

**Quantum computation and information:
Notes for Fall 2018 TAMU class**

J.M. Landsberg

Contents

Chapter 1. Classical and probabilistic computing	1
§1.1. 2025	1
§1.2. Surprising algorithms	2
§1.3. Notation, probability and linear algebra	6
§1.4. Classical Complexity	9
§1.5. Probabilistic computing	12
§1.6. Computation via linear algebra	13
Chapter 2. Quantum mechanics for quantum computing	17
§2.1. Quantum mechanics via probability	17
§2.2. Postulates of quantum mechanics and relevant linear algebra	21
§2.3. Super-dense coding	26
§2.4. Quantum teleportation	27
§2.5. Bell's game	28
Chapter 3. Algorithms	31
§3.1. Primality testing	31
§3.2. Grover's search algorithm	35
§3.3. Simons' algorithm	36
§3.4. Quantum gate sets	38
§3.5. Shor's algorithm	41
§3.6. A unified perspective on quantum algorithms: the hidden subgroup problem	51
§3.7. What is a quantum computer?	51

§3.8. Appendix: review of basic information on groups and rings	52
Chapter 4. Classical information theory	55
§4.1. Data compression: noiseless channels	55
§4.2. Entropy, i.e., uncertainty	58
§4.3. Shannon's noiseless channel theorem	60
§4.4. Transmission over noisy channels	61
Chapter 5. Quantum information	67
§5.1. Reformulation of quantum mechanics	67
§5.2. Distances between classical and quantum probability distributions	74
§5.3. The quantum noiseless channel theorem	76
§5.4. Properties of von Neumann entropy	80
§5.5. Entanglement and LOCC	88
Chapter 6. Representation theory and Quantum information	97
§6.1. Representation theory	97
§6.2. Projections onto isotypic subspaces of $\mathcal{H}^{\otimes d}$	102
Hints and Answers to Selected Exercises	111
Bibliography	113
Index	117

Classical and probabilistic computing

1.1. 2025

In January 2016, and in more detail in October, the NSA released a document warning the world that current encryption algorithms will be no longer secure as early as 2025¹. At that point in time there may be operational quantum computers. What is all the fuss about?

The main way banks, governments, etc. communicate securely now is using the RSA cryptosystem. RSA relies on the assumption that it is difficult to factor a large number N into its prime factors. In 1994 [Sho94] (also see [Sho97]) P. Shor described an algorithm to factor numbers quickly on a “quantum computer”.

Why can’t we factor numbers quickly already? What is a quantum computer?

Before addressing these questions, we need to address more basic ones:

What computations *can* we do quickly on a computer? What is a classical computer and what can it do?

¹See <http://www.math.tamu.edu/~jml/CNSA-Suite-and-Quantum-Computing-FAQ.pdf>

1.2. Surprising algorithms

1.2.1. Logarithms: fast multiplication of numbers. Until the 1600’s, when people had to do astronomical predictions (the king was very interested in knowing his horoscope, see [Lyo09]), a difficult step was the multiplication of large numbers. In 1614 John Napier revolutionized computation by writing a book of lists of numbers to implement a transform that swaps multiplication for addition: the logarithm. Kepler used Naiper’s book to make astronomical tables on the order of 30 times more accurate of previous tables [Gle11, p87].

Even in this example, there is something modern to learn: if the difficult step of a calculation (in this case taking logs and exponentiation) can be precomputed and stored in a database, it becomes essentially “free”.

1.2.2. The DFT: Fast multiplication of polynomials. Say $a(x), b(x)$ are polynomials of degree at most d with complex coefficients. Write $a(x) = \sum_{i=0}^d a_i x^i$, $b(x) = \sum_{j=0}^d b_j x^j$. Write $\bar{a} = (a_0, \dots, a_d)$ and similarly for other coefficients. Writing $a(x)b(x) = \sum_{k=0}^{2d} c_k x^k$, one has

$$(1.2.1) \quad c_k = \sum_{i+j=k} a_i b_j.$$

(One says \bar{c} is the *convolution* of \bar{a} and \bar{b} .) To obtain the coefficient vector \bar{c} by this standard method, one needs to perform on the order of d^2 arithmetic operations (i.e., +’s and *’s). In this situation, we will write $O(d^2)$ arithmetic operations, see 1.3.1 for the precise definition of $O(d^2)$.

Quantum algorithms will be expressed as a sequence of matrix vector multiplications, and we may do so here as well to facilitate comparisons.

To express this calculation in terms of matrix-vector multiplication, note that the vector \bar{c} is the product

$$\begin{pmatrix} a_0 & 0 & \dots & 0 \\ a_1 & a_0 & 0 & \dots & 0 \\ a_2 & a_1 & a_0 & \ddots & \\ \vdots & & \vdots & & \\ a_d & a_{d-1} & & \dots & a_0 \\ 0 & a_d & a_{d-1} & \dots & 0 \\ \vdots & & \ddots & & \\ 0 & \dots & & 0 & a_d \end{pmatrix} \begin{pmatrix} b_0 \\ b_1 \\ \vdots \\ b_d \end{pmatrix}.$$

Here we have broken the symmetry between $a(x)$ and $b(x)$. The symmetry will be restored momentarily.

Now we explain a trick to reduce the amount of computation. Pay attention as a variant of this trick will be critical to Shor's quantum algorithm for factoring. As with the multiplication of numbers, the key will be to do a transformation that re-organizes the input data of the two polynomials.

Since $\deg(ab) \leq 2d$, instead of working in the space of all polynomials, we can work in the ring $\mathbb{C}[x]/(x^N - 1)$ of polynomials quotiented by the ideal generated by the polynomial $x^N - 1$ for any $N > 2d$. For the moment, to fix ideas set $N = 2d + 1$, but later we will take N to be a power of two. We can then write a $(2d + 1) \times (2d + 1)$ matrix for $a(x)$ (allowing it now to have larger degree) as

$$\begin{pmatrix} a_0 & a_{2d} & a_{2d-1} & \cdots & a_2 & a_1 \\ a_1 & a_0 & a_{2d} & \cdots & a_3 & a_2 \\ a_2 & a_1 & a_0 & \ddots & & \\ \vdots & & \vdots & & & \\ a_d & a_{d-1} & & \cdots & a_{d+2} & a_{d+1} \\ a_{d+1} & a_d & a_{d-1} & \cdots & a_{d+3} & a_{d+2} \\ \vdots & & \ddots & & & \\ a_{2d} & \cdots & & & a_1 & a_0 \end{pmatrix}$$

and similarly for $b(x)$ (although we only need the first column of the product).

Note that the first $d + 1$ columns of this matrix is our old matrix. This looks like we are making our problem more complicated. However, now that we have a square matrix we can diagonalize it. At first glance, this seems like a very bad idea: the cost of a change of basis is worse than $O(d^2)$. However, one can use the same change of basis matrix for all polynomials. How could you know this? Because $a(x)b(x) = b(x)a(x)$, in both the usual multiplication and as elements of the $\mathbb{C}[x]/(x^N - 1)$, and if commuting matrices are diagonalizable, they are simultaneously diagonalizable.

Exercise 1.2.1: Show that if two diagonalizable matrices commute, then they are simultaneously diagonalizable.

Diagonalizing the matrix for $b(x)$, we can then perform the matrix product using $2d$ multiplications, instead of $O(d^2)$.

Here we can just construct a linear map that sends the coefficient vector of a polynomial of degree at most N to the vector consisting of eigenvalues of the corresponding $N \times N$ matrix as above. Let $DFT_N : \mathbb{C}^N \rightarrow \mathbb{C}^N$ denote this linear map. (DFT stands for *discrete Fourier transform*.) Write $\hat{a} = DFT_N \bar{a}$ (where we have padded the coefficient vector of $a(x)$ with zeros to make it have length N), and similarly $\hat{b} = DFT_N \bar{b}$. Given

\hat{a} and \hat{b} , the vector \hat{c} can be computed using N scalar multiplications as $\hat{c}_k = \hat{a}_k \hat{b}_k$. Finally $\bar{c} = DFT_N^{-1} \hat{c}$.

Although we viewed \bar{a} as a matrix in our derivation of the algorithm, when we implement the algorithm we will treat it as a column vector.

Exercise 1.2.2: Show that the matrix DFT_N (independent of $a(x)$) is given by $(DFT_N)_{jk} = (e^{\frac{2\pi i}{N}})^{jk}$ and its inverse is given by $(DFT_N^{-1})_{jk} = \frac{1}{N} (e^{\frac{2\pi i}{N}})^{-jk}$. Here use index ranges $0 \leq j, k \leq N - 1$. (Note that these are the roots of the equation $x^N - 1 = 0$.)

However, to multiply by DFT_N and its inverse, we need to perform six matrix multiplications of size N matrices, so the cost is still $O(N^2) \geq O(d^2)$, so we have not improved anything yet.

Now we come to a great discovery of Gauss in 1810 [**Gau**], rediscovered by several people, including Cooley-Tukey in 1965 [**CT65**], who are responsible for its modern implementation the revolutionized signal processing: the DFT matrix factors as a product of sparse matrices. Explicitly, if $N = 2^k$, DFT_N may be written as a product of k matrices, each with only $2N$ nonzero entries. The cost of matrix-vector multiplication of a sparse matrix with S nonzero entries is $O(S)$, so the cost of performing our DFT is $O(\log_2(N)N)$ instead of $O(N^2)$. Performing three such, plus the diagonal matrix multiplication does not change the order of this total cost.

Explicitly,

$$(1.2.2) \quad DFT_{2M} = \begin{pmatrix} DFT_M & \Delta_M DFT_M \\ DFT_M & -\Delta_M DFT_M \end{pmatrix} \Pi$$

where, setting $\omega = e^{\frac{2\pi i}{2M}}$, $\Delta_M = \text{diag}(1, \omega, \omega^2, \dots, \omega^{M-1})$ and Π is a permutation matrix corresponding to the inverse of the shuffle permutation $(1, \dots, 2M) \mapsto (1, 3, 5, \dots, 2M - 1, 2, 4, 6, \dots, 2M)$.

Exercise 1.2.3: Write DFT_4 as a product of two matrices, each with eight nonzero entries. Write DFT_8 as a product of three matrices, each with 16 nonzero entries. ☺

Exercise 1.2.4: Show that DFT_{2^k} may be factored as a product $S_1 \cdots S_k$ where each S_k has $2(2^k) \ll (2^k)^2$, for a total of $2^{k+1}k$ nonzero entries, and thus multiplication of two polynomials of degree at most $d = 2^{k-1}$ may be computed using $O(k2^{k+1}) = O(\log(d)d)$ arithmetic operations. ☺

Exercise 1.2.5: Verify Equation (1.2.2).

Remark 1.2.6. For those familiar with representation theory, the DFT is the change of basis matrix from the standard basis of the regular functions on \mathbb{Z}_N , denoted $\mathbb{C}[\mathbb{Z}_N]$, to the character basis. For an abelian group, matrix multiplication in the character basis becomes scalar multiplication because




Figure 1.2.1. Graph representing product of three 8×8 matrices that gives DFT_8 . Vertices in each row represent indices from 1 to 8, and edge from i to j at level $k \in \{1, 2, 3\}$ means the (i, j) -th entry of the k -th matrix is nonzero.

all its irreducible representations are one-dimensional. We will see this point is central to quantum algorithms.

Remark 1.2.7. You may have seen Fourier transforms of periodic functions, where convolution in the original space corresponds to multiplication in the transform space. This is the analogous transform when the group is the circle, i.e., the space of functions $\mathbb{C}[S^1]$.

More explicitly, write the unit circle in \mathbb{R}^2 as $S^1 = \{(\cos(\theta), \sin(\theta)) \mid \theta \in [0, 2\pi)\} \subset \mathbb{R}^2$. Introduce complex notation $\mathbb{C} = \mathbb{R}^2$, so $S^1 = \{e^{i\theta} \mid \theta \in [0, 2\pi)\} \subset \mathbb{C}$. Then for (e.g., continuous) functions $f(\theta)$ on the unit circle, we may write

$$f(\theta) = \sum_{n=-\infty}^{\infty} c_n e^{\frac{in\theta}{2}}, \text{ where } c_n = \frac{1}{4\pi} \int_0^{2\pi} f(\theta) e^{-\frac{in\theta}{2}} d\theta.$$

Since nonzero complex numbers form a group under multiplication, and the product of elements of length one is of length one, S^1 is naturally a group. In signal processing, we need to digitize (e.g. sound waves), so we approximate a periodic function by sampling it at say N equally spaced points on the circle, e.g., the points $e^{\frac{k2\pi i}{N}}$, $0 \leq k \leq N-1$. Note that these points form a subgroup, in fact the cyclic group of order N , \mathbb{Z}_N , the same group as $\mathbb{C}[x]/(x^N - 1)$. Tracing through the calculation, the DFT really is the discretization of the Fourier transform on the circle, exactly what one needs in signal processing.

Aside 1.2.8. For those familiar with tensors and their ranks, the structure tensor of $\mathcal{A} = \mathbb{C}[x]/(x^N - 1)$ has minimal tensor rank N , and the DFT is a change of basis that rewrites the structure tensor $T_{\mathcal{A}} \in \mathcal{A}^* \otimes \mathcal{A}^* \otimes \mathcal{A}$ as a sum of rank one tensors.

Aside 1.2.9. One might wonder if there is an even more efficient way of computing the operation $\bar{a} \mapsto DFT_N \bar{a}$. This question, and a path to resolving it, were presented by L. Valiant in [Val77]. There is interesting algebraic geometry related to the question, see [KLPSMN09, GHIL16].

1.2.3. Matrix multiplication. Another surprising algorithm deals with matrix multiplication. The usual algorithm for multiplying two $n \times n$ matrices uses $O(n^3)$ arithmetic operations. Strassen [Str69] discovered an algorithm that uses $O(n^{2.81})$ arithmetic operations and it has been conjectured that as n grows, it becomes nearly as easy to multiply matrices as it is to add them, that is for any $\epsilon > 0$, one can multiply matrices using $O(n^{2+\epsilon})$ arithmetic operations.

1.3. Notation, probability and linear algebra

1.3.1. Big/Little O etc. notation. For functions f, g of a real variable (or integer) x :

$f(x) = O(g(x))$ if there exists a constant $C > 0$ and x_0 such that $|f(x)| \leq C|g(x)|$ for all $x \geq x_0$,

$f(x) = o(g(x))$ if $\lim_{x \rightarrow \infty} \frac{|f(x)|}{|g(x)|} = 0$,

$f(x) = \Omega(g(x))$ if there exists a constant $C > 0$ and x_0 such that $C|f(x)| \geq |g(x)|$ for all $x \geq x_0$,

$f(x) = \omega(g(x))$ if $\lim_{x \rightarrow \infty} \frac{|g(x)|}{|f(x)|} = 0$, and

$f(x) = \Theta(g(x))$ if $f(x) = O(g(x))$ and $f(x) = \Omega(g(x))$.

We write \ln for the natural logarithm and \log for \log_2 .

