exactly for the entropy decrease created in the gas.

The demon must maintain an internal memory storing measurement outcomes. Even a one-bit memory must be reset after each step, and overwriting information necessarily involves erasure. Thus, forgetting information has a thermodynamic cost. If the demon has a large memory, it can temporarily store many observations (effectively recording the initial disorder of the gas). However, to reuse its memory, it must eventually erase this information, and this erasure restores the total entropy balance.

Conversely, possessing information (or redundancy) allows one to extract work: knowing the state of a system (e.g. in Example 2.2, if the molecule is in the left half or the right half of the container) can be used to convert environmental heat energy into free mechanical energy. In this sense, information has an energetic value ("knowledge is energy").

Nevertheless, no perpetuum mobile of the second kind arises. Any work gained through such knowledge must ultimately be paid for when the demon resets its memory, ensuring consistency with Law 2.2.

2.5 Reversible Computing

If, in the course of a computation, information is lost about "which branch the computation came from", then free energy $k_{\mathrm{B}}T\ln(2)$ must be invested, which is dissipated as heat to the environment. In other words, the *logical* irreversibility of a computation ("information is lost") implies its *thermodynamic* irreversibility ("free energy is lost"). Thus, it appears advantageous if a computation does not have such collisions of branches.

As it turns out, it is (at least in principle) possible to carry out any logically reversible computation in a thermodynamically reversible way: Landauer's principle is "tight" in this sense.

Example 2.8. The ordinary AND gate is not reversible:

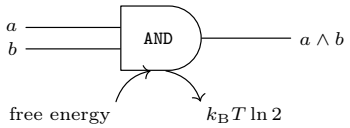

because, from the output 0, one cannot reconstruct whether the input was $(0,0)$, $(0,1)$, or $(1,0)$. The missing information has to be dumped into the environment; in the minimal case this costs the Landauer heat $k_{\mathrm{B}}T\ln 2$ per erased bit.

A first way to avoid this loss of information is to keep (copy) the inputs:

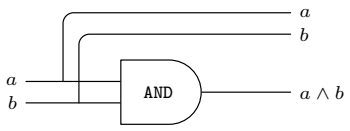

This is now injective, but the gate is still the same, so if before we found out that this gate will produce energy, it will still produce the same amount of energy, even if we keep the inputs. Moreover, we already saw in Example 2.7 that copying, since it creates redundancy, has a work value. Reversible gates should always have the same number of input and output wires. ◀

An ordinary deterministic Turing machine is still logically irreversible in general: "deterministic" merely means that the next configuration is uniquely determined by the present configuration, but it does not mean that the previous configuration can be reconstructed from the next one. Many different histories may lead to the same final output.

One can make such a logically irreversible computation reversible by giving the machine a "history tape" for storing the entire path of the computation.

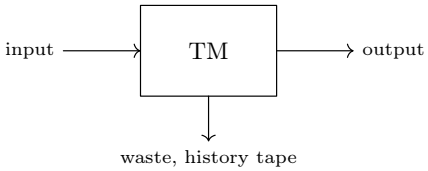

However, keeping the entire "history" of the computation is not "sustainable". The reason is that the original state of that history tape, say the all-0-string, is lost and replaced by the waste that has piled up — just as in Maxwell's demons's brain (see Section 2.4 and Paradox 2.1). (If the history tape was filled with a random string instead of the all-0-string, then ordinary writing to it would overwrite its previous contents. That overwriting is already a logically irreversible erasure, and hence carries a Landauer cost, wasting free energy as heat into the environment.)

Bennett's idea is to get rid of the waste in systematic way, i.e., to uncompute instead of erase it: Do the computation (including writing to the (typically long) history tape), copy the output to a separate (typically much shorter) output tape, and then undo computation step by step in reverse order (similar to Hänsel und Gretel retracing their steps in the forest).

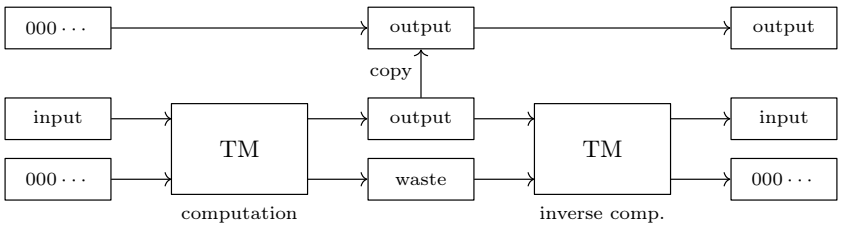

We have shown that every computation can be embedded into a logically reversible one: no information has to be thrown away.

Landauer's principle only tells us that logical irreversibility implies thermodynamic irreversibility, but it does not give us the other direction, i.e., it does not tell us if a logically reversible computation can be implemented in a thermodynamically reversible way.

In 1982, Fredkin and Toffoli showed exactly that, namely that, at least in principle, such logically reversible computations can be realized thermodynamically reversible: in a ballistic computer, bits are represented by moving balls, and gates by elastic collisions. They showed that this device can carry out any computation by a Turing machine — as long as no information is lost (that would be impossible due to the time-reversal symmetry of the laws of classical mechanics), i.e., as long as the computation is logically reversible.

Given these encouraging results, it was a natural question to ask what might be the basic building block of reversible computing (such as the NOT and AND gate for irreversible computing).

2.6 Toffoli gate

The Toffoli gate is obtained by making the AND gate reversible in the strict circuit sense.

Definition 2.5 (Toffoli gate)**.** The Toffoli gate T maps $(x, y, b) \longmapsto \big(x, y, b \oplus (x \wedge y)\big)$.

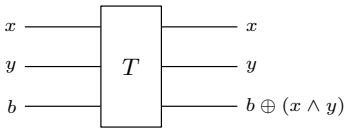

It is reversible because $T^{-1} = T$: applying it twice restores the original triple. ◀

By fixing some wires, the Toffoli gate contains the usual irreversible Boolean operations without becoming irreversible itself:

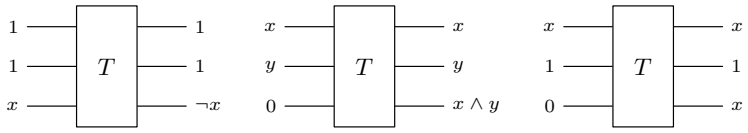

Thus Toffoli gives NOT, AND with a clean target bit, and FAN-OUT (copying of a classical bit). Since NOT and AND are functionally complete for classical computation, an ordinary irreversible circuit can be translated gate by gate into a reversible one using Toffoli gates. The price is that we have to add constant input wires, and at the end these wires generally do not contain constants any more. They contain intermediate results of the computation, analogous to the history tape in Bennett's reversible Turing-machine simulation. And just as in that case, leaving it like that is not yet "sustainable": if this intermediate information were simply reset, it would be erased and Landauer's cost would reappear.

Let $f : \{0,1\}^n \to \{0,1\}^m$ be computed by some irreversible circuit. A reversible circuit cannot in general map only $\boldsymbol{x}$ to $f(\boldsymbol{x})$, since this need not be injective. Instead, it embeds the computation into a bijection by adding ancillary input $\boldsymbol{c}$ and garbage output $J(\boldsymbol{x})$. Then it can be used in either direction:

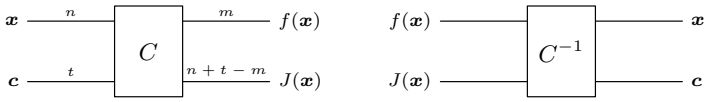

Here $\boldsymbol{c}$ is usually a fixed string of clean ancillas, and $J(\boldsymbol{x})$ is the garbage or junk information that makes the total map bijective. The inverse circuit C^{-1} works only if both $f(\boldsymbol{x})$ and $J(\boldsymbol{x})$ are supplied: the junk is exactly the information needed to reconstruct the input and the ancillas.

Bennett's trick removes this junk without erasing it. First compute with C, producing $f(\boldsymbol{x})$ together with all intermediate information. Then copy the useful output to a clean register, bit by bit, using reversible FAN-OUT/XOR operations. Finally run C^{-1}. Since C^{-1} receives the full pair $f(\boldsymbol{x}), J(\boldsymbol{x})$, it reverses the whole computation and restores the ancillas to their original clean state.

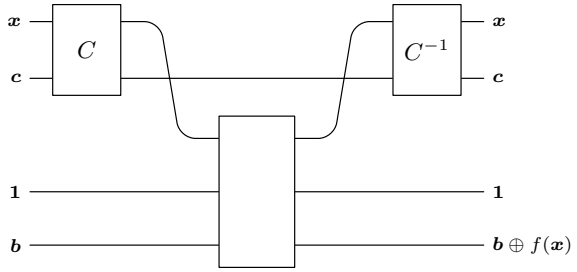

The middle box denotes the reversible copy/XOR operation controlled by the computed value $f(\boldsymbol{x})$. After uncomputation, the only lasting change is that the clean target $\boldsymbol{b}$ has become $\boldsymbol{b} \oplus f(\boldsymbol{x})$; the input $\boldsymbol{x}$ is still present, and the scratch space can be reused without being reset. Thus reversibility costs only a constant-factor overhead: compute, copy, uncompute.

There is one important special case. If f is bijective, then keeping $\boldsymbol{x}$ as part of the final output is not logically necessary: the value $f(\boldsymbol{x})$ already determines $\boldsymbol{x}$. Assume we have reversible circuits C and D implementing f and f^{-1} in the XOR form

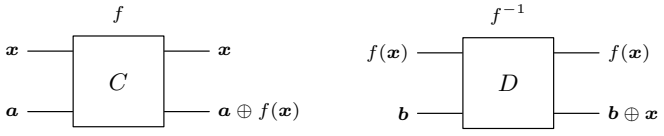

Note that D is *not* the inverse circuit C^{-1}: it is a circuit that computes the inverse function f^{-1} *without* needing the "computation history" $J(\boldsymbol{x})$ as part of the input. Now apply C to $(\boldsymbol{x}, \boldsymbol{0})$, swap the two middle wires, and feed them into D:

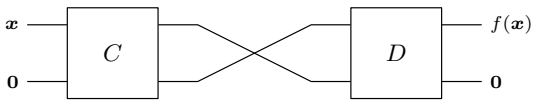

This circuit may sometimes be necessarily much less efficient than the best irreversible circuit for f. In fact, this is the case if and only if f is a "one-way function", a central notion in cryptography: A one-way function is a bijective function that can be computed efficiently, but for which no efficient inversion algorithm exists. Clearly, a function for which this reversible circuit is efficient cannot be one-way since that circuit can be used in both directions (Again, note that this does not apply for those reversible circuits for which $\boldsymbol{x}$ is again part of the output in the form of the junk $J(\boldsymbol{x})$). On the other hand, a function for which the reversible circuit is necessarily inefficient cannot have an efficient inversion algorithm.

3 From Classical to Quantum Physics

3.1 Black-Body Radiation

Consider an idealized cavity of side length l with perfectly absorbing walls, in thermal equilibrium with a heat bath at temperature T:

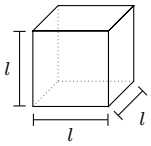

The cavity contains EM radiation. Because the field must satisfy boundary conditions at the walls, only certain standing waves are allowed. A standing wave in the cavity is then a superposition of three standing waves, each corresponding to one spatial dimension. Thus, it is described by

$$\boldsymbol{n} = [n_1, n_2, n_3] \in \mathbb{N}_{>0}^3$$

which specify the number of half-wavelengths that fit into the three spatial directions. For such a mode,

$$\boldsymbol{k} = [k_x, k_y, k_z] = \left[\frac{\pi n_1}{l}, \frac{\pi n_2}{l}, \frac{\pi n_3}{l}\right] = \frac{\pi}{l}\boldsymbol{n}$$

and therefore the wave vector has length $|\boldsymbol{k}| = \frac{\pi}{l}|\boldsymbol{n}|$. Since EM waves satisfy $\omega = c|\boldsymbol{k}|$, their angular frequency is

$$\omega = \frac{c\pi}{l}|\boldsymbol{n}| \tag{3.1}$$

where c denotes the speed of light.

Hence the allowed modes correspond to lattice points $\boldsymbol{n}$ in the first octant of $\mathbb{R}^3$, and the frequency is proportional to the distance of that lattice point from the origin.

Let $N(\omega)$ denote the number of EM modes with frequency at most ω. Then the number of modes in the interval $[\omega, \omega + \mathrm{d}\omega]$ is $\mathrm{d}N$. To estimate $\mathrm{d}N$, we count the lattice points $\boldsymbol{n}$ in a thin spherical shell between radii $|\boldsymbol{n}|$ and $|\boldsymbol{n}| + \mathrm{d}|\boldsymbol{n}|$. For large $|\boldsymbol{n}|$, we approximate this by the volume of that shell in $\boldsymbol{n}$-space. Since only the first octant is allowed, the shell volume is

$$\frac{1}{8} \cdot 4\pi|\boldsymbol{n}|^2 \,\mathrm{d}|\boldsymbol{n}| = \frac{\pi}{2}|\boldsymbol{n}|^2 \,\mathrm{d}|\boldsymbol{n}|$$

Moreover, each wave vector corresponds to two independent polarization states (the two transverse polarizations of the EM field), so the number of EM modes is

$$\mathrm{d}N = 2 \cdot \frac{\pi}{2}|\boldsymbol{n}|^2 \,\mathrm{d}|\boldsymbol{n}| = \pi|\boldsymbol{n}|^2 \,\mathrm{d}|\boldsymbol{n}| \tag{3.2}$$

From (3.1), we have $|\boldsymbol{n}| = \frac{l}{\pi c}\omega$ and $\mathrm{d}|\boldsymbol{n}| = \frac{l}{\pi c}\,\mathrm{d}\omega$. Substituting these into (3.2), we obtain

$$\mathrm{d}N = \pi\left(\frac{l}{\pi c}\omega\right)^2\left(\frac{l}{\pi c}\,\mathrm{d}\omega\right) = \frac{l^3}{\pi^2 c^3}\omega^2\,\mathrm{d}\omega$$

Dividing by the cavity volume l^3, we obtain the number of modes per unit volume in the frequency interval $[\omega, \omega + \mathrm{d}\omega]$,

$$\frac{\mathrm{d}N}{l^3} = g(\omega)\,\mathrm{d}\omega = \frac{\omega^2}{\pi^2 c^3}\,\mathrm{d}\omega$$

where $g(\omega) = \frac{\omega^2}{\pi^2 c^3}$ is the density of modes per unit volume per unit angular frequency.

3.1.1 Classical equipartition and the Rayleigh-Jeans law

In classical statistical mechanics, the equipartition theorem says that every quadratic degree of freedom contributes an average energy of $\frac{1}{2}k_\mathrm{B}T$. Each EM mode contributes two such quadratic terms (one electric and one magnetic), so the average energy per mode is of order $k_\mathrm{B}T$. This can also be seen directly from the Boltzmann distribution

$$\mathbb{P}(E) \propto e^{-E/(k_\mathrm{B}T)}$$

Classically, each EM mode behaves like a harmonic oscillator with continuous energy $E \in [0, \infty)$, and one obtains

$$\mathbb{E}[E] = \frac{\int_0^\infty E\, e^{-E/(k_\mathrm{B}T)} \mathrm{d}E}{\int_0^\infty e^{-E/(k_\mathrm{B}T)} \mathrm{d}E} = k_\mathrm{B}T$$

which is independent of the frequency ω. Thus even very high-frequency modes carry the same average energy $k_\mathrm{B}T$.

Therefore, the spectral energy density is obtained by multiplying the mode density by the average energy per mode:

$$u(\omega) = g(\omega) \cdot k_\mathrm{B}T = \frac{k_\mathrm{B}T}{\pi^2 c^3}\, \omega^2 \tag{3.3}$$

which is the Rayleigh-Jeans law. Thus the energy density grows quadratically without bound as the frequency increases.

3.1.2 Ultraviolet catastrophe

The total energy density is obtained by integrating over all frequencies,

$$\int_0^\infty u(\omega)\, \mathrm{d}\omega \overset{(3.3)}{=} \frac{k_\mathrm{B}T}{\pi^2 c^3} \int_0^\infty \omega^2\, \mathrm{d}\omega = \infty$$

Thus classical theory predicts an infinite amount of energy stored in the EM field inside the cavity.

This is the ultraviolet catastrophe. The term refers to the absurd prediction that most of the energy should be concentrated at very high frequencies (ultraviolet, X-ray, and beyond). If that were true, ordinary objects at room temperature would emit enormous amounts of high-energy radiation, which is not observed.

Experimentally, the Rayleigh-Jeans law (fortunately for us) is only correct for low frequencies. For high frequencies, the spectral energy density does *not* continue to grow quadratically; instead, it decreases rapidly.

3.1.3 Quantization

Max Planck resolved this problem by abandoning the classical assumption that energy can be exchanged continuously. He postulated that a mode of frequency ω can absorb or emit energy only in discrete packets,

$$E_n = n\hbar\omega$$

with $n \in \mathbb{N}_0$ and $\hbar = \frac{h}{2\pi} \approx 1.054 \cdot 10^{-34}\,\mathrm{J\,s}$ (reduced Planck constant). For black-body radiation, it is sufficient to consider the thermal excitation energies $n\hbar\omega$, i.e. we ignore the zero-point offset $\frac{1}{2}\hbar\omega$, since it does not affect the frequency-dependent thermal spectrum.

The Boltzmann factor is still the same as in the classical case,

$$\mathbb{P}(E) \propto e^{-E/(k_\mathrm{B}T)}$$

but now the allowed energies are no longer continuous. Instead of all $E \in [0, \infty)$, only the discrete values $0, \hbar\omega, 2\hbar\omega, \ldots$ are possible.

This changes the statistics drastically. For large ω, even the first excited state already has the large energy $\hbar\omega$, so nonzero occupation of the mode is strongly suppressed by the Boltzmann factor. Thus the average energy per mode decreases for large ω.

In fact, the average energy of one mode becomes

$$\mathbb{E}[E] = \frac{\sum_{n=0}^\infty n\hbar\omega\, e^{-n\hbar\omega/(k_\mathrm{B}T)}}{\sum_{n=0}^\infty e^{-n\hbar\omega/(k_\mathrm{B}T)}} = \frac{\hbar\omega}{e^{\hbar\omega/(k_\mathrm{B}T)} - 1}$$

which is *not* independent of ω anymore! For small ω, however, this is still approximately $k_\mathrm{B}T$, so the classical result is recovered. But for large ω it decreases exponentially.

Multiplying again by the density of modes $g(\omega)$ yields the spectral energy density

$$\boxed{u(\omega) = \frac{\hbar\omega^3}{\pi^2 c^3} \frac{1}{e^{\hbar\omega/(k_\mathrm{B}T)} - 1}}$$

which is known as the *Planck law*.

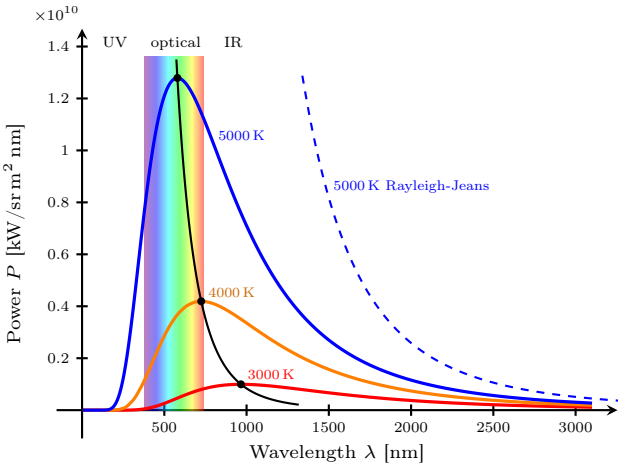

3.2 Photoelectric Effect

In 1887, Heinrich Hertz studied the emission of electrons from a metal surface illuminated by light.

Classically, light is an EM wave whose energy is carried continuously by the wave amplitude. Thus one would expect: Increasing the intensity increases the energy transferred to each electron, therefore the kinetic energy of the emitted electrons should increase with intensity. The color of the light, i.e. its frequency, should play no essential role. Even low-frequency light should eventually eject electrons if one waits long enough (there should be a *lag time*).

But, these expectations are not confirmed experimentally, instead the observed behavior is: Increasing the intensity increases mainly the *number* of emitted electrons. The kinetic energy of the emitted electrons depends on the frequency of the light. Below a certain threshold frequency ω_0, no electrons are emitted at all, regardless of the intensity or duration of illumination. The emission is essentially immediate once the frequency is above the threshold.

Thus the crucial quantity is not the intensity, but the frequency.

In 1905, Einstein explained this by assuming that light is absorbed in discrete quanta, later called *photons*. A photon of angular frequency ω carries the energy

$$\boxed{E_\gamma = \hbar\omega = \frac{hc}{\lambda}} \tag{3.4}$$

where λ is the wavelength of the light. To remove an electron from the metal, one must overcome the *work function*

$$W = \hbar\omega_0$$

where ω_0 is the threshold frequency. Hence emission is possible only if $\hbar\omega \geq W$, regardless of the intensity of the light.

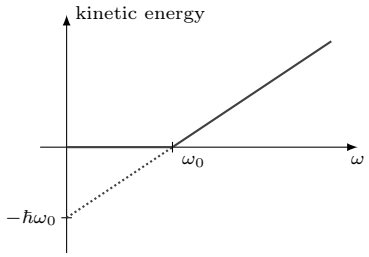

The kinetic energy of the emitted electron is then

$$E_{\text{kin}} = \hbar\omega - W = \hbar(\omega - \omega_0)$$

Thus:

- increasing the frequency increases the energy of each photon and therefore the kinetic energy of the emitted electrons

- increasing the intensity increases only the number of photons and therefore mainly the number of emitted electrons (if the frequency is above the threshold)

The photoelectric effect shows that light cannot be understood as a purely classical wave. Its energy is exchanged in discrete packets. Together with black-body radiation, this was one of the key pieces of evidence for quantization.

At the same time, light still exhibits interference and polarization phenomena. Thus quantum theory combines both aspects:

- waves behave like particles
- particles behave like waves

3.3 Wave-Particle Dualism

plane wave:

$$\psi(\boldsymbol{x}, t) = C \cdot e^{\imath(\boldsymbol{k}\cdot\boldsymbol{x} - \omega t)} \tag{3.5}$$

- $\boldsymbol{k}$: wave vector, perpendicular to the wavefronts
- $\lambda = 2\pi/|\boldsymbol{k}|$: wavelength (distance between wavefronts)
- $\Delta t = 2\pi/\omega$: period (time between passage of two wavefronts)
- $v = \omega/|\boldsymbol{k}|$: phase velocity (speed of wavefronts)

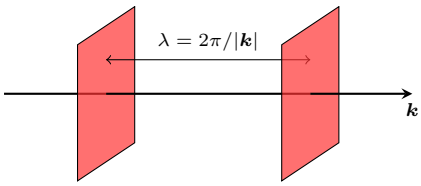

Interlude 3.1 (Hilbert spaces)**.** The state spaces of quantum-mechanical systems are generally *Hilbert spaces*, i.e., complete vector space over the complex numbers with an inner product. *Quantum informatics* happens in *finite*-dimensional Hilbert spaces, but in general, they can be (uncountably) infinite-dimensional. An important result states that, up to isomorphism, Hilbert spaces are classified by the dimension of an orthonormal basis.

A normalized wave packet ψ is an element of an infinite-dimensional Hilbert space, and the inner product is defined as

$$\langle \psi_1 | \psi_2 \rangle \coloneqq \int_{\mathbb{R}^3} \psi_1^*(\boldsymbol{x}, t)\psi_2(\boldsymbol{x}, t)\,\mathrm{d}\boldsymbol{x} \tag{3.6}$$

For the integral to be well-defined, we have to restrict to square-integrable functions. In order to allow for normalization, the factor C in (3.5) has to be an envelope that makes the integral finite. This yields a wave packet that is then again a superposition of plane waves. ◀

Now we want to relate the parameters of (3.5) to particle properties:

- Louis de Broglie had the idea to associate the momentum of a particle with its quantum mechanical wave vector as
  $$\boldsymbol{p} = \hbar\boldsymbol{k}$$
  and, thus: $\boldsymbol{k} = \boldsymbol{p}/\hbar$
- We already introduced the energy in (3.4). As in classical, non-relativistic mechanics, we relate it to the momentum as
  $$E = \hbar\omega = \frac{\|\boldsymbol{p}\|^2}{2m}$$
  and, thus: $\omega = \|\boldsymbol{p}\|^2/(2m\hbar)$

matter wave of a free particle (no potential, no walls) with mass m and momentum $\boldsymbol{p}$:

$$\boxed{\psi(\boldsymbol{x}, t) = C(\boldsymbol{x}, t) \cdot e^{\imath(\boldsymbol{p}\cdot\boldsymbol{x} - \frac{\|\boldsymbol{p}\|^2}{2m}t)/\hbar}} \tag{3.7}$$

Remark 3.2. If C in (3.7) does not depend on $\boldsymbol{x}$ and t, then $\psi(\boldsymbol{x}, t)$ is a plane wave, i.e. an idealized free-particle state that is spatially extended infinitely in all directions. That corresponds to a particle that is not localized in space at all, so it kind of everywhere equally. But when the idea came up that the wave characterizes the probability of finding the particle at a certain position, it was realized that such a plane wave is not a normalizable physical state, since there is no uniform probability distribution over the whole space that is normalized. Therefore, $C(\boldsymbol{x}, t)$ is usually some kind of a Gaussian or other envelope that limits the spatial extent of the wave. ◀

What possible equation of motion would lead to a solution of the form (3.7)?

People were looking for a linear equation, because they observed interference effects (e.g. in double-slit experiments). Interference effects are characteristic for linear behaviors.

For simplicity, we assume that C is constant and analyze the time evolution:

$$\begin{aligned}\frac{\partial}{\partial t}\psi(\boldsymbol{x}, t) &= \psi(\boldsymbol{x}, t) \cdot \frac{\imath}{\hbar} \cdot \left(-\frac{\|\boldsymbol{p}\|^2}{2m}\right) \\ &= \left(p_x^2 + p_y^2 + p_z^2\right) \cdot \psi(\boldsymbol{x}, t) \cdot \left(\frac{-\imath}{2m\hbar}\right) \\ &= -\hbar^2 \cdot \left(\tfrac{\partial^2}{\partial x^2} + \tfrac{\partial^2}{\partial y^2} + \tfrac{\partial^2}{\partial z^2}\right) \psi(\boldsymbol{x}, t) \cdot \left(\frac{-\imath}{2m\hbar}\right) \\ &= \hbar \cdot \Delta\psi(\boldsymbol{x}, t) \cdot \frac{\imath}{2m}\end{aligned}$$

what linear operation could we also apply to ψ that would also lead to these factors?

Δ is the Laplace operator

which directly corresponds to the Schrödinger equation for a free particle:

$$\boxed{\imath\hbar \cdot \frac{\partial}{\partial t}\psi(\boldsymbol{x}, t) = -\frac{\hbar^2}{2m} \cdot \Delta\psi(\boldsymbol{x}, t)} \tag{3.8}$$

The left hand side of the (non-relativistic) Schrödinger equation always looks like in (3.8). The right hand side changes if we have a particle that is not free, e.g. in a potential $V(\boldsymbol{x})$.

Interlude 3.1 (continuing from p. 21)**.** The scalar product (3.6) defines a norm

$$\|\psi\| = \sqrt{\langle\psi|\psi\rangle} = \left(\int_{\mathbb{R}^3} |\psi(\boldsymbol{x}, t)|^2 \, \mathrm{d}\boldsymbol{x}\right)^{1/2} \tag{3.9}$$

on the Hilbert space $\mathcal{H}$. ◀

Properties:

- linearity: if ψ_1 and ψ_2 are solutions, then any linear combination $\alpha\psi_1 + \beta\psi_2$ is also a solution.
  - not surprising, since we derived it to be linear. this makes quantum mechanics a linear theory.
  - entanglement (Section 1.5) comes out of this: if we have two Qbits, i.e. $|00\rangle$ and $|11\rangle$ are both possible states of the system, then $\alpha\,|00\rangle + \beta\,|11\rangle$ is also a possible state, and now we have an entangled state.
- preserves inner product: A solution can be written as $\psi(\boldsymbol{x}, t) = U(t)\psi(\boldsymbol{x}, 0)$. Define $\psi_0 = \psi(\cdot, 0)$. Then the norm square can be written as

  $$\|\psi(\cdot, t)\|^2 = \langle U(t)\psi_0 | U(t)\psi_0\rangle = \langle\psi_0 | U(t)^* U(t)\psi_0\rangle \overset{!}{=} \langle\psi_0|\psi_0\rangle = 1$$

  where the penultimate step holds iff $U(t)$ is unitary, i.e. $U(t)^*U(t) = \mathrm{id}$. Thus, the time evolution operator $U(t)$ is required to be unitary. This also implies that time evolution in quantum mechanics is reversible with $U(t)^{-1} = U(t)^*$.
  - $\int_{\mathbb{R}^3} |\psi(\boldsymbol{x}, 0)|^2 \, \mathrm{d}\boldsymbol{x} = 1$ implies $\int_{\mathbb{R}^3} |\psi(\boldsymbol{x}, t)|^2 \, \mathrm{d}\boldsymbol{x} = 1$ for all $t \in \mathbb{R}$.