1.3.2. Probability. Let $\mathcal{X} = \{a_1, a_2, \dots\}$ be a countable set and let $p : \mathcal{X} \rightarrow [0, 1]$ be a function such that $\sum_j p(a_j) = 1$. Such p is called a *discrete probability distribution* on \mathcal{X} . A function $X : \mathcal{X} \rightarrow \mathbb{R}$ is called a *discrete random variable* and it defines a probability distribution with discrete support on \mathbb{R} by $p_X(z) = \sum_{j|X(a_j)=z} p(a_j)$ so $p_X(z) = 0$ if $z \notin X(\mathcal{X})$. Similarly, random variables X, Y define a probability distribution $p_{X,Y}(x, y)$ with discrete support on $\mathbb{R} \times \mathbb{R}$, and similarly n random variables define a probability distribution with discrete support on \mathbb{R}^n . If $f : \mathbb{R} \rightarrow \mathbb{R}$ is a function, then $f \circ X$ is also a random variable.

The *expectation* (or *average*) of a random variable on a countable set \mathcal{X} equipped with a probability distribution p is

$$(1.3.1) \quad E[X] = \sum_{a_j \in \mathcal{X}} X(a_j)p(a_j).$$

Random variables X, Y are said to be *independent* if $p_{X,Y}(x, y) = p_X(x)p_Y(y)$. They are *identically distributed* if they define the same probability distributions. We write X_1, \dots, X_n are *iid* if they are independent and identically distributed.

For example, if $\mathcal{X} = \{H, T\}$ are the possible outcomes of flipping a biased coin which lands heads (H) with probability p and tails with probability $1 - p$, and $X(H) = 1, X(T) = -1$, then $E[X] = 2p - 1$, which is zero if the coin is fair. We will often be concerned with repeating an experiment many times. A typical situation is to define random variables X_j where X_j is 1 if the outcome of the j -th toss is heads and $X_j = -1$ if the outcome of the j -th toss is tails. Then the X_j are iid.

Note that $E[X] \in [-\infty, \infty]$. The *law of large numbers* implies that the name “expectation” is reasonable, that is, if one makes repeated experiments (e.g., as with the coin flips above) and averages the outcomes, the averages limit towards the expectation.

More precisely, the *weak law of large numbers* states that for X, X_1, X_2, \dots independent identically distributed random variables, and for any $\epsilon > 0$,

$$(1.3.2) \quad \lim_{n \rightarrow \infty} \Pr \left(\left| \frac{X_1 + \dots + X_n}{n} - E[X] \right| \geq \epsilon \right) = 0,$$

and the *strong law of large numbers* states moreover that

$$(1.3.3) \quad \Pr \left(\lim_{n \rightarrow \infty} \frac{X_1 + \dots + X_n}{n} = E[X] \right) = 1.$$

Here and throughout $\Pr(Z)$ denotes the probability of the event Z occurring with respect to some understood distribution.

However, individual outcomes can be far from the expectation. A first measurement of how far one can expect to be from the expectation is the *variance*: The variance of X is

$$(1.3.4) \quad \text{var}(X) = E[(X - E[X])^2]$$

$$(1.3.5) \quad = E[X^2] - E[X]^2$$

Exercise 1.3.1: Verify that (1.3.5)=(1.3.4).

Often one deals with the square-root of the variance, called the *standard deviation*, $\sigma(X) = \sqrt{\text{var}(X)}$.

If P is a probability distribution on $\mathcal{X} \times \mathcal{X}'$, one defines the *marginals* by $P_{\mathcal{X}}(x) = \sum_{y \in \mathcal{X}'} P(x, y)$ and $P_{\mathcal{X}'}(y) = \sum_{x \in \mathcal{X}} P(x, y)$, which are probability distributions on $\mathcal{X}, \mathcal{X}'$ respectively.

Let x_j be iid random variables. The string $x_1 \dots x_n =: \bar{x}^n$ is iid. Say $\mathcal{X} = \{1, \dots, d\}$ with $\Pr(j) = p_j$. The probability of any given string occurring depends only on the number of 1's 2's etc.. in the string and not on their

order. A string with c_j j 's occurs with probability $p_1^{c_1} \cdots p_d^{c_d}$. (Note that $c_1 + \cdots + c_d = n$.) The number of strings with this probability is

$$\binom{n}{c_1, \dots, c_d} := \frac{n!}{c_1! \cdots c_d!}$$

and we will need to estimate this quantity.

1.3.3. Detour on estimating multinomial coefficients. Stirling's formula implies

$$(1.3.6) \quad \ln(n!) = n \ln(n) - n + O(\ln(n)),$$

$$(1.3.7) \quad \log(n!) = n \log(n) - \log(e)n + O(\log(n)).$$

It is often proved using a contour integral of the Gamma function (see, e.g., [Ahl78, §5.2.5]). To see why it is plausible, write $\ln(n!) = \ln(1) + \cdots + \ln(n)$. This quantity may be estimated by

$$\int_1^n \ln(x) dx = [x \ln(x) - x]_1^n = n \ln n - n + 1,$$

giving intuition to (1.3.6).

In particular, for $0 < \beta < 1$ such that $\beta n \in \mathbb{Z}$,

$$(1.3.8) \quad \log \binom{n}{\beta n} = \log \frac{n!}{(\beta n)!((1-\beta)n)!} \\ = n[-\beta \log(\beta) - (1-\beta) \log(1-\beta)] + O(\log(n))$$

Let $H(\beta) = -\beta \log(\beta) - (1-\beta) \log(1-\beta)$ and more generally, for $\bar{p} = (p_1, \dots, p_d)$, let $H(\bar{p}) = -\sum_{i=1}^d p_i \log(p_i)$, called the *Shannon entropy* of \bar{p} .

It will play a central role in information theory.

Exercise 1.3.2: Show that similarly, for the multinomial coefficient

$$\binom{n}{p_1 n, \dots, p_d n} = \frac{n!}{(p_1 n)! \cdots (p_d n)!},$$

where $p_1 + \cdots + p_d = 1$, we have

$$(1.3.9) \quad \log \binom{n}{p_1 n, \dots, p_d n} = nH(\bar{p}) + O(\log(n)).$$

1.3.4. Linear algebra. Terms such as vector space, linear map etc.. will be assumed. For a vector space V , over a field \mathbb{F} , recall the *dual space* $V^* := \{f : V \rightarrow \mathbb{F} \mid f \text{ is linear}\}$. If $V = \mathbb{F}^n$ is the space of column vectors, then V^* may be identified with the space of row vectors. If V is finite dimensional, there is a canonical isomorphism $V \rightarrow (V^*)^*$, so we may also think of V as the space of linear maps $V^* \rightarrow \mathbb{F}$. Following physics convention,

we will usually denote elements of V by $|v\rangle$ and elements of V^* by $\langle\alpha|$, and their pairing by $\langle\alpha|v\rangle$. In bases, if

$$|v\rangle = \begin{pmatrix} v_1 \\ \vdots \\ v_n \end{pmatrix},$$

we may write $\langle\alpha| = (\alpha^1 \cdots \alpha^n)$, and $\langle\alpha|v\rangle = \sum_j \alpha^j v_j$ is row-column matrix multiplication. Let $\text{End}(V)$ denote the space of linear maps $V \rightarrow V$.

Define the *tensor product* $V \otimes W$ of vector spaces V and W to be the space of bi-linear maps $V^* \times W^* \rightarrow \mathbb{F}$, and more generally for a collection of vector spaces V_1, \dots, V_m , $V_1 \otimes \cdots \otimes V_m$ is the space of m -linear maps $V_1^* \times \cdots \times V_m^* \rightarrow \mathbb{F}$. If we work in bases, and $\dim V_j = \mathbf{v}_j$, then $V_1 \otimes V_2$ is the space of $\mathbf{v}_1 \times \mathbf{v}_2$ -matrices and $V_1 \otimes V_2 \otimes V_3$ may be visualized as the space of $\mathbf{v}_1 \times \mathbf{v}_2 \times \mathbf{v}_3$ “three dimensional matrices”.

Given a linear map $f : V \rightarrow W$, we may define a second linear map $f^t : W^* \rightarrow V^*$, by, for $\beta \in W^*$, $f^t(\beta)(v) = \beta(f(v))$. This is the coordinate free definition of the transpose of a matrix. One may also define a bilinear map $W^* \times V \rightarrow \mathbb{C}$, by $(\beta, v) \mapsto \beta(f(v))$. And as indicated in the exercises, this extends to a linear map $W^* \otimes V \rightarrow \mathbb{C}$. Thus we may also think of $V \otimes W$ as the set of bilinear maps $V^* \times W^* \rightarrow \mathbb{C}$. Consider a 2×3 matrix and its roles respectively as a linear map $\mathbb{C}^3 \rightarrow \mathbb{C}^2$, a linear map $\mathbb{C}^{2*} \rightarrow \mathbb{C}^{3*}$ and a bilinear map $\mathbb{C}^{2*} \times \mathbb{C}^3 \rightarrow \mathbb{C}$:

$$\begin{pmatrix} a & b & c \\ d & e & f \end{pmatrix} \begin{pmatrix} x \\ y \\ z \end{pmatrix} = \begin{pmatrix} ax + by + cz \\ dx + ey + fz \end{pmatrix}, \quad (s \quad t) \begin{pmatrix} a & b & c \\ d & e & f \end{pmatrix} = \begin{pmatrix} sa + td \\ sb + te \\ sc + tf \end{pmatrix},$$

$$(s \quad t) \begin{pmatrix} a & b & c \\ d & e & f \end{pmatrix} \begin{pmatrix} x \\ y \\ z \end{pmatrix} = sax + tdx + sby + tey + scz + tfz.$$

1.4. Classical Complexity

Classical complexity works in binary: one deals with strings of 0’s and 1’s. The set $\{0, 1\}$ is called a *bit*: it can encode “one bit” of information.

1.4.1. Circuits. We will mostly deal with *circuits*: Boolean circuits for classical computation, Boolean circuits with access to randomness for probabilistic computation, and quantum circuits for quantum computation.


Let \mathbb{F}_2 denote the field with two elements $\{0, 1\}$. A *Boolean function* is a map $f : \mathbb{F}_2^n \rightarrow \mathbb{F}_2$, or more generally $\mathbb{F}_2^n \rightarrow \mathbb{F}_2^m$. We agree on some basic Boolean functions, whose complexity is designated as having unit cost, e.g., addition \oplus (also called *XOR*) where $a \oplus b$ is addition in \mathbb{F}_2 (i.e., $0 \oplus 0 =$

$1 \oplus 1 = 0$ and $0 \oplus 1 = 1 \oplus 0 = 1$), \neg (NOT) negation, which swaps 0 and 1, (OR) $a \vee b$ where $0 \vee 0 = 0$ and all other $a \vee b = 1$, (AND= multiplication in \mathbb{F}_2) $a \wedge b = ab$, where $1 \wedge 1 = 1$ and all other $a \wedge b = 0$. I will call such a collection a (*logic*) *gate set*.

A *Boolean circuit* is a representation of a Boolean function $f : \mathbb{F}_2^n \rightarrow \mathbb{F}_2^m$ as a directed graph with n input edges, vertices labeled by elements of some fixed gate set, with edges going in and out, and m output edges. The *size* of a circuit is the number of edges in it.

Call a gate set a *universal gate set* if any Boolean function can be computed with a circuit whose vertices are labeled with gate set elements.²

Figure 1.4.1 depicts a Boolean circuit for the addition of two two digit (in binary) numbers:



We began by saying factorization is not known to have an efficient algorithm. We can now make that precise: a classical algorithm for a task (such as factoring) is *efficient* if there exists a polynomial p , such that if the input (in the case of factoring, the number to be factored N expressed in binary) is of size $M = \log N$ (in the case of factoring, the expression of N in binary has at most M digits), then there exists a Boolean circuit of size $p(M)$ that accomplishes the task. We now rephrase this more formally:

1.4.2. $\mathbf{P}/poly$, \mathbf{P} and \mathbf{NP} . Fix a universal gate set G , a natural complexity measure for a Boolean function is then the minimal size of a Boolean circuit that computes it. Let $p(n)$ be a polynomial and let $F_n : \mathbb{F}_2^n \rightarrow \mathbb{F}_2^{p(n)}$ be a sequence of functions (F_n). We consider the the growth with n of the size of a circuit needed to compute F_n . The critical issue, according to complexity theorists, is whether or not this growth is bounded by a polynomial. If it grows like a polynomial we say the sequence (F_n) is in the class $\mathbf{P}/poly$ with respect to G .

Exercise 1.4.1: Show that membership in $\mathbf{P}/poly$ with respect to G is independent of the choice of the finite universal gate set G .

²In some of the literature a gate set is sometimes called a “basis” (despite being unrelated to bases of vector spaces) and a universal gate set is called a “complete basis”.

Thus we will just say that the sequence (F_n) is in the class $\mathbf{P}/poly$.

The famous complexity class \mathbf{P} is the standard model for feasible computations, and it is unfortunate that $\mathbf{P}/poly$, with a simple description, is not \mathbf{P} . It is strictly larger, but not too much larger, and it can be used as a substitute for \mathbf{P} , see [AB09, Chap. 6].

The class \mathbf{P} is usually defined in terms of a different model of computation, namely *Turing machines*. We will avoid defining them, and assume the reader has at least a passing familiarity with them. A function F is in \mathbf{P} if it is in $\mathbf{P}/poly$ and there exists a Turing Machine TM such that the circuits C_n computing F_n are constructed by TM in time $poly(n)$, see [KSV02, Thm. 2.3].

The famous class \mathbf{NP} essentially consists of problems whose proposed solutions can be verified quickly, i.e., in polynomial time. For example the traveling salesman problem, where if someone claims to find a route to visit 30 cities traveling less than 2000 miles, it is easy to verify the claim, but the only known way of finding such a route is essentially by a brute force search. Another problem in \mathbf{NP} is “SAT”: one is handed a Boolean circuit and wants to know if it ever outputs 1. (If it does, to convince you it does, someone just needs to hand you an input that works, and you can quickly check if it outputs 1.) SAT is \mathbf{NP} -complete, which means one could define \mathbf{NP} to be the collection of problems that can be *reduced* (in polynomial time) to SAT, see, e.g., [AB09, Chap. 2]. In other words, there is a polynomial time algorithm for SAT if and only if $\mathbf{P} = \mathbf{NP}$.

1.4.3. How does a (classical) computer work? One can build mechanical devices that implement the classical gates. In our computers, logic gates are made out of electrical circuits. Input is either a 5 volt impulse for 1 and no impulse for 0. For example, the NOT gate is realized by the following diagram **** from top to bottom, there is a voltage source, a connection to an output wire the input source, and a ground

the NAND gate is realized by the following diagram (p15 Suil book) A voltage source is connected

1.4.4. Reversible classical computation. We will see that the gates of a quantum circuit (other than the measurements) must be reversible. Long before quantum computing, researchers were concerned that the second law of thermodynamics would have the consequence that as computers got more powerful, they would generate too much heat (entropy) from erasing bits. One way out of this would be to have reversible computation, see [Lan61], so one could argue for it independent of quantum computation. Of the gates we saw, NOT is clearly reversible as $\neg\neg x = x$. At the cost of adding an extra bit, one can make addition and multiplication reversible. Consider the

following gate, called to *Toffoli gate* Tof :

$$(1.4.1) \quad |x, y, z\rangle \mapsto |x, y, z \oplus (x * y)\rangle = |x, y, z \oplus (x \wedge y)\rangle$$

Note that if we send in $|x, y, 0\rangle$ we obtain $x * y$ in the third slot (register) and if we send in $|x, 1, y\rangle$ we obtain $x \oplus y$.