Definition 3.1 (adjoint operator)**.** The adjoint A^* of a linear operator A is defined as

$$\langle\psi_1 | A\psi_2\rangle = \langle A^*\psi_1 | \psi_2\rangle$$

where $\langle\cdot|\cdot\rangle$ denotes the inner product defined in (3.6). ◀

Fact 3.1 (Properties of Adjoint Operators)**.**

- $|\psi'\rangle = \hat{A}\,|\psi\rangle \;\longleftrightarrow\; \langle\psi'| = \langle\psi|\,\hat{A}^*$
- $(\hat{A}^*)^* = \hat{A}$
- $(a\hat{A})^* = a^*\hat{A}^*$ for $a \in \mathbb{C}$
- $(\hat{A} + \hat{B})^* = \hat{A}^* + \hat{B}^*$
- $(\hat{A}\hat{B})^* = \hat{B}^*\hat{A}^*$ ◁

Definition 3.2. a linear operator A is self-adjoint or Hermitian iff $A^* = A$ ◀

Definition 3.3. a linear operator U is unitary iff $UU^* = \mathrm{id}$, or, equivalently, $U^* = U^{-1}$ ◀

Theorem 3.2 (Unitary operators preserve inner products)**.** If U is unitary, then

$$\langle U\psi_1 | U\psi_2\rangle = \langle\psi_1|\psi_2\rangle \tag{3.10}$$

for all ψ_1, ψ_2. ◁

Proof. By Definition 3.1 and Definition 3.3, $\langle U\psi_1|U\psi_2\rangle = \langle\psi_1|U^*U\psi_2\rangle = \langle\psi_1|\psi_2\rangle$. □

3.4 Observables

Correspondence principle: measurable physical quantities in classical mechanics can be translated to quantum mechanics by using the corresponding self-adjoint (Hermitian) operator, so-called *observable*. The probability density for finding a particle at position $\boldsymbol{x}$ at time t is $|\psi(\boldsymbol{x},t)|^2$, and expectation values are computed by integrating this density against the corresponding operator. We use (3.7) to derive some important observables:

Position:

$$\mathbb{E}_\psi[\boldsymbol{x}] = \int_{\mathbb{R}^3} |\psi(\boldsymbol{x},t)|^2\,\boldsymbol{x}\,\mathrm{d}\boldsymbol{x} = \int_{\mathbb{R}^3} \psi^*(\boldsymbol{x},t)\,\boldsymbol{x}\,\psi(\boldsymbol{x},t)\,\mathrm{d}\boldsymbol{x} = \langle\psi|\boldsymbol{x}\psi\rangle$$

operator: $\psi \mapsto \boldsymbol{x}\psi$

Momentum:

For the plane wave (3.7), $\nabla\psi(\boldsymbol{x},t) = \frac{\imath}{\hbar}\,\boldsymbol{p}\,\psi(\boldsymbol{x},t)$ and thus $\frac{\hbar}{\imath}\nabla\psi(\boldsymbol{x},t) = \boldsymbol{p}\,\psi(\boldsymbol{x},t)$. This motivates the momentum operator

$$\psi \mapsto \tfrac{\hbar}{\imath}\nabla\psi$$

and the expectation value

$$\mathbb{E}_\psi[\boldsymbol{p}] = \int_{\mathbb{R}^3} \psi^*(\boldsymbol{x},t)\cdot\tfrac{\hbar}{\imath}\nabla\psi(\boldsymbol{x},t)\,\mathrm{d}\boldsymbol{x} = \langle\psi|\tfrac{\hbar}{\imath}\nabla\psi\rangle$$

Energy:

For the plane wave (3.7), $\frac{\partial}{\partial t}\psi(\boldsymbol{x},t) = -\frac{\imath E}{\hbar}\,\psi(\boldsymbol{x},t)$ and thus $\imath\hbar\frac{\partial}{\partial t}\psi(\boldsymbol{x},t) = E\,\psi(\boldsymbol{x},t)$. This motivates the energy operator

$$\psi \mapsto \imath\hbar\tfrac{\partial}{\partial t}\psi$$

and the expectation value

$$\mathbb{E}_\psi[E] = \int_{\mathbb{R}^3} \psi^*(\boldsymbol{x},t)\cdot\imath\hbar\tfrac{\partial}{\partial t}\psi(\boldsymbol{x},t)\,\mathrm{d}\boldsymbol{x} = \langle\psi|\imath\hbar\tfrac{\partial}{\partial t}\psi\rangle$$

In classical mechanics, the energy of a free particle is $E = \|\boldsymbol{p}\|^2/(2m)$. If we replace the left side with $E = \imath\hbar\frac{\partial}{\partial t}$ and the right side with $(\frac{\hbar}{\imath}\nabla)^2/(2m)$, we return to the Schrödinger equation (3.8) for a free particle.

The general, non-relativistic, time-dependent Schrödinger equation for a particle in a potential $V(\boldsymbol{x},t)$ is

$$\boxed{\imath\hbar\frac{\partial}{\partial t}\psi(\boldsymbol{x},t) = \left(-\frac{\hbar^2}{2m}\Delta + V(\boldsymbol{x},t)\right)\psi(\boldsymbol{x},t)} \tag{3.11}$$

which can be written as

$$\imath\hbar\frac{\partial}{\partial t}\psi(\boldsymbol{x},t) = H\psi(\boldsymbol{x},t)$$

where

$$H = -\frac{\hbar^2}{2m}\Delta + V(\boldsymbol{x}, t) \tag{3.12}$$

is the Hamiltonian operator corresponding to the total (kinetic + potential) non-relativistic energy of the system. For fixed t, H is Hermitian.

For a free particle, $V(\boldsymbol{x}, t) = 0$, and (3.11) reduces to (3.8). In that case, $H = -\frac{\hbar^2}{2m}\Delta$ and the general solution is

$$\psi(\boldsymbol{x}, t) = e^{-\imath Ht/\hbar}\psi(\boldsymbol{x}, 0)$$

where $e^A \coloneqq \sum_{n=0}^{\infty} \frac{A^n}{n!}$ (convergent for some assumptions on the operator A). This yields the time evolution operator $U(t) = e^{-\imath Ht/\hbar}$. Since H is Hermitian (Definition 3.2), we have

$$e^{-\imath Ht/\hbar}\left(e^{-\imath Ht/\hbar}\right)^* = e^{-\imath Ht/\hbar}e^{\imath H^* t/\hbar} = e^{\imath(H^*-H)t/\hbar} = e^0 = \mathrm{id}$$

and hence $U(t)$ is unitary and therefore preserves the norm of the state vector. In general, if A is Hermitian, then $e^{\imath A}$ is unitary. Furthermore, we have the properties

- $U(t_1) \cdot U(t_2) = U(t_1 + t_2)$
- $U(0) = \mathrm{id}$
- $U(t)^{-1} = U(t)^* = U(-t)$

4 Digression on Operators

4.1 Bounded Operators

Definition 4.1. Suppose X and Y are Banach spaces and U is the open unit ball in X. A linear map $T\colon X \to Y$ is said to be *compact* if the closure of $T(U)$ is compact in Y. It is clear that T is then also bounded. Thus $T \in B(X, Y)$. ◀

Equivalently, T is compact if and only if every bounded sequence $(x_n)_n$ in X contains a subsequence $(x_{n_i})_i$ such that $(Tx_{n_i})_i$ converges to a point of Y.

Compact operators are as similar to linear operators on finite-dimensional spaces as one has any right to expect from operators on infinite-dimensional spaces. In particular, these similarities show up in their spectral properties.

Definition 4.2. Suppose X is a Banach space. Then $B(X) \coloneqq B(X, X)$ is not merely a Banach space but also an algebra: If $S, T \in B(X)$, one defines $ST \in B(X)$ by $(ST)x = S(T(x))$, $x \in X$. The inequality $\|ST\| \leq \|S\|\|T\|$ is easy to verify. In particular, powers of $T \in B(X)$ can be defined: $T^0 = \mathrm{id}$, the identity mapping on X, and $T^n = T^{n-1}T$ for $n \geq 1$. ◀

Definition 4.3. An operator $T \in B(X)$ is said to be *invertible* if there exists $S \in B(X)$ such that $ST = TS = \mathrm{id}$. In this case, we write $S = T^{-1}$. By the open mapping theorem, this happens if and only if $N(T) = \{0\}$ and $R(T) = X$. ◀

Definition 4.4. The *spectrum* $\sigma(T)$ of an operator $T \in B(X)$ is the set of all scalars $\lambda \in \mathbb{C}$ such that $T - \lambda\,\mathrm{id}$ is not invertible. Thus, $\lambda \in \sigma(T)$ if and only if at least one of the following statements is true:

(i) The range of $T - \lambda\,\mathrm{id}$ is not all of X.

(ii) $T - \lambda\,\mathrm{id}$ is not one-to-one (i.e., it is not injective, i.e. it has a non-trivial kernel).

If (ii) holds, λ is said to be an *eigenvalue* of T; the corresponding *eigenspace* is $N(T - \lambda\,\mathrm{id})$; each $x \in N(T - \lambda\,\mathrm{id})$ except $x = 0$ is an *eigenvector* of T corresponding to λ; it satisfies $Tx = \lambda x$. ◀

Theorem 4.1. Let X and Y be Banach spaces.

(a) If $T \in B(X, Y)$ and $\dim R(T) < \infty$, then T is compact.

(b) If $T \in B(X, Y)$, T is compact, and $R(T)$ is closed, then $\dim R(T) < \infty$.

(c) The compact operators form a closed subspace of $B(X, Y)$ in its norm topology.

(d) If $T \in B(X)$, T is compact, and $\lambda \neq 0$, then $\dim N(T - \lambda\,\mathrm{id}) < \infty$.

(e) If $\dim X = \infty$, $T \in B(X)$, and T is compact, then $0 \in \sigma(T)$.

(f) If $S \in B(X)$, $T \in B(X)$, and T is compact, so are ST and TS. ◁

Definition 4.5. an operator $T \in B(\mathcal{H})$ is said to be

(a) *normal* if $T^*T = TT^*$

(b) *self-adjoint* (or *Hermitian*) if $T^* = T$

(c) *unitary* if $T^*T = \mathrm{id} = TT^*$, where id is the identity operator on $\mathcal{H}$

(d) a *projection* if $T^2 = T$ ◀

It is clear that self-adjoint operators and unitary operators are normal. Many results are about normal operators. This algebraic requirement, namely that T commutes with its adjoint, has strong analytic and geometric consequences.

Theorem 4.2. An operator $T \in B(\mathcal{H})$ is normal if and only if $\|Tx\| = \|T^*x\|$ for every $x \in \mathcal{H}$. If T is normal, then the following properties hold:

(a) $N(T) = N(T^*)$

(b) $R(T)$ is dense in $\mathcal{H}$ if and only if T is one-to-one

(c) T is invertible if and only if there exists $\delta > 0$ such that $\|Tx\| \geq \delta\|x\|$ for every $x \in \mathcal{H}$

(d) If $Tx = \alpha x$ for some $x \in \mathcal{H}$ and $\alpha \in \mathbb{C}$, then $T^*x = \alpha^* x$

(e) If α and β are distinct eigenvalues of T, then the corresponding eigenspaces are orthogonal to each other. ◁

Theorem 4.3. Each of the following four properties of a projection $P \in B(\mathcal{H})$ (Definition 4.5) implies the other three:

- P is self-adjoint
- P is normal
- $R(P) = N(P)^\perp$
- $\langle Px|x\rangle = \|Px\|^2$ for every $x \in \mathcal{H}$

Moreover, two self-adjoint projections P and Q have $R(P) \perp R(Q)$ if and only if $PQ = 0$. ◁

Theorem 4.4. If $M \subseteq \mathcal{H}$ is a closed subspace, then there exists a unique self-adjoint projection $P \in B(\mathcal{H})$ such that $R(P) = M$. Moreover, $N(P) = M^\perp$. ◁

Theorem 4.5. (a) If U is unitary and $\lambda \in \sigma(U)$, then $|\lambda| = 1$.

(b) If S is hermitian (self-adjoint) and $\lambda \in \sigma(S)$, then λ is real. ◁

Definition 4.6. Let $\mathfrak{M}$ be a σ-algebra on a set Ω, and let $\mathcal{H}$ be a Hilbert space. A *resolution of the identity* (on $\mathfrak{M}$) is a map

$$E\colon \mathfrak{M} \to B(\mathcal{H})$$

with the following properties:

(a) $E(\emptyset) = 0$, $E(\Omega) = \mathrm{id}$

(b) Each $E(\omega)$ is a self-adjoint projection (see Definition 4.5)

(c) $E(\omega_1 \cap \omega_2) = E(\omega_1)E(\omega_2)$

(d) If $\omega_1 \cap \omega_2 = \emptyset$, then $E(\omega_1 \cup \omega_2) = E(\omega_1) + E(\omega_2)$

(e) For every $x, y \in \mathcal{H}$, the set function $E_{x,y}$ defined by

$$E_{x,y}(\omega) = \langle E(\omega)x|y\rangle \tag{4.1}$$

is a complex measure on $\mathfrak{M}$. ◀

When $\mathfrak{M}$ is the σ-algebra of all Borel sets on a compact or locally compact Hausdorff space, it is customary to add another requirement to (e): Each $E_{x,y}$ should be a *regular* Borel measure (this is automatically satisfied on compact metric spaces).

A resolution of the identity is also called a *projection-valued measure* (PVM) or *spectral measure*. In quantum mechanics, PVMs are the mathematical description of projective measurements (see Section 5.3). They are generalized by *positive operator-valued measures* (POVMs), in the same sense that a mixed state (Definition 5.4) or density matrix (Definition 5.3) generalizes the notion of a pure state (Definition 5.5).

Some immediate consequences of the properties in Definition 4.6:

- By (b) and (e), we have

$$E_{x,x}(\omega) = \langle E(\omega)x|x\rangle = \langle E(\omega)^2 x|x\rangle = \langle E(\omega)x|E(\omega)x\rangle = \|E(\omega)x\|^2 \tag{4.2}$$

  for every $x \in \mathcal{H}$ so that each $E_{x,x}$ is a positive measure on $\mathfrak{M}$ with total variation $\|E_{x,x}\| = E_{x,x}(\Omega) = \|x\|^2$.[1]

- By (c), any two of the projections $E(\omega)$ commute with each other.

- If $\omega_1 \cap \omega_2 = \emptyset$, (a) and (c) show that the ranges of $E(\omega_1)$ and $E(\omega_2)$ are orthogonal to each other (Theorem 4.3).

Proposition 4.6. If E is a resolution of the identity, and if $x \in \mathcal{H}$, then

$$\omega \mapsto E(\omega)x$$

is countably additive $\mathcal{H}$-valued measure on $\mathfrak{M}$. ◁

As in the case of ordinary measures, it is possible to integrate functions with respect to a PVM (Definition 4.6), the result being a linear operator on $\mathcal{H}$.

[1] In particular, if $\|x\| = 1$, then $E_{x,x}$ is a probability measure on $\mathfrak{M}$, which will be important in Section 5.3.

Theorem 4.7. If E is a resolution of the identity, then there exists an isometric *-isomorphism Ψ of the Banach algebra $L^\infty(E)$ onto a closed normal subalgebra $\mathcal{A}$ of $B(\mathcal{H})$, which is related to E by

$$\langle \Psi(f)x|y\rangle = \int_\Omega f \,\mathrm{d}E_{x,y}$$

for every $f \in L^\infty(E)$ and every $x, y \in \mathcal{H}$. This justifies the notation $\Psi(f) = \int_\Omega f \,\mathrm{d}E$. Moreover,

$$\|\Psi(f)x\|^2 = \int_\Omega |f|^2 \,\mathrm{d}E_{x,x}$$

for every $f \in L^\infty(E)$ and every $x \in \mathcal{H}$ and an operator $Q \in B(\mathcal{H})$ commutes with every $E(\omega)$ if and only if Q commutes with every $\Psi(f)$.

Recall that a *normal* subalgebra $\mathcal{A}$ of $B(\mathcal{H})$ is a commutative one which contains T^* for every $T \in \mathcal{A}$. To say that Ψ is a *-isomorphism means that Ψ is one-to-one, linear, multiplicative and that $\Psi(f^*) = \Psi(f)^*$ for every $f \in L^\infty(E)$. ◁

The principal assertion of Theorem 4.8 is that every bounded normal operator T on a Hilbert space induces (in a canonical way) a resolution E of the identity on the Borel subsets of its spectrum $\sigma(T)$ and that T can be reconstructed from E by an integral of the type discussed in Theorem 4.7. A large part of the theory of normal operators depends on this fact.

Theorem 4.8 (Spectral theorem)**.** If $T \in B(\mathcal{H})$ and T is normal (Definition 4.5), then there exists a unique resolution of the identity E (Definition 4.6) on the Borel subsets of $\sigma(T)$ (Definition 4.4) such that

$$T = \int_{\sigma(T)} \lambda \,\mathrm{d}E(\lambda) \tag{4.3}$$

Moreover, every projection $E(\omega)$ commutes with every $S \in B(\mathcal{H})$ which commutes with T. ◁

We call E the *spectral decomposition* of T.

Sometimes it is convenient to think of E as being defined for all Borel sets in $\mathbb{C}$; to achieve this put $E(\omega) = 0$ if $\omega \cap \sigma(T) = \emptyset$.

Notation 4.7 (Symbolic Calculus for Normal Operators)**.** If E is the spectral decomposition of a normal operator $T \in B(\mathcal{H})$, and if f is a bounded Borel function on $\sigma(T)$, it is customary to denote the operator $\Psi(f) = \int_{\sigma(T)} f \,\mathrm{d}E$ by $f(T)$. ◀

Theorem 4.9. A normal $T \in B(\mathcal{H})$ is

(a) hermitian if and only if $\sigma(T)$ lies on the real axis

(b) unitary if and only if $\sigma(T)$ lies on the unit circle ◁

4.2 Unbounded operators

Let $\mathcal{H}$ be a Hilbert space. By an *operator* in $\mathcal{H}$ we mean a linear map $T\colon D(T) \to \mathcal{H}$, where the *domain* $D(T)$ is a linear subspace of $\mathcal{H}$. If $D(T)$ is dense in $\mathcal{H}$, we say that T is densely defined.

The *graph* of an operator T is $G(T) \coloneqq \{(x, Tx) : x \in D(T)\} \subseteq \mathcal{H} \times \mathcal{H}$. The operator T is said to be *closed* if $G(T)$ is closed in $\mathcal{H} \times \mathcal{H}$. It is said to be *closable* if the closure of $G(T)$ is itself the graph of an operator.

Equivalently, T is closed if whenever $x_n \in D(T)$, $x_n \to x$ in $\mathcal{H}$, and $Tx_n \to y$ in $\mathcal{H}$, then $x \in D(T)$ and $Tx = y$.

Suppose T is densely defined. An element $y \in \mathcal{H}$ is said to lie in the domain of the *adjoint* T^* if the map $x \mapsto \langle Tx|y\rangle$, $x \in D(T)$, is a bounded linear functional on $D(T)$ (with the norm inherited from $\mathcal{H}$). By Hahn-Banach and the Riesz representation theorem, there is then a unique vector $T^*y \in \mathcal{H}$ such that

$$\langle Tx|y\rangle = \langle x|T^*y\rangle$$

for every $x \in D(T)$. Thus

$$D(T^*) \coloneqq \{y \in \mathcal{H} : x \mapsto \langle Tx|y\rangle \text{ is bounded on } D(T)\}$$

and $T^*\colon D(T^*) \to \mathcal{H}$ is an operator in $\mathcal{H}$. If $T \in B(\mathcal{H})$, this agrees with the usual adjoint of a bounded operator.

Definition 4.8. A densely defined operator T is called

(a) *symmetric* if $\langle Tx|y\rangle = \langle x|Ty\rangle$ for every $x, y \in D(T)$, or, equivalently, $T \subseteq T^*$

(b) *self-adjoint* if $T = T^*$ ◀

Unlike for bounded operators, symmetry and self-adjointness are not the same for unbounded operators.

If S and T are operators in $\mathcal{H}$, their sum and product are defined on the natural domains

$$D(S+T) = D(S) \cap D(T) \qquad D(ST) = \{x \in D(T) : Tx \in D(S)\}$$

With these definitions, the associative laws remain valid, but one must always watch domains carefully.

Theorem 4.10. If T is densely defined, then T^* is closed. In particular, every self-adjoint operator is closed. ◁

Theorem 4.11. If T is densely defined and closed, then

(i) $D(T^*)$ is dense in $\mathcal{H}$ and $T^{**} = T$.

(ii) $I + T^*T$ is a one-to-one mapping of $D(T^*T) = \{x \in D(T) : Tx \in D(T^*)\}$ onto $\mathcal{H}$. Its inverse $(I + T^*T)^{-1}$ belongs to $B(\mathcal{H})$.

(iii) T^*T is self-adjoint and positive, i.e., $\langle T^*Tx|x\rangle = \|Tx\|^2 \geq 0$ for all $x \in D(T^*T)$. ◁

Definition 4.9. The *resolvent set* of an operator T is the set of all $\lambda \in \mathbb{C}$ such that $T - \lambda I$ is one-to-one, maps $D(T)$ onto $\mathcal{H}$, and has bounded inverse in $B(\mathcal{H})$. Its complement is the *spectrum* $\sigma(T)$. ◀

Definition 4.10. Let E be a resolution of the identity on a measurable space $(\Omega, \mathfrak{M})$, and let $f : \Omega \to \mathbb{C}$ be measurable. Define $D_f := \{x \in \mathcal{H} : \int_\Omega |f|^2 \, \mathrm{d}E_{x,x} < \infty\}$. Then D_f is a dense subspace of $\mathcal{H}$. The operator $\Psi(f)$ is defined by the requirement

$$\langle \Psi(f)x|y\rangle = \int_\Omega f \, \mathrm{d}E_{x,y}$$

for every $x \in D_f$ and every $y \in \mathcal{H}$. ◀

Theorem 4.12. With the notation above, $\Psi(f)$ is a densely defined closed operator with domain D_f, and

$$\|\Psi(f)x\|^2 = \int_\Omega |f|^2 \, \mathrm{d}E_{x,x}$$

for every $x \in D_f$. If f and g are measurable, then $\Psi(f)\Psi(g) \subseteq \Psi(fg)$, and $\Psi(f)^* = \Psi(f^*)$. Hence $\Psi(f)^*\Psi(f) = \Psi(|f|^2) = \Psi(f)\Psi(f)^*$. So every operator of the form $\Psi(f)$ is normal. ◁

Definition 4.11. A (not necessarily bounded) linear operator T in $\mathcal{H}$ is said to be *normal* if T is closed, densely defined, and $T^*T = TT^*$. ◀

Theorem 4.13. If N is normal, then

(a) $D(N) = D(N^*)$

(b) $\|Nx\| = \|N^*x\|$ for every $x \in D(N)$

(c) N is *maximally normal*: if M is normal and $N \subseteq M$, then $M = N$ ◁

Theorem 4.9 can be extended to unbounded normal operators as follows:

Theorem 4.14. A densely defined normal operator T is

(a) self-adjoint if and only if $\sigma(T)$ lies on the real axis

(b) unitary if and only if $\sigma(T)$ lies on the unit circle

Note that if (b) holds, then necessarily $D(T) = \mathcal{H}$, so T is bounded. ◁

Definition 4.12. An operator T in $\mathcal{H}$ is called an *involution* if $D(T^2) = D(T)$ and

$$T^2x = x$$

for every $x \in D(T)$. ◀

Theorem 4.15. Let T be an operator in $\mathcal{H}$. Then any two of the following properties imply the third:

(a) $T^* = T^{-1}$ (T is unitary) ☞ Definition 4.5 (c)

(b) $T^* = T$ (T is self-adjoint) ☞ Definition 4.8 (b)

(c) $T^2 = I$ (T is an involution) ☞ Definition 4.12

In particular, if one of these three implications involves unitarity, then necessarily $D(T) = \mathcal{H}$, so T is bounded. ◁

Proof. If T is unitary and self-adjoint, then $T^{-1} = T^* = T$, hence $T^2 = I$, so T is an involution. If T is unitary and an involution, then $T^* = T^{-1} = T$, so T is self-adjoint. Finally, assume that T is self-adjoint and an involution. Then T is closed, and for every $x \in D(T)$,

$$\|Tx\|^2 = \langle Tx|Tx\rangle = \langle x|T^2x\rangle = \|x\|^2$$

so T is an isometry on $D(T)$. Moreover,

$$\|x\|_T^2 = \|x\|^2 + \|Tx\|^2 = 2\|x\|^2$$

for every $x \in D(T)$. Hence $D(T)$ is complete in the graph norm and therefore also in the Hilbert norm. Since a self-adjoint operator is densely defined, $D(T)$ is a dense complete subspace of $\mathcal{H}$, hence $D(T) = \mathcal{H}$. Thus T is everywhere defined, isometric, and surjective because $T^2 = I$, so T is unitary. □

Notation 4.13. If $M \subseteq \mathcal{H}$ is a closed subspace, we write P_M for the orthogonal projection onto M, i.e. the unique self-adjoint projection with range $R(P_M) = M$ (see Theorem 4.4). ◀

Proposition 4.16. Let A be self-adjoint with spectral measure E_A. Then for every eigenvalue λ of A,

$$E_A(\{\lambda\}) = P_{N(A-\lambda\,\mathrm{id})}$$

where $P_{N(A-\lambda\,\mathrm{id})}$ is the orthogonal projection onto the eigenspace

$$N(A - \lambda\,\mathrm{id}) = \{x \in \mathcal{H} : Ax = \lambda x\}$$

This notation is justified by Theorem 4.4, since eigenspaces of self-adjoint operators are closed subspaces. ◁

Theorem 4.17 (Spectral theorem)**.** Every normal operator T in $\mathcal{H}$ has a unique spectral decomposition[2] E, such that

$$\langle Tx|y\rangle = \int_{\sigma(T)} \lambda \,\mathrm{d}E_{x,y}(\lambda) \tag{4.4}$$

for every $x \in D(T) \overset{4.10}{=} \{x \in \mathcal{H} : \int_{\sigma(T)} |\lambda|^2 \,\mathrm{d}E_{x,x}(\lambda) < \infty\}$ and every $y \in \mathcal{H}$. Equivalently:

$$T = \int_{\sigma(T)} \lambda \,\mathrm{d}E(\lambda) \tag{4.5}$$

Moreover, $E(\omega)S = SE(\omega)$ for every Borel set $\omega \subseteq \sigma(T)$ and every $S \in B(\mathcal{H})$ which commutes with T in the sense that $ST \subseteq TS$. ◁

[2] In the unbounded case, a resolution of the identity means exactly the same thing as before (Definition 4.6): a projection-valued measure (PVM) $E : \mathfrak{M} \to B(\mathcal{H})$. What is new is that for measurable f, the operator $\Psi(f)$ need not be bounded; its domain is $D_f = \{x \in \mathcal{H} : \int |f|^2 \,\mathrm{d}E_{x,x} < \infty\}$ (see Definition 4.10).

5 Postulates of Quantum Mechanics

5.1 The state

To an isolated quantum system we associate a complex Hilbert space $\mathcal{H}$. A *state* of the system is represented by a normalized vector $\psi \in \mathcal{H}$ with $\|\psi\| = 1$. If $\mathcal{H}$ is realized as a space of wave functions, for example $\mathcal{H} = L^2(\mathbb{R}^3)$, then, as already introduced in (3.6), the scalar product is

$$\langle \psi_1 | \psi_2 \rangle \coloneqq \int_{\mathbb{R}^3} \psi_1(\boldsymbol{x}, t)^* \, \psi_2(\boldsymbol{x}, t) \, \mathrm{d}\boldsymbol{x}$$

and the normalization condition becomes $\|\psi\|^2 = \langle \psi | \psi \rangle = \int_{\mathbb{R}^3} |\psi(\boldsymbol{x}, t)|^2 \, \mathrm{d}\boldsymbol{x} = 1$. In that case, $|\psi(\boldsymbol{x}, t)|^2$ is interpreted as probability density.

The linear structure of $\mathcal{H}$ gives the *superposition principle*: if $\psi_1, \psi_2 \in \mathcal{H}$ and $\alpha, \beta \in \mathbb{C}$, then also $\alpha\psi_1 + \beta\psi_2 \in \mathcal{H}$. If, in addition, $\|\alpha\psi_1 + \beta\psi_2\| = 1$, then this superposition again represents a physical state. Thus quantum states behave like waves: linear combinations of states are again states.

Digression 5.1 (Position Basis)**.** Physicists often describe position by a family of kets $|\boldsymbol{x}\rangle$, $\boldsymbol{x} \in \mathbb{R}^3$, and write

$$|\psi\rangle = \int_{\mathbb{R}^3} \mathrm{d}\boldsymbol{x} \, \psi(\boldsymbol{x}) \, |\boldsymbol{x}\rangle \tag{5.1}$$

This is analogous to the expansion of $|\psi\rangle$ in a finite-dimensional orthonormal basis $\{|e_i\rangle\}$:

$$|\psi\rangle = \sum_i c_i \, |e_i\rangle$$

Thus the wave function $\psi(\boldsymbol{x})$ plays the role of a *continuous coefficient* of the state in the position representation.

Formally, the coefficient of $|\boldsymbol{x}\rangle$ is obtained by projection:

$$\psi(\boldsymbol{x}) = \langle \boldsymbol{x} | \psi \rangle$$

Thus $\psi(\boldsymbol{x})$ is the probability amplitude for the position outcome $\boldsymbol{x}$. Hence the Born rule in the position representation becomes

$$\mathbb{P}(\text{particle in } \omega) = \int_{\omega} |\psi(\boldsymbol{x})|^2 \, \mathrm{d}\boldsymbol{x}$$

for every Borel set $\omega \subseteq \mathbb{R}^3$.

The formal vectors $|\boldsymbol{x}\rangle$ satisfy the relations $\langle \boldsymbol{x}_1 | \boldsymbol{x}_2 \rangle = \delta(\boldsymbol{x}_1 - \boldsymbol{x}_2)$ and $\mathrm{id} = \int_{\mathbb{R}^3} |\boldsymbol{x}\rangle \langle \boldsymbol{x}| \, \mathrm{d}\boldsymbol{x}$. The former is the continuous analogue of $\langle e_i | e_j \rangle = \delta_{ij}$, while the latter is the continuous analogue of $\mathrm{id} = \sum_i |e_i\rangle \langle e_i|$.

Applying that continuous resolution of the identity to a state $|\psi\rangle$ yields

$$|\psi\rangle = \left(\int_{\mathbb{R}^3} |\boldsymbol{x}\rangle \langle \boldsymbol{x}| \, \mathrm{d}\boldsymbol{x} \right) |\psi\rangle = \int_{\mathbb{R}^3} |\boldsymbol{x}\rangle \langle \boldsymbol{x} | \psi \rangle \, \mathrm{d}\boldsymbol{x}$$

which is exactly (5.1).

Moreover, for two states $\psi_1, \psi_2 \in L^2(\mathbb{R}^3)$, one may write formally

$$\langle \psi_1 | = \int_{\mathbb{R}^3} \psi_1(\boldsymbol{x})^* \, \langle \boldsymbol{x}| \, \mathrm{d}\boldsymbol{x} \qquad |\psi_2\rangle = \int_{\mathbb{R}^3} \psi_2(\boldsymbol{y}) \, |\boldsymbol{y}\rangle \, \mathrm{d}\boldsymbol{y}$$

and therefore

$$\begin{aligned}
\langle \psi_1 | \psi_2 \rangle &= \left(\int_{\mathbb{R}^3} \psi_1(\boldsymbol{x})^* \, \langle \boldsymbol{x}| \, \mathrm{d}\boldsymbol{x} \right) \left(\int_{\mathbb{R}^3} \psi_2(\boldsymbol{y}) \, |\boldsymbol{y}\rangle \, \mathrm{d}\boldsymbol{y} \right) \\
&= \int_{\mathbb{R}^3} \int_{\mathbb{R}^3} \psi_1(\boldsymbol{x})^* \psi_2(\boldsymbol{y}) \, \langle \boldsymbol{x} | \boldsymbol{y} \rangle \, \mathrm{d}\boldsymbol{x} \, \mathrm{d}\boldsymbol{y} \\
&= \int_{\mathbb{R}^3} \int_{\mathbb{R}^3} \psi_1(\boldsymbol{x})^* \psi_2(\boldsymbol{y}) \delta(\boldsymbol{x} - \boldsymbol{y}) \, \mathrm{d}\boldsymbol{x} \, \mathrm{d}\boldsymbol{y} \\
&= \int_{\mathbb{R}^3} \psi_1(\boldsymbol{x})^* \psi_2(\boldsymbol{x}) \, \mathrm{d}\boldsymbol{x}
\end{aligned}$$

which recovers (3.6), the usual inner product on $L^2(\mathbb{R}^3)$.

The family $\{|\boldsymbol{x}\rangle : \boldsymbol{x} \in \mathbb{R}^3\}$ is often called a *continuous orthonormal basis*. This characterization is standard and useful, but not entirely accurate. The objects $|\boldsymbol{x}\rangle$ are not even honest vectors of $L^2(\mathbb{R}^3)$, since $\langle \boldsymbol{x}|\boldsymbol{x}\rangle = \delta(\mathbf{0})$ is not a finite number. Rather, they are *generalized eigenvectors* of the position operator, more precisely distribution-like objects in the sense of a *rigged Hilbert space* [5]

$$\Phi \subset \mathcal{H} \subset \Phi'$$

where $\mathcal{H} = L^2(\mathbb{R}^3)$, Φ is a space of sufficiently nice test functions (for example Schwartz functions), and Φ' is its dual space of continuous linear functionals (distributions). The generalized position eigenstates are then not elements of $\mathcal{H}$, but distribution-like objects in Φ'. In this sense, the notation $|\boldsymbol{x}\rangle$ is a formal Dirac ket corresponding to the evaluation functional at $\boldsymbol{x}$, and the relation $\langle \boldsymbol{x}|\psi\rangle = \psi(\boldsymbol{x})$ should be understood in this generalized sense.