Exercise 1.4.2: Show that $Tof \circ Tof = \text{Id}$, so Tof is indeed reversible.

The gate set $\{Tof, \neg\}$, is universal and reversible, so there is no loss in computing power restricting to reversible classical computation.

1.5. Probabilistic computing

We will develop quantum mechanics as a generalization of probability, and we will view quantum computing as a generalization of probabilistic computing.

We will want to see the improvement of quantum computing to classical computing, so we should understand the what can be computed efficiently on a computer with access to randomness. Quantum computing itself is probabilistic, so we will need to implement notions from probability. Rather than introduce both the quantum-ness and the probabilistic nature at the same time, it will be easier to digest them one at a time.

1.5.1. BPP. It might increase our computational power if we exploit randomness. (Assuming we have a method to generate random numbers - more on this later.) For example, if someone hands you a complicated expression for a polynomial, e.g., in terms of an (algebraic) circuit, it can be very difficult to determine if the polynomial is just the zero polynomial in disguise. If we test the polynomial at a point, and its evaluation is non-zero, then we know it is not the zero polynomial. If it does evaluate to zero, then we have no information. For a polynomial of degree d in one variable, it is sufficient to test $d + 1$ distinct points, but as the number of variables grows, the number of points one needs to check grows exponentially. However, if we are allowed to test at a random point and it evaluates to zero, then with high probability over finite fields and probability one over \mathbb{Z} , the polynomial is the zero polynomial. Over finite fields, we can make this probability as high as we want by testing on several random points.

These observations motivate the class **BPP** (short for “bounded-error probabilistic polynomial time”), where one works with a Turing machine with access to randomness, and instead of asking for a correct answer on any input in polynomial time, one asks for a correct answer with probability strictly greater than $\frac{1}{2}$ on any input in polynomial time. (So that if one runs the program enough times, one can get a correct answer on any input with probability as high as one wants.)

Another motivation for probabilistic computation is that physical computers sometimes make mistakes (e.g. short circuit, input misread), so in the real world we are never completely sure of our answers.

Remark 1.5.1. It is actually subtle to know if one has a random sequence of numbers (e.g., take the last digit of the temperature in binary or similar). For example, the first digit of the number 2^n is far from a random element of $\{1, \dots, 9\}$, see [Ad89, §16, Ex. 4]. It is a subtle problem to make a machine to generate random numbers for us. Fortunately, for most situations, *pseudo-random* numbers suffice, see [AB09, §9.2.3].

Aside 1.5.2. If we are given additional information about the polynomial, then under certain circumstances one can test if the polynomial is zero by testing a reasonable number of points. This subject PIT (polynomial identity testing) is an active area of research, see [AB09, §7.2.3]. For a geometric perspective see [Lan17, §7.7].

Probabilistic computation however *cannot* be made reversible on a classical computer, as we will see in §1.6.2.

1.6. Computation via linear algebra

(Following [AB09, Exercise 10.4])

1.6.1. Reversible classical computation. Say $f : \mathbb{F}_2^n \rightarrow \mathbb{F}_2^m$ can be computed by a reversible Boolean circuit C . We describe how to rephrase the computation as a sequence of restricted linear operations on a vector space containing \mathbb{R}^{2^n} in anticipation of what will come in quantum computation. Give \mathbb{R}^2 basis $|0\rangle, |1\rangle$, which induces the basis $|i\rangle \otimes |j\rangle$, $i, j \in \{0, 1\}$ of $(\mathbb{R}^2)^{\otimes 2}$ and $|I\rangle := |i_1\rangle \otimes \dots \otimes |i_N\rangle$ of $(\mathbb{R}^2)^{\otimes N}$, $i_\alpha \in \{0, 1\}$, $1 \leq \alpha \leq N$. E.g., if our bit string is 00101100, we represent it by the vector $|00101100\rangle \in \mathbb{R}^{2^8}$. The restrictions will be:

- (1) Each linear map must be invertible and take a vector representing a sequence of bits to a sequence of bits. Such matrices are *permutation matrices*.
- (2) In order to deal with finite gate sets, we will require that each linear map only alters a small number of entries. For simplicity we assume it alters at most three entries, i.e., it acts on at most \mathbb{R}^{2^3} and is the identity on all other factors in the tensor product.

Each map will imitate some Boolean gate. For example, say we want to effect the Toffoli gate,

$$|x, y, z\rangle \mapsto |x, y, z \oplus (x * y)\rangle = |x, y, z \oplus (x \wedge y)\rangle$$

and act as the identity on all other basis vectors (sometimes called *registers*). Here, if the Toffoli gate is to compute $x*y$ z will represent “workspace bits”: x, y will come from the input to the problem and z will be set to 0 in the input. In the basis $|000\rangle, |001\rangle, |010\rangle, |100\rangle, |011\rangle, |101\rangle, |110\rangle, |111\rangle$, of \mathbb{R}^8 , the matrix is

$$(1.6.1) \quad \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \end{pmatrix}.$$

Call this matrix the *Toffoli matrix*.

The negation gate \neg may be defined by the linear map $\mathbb{C}^2 \rightarrow \mathbb{C}^2$ given by the matrix

$$\sigma_x = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}.$$

Exercise 1.6.1: Write matrices for $(x, y, z) \mapsto (x, y, z \oplus (x \oplus y))$ and $(x, y, z) \mapsto (x, y, z \oplus (x \vee y))$.

1.6.2. Probabilistic computation via linear algebra. If on given input, a probabilistic computation outputs 0 with probability p and 1 with probability $1 - p$, we could encode this with the vector $p|0\rangle + (1 - p)|1\rangle$, and then obtain either 0 or 1 by flipping a biased coin that gives heads with probability p .

Say $f : \mathbb{F}_2^n \rightarrow \mathbb{F}_2$ can be computed correctly with probability at least $\frac{1}{2}$ by a Boolean circuit C that is allowed to access randomness. (In particular, we can compute f correctly with probability as close as we want to one by repeating the computation enough times.) If we want to represent this in terms of linear algebra, we have to introduce a matrix for the coin flip:

$$\begin{pmatrix} \frac{1}{2} & \frac{1}{2} \\ \frac{1}{2} & \frac{1}{2} \end{pmatrix}.$$

Here we see that probabilistic computation cannot be made reversible, as this matrix is not invertible. So for probabilistic computation via linear algebra we will require the following of our matrices:

- (1) Each linear map must take probability distributions to probability distributions. This implies the matrices are *stochastic*: the entries are non-negative and each column sums to 1.

- (2) In order to deal with finite gate sets, we will require that each linear map only alters a small number of entries. For simplicity we assume it alters at most three entries, i.e., it acts on at most \mathbb{R}^{2^3} and is the identity on all other factors in the tensor product.

Consider $\{0, 1\}^m \subset \mathbb{R}^{2^m}$. A probability distribution on $\{0, 1\}^m$ may be encoded as a vector in \mathbb{R}^{2^m} : Give \mathbb{R}^2 basis $|0\rangle, |1\rangle$ and $(\mathbb{R}^2)^{\otimes m} = \mathbb{R}^{2^m}$ basis $|I\rangle$ where $I \in \{0, 1\}^m$. If the probability distribution assigns probability p_I to $I \in \{0, 1\}^m$, assign to the distribution the vector $v = \sum_I p_I |I\rangle \in \mathbb{R}^{2^m}$.

We will work with $\mathbb{R}^{2^{n+s+r}}$ where r is the number of times we want to access a random choice and s is the number of gates a circuit computing f would need.

Exercise 1.6.2: In probabilistic algorithms we will need to choose an element uniformly at random from a set of M elements. How can we realize this choice with the above matrices?

A probabilistic computation, viewed this way, starts with $|x0^{r+s}\rangle$, where $x \in \mathbb{F}_2^n$ is the input. One then applies a sequence of admissible stochastic linear maps to it, and ends with a vector that encodes a probability distribution on $\{0, 1\}^{n+s+r}$. One then restricts this to $\{0, 1\}^{p(n)}$, that is, one takes the vector and throws away all but the first $p(n)$ entries. This vector encodes a probability sub-distribution, i.e., all coefficients are non-negative and they sum to a number between zero and one. One then renormalizes (dividing each entry by the sum of the entries) to obtain a vector encoding a probability distribution on $\{0, 1\}^{p(n)}$ and then outputs the answer according to this distribution. Note that even if our calculation was “feasible” (i.e., polynomial in n size circuit), to write out the original output vector that we truncate would be exponential in cost. A stronger variant of this phenomenon will occur with quantum computing, where the result will be obtained with a polynomial size calculation, but one does not have access to the vector created, even using an exponential amount of computation.

To further prepare for the analogy with quantum computation, define a probabilistic bit (a *pbit*) to be the set

$$\{p_0|0\rangle + p_1|1\rangle \mid p_j \in [0, 1] \text{ and } p_0 + p_1 = 1\} \subset \mathbb{R}^2.$$

Note that the set of pbits is a convex set, and the basis vectors are the extremal points of this convex set.

Exercise 1.6.3: Show that if we have two problems to solve, one in $(\mathbb{R}^2)^{\otimes m}$ and another in $(\mathbb{R}^2)^{\otimes n}$, and we want to solve them simultaneously via linear algebra, then we should work in $(\mathbb{R}^2)^{\otimes m} \otimes (\mathbb{R}^2)^{\otimes n} = (\mathbb{R}^2)^{\otimes n+m}$.

1.6.3. What is known. $\mathbf{P} \subseteq \mathbf{BPP} \subseteq \mathbf{P}/\mathbf{Poly} = \mathbf{BPP}/\mathbf{Poly}$.

The inclusion $\mathbf{BPP} \subseteq \mathbf{P}/Poly$ is Adelman's theorem [Ad178]. The key observation is that “off-line” computations are not counted in the complexity assessment. So one can create, for any given n , a library of “random” a 's to test on. For example, to correctly determine the primality of 32-bit numbers, it is enough to test $a = 2, 7$, and 61 .

Does randomness really help? At the moment, we don't know. See [AB09, Chap. 20] for a discussion.

1.6.4. BQP. We are not yet in a position to define it, but the class \mathbf{BQP} will be the quantum analog of \mathbf{BPP} , the problems that can be solved efficiently, with high probability, on a quantum computer. Pbits will be replaced by *qubits*, which are unit vectors in \mathbb{C}^2 subject to an equivalence relation. The matrices will be allowed to have complex entries, they will be required to be unitary instead of stochastic, and at the end of the computation, one will not have the resulting vector in hand, but the result of a projection operator applied to it. The probability of obtaining I_0 from $\sum_I z_I |I\rangle$ will be $|z_{I_0}|^2$. There is no analog of \mathbf{P} for a quantum computer as answers will always have a probability of being incorrect.

A subtlety about quantum gate sets is that the notion of a “universal quantum gate set” will have a different meaning, namely that one can approximate any unitary map arbitrarily closely by elements of the gate set, not that one can perform the map exactly.

1.6.5. The Church-Turing theses. The Church-Turing thesis (made explicitly by Church in [Chu36]) is:

Any algorithm can be realized by a Turing machine.

So far there has been no challenge to this - any computation that can be done on a quantum computer can in principle be done with a sufficiently large Turing machine.

The quantitative (sometimes called *strong*) Church-Turing thesis [VSD86] is:

Any algorithmic process can be simulated efficiently by a Turing machine

or

Any algorithmic process can be simulated efficiently by a probabilistic Turing machine

Shor's algorithm challenges this thesis. On the other hand, there are experts who think that factoring could be in \mathbf{P} , because unlike, say SAT or the traveling salesman problem, the problem is highly structured.

Quantum mechanics for quantum computing

This chapter covers basic quantum mechanics needed for quantum computing. I present quantum mechanics as a generalization of probability and quantum computing will be viewed as a generalization of probabilistic computing.

put somewhere - stochastic matrices will be replaced by completely positive and trace preserving operators.

2.1. Quantum mechanics via probability

2.1.1. A wish list. In §1.6 we saw that any $f_n : \{0,1\}^n \rightarrow \{0,1\}^{p(n)}$ that could be computed correctly with probability say at least $\frac{2}{3}$ on any $I \in \{0,1\}^n$ with a circuit of size s and r coin flips, could be computed with the same probability via a sequence of linear operators on $(\mathbb{R}^2)^{\otimes n+r+s}$. Each linear operator was stochastic, so it took probability distributions to probability distributions, and acted on at most three registers via the action of one of the gates from the gate set used to construct the circuit. To get the output, after performing the linear operations, one throws away all but the first $p(n)$ entries of the output vector. The resulting vector encodes a non-normalized probability distribution, i.e., is of the form $|v\rangle = \sum_{|I|=p(n)} q_I |I\rangle$ with $q_I \geq 0$ and $\sum q_I \leq 1$. One then renormalizes, dividing each coefficient by $\sum q_I$, to obtain a vector $\sum_{|I|=p(n)} p_I |I\rangle$ with $p_I \geq 0$ and $\sum p_I = 1$. Then the algorithm outputs I with probability p_I .

Here is a wish list for how one might want to improve upon this set-up:

- (1) Allow more general kinds of linear maps to get more computing power, while keeping the maps easy to compute.
- (2) Have reversible computation: we saw that classical computation can be made reversible, but the coin flip was not. This property is motivated by physics, where many physical theories require time reversibility.
- (3) Again motivated by physics, one would like to have a continuous evolution of the probability vector, more precisely, one would like the probability vector to depend on a continuous parameter t such that if $|\psi_{t_1}\rangle = X|\psi_{t_0}\rangle$, then there exist admissible matrices Y, Z such that $|\psi_{t_0+\frac{1}{2}t_1}\rangle = Y|\psi_{t_0}\rangle$ and $|\psi_{t_1}\rangle = Z|\psi_{t_0+\frac{1}{2}t_1}\rangle$ and $X = ZY$. A physicist would say “time evolution is described by a semi-group”.

Let’s start with wish (2). One way to make the coin flip reversible is, instead of making the probability distribution be determined by the sum of the coefficients, one could take the sum of the squares. If we do this, there is no harm in allowing the entries of the output vectors to become negative, and one could use

$$H := \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}$$

for the coin flip applied to $|0\rangle$. The matrix H is called the *Hadamard matrix* or *Hadamard gate* in the quantum computing literature. It could just as well be called the *quantum coin flip*. If we made this change, we would obtain our second wish, and moreover have many operations be “continuous”, because the set of matrices preserving the L_2 -norm of a real-valued vector is the *orthogonal group* $O(n) = \{A \in \text{Mat}_{n \times n} \mid AA^T = \text{Id}\}$. So for example, any rotation has a square root.

As an indication that generalized probability may be related to quantum mechanics, the interference patterns observed in the famous two slit experiments is manifested in generalized probability: We obtain a “random bit” by applying H to $|0\rangle$: $H|0\rangle = \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$. However, if we apply a second quantum coin flip to the vector, we lose the randomness as $H^2|0\rangle = |1\rangle$, which, as pointed out in [Aar13], could be interpreted as a manifestation of interference.

However our third property will not be completely satisfied, as the matrix

$$\begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}$$

which represents a reflection, does not have a square root in $O(2)$.