To see how the position-basis connects to the spectral theory from Section 4 and Section 5.3, recall that by Theorem 4.17 every self-adjoint observable A has a spectral measure E_A, a PVM on $\sigma(A)$. For the position observable, the corresponding PVM $E_{\boldsymbol{X}}$ on the Borel subsets of $\mathbb{R}^3$ acts by

$$(E_{\boldsymbol{X}}(\omega)\psi)(\boldsymbol{x}) = \chi_\omega(\boldsymbol{x})\,\psi(\boldsymbol{x})$$

for each Borel set $\omega \subseteq \mathbb{R}^3$, where $\chi_\omega(\boldsymbol{x}) = 1$ if $\boldsymbol{x} \in \omega$ and $\chi_\omega(\boldsymbol{x}) = 0$ otherwise. So $E_{\boldsymbol{X}}(\omega)$ simply projects a wave function onto the part supported in ω. Hence the Born rule becomes

$$\begin{aligned}
\mathbb{P}^{\psi}_{\boldsymbol{X}}(\omega) &= \langle \psi|E_{\boldsymbol{X}}(\omega)|\psi\rangle = \langle \psi|E_{\boldsymbol{X}}(\omega)\psi\rangle \\
&= \int_{\mathbb{R}^3} \psi(\boldsymbol{x})^*\,(E_{\boldsymbol{X}}(\omega)\psi)(\boldsymbol{x})\,\mathrm{d}\boldsymbol{x} \\
&= \int_{\mathbb{R}^3} \psi(\boldsymbol{x})^*\,\chi_\omega(\boldsymbol{x})\,\psi(\boldsymbol{x})\,\mathrm{d}\boldsymbol{x} \\
&= \int_{\omega} |\psi(\boldsymbol{x})|^2\,\mathrm{d}\boldsymbol{x}
\end{aligned}$$

In Dirac notation, one writes the corresponding spectral projection as

$$E_{\boldsymbol{X}}(\omega) = \int_\omega |\boldsymbol{x}\rangle\,\langle \boldsymbol{x}|\;\mathrm{d}\boldsymbol{x}$$

Then one obtaines the same result as

$$\begin{aligned}
\mathbb{P}^{\psi}_{\boldsymbol{X}}(\omega) &= \langle \psi|E_{\boldsymbol{X}}(\omega)|\psi\rangle \\
&= \left\langle \psi \,\middle|\, \int_\omega |\boldsymbol{x}\rangle\langle \boldsymbol{x}|\,\mathrm{d}\boldsymbol{x} \,\middle|\, \psi \right\rangle \\
&= \int_\omega \langle \psi|\boldsymbol{x}\rangle\,\langle \boldsymbol{x}|\psi\rangle\;\mathrm{d}\boldsymbol{x} \\
&= \int_\omega |\psi(\boldsymbol{x})|^2\,\mathrm{d}\boldsymbol{x}
\end{aligned}$$

so the "position basis" is the Dirac-notation form of the spectral decomposition of the position observable. ◀

5.2 The time evolution

As already introduced in (3.11), the time evolution of a quantum state is governed by the Schrödinger equation

$$\boxed{\imath\hbar\frac{\partial}{\partial t}\psi(\boldsymbol{x},t) = H\psi(\boldsymbol{x},t)}$$

where H is the Hamiltonian operator (3.12). It is linear: any linear combination of solutions is again a solution.

We assume the Hamiltonian H to be self-adjoint (see Definition 3.2, 4.5 (b), 4.8). If H is time-independent, then the solutions can be written as

$$\psi(t) = e^{-\imath Ht/\hbar}\psi(0)$$

Thus the time evolution operator is $U(t) = e^{-\imath Ht/\hbar}$. As discussed before, $e^{\imath A}$ is unitary whenever A is Hermitian, so $U(t)$ is unitary (see Definition 3.3, 4.5 (c)). Hence time evolution preserves the inner product and in particular the norm: if $\langle \psi(0)|\psi(0)\rangle = 1$, then $\langle \psi(t)|\psi(t)\rangle = 1$ for all $t \in \mathbb{R}$.

5.3 Observables

A measurable physical quantity is represented by an observable, that is, by a self-adjoint operator $A = A^*$ (see Definitions 3.2, 4.5 (b), 4.8, Theorems 4.8, 4.17). The expectation value of A in a state ψ is

$$\langle A \rangle_\psi = \langle \psi | A | \psi \rangle \coloneqq \langle \psi | A\psi \rangle \tag{5.2}$$

provided $\psi \in D(A)$. The possible outcomes of a measurement of A are given by the spectrum (see Definition 4.9) $\sigma(A) = \sigma_p(A) \cup \sigma_c(A)$ (for self-adjoint (wether bounded or unbounded) operators, the residual spectrum $\sigma_r(A)$ is always empty) of A.

In a general Hilbert space $\mathcal{H}$, every self-adjoint operator A admits a spectral decomposition (see Section 4, Definition 4.9 and Theorem 4.17). Denote the spectrum of A by $\sigma_A \coloneqq \sigma(A)$. By Theorem 4.17, for every self-adjoint operator there exists a unique projection-valued measure (PVM)

$$E_A \colon \mathcal{B}(\sigma_A) \to B(\mathcal{H})$$

on the Borel sets of σ_A such that $A = \int_{\sigma_A} \lambda \, \mathrm{d}E_A(\lambda)$. Here each $E_A(\omega)$ is a self-adjoint projection, and by Definition 4.6 and the remarks following it,

$$\langle \psi | E_A(\omega) \psi \rangle \overset{(4.2)}{=} \| E_A(\omega)\psi \|^2 \geq 0$$

for every Borel set $\omega \subseteq \sigma_A$. Moreover, $E_A(\sigma_A) = \mathrm{id}$ and $E_A(\emptyset) = 0$, so for normalized ψ we have $\langle \psi | E_A(\sigma_A) \psi \rangle = \langle \psi | \psi \rangle = 1$. Therefore, for every normalized state ψ, the map

$$\mu_\psi^A(\omega) \coloneqq \langle \psi | E_A(\omega)\psi \rangle \overset{4.6.(\mathrm{b})}{=} \langle E_A(\omega)\psi | \psi \rangle \overset{(4.1)}{=} E_{\psi,\psi}(\omega) \tag{5.3}$$

for $\omega \in \mathcal{B}(\sigma_A)$, defines a probability measure $\mathbb{P}_A^\psi(\cdot) \coloneqq \mu_\psi^A(\cdot)$ on σ_A. This is the probability distribution of the measurement outcomes of A in the state ψ. This is called the *Born rule*.

With this notation, the expectation value is obtained as

$$\begin{aligned} \langle A \rangle_\psi = \langle \psi | A\psi \rangle &\overset{4.8.(\mathrm{b})}{=} \langle A\psi | \psi \rangle \\ &\overset{(4.4)}{=} \int_{\sigma_A} \lambda \, \mathrm{d}E_{\psi,\psi}(\lambda) \\ &\overset{(5.3)}{=} \int_{\sigma_A} \lambda \, \mathrm{d}\mu_\psi^A(\lambda) \end{aligned}$$

which is precisely the probabilistic expectation of the measurement outcomes with respect to the spectral probability measure. In case the spectrum has both discrete and continuous parts, then μ_ψ^A automatically contains both.

After an ideal measurement, the state updates by projection. If one only knows that the outcome lies in a Borel set ω, then the post-measurement state is

$$\psi \longmapsto \frac{E_A(\omega)\psi}{\| E_A(\omega)\psi \|} \overset{(4.2)}{=} \frac{E_A(\omega)\psi}{\sqrt{\langle \psi | E_A(\omega) | \psi \rangle}} \tag{5.4}$$

provided $E_A(\omega)\psi \neq 0$.

Example 5.2 (Finite-dimensional case)**.** If $A \colon \mathbb{C}^n \to \mathbb{C}^n$ is self-adjoint, then by the (finite-dimensional) spectral theorem there exist distinct real eigenvalues. Since this is explicitly the finite-dimensional case, all sums in this example are finite. $\lambda_1, \ldots, \lambda_r$ and pairwise orthogonal eigenspaces

$$M_j \coloneqq N(A - \lambda_j \, \mathrm{id})$$

for $j = 1, \ldots, r$, such that

$$\mathbb{C}^n = \bigoplus_{j=1}^r M_j$$

and

$$\sigma(A) = \{\lambda_1, \ldots, \lambda_r\}$$

Let E_A be the spectral measure of A from Theorem 4.17. Since the spectrum is finite, the spectral integral reduces to a finite sum over the atoms $\{\lambda_j\}$:

$$A = \int_{\sigma(A)} \lambda \, \mathrm{d}E_A(\lambda) = \sum_{j=1}^r \lambda_j \, E_A(\{\lambda_j\}) \tag{5.5}$$

By Proposition 4.16, $E_A(\{\lambda_j\}) = P_{M_j}$, where P_{M_j} denotes the orthogonal projection onto the eigenspace M_j. Hence

$$A = \sum_{j=1}^{r} \lambda_j P_{M_j} \tag{5.6}$$

If we choose an orthonormal basis of eigenvectors in each eigenspace M_j,

$$\{\varphi_{j,k} : 1 \le j \le r,\ 1 \le k \le d_j\}$$

with $d_j = \dim M_j$, $M_j = \operatorname{span}\{\varphi_{j,1}, \dots, \varphi_{j,d_j}\}$, the orthogonal projection onto M_j is the sum of the one-dimensional projections onto these basis vectors. Hence

$$P_{M_j}\psi = \sum_{k=1}^{d_j} P_{\varphi_{j,k}}\psi = \sum_{k=1}^{d_j} \langle \varphi_{j,k}|\psi\rangle\, \varphi_{j,k}$$

and, by orthonormality,

$$\|P_{M_j}\psi\|^2 = \sum_{k=1}^{d_j} |\, \langle \varphi_{j,k}|\psi\rangle\, |^2 \tag{5.7}$$

Substituting (5.6) into the expectation value gives

$$\begin{aligned}
\langle A\rangle_\psi &= \langle \psi | A\psi\rangle \\
&= \left\langle \psi \middle| \sum_{j=1}^{r} \lambda_j P_{M_j}\psi \right\rangle \\
&= \sum_{j=1}^{r} \lambda_j\ \langle \psi | P_{M_j}\psi\rangle \\
&= \sum_{j=1}^{r} \lambda_j\ \|P_{M_j}\psi\|^2
\end{aligned}$$

Thus the numbers $p_j \coloneqq \|P_{M_j}\psi\|^2$ are precisely the probabilities of obtaining the outcomes λ_j. So in the finite-dimensional case the expectation value has the familiar form $\langle A\rangle_\psi = \sum_{j=1}^{r} \lambda_j p_j$.

The general formulas from the main text hence reduce to these familiar finite-dimensional discrete ones. In particular, (5.5) shows explicitly how the spectral integral collapses to a finite sum, and Proposition 4.16 identifies the spectral projections with the orthogonal projections onto the eigenspaces. Indeed, given a Borel set $\omega \subseteq \sigma(A)$, we have

$$E_A(\omega) = E_A\left(\bigcup_{\lambda_j \in \omega} \{\lambda_j\}\right) \overset{4.6.(\mathrm{d})}{=} \sum_{\lambda_j \in \omega} E_A(\{\lambda_j\}) \overset{4.16}{=} \sum_{\lambda_j \in \omega} P_{M_j} \tag{5.8}$$

Here the sum is finite because $\sigma(A) = \{\lambda_1, \dots, \lambda_r\}$ is finite and hence $\omega \subseteq \sigma(A)$ is so as well. Therefore

$$\begin{aligned}
\mu_\psi^A(\omega) &\overset{(5.3)}{=} \langle \psi | E_A(\omega)\psi\rangle \\
&\overset{(5.8)}{=} \left\langle \psi \middle| \sum_{\lambda_j \in \omega} P_{M_j}\psi \right\rangle \\
&= \sum_{\lambda_j \in \omega} \langle \psi | P_{M_j}\psi\rangle \\
&\overset{4.3}{=} \sum_{\lambda_j \in \omega} \|P_{M_j}\psi\|^2 \\
&\overset{(5.7)}{=} \sum_{\lambda_j \in \omega} \sum_{k=1}^{d_j} |\, \langle \varphi_{j,k}|\psi\rangle\, |^2
\end{aligned}$$

which recovers exactly the usual formula. In particular,

$$\mathbb{P}_A^\psi(\{\lambda_j\}) = \|P_{M_j}\psi\|^2 = \sum_{k=1}^{d_j} |\, \langle \varphi_{j,k}|\psi\rangle\, |^2 \tag{5.9}$$

If the measurement outcome is the eigenvalue λ_j, then the post-measurement state from (5.4) becomes

$$\psi \longmapsto \frac{P_{M_j}\psi}{\|P_{M_j}\psi\|}$$

provided $P_{M_j}\psi \neq 0$. Thus the state collapses to the normalized component of ψ in the eigenspace M_j. ◀

5.4 Joint systems and composition

If two quantum systems are described by Hilbert spaces $\mathcal{H}_A$ and $\mathcal{H}_B$, then the composite system is described by the tensor product Hilbert space $\mathcal{H}_A \otimes \mathcal{H}_B$.

The tensor product is linear in each factor, so it is compatible with the superposition principle: superpositions in the subsystems induce corresponding superpositions in the composite system.

If the two subsystems are in pure states $\psi_A \in \mathcal{H}_A$ and $\psi_B \in \mathcal{H}_B$, then the corresponding product state of the joint system is

$$\psi_A \otimes \psi_B \in \mathcal{H}_A \otimes \mathcal{H}_B$$

More generally, the composite Hilbert space also contains superpositions of product states. Not every state in $\mathcal{H}_A \otimes \mathcal{H}_B$ can be written in the form $\psi_A \otimes \psi_B$. States that cannot be written in this form are called *entangled*.

5.5 Abstraction and simplification

From now on, we restrict to the case where the Hilbert space is finite-dimensional. Then

$$\mathcal{H} \cong \mathbb{C}^n$$

for some $n \in \mathbb{N}$, i.e. $\mathcal{H}$ is isomorphic to the standard n-dimensional complex vector space.

From this point on, all basis expansions are therefore finite sums. This allows us to represent states by coordinates with respect to some fixed orthonormal basis and hence to work with vectors and matrices instead of general Hilbert-space operators.

Theorem 5.1. Let $\underline{\boldsymbol{A}} \in \mathbb{C}^{n\times n}$ be Hermitian. Then there exists a unitary matrix $\underline{\boldsymbol{U}}$ and a real diagonal matrix $\underline{\boldsymbol{\Lambda}}$ such that

$$\underline{\boldsymbol{A}} = \underline{\boldsymbol{U}}\,\underline{\boldsymbol{\Lambda}}\,\underline{\boldsymbol{U}}^* \tag{5.10}$$

where $\underline{\boldsymbol{\Lambda}} = \operatorname{diag}(\lambda_1, \ldots, \lambda_n)$ is a diagonal matrix of eigenvalues listed with algebraic multiplicity. Equivalently, $\underline{\boldsymbol{A}}$ admits an orthonormal basis of eigenvectors. ◁

For $\psi \in \mathcal{H} \cong \mathbb{C}^n$, the corresponding column vector

$$|\psi\rangle = \begin{bmatrix} \psi_1 \\ \vdots \\ \psi_n \end{bmatrix}$$

is called a *ket*. Its adjoint (complex conjugated transpose) is the row vector

$$\langle\psi| = |\psi\rangle^* = \begin{bmatrix} \psi_1^*, \ldots, \psi_n^* \end{bmatrix}$$

called a *bra*. The inner product becomes

$$\langle\psi_A|\psi_B\rangle = \langle\psi_A|\,|\psi_B\rangle = \sum_{i=1}^{n} \psi_{A,i}^* \psi_{B,i} \in \mathbb{C}$$

called a *bra-ket*.

The notation is also useful for expressing Hermitian observables. As seen in Example 5.2, if A is self-adjoint, we have

$$A = \sum_{j=1}^{r} \lambda_j P_{M_j} = \sum_{j=1}^{r} \lambda_j \sum_{k=1}^{d_j} P_{\varphi_{j,k}}$$

and therefore

$$\begin{aligned} |A\psi\rangle &= \sum_{j=1}^{r} \lambda_j \sum_{k=1}^{d_j} \left|P_{\varphi_{j,k}}\psi\right\rangle \\ &= \sum_{j=1}^{r} \lambda_j \sum_{k=1}^{d_j} \langle\varphi_{j,k}|\psi\rangle\,|\varphi_{j,k}\rangle \\ &= \sum_{j=1}^{r} \lambda_j \sum_{k=1}^{d_j} |\varphi_{j,k}\rangle\,\langle\varphi_{j,k}|\,|\psi\rangle \end{aligned}$$

whereby we obtained the common way to represent projectors as

$$P_\varphi = |\varphi\rangle \langle\varphi|$$

and thus the observable A itself as $A = \sum_{j=1}^{r} \lambda_j \sum_{k=1}^{d_j} |\varphi_{j,k}\rangle \langle\varphi_{j,k}|$.

Fix an orthonormal basis $\{|\varphi_i\rangle\}_{i=1}^n$ of $\mathcal{H}$. In the following, we abbreviate the basis vectors with their indices, i.e. $|\varphi_i\rangle =: |i\rangle$.

A common trick when working with an orthonormal basis in braket notation is to insert the identity operator

$$\mathrm{id} = \sum_{i=1}^{n} |i\rangle \langle i| \tag{5.11}$$

and then rearrange the terms. Note that this is a discrete instance of the resolution of the identity from Definition 4.6. For example, we can write a linear operator A as

$$A = \left(\sum_{k=1}^{n} |k\rangle \langle k|\right) A \left(\sum_{l=1}^{n} |l\rangle \langle l|\right) = \sum_{k=1}^{n}\sum_{l=1}^{n} |k\rangle \langle k|A|l\rangle \langle l| = \sum_{k=1}^{n}\sum_{l=1}^{n} \underbrace{\langle k|A|l\rangle}_{[\underline{\boldsymbol{A}}]_{kl}} |k\rangle \langle l|$$

hence, every linear operator A can be represented as a matrix $\underline{\boldsymbol{A}}$. The entry in the k^{th} row and l^{th} column is $[\underline{\boldsymbol{A}}]_{kl} = \langle k|A|l\rangle$.

The way in which it is written here is rather an expansion in a basis of the vector space of linear operators $\mathrm{End}(\mathcal{H})$, and $\langle k|A|l\rangle$ is the coefficient before the basis vector $|k\rangle \langle l|$, i.e., a matrix with a one in the k^{th} row and the l^{th} column.

5.6 The trace

Definition 5.1. For an operator $A \in \mathrm{End}(\mathcal{H})$ on a finite-dimensional Hilbert space, the trace is defined by

$$\mathrm{tr}(A) \coloneqq \sum_{i=1}^{n} \langle i|A|i\rangle \tag{5.12}$$

where $\{|i\rangle\}_{i=1}^n$ is any orthonormal basis. ◀

The trace can be generalized to infinite-dimensional spaces, at least for certain, so called *trace-class*, operators:

Definition 5.2 (Trace Class)**.** Let $\mathcal{H}$ be a separable Hilbert space with orthonormal basis $\{|i\rangle\}_{i=1}^\infty$. A bounded operator $A \in B(\mathcal{H})$ is called *trace class* if

$$\sum_{i=1}^{\infty} \langle i| \, |A| \, |i\rangle < \infty$$

where $|A| \coloneqq \sqrt{A^* A}$ denotes the positive-semidefinite Hermitian square root, i.e., the unique positive operator B such that $B^2 = A^* A$. In that case, the trace of A is defined by

$$\mathrm{tr}(A) \coloneqq \sum_{i=1}^{\infty} \langle i|A|i\rangle$$

and this series converges absolutely and is independent of the choice of orthonormal basis. ◀

In particular, finite-rank operators are trace class. Hence rank-one operators such as

$$|\psi_1\rangle \langle\psi_2|$$

are trace class, which is why the density-matrix formalism still works in infinite-dimensional Hilbert spaces for pure states and, more generally, for trace-class density operators.

Interlude 5.3 (Change of Basis)**.** Fix an "old" orthonormal basis $B = (\boldsymbol{e}_1, \ldots, \boldsymbol{e}_n)$ and a "new" orthonormal basis $\widetilde{B} = (\tilde{\boldsymbol{e}}_1, \ldots, \tilde{\boldsymbol{e}}_n)$ of $\mathcal{H}$. Let $\boldsymbol{x}$ be the coordinate column of a vector $v \in \mathcal{H}$ with respect to B, and let $\tilde{\boldsymbol{x}}$ be the coordinate column of the same vector with respect to $\widetilde{B}$. If $\underline{\boldsymbol{U}}$ denotes the change-of-basis matrix from B to $\widetilde{B}$, then

$$\tilde{\boldsymbol{x}} = \underline{\boldsymbol{U}} \boldsymbol{x}$$

where the columns of $\underline{\boldsymbol{U}}$ are the old basis vectors written in the new basis. Equivalently, the columns of $\underline{\boldsymbol{U}}^{-1}$ are the new basis vectors written in the old basis. If a linear operator $\hat{A}\colon \mathcal{H} \to \mathcal{H}$ has matrix $\underline{\boldsymbol{A}} \in \mathbb{C}^{n\times n}$ in the old basis and matrix $\underline{\tilde{\boldsymbol{A}}} \in \mathbb{C}^{n\times n}$ in the new basis, then

$$\underline{\tilde{\boldsymbol{A}}} = \underline{\boldsymbol{U}}\,\underline{\boldsymbol{A}}\,\underline{\boldsymbol{U}}^{-1}$$

where $\underline{\boldsymbol{U}}$ is the change-of-basis matrix from the old basis B to the new basis $\widetilde{B}$.

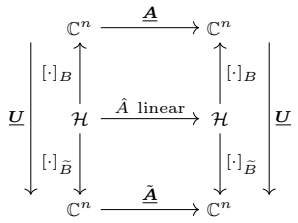

So matrices of the same operator in different bases are related by conjugation. In the orthonormal case, $\underline{\boldsymbol{U}}$ is unitary, hence $\underline{\boldsymbol{U}}^{-1} = \underline{\boldsymbol{U}}^*$. ◀

The trace is basis-independent. To see this, recall that $\operatorname{tr}(AB) = \operatorname{tr}(BA)$. Hence for a unitary change of basis U,

$$\operatorname{tr}(UAU^*) = \operatorname{tr}(U^*UA) = \operatorname{tr}(A) \tag{5.13}$$

This can also be seen using the trick mentioned above:

$$\begin{aligned}
\operatorname{tr}(A) &= \sum_{i=1}^{n} \langle i|A|i\rangle = \sum_{i=1}^{n} \langle i|\,\mathrm{id}\,A\,\mathrm{id}\,|i\rangle \\
&= \sum_{i=1}^{n} \langle i| \left(\sum_{\tilde{u}=1}^{n} |\tilde{u}\rangle\langle\tilde{u}|\right) A \left(\sum_{\tilde{v}=1}^{n} |\tilde{v}\rangle\langle\tilde{v}|\right) |i\rangle \\
&= \sum_{i=1}^{n}\sum_{\tilde{u}=1}^{n}\sum_{\tilde{v}=1}^{n} \langle i|\tilde{u}\rangle\langle\tilde{u}|A|\tilde{v}\rangle\langle\tilde{v}|i\rangle \\
&= \sum_{\tilde{u}=1}^{n}\sum_{\tilde{v}=1}^{n}\sum_{i=1}^{n} \langle\tilde{v}|i\rangle\langle i|\tilde{u}\rangle\langle\tilde{u}|A|\tilde{v}\rangle \\
&= \sum_{\tilde{u}=1}^{n}\sum_{\tilde{v}=1}^{n} \langle\tilde{v}| \underbrace{\left(\sum_{i=1}^{n} |i\rangle\langle i|\right)}_{\mathrm{id}} |\tilde{u}\rangle\langle\tilde{u}|A|\tilde{v}\rangle \\
&= \sum_{\tilde{u}=1}^{n}\sum_{\tilde{v}=1}^{n} \underbrace{\langle\tilde{v}|\tilde{u}\rangle}_{\delta_{\tilde{v},\tilde{u}}} \langle\tilde{u}|A|\tilde{v}\rangle \\
&= \sum_{\tilde{u}=1}^{n} \langle\tilde{u}|A|\tilde{u}\rangle
\end{aligned}$$

For a pure state $|\psi\rangle$ and an observable A, the expectation value can be written as

$$\begin{aligned}
\langle A\rangle_\psi &\overset{(5.2)}{=} \langle\psi|A|\psi\rangle \\
&= \langle\psi|\,A \left(\sum_{i=1}^{n} |i\rangle\langle i|\right) |\psi\rangle \\
&= \sum_{i=1}^{n} \langle\psi|A|i\rangle\langle i|\psi\rangle \\
&= \sum_{i=1}^{n} \langle i|\,(|\psi\rangle\langle\psi|\,A)\,|i\rangle \\
&\overset{(5.12)}{=} \operatorname{tr}\big(|\psi\rangle\langle\psi|\,A\big)
\end{aligned}$$

which is the form that extends naturally to mixed states.

5.7 Density matrices

So far, we have described quantum states by normalized vectors $\psi \in \mathcal{H}$, i.e. by *pure states*. In practice, however, one often does not know the actual pure state exactly, but only a probability distribution over possible pure states. Such a *statistical mixture* is **not** the same thing as a superposition. While the vector representation is sufficient for pure states, as soon as we want to describe statistical mixtures, it is more convenient to use an equivalent representation in terms of density matrices.

Definition 5.3 (Density Matrix)**.** A density matrix on a finite-dimensional Hilbert space $\mathcal{H}$ is an operator $\rho \in \mathrm{End}(\mathcal{H})$ such that

(a) $\rho \geq 0$, i.e. $\langle\psi|\rho|\psi\rangle \geq 0$ for every $\psi \in \mathcal{H}$

(b) $\mathrm{tr}(\rho) = 1$ ◀

Since ρ is positive, it is in particular self-adjoint, and by the spectral theorem (Theorem 4.17; in the finite-dimensional case compare also Example 5.2) it admits a decomposition

$$\rho = \sum_{j=1}^{r} \lambda_j P_{M_j} = \sum_{j=1}^{r} \lambda_j \sum_{k=1}^{d_j} |\varphi_{j,k}\rangle \langle\varphi_{j,k}|$$

where $\lambda_1, \ldots, \lambda_r \geq 0$ are the distinct eigenvalues of ρ, M_j is the eigenspace corresponding to λ_j,

$$P_{M_j} = \sum_{k=1}^{d_j} P_{\varphi_{j,k}} = \sum_{k=1}^{d_j} |\varphi_{j,k}\rangle \langle\varphi_{j,k}| \tag{5.14}$$

is the orthogonal projection onto M_j, and $\{\varphi_{j,1}, \ldots, \varphi_{j,d_j}\}$ is an orthonormal basis of M_j, with $d_j = \dim M_j$. Since

$$\mathrm{tr}(\rho) = \sum_{j=1}^{r} \lambda_j \,\mathrm{tr}(P_{M_j}) = \sum_{j=1}^{r} d_j \lambda_j = 1$$

the numbers $d_j \lambda_j$ add up to one. Equivalently, if one lists eigenvalues with multiplicity, then

$$\rho = \sum_{i=1}^{n} \lambda_i |\varphi_i\rangle \langle\varphi_i|$$

with $\lambda_i \geq 0$ and $\sum_{i=1}^{n} \lambda_i = 1$ for some orthonormal basis of eigenvectors $\{\varphi_i\}_{i=1}^{n}$ (compare also Example 5.2). Thus the eigenvalues, counted with multiplicity, can be interpreted as probabilities of the system to be in the state of the corresponding eigenvector.

Definition 5.4 (Mixed state)**.** If the system is in the pure state $|\varphi_i\rangle$ with probability p_i, then the corresponding density matrix is exactly

$$\rho = \sum_i p_i |\varphi_i\rangle \langle\varphi_i| \tag{5.15}$$

where $p_i \geq 0$ and $\sum_i p_i = 1$. This is called a *mixed state* or *statistical mixture*. ◀

Caution 5.4. A normalized superposition $\alpha|\psi_1\rangle + \beta|\psi_2\rangle$ with $|\alpha|^2 + |\beta|^2 = 1$ is again a pure state. By contrast, $p|\psi_1\rangle\langle\psi_1| + (1-p)|\psi_2\rangle\langle\psi_2|$ is generally a mixed state. So convex combinations of projectors are fundamentally different from linear combinations of vectors. ◀

The time evolution of density matrices is induced by the unitary time evolution $U(t) = e^{-\imath Ht/\hbar}$ of pure states (Section 5.2):

$$\rho(t) = U(t)\,\rho(0)\,U(t)^* \tag{5.16}$$

For an observable A, the probability of obtaining a measurement outcome in a Borel set $\omega \subseteq \sigma(A)$ is

$$\boxed{\mathbb{P}_A^\rho(\omega) = \mathrm{tr}\big(E_A(\omega)\rho\big)} \tag{5.17}$$

where E_A is the spectral measure of A from Theorem 4.17. In the finite-dimensional case, if

$$A = \sum_{j=1}^{r} \lambda_j P_{M_j}$$

with distinct eigenvalues λ_j and eigenspaces M_j, then

$$\mathbb{P}_A^\rho(\{\lambda_j\}) = \operatorname{tr}\left(P_{M_j}\rho\right) \overset{(5.14)}{=} \operatorname{tr}\left(\sum_{k=1}^{d_j} |\varphi_{j,k}\rangle\langle\varphi_{j,k}|\,\rho\right) \tag{5.18}$$

So in the degenerate case one must use the projection onto the whole eigenspace, not merely a (rank-one) projector onto a single eigenvector.

Definition 5.5 (Pure State as a Density Matrix)**.** A density matrix ρ is called *pure* if

$$\rho = |\psi\rangle\langle\psi|$$

for some normalized $\psi \in \mathcal{H}$. ◀

Equivalently, ρ is pure if and only if one eigenvalue equals 1 and all others are 0, or, equivalently, if and only if ρ has rank 1. Thus pure states are precisely the rank-one orthogonal projections (compare also Equation 5.4 and Example 5.2).

Overview 5.5 (pure vs. mixed states)**.** A density operator ρ is

- pure iff $\rho^2 = \rho$ (i.e. ρ is a rank-1 projection)
- mixed iff $\rho^2 \neq \rho$

Equivalently,

- pure iff $\operatorname{tr}(\rho^2) = 1$
- mixed iff $\operatorname{tr}(\rho^2) < 1$

In terms of the eigenvalues of ρ:

- pure iff ρ has exactly one eigenvalue equal to 1 and all others equal to 0
- mixed iff ρ has more than one nonzero eigenvalue ◀

For the density matrix of a pure state $\rho = |\psi\rangle\langle\psi|$, (5.17) becomes

$$\mathbb{P}_A^\rho(\omega) \overset{(5.12)}{=} \sum_{i=1}^{n} \langle i|\,E_A(\omega)\,|\psi\rangle\langle\psi|\,|i\rangle = \langle\psi|\left(\sum_{i=1}^{n}|i\rangle\langle i|\right)E_A(\omega)\,|\psi\rangle = \langle\psi|\,E_A(\omega)\,|\psi\rangle$$

which is consistent with (5.3).

For a pure state $|\psi\rangle$ and an observable A, the expectation value can be written as

$$\langle A\rangle_\psi = \operatorname{tr}\left(|\psi\rangle\langle\psi|\,A\right)$$

as seen in Section 5.6. This motivates the general formula (compare also Equation 5.2)

$$\boxed{\langle A\rangle_\rho = \operatorname{tr}(\rho A)} \tag{5.19}$$

for arbitrary density matrices ρ.

Definition 5.6 (Separability and Entanglement)**.** Let $\mathcal{H}_A$ and $\mathcal{H}_B$ be Hilbert spaces. A pure state $\psi \in \mathcal{H}_A \otimes \mathcal{H}_B$ is called *separable* if it can be written as

$$\psi = \varphi_A \otimes \varphi_B$$

for some $\varphi_A \in \mathcal{H}_A$ and $\varphi_B \in \mathcal{H}_B$. Otherwise it is called *entangled* (compare Section 1.5). ◀

Example 5.6. The singlet state $|\Psi^-\rangle$ from Section 1.5 is entangled. It is a superposition (normalized linear combination) of two pure states and hence a pure state itself. ◀

Digression 5.7 (Tensor products)**.** The tensor product combines two vector spaces V and W into a larger space

$$V \otimes W$$

whose elements are finite linear combinations of simple tensors

$$|v\rangle \otimes |w\rangle$$

with $|v\rangle \in V$, $|w\rangle \in W$. If $\dim V = m$ and $\dim W = n$, then

$$\dim(V \otimes W) = mn$$

Moreover, if $\{|i\rangle\}_i$ is a basis of V and $\{|j\rangle\}_j$ is a basis of W, then

$$\{|i\rangle \otimes |j\rangle\}_{i,j}$$

is a basis of $V \otimes W$. In Dirac notation one often abbreviates

$$|v\rangle \otimes |w\rangle = |v\rangle\,|w\rangle = |v, w\rangle = |vw\rangle$$

The tensor product is bilinear, i.e.

$$z(|v\rangle \otimes |w\rangle) = (z\,|v\rangle) \otimes |w\rangle = |v\rangle \otimes (z\,|w\rangle)$$
$$(|v_1\rangle + |v_2\rangle) \otimes |w\rangle = |v_1\rangle \otimes |w\rangle + |v_2\rangle \otimes |w\rangle$$
$$|v\rangle \otimes (|w_1\rangle + |w_2\rangle) = |v\rangle \otimes |w_1\rangle + |v\rangle \otimes |w_2\rangle$$

If $A\colon V \to V'$ and $B\colon W \to W'$ are linear maps, their tensor product

$$A \otimes B\colon V \otimes W \to V' \otimes W'$$

is defined by

$$(A \otimes B)(|v\rangle \otimes |w\rangle) = A\,|v\rangle \otimes B\,|w\rangle$$

and extended linearly to all of $V \otimes W$.

If V and W are Hilbert spaces, then $V \otimes W$ becomes a Hilbert space with inner product determined by

$$\boxed{\langle v_1 \otimes w_1 \,|\, v_2 \otimes w_2\rangle = \langle v_1|v_2\rangle\ \langle w_1|w_2\rangle} \tag{5.20}$$

and extended sesquilinearly. Hence, if $\{|i\rangle\}_i$ and $\{|j\rangle\}_j$ are orthonormal bases of V and W, then $\{|i\rangle \otimes |j\rangle\}_{i,j}$ is an orthonormal basis of $V \otimes W$.

For the trace of the tensor product of two operators $A \in \mathrm{End}(V)$ and $B \in \mathrm{End}(W)$ one has

$$\begin{aligned}
\mathrm{tr}(A \otimes B) &= \sum_{i,j} \langle i \otimes j|A \otimes B|i \otimes j\rangle \\
&= \sum_{i,j} \langle i|A|i\rangle\ \langle j|B|j\rangle \\
&= \left(\sum_i \langle i|A|i\rangle\right)\left(\sum_j \langle j|B|j\rangle\right) \\
&= \mathrm{tr}(A)\ \mathrm{tr}(B)
\end{aligned}$$

In matrix form, the tensor product is represented by the Kronecker product. For matrices $\underline{\boldsymbol{A}} = (a_{ij})$, $\underline{\boldsymbol{B}}$, one has

$$\underline{\boldsymbol{A}} \otimes \underline{\boldsymbol{B}} = \begin{bmatrix} a_{11}\underline{\boldsymbol{B}} & \cdots & a_{1n}\underline{\boldsymbol{B}} \\ \vdots & \ddots & \vdots \\ a_{m1}\underline{\boldsymbol{B}} & \cdots & a_{mn}\underline{\boldsymbol{B}} \end{bmatrix}$$

Useful identities include $(\underline{\boldsymbol{A}} \otimes \underline{\boldsymbol{B}})^* = \underline{\boldsymbol{A}}^* \otimes \underline{\boldsymbol{B}}^*$. In particular, if $\underline{\boldsymbol{A}}$ and $\underline{\boldsymbol{B}}$ are unitary, Hermitian, positive, or projectors, then $\underline{\boldsymbol{A}} \otimes \underline{\boldsymbol{B}}$ has the same property. ◀

However, parts of an entangled state are mixed states. To understand what Alice's part of the entangled system from Example 5.6 looks like, we introduce the

Definition 5.7 (Partial Trace)**.** To describe the state of only one subsystem, we use the linear map

$$\mathrm{tr}_B\colon \mathrm{End}(\mathcal{H}_A \otimes \mathcal{H}_B) \to \mathrm{End}(\mathcal{H}_A)$$

defined on simple tensors by

$$\mathrm{tr}_B(S_A \otimes S_B) \coloneqq \mathrm{tr}(S_B)\, S_A$$

and extended linearly. Analogously one defines tr_A. ◀

Example 5.8. The density matrix of the singlet state is

$$\begin{aligned}\rho = |\Psi^-\rangle\langle\Psi^-| &= \frac{1}{\sqrt{2}}(|01\rangle - |10\rangle)\frac{1}{\sqrt{2}}(\langle 01| - \langle 10|) \\ &= \frac{1}{2}\big(|01\rangle\langle 01| - |01\rangle\langle 10| - |10\rangle\langle 01| + |10\rangle\langle 10|\big)\end{aligned}$$

Using

$$\operatorname{tr}(|i\rangle\langle j|) \overset{(5.12)}{=} \sum_k \langle k|i\rangle\langle j|k\rangle = \langle j|\left(\sum_k |k\rangle\langle k|\right)|i\rangle = \langle j|i\rangle = \delta_{ji}$$

we get

$$\begin{aligned}&\operatorname{tr}_B\big(|\Psi^-\rangle\langle\Psi^-|\big) \\ &= \frac{1}{2}\Big(\operatorname{tr}_B(|01\rangle\langle 01|) - \operatorname{tr}_B(|01\rangle\langle 10|) - \operatorname{tr}_B(|10\rangle\langle 01|) + \operatorname{tr}_B(|10\rangle\langle 10|)\Big) \\ &= \frac{1}{2}\Big(\operatorname{tr}_B(|0\rangle\langle 0|\otimes|1\rangle\langle 1|) - \operatorname{tr}_B(|0\rangle\langle 1|\otimes|1\rangle\langle 0|) - \operatorname{tr}_B(|1\rangle\langle 0|\otimes|0\rangle\langle 1|) + \operatorname{tr}_B(|1\rangle\langle 1|\otimes|0\rangle\langle 0|)\Big) \\ &= \frac{1}{2}\Big(\operatorname{tr}(|1\rangle\langle 1|)\ |0\rangle\langle 0| - \operatorname{tr}(|1\rangle\langle 0|)\ |0\rangle\langle 1| - \operatorname{tr}(|0\rangle\langle 1|)\ |1\rangle\langle 0| + \operatorname{tr}(|0\rangle\langle 0|)\ |1\rangle\langle 1|\Big) \\ &= \frac{1}{2}\big(|0\rangle\langle 0| + |1\rangle\langle 1|\big) = \frac{\mathrm{id}_A}{2}\end{aligned}$$

and similarly $\operatorname{tr}_A\big(|\Psi^-\rangle\langle\Psi^-|\big) = \frac{1}{2}\big(|0\rangle\langle 0| + |1\rangle\langle 1|\big) = \frac{\mathrm{id}_B}{2}$. So each subsystem is in the *maximally mixed state*.

Importantly, the partial trace contains all information about local measurements. If Alice performs a measurement corresponding to an observable A on her subsystem and E_A denotes its spectral measure from Theorem 4.17, then the event that Alice's measurement outcome lies in a Borel set $\omega \subseteq \sigma(A)$ is represented on the joint system by the projection $E_A(\omega)\otimes \mathrm{id}_B$ since the measurement acts on Alice's subsystem and trivially (meaning as the identity, i.e. Bob's subsystem is left unchanged) on Bob's. Therefore, by the Born rule (5.17),

$$\mathbb{P}^\rho_A(\omega) = \operatorname{tr}\big((E_A(\omega)\otimes \mathrm{id}_B)\rho\big)$$

which can be split into a partial trace over Bob's subsystem and a subsequent trace over Alice's subsystem,

$$\mathbb{P}^\rho_A(\omega) = \operatorname{tr}\big(E_A(\omega)\operatorname{tr}_B(\rho)\big)$$

where we used Theorem 5.2. Hence Bob's subsystem influences Alice's local statistics only through the reduced density matrix $\operatorname{tr}_B(\rho)$. In particular, local operations on Bob's side alone cannot be used to transmit information to Alice, saving us from problems with relativity. ◀

Theorem 5.2 (Partial trace and local observables)**.** Let $\mathcal{H}_A$ and $\mathcal{H}_B$ be finite-dimensional Hilbert spaces. Let $\rho \in \operatorname{End}(\mathcal{H}_A\otimes\mathcal{H}_B)$ and $A \in \operatorname{End}(\mathcal{H}_A)$. Then

- $\operatorname{tr}_B\big((A\otimes\mathrm{id}_B)\rho\big) = A\ \operatorname{tr}_B(\rho)$
- $\operatorname{tr}\big((A\otimes\mathrm{id}_B)\rho\big) = \operatorname{tr}_A\big(A\ \operatorname{tr}_B(\rho)\big)$
- $\operatorname{tr}_B\big(\rho(A\otimes\mathrm{id}_B)\big) = \operatorname{tr}_B(\rho)\,A$ ◁

6 Qbits

We now consider finite-dimensional quantum systems. Since quantum theory is linear, restricting to such spaces does not alter the underlying physics. The basic unit is the quantum bit, or Qbit, corresponding to a 2-dimensional Hilbert space. Unlike classical bits, however, understanding a single Qbit is not enough to understand systems of several Qbits, because qualitatively new phenomena arise there, most importantly entanglement, which lies at the heart of the power of quantum information processing.

6.1 One Qbit: $\mathcal{H} = \mathbb{C}^2$

A quantum bit, or Qbit, or Qbit is a quantum system whose state space is

$$\mathcal{H} = \mathbb{C}^2$$

We choose the standard orthonormal basis

$$|0\rangle = \begin{bmatrix} 1 \\ 0 \end{bmatrix} \qquad |1\rangle = \begin{bmatrix} 0 \\ 1 \end{bmatrix}$$

called the *computational basis*.

A general pure state of one Qbit is

$$|\psi\rangle = \alpha\,|0\rangle + \beta\,|1\rangle$$

with $\alpha, \beta \in \mathbb{C}$, $|\alpha|^2 + |\beta|^2 = 1$.

At first sight this seems to involve four real parameters minus one normalization condition, hence three real degrees of freedom. But one more degree of freedom is physically irrelevant: If $c \in \mathbb{C}$, $|c| = 1$, then

$$(c\,|\psi\rangle)(c\,|\psi\rangle)^* = c\,|\psi\rangle\,\langle\psi|\,\overline{c} = |\psi\rangle\,\langle\psi|$$

Thus $|\psi\rangle$ and $c\,|\psi\rangle$ represent the same physical pure state. Global phase factors cannot be observed.

Mathematically, two vectors are said to be equivalent if and only of they differ simply by a global phase factor. Quantum states then correspond to equivalence relations with respect to that relation, sometimes called "unit rays".

So a pure one-Qbit state is really determined by only two real parameters. A convenient parametrization is

$$\begin{aligned} |\psi\rangle &= \alpha\,|0\rangle + \beta\,|1\rangle \\ &= e^{\imath\varphi_\alpha}|\alpha|\,|0\rangle + e^{\imath\varphi_\beta}|\beta|\,|1\rangle \\ &= e^{\imath\varphi_\alpha}\left(|\alpha|\,|0\rangle + e^{\imath(\varphi_\beta-\varphi_\alpha)}|\beta|\,|1\rangle\right) \\ &\sim |\alpha|\,|0\rangle + e^{\imath(\varphi_\beta-\varphi_\alpha)}|\beta|\,|1\rangle \\ &= \cos\frac{\Theta}{2}\,|0\rangle + e^{\imath\varphi}\sin\frac{\Theta}{2}\,|1\rangle \end{aligned}$$

with $\Theta \in [0, \pi]$ and *relative phase* $\varphi \in [0, 2\pi)$. A representation of states which takes into account the irrelevance of global phases is the Bloch sphere. In this representation of the unit rays, or projectors, "orthogonal" becomes "antipodal".

This equals the coordinate system on earth (including the fact that the poles do not have a well-defined longitude).

Orientation on the Bloch sphere:

- Standard / Computational / Z Basis: $|0\rangle$, $|1\rangle$
- Diagonal / Hadamard / X Basis: $|+\rangle = \frac{|0\rangle+|1\rangle}{\sqrt{2}}$, $|-\rangle = \frac{|0\rangle-|1\rangle}{\sqrt{2}}$
- Circular / Y Basis: $|\imath\rangle = \frac{|0\rangle+\imath|1\rangle}{\sqrt{2}}$, $|-\imath\rangle = \frac{|0\rangle-\imath|1\rangle}{\sqrt{2}}$

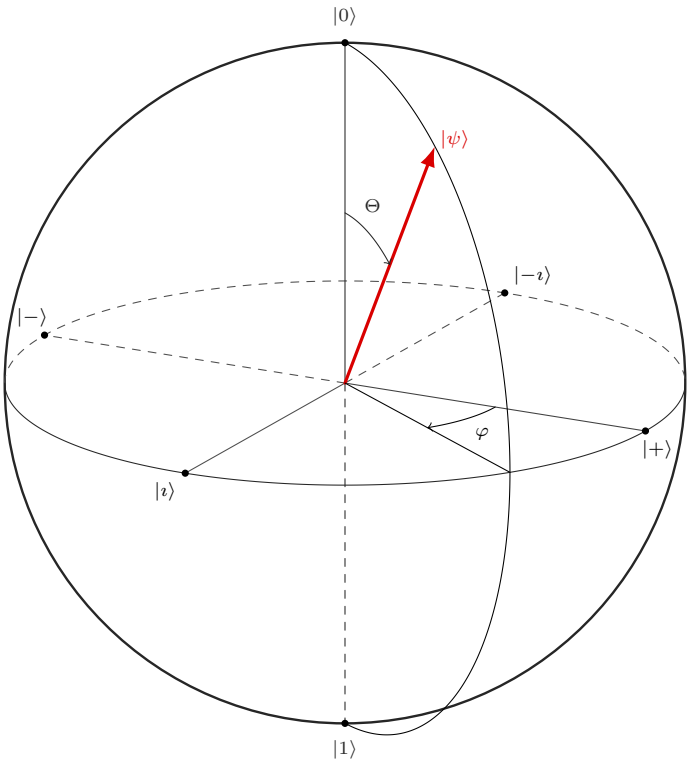

While the surface of the Bloch sphere represents pure states, the interior corresponds to mixed states [7]. To see this, recall that a mixed state is represented by a density operator ρ. Any two-dimensional density operator ρ can be expanded using the identity $\underline{\boldsymbol{I}}$ and the Hermitian, traceless Pauli matrices $\boldsymbol{\sigma} = [\underline{\boldsymbol{\sigma}}_x, \underline{\boldsymbol{\sigma}}_y, \underline{\boldsymbol{\sigma}}_z]$,

$$\begin{aligned}\rho &= \frac{1}{2}(\underline{\boldsymbol{I}} + \boldsymbol{a}\cdot\boldsymbol{\sigma}) \\ &= \frac{1}{2}\begin{bmatrix}1 & 0\\ 0 & 1\end{bmatrix} + \frac{a_x}{2}\begin{bmatrix}0 & 1\\ 1 & 0\end{bmatrix} + \frac{a_y}{2}\begin{bmatrix}0 & -\imath\\ \imath & 0\end{bmatrix} + \frac{a_z}{2}\begin{bmatrix}1 & 0\\ 0 & -1\end{bmatrix} \\ &= \frac{1}{2}\begin{bmatrix}1+a_z & a_x - \imath a_y\\ a_x + \imath a_y & 1-a_z\end{bmatrix}\end{aligned}$$

where $\boldsymbol{a} = [a_x, a_y, a_z] \in \mathbb{R}^3$ is called the Bloch vector.

Digression 6.1. This is because $\{\underline{\boldsymbol{I}}, \underline{\boldsymbol{\sigma}}_x, \underline{\boldsymbol{\sigma}}_y, \underline{\boldsymbol{\sigma}}_z\}$ is a basis of the vector space of 2×2 Hermitian matrices. This space has 4 real dimensions, since a general Hermitian 2×2 matrix can be written as $\rho = c\underline{\boldsymbol{I}} + a_x\underline{\boldsymbol{\sigma}}_x + a_y\underline{\boldsymbol{\sigma}}_y + a_z\underline{\boldsymbol{\sigma}}_z$.

For density matrices, however, the normalization condition $\operatorname{tr}(\rho) = 1$ removes one degree of freedom. Since the Pauli matrices are traceless, $\operatorname{tr}(\rho) = \operatorname{tr}(c\underline{\boldsymbol{I}}) = 2c$, so $c = \frac{1}{2}$.

Thus the identity coefficient is not free, but fixed by normalization. The remaining freedom is exactly the 3 real coordinates (a_x, a_y, a_z).

Equivalently, the trace-one Hermitian matrices form a 3-dimensional affine slice of the 4-dimensional real vector space of Hermitian matrices. The Bloch ball lives in that affine slice. ◀

It is this Bloch vector $\boldsymbol{a}$ that indicates the point within the sphere that corresponds to a given mixed state. It can be shown that the eigenvalues of ρ are $\lambda_\pm = \frac{1}{2}(1 \pm \|\boldsymbol{a}\|)$. Density operators must be positive-semidefinite, so it follows that $\|\boldsymbol{a}\| \le 1$.

We have

$$\operatorname{tr}(\rho^2) \overset{(5.10)}{=} \operatorname{tr}(\underline{\boldsymbol{U}}\underline{\boldsymbol{\Lambda}}\underline{\boldsymbol{U}}^*\underline{\boldsymbol{U}}\underline{\boldsymbol{\Lambda}}\underline{\boldsymbol{U}}^*) \overset{(5.13)}{=} \operatorname{tr}(\underline{\boldsymbol{\Lambda}}^2) = \lambda_+^2 + \lambda_-^2 = \frac{1}{2}(1 + \|\boldsymbol{a}\|^2)$$

which, by Overview 5.5, is equal to 1 if and only if ρ is pure. Hence $\|\boldsymbol{a}\| = 1$ if and only if ρ is pure.

For a pure state $|\psi\rangle = [\cos\frac{\Theta}{2}, e^{\imath\varphi}\sin\frac{\Theta}{2}]^\top$ the corresponding density matrix is

$$\rho = |\psi\rangle\langle\psi| = \frac{1}{2}\begin{bmatrix} 1+\cos\Theta & \sin\Theta\cos\varphi - \imath\sin\Theta\sin\varphi \\ \sin\Theta\cos\varphi + \imath\sin\Theta\sin\varphi & 1-\cos\Theta \end{bmatrix}$$

Comparing this with the expansion of ρ in the Pauli basis above, we obtain

$$\boldsymbol{a} = \begin{bmatrix} \sin\Theta\cos\varphi \\ \sin\Theta\sin\varphi \\ \cos\Theta \end{bmatrix}$$

which is exactly the point on the unit sphere introduced above. Together with with the previous calculation this shows that the surface of the Bloch sphere represents precisely the pure states, whereas its interior corresponds to mixed states.

Operations on one Qbit are unitary 2×2-matrices.

Example 6.2. A central example is the Hadamard gate H. Its coordinate matrix in the computational basis is $\underline{\boldsymbol{H}} = \frac{1}{\sqrt{2}}\left[\begin{smallmatrix} 1 & 1 \\ 1 & -1 \end{smallmatrix}\right]$, introduced in Section 1.4. It satisfies $H^2 = \mathrm{id}$, and, hence, is not only unitary but also an involution (and, therefore, by Theorem 4.15, also self-adjoint). For $b \in \{0,1\}$,

$$H|b\rangle = \frac{|0\rangle + (-1)^b|1\rangle}{\sqrt{2}}$$

i.e. the Hadamard maps $|0\rangle \mapsto |+\rangle$, $|1\rangle \mapsto |-\rangle$, so its action on the Bloch sphere includes a rotation of $\pi/2$ around the axis connecting the circular polarizations $|\imath\rangle, |-\imath\rangle$. However, it also maps $|\imath\rangle \mapsto |-\imath\rangle$ and $|-\imath\rangle \mapsto |\imath\rangle$, i.e. it exchanges the two circular basis states. So the first rotation must be combined with a second rotation of π around the axis connecting the diagonal states $|+\rangle, |-\rangle$. ◀

Since arbitrary measurements in any orthonormal basis can be reduced to a unitary basis change followed by a standard-basis measurement, as discussed in Section 5.3, it is enough to understand measurements in the basis $\{|0\rangle, |1\rangle\}$. For instance, a measurement in the diagonal basis is a Hadamard transform followed by a standard measurement.

6.2 Two Qbits: $\mathcal{H} = \mathbb{C}^4$

Definition 6.1 (Product State)**.** A product state has the form $|\psi_1\rangle \otimes |\psi_2\rangle$. ◀

Not every state of two Qbits is of this form.

Example 6.3. The Bell states, which are the following maximally entangled two-Qbit states,

$$|\Phi^+\rangle = \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle)$$
$$|\Phi^-\rangle = \frac{1}{\sqrt{2}}(|00\rangle - |11\rangle)$$
$$|\Psi^+\rangle = \frac{1}{\sqrt{2}}(|01\rangle + |10\rangle)$$
$$|\Psi^-\rangle = \frac{1}{\sqrt{2}}(|01\rangle - |10\rangle)$$

cannot be written as a product state, see Section 1.5. They are also called EPR (Einstein-Podolsky-Rosen) pairs. They form an orthonormal basis of the two-Qbit Hilbert space, hence together they are often referred to as the Bell basis. We will see how to translate between the computational basis and the Bell basis in Section 6.3. ◀

Remark 6.4. A general two-Qbit pure state

$$|\psi\rangle = \alpha|00\rangle + \beta|01\rangle + \gamma|10\rangle + \delta|11\rangle$$

is a product state if and only if there exist $a, b, c, d \in \mathbb{C}$ such that

$$|\psi\rangle = (a|0\rangle + b|1\rangle) \otimes (c|0\rangle + d|1\rangle)$$

and expanding this gives $\alpha = ac$, $\beta = ad$, $\gamma = bc$, $\delta = bd$. Thus the coefficient matrix

$$\underline{\boldsymbol{A}} \coloneqq \begin{bmatrix} \alpha & \beta \\ \gamma & \delta \end{bmatrix} = \begin{bmatrix} a \\ b \end{bmatrix}\begin{bmatrix} c & d \end{bmatrix}$$

must have rank 1. Conversely, every rank-1 matrix admits such a factorization, hence determines a product state. ◀

The state space of a Qbit pair is the tensor product of the individual spaces, which is the span of the set of product states

$$\operatorname{span}\{|\psi_1\psi_2\rangle\} = \operatorname{span}\{|00\rangle, |01\rangle, |10\rangle, |11\rangle\} = \mathbb{C}^2 \otimes \mathbb{C}^2 \cong \mathbb{C}^4$$

Analogously, for n Qbits, the state space is

$$\operatorname{span}(\{|\boldsymbol{x}\rangle : \boldsymbol{x} \in \{0,1\}^n\}) = \left(\mathbb{C}^2\right)^{\otimes n} \cong \mathbb{C}^{2^n}$$

with computational basis indexed by classical n-bit strings.

On the level of coordinates, the tensor product is given by the Kronecker product:

$$\begin{bmatrix} a \\ b \end{bmatrix} \otimes \begin{bmatrix} c \\ d \end{bmatrix} = \begin{bmatrix} ac \\ ad \\ bc \\ bd \end{bmatrix}$$

Inner Product

The inner product on the tensor-product space is characterized by the requirement that tensor products of orthonormal basis vectors again form an orthonormal basis, i.e., if $\{|\varphi_{A,i}\rangle\}_i$ and $\{|\varphi_{B,i}\rangle\}_i$ are orthonormal bases of $\mathcal{H}_A$ and $\mathcal{H}_B$, then $\{|\varphi_{A,i_A}\rangle \otimes |\varphi_{B,i_B}\rangle\}_{i_A,i_B}$ should be an orthonormal basis of $\mathcal{H}_A \otimes \mathcal{H}_B$. Equivalently,

$$\langle \varphi_{A,i_{A,1}} \otimes \varphi_{B,i_{B,1}} | \varphi_{A,i_{A,2}} \otimes \varphi_{B,i_{B,2}} \rangle \overset{!}{=} \delta_{i_{A,1},i_{A,2}} \delta_{i_{B,1},i_{B,2}}$$

for all basis vectors. Since every vector in $\mathcal{H}_A \otimes \mathcal{H}_B$ is a linear combination of such product basis vectors, this uniquely determines the inner product on the whole tensor-product space.

The inner product therefore behaves multiplicatively on tensor products:

$$\langle \psi_{A,1} \otimes \psi_{B,1} | \psi_{A,2} \otimes \psi_{B,2} \rangle = \langle \psi_{A,1} | \psi_{A,2} \rangle \langle \psi_{B,1} | \psi_{B,2} \rangle$$

as mentioned in Digression 5.7.

In particular, two tensor products are orthogonal if and only if in at least one factor the corresponding states are orthogonal.

This is consistent with the coordinate level, i.e. let

$$|\psi_{A,1}\rangle = \begin{bmatrix} a_1 \\ b_1 \end{bmatrix} \qquad |\psi_{B,1}\rangle = \begin{bmatrix} c_1 \\ d_1 \end{bmatrix} \qquad |\psi_{A,2}\rangle = \begin{bmatrix} a_2 \\ b_2 \end{bmatrix} \qquad |\psi_{B,2}\rangle = \begin{bmatrix} c_2 \\ d_2 \end{bmatrix}$$

Then

$$\begin{aligned}
\langle \psi_{A,1} \otimes \psi_{B,1} | \psi_{A,2} \otimes \psi_{B,2} \rangle &= (|\psi_{A,1}\rangle \otimes |\psi_{B,1}\rangle)^* (|\psi_{A,2}\rangle \otimes |\psi_{B,2}\rangle) \\
&= \begin{bmatrix} a_1^* c_1^* & a_1^* d_1^* & b_1^* c_1^* & b_1^* d_1^* \end{bmatrix} \begin{bmatrix} a_2 c_2 \\ a_2 d_2 \\ b_2 c_2 \\ b_2 d_2 \end{bmatrix} \\
&= a_1^* a_2 c_1^* c_2 + a_1^* a_2 d_1^* d_2 + b_1^* b_2 c_1^* c_2 + b_1^* b_2 d_1^* d_2 \\
&= (a_1^* a_2 + b_1^* b_2)(c_1^* c_2 + d_1^* d_2) \\
&= \langle \psi_{A,1} | \psi_{A,2} \rangle \langle \psi_{B,1} | \psi_{B,2} \rangle
\end{aligned}$$

Operations

An example of an operation that can be carried out on a pair of Qbits is a product operation. Let us define it on the set of product states, which is a generating system of the full space. The operation

$$(U \otimes V)(|\psi_A\rangle \otimes |\psi_B\rangle) \coloneqq U |\psi_A\rangle \otimes V |\psi_B\rangle \tag{6.1}$$

is thus uniquely defined by linearity, as already mentioned in Digression 5.7.

This map is unitary as long as both U and V are: It is easy to see that orthogonal tensor products are mapped to orthogonal tensor product along our understanding of orthogonality above.

Example 6.5. The matrix Hadamard on two Qbits can again be computed through the Kronecker product:

$$\underline{\boldsymbol{H}} \otimes \underline{\boldsymbol{H}} = \frac{1}{2}\left[\begin{array}{c:c} 1 \cdot \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix} & 1 \cdot \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix} \\ \hdashline 1 \cdot \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix} & -1 \cdot \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix} \end{array}\right] = \left[\begin{array}{cc:cc} 1 & 1 & 1 & 1 \\ 1 & -1 & 1 & -1 \\ \hdashline 1 & 1 & -1 & -1 \\ 1 & -1 & -1 & 1 \end{array}\right]$$

This representation is not very intuitive or transparent. A better understanding can be obtained from the following formula (see also Example 6.2): For $b_1, b_2 \in \{0, 1\}$,

$$\begin{aligned} H^{\otimes 2} |b_1 b_2\rangle &= H |b_1\rangle \otimes H |b_2\rangle \\ &= \frac{1}{2}\left(\left(|0\rangle + (-1)^{b_1} |1\rangle\right) \otimes \left(|0\rangle + (-1)^{b_2} |1\rangle\right)\right) \\ &= \frac{1}{2}\left(|00\rangle + (-1)^{b_2} |01\rangle + (-1)^{b_1} |10\rangle + (-1)^{b_1 \oplus b_2} |11\rangle\right) \end{aligned}$$

where $\oplus$ denotes addition modulo 2 (i.e., XOR). ◀

Measurements

Measurements can also be performed on only one part of a bipartite system, in accordance with the general formalism from Section 5.3. If

$$|\Psi\rangle = \alpha |00\rangle + \beta |01\rangle + \gamma |10\rangle + \delta |11\rangle$$

and one measures only the first Qbit in the standard basis, then the probabilities are $p_0 = |\alpha|^2 + |\beta|^2$, $p_1 = |\gamma|^2 + |\delta|^2$. Rewrite the state as

$$\begin{aligned} |\Psi\rangle &= |0\rangle \otimes (\alpha |0\rangle + \beta |1\rangle) + |1\rangle \otimes (\gamma |0\rangle + \delta |1\rangle) \\ &= \sqrt{p_0}\, |0\rangle \otimes \frac{\alpha |0\rangle + \beta |1\rangle}{\sqrt{p_0}} + \sqrt{p_1}\, |1\rangle \otimes \frac{\gamma |0\rangle + \delta |1\rangle}{\sqrt{p_1}} \\ &= \sqrt{p_0}\ |0\rangle \otimes |\psi_0\rangle + \sqrt{p_1}\ |1\rangle \otimes |\psi_1\rangle \end{aligned}$$

provided $p_0, p_1 \neq 0$. Thus, conditioned on the outcome 0 or 1, the second Qbit is left in the corresponding pure state $|\psi_0\rangle$ or $|\psi_1\rangle$.

If we ignore the measurement outcome on the first Qbit, or we don't perform any measurement at all on the first Qbit, then the second Qbit is described by the mixture

$$p_0 P_{|\psi_0\rangle} + p_1 P_{|\psi_1\rangle} = p_0 |\psi_0\rangle \langle\psi_0| + p_1 |\psi_1\rangle \langle\psi_1|$$

which is exactly the reduced state obtained by the partial trace, i.e. when the first Qbit is "traced out" (Section 5.7, Definition 5.7 and Example 5.8). Once again we see: parts of entangled pure states are mixed.