To have the third wish satisfied, we will allow ourselves vectors with *complex* entries. From now on, set $i = \sqrt{-1}$. For a complex number $z = x + iy$, let $\bar{z} = x - iy$ denote its complex conjugate and $|z|^2 = z\bar{z}$ the square of its norm.

****picture of sphere here****

So we go from pbits, $\{p|0\rangle + q|1\rangle \mid p, q \geq 0 \text{ and } p + q = 1\}$ to *qubits*

$$\{\alpha|0\rangle + \beta|1\rangle \mid \alpha, \beta \in \mathbb{C} \text{ and } |\alpha|^2 + |\beta|^2 = 1\}.$$

The set of pbits is given in Figure 2.1.2




Figure 2.1.1. set of bits are just two points




Figure 2.1.2. set of pbits is a face of the unit simplex in the ℓ_1 -norm, extremal points correspond to classical bits

The set of qubits, considered in terms of real parameters, looks at first like the 3-sphere S^3 in $\mathbb{R}^4 \simeq \mathbb{C}^2$. However, the probability distributions induced by $|\psi\rangle$ and $e^{i\theta}|\psi\rangle$ are the same so it is really S^3/S^1 (the Hopf fibration), i.e., the two-sphere S^2 . Physicists call this S^2 the “Bloch sphere”. Geometrically, it would be more natural (especially since we have already seen the need to re-normalize in probabilistic computation) to work with projective space $\mathbb{C}P^1 \simeq S^2$ as our space of qubits, instead of a subset of \mathbb{C}^2 . For $v = (v_1, \dots, v_n) \in \mathbb{C}^n$, write $|v|^2 = |v_1|^2 + \dots + |v_n|^2$. The norm induces a Hermitian inner product $\langle v|w \rangle := \bar{v}_1 w_1 + \dots + \bar{v}_n w_n$. Note the physicist convention (which I use in this book) is the reverse of the mathematician

one, where the product is conjugate linear in the first factor and linear in the second.

The set of stochastic matrices is now replaced by the set of matrices

$$\mathbf{U}(n) := \{A \in \text{Mat}_{n \times n}(\mathbb{C}) \mid |Av| = |v| \ \forall v \in \mathbb{C}^n\},$$

which is called the *unitary group*.

Exercise 2.1.1: Show that $\mathbf{U}(n) = \{A \in \text{Mat}_{n \times n}(\mathbb{C}) \mid \overline{A}^T A = \text{Id}\}$

Claim: $\mathbf{U}(n)$ satisfies the third wish on the list. More precisely:

Proposition 2.1.2. For all $A \in \mathbf{U}(n)$, there exists a matrix $B \in \mathbf{U}(n)$ satisfying $B^2 = A$.

The proof is given below. First let's examine wish 1: it is an open question! However we can at least see that our generalized probabilistic computation includes our old probabilistic computation by the following easy exercise:

Exercise 2.1.3: Show that quantum coin flip H , the not (\neg) matrix and the Toffoli matrix (1.6.1) are unitary.

Proof of Proposition 2.1.2. Let A be a unitary matrix and let $|v\rangle$ be an eigenvector for A with eigenvalue λ . Since $|v| = |Av| = |\lambda v| = |\lambda||v|$, we see $|\lambda| = 1$, i.e., $\lambda = e^{i\theta}$ for some $\theta \in \mathbb{R}$. Note that A must have a basis of eigenvectors, $|v_1\rangle, \dots, |v_n\rangle$, as otherwise, let $|w\rangle$ be a putative generalized eigenvector, i.e., $A|w\rangle = \lambda|w\rangle + |u\rangle$. Since $|\lambda| = 1$, $|Aw| \neq |w|$. (Alternatively, just observe $\begin{pmatrix} e^{i\theta} & 1 \\ 0 & e^{i\theta} \end{pmatrix} \notin U(2)$.)

So let $|v_1\rangle, \dots, |v_n\rangle$ be an eigen-basis of A where $|v_j\rangle$ has eigenvalue $e^{i\theta_j}$. Let $B : \mathbb{C}^n \rightarrow \mathbb{C}^n$ be the matrix with the property that $B|v_j\rangle = e^{i\frac{\theta_j}{2}}|v_j\rangle$. Then B preserves the lengths of the eigenvectors, and thus of all vectors since the eigenvectors form an orthogonal basis by Exercise 2.1.5 below, and is therefore unitary, and clearly satisfies $B^2 = A$. \square

Exercise 2.1.4: Show that if $A \in \mathbf{U}(n)$, then $\langle v|w\rangle = \langle Av|Aw\rangle$ for all $v, w \in \mathbb{C}^n$.

Exercise 2.1.5: Show that if A is unitary, eigenvectors corresponding to distinct eigenvalues are orthogonal. \odot

2.1.2. Quantum mechanics from probability via four properties.

Consider the following properties of classical probability:

- (1) The law of large numbers is satisfied: that is relative frequencies of outcomes of measurements tend to the same value (the probability)

when a measurement is performed on an ensemble of n systems prepared in the same way, in the limit as n goes to infinity.

- (2) Let d denote the number of real parameters required to specify a state, and let N denote the maximum number of states that can be reliably distinguished from one another in a single measurement. Then $d = N^c$ for some natural number c which is chosen to be minimal.
- (3) A system whose state is constrained to belong to an M dimensional subspace of an N dimensional space behaves like a system of dimension M .
- (4) A composite system consisting of subsystems A and B satisfies $N = N_A N_B$ and $d = d_A d_B$.
- (5) There exists a reversible transform on a system between any two extremal points of the convex set of states (in our situation, the coordinate vectors).

Hardy [Har01] proved that any theory satisfying the above properties of probability must be classical probability expressed as stochastic matrices. In condition 2, one obtains $c = 1$. He also showed that if one adds the requirement in condition 5 that any transform can be written as a product of transforms that are arbitrarily close to the identity transform, one obtains $d = N^2$ and the axioms of quantum mechanics. We will not go through Hardy's proof, but at least we will verify that the standard axioms of quantum mechanics are compatible with Hardy's generalized probability. Later Shack [Sch03] showed that the first axiom was implied by the other four in both cases.

2.2. Postulates of quantum mechanics and relevant linear algebra

Here are the standard postulates of quantum mechanics and relevant definitions from linear algebra.

2.2.1. Postulate 1: State space. The first postulate describes the space one works in:

P1. Associated to any isolated physical system is a Hilbert space \mathcal{H} , called the *state space*. The system is completely described at a given moment by a unit vector $|\psi\rangle \in \mathcal{H}$, called its *state vector*, which is well defined up to a phase $e^{i\theta}$ with $\theta \in \mathbb{R}$. Alternatively one may work in projective space $\mathbb{P}\mathcal{H}$.

Definition 2.2.1. A *Hilbert space* \mathcal{H} is a complex vector space endowed with a non-degenerate Hermitian inner-product, $h : \mathcal{H} \times \mathcal{H} \rightarrow \mathbb{C}$, where by

definition h is conjugate linear in the first factor and linear in the second, $h(|v\rangle, |w\rangle) = \overline{h(|w\rangle, |v\rangle)}$, and $h(|v\rangle, |v\rangle) > 0$ for all $|v\rangle \neq 0$. (This is the physicists' convention, mathematicians generally require linearity in the first factor and conjugate linearity in the second.)

The Hermitian inner-product h allows an identification of \mathcal{H} with \mathcal{H}^* by $|w\rangle \mapsto \langle w| := h(\cdot, |w\rangle)$. This identification will be used repeatedly. We write $h(|v\rangle, |w\rangle) = \langle w|v\rangle$ and $|v| = \sqrt{\langle v|v\rangle}$ for the *length* of $|v\rangle$.

If $\mathcal{H} = \mathbb{C}^n$ with its standard basis, where $|v\rangle = (v_1, \dots, v_n)^T$, the *standard Hermitian inner-product* on \mathbb{C}^n is $\langle w|v\rangle = \sum_{j=1}^n \overline{w_j} v_j$. We will always assume \mathbb{C}^n is equipped with its standard Hermitian inner-product.

Remark 2.2.2. In quantum mechanics in general one needs to deal with infinite dimensional Hilbert spaces, but fortunately this is not necessary in quantum computing and quantum information theory.

Remark 2.2.3. Note the first postulate is identical to what one gets with generalized probability.

2.2.2. Postulate 2: Evolution. The second postulate describes how a state vector evolves over time:

P2. The state of an isolated system evolves with time according to the *Schrödinger equation*

$$i\hbar \frac{d|\psi\rangle}{dt} = X|\psi\rangle$$

where \hbar is a constant (*Planck's constant*) and X is a fixed *Hermitian operator*, called the *Hamiltonian* of the system.

Explanations. To define Hermitian operators, first define the *adjoint* of an operator $X \in \text{End}(\mathcal{H})$, to be the operator $X^\dagger \in \text{End}(\mathcal{H})$ such that $h(|X^\dagger v\rangle, |w\rangle) = h(|v\rangle, |Xw\rangle)$, i.e., $\langle X^\dagger v|w\rangle = \langle v|Xw\rangle$. Call X *Hermitian* if $X = X^\dagger$. When $\mathcal{H} = \mathbb{C}^n$, so $\text{End}(\mathcal{H})$ is the space of $n \times n$ matrices, then $X^\dagger = \overline{X}^t$, where t denotes transpose.

Exercise 2.2.4: Show that the eigenvalues of a Hermitian matrix are real.

Remark 2.2.5. Physicists tend to use the letter H for the Hamiltonian, but since we already use H for the Hadamard matrix, I do not adopt this convention.

Relation to generalized probability Recall that in generalized probability theory, transformations are unitary. For a general Hilbert space, define the *Unitary group*

$$\mathbf{U}(\mathcal{H}) := \{U \in \text{End}(\mathcal{H}) \mid |Uv\rangle = |v\rangle \ \forall |v\rangle \in \mathcal{H}\}.$$

When $\mathcal{H} = \mathbb{C}^n$ we have $\mathbf{U}(\mathbb{C}^n) = \mathbf{U}(n)$.

How do unitary operators from generalized probability lead to Schrödinger's equation? Recall that in generalized probability we are allowed to break up our action of an element $U \in \mathbf{U}(\mathcal{H})$ into a product of elements of $\mathbf{U}(\mathcal{H})$. More precisely, for each $\epsilon > 0$, there exists $k = k(\epsilon, U)$, such that $U = U_1 \cdots U_k$ with each U_j a distance at most ϵ from the identity. Similarly, we may find a curve from the identity to U in $\mathbf{U}(\mathcal{H})$.

Now say we have a smooth curve $U(t) \subset \mathbf{U}(\mathcal{H})$ with $U(0) = \text{Id}$. Write $U'(0) = \frac{d}{dt}|_{t=0}U(t)$. Consider

$$\begin{aligned} 0 &= \frac{d}{dt}|_{t=0}\langle v|w \rangle \\ &= \frac{d}{dt}|_{t=0}\langle U(t)v|U(t)w \rangle \\ &= \langle U'(0)v|w \rangle + \langle v|U'(0)w \rangle. \end{aligned}$$

Remark 2.2.6. The trick of writing 0 as the derivative of a constant function is ubiquitous in differential geometry.

Thus $U'(0)$ behaves almost like a Hermitian operator, which instead satisfies $0 = \langle Xv|w \rangle - \langle v|Xw \rangle$.

Exercise 2.2.7: Show that $iU'(0)$ is Hermitian.

We are almost at Schrödinger's equation.

Let $\mathfrak{u}(\mathcal{H}) \subset \text{End}(\mathcal{H})$ be the set of endomorphisms of the form $U'(0)$ for some curve as above, in other words $\mathfrak{u}(\mathcal{H}) = T_{\text{Id}}\mathbf{U}(\mathcal{H})$, the tangent space to the unitary group at the identity. The vector space $\mathfrak{u}(\mathcal{H})$ is called the *Lie algebra* of $\mathbf{U}(\mathcal{H})$. Note that $\mathfrak{u}(\mathcal{H})$ is a *real* vector space, not a complex one, because complex conjugation is not a complex linear map.

Then $i\mathfrak{u}(\mathcal{H}) \subset \text{End}(\mathcal{H})$ is a subspace of Hermitian endomorphisms and since both spaces have (real) dimension n^2 , it equals the space of Hermitian endomorphisms.

Exercise 2.2.8: Write $\mathfrak{u}(n) = \mathfrak{u}(\mathbb{C}^n)$. Verify that both $\mathfrak{u}(n)$ and the space of Hermitian matrices have (real) dimension n^2 .

For $X \in \text{End}(\mathcal{H})$, write $X^k \in \text{End}(\mathcal{H})$ for $X \cdots X$ applied k times. Write $e^X := \sum_{k=0}^{\infty} \frac{1}{k!} X^k$. This sum converges to a fixed matrix, essentially for the same reason it does in the $\dim \mathcal{H} = 1$ case.

Exercise 2.2.9: Show that the sum indeed converges, assuming the scalar case. \odot

Proposition 2.2.10. *If X is Hermitian, then $e^{iX} \in \mathbf{U}(\mathcal{H})$.*

Exercise 2.2.11: Prove Proposition 2.2.10. \odot

Postulate 2 implies the system will evolve unitarily, by (assuming we start at $t = 0$), $|\psi_t\rangle = U(t)|\psi_0\rangle$, where

$$U(t) = e^{-\frac{itX}{\hbar}}.$$

We conclude Postulate 2 is indeed predicted by generalized probability.

2.2.3. Postulate 3: measurements. In our first two postulates we dealt with isolated systems. In reality, no system is isolated and the whole universe is modeled by one enormous Hilbert space. In practice, parts of the system are sufficiently isolated that they can be treated as isolated systems. However, they are occasionally acted upon by the outside world, and we need a way to describe this outside interference. For our purposes, the isolated systems will be the Hilbert space attached to the input in a quantum algorithm and the outside interference will be the measurement at the end. That is, after a sequence of unitary operations one obtains a vector $|\psi\rangle = \sum z_j|j\rangle$ and as in generalized probability:

P3. If $|\psi\rangle = \sum_j z_j|j\rangle$, and a measurement is taken, the output is j with probability $|z_j|^2$.

2.2.4. Postulate 4: composite systems. A typical situation in quantum mechanics and quantum computing is that there are two or more isolated systems, say $\mathcal{H}_A, \mathcal{H}_B$ that are brought together (i.e., allowed to interact with each other) to form a larger isolated system \mathcal{H}_{AB} . The larger system is called the *composite system*. In classical probability, the composite space is $\{0, 1\}^{N_A} \times \{0, 1\}^{N_B}$. We have already seen in our generalized probability, the correct composite space is $(\mathbb{C}^2)^{\otimes N_A} \otimes (\mathbb{C}^2)^{\otimes N_B} = (\mathbb{C}^2)^{\otimes N_A + N_B}$ (Exercise 1.6.3).

P4. The state of a composite system \mathcal{H}_{AB} is the tensor product of the state spaces of the component physical systems $\mathcal{H}_A, \mathcal{H}_B$: $\mathcal{H}_{AB} = \mathcal{H}_A \otimes \mathcal{H}_B$.

When dealing with composite systems, we will allow partial measurements whose outcomes are of the form $|I\rangle \otimes \phi$.