6.3 The CNOT gate

Similarly to states, also operations on multiple Qbits are not necessarily just products of operations on the individual Qbits.

An important two-Qbit gate is the controlled NOT gate, abbreviated CNOT. It can be seen as a "made-reversible" XOR gate.

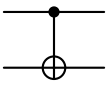

As mentioned above, there exist unitaries on multiple Qbits that cannot be written as a product of unitaries on the individual Qbits. Remarkably, it is enough to have the CNOT plus unary (one-bit unitary) operations, and every unitary on an arbitrary number of Qbits becomes possible.

Its action on the computational basis is

$$|x\rangle \otimes |b\rangle \longmapsto |x\rangle \otimes |b \oplus x\rangle$$

where $x, b \in \{0, 1\}$ and $\oplus$ denotes addition modulo 2 (i.e., XOR). That is, the first Qbit is the *control* and the second is the *target*: the second bit is flipped exactly if the first bit equals 1.

In the basis $\{|00\rangle, |01\rangle, |10\rangle, |11\rangle\}$, the coordinate matrix of the CNOT gate is

$$\underline{\boldsymbol{U}}_{\mathrm{CNOT}} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{bmatrix}$$

Since the gate is linear, its action on arbitrary superpositions is determined automatically.

A first surprise occurs if one works in the diagonal basis. To understand the action of the CNOT on the diagonal basis, we first express the diagonal basis states in terms of the computational basis:

$$\begin{aligned} |{+}{+}\rangle &= \frac{|0\rangle + |1\rangle}{\sqrt{2}} \otimes \frac{|0\rangle + |1\rangle}{\sqrt{2}} = \frac{1}{2}(|00\rangle + |01\rangle + |10\rangle + |11\rangle) \\ |{+}{-}\rangle &= \frac{|0\rangle + |1\rangle}{\sqrt{2}} \otimes \frac{|0\rangle - |1\rangle}{\sqrt{2}} = \frac{1}{2}(|00\rangle - |01\rangle + |10\rangle - |11\rangle) \\ |{-}{+}\rangle &= \frac{|0\rangle - |1\rangle}{\sqrt{2}} \otimes \frac{|0\rangle + |1\rangle}{\sqrt{2}} = \frac{1}{2}(|00\rangle + |01\rangle - |10\rangle - |11\rangle) \\ |{-}{-}\rangle &= \frac{|0\rangle - |1\rangle}{\sqrt{2}} \otimes \frac{|0\rangle - |1\rangle}{\sqrt{2}} = \frac{1}{2}(|00\rangle - |01\rangle - |10\rangle + |11\rangle) \end{aligned}$$

A straightforward calculation reveals that, on a product of diagonal states (one of the 4 possible from above), the CNOT again yields a product of diagonal states, but the roles of control and target are exchanged; the second Qbit acts as the control, and whenever it is in the state $|-\rangle$, the first Qbit is flipped between $|+\rangle$ and $|-\rangle$.

Equivalently,

$$(H \otimes H)\,\mathrm{CNOT}\,(H \otimes H)$$

is a CNOT with swapped control and target:

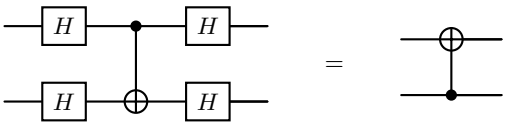

More interestingly even, the CNOT can generate entanglement, i.e., map product states to non-product states:

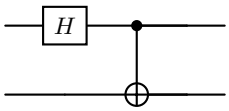

Indeed,

$$\begin{aligned} |00\rangle &\overset{H\otimes\mathrm{id}}{\longmapsto} \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)\,|0\rangle \overset{\mathrm{CNOT}}{\longmapsto} \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle) = |\Phi^+\rangle \\ |01\rangle &\overset{H\otimes\mathrm{id}}{\longmapsto} \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)\,|1\rangle \overset{\mathrm{CNOT}}{\longmapsto} \frac{1}{\sqrt{2}}(|01\rangle + |10\rangle) = |\Psi^+\rangle \\ |10\rangle &\overset{H\otimes\mathrm{id}}{\longmapsto} \frac{1}{\sqrt{2}}(|0\rangle - |1\rangle)\,|0\rangle \overset{\mathrm{CNOT}}{\longmapsto} \frac{1}{\sqrt{2}}(|00\rangle - |11\rangle) = |\Phi^-\rangle \\ |11\rangle &\overset{H\otimes\mathrm{id}}{\longmapsto} \frac{1}{\sqrt{2}}(|0\rangle - |1\rangle)\,|1\rangle \overset{\mathrm{CNOT}}{\longmapsto} \frac{1}{\sqrt{2}}(|01\rangle - |10\rangle) = |\Psi^-\rangle \end{aligned}$$

Thus a Hadamard on the control followed by a CNOT maps computational basis states to Bell states from Section 1.5 and Example 6.3. So this combination implements the basis change between the computational basis and the Bell basis.

Consequently, a Bell-basis measurement can be implemented by reversing this basis change (recall that $H^{-1} = H$) and then measuring in the computational basis:

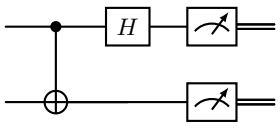

Classical universality

The CNOT is universal together with one-Qbit gates. For classical universality, it is enough to obtain a Toffoli gate (Definition 2.5), also called CCNOT (controlled-controlled NOT), i.e., the gate that maps

$$|x\rangle \otimes |y\rangle \otimes |b\rangle \longmapsto |x\rangle \otimes |y\rangle \otimes |b \oplus (x \wedge y)\rangle$$

since this gives reversible implementations of arbitrary classical computations. Let R be the rotation by 22.5°. Its coordinate matrix is

$$\underline{\boldsymbol{R}} = \begin{bmatrix} \cos 22.5^\circ & -\sin 22.5^\circ \\ \sin 22.5^\circ & \cos 22.5^\circ \end{bmatrix}$$

Let us first consider a CNOT with a classical control $|x\rangle \in \{|0\rangle, |1\rangle\}$ and an arbitrary second Qbit $|\varphi\rangle = \alpha|0\rangle + \beta|1\rangle$. If $x = 0$, the operation is the identity; if $x = 1$, the second Qbit is negated:

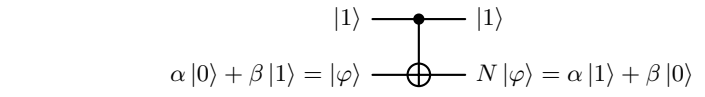

where the coordinate matrix of N is $\underline{\boldsymbol{N}} = \left[\begin{smallmatrix} 0 & 1 \\ 1 & 0 \end{smallmatrix}\right]$. Geometrically, the case $x = 1$ corresponds to a reflection of the second Qbit around the positive diagonal.

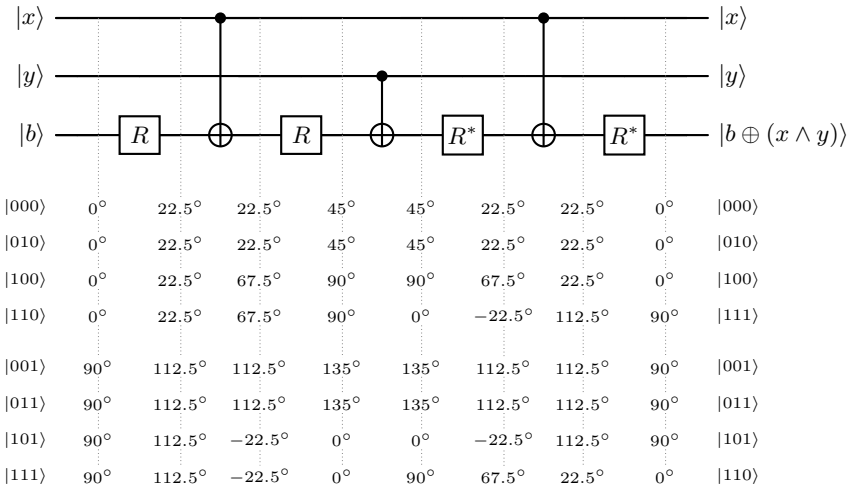

$\|000\rangle$	0°	22.5°	22.5°	45°	45°	22.5°	22.5°	0°	$\|000\rangle$
$\|010\rangle$	0°	22.5°	22.5°	45°	45°	22.5°	22.5°	0°	$\|010\rangle$
$\|100\rangle$	0°	22.5°	67.5°	90°	90°	67.5°	22.5°	0°	$\|100\rangle$
$\|110\rangle$	0°	22.5°	67.5°	90°	0°	−22.5°	112.5°	90°	$\|111\rangle$
$\|001\rangle$	90°	112.5°	112.5°	135°	135°	112.5°	112.5°	90°	$\|001\rangle$
$\|011\rangle$	90°	112.5°	112.5°	135°	135°	112.5°	112.5°	90°	$\|011\rangle$
$\|101\rangle$	90°	112.5°	−22.5°	0°	0°	−22.5°	112.5°	90°	$\|101\rangle$
$\|111\rangle$	90°	112.5°	−22.5°	0°	90°	67.5°	22.5°	0°	$\|110\rangle$

In analogy to the CNOT, the Toffoli/CCNOT gate is drawn as

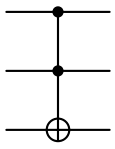

and its coordinate matrix in the computational basis is

$$\underline{\boldsymbol{U}}_{\text{CCNOT}} = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \end{bmatrix}$$

We can also obtain a SWAP gate from CNOTs, which swaps two Qbits:

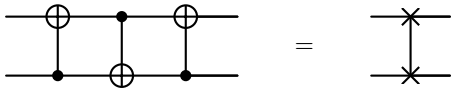

6.4 Cloning, pseudo-cloning, and pseudo-measurements

Classical bits can be copied with a CNOT by initializing the target (source) to 0. For quantum states this cannot be extended to arbitrary inputs.

Theorem 6.1 (No-cloning)**.** There is no unitary C on $\mathcal{H} \otimes \mathcal{H}$ such that for all normalized $|\psi\rangle \in \mathcal{H}$ and some initial normalized $|e\rangle \in \mathcal{H}$, $C(|\psi\rangle \otimes |e\rangle) = c_\psi |\psi\rangle \otimes |\psi\rangle$ for some $c_\psi \in \mathbb{C}$ depending on $|\psi\rangle$ with $|c_\psi| = 1$. ◁

Proof. Consider two arbitrary normalized states $|\psi_1\rangle$ and $|\psi_2\rangle$ in $\mathcal{H}$. If there were a unitary operation C that cloned both $|\psi_1\rangle$ and $|\psi_2\rangle$ in the same way as the CNOT clones classical bits, i.e.,

$$C(|\psi_i\rangle \otimes |e\rangle) = c_i |\psi_i\rangle \otimes |\psi_i\rangle , \qquad |c_i| = 1 \tag{6.2}$$

for $i = 1, 2$, then preservation of inner products by unitaries (Theorem 3.2, Equation 3.10) would give

$$\begin{aligned} \langle\psi_1|\psi_2\rangle \langle e|e\rangle &\overset{(5.20)}{=} \langle\psi_1 e|\psi_2 e\rangle \\ &\overset{(3.10)}{=} \langle C|\psi_1 e\rangle |C|\psi_2 e\rangle\rangle \\ &\overset{(6.2)}{=} \overline{c_1} c_2 \langle\psi_1\psi_1|\psi_2\psi_2\rangle \\ &\overset{(5.20)}{=} \overline{c_1} c_2 \langle\psi_1|\psi_2\rangle \langle\psi_1|\psi_2\rangle \end{aligned}$$

Taking absolute values gives

$$|\langle\psi_1|\psi_2\rangle| = |\langle\psi_1|\psi_2\rangle|^2$$

so $|\langle\psi_1|\psi_2\rangle| \in \{0, 1\}$. Thus only orthogonal states or parallel states, i.e. states representing the same physical state up to global phase, can be cloned by a single unitary operation. Therefore no unitary cloning operation can work for arbitrary inputs. □

Instead, for a general input $|\varphi\rangle = \alpha|0\rangle + \beta|1\rangle$ on the control wire and target $|0\rangle$, CNOT gives, by linearity,

$$\alpha|00\rangle + \beta|11\rangle$$

which is generally entangled, not the product of two copies, which would be

$$\alpha^2|00\rangle + \alpha\beta|01\rangle + \alpha\beta|10\rangle + \beta^2|11\rangle$$

Therefore, if the second Qbit is then measured in the standard basis, the observed statistics are actually the same as if the input had been cloned:

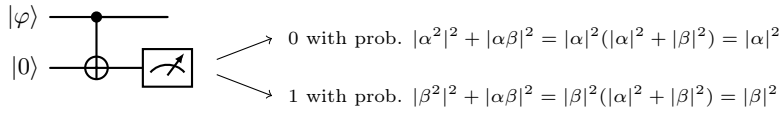

The same standard-basis statistics are obtained on the first Qbit, whose reduced state is

$$|\alpha|^2 P_{|0\rangle} + |\beta|^2 P_{|1\rangle}$$

The two measurement outcomes are perfectly correlated, because the joint state is $\alpha|00\rangle + \beta|11\rangle$. This is why the action is sometimes called *pseudo-cloning*.

If the measurement on the second wire is not done, this should not change the reduced state of the first wire. CNOT leads to the same transition of the first Qbit as if a standard-basis measurement on the first Qbit had happened and the outcome had been forgotten. The operation is not an actual measurement, since it has no outcome and is reversible: CNOT is self-inverse,

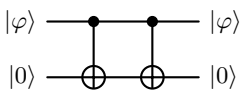

which motivates the term *pseudo-measurement*.

More generally, a measurement can be seen as replacing the second Qbit with a large system containing the apparatus, observer, laboratory, and environment, and replacing the CNOT by a controlled unitary:

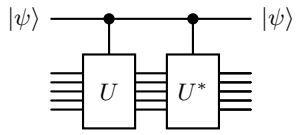

In this interpretation, the measurement would always be reversible and hence not induce a "collapse", but merely *decoherence*.

The same mechanism also illustrates *disturbance*: if a quantum computation couples to even a single uncontrolled binary degree of freedom that should not be part of the computation, then, once that degree of freedom escapes, the reduced state of the computer becomes a mixture. The process is then effectively irreversible, and the computation may fail.

Fact 6.2 (Deferred measurement principle)**.** Measurements can always be moved from an intermediate stage of a quantum circuit to the end of the circuit; if the measurement results are used at any stage of the circuit then the classically controlled operations can be replaced by conditional quantum operations [18]. ◁

The idea is to keep the measurement outcome coherent. Instead of measuring a Qbit immediately, copy its computational-basis value into an ancilla by a CNOT and keep that ancilla as a control wire. Every later gate that would have been classically conditioned on the measurement result is replaced by the corresponding quantum-controlled gate. At the end, the ancilla and the output wires are measured. The joint probability distribution of the final classical outcomes is the same as in the original circuit.

Thus, for complexity-theoretic purposes, a quantum circuit can usually be treated as a unitary computation followed by one final measurement. We will use this in Section 12.

6.5 n Qbits: $\mathcal{H} = \mathbb{C}^{2^n}$

The state space of n Qbits is

$$\mathcal{H}_n = \left(\mathbb{C}^2\right)^{\otimes n} \cong \mathbb{C}^{2^n}$$

with computational basis $\{|\boldsymbol{x}\rangle : \boldsymbol{x} \in \{0,1\}^n\}$. Thus the finite-basis identity insertion from (5.11) becomes

$$\mathrm{id}_{\mathcal{H}_n} = \sum_{\boldsymbol{x}\in\{0,1\}^n} |\boldsymbol{x}\rangle\langle\boldsymbol{x}| \tag{6.3}$$

Example 6.6. Generalizing Example 6.5 to n Qbits

$$|\boldsymbol{b}\rangle = \bigotimes_{j=1}^{n} |b_j\rangle$$

we have

$$H^{\otimes n}|\boldsymbol{b}\rangle = \frac{1}{2^{n/2}} \sum_{\boldsymbol{x}\in\{0,1\}^n} (-1)^{\boldsymbol{b}\cdot\boldsymbol{x}} |\boldsymbol{x}\rangle \tag{6.4}$$

where

$$\boldsymbol{b}\cdot\boldsymbol{x} \coloneqq \bigoplus_{j=1}^{n} (b_j \wedge x_j) \tag{6.5}$$

denotes the XOR of the bitswise AND of $\boldsymbol{b}$ and $\boldsymbol{x}$.

In particular,

$$H^{\otimes n}|\boldsymbol{0}\rangle = \frac{1}{2^{n/2}} \sum_{\boldsymbol{x}\in\{0,1\}^n} |\boldsymbol{x}\rangle \tag{6.6}$$

will be the initial state for most quantum algorithms, and as a equal superposition of all classical inputs, the basis for "Quantum parallelism". ◀

The logical scalar product is linear, i.e.,

$$(\boldsymbol{x} \oplus \boldsymbol{s}) \cdot \boldsymbol{z} = (\boldsymbol{x} \cdot \boldsymbol{z}) \oplus (\boldsymbol{s} \cdot \boldsymbol{z}) \tag{6.7}$$

for all $\boldsymbol{x}, \boldsymbol{s}, \boldsymbol{z} \in \{0,1\}^n$, because

$$\begin{aligned}
(\boldsymbol{x} \oplus \boldsymbol{s}) \cdot \boldsymbol{z} &= \bigoplus_{j=1}^{n} \big((x_j \oplus s_j) \wedge z_j\big) \\
&= \bigoplus_{j=1}^{n} \big((x_j \wedge z_j) \oplus (s_j \wedge z_j)\big) \\
&= \left(\bigoplus_{j=1}^{n} x_j \wedge z_j\right) \oplus \left(\bigoplus_{j=1}^{n} s_j \wedge z_j\right) \\
&= (\boldsymbol{x} \cdot \boldsymbol{z}) \oplus (\boldsymbol{s} \cdot \boldsymbol{z})
\end{aligned}$$

7 Quantum Communication

7.1 Teleportation

Quantum teleportation transfers an unknown quantum state from Alice to Bob without sending the physical system that originally carried it. The transfer is neither instantaneous nor a way of copying the state: Alice's original state is destroyed during the protocol, in agreement with the no-cloning theorem (Theorem 6.1), and Bob can recover it only after receiving a classical message from Alice.

The essential resource is an EPR pair, such as the Bell state $|\Phi^+\rangle = \frac{|00\rangle+|11\rangle}{\sqrt{2}}$ from Example 6.3, which Alice and Bob distribute before the state to be transmitted is available:

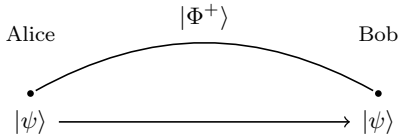

Alice keeps one half of the pair and Bob keeps the other. In this sense, teleportation moves the need for a quantum channel to an earlier time: entanglement is distributed while such a channel is available, and the unknown state can later be transmitted using only a classical channel. When Alice later receives the unknown state $|\psi\rangle$, she performs a Bell-basis measurement on it together with her half of the EPR pair. As explained in Section 6.3, this measurement can be implemented by a CNOT, a Hadamard gate, and two standard-basis measurements.

Let $|\psi\rangle = \alpha\,|0\rangle + \beta\,|1\rangle$ be the normalized unknown state to be teleported. Writing X for negation and Z for the phase flip, the complete circuit is

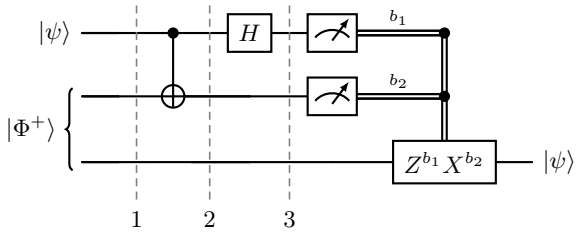

At position 1, the unknown Qbit and the EPR pair form the product state

$$|\psi\rangle \otimes |\Phi^+\rangle = (\alpha\,|0\rangle + \beta\,|1\rangle) \otimes \tfrac{|00\rangle+|11\rangle}{\sqrt{2}} = \tfrac{1}{\sqrt{2}}\,(\alpha\,|000\rangle + \alpha\,|011\rangle + \beta\,|100\rangle + \beta\,|111\rangle)$$

At position 2, after the CNOT from the unknown Qbit to Alice's half of the EPR pair, the state is

$$\tfrac{1}{\sqrt{2}}\,(\alpha\,|000\rangle + \alpha\,|011\rangle + \beta\,|110\rangle + \beta\,|101\rangle) = \alpha\,|0\rangle \otimes \tfrac{|00\rangle+|11\rangle}{\sqrt{2}} + \beta\,|1\rangle \otimes \tfrac{|10\rangle+|01\rangle}{\sqrt{2}}$$

This is generally a genuinely entangled three-Qbit state.

At position 3, after applying the Hadamard to the first Qbit, the joint state can be expanded and then regrouped by Alice's two computational-basis states:

$$\begin{aligned}
&\alpha H\,|0\rangle \otimes \tfrac{|00\rangle+|11\rangle}{\sqrt{2}} + \beta H\,|1\rangle \otimes \tfrac{|10\rangle+|01\rangle}{\sqrt{2}} \\
&= \alpha \tfrac{|0\rangle+|1\rangle}{\sqrt{2}} \otimes \tfrac{|00\rangle+|11\rangle}{\sqrt{2}} + \beta \tfrac{|0\rangle-|1\rangle}{\sqrt{2}} \otimes \tfrac{|10\rangle+|01\rangle}{\sqrt{2}} \\
&= \tfrac{\alpha}{2}\,(|000\rangle + |011\rangle + |100\rangle + |111\rangle) + \tfrac{\beta}{2}\,(|010\rangle + |001\rangle - |110\rangle - |101\rangle) \\
&= \tfrac{|00\rangle\otimes(\alpha|0\rangle+\beta|1\rangle)+|01\rangle\otimes(\alpha|1\rangle+\beta|0\rangle)+|10\rangle\otimes(\alpha|0\rangle-\beta|1\rangle)+|11\rangle\otimes(\alpha|1\rangle-\beta|0\rangle)}{2} \\
&= \tfrac{1}{2}\Big(|00\rangle \otimes |\psi\rangle + |01\rangle \otimes X\,|\psi\rangle + |10\rangle \otimes Z\,|\psi\rangle + |11\rangle \otimes XZ\,|\psi\rangle\Big)
\end{aligned}$$

where the coordinate matrices of X and Z are $\underline{\underline{\boldsymbol{X}}} = \left[\begin{smallmatrix} 0 & 1 \\ 1 & 0 \end{smallmatrix}\right]$ and $\underline{\underline{\boldsymbol{Z}}} = \left[\begin{smallmatrix} 1 & 0 \\ 0 & -1 \end{smallmatrix}\right]$.

Alice's measurement produces two classical bits b_1, b_2, which she sends to Bob. Their values tell Bob which of four possible Pauli corrections to apply to his Qbit; after the correction, his Qbit is in the original state $|\psi\rangle$ in all four cases. The protocol therefore moves the state rather than the underlying particle.

The classical message is indispensable. Before it arrives, Bob does not know which correction is required, and his Qbit is described by the maximally mixed state

$$\begin{aligned}
\rho &= \frac{1}{4}\left(P_{|\psi\rangle} + P_{X|\psi\rangle} + P_{Z|\psi\rangle} + P_{XZ|\psi\rangle}\right) \\
&= \frac{1}{4}\left(P_{\alpha|0\rangle+\beta|1\rangle} + P_{\alpha|1\rangle+\beta|0\rangle} + P_{\alpha|0\rangle-\beta|1\rangle} + P_{\alpha|1\rangle-\beta|0\rangle}\right) \\
&= \frac{1}{4}\left(\begin{bmatrix} |\alpha|^2 & \alpha\overline{\beta} \\ \overline{\alpha}\beta & |\beta|^2 \end{bmatrix} + \begin{bmatrix} |\beta|^2 & \beta\overline{\alpha} \\ \overline{\beta}\alpha & |\alpha|^2 \end{bmatrix} + \begin{bmatrix} |\alpha|^2 & -\alpha\overline{\beta} \\ -\overline{\alpha}\beta & |\beta|^2 \end{bmatrix} + \begin{bmatrix} |\beta|^2 & -\beta\overline{\alpha} \\ -\overline{\beta}\alpha & |\alpha|^2 \end{bmatrix}\right) \\
&= \frac{1}{4}\begin{bmatrix} 2|\alpha|^2 + 2|\beta|^2 & 0 \\ 0 & 2|\alpha|^2 + 2|\beta|^2 \end{bmatrix} \\
&= \begin{bmatrix} \frac{1}{2} & 0 \\ 0 & \frac{1}{2} \end{bmatrix}
\end{aligned}$$

where we used Definition 5.4 (5.15). Since this (maximally) mixed state is independent of $|\psi\rangle$, teleportation cannot be used for faster-than-light communication, saving us from trouble with relativity theory.

An application is a *quantum repeater*. Directly distributing entanglement over a large distance is difficult because interaction with the environment causes decoherence. Instead, intermediate stations can share shorter EPR pairs with their neighbours. An intermediate station performs a Bell measurement on its two local Qbits, thereby *swapping* the entanglement so that the distant endpoints become entangled. Repeating this procedure extends entanglement across the network, although a repeater used in a cryptographic protocol may have to be trusted.

7.2 Superdense Coding

Superdense coding reverses the resource trade of teleportation. Without shared entanglement, transmitting one Qbit can convey at most one classical bit. If Alice and Bob already share an EPR pair, Alice can instead communicate two classical bits b_1, b_2 by sending only her half of the pair:

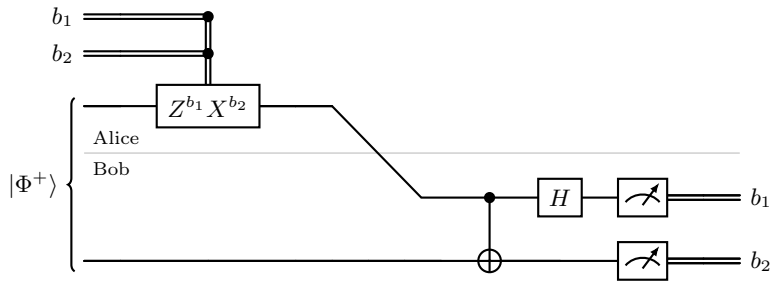

Alice encodes the pair b_1, b_2 by applying $Z^{b_1}X^{b_2}$ to her half of the EPR pair. The four possible encodings turn the shared pair into the four mutually orthogonal Bell states. Alice then sends her Qbit to Bob, who performs the Bell-basis measurement from Section 6.3 and thereby identifies both classical bits.

Although entanglement alone carries no message, it changes what the later Qbit transmission can accomplish. This activation of one resource by another is the central surprise of superdense coding. The two protocols use the same gates in opposite order: teleportation consumes one EPR pair and two transmitted classical bits to communicate one Qbit, whereas superdense coding consumes one EPR pair and one transmitted Qbit to communicate two classical bits.

8 Simple Algorithms

We recall from Section 6.5 that $H^{\otimes n}|\mathbf{0}\rangle$ is the equal superposition of all classical inputs.

8.0.1 Reversible oracles and quantum parallelism

A classical function $f : \{0,1\}^n \to \{0,1\}^m$ is embedded into the reversible oracle

$$U_f|\boldsymbol{x}\rangle|\boldsymbol{y}\rangle = |\boldsymbol{x}\rangle|\boldsymbol{y} \oplus f(\boldsymbol{x})\rangle \tag{8.1}$$

where the XOR is applied bitwise. Since this is a permutation of the computational basis, it extends linearly to a unitary operation. The corresponding circuit is

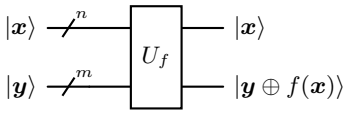

Applying U_f after preparing an equal superposition (see Section 6.5) yields

$$|\mathbf{0}_n\rangle|\mathbf{0}_m\rangle \xmapsto{(H^{\otimes n})\otimes(\mathrm{id}^{\otimes m})} \frac{1}{2^{n/2}}\sum_{\boldsymbol{x}}|\boldsymbol{x}\rangle|\mathbf{0}_m\rangle \xmapsto{U_f} \frac{1}{2^{n/2}}\sum_{\boldsymbol{x}}|\boldsymbol{x}\rangle|f(\boldsymbol{x})\rangle$$

which contains all function values at once. Nevertheless, a standard-basis measurement returns only one random pair $(\boldsymbol{x}, f(\boldsymbol{x}))$, so parallelism by itself gives no speedup.

8.0.2 Phase kickback

For a one-bit output, initialize the result Qbit in the diagonal state $|-\rangle = H|1\rangle = \frac{1}{\sqrt{2}}(|0\rangle - |1\rangle)$. Then

$$U_f|\boldsymbol{x}\rangle|-\rangle = |\boldsymbol{x}\rangle\frac{|f(\boldsymbol{x})\rangle - |1 \oplus f(\boldsymbol{x})\rangle}{\sqrt{2}} = (-1)^{f(\boldsymbol{x})}|\boldsymbol{x}\rangle|-\rangle \tag{8.2}$$

The output Qbit is unchanged, while the function value is "kicked back" as a phase on the input register:

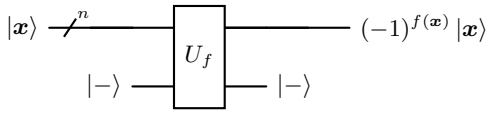

By linearity, the same transformation applies simultaneously to every term of a superposition:

$$|\mathbf{0}_n\rangle|0\rangle \xmapsto{(H^{\otimes n})\otimes(HX)} \frac{1}{2^{n/2}}\sum_{\boldsymbol{x}}|\boldsymbol{x}\rangle|-\rangle \xmapsto{U_f} \frac{1}{2^{n/2}}\sum_{\boldsymbol{x}}(-1)^{f(\boldsymbol{x})}|\boldsymbol{x}\rangle|-\rangle$$

More generally, phase kickback occurs whenever the target register is prepared in an eigenstate of the controlled operation. Let a controlled unitary U_C act as

$$U_C|0\rangle|u\rangle = |0\rangle|u\rangle, \qquad U_C|1\rangle|u\rangle = |1\rangle U|u\rangle$$

and assume that the target state $|u\rangle$ is an eigenvector of U with eigenvalue $e^{i\phi}$:

$$U|u\rangle = e^{i\phi}|u\rangle$$

Then, by linearity,

$$\begin{aligned} U_C((\alpha|0\rangle + \beta|1\rangle)\otimes|u\rangle) &= \alpha U_C|0\rangle|u\rangle + \beta U_C|1\rangle|u\rangle \\ &= \alpha|0\rangle|u\rangle + \beta|1\rangle U|u\rangle \\ &= \alpha|0\rangle|u\rangle + \beta e^{i\phi}|1\rangle|u\rangle \\ &= (\alpha|0\rangle + e^{i\phi}\beta|1\rangle)\otimes|u\rangle \end{aligned}$$

Thus the target state remains unchanged, while the phase $e^{i\phi}$ is kicked back onto the control branch. Equation (8.2) is the special case where $U = X$ and

$$X|-\rangle = -|-\rangle$$

so the eigenvalue is $e^{i\phi} = -1$.

The following algorithms arrange a final Hadamard transform so that the relative phases interfere into measurable information.