This tensor product structure gives rise to the notion of *entanglement*, which, in the next few sections, we will see accounts for phenomenon outside of our classical intuition.

Definition 2.2.12. A state $|\psi\rangle \in \mathcal{H}_1 \otimes \cdots \otimes \mathcal{H}_n$ is called *separable* if it corresponds to a rank one tensor, i.e., $|\psi\rangle = |v_1\rangle \otimes \cdots \otimes |v_n\rangle$ with each $|v_j\rangle \in \mathcal{H}_j$. Otherwise it is *entangled*.

2.2.5. Generalized probability compared to the postulates. So far, Hardy's generalized probability has been shown to be completely compatible with the postulates of quantum mechanics. What we have not yet seen,

is why $d = N^2$ in the Hardy set-up. We could do this now, but it will be much easier after we reformulate quantum mechanics in §5.1, so I wait until then to explain it. The reformulation will be a logically equivalent theory, but will be easier to work with, especially regarding information theoretic questions. At that time the relation between measurements and other admissible operations will become clearer as well. In particular, the different gates allowed in computation is hard to extract from the above postulates.

2.2.6. Further Exercises. For $X, Y \in \text{End}(\mathcal{H})$, let $[X, Y] := XY - YX \in \text{End}(\mathcal{H})$ denote their commutator.

Exercise 2.2.13: For $X, Y \in \mathfrak{u}(\mathcal{H})$ show that $[X, Y] \in \mathfrak{u}(\mathcal{H})$, showing that $\mathfrak{u}(\mathcal{H})$ is indeed an algebra with the multiplication given by the commutator.

Exercise 2.2.14: Show that if $\mathcal{H} = \mathbb{C}^n$, then $X^\dagger = \overline{X}^T$, where the T denotes transpose.

Exercise 2.2.15: Show that if $Y \in \text{End}(\mathcal{H})$ is arbitrary, then YY^\dagger and $Y^\dagger Y$ are Hermitian.

Exercise 2.2.16: Show that the eigenvalues of a Hermitian operator are real.

Exercise 2.2.17: Prove the *spectral decomposition theorem* for Hermitian operators: Hermitian operators are diagonalizable and the eigenspaces of a Hermitian operator M are orthogonal. In particular we may write $M = \sum_\lambda \lambda P_\lambda$ where λ are the eigenvalues of M and the P_λ are commuting projection operators: $P_\lambda P_\mu = P_\mu P_\lambda$ and $P_\lambda^2 = P_\lambda$. Hint: differentiate $U(t)v(t)$ where $v(t)$ is an eigenvalue of $U(t)$, and $U(0) = \text{Id}$.

Exercise 2.2.18: Show that

$$\mathbf{U}(\mathcal{H}) = \{U \in \text{End}(\mathcal{H}) \mid \langle Uv|Uw \rangle = \langle v|w \rangle \forall |v\rangle, |w\rangle \in \mathcal{H}\},$$

and that if $U \in \mathbf{U}(\mathcal{H})$, then $U^{-1} = U^\dagger$.

Exercise 2.2.19: Show that $\mathbf{U}(2)$ acts transitively on lines in \mathbb{C}^2 , i.e, given any nonzero $v, w \in \mathbb{C}^2$ there exists $U \in \mathbf{U}(2)$ such that $U|v\rangle = \lambda|w\rangle$ for some $\lambda \in \mathbb{C}^*$. Hint: it suffices to do the case $w = \begin{pmatrix} 1 \\ 0 \end{pmatrix}$.

A *reflection* in a hyperplane $\mathbb{C}^{n-1} \subset \mathbb{C}^n$ is the linear map that, writing $|v\rangle \in \mathbb{C}^n$ as $|v\rangle = |v_1\rangle + |v_2\rangle$ with $|v_1\rangle \in \mathbb{C}^{n-1}$ and $|v_2\rangle \perp \mathbb{C}^{n-1}$, sends $|v\rangle \mapsto |v_1\rangle - |v_2\rangle$.

Exercise 2.2.20: Show that $\mathbf{U}(n)$ contains the reflections.

Exercise 2.2.21: Show that the product of two reflections is a rotation. More precisely, show that if $|v\rangle, |w\rangle$ are vectors in \mathbb{C}^n , the composition of a reflection in the hyperplane perpendicular to $|v\rangle$, followed by a reflection in the hyperplane perpendicular to $|w\rangle$, is a rotation in the $|v\rangle, |w\rangle$ plane by an angle equal to twice the angle between $|v\rangle$ and $|w\rangle$ (and the identity elsewhere).

2.3. Super-dense coding

In this section, we show that with a shared entangled state one can transmit two bits of classical information by transmitting a vector in just one qubit, which has led to the term “super¹-dense coding”. Super-dense coding was introduced in [BW92].

Physicists describe their experiments in terms of two characters, Alice and Bob. We generally follow this convention.

Let $\mathcal{H} = \mathbb{C}^2 \otimes \mathbb{C}^2 = \mathcal{H}_A \otimes \mathcal{H}_B$, and let $|epr\rangle = \frac{|00\rangle + |11\rangle}{\sqrt{2}}$ (called the *EPR state* in the physics literature, named after Einstein-Podolsky-Rosen) Assume this state has been created, both Alice and Bob are aware of it, Alice is in possession of (i.e., can manipulate) the first qubit, and Bob the second. This all happens before the experiment begins. They are allowed to agree on a protocol in advance. Then they are separated, but have a “quantum channel” along which they can transmit qubits. (Such will be explained in ***)

Now say Alice wants to transmit a two classical bit message to Bob, i.e., one of 00, 01, 10, 11. She is allowed to act on her half of $|epr\rangle$ by unitary transformations and then send it to Bob. (Later we will establish a gate set she must choose from, but it will include the gates we need below.) They agree in advance that once Bob is in possession of it, he will act on the four-dimensional space $\mathcal{H}_A \otimes \mathcal{H}_B$ by the unitary operator that performs the following change of basis:

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

and then will measure.

¹Physicists use the word “super” in the same way American teenagers use the word “like”.

If Alice wants to send 00, she just does nothing as then when Bob measures he will get $|00\rangle$ with probability one. Similarly, if she wants to send 01, she acts by

$$\sigma_x := \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$$

so Bob will be in possession of the state $\frac{1}{\sqrt{2}}(|00\rangle - |11\rangle)$, so when he performs the change of basis and measures, he will get $|01\rangle$ with probability one.

Exercise 2.3.1: What are the other two matrices Alice should act by to transmit the other two-bit messages?

In summary, with preparation of an EPR state in advance, plus transmission of a single qubit, one can transmit two classical bits of information.

2.4. Quantum teleportation

A similar phenomenon is *quantum teleportation*, where again Alice and Bob share half of an EPR state. This time Alice is in possession of a qubit $|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$, and wants to “send” $|\psi\rangle$ to Bob. However Alice only has access to a classical channel that sends bits to Bob. Can she transmit $|\psi\rangle$ to Bob, and if so, how many classical bits does she need to transmit to do so? Write the state of the system as

$$\frac{1}{\sqrt{2}} [\alpha|0\rangle \otimes (|00\rangle + |11\rangle) + \beta|1\rangle \otimes (|00\rangle + |11\rangle)]$$

where Alice can operate on the first two qubits.

Exercise 2.4.1: Show that if Alice acts on the first two qubits by $\text{Id}_1 \otimes \sigma_x = \text{Id}_1 \otimes \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$ then $H \otimes \text{Id}_2 = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix} \otimes \text{Id}_2$ (where the subscripts on Id indicate which factor the identity acts). She obtains

$$\frac{1}{2} [|00\rangle \otimes (\alpha|0\rangle + \beta|1\rangle) + |01\rangle \otimes (\alpha|1\rangle + \beta|0\rangle) + |10\rangle \otimes (\alpha|0\rangle - \beta|1\rangle) + |11\rangle \otimes (\alpha|1\rangle - \beta|0\rangle)].$$

Notice that Bob’s coefficient of Alice’s $|00\rangle$ is the state ψ that is to be transmitted. Alice performs a measurement. If she has the good luck to obtain $|00\rangle$, then she knows Bob has $|\psi\rangle$ and she can tell him classically that he is in possession of $|\psi\rangle$. But say she obtains the state $|01\rangle$: the situation is still good, she knows Bob is in possession of a state such that, if he acts on it with $\sigma_x = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$, he will obtain the state $|\psi\rangle$, so she just needs to tell him classically to apply σ_x . Since they had communicated the algorithm in the past, all Alice really needs to tell Bob in the first case is the classical message 00 and in the second case the message 01.

Exercise 2.4.2: Write out the other two different actions Bob should take depending on the possible bit pairs Alice could send him.

In summary, a shared EPR pair plus sending two classical bits of information allows one to transmit one qubit.

Remark 2.4.3. The name “teleportation” is misleading because information is transmitted at a speed slower than the speed of light.

2.5. Bell’s game

The 1934 Einstein-Podolsky-Rosen paper [EPR35] challenged quantum mechanics with the following thought experiment that they believed implied instantaneous communication across distances, in violation of principles of relativity: Alice and Bob prepare $|epr\rangle = \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle)$, then travel far apart. Alice measures her bit. If she gets 0, then she can predict with certainty that Bob will get 0 in his measurement, even if his measurement is taken a second later and they are a light year apart. (The essential property of the state is that Alice’s measurement makes Bob’s state classical as well.)

Ironically, this thought experiment has been made into an actual experiment designed by Bell [Bel64] and realized. The modern interpretation is that there is no paradox because the system does not transmit information faster than the speed of light, but rather they are acting on information that has already been shared. What follows is a version from [CHSH69], adapted from the presentation in [AB09].

The experiment can be described in a game, where Alice and Bob are on the same team and Charlie is a referee. Charlie chooses $x, y \in \{0, 1\}$ at random and sends x to Alice and y to Bob. Based on this information, Alice and Bob, without communicating with each other, get to choose bits a, b and send them to Charlie. They win if $a \oplus b = x \wedge y$, i.e., either $(x, y) \neq (1, 1)$ and $a = b$ or $(x, y) = (1, 1)$ and $a \neq b$.

2.5.1. Classical version. Note that if Alice and Bob both always choose 0, they win with probability $\frac{3}{4}$.

Theorem 2.5.1. [Bel64] *Regardless of the classical or probabilistic strategy Alice and Bob use, they never win with probability greater than $\frac{3}{4}$.*

The idea of the proof is that one first reduces a probabilistic strategy to a classical one, because after repeated rounds of the game, one can just adopt the most frequent choice. Then there are only 2^4 possible strategies and each can be analyzed. See, e.g., [AB09, Thm 20.2] for more detail.

2.5.2. Quantum version. Alice and Bob prepare $|epr\rangle = \frac{|00\rangle + |11\rangle}{\sqrt{2}}$ in advance, and Alice takes the first qubit and Bob the second. When Alice gets

x from Charlie, if $x = 1$, she applies a rotation by $\frac{\pi}{8}$ to her qubit, and if $x = 0$ she does nothing. When Bob gets y from Charlie, he applies a rotation by $-\frac{\pi}{8}$ to his qubit if $y = 1$ and if $y = 0$ he does nothing. (The order these rotations are applied does not matter because the operators on $(\mathbb{C}^2)^{\otimes 2}$ commute.) Both of them measure their respective qubits (again, the order will not matter) and send the values obtained to Charlie.

Theorem 2.5.2. *With this strategy, Alice and Bob win with probability at least $\frac{4}{5}$.*

The idea behind the strategy is that when $(x, y) \neq (1, 1)$, the states of the two qubits will have an angle at most $\frac{\pi}{8}$ between them, but when $(x, y) = (1, 1)$, the angle will be $\frac{\pi}{4}$.

Proof. If $(x, y) = (0, 0)$, then they are measuring $|epr\rangle$, so the measurement either yields 0 for both or 1 for both and $0 \oplus 0 = 1 \oplus 1 = 0 = 0 \wedge 0$, so they always win in this case.

If $(x, y) = (1, 0)$, then they are measuring

$$\frac{1}{\sqrt{2}} \left(\cos\left(\frac{\pi}{8}\right)|00\rangle + \sin\left(\frac{\pi}{8}\right)|10\rangle - \sin\left(\frac{\pi}{8}\right)|01\rangle + \cos\left(\frac{\pi}{8}\right)|11\rangle \right),$$

and the outputs are equal with probability $(\frac{1}{2} + \frac{1}{2}) \cos^2\left(\frac{\pi}{8}\right) \geq \frac{17}{20}$, and similarly if $(x, y) = (0, 1)$.

If $(x, y) = (1, 1)$, then they are measuring

$$\begin{aligned} & \frac{1}{\sqrt{2}} \left[\cos\left(\frac{\pi}{8}\right) \left(\cos\left(-\frac{\pi}{8}\right)|00\rangle + \sin\left(-\frac{\pi}{8}\right)|01\rangle \right) + \cos\left(\frac{\pi}{8}\right) \left(-\sin\left(-\frac{\pi}{8}\right)|00\rangle + \cos\left(-\frac{\pi}{8}\right)|01\rangle \right) \right. \\ & \quad + \sin\left(\frac{\pi}{8}\right) \left(\cos\left(-\frac{\pi}{8}\right)|10\rangle + \sin\left(-\frac{\pi}{8}\right)|11\rangle \right) + \sin\left(\frac{\pi}{8}\right) \left(-\sin\left(-\frac{\pi}{8}\right)|10\rangle + \cos\left(-\frac{\pi}{8}\right)|11\rangle \right) \\ & \quad - \sin\left(\frac{\pi}{8}\right) \left(\cos\left(-\frac{\pi}{8}\right)|00\rangle + \sin\left(-\frac{\pi}{8}\right)|01\rangle \right) - \sin\left(\frac{\pi}{8}\right) \left(-\sin\left(-\frac{\pi}{8}\right)|00\rangle + \cos\left(-\frac{\pi}{8}\right)|01\rangle \right) \\ & \quad \left. + \cos\left(\frac{\pi}{8}\right) \left(\cos\left(-\frac{\pi}{8}\right)|10\rangle + \sin\left(-\frac{\pi}{8}\right)|11\rangle \right) + \cos\left(\frac{\pi}{8}\right) \left(-\sin\left(-\frac{\pi}{8}\right)|10\rangle + \cos\left(-\frac{\pi}{8}\right)|11\rangle \right) \right] \\ & = \frac{1}{2} [|00\rangle + |01\rangle + |10\rangle + |11\rangle], \end{aligned}$$

so they win with probability $\frac{1}{2}$, as all coefficients have the same norm.

In sum, the overall chance of winning is at least $\frac{1}{4}(1) + \frac{1}{4}\left(\frac{17}{20}\right) + \frac{1}{4}\left(\frac{17}{20}\right) + \frac{1}{4}\left(\frac{1}{2}\right) = \frac{4}{5}$. \square

Exercise 2.5.3: Show that this strategy can be improved. What is its limit?

©

Algorithms

This chapter covers the basics of quantum computing, and the standard quantum algorithms. We begin with a probabilistic algorithm, the Miller-Rabin primality test, as ideas from its proof appear in Shor's algorithm. We next present the algorithms of Grover and Simons. We then discuss admissible quantum gates, and then, after considerable preliminaries, present Shor's algorithm. For those not familiar with basic facts regarding groups and rings, I suggest starting with the Appendix §3.8.