8.1 Deutsch's algorithm

Deutsch's algorithm is the $n = 1$ case of the Deutsch–Jozsa algorithm. Given

$$f : \{0,1\} \to \{0,1\}$$

promised to be either *constant* or *balanced*, decide which case holds. Equivalently, determine the bit

$$f(0) \oplus f(1) = \begin{cases} 0 & f \text{ constant} \\ 1 & f \text{ balanced} \end{cases}$$

A deterministic classical algorithm needs two oracle calls. Deutsch's quantum algorithm needs one:

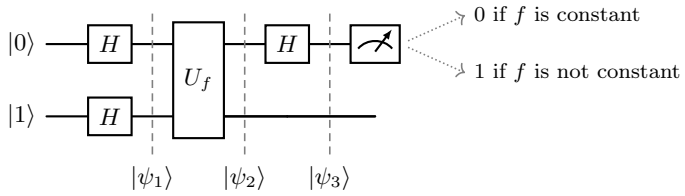

The state after the first Hadamards is

$$|\psi_1\rangle = \frac{|0\rangle + |1\rangle}{\sqrt{2}} \otimes \frac{|0\rangle - |1\rangle}{\sqrt{2}}$$

Phase kickback from (8.2) gives

$$|\psi_2\rangle = \frac{(-1)^{f(0)}\,|0\rangle + (-1)^{f(1)}\,|1\rangle}{\sqrt{2}} \otimes |-\rangle$$

Applying the final Hadamard to the first Qbit yields

$$\begin{aligned} |\psi_3\rangle &= \frac{1}{2}\left((-1)^{f(0)}(|0\rangle + |1\rangle) + (-1)^{f(1)}(|0\rangle - |1\rangle)\right) \otimes |-\rangle \\ &= \frac{1}{2}\left(\left((-1)^{f(0)} + (-1)^{f(1)}\right)|0\rangle + \left((-1)^{f(0)} - (-1)^{f(1)}\right)|1\rangle\right) \otimes |-\rangle \\ &= \begin{cases} \pm\,|0\rangle \otimes |-\rangle & \text{if } f \text{ is constant} \\ \pm\,|1\rangle \otimes |-\rangle & \text{if } f \text{ is balanced} \end{cases} \end{aligned}$$

and hence the measured bit is exactly $f(0) \oplus f(1)$.

8.2 Deutsch–Jozsa algorithm

Let

$$f : \{0,1\}^n \to \{0,1\}$$

with the promise that f is either constant or balanced, meaning

$$|\{\boldsymbol{x} : f(\boldsymbol{x}) = 0\}| = |\{\boldsymbol{x} : f(\boldsymbol{x}) = 1\}| = 2^{n-1}$$

A deterministic classical algorithm needs $2^{n-1} + 1$ oracle calls in the worst case: after 2^{n-1} identical values, the unseen inputs could still all have the opposite value. The Deutsch–Jozsa algorithm uses one oracle call:

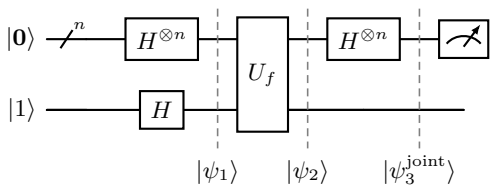

After the first Hadamards, the joint state is

$$|\psi_1\rangle = \left(H^{\otimes n}\,|\mathbf{0}\rangle\right) \otimes \left(H\,|1\rangle\right) \overset{(6.6)}{=} \frac{1}{2^{n/2}} \sum_{\boldsymbol{x} \in \{0,1\}^n} |\boldsymbol{x}\rangle \otimes |-\rangle$$

Phase kickback from (8.2) gives

$$|\psi_2\rangle = U_f\,|\psi_1\rangle = \frac{1}{2^{n/2}} \sum_{\boldsymbol{x}\in\{0,1\}^n} U_f(|\boldsymbol{x}\rangle \otimes |-\rangle) \overset{(8.2)}{=} \frac{1}{2^{n/2}} \sum_{\boldsymbol{x}\in\{0,1\}^n} (-1)^{f(\boldsymbol{x})}\,|\boldsymbol{x}\rangle \otimes |-\rangle$$

We denote the state after the final Hadamards $|\psi_3^{\text{joint}}\rangle = |\psi_3\rangle \otimes |-\rangle$ where $|\psi_3\rangle$ is the state of the first n Qbits. We don't need the last Qbit anymore, so we ignore it for the sake of readability. Using (6.4),

$$\begin{aligned}
|\psi_3\rangle &= \frac{1}{2^{n/2}} \sum_{\boldsymbol{x}\in\{0,1\}^n} (-1)^{f(\boldsymbol{x})} H^{\otimes n}\,|\boldsymbol{x}\rangle \\
&\overset{(6.4)}{=} \frac{1}{2^{n/2}} \sum_{\boldsymbol{x}\in\{0,1\}^n} (-1)^{f(\boldsymbol{x})} \frac{1}{2^{n/2}} \sum_{\boldsymbol{z}\in\{0,1\}^n} (-1)^{\boldsymbol{x}\cdot\boldsymbol{z}}\,|\boldsymbol{z}\rangle \\
&= \sum_{\boldsymbol{z}\in\{0,1\}^n} \left(\frac{1}{2^n} \sum_{\boldsymbol{x}\in\{0,1\}^n} (-1)^{f(\boldsymbol{x})+\boldsymbol{x}\cdot\boldsymbol{z}} \right) |\boldsymbol{z}\rangle
\end{aligned}$$

Only the amplitude of $|\mathbf{0}\rangle$ is needed:

$$\langle \mathbf{0}|\psi_3\rangle = \frac{1}{2^n} \sum_{\boldsymbol{x}} (-1)^{f(\boldsymbol{x})} = \frac{|\{\boldsymbol{x} : f(\boldsymbol{x}) = 0\}| - |\{\boldsymbol{x} : f(\boldsymbol{x}) = 1\}|}{2^n}$$

For a constant function this amplitude is 1 or -1, so the outcome is certainly $\mathbf{0}$. For a balanced function it is 0, so the outcome can never be $\mathbf{0}$. Consequently,

$$\text{measure } \mathbf{0} \quad \Longleftrightarrow \quad f \text{ is constant}$$

The algorithm does not determine which constant value f assumes, because the two constant functions differ only by a global phase.

8.3 The secret mask: Bernstein–Vazirani algorithm

Let a secret string $\boldsymbol{s} \in \{0,1\}^n$ define

$$f_{\boldsymbol{s}}(\boldsymbol{x}) = \boldsymbol{s} \cdot \boldsymbol{x} = \bigoplus_{j=1}^{n} (s_j \wedge x_j)$$

Classically, querying the standard basis vectors reveals one bit of $\boldsymbol{s}$ per call and therefore requires n calls:

$$f_{\boldsymbol{s}}(\boldsymbol{e}_j) = s_j$$

The quantum oracle has the reversible action

$$U_{f_{\boldsymbol{s}}}\,|\boldsymbol{x}\rangle\,|b\rangle = |\boldsymbol{x}\rangle\,|b \oplus \boldsymbol{s}\cdot\boldsymbol{x}\rangle$$

and the Bernstein–Vazirani circuit is

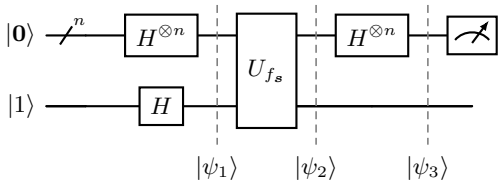

Immediately before the oracle, the state is

$$|\psi_1\rangle = \frac{1}{2^{n/2}} \sum_{\boldsymbol{x}} |\boldsymbol{x}\rangle \otimes |-\rangle$$

After phase kickback,

$$\begin{aligned}
|\psi_2\rangle &= \frac{1}{2^{n/2}} \sum_{\boldsymbol{x}} (-1)^{f_{\boldsymbol{s}}(\boldsymbol{x})}\,|\boldsymbol{x}\rangle \otimes |-\rangle \\
&= \frac{1}{2^{n/2}} \sum_{\boldsymbol{x}} (-1)^{\boldsymbol{s}\cdot\boldsymbol{x}}\,|\boldsymbol{x}\rangle \otimes |-\rangle \\
&= H^{\otimes n}\,|\boldsymbol{s}\rangle \otimes |-\rangle
\end{aligned}$$

Since $H^{-1} = H$, the last Hadamard transform recovers

$$|\psi_3\rangle = |\boldsymbol{s}\rangle \otimes |-\rangle$$

and measuring the first register returns the complete secret string with certainty after one oracle call.

Alternative interpretation

An illuminating view of the algorithm is the following.

From Section 6.3, we recall that conjugating a CNOT by Hadamards on both wires reverses its direction.

The oracle $U_{f_{\boldsymbol{s}}}$ for the function $f_{\boldsymbol{s}}(\boldsymbol{x}) = \boldsymbol{s} \cdot \boldsymbol{x}$ can be written as a sequence of CNOTs: for every position with $s_j = 1$, there is a CNOT from the j-th input Qbit to the output Qbit; for every position with $s_j = 0$, there is no such gate.

Using the Hadamards at the beginning and at the end of the algorithm, and inserting cancelling pairs of Hadamards between consecutive CNOTs, one may equivalently view these CNOTs as pointing in the opposite direction: from the output Qbit back to the input Qbits. In this reversed-CNOT picture, the output wire carries the value $|1\rangle$. Therefore, precisely those input wires for which $s_j = 1$ are flipped, while the others are left unchanged. Starting from $|\mathbf{0}\rangle$, the input register is thus transformed into $|\boldsymbol{s}\rangle$.

8.4 Simon's algorithm

Simon's problem is a collision-search problem with additional algebraic structure. Let

$$f : \{0,1\}^n \to \{0,1\}^{n-1}$$

and suppose there is a nonzero secret string $\boldsymbol{s} \in \{0,1\}^n$ such that

$$f(\boldsymbol{x}) = f(\boldsymbol{y}) \quad \Longleftrightarrow \quad \boldsymbol{x} \oplus \boldsymbol{y} \in \{\mathbf{0}, \boldsymbol{s}\} \tag{8.3}$$

Since $\boldsymbol{x} \oplus \boldsymbol{y} = \boldsymbol{s}$ is equivalent to $\boldsymbol{y} = \boldsymbol{x} \oplus \boldsymbol{s}$, the promise can be restated as

$$f(\boldsymbol{x}) = f(\boldsymbol{x} \oplus \boldsymbol{s})$$

for all $\boldsymbol{x} \in \{0,1\}^n$ and every function value has exactly two preimages. Hence, $\boldsymbol{s}$ can be seen as the period of the function f. The task is to determine $\boldsymbol{s}$.

Classical query complexity

Suppose a classical randomized algorithm makes T distinct queries. The secret becomes known as soon as it finds a collision, because then $\boldsymbol{s} = \boldsymbol{x}_i \oplus \boldsymbol{x}_j$. There are at most $\binom{T}{2}$ tested differences (pairwise XORs) among the $2^n - 1$ possible nonzero secrets, so

$$\mathbb{P}(\text{collision after } T \text{ queries}) \leq \frac{\binom{T}{2}}{2^n - 1} = \mathcal{O}\left(\frac{T^2}{2^n}\right)$$

A constant success probability therefore requires

$$T = \Omega\left(2^{n/2}\right)$$

which is the birthday bound.

Quantum algorithm

The reversible oracle acts as

$$U_f \, |\boldsymbol{x}\rangle \, |\boldsymbol{y}\rangle = |\boldsymbol{x}\rangle \, |\boldsymbol{y} \oplus f(\boldsymbol{x})\rangle$$

One run of Simon's quantum subroutine is

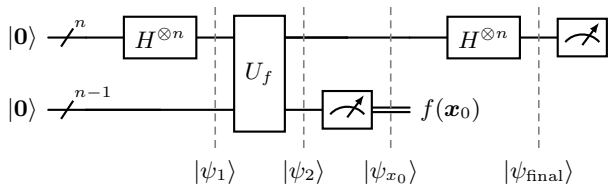

After the first Hadamard,

$$|\psi_1\rangle \overset{(6.6)}{=} \frac{1}{2^{n/2}} \sum_{\boldsymbol{x} \in \{0,1\}^n} |\boldsymbol{x}\rangle \, |\mathbf{0}\rangle$$

After the oracle,

$$|\psi_2\rangle \overset{(8.1)}{=} \frac{1}{2^{n/2}} \sum_{\boldsymbol{x}\in\{0,1\}^n} |\boldsymbol{x}\rangle\,|f(\boldsymbol{x})\rangle$$

Let's say we measure the second register and obtain $f(\boldsymbol{x}_0)$. Measuring the second register and obtaining $f(\boldsymbol{x}_0)$ collapses the first register to

$$|\psi_{\boldsymbol{x}_0}\rangle = \frac{|\boldsymbol{x}_0\rangle + |\boldsymbol{x}_0 \oplus \boldsymbol{s}\rangle}{\sqrt{2}}$$

Applying the final Hadamard,

$$\begin{aligned}
|\psi_{\text{final}}\rangle = H^{\otimes n}\,|\psi_{\boldsymbol{x}_0}\rangle &\overset{(6.4)}{=} \frac{1}{\sqrt{2}}\left(\frac{1}{2^{n/2}}\sum_{\boldsymbol{z}}(-1)^{\boldsymbol{x}_0\cdot\boldsymbol{z}}\,|\boldsymbol{z}\rangle + \frac{1}{2^{n/2}}\sum_{\boldsymbol{z}}(-1)^{(\boldsymbol{x}_0\oplus\boldsymbol{s})\cdot\boldsymbol{z}}\,|\boldsymbol{z}\rangle\right) \\
&= \frac{1}{2^{(n+1)/2}}\sum_{\boldsymbol{z}}\left((-1)^{\boldsymbol{x}_0\cdot\boldsymbol{z}} + (-1)^{(\boldsymbol{x}_0\oplus\boldsymbol{s})\cdot\boldsymbol{z}}\right)|\boldsymbol{z}\rangle \\
&\overset{(6.7)}{=} \frac{1}{2^{(n+1)/2}}\sum_{\boldsymbol{z}}(-1)^{\boldsymbol{x}_0\cdot\boldsymbol{z}}\left(1 + (-1)^{\boldsymbol{s}\cdot\boldsymbol{z}}\right)|\boldsymbol{z}\rangle \\
&= \frac{1}{2^{(n-1)/2}}\sum_{\boldsymbol{z}:\boldsymbol{s}\cdot\boldsymbol{z}=0}(-1)^{\boldsymbol{x}_0\cdot\boldsymbol{z}}\,|\boldsymbol{z}\rangle
\end{aligned}$$

Hence each run returns a uniformly random vector $\boldsymbol{z}$ satisfying the linear equation

$$\boldsymbol{s}\cdot\boldsymbol{z} = 0$$

over $\mathbb{F}_2$. The 2^{n-1} possible outcomes form the orthogonal subspace

$$\boldsymbol{s}^{\perp} = \{\boldsymbol{z}\in\mathbb{F}_2^n : \boldsymbol{s}\cdot\boldsymbol{z} = 0\}$$

Collect outcomes until they span this $(n-1)$-dimensional subspace $\boldsymbol{s}^{\perp}$, meaning that every vector $\boldsymbol{z}\in\boldsymbol{s}^{\perp}$ can be written as a componentwise XOR of some of the sampled vectors:

$$\boldsymbol{z} = \bigoplus_{i=1}^{n-1} c_i\boldsymbol{z}_i$$

with $c_i\in\{0,1\}$. Then solve the homogeneous linear system

$$\begin{cases} \boldsymbol{s}\cdot\boldsymbol{z}_1 = 0 \\ \quad\vdots \\ \boldsymbol{s}\cdot\boldsymbol{z}_{n-1} = 0 \end{cases}$$

over $\mathbb{F}_2$. Once the sampled vectors span $\boldsymbol{s}^{\perp}$, the orthogonal complement

$$(\boldsymbol{s}^{\perp})^{\perp} = \operatorname{span}\{\boldsymbol{s}\}$$

is one-dimensional, so the unique nonzero solution is the secret string.

If the current sampled vectors have rank $k < n-1$, the probability that the next vector increases the rank is

$$1 - \frac{2^k}{2^{n-1}} \geq \frac{1}{2}$$

because the current span contains 2^k vectors, whereas $\boldsymbol{s}^{\perp}$ contains 2^{n-1} vectors.

Thus each missing independent equation takes expected at most 2 samples, and there are $n-1$ such equations. Therefore the expected number of oracle calls is at most $2(n-1)$.

More precisely, after m samples, a fixed wrong candidate survives with probability 2^{-m}, since each sample is orthogonal to it with probability $1/2$. Since there are fewer than 2^n wrong candidates,

$$\begin{aligned}
\mathbb{P}(\text{failure after } m \text{ samples}) &= \mathbb{P}(\text{some wrong candidate survives}) \\
&= \mathbb{P}\left(\bigcup_{\text{wrong } \boldsymbol{a}} \{\boldsymbol{a} \text{ survives}\}\right) \\
&\leq \sum_{\text{wrong } \boldsymbol{a}} \mathbb{P}(\boldsymbol{a} \text{ survives}) \\
&< 2^n 2^{-m}
\end{aligned}$$

For $m = n + t$, this gives

$$\mathbb{P}(\text{failure}) < 2^{-t}$$

Simon's algorithm is therefore probabilistic but in expectation uses only $O(n)$ oracle calls, exponentially fewer than the classical birthday bound.

9 Intermezzo: Pseudo-telepathy

Pseudo-telepathy games expose nonlocal correlations as cooperative games. Several separated parties may agree on a strategy and share entanglement before the game, but after receiving their individual questions they may not communicate. A game exhibits pseudo-telepathy when entanglement permits a perfect strategy although no perfect classical strategy exists. As with the Bell experiments discussed in Section 1.5, entanglement correlates the answers but cannot transmit a message.

9.1 Mermin's three-party game

Alice, Bob, and Charlie receive input bits a, b, c and return output bits x, y, z. The inputs satisfy the promise

$$a \oplus b \oplus c = 1$$

and the players win if

$$x \oplus y \oplus z = a \wedge b \wedge c$$

Classical limit

Shared randomness cannot improve the best success probability beyond the best deterministic strategy, because it merely produces a convex mixture of deterministic strategies. A deterministic strategy fixes six bits $x_0, x_1, y_0, y_1, z_0, z_1$. Perfect success would require

$$\begin{aligned} x_0 \oplus y_0 \oplus z_1 &= 0 \\ x_0 \oplus y_1 \oplus z_0 &= 0 \\ x_1 \oplus y_0 \oplus z_0 &= 0 \\ x_1 \oplus y_1 \oplus z_1 &= 1 \end{aligned}$$

XORing the four left-hand sides cancels every variable twice and gives 0, while XORing the right-hand sides gives 1. Therefore no perfect classical strategy exists. For uniformly distributed promised inputs, any classical strategy wins at most three of the four cases:

$$p_{\text{win}}^{\text{classical}} \leq \frac{3}{4}$$

Quantum strategy

The players share one Qbit each of the symmetric entangled state

$$\begin{aligned} |\Psi\rangle &= \frac{1}{2\sqrt{2}} \sum_{i,j,k \in \{0,1\}} (-1)^{\text{maj}(i,j,k)} \, |ijk\rangle \\ &= \frac{1}{2\sqrt{2}} \left(|000\rangle + |001\rangle + |010\rangle - |011\rangle + |100\rangle - |101\rangle - |110\rangle - |111\rangle \right) \end{aligned} \tag{9.1}$$

On input 0, a player measures directly in the standard basis. On input 1, the player first applies a Hadamard and then measures:

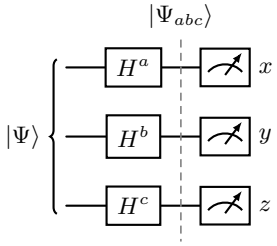

For an input with exactly one 1, every basis vector in the superposition $|\Psi_{abc}\rangle$ has even parity, so $x \oplus y \oplus z = 0$, as required because $a \wedge b \wedge c = 0$.

For $a = b = c = 1$, all players apply Hadamards:

$$|\Psi_{abc}\rangle = H^{\otimes 3} \, |\Psi\rangle = \frac{1}{2} \left(|001\rangle + |010\rangle + |100\rangle - |111\rangle \right)$$

Every possible outcome now has odd parity, so $x \oplus y \oplus z = 1$, as required because $a \wedge b \wedge c = 1$. The quantum strategy therefore wins with certainty.

9.2 The Deutsch–Jozsa game

Let $N = 2^n$. Alice receives $\boldsymbol{a} \in \{0,1\}^N$, Bob receives $\boldsymbol{b} \in \{0,1\}^N$, and they are promised that

$$d_{\mathrm{Ham}}(\boldsymbol{a}, \boldsymbol{b}) \in \{0, N/2\}$$

where $d_{\mathrm{Ham}}(\boldsymbol{a}, \boldsymbol{b}) \coloneqq |\{i \in \{1, \dots, N\} : a_i \neq b_i\}|$ denotes the Hamming distance.

They must output $\boldsymbol{x}, \boldsymbol{y} \in \{0,1\}^n$ satisfying

$$\boldsymbol{x} = \boldsymbol{y} \quad \Longleftrightarrow \quad \boldsymbol{a} = \boldsymbol{b}$$

A deterministic classical perfect strategy would be a map

$$g : \{0,1\}^N \to \{0,1\}^n$$

that assigns equal outputs to equal inputs and different outputs to every pair at Hamming distance $N/2$. Equivalently, it would color the corresponding Hadamard graph with only N colors. Such a coloring does not exist in the pseudo-telepathy instances, including $N = 16$, so no perfect classical strategy exists.

Quantum strategy

Alice and Bob share the maximally entangled state of two n-Qbit registers

$$|\Phi_N\rangle = \frac{1}{\sqrt{N}} \sum_{\boldsymbol{i} \in \{0,1\}^n} |\boldsymbol{i}\rangle\,|\boldsymbol{i}\rangle$$

Each player first applies an input-dependent phase, then a bitwise Hadamard, and finally measures:

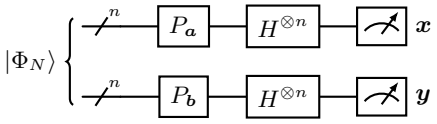

where

$$P_{\boldsymbol{a}}\,|\boldsymbol{i}\rangle = (-1)^{a_{\boldsymbol{i}}}\,|\boldsymbol{i}\rangle \qquad P_{\boldsymbol{b}}\,|\boldsymbol{i}\rangle = (-1)^{b_{\boldsymbol{i}}}\,|\boldsymbol{i}\rangle$$

After the phase gates, the joint state is

$$\frac{1}{\sqrt{N}} \sum_{\boldsymbol{i}} (-1)^{a_{\boldsymbol{i}} \oplus b_{\boldsymbol{i}}}\,|\boldsymbol{i}\rangle\,|\boldsymbol{i}\rangle$$

After both Hadamard transforms, using (6.4) on both registers, the state becomes

$$\begin{aligned}
&\frac{1}{\sqrt{N}} \sum_{\boldsymbol{i}} (-1)^{a_{\boldsymbol{i}} \oplus b_{\boldsymbol{i}}} \left(H^{\otimes n}\,|\boldsymbol{i}\rangle\right)\left(H^{\otimes n}\,|\boldsymbol{i}\rangle\right) \\
&\quad = \frac{1}{\sqrt{N}} \sum_{\boldsymbol{i}} (-1)^{a_{\boldsymbol{i}} \oplus b_{\boldsymbol{i}}} \left(\frac{1}{\sqrt{N}} \sum_{\boldsymbol{x} \in \{0,1\}^n} (-1)^{\boldsymbol{i} \cdot \boldsymbol{x}}\,|\boldsymbol{x}\rangle\right)\left(\frac{1}{\sqrt{N}} \sum_{\boldsymbol{y} \in \{0,1\}^n} (-1)^{\boldsymbol{i} \cdot \boldsymbol{y}}\,|\boldsymbol{y}\rangle\right) \\
&\quad = \frac{1}{\sqrt{N^3}} \sum_{\boldsymbol{i}} \sum_{\boldsymbol{x}} \sum_{\boldsymbol{y}} (-1)^{a_{\boldsymbol{i}} \oplus b_{\boldsymbol{i}}} (-1)^{\boldsymbol{i} \cdot \boldsymbol{x}} (-1)^{\boldsymbol{i} \cdot \boldsymbol{y}}\,|\boldsymbol{x}\rangle\,|\boldsymbol{y}\rangle \\
&\quad = \frac{1}{\sqrt{N^3}} \sum_{\boldsymbol{x}} \sum_{\boldsymbol{y}} \left(\sum_{\boldsymbol{i}} (-1)^{a_{\boldsymbol{i}} \oplus b_{\boldsymbol{i}}} (-1)^{\boldsymbol{i} \cdot (\boldsymbol{x} \oplus \boldsymbol{y})}\right) |\boldsymbol{x}\rangle\,|\boldsymbol{y}\rangle
\end{aligned}$$

If $\boldsymbol{a} = \boldsymbol{b}$, then $a_{\boldsymbol{i}} \oplus b_{\boldsymbol{i}} = 0$ for every $\boldsymbol{i}$, so the state is

$$\begin{aligned}
\frac{1}{\sqrt{N^3}} \sum_{\boldsymbol{x}} \sum_{\boldsymbol{y}} \left(\sum_{\boldsymbol{i}} (-1)^{\boldsymbol{i} \cdot (\boldsymbol{x} \oplus \boldsymbol{y})}\right) |\boldsymbol{x}\rangle\,|\boldsymbol{y}\rangle &= \frac{1}{\sqrt{N^3}} \sum_{\boldsymbol{x}} \sum_{\boldsymbol{y}} N \delta_{\boldsymbol{x},\boldsymbol{y}}\,|\boldsymbol{x}\rangle\,|\boldsymbol{y}\rangle \\
&= \frac{1}{\sqrt{N}} \sum_{\boldsymbol{x}} |\boldsymbol{x}\rangle\,|\boldsymbol{x}\rangle
\end{aligned}$$

Thus the outputs are always equal.

If $d_{\mathrm{Ham}}(\boldsymbol{a}, \boldsymbol{b}) = N/2$, then exactly half of the values $(-1)^{a_{\boldsymbol{i}} \oplus b_{\boldsymbol{i}}}$ are $+1$, and exactly half are -1. For equal outputs $\boldsymbol{x} = \boldsymbol{y}$, the corresponding part of the state has coefficient

$$\frac{1}{\sqrt{N^3}} \sum_{\boldsymbol{i}} (-1)^{a_{\boldsymbol{i}} \oplus b_{\boldsymbol{i}}} (-1)^{\boldsymbol{i} \cdot (\boldsymbol{x} \oplus \boldsymbol{x})} = \frac{1}{\sqrt{N^3}} \sum_{\boldsymbol{i}} (-1)^{a_{\boldsymbol{i}} \oplus b_{\boldsymbol{i}}} = 0$$

because $\boldsymbol{x} \oplus \boldsymbol{x} = \boldsymbol{0}$. Thus equal outputs are impossible in the second promised case, and the entangled strategy wins with certainty.

9.3 Kochen–Specker theorem

Ernst Specker argued that classical logic is inadequate for describing quantum mechanics [8]. Like Bell's theorem (Claim 1.2), discussed in Section 1.5, his argument rules out a classical deterministic explanation. The common structure is that the result should not depend on a choice that, classically, seems irrelevant. In Bell's theorem, Alice's predetermined answer should not depend on Bob's distant measurement choice; this is the locality assumption. In the Kochen–Specker theorem, the predetermined value of a ray should not depend on which other orthogonal rays are measured with it; this is the noncontextuality assumption. Thus Claim 1.2 rules out deterministic hidden variables plus *locality*, whereas Theorem 9.1 rules out deterministic hidden variables plus *noncontextuality*. In Specker's words, "quantum mechanics is bad for seers": a prophet who claims to know all future measurement outcomes does not merely risk being contradicted by a later experiment; the full table of predictions is already internally inconsistent.

The parable of the overprotective seer

Specker's 1960 article [8] begins with a parable. A seer presents three boxes to each suitor of his daughter. The suitor may choose and open any two boxes, and exactly one of those two contains a jewel, but the remaining box can never be opened:

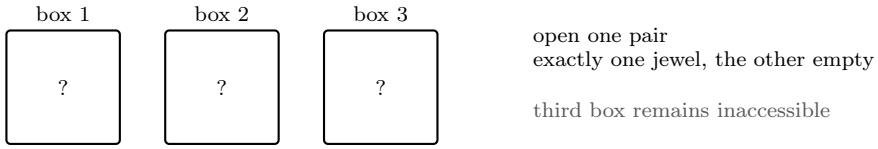

A classical preassignment $x_i \in \{0, 1\}$, where $x_i = 1$ means that box i contains a jewel, would have to satisfy

$$x_1 + x_2 = x_2 + x_3 = x_1 + x_3 = 1$$

Adding the three equations gives

$$2(x_1 + x_2 + x_3) = 3$$

which is impossible. The content of a box would therefore have to depend on which other box is opened: it would be *contextual*.

This magic-box behaviour is not itself an ordinary quantum-mechanical phenomenon. For sharp quantum yes/no measurements, pairwise joint measurability already implies joint measurability of all three questions. Equivalently, if quantum mechanics lets us answer any two of three such questions jointly, then those answers can be embedded into a joint answer for all three. The parable is therefore not the theorem itself, but a motivating foil for the genuinely quantum obstruction below.

Noncontextual predictions

As explained in Section 5.3, a measurement in an orthonormal basis

$$B = \{|v_1\rangle, \ldots, |v_d\rangle\}$$

has exactly one outcome.

Each ray in the basis also represents a yes/no property of the system. The ray $|v\rangle$ corresponds to the projector

$$P_v = |v\rangle \langle v|$$

and hence to the question whether the measurement outcome lies in that ray. In this sense, the proposition associated with P_v is: "if we test for the ray v, do we get yes?" Quantum mechanics gives only the probabilities of these outcomes, but a hidden-variable explanation would say that the individual system carries an additional hidden state λ. Many systems may have the same quantum state $|\psi\rangle$, while having different hidden states λ. The observed randomness would then be ignorance about λ, and each hidden state would give definite answers. The Kochen–Specker question is whether every possible yes/no quantum property P_v can have such a pre-existing truth value. Equivalently, it asks whether a single deterministic hidden answer sheet λ is even logically possible. Such a sheet would have to be prepared before the measurement context is known: it must answer all possible basis choices at once. For fixed λ, one would assign a value

$$c : \mathcal{V} \to \{0, 1\}$$

to each relevant ray, where $c(v) = 1$ means that the outcome corresponding to $|v\rangle$ is the answer selected by the hidden state λ if a basis containing $|v\rangle$ is measured.

The assignment must satisfy:

(i) if $\langle v|w\rangle = 0$, then $c(v) + c(w) \leq 1$, because two orthogonal rays can occur in the same measurement basis and cannot both be the unique outcome

(ii) for every orthonormal basis $B \subseteq \mathcal{V}$, $\sum_{|v\rangle \in B} c(v) = 1$, because a complete basis measurement asks which one of the mutually exclusive alternatives occurs, and exactly one outcome occurs

The value $c(v)$ is required to depend only on the ray $|v\rangle$, not on the choice of the other orthogonal rays used to complete it to a basis. Indeed, if these yes/no properties were ordinary classical properties, then the truth value of P_v should not depend on which other compatible properties are measured alongside it. This basis-independence is the *noncontextuality* assumption.

For a Qbit

$$|\psi\rangle = \alpha\,|0\rangle + \beta\,|1\rangle$$

such a coloring is possible: orthogonal pure states are antipodal points of the Bloch sphere, and one may select exactly one point of each antipodal pair. For example, choose one half of the sphere, with an arbitrary boundary convention, and color every selected ray green and its antipodal ray red. This does not mean that the quantum state $|\psi\rangle$ determines the outcome with certainty. It only means that there is no logical contradiction in imagining additional hidden variables λ that determine the outcome of each individual run.

For a Qtrit

$$|\psi\rangle = \alpha\,|0\rangle + \beta\,|1\rangle + \gamma\,|2\rangle$$

even the real rays inside the Hilbert space cannot be colored so that every orthonormal tripod contains exactly one selected ray. The direct analogue of the Qbit coloring would be a coloring of the real sphere in which every orthogonal triple contains exactly one green ray. Kochen–Specker says that this apparently local requirement cannot be satisfied globally [9]:

Theorem 9.1 (Kochen–Specker, 1967)**.** There is a finite set $\mathcal{V}$ of rays in $\mathbb{R}^3$ that admits no noncontextual coloring satisfying the two conditions (i) and (ii) above. ◁

Specker already announced the geometric obstruction in 1960, claiming that a finite set of rays on the sphere already suffices, but did not include the construction in that article [8, 11]. Kochen and Specker gave an explicit finite construction in 1967 using more than one hundred rays [9]. Thus a Qtrit cannot carry context-independent predetermined answers for all projective measurements: one may choose an outcome separately for each basis, but not by assigning fixed values to the rays themselves.

Theorem 9.1 is the precise sense in which quantum mechanics is "bad for seers/prophets": even in principle, one cannot assign context-independent predetermined answers to all these yes/no questions.[3]

A Kochen–Specker pseudo-telepathy game

Following Renner and Wolf's Kochen–Specker pseudo-telepathy formulation [10], let $\mathcal{B}$ be a finite family of real orthonormal bases whose union of rays is a Kochen–Specker set, i.e. it admits no noncontextual coloring as in Theorem 9.1. Alice receives $B_{\mathrm{A}} \in \mathcal{B}$, Bob receives $B_{\mathrm{B}} \in \mathcal{B}$, and each must output one vector from the received basis. They win exactly when their vectors are not orthogonal:

$$\langle v_{\mathrm{A}}|v_{\mathrm{B}}\rangle \neq 0$$

where $v_{\mathrm{A}} \in B_{\mathrm{A}}$ and $v_{\mathrm{B}} \in B_{\mathrm{B}}$ are the outputs of Alice and Bob, respectively.

A perfect deterministic classical strategy would select one vector from every basis while never selecting two orthogonal vectors. If the same ray occurs in two bases, the strategy must select it in both or in neither, because any different vector in the same basis is orthogonal to it. Thus coloring the selected rays by 1 and all other rays by 0 would give precisely the forbidden Kochen–Specker coloring. Hence no perfect classical strategy exists.

For the quantum strategy, Alice and Bob share the maximally entangled pair of Qtrits

$$|\Psi\rangle = \frac{1}{\sqrt{3}}\,(|00\rangle + |11\rangle + |22\rangle)$$

[3] This is similar in spirit to Claim 1.2: Alice and Bob cannot consistently prepare a table of answers for all possible measurement choices in advance. Bell rules out such tables under locality assumptions, while Kochen–Specker rules out context-independent value assignments already for a single system.

If U_A and U_B are the unitaries mapping the standard basis to B_A and B_B, each player measures in the received basis:

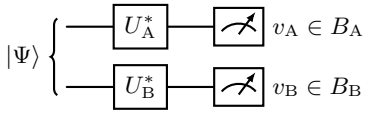

For real vectors $|v_\mathrm{A}\rangle, |v_\mathrm{B}\rangle \in \mathbb{R}^3$, the probability of the joint outcome $(v_\mathrm{A}, v_\mathrm{B})$ is

$$\begin{aligned}
\mathbb{P}(v_\mathrm{A}, v_\mathrm{B}) &\overset{(5.9)}{=} \left\| P_{\mathrm{span}\{|v_\mathrm{A}, v_\mathrm{B}\rangle\}} |\Psi\rangle \right\|^2 \overset{(5.14)}{=} \left\| |v_\mathrm{A}, v_\mathrm{B}\rangle \langle v_\mathrm{A}, v_\mathrm{B}| \, |\Psi\rangle \right\|^2 = \left\| \langle v_\mathrm{A}, v_\mathrm{B}|\Psi\rangle \, |v_\mathrm{A}, v_\mathrm{B}\rangle \right\|^2 \\
&= |\langle v_\mathrm{A}, v_\mathrm{B}|\Psi\rangle|^2 = |\langle \Psi | v_\mathrm{A}, v_\mathrm{B}\rangle|^2 \\
&\overset{(5.20)}{=} \left| \frac{1}{\sqrt{3}} \left(\langle 0|v_\mathrm{A}\rangle \langle 0|v_\mathrm{B}\rangle + \langle 1|v_\mathrm{A}\rangle \langle 1|v_\mathrm{B}\rangle + \langle 2|v_\mathrm{A}\rangle \langle 2|v_\mathrm{B}\rangle \right) \right|^2 \\
&= \frac{1}{3} \left| (v_\mathrm{A})_0 (v_\mathrm{B})_0 + (v_\mathrm{A})_1 (v_\mathrm{B})_1 + (v_\mathrm{A})_2 (v_\mathrm{B})_2 \right|^2 \\
&= \frac{1}{3} |\langle v_\mathrm{A}|v_\mathrm{B}\rangle|^2
\end{aligned}$$

In particular, if $\langle v_\mathrm{A}|v_\mathrm{B}\rangle = 0$ then $\mathbb{P}(v_\mathrm{A}, v_\mathrm{B}) = 0$, so Alice and Bob never output orthogonal vectors and win with certainty. When $B_\mathrm{A} = B_\mathrm{B}$, all distinct vectors in the common basis are orthogonal, so both players necessarily obtain the same vector. The entangled pair therefore behaves like two consistent measurements of the same Qtrit without communication.

The game uses local dimension 3×3: Alice and Bob each hold one qtrit, so the shared state lives in $\mathbb{C}^3 \otimes \mathbb{C}^3$. This is minimal for two-player pseudo-telepathy: local dimension 2×2 cannot support a perfect quantum strategy without a perfect classical one [10].

10 The Needle in the Haystack: Grover's algorithm

Grover's algorithm solves the unstructured search problem: given a Boolean function

$$f : \{0,1\}^n \to \{0,1\}$$

with a small number M of marked inputs satisfying $f(\boldsymbol{x}) = 1$, find such an input. Classically this requires $\Theta(N/M)$ queries on average, where $N = 2^n$. Grover reduces this to $\Theta(\sqrt{N/M})$ queries, which is a quadratic speed-up.

Let $|\Psi\rangle = H^{\otimes n}|\mathbf{0}\rangle$ as always (see Example 6.6, Equation 6.6). The two ingredients of Algorithm 1 are:

(i) The *oracle* O_f that implements the function f in the usual way (see Section 8.0.2):

$$O_f|\boldsymbol{x}\rangle = (-1)^{f(\boldsymbol{x})}|\boldsymbol{x}\rangle$$

and flips the phase of marked states.

(ii) The *diffusion operator*

$$D = H^{\otimes n}(2|\mathbf{0}\rangle\langle\mathbf{0}| - \mathrm{id})H^{\otimes n}$$

which reflects amplitudes about their average.

The product

$$G = DO_f$$

is called the *Grover iterate*.

Both operators really are reflections. For any state $|\varphi\rangle = a|\alpha\rangle + b|\beta\rangle$ in the plane introduced below,

$$O_f|\varphi\rangle = a|\alpha\rangle - b|\beta\rangle$$

so O_f reflects at the non-solution axis span($|\alpha\rangle$). Moreover, if $|\phi_\parallel\rangle$ denotes the part of $|\phi\rangle$ parallel to $|\Psi\rangle$ (i.e., the projection of $|\phi\rangle$ onto $|\Psi\rangle$) and $|\phi_\perp\rangle$ denotes the part of $|\phi\rangle$ perpendicular to $|\Psi\rangle$, so that $|\phi\rangle = |\phi_\parallel\rangle + |\phi_\perp\rangle$, then

$$\begin{aligned}
|\phi_{\text{reflected}}\rangle &= |\phi_\parallel\rangle - |\phi_\perp\rangle \\
&= |\phi_\parallel\rangle - (|\phi\rangle - |\phi_\parallel\rangle) \\
&= 2|\phi_\parallel\rangle - |\phi\rangle \\
&= 2\frac{\langle\Psi|\phi\rangle}{\langle\Psi|\Psi\rangle}|\Psi\rangle - |\phi\rangle \\
&= 2|\Psi\rangle\langle\Psi|\,|\phi\rangle - |\phi\rangle \\
&= (2|\Psi\rangle\langle\Psi| - \mathrm{id})|\phi\rangle \\
&= (2(H^{\otimes n}|\mathbf{0}\rangle)(\langle\mathbf{0}|H^{\otimes n}) - H^{\otimes n}H^{\otimes n})|\phi\rangle \\
&= \underbrace{H^{\otimes n}(2|\mathbf{0}\rangle\langle\mathbf{0}| - \mathrm{id})H^{\otimes n}}_{D}|\phi\rangle
\end{aligned}$$

so D reflects at span($|\Psi\rangle$).

The name *diffusion operator* comes from its action on amplitudes. Writing $|\varphi\rangle = \sum_{\boldsymbol{x}} a_{\boldsymbol{x}}|\boldsymbol{x}\rangle$ and $\bar{a} = \frac{1}{N}\sum_{\boldsymbol{x}} a_{\boldsymbol{x}}$, we obtain

$$\begin{aligned}
D|\varphi\rangle &= (2|\Psi\rangle\langle\Psi| - \mathrm{id})|\varphi\rangle \\
&= \left(2\left(\frac{1}{\sqrt{N}}\sum_{\boldsymbol{x}}|\boldsymbol{x}\rangle\right)\left(\frac{1}{\sqrt{N}}\sum_{\boldsymbol{y}}\langle\boldsymbol{y}|\right) - \mathrm{id}\right)\sum_{\boldsymbol{z}} a_{\boldsymbol{z}}|\boldsymbol{z}\rangle \\
&= 2\left(\frac{1}{\sqrt{N}}\sum_{\boldsymbol{x}}|\boldsymbol{x}\rangle\right)\left(\frac{1}{\sqrt{N}}\sum_{\boldsymbol{y}}\sum_{\boldsymbol{z}} a_{\boldsymbol{z}}\langle\boldsymbol{y}|\boldsymbol{z}\rangle\right) - \sum_{\boldsymbol{z}} a_{\boldsymbol{z}}|\boldsymbol{z}\rangle \\
&= \left(\sum_{\boldsymbol{x}}|\boldsymbol{x}\rangle\right)\left(2\frac{1}{N}\sum_{\boldsymbol{y}} a_{\boldsymbol{y}}\right) - \sum_{\boldsymbol{z}} a_{\boldsymbol{z}}|\boldsymbol{z}\rangle \\
&= \left(\sum_{\boldsymbol{x}}|\boldsymbol{x}\rangle\right)2\bar{a} - \sum_{\boldsymbol{z}} a_{\boldsymbol{z}}|\boldsymbol{z}\rangle \\
&= \sum_{\boldsymbol{x}}(2\bar{a} - a_{\boldsymbol{x}})|\boldsymbol{x}\rangle
\end{aligned}$$

Thus every amplitude is reflected about the average amplitude.

10.1 Geometric picture

Let

$$|\alpha\rangle = \frac{1}{\sqrt{N-M}} \sum_{f(\boldsymbol{x})=0} |\boldsymbol{x}\rangle \qquad |\beta\rangle = \frac{1}{\sqrt{M}} \sum_{f(\boldsymbol{x})=1} |\boldsymbol{x}\rangle$$

be the normalized superpositions of non-solutions and solutions. Then

$$|\Psi\rangle = \sqrt{\frac{N-M}{N}}\,|\alpha\rangle + \sqrt{\frac{M}{N}}\,|\beta\rangle$$

so all dynamics take place in the two-dimensional plane $\operatorname{span}\{|\alpha\rangle, |\beta\rangle\}$. Define

$$\sin\theta = \sqrt{\frac{M}{N}} \qquad \cos\theta = \sqrt{\frac{N-M}{N}} \tag{10.1}$$

Then

$$|\Psi\rangle = \cos\theta\,|\alpha\rangle + \sin\theta\,|\beta\rangle \tag{10.2}$$

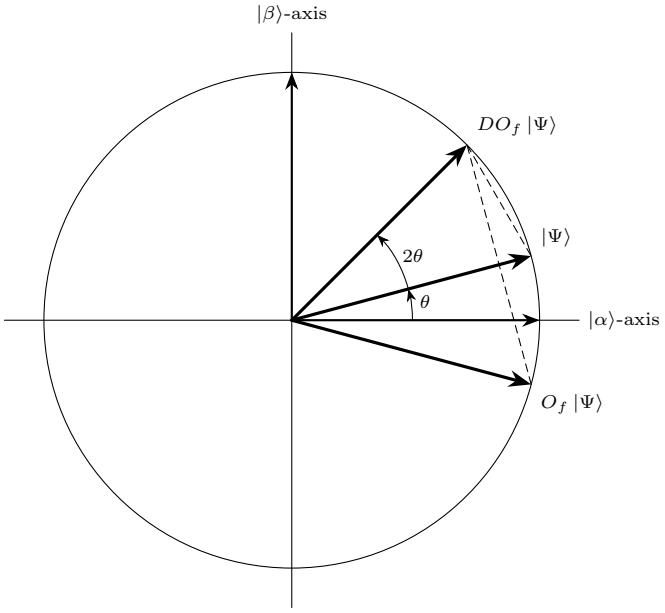

The oracle reflects at $|\alpha\rangle$, while the diffusion operator reflects at $|\Psi\rangle$. Two reflections compose to a rotation, hence each Grover iteration rotates the state toward $|\beta\rangle$, thereby amplifying the probability of observing a marked state.

After k iterations the exact state is

$$G^k\,|\Psi\rangle = \cos((2k+1)\theta)\,|\alpha\rangle + \sin((2k+1)\theta)\,|\beta\rangle \tag{10.3}$$

and hence the success probability is

$$\mathbb{P}[\text{marked after } k \text{ iterations}] = \sin^2((2k+1)\theta) \tag{10.4}$$

The optimal integer number of iterations is therefore

$$k_{\text{opt}} = \left\lfloor \frac{\pi}{4\theta} - \frac{1}{2} \right\rfloor \tag{10.5}$$

For $M \ll N$,

$$\theta \overset{(10.1)}{=} \arcsin\sqrt{\frac{M}{N}} \approx \sqrt{\frac{M}{N}}$$

and therefore

$$k_{\text{opt}} \approx \frac{\pi}{4}\sqrt{\frac{N}{M}}$$

After this many iterations a measurement returns a solution with high probability. Continuing past this point rotates the state away from the solution axis again, which is known as the *soufflé problem*: If you do not take it out of the oven in time, it collapses again.

If M is unknown, choose the number T of iterations uniformly from $\{0, \ldots, m-1\}$. Trying geometrically increasing values m retains the expected query complexity $O(\sqrt{N/M})$ without knowing M.

10.2 Circuit

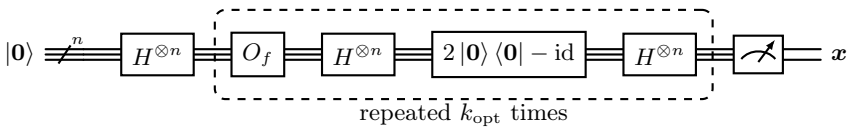

The dashed block consists of O_f and the three gates implementing $D = H^{\otimes n}(2\,|\mathbf{0}\rangle\langle\mathbf{0}| - \mathrm{id})H^{\otimes n}$ and it is repeated k_{opt} times.

Algorithm 1 Grover algorithm

Require: Oracle $O_f\,|\boldsymbol{x}\rangle = (-1)^{f(\boldsymbol{x})}\,|\boldsymbol{x}\rangle$, number of solutions M, search space size $N = 2^n$

$|\psi\rangle \leftarrow |\mathbf{0}\rangle$
$|\psi\rangle \leftarrow H^{\otimes n}\,|\psi\rangle$ ▷ equal superposition over all candidates
$\theta \leftarrow \arcsin\sqrt{M/N}$
$k \leftarrow \left\lfloor \frac{\pi}{4\theta} - \frac{1}{2} \right\rfloor$
for $j = 1, \ldots, k$ **do**
    $|\psi\rangle \leftarrow O_f\,|\psi\rangle$ ▷ flip the phase of marked states
    $|\psi\rangle \leftarrow D\,|\psi\rangle$ ▷ diffusion / reflection about the average
Measure $|\psi\rangle$ in the computational basis
return measured bit string $\boldsymbol{x}$

11 Shor's algorithm

11.1 Quantum Fourier transform

Let $N = 2^n$. The quantum Fourier transform (QFT) maps

$$\sum_{\boldsymbol{j}\in\{0,1\}^n} f_{\boldsymbol{j}} \,|\boldsymbol{j}\rangle \quad \overset{\text{QFT}}{\longmapsto} \quad \sum_{\boldsymbol{j}\in\{0,1\}^n} \hat{f}_{\boldsymbol{j}} \,|\boldsymbol{j}\rangle$$

where $\hat{\boldsymbol{f}} = [\hat{f}_0, \ldots, \hat{f}_{N-1}]$ is the discrete Fourier transform of $\boldsymbol{f} = [f_0, \ldots, f_{N-1}]$, i.e.

$$\hat{\boldsymbol{f}} = \underline{\boldsymbol{F}}_N \boldsymbol{f}$$

where $\underline{\boldsymbol{F}}_N$ is the $N \times N$ Fourier matrix with entries $[\underline{\boldsymbol{F}}_N]_{kl} = \frac{1}{\sqrt{N}}\omega^{kl}$ where $\omega \coloneqq e^{2\pi\imath/N}$ is the N-th root of unity.

Let $\boldsymbol{x} = x_1 \cdots x_n$ be an n-bit string and let $x = [\boldsymbol{x}]_2$ denote the integer represented by that string. The QFT acts on a computational basis state as

$$\text{QFT}\,|\boldsymbol{x}\rangle = \frac{1}{\sqrt{N}} \sum_{y=0}^{N-1} \omega^{xy} \,|\boldsymbol{y}\rangle = \frac{1}{2^{n/2}} \sum_{y=0}^{2^n-1} e^{2\pi\imath\, xy/2^n} \,|\boldsymbol{y}\rangle \tag{11.1}$$

where $\boldsymbol{y} \in \{0,1\}^n$ is the binary representation of y.

Writing

$$y = [\boldsymbol{y}]_2 = \sum_{k=1}^{n} y_k 2^{n-k}$$

we obtain

$$\begin{aligned}
\text{QFT}\,|\boldsymbol{x}\rangle &= \frac{1}{2^{n/2}} \sum_{y=0}^{2^n-1} e^{2\pi\imath\, xy/2^n} \,|\boldsymbol{y}\rangle \\
&= \frac{1}{2^{n/2}} \sum_{y=0}^{2^n-1} e^{2\pi\imath\, x\left(\sum_{k=1}^{n} y_k 2^{n-k}\right)/2^n} \,|\boldsymbol{y}\rangle \\
&= \frac{1}{2^{n/2}} \sum_{y=0}^{2^n-1} \prod_{k=1}^{n} \left(e^{2\pi\imath\, x/2^k}\right)^{y_k} |\boldsymbol{y}\rangle \\
&= \frac{1}{2^{n/2}} \sum_{y=0}^{2^n-1} \bigotimes_{k=1}^{n} \left(\left(e^{2\pi\imath\, x/2^k}\right)^{y_k} |y_k\rangle\right) \\
&= \frac{1}{2^{n/2}} \sum_{\boldsymbol{y}\in\{0,1\}^n} \bigotimes_{k=1}^{n} \left(\left(e^{2\pi\imath\, x/2^k}\right)^{y_k} |y_k\rangle\right) \\
&= \frac{1}{2^{n/2}} \bigotimes_{k=1}^{n} \left(\sum_{y_k\in\{0,1\}} \left(e^{2\pi\imath\, x/2^k}\right)^{y_k} |y_k\rangle \right) \\
&= \frac{1}{2^{n/2}} \bigotimes_{k=1}^{n} \left(|0\rangle + e^{2\pi\imath\, x/2^k} \,|1\rangle\right)
\end{aligned}$$

Since

$$\frac{x}{2^k} = \text{integer} + 0.x_{n-k+1} \cdots x_n$$

in binary notation, the integer part contributes no phase because $e^{2\pi\imath m} = 1$ for every integer m. Hence

$$\text{QFT}\,|\boldsymbol{x}\rangle = \frac{1}{2^{n/2}} \bigotimes_{k=1}^{n} \left(|0\rangle + e^{2\pi\imath\, 0.x_{n-k+1}\cdots x_n} \,|1\rangle\right)$$

This is the QFT output in the standard computational-basis order. In circuit implementations, the tensor factors are often produced in the opposite wire order; one then adds final SWAP gates to reverse the order of the output bits.

Circuit

Define the phase gates

$$R_k \coloneqq \begin{bmatrix} 1 & 0 \\ 0 & e^{2\pi\imath/2^k} \end{bmatrix} \tag{11.2}$$

A Hadamard on the j-th Qbit, followed by controlled phase gates from the less significant Qbits, prepares the factor with phase $0.x_j x_{j+1} \cdots x_n$. Thus the circuit produces the tensor factors in the opposite order compared with the product representation above; the final SWAP gates restore the standard computational-basis order. The general n-Qbit circuit is

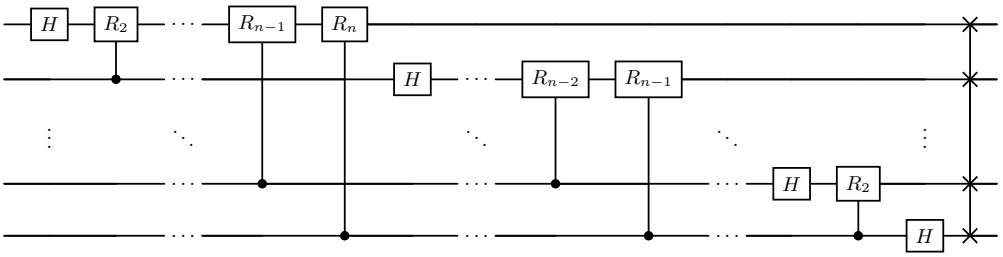

There are $n(n-1)/2$ controlled rotations, n Hadamards, and $\lfloor n/2 \rfloor$ swaps, hence the exact QFT uses $O(n^2)$ gates [12].

11.2 Phase estimation

Let U be unitary and suppose an eigenstate $|u\rangle$ is available with

$$U|u\rangle = e^{2\pi\imath\varphi}|u\rangle \qquad \varphi \in [0,1) \tag{11.3}$$

The goal is to determine the first t binary digits of φ. Controlled powers of U transfer the phase to a t-Qbit control register by phase kickback:

$$|j\rangle|u\rangle \overset{C_U}{\longmapsto} |j\rangle U^j|u\rangle = e^{2\pi\imath j\varphi}|j\rangle|u\rangle \tag{11.4}$$

Let n_u denote the number of Qbits in the register containing $|u\rangle$. In compact register-level notation, phase estimation is

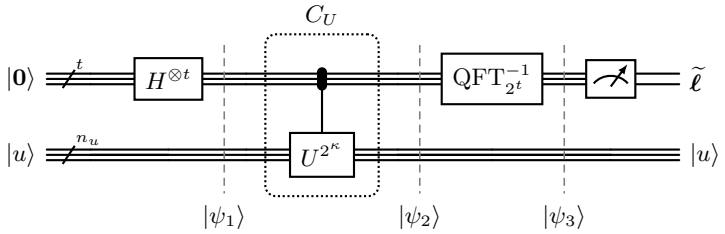

The slices in the circuit mark the following joint states. Initially,

$$|\psi_0\rangle = |\mathbf{0}\rangle|u\rangle \tag{11.5}$$

After the Hadamards, the control register is the usual equal superposition:

$$|\psi_1\rangle = \left(H^{\otimes t} \otimes \mathrm{id}\right)|\psi_0\rangle = \frac{1}{\sqrt{2^t}} \sum_{j=0}^{2^t-1} |j\rangle|u\rangle \tag{11.6}$$

The compact controlled block represents the unitary

$$C_U \coloneqq \sum_{j=0}^{2^t-1} |j\rangle\langle j| \otimes U^j \tag{11.7}$$

Indeed, the binary expansion $j = \sum_{k=0}^{t-1} j_k 2^k$ selects precisely the powers U^{2^k} whose control bits satisfy $j_k = 1$, and their product is

$$\prod_{k=0}^{t-1} \left(U^{2^k}\right)^{j_k} = U^{\sum_{k=0}^{t-1} j_k 2^k} = U^j$$

Applying this block to $|\psi_1\rangle$ and using $U^j |u\rangle = e^{2\pi\imath j\varphi} |u\rangle$ gives

$$|\psi_2\rangle = C_U |\psi_1\rangle = \frac{1}{\sqrt{2^t}} \sum_{j=0}^{2^t-1} |j\rangle U^j |u\rangle = \frac{1}{\sqrt{2^t}} \sum_{j=0}^{2^t-1} e^{2\pi\imath j\varphi} |j\rangle |u\rangle \tag{11.8}$$

Suppose first that the phase has an exact t-bit representation $\varphi = \frac{\ell}{2^t}$, $\ell \in \{0, \ldots, 2^t - 1\}$ (ideal case). Substitution into Equation 11.8 gives

$$|\psi_2\rangle = \frac{1}{\sqrt{2^t}} \sum_{j=0}^{2^t-1} e^{2\pi\imath j\ell/2^t} |j\rangle |u\rangle \overset{(11.1)}{=} (\mathrm{QFT}_{2^t} |\ell\rangle) |u\rangle \tag{11.9}$$

where the second equality is exactly the definition of the QFT from Section 11.1. Consequently, the inverse QFT produces

$$|\psi_3\rangle = \left(\mathrm{QFT}_{2^t}^{-1} \otimes \mathrm{id}\right) |\psi_2\rangle = |\ell\rangle |u\rangle \tag{11.10}$$

and measurement returns ℓ with certainty.

In the non-ideal case, we can extend the circuit: Adding s extra Qbits and discarding their measured bits gives the desired t-bit approximation with probability at least $1 - 2^{-s}$.

11.3 Number Theory

Definition 11.1 (Group)**.** A *group* is a set G together with a binary operation

$$\circ : G \times G \to G$$

such that:

1. Associativity:	$\forall a, b, c \in G$	$(a \circ b) \circ c = a \circ (b \circ c)$
2. Identity element:	$\exists e \in G\ \forall a \in G$	$e \circ a = a \circ e = a$
3. Inverse element:	$\forall a \in G\ \exists a^{-1} \in G$	$a \circ a^{-1} = a^{-1} \circ a = e$

◀

Remark 11.1. Uniqueness of the identity and uniqueness of inverse elements are not part of the axioms; they are consequences of the three axioms. ◀

Definition 11.2 (Cyclic group)**.** A group G is called *cyclic* if there exists an element $g \in G$ such that every element of G can be written as a power of g, i.e.

$$G = \{g^n : n \in \mathbb{Z}\}$$

The element g is called a *generator* of G, and we write $G = \langle g \rangle$. ◀

Definition 11.3 (Residue classes and units)**.** The integers modulo N form

$$\mathbb{Z}_N \coloneqq \{0, \ldots, N-1\}$$

with addition and multiplication modulo N. Their multiplicative group of units is

$$\mathbb{Z}_N^* \coloneqq \{x \in \mathbb{Z}_N : \gcd(x, N) = 1\}$$

which contains the numbers coprime to N. ◀

Definition 11.4 (Order)**.** The order of an element a in a finite group is

$$\mathrm{ord}(a) \coloneqq \min\{r > 0 : a^r = e\}$$

For $a \in \mathbb{Z}_N^*$, this means $\mathrm{ord}_N(a) \coloneqq \min\{r > 0 : a^r \equiv 1 \pmod N\}$. ◀

Remark 11.2. If $N = pq$ for distinct primes, then $\mathbb{Z}_N^*$ consists of the residue classes modulo N that are not divisible by p or by q:

$$\mathbb{Z}_N^* = \{x \in \{0, \ldots, N-1\} : p \nmid x \text{ and } q \nmid x\}$$

Thus, to obtain $\mathbb{Z}_N^*$ from $\mathbb{Z}_N$, one has to remove all multiples of p and all multiples of q. There are q multiples of p modulo N and p multiples of q. The residue 0 is counted twice, so the number of non-units is $p + q - 1$. Therefore

$$|\mathbb{Z}_N^*| = pq - (p + q - 1) = (p-1)(q-1)$$ ◀

11.4 Order finding

The following order-finding routine is the quantum subroutine behind Shor's factoring algorithm [14, 13].

For $a \in \mathbb{Z}_N^*$, define the modular multiplication unitary

$$\boxed{U_a \left|y\right\rangle = \left|ay \bmod N\right\rangle} \tag{11.11}$$

on basis states $0 \le y < N$, extending it as the identity on unused computational-basis states. It is unitary because $a \in \mathbb{Z}_N^*$, so multiplication by a is a permutation of $\mathbb{Z}_N$.

A controlled U^{2^κ} can be implemented efficiently by translating the classical "repeated squaring" method in time $O(\log N)^3 = O(n^3)$, which is the asymptotic running time of Algorithm 2. (This can be improved to essentially $O(\log N)^2$ by using an asymptotically faster multiplication algorithm — based also on the discrete Fourier transform.)

If $r = \operatorname{ord}_N(a)$, then for $0 \le s < r$,

$$\left|u_s\right\rangle \coloneqq \frac{1}{\sqrt{r}} \sum_{k=0}^{r-1} e^{-2\pi\imath sk/r} \left|a^k \bmod N\right\rangle \tag{11.12}$$

is an eigenstate of U_a with eigenvalue $e^{2\pi\imath s/r}$, because

$$\begin{aligned} U_a \left|u_s\right\rangle &= \frac{1}{\sqrt{r}} \sum_{k=0}^{r-1} e^{-2\pi\imath sk/r} \left|a^{k+1} \bmod N\right\rangle \\ &= e^{2\pi\imath s/r} \frac{1}{\sqrt{r}} \sum_{k=0}^{r-1} e^{-2\pi\imath sk/r} \left|a^k \bmod N\right\rangle \\ &= e^{2\pi\imath s/r} \left|u_s\right\rangle \end{aligned}$$

where we used the r-periodicity of both functions. We conclude that phase estimation yields digits of the phase $\varphi = \frac{s}{r}$, and given that we know sufficiently many, i.e., $O(\log N)$, digits, we can determine the period and, hence, the rational number, in particular r, which is the unknown.