3.1. Primality testing

Although the complexity of factoring a number is not known, testing if it is prime has been known to belong to **BPP** since 1980 thanks to the *Miller-Rabin test* [Rab80]. I present the proof because parts of the proof will be used for Shor's algorithm.

Let $\mathbb{Z}/N\mathbb{Z}$ denote the ring of integers mod N , write $m \bmod N$ for the equivalence class of m .

The *Chinese remainder theorem* asserts that, for primes p, q , there is a ring isomorphism $\mathbb{Z}/pq\mathbb{Z} \simeq \mathbb{Z}/p\mathbb{Z} \times \mathbb{Z}/q\mathbb{Z}$.

Exercise 3.1.1: Verify the map $m \bmod pq \mapsto (m \bmod p, m \bmod q)$ defines a ring isomorphism.

More generally if $N = p_1^{a_1} \cdots p_k^{a_k}$ with p_j distinct primes, there is a ring isomorphism $(\mathbb{Z}/N\mathbb{Z})^* = (\mathbb{Z}/p_1^{a_1}\mathbb{Z})^* \times \cdots \times (\mathbb{Z}/p_k^{a_k}\mathbb{Z})^*$. For a ring R , let R^* denote its invertible elements under multiplication, which form a group. We also have $(\mathbb{Z}/N\mathbb{Z})^* = (\mathbb{Z}/p_1^{a_1}\mathbb{Z})^* \times \cdots \times (\mathbb{Z}/p_k^{a_k}\mathbb{Z})^*$.

Here is a warm-up: an inconclusive test to see if N is prime. Recall that if p is prime, then the multiplicative group $(\mathbb{Z}/p\mathbb{Z})^*$ is a cyclic group of order $p - 1$. As a consequence, if $x \not\equiv 0 \pmod{p}$ then $x^{p-1} \equiv 1 = x^0 \pmod{p}$ (the little Fermat theorem). In other words, if we find x such that $x^{N-1} \not\equiv 1 \pmod{N}$, then we know N is composite.

Call the following probabilistic algorithm the *Fermat test*: Choose a uniformly at random from $\{2, \dots, N - 1\}$ and compute $a^{N-1} \pmod{N}$. It will be clear that for this and the algorithm that follows, the tests will always report that N is prime when it is prime, so say N is composite. Under what circumstances do we correctly determine compositeness with probability at least $\frac{1}{2}$? Consider the following two cases:

- (1) $\gcd(a, N) = d \neq 1$. (This occurs with low probability.) Then the test detects that N is composite as in this situation $a \equiv 0 \pmod{d}$, and hence $a^{N-1} \not\equiv 1 \pmod{N}$.
- (2) $\gcd(a, N) = 1$, so $a \in (\mathbb{Z}/N\mathbb{Z})^*$.

Since we must account for the worst case scenario, assume we are in the second case:

Lemma 3.1.2. *If there exists $a \in (\mathbb{Z}/N\mathbb{Z})^*$, such that $a^{N-1} \not\equiv 1 \pmod{N}$, then the Fermat test detects the compositeness of N with probability $\geq \frac{1}{2}$.*

Before giving the proof, introduce the following groups associated to an abelian group G : For any natural number m , consider the group homomorphism $\phi_m : G \rightarrow G$, $\phi_m(x) = x^m$, and let

$$(3.1.1) \quad G^{(m)} = \text{Image } \phi_m \text{ and } G_{(m)} = \ker \phi_m,$$

both of which are abelian groups. In this language, the a 's for which the Fermat test fails are those in $(\mathbb{Z}/N\mathbb{Z})_{(N-1)}^*$.

Proof. The hypothesis is that $(\mathbb{Z}/N\mathbb{Z})_{(N-1)}^* \neq (\mathbb{Z}/N\mathbb{Z})^*$. Since $N > 3$, the quotient $(\mathbb{Z}/N\mathbb{Z})^*/(\mathbb{Z}/N\mathbb{Z})_{(N-1)}^*$ has cardinality at least 2. Thus $a^{N-1} \not\equiv 1 \pmod{N}$ for at least half of the elements of $(\mathbb{Z}/N\mathbb{Z})^*$ and we conclude. \square

It is possible that $a^{N-1} \equiv 1 \pmod{N}$ for all $a \in (\mathbb{Z}/N\mathbb{Z})^*$. So we will need an additional test to apply when the Fermat test fails to get our desired algorithm.

Exercise 3.1.3: Show that $N = 561 = 3 * 11 * 17$ is such that $a^{N-1} \equiv 1 \pmod{N}$ for all $a \in (\mathbb{Z}/N\mathbb{Z})^*$.

The second test uses the following proposition:

Proposition 3.1.4. *If there exists a natural number b such that $b^2 \equiv 1 \pmod{N}$ and $b \not\equiv \pm 1 \pmod{N}$, then N is composite with nontrivial factors in common with both $b + 1$ and $b - 1$.*

Proof. In this case $b^2 - 1 = (b - 1)(b + 1)$ is a multiple of N but $b - 1, b + 1$ are not, so N must have nontrivial factors in common with both $b + 1$ and $b - 1$. \square

Here is the Miller-Rabin algorithm: to avoid trivialities, assume N is odd.

Choose $a \in \{2, \dots, N - 2\}$ uniformly at random.

Step 1: Test if $a^{N-1} \not\equiv 1 \pmod{N}$. If so, then N is composite by the Little Fermat theorem and one concludes. Otherwise go to step 2:

Step 2: Let 2^k be the largest power of 2 that divides $N - 1$ and write $N - 1 = 2^k l$. Compute the sequence $a^l, a^{2l}, a^{4l}, \dots, a^{2^{k-1}l}$, all mod N . If this sequence contains a 1 preceded by anything except ± 1 , i.e., if there exists j such that $a^{2^j l} \not\equiv \pm 1 \pmod{N}$ and $(a^{2^j l})^2 \equiv 1 \pmod{N}$, then N is composite by Proposition 3.1.4. Otherwise the algorithm replies “ N is prime”.

One can check that the total circuit size of this algorithm is $O(\log(N)^3)$. The only subtlety is that taking exponentially many powers of a would violate this size, but we are only taking powers mod N .

Exercise 3.1.5: Prove that for $k \in \{0, \dots, N - 1\}$, $a^k \pmod{N}$ can be computed by a reversible classical circuit of size $\text{poly}(\log(N))$. \odot

Proposition 3.1.6. *The Miller-Rabin algorithm succeeds on any input with probability at least $\frac{1}{2}$.*

Proof. It is clear that if N is prime, the algorithm always indicates that it is prime, so assume N is composite and odd. Start the algorithm, get some $a \in \{2, \dots, N - 2\}$ chosen uniformly at random. If $\gcd(a, N) > 1$, then step 1 shows that N is composite, so assume this does not happen, which implies a is uniformly distributed over $(\mathbb{Z}/N\mathbb{Z})^*$. (This last assertion holds because for any group homomorphism of finite groups $f : G \rightarrow H$, all fibers have the same cardinality. We will use this repeatedly in what follows.)

In order for step 1 to work with probability at least $\frac{1}{2}$, it is enough that there is one $x \in \{2, \dots, N - 1\}$ such that $x^{N-1} \not\equiv 1 \pmod{N}$.

Exercise 3.1.7: Show that if $N = p^c$ for some prime p , then taking $a = p^c + 1 - p^{c-1}$, then $a^{N-1} \not\equiv 1 \pmod{N}$ because $a^{N-1} \equiv p^{c-1} + 1 \pmod{N}$ and $p^{c-1} + 1 \not\equiv 1 \pmod{N}$.

By Exercise 3.1.7, we may assume $N = uv$, where u, v are odd, $u, v > 1$, and $\gcd(u, v) = 1$. By the Chinese remainder theorem $(\mathbb{Z}/N\mathbb{Z})^* \simeq (\mathbb{Z}/u\mathbb{Z})^* \times (\mathbb{Z}/v\mathbb{Z})^*$.

If a is uniformly distributed over an abelian group G then a^m is uniformly distributed over $G^{(m)}$ defined in (3.1.1).

Now $(\mathbb{Z}/N\mathbb{Z})^{*(m)} \simeq (\mathbb{Z}/u\mathbb{Z})^{*(m)} \times (\mathbb{Z}/v\mathbb{Z})^{*(m)}$. If either $(\mathbb{Z}/u\mathbb{Z})^{*(N-1)}$ or $(\mathbb{Z}/v\mathbb{Z})^{*(N-1)}$ is non-trivial, step 1 will detect compositeness with probability at least $\frac{1}{2}$, so assume both are trivial. To apply step 2, we need to consider the powers $a^{2^j l} \bmod N$ and show there exists j such that $a^{2^j l} \not\equiv \pm 1 \pmod N$ but $(a^{2^j l})^2 = a^{2^{j+1} l} \equiv 1 \pmod N$ with probability at least $\frac{1}{2}$.

Let j_0 be the largest value such that $(\mathbb{Z}/N\mathbb{Z})^{*(2^{j_0 l})} \neq \{1\}$ but $(\mathbb{Z}/N\mathbb{Z})^{*(2^{j_0+1} l)} = \{1\}$. Use the isomorphism $(\mathbb{Z}/N\mathbb{Z})^{*(2^{j_0 l})} \simeq (\mathbb{Z}/u\mathbb{Z})^{*(2^{j_0 l})} \times (\mathbb{Z}/v\mathbb{Z})^{*(2^{j_0 l})}$: both the factors cannot be trivial by assumption. If one of the two factors is trivial, since $-1 \mapsto (-1, -1)$ under the Chinese remainder theorem map, we could only fail if $a^{2^{j_0 l}}$ maps to $(1, 1)$, but this will happen for the nontrivial factor with probability at most $\frac{1}{2}$. Now assume both factors are nontrivial, say of cardinalities c_u, c_v . In this case, the image of $a^{2^{j_0 l}}$ in the first factor is 1 with probability $\frac{1}{c_u}$, and is 1 in the second factor with probability $\frac{1}{c_v}$, and these events are independent (again by the Chinese remainder theorem). Thus the probability $a^{2^{j_0 l}} \equiv 1 \pmod N$ is $\frac{1}{c_u c_v}$. For similar reasons the probability $a^{2^{j_0 l}} \equiv -1 \pmod N$ is either $\frac{1}{c_u c_v}$ or zero. Thus the probability of failure is at most $\frac{2}{c_u c_v} \leq \frac{1}{2}$. \square

To get an algorithm that works with probability greater than $\frac{1}{2}$, apply the test twice.

But this is not the end of the story:

Theorem 3.1.8. [AKS04] *Primality testing is in P.*

The core of the proof is a variant of the little Fermat theorem: Let a, N be relatively prime integers with $N > 2$, Then N is prime if and only if $(x + a)^N \equiv x^N + a \pmod N$.

Exercise 3.1.9: Prove the assertion. \odot

The bulk of the work is reducing the number of coefficients one needs to check in the expansion of the left hand side.

The lesson to be drawn here is that we should not make any assumptions regarding the difficulty of a problem until we have a proof.

3.2. Grover's search algorithm

The problem: given $F_n : \mathbb{F}_2^n \rightarrow \mathbb{F}_2$, computable by a $\text{poly}(n)$ -size classical circuit, find a such that $F_n(a) = 1$ if such a exists.

Grover found a quantum circuit of size $\text{poly}(n)2^{\frac{n}{2}}$ that solves this problem (with high probability). Compare this with a brute force search, which requires a circuit of size $\text{poly}(n)2^n$. No classical or probabilistic algorithm is known that does better than $\text{poly}(n)2^n$. Note that it also gives a size $\text{poly}(n)2^{\frac{n}{2}}$ probabilistic solution to the **NP**-complete problem SAT (it is stronger, as it not only determines existence of a solution, but finds it).

We will present the algorithm for the following simplified version where one is promised there exists exactly one solution. All essential ideas of the general case are here.

Problem: given $F_n : \mathbb{F}_2^n \rightarrow \mathbb{F}_2$, computable by a $\text{poly}(n)$ -size classical circuit, and the information that F has exactly one solution a , find a .

The idea of the algorithm is to start with a vector equidistant from all possible solutions, and then to incrementally rotate it towards a . What is strange for our classical intuition is that we will be able to rotate towards the solution without knowing what it is, and similarly, we won't "see" the rotation matrix either.

We work in $(\mathbb{C}^2)^{\otimes n+1+s}$ where $s = s(n)$ is the size of the classical circuit needed to compute F_n . We suppress reference to the s "workspace bits" in what follows.

The first step is to construct such a starting vector:

The following vector is the average of all the classical (observable) states:

$$(3.2.1) \quad |av\rangle := \frac{1}{2^{\frac{n}{2}}} \sum_{x \in \{0,1\}^n} |x\rangle.$$

To prepare $|av\rangle$, note that $H|0\rangle = \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$, so applying $H^{\otimes n}$ to $|0 \cdots 0\rangle$ transforms it to $|av\rangle$.

The cost of this is n gates, as $H^{\otimes n}$ is the composition of $H \otimes \text{Id}_{2, \dots, n}$, $\text{Id}_1 \otimes H \otimes \text{Id}_{3, \dots, n}$, \dots , $\text{Id}_{1, \dots, n-1} \otimes H$.

Since $|av\rangle$ is equidistant from all possible solution vectors, we have $\langle av|a\rangle = \frac{1}{2^{\frac{n}{2}}}$. We want to rotate $|av\rangle$ towards the unknown a . Recall that $\cos(\angle(|v\rangle, |w\rangle)) = \frac{\langle v|w\rangle}{|v||w|}$. Write the angle between av and a as $\frac{\pi}{2} - \theta$, so $\sin(\theta) = \frac{1}{2^{\frac{n}{2}}}$.

Recall from Exercise 2.2.21, that a rotation is a product of two reflections. In order to perform the rotation R that moves $|av\rangle$ towards $|a\rangle$, we

first reflect in the hyperplane orthogonal to $|a\rangle$, and then in the hyperplane orthogonal to $|av\rangle$.

Consider the map

$$(3.2.2) \quad |xy\rangle \mapsto |x(y \oplus F(x))\rangle$$

defined on basis vectors and extended linearly. To execute this, we use the s workspace bits corresponding to y to effect s reversible classical gates. We initially set $y = 0$ so that the image is $|x0\rangle$ for $x \neq a$, and $|x1\rangle$ when $x = a$. Next apply the quantum gate $\text{Id} \otimes \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}$ which sends $|x0\rangle \mapsto |x0\rangle$, and $|x1\rangle \mapsto -|x1\rangle$. Finally apply the map $|xy\rangle \mapsto |x(y \oplus F(x))\rangle$ again.

Thus $|a0\rangle \mapsto -|a0\rangle$ and all other basis vectors $|b0\rangle$ are mapped to themselves, which is what we desired.

Next we need to reflect around $|av\rangle$. It is easy to reflect around a classical state, so first perform the map $H^{\otimes n}$ that sends $|av\rangle$ to $|0 \cdots 0\rangle$ (recall that $H = H^{-1}$), then reflect in the hyperplane perpendicular to $|0 \cdots 0\rangle$ using the Boolean function $g : \mathbb{F}_2^n \rightarrow \mathbb{F}_2$ that outputs 1 if and only if its input is $(0, \dots, 0)$, in the role of F for our previous reflection, then apply Hadamard again so the resulting reflection is about $|av\rangle$.