But how do we obtain one of the $\left|u_s\right\rangle$? We cannot: The definition of $\left|u_s\right\rangle$ depends on r. But luckily,

$$\begin{aligned} \frac{1}{\sqrt{r}} \sum_{s=0}^{r-1} \left|u_s\right\rangle &= \frac{1}{r} \sum_{s=0}^{r-1} \sum_{k=0}^{r-1} e^{-\frac{2\pi\imath sk}{r}} \left|a^k \bmod N\right\rangle \\ &= \frac{1}{r} \sum_{k=0}^{r-1} \left(\sum_{s=0}^{r-1} e^{-\frac{2\pi\imath sk}{r}} \right) \left|a^k \bmod N\right\rangle \\ &= \frac{1}{r} \sum_{k=0}^{r-1} r\ \delta_{k,0} \left|a^k \bmod N\right\rangle \\ &= \left|1\right\rangle \end{aligned}$$

so an eigenstate does not need to be prepared explicitly. Let $n = \lceil \log_2 N \rceil$. The order-finding circuit is phase estimation for U_a, written at register level with the same controlled-unitary convention as in Section 11.2:

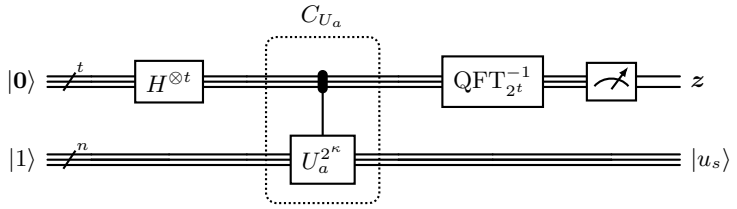

Here the dotted block denotes the controlled operation

$$C_{U_a} = \sum_{j=0}^{2^t-1} \left|j\right\rangle\left\langle j\right| \otimes U_a^j$$

from (11.7). In the expanded circuit, this is implemented by the controlled powers $U_a^{2^\kappa}$ for $0 \le \kappa < t$. Equivalently, on basis states,

$$C_{U_a} \left|j\right\rangle \left|y\right\rangle = \left|j\right\rangle \left|a^j y \bmod N\right\rangle$$

which is the usual coherent modular exponentiation.

The measured bit string $\boldsymbol{z}$ encodes an approximation $[\boldsymbol{z}]_2/2^t \approx s/r$. Repeated squaring makes every controlled power efficient:

$$U_a^{2^\kappa} |y\rangle = |a^{2^\kappa} y \bmod N\rangle$$

and the constants $a^{2^\kappa} \bmod N$ are precomputed classically. Phase estimation (Section 11.2) returns s/r approximately, and continued fractions recover r.

11.5 Integer Factoring

Shor's factoring algorithm combines the quantum order-finding subroutine from Section 11.4 with a classical reduction [14, 13]. This puts factoring in the complexity class BQP, but as we will see in Section 12, this by itself does not prove P $\subsetneq$ BQP. Let N be an odd composite integer that is not a prime power and choose $a \in \{2, \ldots, N-1\}$ uniformly. If $\gcd(a, N) > 1$, a factor has already been found. Otherwise compute $r = \mathrm{ord}_N(a)$.

If r is even, then

$$\begin{aligned} a^r &\equiv 1 \pmod N \\ a^r - 1 &\equiv 0 \pmod N \\ (a^{r/2} - 1)(a^{r/2} + 1) &\equiv 0 \pmod N \end{aligned} \tag{11.13}$$

Minimality of r guarantees $a^{r/2} \not\equiv 1 \pmod N$. If also $a^{r/2} \not\equiv -1 \pmod N$, then neither factor in Equation 11.13 is divisible by all of N, while their product is. Therefore

$$\gcd(a^{r/2} - 1, N) \qquad \gcd(a^{r/2} + 1, N) \tag{11.14}$$

yield non-trivial factors. For an admissible N, a random $a \in \mathbb{Z}_N^*$ has even order and avoids $a^{r/2} \equiv -1 \pmod N$ with probability at least $1/2$.

Algorithm 2 Shor's factoring algorithm

Require: Composite integer N
1: **if** N is even **then**
2: ⌊ **return** 2
3: Check classically whether N is a prime power
4: Choose a uniformly from $\{2, \ldots, N-1\}$ ▷ Random base for the reduction to order finding
5: $d \leftarrow \gcd(a, N)$ ▷ Check whether a is already not coprime to N
6: **if** $d > 1$ **then** ▷ $a \notin \mathbb{Z}_N^*$
7: ⌊ **return** d ▷ The gcd already gives a non-trivial factor
8: $r \leftarrow \mathrm{ord}_N(a)$ using quantum order finding ▷ See Definition 11.4 and Section 11.4
9: **if** r is odd or $a^{r/2} \equiv -1 \pmod N$ **then** ▷ Then (11.13) does not yield useful gcds
10: ⌊ Restart ▷ For a random $a \in \mathbb{Z}_N^*$, success has constant probability
11: **return** $\gcd(a^{r/2} - 1, N)$ ▷ A non-trivial factor from (11.14)

11.6 Discrete Logarithms

Shor's algorithm also solves the discrete logarithm problem [14, 13]. The factoring algorithm above is slightly indirect: the quantum part solves order finding, and the classical part reduces factoring to order finding. Discrete logarithms are cleaner, because the relevant group order is part of the input, so the phase estimation can be made exact with a QFT_q. This is also why the same idea applies not only to $\mathbb{Z}_p^*$, but to any efficiently represented cyclic group, including elliptic-curve groups.

Problem 11.5 (Discrete logarithm)**.** Let

$$G = \langle g \rangle = \left\{1, g, g^2, \ldots, g^{q-1}\right\}$$

be a finite cyclic group of known order q. Given $h = g^\ell$, determine $\ell \in \mathbb{Z}_q$. ◀

Assume that group multiplication and inversion are efficiently computable (i.e. in polynomial time in $\log q$). Define the shift unitaries $U|x\rangle = |gx\rangle$, $V|x\rangle = |hx\rangle$ on group-element basis states. They are unitary because multiplication by a fixed group element is a permutation of G. Since $h = g^\ell$, we have $V = U^\ell$. Thus U and V have the same eigenvectors, and their eigenvalues differ by the unknown factor ℓ.

For $s \in \mathbb{Z}_q$ (Definition 11.3), define

$$|u_s\rangle \coloneqq \frac{1}{\sqrt{q}} \sum_{k=0}^{q-1} e^{-2\pi\imath sk/q} |g^k\rangle$$

Then

$$\begin{aligned} U|u_s\rangle &= \frac{1}{\sqrt{q}} \sum_{k=0}^{q-1} e^{-2\pi\imath sk/q} |g^{k+1}\rangle \\ &= e^{2\pi\imath s/q} \frac{1}{\sqrt{q}} \sum_{k=0}^{q-1} e^{-2\pi\imath sk/q} |g^k\rangle \\ &= e^{2\pi\imath s/q} |u_s\rangle \end{aligned}$$

and therefore

$$V|u_s\rangle = U^\ell |u_s\rangle = e^{2\pi\imath s\ell/q} |u_s\rangle$$

As in Section 11.4, the eigenstates need not be prepared directly. Orthogonality of roots of unity gives $\frac{1}{\sqrt{q}} \sum_{s=0}^{q-1} |u_s\rangle = |1\rangle$.

For an eigenstate $|u_s\rangle$, exact phase estimation (Section 11.2) with QFT_q gives

$$|0\rangle |0\rangle |u_s\rangle \xmapsto{\mathrm{PE}(U),\,\mathrm{PE}(V)} |s\rangle |s\ell \bmod q\rangle |u_s\rangle$$

Starting the lower register in $|1\rangle$ therefore gives a uniformly random s (one out of $q = |\mathbb{Z}_q|$ possible values):

$$|0\rangle |0\rangle |1\rangle \mapsto \frac{1}{\sqrt{q}} \sum_{s=0}^{q-1} |s\rangle |s\ell \bmod q\rangle |u_s\rangle$$

A register-level view of the two phase estimations is:

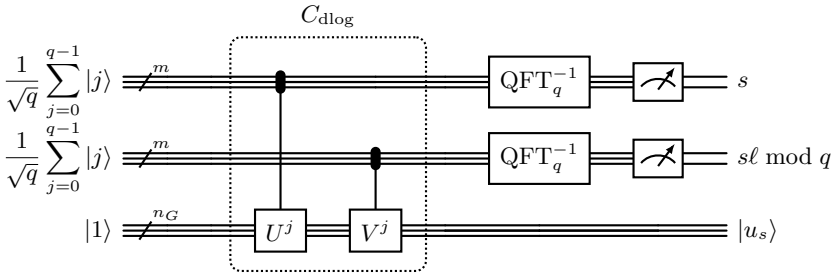

where

$$C_{\mathrm{dlog}} |j_1\rangle |j_2\rangle |z\rangle = |j_1\rangle |j_2\rangle |g^{j_1} h^{j_2} z\rangle$$

and $m = \lceil \log_2 q \rceil$ only indicates the number of Qbits needed to encode q basis states. If q is not a power of two, one can either implement the exact QFT_q on the q-dimensional subspace or use the usual high-precision power-of-two phase estimation as in Section 11.2.

After measurement we get

$$(x, y) = (s, s\ell \bmod q) \in \mathbb{Z}_q^2$$

If $s \in \mathbb{Z}_q^*$, i.e. s is invertible[4] modulo q i.e. s is coprime to q i.e. $\gcd(s, q) = 1$, the extended Euclidean algorithm gives s^{-1}, and

$$\ell \equiv y s^{-1} \pmod{q}$$

If $s \notin \mathbb{Z}_q^*$, then s is not invertible modulo q, so the congruence $y \equiv s\ell \pmod{q}$ cannot be divided by s, because it does not determine ℓ uniquely in that case. In fact, if $d = \gcd(s, q) > 1$, then this congruence determines ℓ only modulo q/d, and hence leaves d possible residue classes modulo q. Thus this measurement gives only partial information about ℓ, not the unique discrete logarithm. We therefore discard this run and repeat until $s \in \mathbb{Z}_q^*$. The success probability of one run is

$$\mathbb{P}(s \perp q) = \mathbb{P}(s \in \mathbb{Z}_q^*) = \frac{|\mathbb{Z}_q^*|}{q}$$

where the numerator is the so-called Euler totient function, which counts the number of integers coprime to q. It is known to grow with $\Omega(q/\log\log q)$, implying that the success probability is $\Omega(1/\log\log q)$ [17].

[4] If $\gcd(s, q) = 1$, then cancellation by s is valid modulo q: $sa \equiv sb \pmod{q} \implies a \equiv b \pmod{q}$. Hence multiplication by s modulo q is injective, and therefore, on the finite set $\mathbb{Z}_q$, a permutation.

Thus the expected number of repetitions is $O(\log \log q)$.

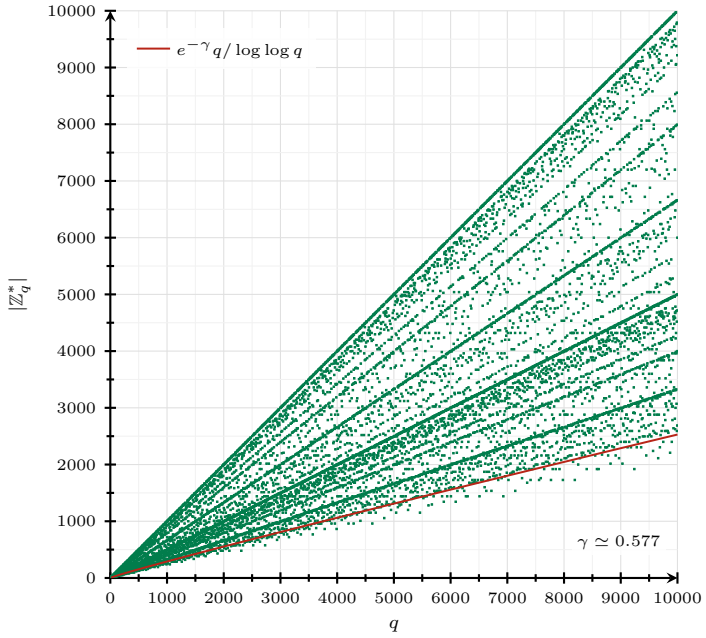

Let $n = \lceil \log_2 q \rceil$ be the input length of the group order, and let $p_{\max}$ be the largest prime divisor of q. Let T_G denote the number of gates needed to implement one group operation coherently. For the repeat-until-$s \in \mathbb{Z}_q^*$ version above, the expected gate count is

$$O\big((T_G \log q + \log^2 q) \log \log q\big)$$

because one run uses $O(T_G \log q + \log^2 q)$ quantum gates and $O(\log \log q)$ expected repetitions [13, 12, 17].

Generic classical algorithms, by contrast, need

$$\Omega(\sqrt{p_{\max}})$$

group operations in Shoup's generic-group model [15]. In the prime-order case $p_{\max} = q$, and this becomes

$$\Omega(\sqrt{q}) = \Omega(2^{n/2})$$

which is exponential in the input length n [15, 16].

12 Quantum Complexity Theory

A decision problem is represented by a language

$$L \subseteq \{0,1\}^* := \bigcup_{n=0}^{\infty} \{0,1\}^n$$

and solving the decision problem means deciding membership $x \in L$. A complexity class is then a collection of languages that can be decided with a prescribed computational model and prescribed resource bounds.

The three knobs are:

- resource: time, space, queries, circuit size
- kind of task: decision, search, optimization, approximation
- computational model: deterministic, randomized, quantum

12.1 Classical complexity classes

Definition 12.1 (Some classical complexity classes)**.** Following [18], for languages $L \subseteq \{0,1\}^*$:

- L is the class decidable using $O(\log n)$ working space. The input tape is read-only; the logarithmic bound applies to the work tape.
- P is the class decidable in polynomial time.
- NP is the class with polynomial-size certificates verifiable in polynomial time.
- PSPACE is the class decidable using polynomial working space, with no polynomial bound on time.
- EXP is the class decidable in exponential time $O(2^{n^k})$ for some constant k. ◀

Claim 12.1. $\mathrm{L} \subseteq \mathrm{P} \subseteq \mathrm{NP} \subseteq \mathrm{PSPACE} \subseteq \mathrm{EXP}$ ◁

Proof. The inclusions have simple resource-counting proofs.

L ⊆ P: a deterministic $O(\log n)$-space machine has only polynomially many configurations: polynomially many input-head positions, polynomially many work-tape contents, polynomially many work-head positions, and constantly many internal states. A decider cannot visit the same configuration twice without looping, so it halts in polynomial time[5].

P ⊆ NP: use the empty certificate and run the polynomial-time decision algorithm.

NP ⊆ PSPACE: if the certificate length is bounded by $p(n)$, enumerate all $2^{p(n)}$ certificates one after another. Each verification uses polynomial space, and the work tape can be erased between trials.

PSPACE ⊆ EXP: a polynomial-space machine has at most exponentially many configurations. If it is a decider, it must halt before repeating a configuration. □

The time hierarchy theorem implies $\mathrm{P} \subsetneq \mathrm{EXP}$ and the space hierarchy theorem implies $\mathrm{L} \subsetneq \mathrm{PSPACE}$. Thus at least one inclusion in the chain of Claim 12.1 must be strict, but it is not known which one(s).

12.2 Bounded-error computation

Randomized and quantum algorithms are usually allowed a small probability of error. The constant is not important as long as it is bounded away from 1/2, because independent repetitions and majority vote amplify the success probability[6]. For example, success probability 3/4 can be amplified to $1 - \varepsilon$ with $O(\log(1/\varepsilon))$ repetitions.

Definition 12.2 (BPP)**.** BPP (Bounded-error Probabilistic Polynomial-time) is the class of languages decidable by a randomized polynomial-time classical algorithm with bounded two-sided error [18]. ◀

[5] An L-machine is a decider, it is not allowed to loop. Therefore it must halt before any repetition happens.
[6] 1/2 is the "random guessing" threshold for a yes/no answer.

Definition 12.3 (BQP)**.** BQP (Bounded-error Quantum Polynomial-time) is the quantum analogue of Definition 12.2 [18]. A language L is in BQP if there exists a *uniform* family of quantum circuits

$$\{C_n\}_{n\in\mathbb{N}}$$

such that:

- C_n has size polynomial in n
- on input $x \in \{0,1\}^n$, the circuit C_n decides whether $x \in L$ by measuring a designated output Qbit
- if $x \in L$, the circuit accepts with probability at least $3/4$
- if $x \notin L$, the circuit rejects with probability at least $3/4$

Uniform means that there is a classical Turing machine which, on input 1^n, outputs a description of C_n in time polynomial in n. ◀

The family $\{C_n\}$ is needed because one fixed circuit only has finitely many input wires. Uniformity rules out hiding an arbitrarily hard language in the choice of the circuits themselves: there must be an efficient classical recipe for constructing the circuit for length n.

Fact 12.2. $\mathrm{P} \subseteq \mathrm{BPP} \subseteq \mathrm{BQP} \subseteq \mathrm{PSPACE}$ ◁

Claim 12.1 and Fact 12.2 are summarized in the following diagram:

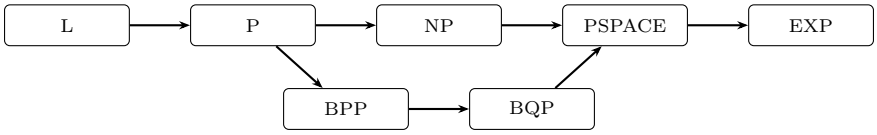

The inclusion $\mathrm{BPP} \subseteq \mathrm{BQP}$ is the easy one: a quantum circuit can simulate a randomized classical polynomial-time computation.

Random bits are generated directly in parallel. To obtain k independent random bits, prepare $|\mathbf{0}\rangle = |0\rangle^{\otimes k}$, apply Hadamards to all k Qbits, and measure in the computational basis. As seen in Example 6.6,

$$H^{\otimes k}\,|\mathbf{0}\rangle \overset{(6.6)}{=} \frac{1}{\sqrt{2^k}} \sum_{\boldsymbol{r}\in\{0,1\}^k} |\boldsymbol{r}\rangle$$

so the Born rule from Section 5.3 gives every outcome $\boldsymbol{r} \in \{0,1\}^k$ probability $1/2^k$. Thus the measurement produces a uniformly random k-bit string.

The deterministic part of the classical computation can then be embedded into a reversible computation as in Section 2.5. At the gate level, Section 2.6 explains why Toffoli gates are enough for reversible classical computation, and the classical-universality part of Section 6.3 explains how CNOT plus one-Qbit gates gives Toffoli.

The nontrivial and important inclusion is

12.3 BQP $\subseteq$ PSPACE

The proof below follows the standard polynomial-space simulation argument in [18]. Let n be the input length. After adding polynomially many ancilla Qbits, the circuit acts on $m(n) = \mathrm{poly}(n)$ Qbits and has $T(n) = \mathrm{poly}(n)$ elementary gates. We assume a fixed finite universal gate set, or more generally elementary gates whose matrix entries can be computed to polynomially many bits using polynomial space.

By Fact 6.2, we may assume that the computation is unitary until a final measurement of one designated output Qbit. Write the unitary part as

$$C = U_T \cdots U_1$$

where each U_t denotes the full $2^m \times 2^m$ unitary on $\mathcal{H}_m = (\mathbb{C}^2)^{\otimes m}$ obtained by applying one elementary one- or two-Qbit gate to its wires and the identity to all other wires.

The point is not that every m-Qbit unitary is counted as one allowed operation. Rather, the circuit model counts local elementary gates. One- and two-Qbit gate sets are universal in the sense

discussed around Section 6.3, but a generic m-Qbit unitary may require exponentially many local gates to implement. Allowing such a unitary as a single U_t would hide the exponential complexity inside one gate.

Let the initial computational-basis state be

$$|\eta_x\rangle = |x\rangle\,|0\rangle^{\otimes a(n)}$$

with $a(n) = m(n) - n$, where the second register contains the ancilla Qbits.

First fix a final basis state $|\boldsymbol{y}\rangle$. To compute its amplitude, insert the computational-basis resolution of the identity from (6.3), with n replaced by m,

$$\sum_{\boldsymbol{z}\in\{0,1\}^m} |\boldsymbol{z}\rangle\,\langle\boldsymbol{z}| = \mathrm{id}_{\mathcal{H}_m}$$

between every pair of gates:

$$\langle\boldsymbol{y}|U_T\cdots U_1|\eta_x\rangle = \sum_{\boldsymbol{z}_1,\ldots,\boldsymbol{z}_{T-1}} \langle\boldsymbol{y}|U_T|\boldsymbol{z}_{T-1}\rangle\,\langle\boldsymbol{z}_{T-1}|U_{T-1}|\boldsymbol{z}_{T-2}\rangle\cdots\langle\boldsymbol{z}_2|U_2|\boldsymbol{z}_1\rangle\,\langle\boldsymbol{z}_1|U_1|\eta_x\rangle \quad (12.1)$$

This is a sum over many computational histories: each intermediate $\boldsymbol{z}_t$ ranges over 2^m basis strings, and there are $T-1$ such intermediate times. Thus the number of summands is $2^{m(T-1)} = 2^{\mathrm{poly}(n)}$, because $m = \mathrm{poly}(n)$ and $T = \mathrm{poly}(n)$. This may take exponential time, but only polynomial space: the classical simulator can run nested loops over $\boldsymbol{z}_1,\ldots,\boldsymbol{z}_{T-1}$, store the current m-bit strings, the current product, and the running sum, and then reuse this workspace after each summand has been added. Although exponentially many summands are added, accuracy $2^{-\mathrm{poly}(n)}$ still requires only polynomially many bits, and is enough to make the final additive error smaller than any fixed constant.

Now use that BQP is a decision class. The simulator does not need to find the most likely output string, nor sample from the whole final distribution. It only needs the probability that the designated output Qbit is 1. If

$$A \coloneqq \{\boldsymbol{y}\in\{0,1\}^m : y_{\mathrm{out}} = 1\}$$

is the set of accepting basis states, then

$$\mathbb{P}(\mathrm{accept}) = \sum_{\boldsymbol{y}\in A} |\langle\boldsymbol{y}|U_T\cdots U_1|\eta_x\rangle|^2 \quad (12.2)$$

This outer sum is exponential in the number m of Qbits: if one output Qbit is designated, then roughly half of the 2^m basis states have $y_{\mathrm{out}} = 1$, so

$$|A| = 2^{m-1} = 2^{\mathrm{poly}(n)}$$

Again this is allowed in PSPACE, because only space is restricted. The simulator enumerates the accepting $\boldsymbol{y}$'s one after another, computes the amplitude for the current $\boldsymbol{y}$ using the history sum (12.1), adds the squared modulus to a running total, and erases the temporary amplitude.

Finally use the bounded-error gap in the definition of BQP. For every input x, the circuit is guaranteed to be in one of the two separated cases

$$x\in L \implies \mathbb{P}(\mathrm{accept}) \geq \frac{3}{4}, \qquad x\notin L \implies \mathbb{P}(\mathrm{accept}) \leq \frac{1}{4}$$

The gap between the two thresholds is $1/2$. Hence the classical simulator does not need the exact acceptance probability.

Remark 12.1. Mathematically, the acceptance probability (12.2) is exact. The error below is not quantum error; it is only numerical approximation error from storing amplitudes and partial sums with finite precision on the classical work tape. ◀

Any additive error strictly smaller than half of this gap is enough. Fix any constant $0 < \varepsilon < 1/4$. If the simulator computes an estimate $\widetilde{p}$ with

$$|\widetilde{p} - \mathbb{P}(\mathrm{accept})| \leq \varepsilon$$

then

$$x\in L \implies \widetilde{p} \geq \frac{3}{4} - \varepsilon > \frac{1}{2}, \qquad x\notin L \implies \widetilde{p} \leq \frac{1}{4} + \varepsilon < \frac{1}{2}$$

Comparing $\widetilde{p}$ with $1/2$ therefore decides whether $x\in L$. Thus a classical Turing machine can simulate the bounded-error quantum circuit using polynomial space, proving BQP $\subseteq$ PSPACE.

Consequences Quantum computers do not enlarge the class of computable functions: their computations can be simulated classically if no efficiency requirement is imposed. What is at stake is efficiency.

12.4 Extended Church–Turing thesis

While the original Church–Turing thesis is about computability, the *extended* (or *strong*) Church–Turing thesis is about efficient computability. It says that every "reasonable" or "physically realizable" model of computation can be simulated by a probabilistic classical machine with at most polynomial overhead. It is less known and less widely accepted than the original Church–Turing thesis.

Thanks to Shor's Algorithm 2 from Section 11, we know that one of the following statements must be false [19]:

1. The extended Church–Turing thesis is true.
2. FACTOR cannot be solved efficiently by a classical randomized algorithm.[7]
3. large-scale universal quantum computers can be built.

Any two of the three statements together imply that the remaining one is false.[8]

12.5 The Sword in the Stone: Quantum Analogues of NP

There is not a single unique quantum analogue of NP (Definition 12.1) [19]. The verifier can be quantum, and the certificate can be classical or quantum. Merlin–Arthur terminology is usually used here: Merlin sends a witness, and Arthur verifies it.

Strictly speaking, these classes are usually defined for *promise problems* [19]. This means that there are two disjoint sets L_{yes} and L_{no}, and the verifier only has to behave correctly on inputs promised to lie in $L_{\text{yes}} \cup L_{\text{no}}$. On inputs outside this promise set, no correctness requirement is imposed.

Definition 12.4 (QMA)**.** A promise problem $L = (L_{\text{yes}}, L_{\text{no}})$ is in QMA if there is a polynomial-size uniform quantum circuit family $\{V_n\}_{n\in\mathbb{N}}$ and polynomials p, q such that, on an input $x \in \{0,1\}^n$, the verifier receives x, a $p(n)$-Qbit witness $|\psi\rangle$, and $q(n)$ ancilla Qbits initialized to $|0\rangle^{\otimes q(n)}$. After applying V_n, one designated output Qbit is measured in the computational basis. The requirements are:

- Completeness (YES-case): if $x \in L_{\text{yes}}$, then there exists a quantum witness $|\psi\rangle$ that makes the verifier accept with probability at least 3/4
- Soundness (NO-case): if $x \in L_{\text{no}}$, then every quantum witness $|\psi\rangle$ makes the verifier accept with probability at most 1/4 ◀

The constants 3/4 and 1/4 are not important. As in BPP (Definition 12.2) and BQP (Definition 12.3), the precise constants are not part of the class definition. For QMA, this robustness follows from standard QMA error-reduction results. It is also enough to quantify over pure witnesses: a mixed witness is a convex combination of pure states, and the verifier's acceptance probability is linear in the density matrix. Therefore the best mixed witness cannot do better than some pure witness.

Definition 12.5 (QCMA)**.** QCMA is defined like QMA, except that Merlin's witness is an ordinary polynomial-size classical bit string (i.e. a classical "certificate"). ◀

MA is the classical Merlin–Arthur class: Merlin sends a classical witness, and Arthur checks it using a randomized polynomial-time classical verifier. MA fits between

$$\text{NP} \subseteq \text{MA} \subseteq \text{QCMA}$$

The first inclusion holds because an MA verifier is allowed but not forced to use randomness, so it can simply run the deterministic NP verifier. The second inclusion holds because a quantum verifier can simulate the classical randomized verifier and generate random bits (same story as in the proof of BPP $\subseteq$ BQP in Section 12.2).

[7] i.e., suitable decision versions of FACTOR are not in BPP

[8] This comparison needs some care: the first two claims are asymptotic complexity statements, while the third concerns finite physical circuits where one does not necessarily care about very large input sizes.

We have [19]

$$\mathrm{BQP} \subseteq \mathrm{QCMA} \subseteq \mathrm{QMA} \subseteq \mathrm{PSPACE}$$

The first inclusion uses an empty classical witness: Arthur ignores the witness register and simply runs the BQP verifier from Definition 12.3 on the input x (same story as for $\mathrm{P} \subseteq \mathrm{NP}$).

The second inclusion holds because every classical certificate w can be encoded as the computational-basis quantum state $|w\rangle$. The QMA verifier first measures the witness register in the computational basis, obtaining a classical string, and then runs the QCMA verifier on that string. Hence honest QCMA witnesses (satisfying the completeness condition in a YES case) are valid QMA witnesses after this encoding. If Merlin sends an arbitrary quantum state instead, the first measurement only produces a probability distribution over classical strings, so the acceptance probability is a convex combination of the QCMA acceptance probabilities.

The last inclusion is analogous in spirit to Section 12.3. For each fixed input x, the verifier circuit V_x defines an acceptance operator A_x acting on Merlin's witness register. For every pure witness $|\psi\rangle$, the acceptance probability has the form $p_{\mathrm{acc}}(\psi) = \langle\psi|A_x|\psi\rangle$. Hence the best possible witness is obtained by $\max_{|\psi\rangle} \langle\psi|A_x|\psi\rangle = \lambda_{\max}(A_x)$, where $\lambda_{\max}(A_x)$ is the largest eigenvalue of A_x. Although A_x is exponentially large, its entries can be computed on demand from the verifier circuit V_x, and $\lambda_{\max}(A_x)$ can therefore be approximated using polynomial space.

The quantum analogue of SAT is the *Local Hamiltonian* problem [19]. Recall from (3.12) and Section 5.3 that a Hamiltonian is a Hermitian operator whose eigenvalues are interpreted as energy levels. For a Hamiltonian H, write $\lambda_{\min}(H)$ for its smallest eigenvalue, i.e. its ground-state energy. The analogy is that a satisfying assignment is a classical witness for SAT, whereas a low-energy quantum state is a quantum witness for LH. Arthur's task is to check whether the supplied state has low expected energy $\langle\psi|H|\psi\rangle$ with respect to H.

Problem 12.6 (k-LH)**.** Given a Hamiltonian $H = \sum_j H_j$ on n Qbits, where each term H_j acts nontrivially on at most k Qbits, and two thresholds $a < b$ with $b - a \geq 1/\operatorname{poly}(n)$. Here the number of terms is polynomial in n, each H_j is specified by polynomially many bits, and the operator norms are bounded by $\operatorname{poly}(n)$. Decide whether

$$\lambda_{\min}(H) \leq a \qquad \text{or} \qquad \lambda_{\min}(H) \geq b$$

promised that one of the two cases holds. ◀

Kitaev's quantum Cook–Levin theorem says that k-LH is QMA-complete for $k \geq 5$. QMA-complete means that the problem is in QMA, and every problem in QMA can be reduced to it by a polynomial-time classical reduction. Thus LH plays for QMA a similare role as SAT plays for NP [19].

12.6 How optimistic should we be?

Section 10 gives a provable query-complexity advantage, while Section 11 gives polynomial-time quantum algorithms for important arithmetic problems [18]. These results suggest that quantum computation can be more efficient than classical randomized computation for some natural problems.

However, they are still not unconditional complexity-class separations: the Shor-type algorithms from Section 11 put suitable decision versions of factoring and discrete logarithms in BQP, but they do not prove $\mathrm{P} \subsetneq \mathrm{BQP}$ or $\mathrm{BPP} \subsetneq \mathrm{BQP}$.

The inclusions in Fact 12.2 explain why such a proof would be a major breakthrough:

- $\mathrm{P} \subsetneq \mathrm{BQP} \implies \mathrm{P} \neq \mathrm{PSPACE}$
- $\mathrm{BPP} \subsetneq \mathrm{BQP} \implies \mathrm{BPP} \neq \mathrm{PSPACE}$

There are different ways to look at this situation [18].

Optimistic. Quantum algorithms might provide new structural ideas for proving separations such as $\mathrm{P} \subsetneq \mathrm{PSPACE}$, or a natural separation such as $\mathrm{BPP} \subsetneq \mathrm{BQP}$.

Pessimistic. Because such separations imply major open lower bounds in classical complexity theory, proving that quantum computers can efficiently solve natural problems that are intractable for randomized classical polynomial-time algorithms might be very difficult.

More pessimistic. If $\mathrm{P} = \mathrm{PSPACE}$, then $\mathrm{BQP} = \mathrm{P}$ as decision-problem classes, so polynomial-time quantum computers would offer no asymptotic decision-problem advantage over deterministic classical computers. However, this equality is considered very unlikely.