The composition of these two reflections is the desired R .

Exercise 3.2.1: What is the probability that a measurement of $R|av\rangle$ will produce $|a\rangle$?

As mentioned above, the vector $R|av\rangle$ is not useful, but if we instead compose this map with itself $O(\frac{1}{\delta})$ times, we obtain a vector much closer to $|a\rangle$.

Exercise 3.2.2: Show that applying the procedure $2^{\frac{n}{2}}$ times, one obtains a vector such that the probability of it being in state $|a\rangle$ after a measurement is greater than $\frac{1}{2}$.

3.3. Simons' algorithm

The problem: given $F : \mathbb{F}_2^n \rightarrow \mathbb{F}_2^n$, computable by a Boolean circuit of size polynomial in n , such that there exists $a \in \mathbb{F}_2^n$ satisfying for all $x, y \in \mathbb{F}_2^n$, $F(x) = F(y)$ if and only if $x = y \oplus a$, find a . For simplicity of exposition, assume we know $a \neq (0, \dots, 0)$ as well.

Simons gives a $poly(n)$ size quantum circuit that obtains the solution.

Remark 3.3.1. Although this problem may look unnatural, the resulting algorithm inspired Shor's algorithm and its generalizations, and it fits into a larger framework of problems that allow for an exponential quantum speedup over known probabilistic algorithms.

Remark 3.3.2. This problem is expected to be hard on a classical computer. Consider the following variant where F is allowed to be difficult to compute, but we are handed a black box that will compute it for us at unit cost. If a and F are chosen at random subject to the condition that $F(x) = F(y)$ if and only if $x = y \oplus a$, then classically one would need to use the black box $2^{\frac{n}{2}}$ times before having any information at all, as with fewer calls, it is likely that one never gets the same answer twice. On the other hand, Simons' algorithm still gives a $\text{poly}(n)$ -size solution in this setting.

Work in $(\mathbb{C}^2)^{\otimes 2n+s}$, where s is the size of a reversible Boolean circuit needed to compute F . We suppress reference to the s workspace qubits in what follows. As with Grover's algorithm, we will construct a vector that "sees" the answer a , but we will not be able to see the vector, so instead we manipulate it to get information about the solution. Also as before, first prepare $|av\rangle = \frac{1}{2^{\frac{n}{2}}} \sum_{x \in \{0,1\}^n} |x\rangle$. Then apply the operation $|xz\rangle \mapsto |x(z \oplus F(x))\rangle$ to $|av, 0^n\rangle$, to obtain

$$\frac{1}{2^{\frac{n}{2}}} \sum_{x \in \{0,1\}^n} |x\rangle \otimes |F(x)\rangle$$

Now measure the second n bits of the register to put the second n qubits into some classical state z_0 :

$$\frac{1}{2^{\frac{n}{2}}} \sum_{\{x|F(x)=z_0\}} |x\rangle \otimes |z_0\rangle.$$

Say $F(x_0) = z_0$, then (assuming $a \neq 0^n$) our sum collapses to

$$\frac{1}{2^{\frac{n}{2}}} (|x_0\rangle + |x_0 \oplus a\rangle) \otimes |z_0\rangle.$$

We want to manipulate this vector to gain information about a .

For $x, y \in \mathbb{F}_2^n$, let $x \cdot y := \bigoplus_{j=1}^n x_j y_j \in \mathbb{F}_2$ denote their inner product. Now perform the Hadamard operation on the first n bits again.

Exercise 3.3.3: Show that $H^{\otimes n}|x\rangle = \frac{1}{2^{\frac{n}{2}}} \sum_y (-1)^{x \cdot y} |y\rangle$.

We obtain

$$(3.3.1) \quad \sum_{y \in \mathbb{F}_2^n} \left((-1)^{x_0 \cdot y} + (-1)^{(x_0 \oplus a) \cdot y} \right) |y f(x_0)\rangle$$

Since $(x \oplus a) \cdot y = x \cdot y \oplus a \cdot y$, the term in the summand for any given y is non-zero if and only if $a \cdot y = 0$.

Thus when we measure the vector, we obtain a $y \in \mathbb{F}_2^n$ that is orthogonal to a with respect to the inner-product on \mathbb{F}_2^n (and is chosen uniformly at random among such). So we may restrict our search for a to the hyperplane in \mathbb{F}_2^n perpendicular to y . If we continue, with good luck, we could find

a after $n - 1$ runs of the algorithm. However, we have no guarantee we do not end up with linearly dependent y 's. If we run the algorithm more than $2n$ times, then there will be $n - 1$ independent hyperplanes with high probability.

Exercise 3.3.4: Prove that for any $a \in \mathbb{F}_2^n$, if vectors y_1, \dots, y_{2n} are chosen uniformly at random subject to $a \cdot y_j = 0$ for every $j \in [n - 1]$, then with probability at least $\frac{9}{10}$ a subset of $n - 1$ of them are linearly independent.

3.4. Quantum gate sets

Previously we only need the quantum gates corresponding to the Toffoli gate, the swap, and H . For general computation one needs more general gates. In this section we discuss admissible quantum gate sets. There are two issues we need to deal with: locality and approximation. We will first see, in §3.4.1 that locality is not a problem. Approximation is more subtle. We will first need to a way to measure how close an approximation is to the desired gate set. This will be done via the *operator norm* defined in §3.4.2. I introduce the “standard quantum gate set” in §3.4.4. Before that, in §3.4.3, I briefly discuss “controlled gates” which encode classical quantum interaction.

3.4.1. Locality. We would like to execute arbitrary unitary operations but we expect to only be able implement local quantum gates, as it is likely that one can only create *entanglement* in a laboratory on qubits that are physically close to one another. Thanks to the following unitary version of the classical Cartan-Dieudonné theorem, this issue is not a problem:

Lemma 3.4.1. *Any $U \in \mathbf{U}(n)$ may be written as a product of at most $\binom{n}{2}$ elements, each of which acts on some $\mathbb{C}\{e_i, e_{i+1}\}$ as an element of $U(2)$ and is the identity on the span of $e_1, \dots, e_{i-1}, e_{i+2}, \dots, e_n$, where e_1, \dots, e_n is the standard basis of \mathbb{C}^n .*

Proof. Let $\mathbf{U}(2)_i \subset \mathbf{U}(n)$ be the copy of $\mathbf{U}(2)$ acting only on $\mathbb{C}\{e_i, e_{i+1}\}$. By Exercise 2.2.19, $\mathbf{U}(2)$ acts transitively on lines in \mathbb{C}^2 . Moreover, it can send any $|v\rangle \in \mathbb{C}^2$ to $\begin{pmatrix} |v| \\ 0 \end{pmatrix}$. Thus for any unit vector $|\psi\rangle \in \mathbb{C}^n$, there exist $U_j \in \mathbf{U}(2)_j$ such that $U_1 \cdots U_n |\psi\rangle = e_1$. Write the columns of U^{-1} as $|\psi_1\rangle, \dots, |\psi_n\rangle$ where $|\psi_j| = 1$. We may find $U_{1,1}, \dots, U_{1,n-1}$ such that $U_{1,1} \cdots U_{1,n-1} |\psi_1\rangle = e_1$. Note that their effect on the remaining columns will make them orthogonal to e_1 . Next we may find $U_{2,2}, \dots, U_{2,n-1}$ with $U_{2,j} \in \mathbf{U}(2)_j$ such that their product applied to $U_{1,1} \cdots U_{1,n-1} |\psi_2\rangle$ is e_2 . Continuing, we obtain $U = U_{n-1,n-1} U_{n-2,n-2} U_{n-3,n-2} \cdots U_{2,2} \cdots U_{2,n-1} U_{1,1} \cdots U_{1,n-1}$. \square

3.4.2. Approximation. Ideally one would like to work with a universal gate set as in the classical case, but that will not be possible with a finite (or even countable) set of gates. If we start with a finite gate set, we cannot generate the entire unitary group. This results in quantum computation in general being not only probabilistic, as we have already discussed, but also *approximate*. This has not been an issue so far, but will be when we do Shor's algorithm.

We would like some assurance that we can get arbitrarily “close” to any $U \in \mathbf{U}(\mathcal{H})$ with polynomially many elements from our finite gate set.

First we need to make the meaning of “close” precise, i.e., we need to choose a norm on $\text{End}(\mathcal{H})$. We work with a norm on $\text{End}(\mathcal{H})$ rather than a distance function on $\mathbf{U}(\mathcal{H})$ because we will need to measure the norm of the difference of two unitary operators, which in general is not unitary. Moreover, $\text{End}(\mathcal{H})$ has the advantage of being a linear space where distances are easier to work with.

Let V be a vector space. A *norm* on V is a function $\|\cdot\| : V \rightarrow \mathbb{R}_{\geq 0}$ satisfying, for all $v, w \in V$ and all $c \in \mathbb{C}$:

$$\begin{aligned} \|v\| &\geq 0 \text{ with equality iff } v = 0, \\ \|v + w\| &\leq \|v\| + \|w\|, \\ \|cv\| &= |c| \|v\|. \end{aligned}$$

In our case, we have $V = \text{End}(\mathcal{H})$, and we take a norm that reflects this additional structure of our vector space as a space of operators, called the *operator norm*. It is particularly convenient for unitary operators.

Definition 3.4.2. For $X \in \text{End}(\mathcal{H})$, define the *operator norm* of X ,

$$\|X\| := \sup_{|\xi\rangle \neq 0} \frac{|X|\xi\rangle|}{|\xi\rangle|}$$

where $|X|\psi\rangle|$ is the usual Hermitian norm of the vector $X|\psi\rangle$.

Exercise 3.4.3: Verify that the operator norm is indeed a norm.

Exercise 3.4.4: Prove that when \mathcal{H} is finite dimensional, that $\|X\|^2$ is the largest eigenvalue of $X^\dagger X$.

For $X \in \text{End}(\mathcal{H}_A)$, $Z \in \text{End}(\mathcal{H}_B)$, we may consider $X \otimes Z \in \text{End}(\mathcal{H}_A \otimes \mathcal{H}_B)$. The operator norm has the following additional properties:

$$\begin{aligned} \|XY\| &\leq \|X\| \|Y\|, \\ \|X^\dagger\| &= \|X\|, \\ \|X \otimes Z\| &= \|X\| \|Z\|, \\ \|U\| &= 1 \quad \forall U \in \mathbf{U}(\mathcal{H}). \end{aligned}$$

Exercise 3.4.5: Verify these additional properties.

Definition 3.4.6. An operator $\tilde{U} \in \mathbf{U}(\mathcal{H})$ approximates $U \in \mathbf{U}(\mathcal{H})$ with precision δ if $\|\tilde{U} - U\| \leq \delta$.

Exercise 3.4.7: Show that $\tilde{U} \in \mathbf{U}(\mathcal{H})$ approximates $U \in \mathbf{U}(\mathcal{H})$ with precision δ if and only if \tilde{U}^{-1} approximates U^{-1} with precision δ .

We will not be concerned with a single unitary transformation, but the product of many such, so we need to examine how errors grow under composition of maps. The key property of the operator norm is that for unitary transformations, errors accumulate linearly:

Proposition 3.4.8. Say $U = U_L \cdots U_2 U_1$ with $U, U_j \in \mathbf{U}(\mathcal{H})$, and that U_j is approximated by $\tilde{U}_j \in \mathbf{U}(\mathcal{H})$ with precision δ_j . Then $\tilde{U} := \tilde{U}_L \cdots \tilde{U}_2 \tilde{U}_1$ approximates U with precision $\sum_{j=1}^L \delta_j$.

Proof. By induction, it will be sufficient to prove the case $L = 2$. We have

$$\begin{aligned} \|\tilde{U}_2 \tilde{U}_1 - U_2 U_1\| &= \|\tilde{U}_2(\tilde{U}_1 - U_1) + (\tilde{U}_2 - U_2)U_1\| \\ &\leq \|\tilde{U}_2(\tilde{U}_1 - U_1)\| + \|(\tilde{U}_2 - U_2)U_1\| \\ &\leq \|\tilde{U}_2\| \|\tilde{U}_1 - U_1\| + \|\tilde{U}_2 - U_2\| \|U_1\| \\ &\leq \|\tilde{U}_1 - U_1\| + \|\tilde{U}_2 - U_2\|. \end{aligned}$$

□

This linear accumulation of errors allows for good approximation as we will see in Theorem 3.4.12 below.

3.4.3. Classical-quantum gates. Recall that in Grover's and Simons' algorithms, we used the notation (3.2.2) as an abbreviation for a reversible classical computation embedded into a quantum circuit. We now introduce notation for these classical “controls” in a quantum circuit.

For $U \in \mathbf{U}(n)$, introduce k -controlled U , $\Lambda^k(U) : \mathbb{C}^k \otimes \mathbb{C}^n \rightarrow \mathbb{C}^k \otimes \mathbb{C}^n$ by

$$\Lambda^k(U)(|x_1, \dots, x_k\rangle \otimes |\xi\rangle) = \begin{cases} |x_1, \dots, x_k\rangle \otimes |\xi\rangle & \text{if } x_1 \cdots x_k = 0 \\ |x_1, \dots, x_k\rangle \otimes U|\xi\rangle & \text{if } x_1 \cdots x_k = 1. \end{cases}$$

When acting by these controlling bits, we will allow violation of strict locality for the controlling bits (they will be the “last” s bits, as was the situation with Grover's and Simons' algorithms). This is physically acceptable, because it will correspond to interfering with the quantum system from “outside”.

Introduce the notation (taken from physics) $\sigma_x := \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$.

Exercise 3.4.9: Show that $\Lambda^1(\sigma_x)|ab\rangle = |a, a \oplus b\rangle$.

Exercise 3.4.10: Show the Toffoli gate is $\Lambda^2(\sigma_x)$.

It is exactly in implementing these classical-quantum gates that we will allow violation of strict locality- the Hilbert spaces corresponding to the classical bits need not be adjacent to the Hilbert spaces corresponding to the quantum bits in these gates.

3.4.4. The standard quantum gate set.

Definition 3.4.11. The quantum gate set

$$(3.4.1) \quad H := \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix} \text{ (the Hadamard gate),}$$

$$(3.4.2) \quad K := \begin{pmatrix} 1 & 0 \\ 0 & i \end{pmatrix},$$

$$(3.4.3) \quad K^{-1},$$

$$(3.4.4) \quad \Lambda^1(\sigma_x) \text{ where } \sigma_x = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix},$$

$$(3.4.5) \quad \Lambda^2(\sigma_x) \text{ the Toffoli gate}$$

is called *standard*.

Theorem 3.4.12. Any $U \in \mathbf{U}(n)$ can be realized with precision δ by a $\text{poly}(\log(\frac{1}{\delta}))$ -size circuit over the standard basis, using workspace bits (i.e., working with elements of $\mathbf{U}(n+s)$). Moreover, there exists a polynomial size algorithm constructing the circuit.

Theorem 3.4.12 justifies the assertion that one can achieve good approximate quantum algorithms from a fixed gate set. For the proof, see [KSV02, Thm. 13.5].

**add exercise how to approximate elements of $\mathbf{U}(2)$ **

3.5. Shor's algorithm

Shor's algorithm involves a classical part and a quantum part. The quantum part is: given a randomly chosen $a \in (\mathbb{Z}/N\mathbb{Z})^*$, find the order of a in $(\mathbb{Z}/N\mathbb{Z})^*$. The classical part uses the quantum algorithm as a black box and finds a factor of N . The size of the quantum circuit will be $\text{poly}(\log(N))$. Since there are at most $\log(N)$ factors of N , it will still be $\text{poly}(\log(N))$ -operations to factor N completely.

The quantum part will hinge on i) there being "enough" prime numbers less than N , and ii) the ability to "closely" approximate a rational number by other rational numbers. After explaining the classical part, I address these two issues.

3.5.1. The classical part. Here is the algorithm that, given N , finds a nontrivial divisor with probability at least $\frac{1}{2}$.

If N is even, we are done. Otherwise, choose a uniformly at random from $\{2, \dots, N-1\}$. If $\gcd(a, N) > 1$ then we are done. Otherwise, we may consider a as an element of $(\mathbb{Z}/N\mathbb{Z})^*$, and we call the quantum algorithm to compute the order of a in $(\mathbb{Z}/N\mathbb{Z})^*$. Say this order is r .

If r is odd, the algorithm fails.

If r is even, compute $\gcd(a^{\frac{r}{2}} - 1, N)$. If it is greater than one, we are done, otherwise, the algorithm fails.

There are two bad cases to analyze: i) when r is odd and ii) when r is even and $\gcd(a^{\frac{r}{2}} - 1, N) = 1$.

Proposition 3.5.1. *The above algorithm succeeds with probability at least $1 - \frac{1}{2^{k-1}}$ where k is the number of distinct prime divisors of N .*

Proof. Introduce the following notation for the proof:

$$\begin{aligned} N &= p_1^{\alpha_1} \cdots p_k^{\alpha_k} \text{ with } p_j \text{ prime} \\ a_j &\text{ is } a \text{ reduced mod } p_j^{\alpha_j} \\ r_j &= \text{order of } a_j \text{ in } (\mathbb{Z}/p_j^{\alpha_j}\mathbb{Z})^*. \end{aligned}$$

Recall the Chinese remainder theorem (CRT) implies that $(\mathbb{Z}/N\mathbb{Z})^* = (\mathbb{Z}/p_1^{\alpha_1}\mathbb{Z})^* \times \cdots \times (\mathbb{Z}/p_k^{\alpha_k}\mathbb{Z})^*$, so in particular $r = \text{lcm}(r_1, \dots, r_k)$.

The first chance of failure happens when all r_j are odd. When will the second failure occur?

Write $r_j = 2^{s_j} r'_j$ with r'_j odd and $r = 2^s r'$ where $s = \max\{s_1, \dots, s_k\}$ and r' is odd. To succeed, we just need one j such that $a_j^{\frac{r}{2}} \equiv 1 \pmod{p_j^{\alpha_j}}$ (as then $a^{\frac{r}{2}} - 1 | N$).

Claim: the algorithm fails if and only if $s_1 = s_2 = \cdots = s_k$.

If there exists some $s_j < s$ and $a_j^{\frac{r}{2}} \equiv 1 \pmod{p_j^{\alpha_j}}$, since $-1 \mapsto (-1, \dots, -1)$ under the CRT isomorphism, we see $a^{\frac{r}{2}} \not\equiv -1 \pmod{N}$ and we obtain a nontrivial divisor, showing if the s_i are not all equal, the algorithm succeeds.

If $s_1 = \cdots = s_k = 0$, then r is odd and the algorithm fails at the first chance.

So assume $s_1 = \cdots = s_k \geq 1$. By the CRT $a_j^{r_j} \equiv 1 \pmod{p_j^{\alpha_j}}$, i.e.,

$$a_j^{r_j} - 1 = (a_j^{r_j/2} - 1)(a_j^{r_j/2} + 1) \equiv 0 \pmod{p_j^{\alpha_j}}.$$

Now recall Proposition 3.1.4 which asserts that if $b^2 \equiv 1 \pmod{M}$ and $b \not\equiv \pm 1 \pmod{M}$, then $b+1, b-1$ have nontrivial factors in common with M . Since

each r_j is minimal, we cannot have $(a_j^{r_j/2} - 1) \equiv 0 \pmod{p_j^{\alpha_j}}$. Thus, since p_j is a prime greater than 2, $\gcd(a_j^{r_j/2} - 1, p_j^{\alpha_j}) = 1$, so $a_j^{r_j/2} + 1 \equiv 0 \pmod{p_j^{\alpha_j}}$ for all j , i.e.,

$$a_j^{r_j/2} \equiv -1 \pmod{p_j^{\alpha_j}} \quad \forall j.$$

Again by the CRT

$$a^{\frac{r}{2}} \equiv -1 \pmod{N}$$

and the algorithm fails.

It remains to show that the probability of failure is small, i.e., that the probability that $s_1 = \dots = s_k$ is at most $\frac{1}{2^{k-1}}$.

Since a has been chosen uniformly at random in $(\mathbb{Z}/N\mathbb{Z})^*$, the CRT implies the a_j are chosen uniformly at random as elements of $(\mathbb{Z}/p_j^{\alpha_j}\mathbb{Z})^*$.

Fix j and any $s \geq 0$. We claim that the probability that $s_j = s$ is at most $\frac{1}{2}$. Once we have proven the claim, the theorem will follow.

Exercise 3.5.2: Show that for p prime, $|(\mathbb{Z}/p^\alpha\mathbb{Z})^*| = p^\alpha - p^{\alpha-1}$.

Write $p^\alpha - p^{\alpha-1} = 2^t q$ where q is odd and note that $t > 0$. Let g be a generator of $(\mathbb{Z}/p^\alpha\mathbb{Z})^*$. Write the order of a as $2^{s_a} r'_a$ where r'_a is odd.

$$\begin{aligned} \#\{a \in (\mathbb{Z}/p^\alpha\mathbb{Z})^* \mid s_a = s\} &= \#\{\text{powers } g^{2^{t-s}m} \text{ with } m \text{ odd}\} \\ &= \begin{cases} q & s = 0 \\ (2^s - 2^{s-1})q & s = 1, \dots, t \\ 0 & s > t \end{cases}. \end{aligned}$$

All these numbers are at most $\frac{1}{2}(p^\alpha - p^{\alpha-1})$ because $t > 0$, so they have probability at most $\frac{1}{2}$ of being selected, so we conclude the probability of failure is at most $\frac{1}{2^{k-1}}$. \square

3.5.2. Preliminaries I: the number of primes. Let $\pi(x)$ denote the number of prime numbers that are less than or equal to x . The *prime number theorem* states $\pi(x) \sim \frac{x}{\ln(x)}$, more precisely,

$$\lim_{x \rightarrow \infty} \frac{\pi(x)}{x/\ln(x)} = 1.$$

The proof is not so simple, but for our purposes, the following result will suffice:

Theorem 3.5.3. For all $x \geq 2$,

$$\pi(x) \geq \frac{x}{2 \log(x)}.$$

We give the proof after several preliminary lemmas. Let $v_p(m)$ denote the largest power of p that divides m .

Lemma 3.5.4.

$$v_p(n!) = \sum_{m \geq 1} \lfloor \frac{n}{p^m} \rfloor.$$

Proof. Among $1, 2, \dots, n$, exactly $\lfloor \frac{n}{p} \rfloor$ are multiples of p , contributing $\lfloor \frac{n}{p} \rfloor$ to the summation, exactly $\lfloor \frac{n}{p^2} \rfloor$ are multiples of p^2 , contributing an additional $\lfloor \frac{n}{p^2} \rfloor$ to the summation (and thus multiples of p^2 are counted two times). Continuing, one gets the result. \square

Recall that $\binom{m}{q} = \frac{m!}{q!(m-q)!}$. Note that

$$v_p\left(\binom{2n}{n}\right) = v_p(2n) - 2v_p(n) = \sum_{m \geq 1} \lfloor \frac{2n}{p^m} \rfloor - 2\lfloor \frac{n}{p^m} \rfloor.$$

Exercise 3.5.5: Show that for all $x \in \mathbb{R}$, $\lfloor 2x \rfloor - 2\lfloor x \rfloor \in \{0, 1\}$. \odot

Note that if $\frac{2n}{p^m} < 1$, i.e., $m > \frac{\log(2n)}{\log(p)}$, then $\lfloor \frac{2n}{p^m} \rfloor - 2\lfloor \frac{n}{p^m} \rfloor = 0$, and thus $v_p\left(\binom{2n}{n}\right) \leq \frac{\log(2n)}{\log(p)}$.

Exercise 3.5.6: Show that $\binom{2n}{n} \geq \frac{2^{2n}}{2n}$. \odot

By Exercise 3.5.6

$$\begin{aligned} 2n \log(2) - \log(2n) &\leq \log\left(\binom{2n}{n}\right) \\ &\leq \sum_{p \leq 2n} \lfloor \frac{\log(2n)}{\log(p)} \rfloor \log(p), \text{ as } N = \prod_p p^{v_p(N)} \\ &\leq \sum_{p \leq 2n} \log(2n) = \pi(2n) \log(2n). \end{aligned}$$

Thus

$$(3.5.1) \quad \pi(2n) \geq \frac{2n}{\log(2n)} - 1.$$

proof of Theorem 3.5.3. For $x \leq 16$, one can check the result by hand. So assume $x > 16$ and take n such that $16 \leq 2n \leq x \leq 2n + 2$. Then

$$\begin{aligned} \pi(x) &\geq \pi(2n) \geq \frac{2n}{\log(2n)} - 1 = \frac{1}{2 \log(2n)} \left[4n - \frac{1}{2 \log(2n)} \right] \\ &\geq \frac{1}{2 \log(x)} \left[4n - \frac{1}{2 \log(2n)} \right] \\ &\geq \frac{x}{2 \log(x)}. \end{aligned}$$

The last line holds because we have $x \leq 2n + 2 \leq 4n - \frac{1}{2\log(2n)}$, since $n \geq 8$. \square

We will use Theorem 3.5.3 in the form:

Corollary 3.5.7. *For every natural number r , there are at least $\Omega(\frac{r}{\log r})$ numbers in $\{1, \dots, \frac{r}{10}\}$ relatively prime to r .*

The corollary follows as r has at most $\log(r)$ prime factors and there are at least $\frac{r}{20(\log(r)-\log(10))}$ prime numbers less than $\frac{r}{10}$.

3.5.3. Preliminaries II: Continued Fractions. We will need to approximate real numbers by sequences of rational numbers. A first idea would simply be to use the decimal expansion. The following scheme has better convergence properties for our purposes. (E.g., consider the decimal expansion of $\frac{1}{3}$ compared with the results below.) Given $\alpha \in \mathbb{R}$, consider the expansion

$$\alpha = \alpha_0 + \frac{1}{\alpha_1 + \frac{1}{\alpha_2 + \frac{1}{\alpha_3 + \frac{1}{\alpha_4 + \dots}}}}$$

where $\alpha_0 = \lfloor \alpha \rfloor$, and

$$\frac{1}{\alpha - \alpha_0} = \alpha_1 + \left\lfloor \frac{1}{\alpha - \alpha_0} \right\rfloor$$

etc... I.e., α_k is the integer part of the reciprocal of the error term in the previous estimate. Write $\frac{p_n}{q_n} = [\alpha_0, \dots, \alpha_n]$ for the rational number obtained after the n -th step, which is an approximation to the real number α .

For example, taking the continued fraction expansion of π , one obtains $\frac{22}{7} = [3, 7]$, $\frac{333}{106} = [3, 7, 15]$, $\frac{355}{113} = [3, 5, 15, 1]$.

If α is rational, we will see momentarily that the algorithm reproduces α . For example $\frac{11}{9} = 1 + \frac{2}{9}$, so $\alpha_0 = 1$. Since $\frac{9}{2} = 1 + \frac{1}{4}$, we have $\alpha_1 = 4$, and since $2 = 2 + 0$, we have $\alpha_2 = 2$ and $\frac{11}{9} = [1, 4, 2]$.

Proposition 3.5.8. *In the continued fraction expansion of $\alpha \in \mathbb{R}_{>0}$, if it does not converge at the n -th step, then after the n -th step one obtains $[\alpha_0, \dots, \alpha_n] = \frac{a}{b} \in \mathbb{Q}$ with $b \geq 2^{\frac{n}{2}}$.*

Proof. Write $[\alpha_0, \dots, \alpha_n] = \frac{p_n}{q_n}$. Then $p_0 = \alpha_0$, $q_0 = 1$, $p_1 = 1 + p_0\alpha_1$, $q_1 = \alpha_1$ and

$$(3.5.2) \quad \begin{aligned} p_k &= \alpha_k p_{k-1} + p_{k-2} \\ q_k &= \alpha_k q_{k-1} + q_{k-2}. \end{aligned}$$

Exercise 3.5.9: Verify the equalities (3.5.2).

Since $\alpha_j > 0$, we have $p_j \geq 2p_{j-2}$ and $q_j \geq 2q_{j-2}$, so $p_{n-1}, q_{n-1} \geq 2^{\lfloor \frac{n}{2} \rfloor}$, proving the lower bound on b . \square

Theorem 3.5.10. *Let $\alpha, \frac{s}{r} \in \mathbb{Q}$ be such that $|\frac{s}{r} - \alpha| \leq \frac{1}{2r^2}$. Then $\frac{s}{r}$ appears in the continued fraction expansion for α .*

Proof. Take the continued fraction expansion of $\frac{s}{r} = [\beta_0, \dots, \beta_n]$, and write $\frac{p_j}{q_j} = [\beta_0, \dots, \beta_j]$. Define δ by the equation

$$\alpha = \frac{s}{r} + \frac{\delta}{2r^2} = \frac{p_n}{q_n} + \frac{\delta}{2q_n^2}.$$

Note that δ is a measure of the failure of $\frac{p_n}{q_n}$ to be equal to α . The hypothesis implies $|\delta| < 1$. Set

$$\lambda = 2 \left(\frac{q_n p_{n-1} - p_n q_{n-1}}{\delta} \right) - \frac{q_{n-1}}{q_n}.$$

Exercise 3.5.11: Show that $\alpha = \frac{\lambda p_n + p_{n-1}}{\lambda q_n + q_{n-1}}$.

Slightly abusing notation (as $\lambda \notin \mathbb{Z}$) $\alpha = [\beta_0, \dots, \beta_n, \lambda]$. Since $\lambda \in \mathbb{Q}$, we may write $\lambda = [\gamma_0, \dots, \gamma_m]$, so $\alpha = [\beta_0, \dots, \beta_n, \gamma_0, \dots, \gamma_m]$ and $\frac{s}{r}$ appears at the n -th step. \square

delete or clean this Punch line: Given $\alpha \in \mathbb{R}$ and $N \in \mathbb{Z}$, the continued fraction algorithm, in $\text{poly}(\log(N))$ steps finds $\frac{a}{b} \in \mathbb{Q}$ such that $b \leq 16N$ and $\frac{a}{b}$ approximates α better than any other rational number with denominator at most b . **more details here***

3.5.4. Preliminaries III: the quantum Discrete Fourier Transform.

Recall the DFT for $\mathbb{Z}/M\mathbb{Z}$, which we may write in vector notation, for $j \in \mathbb{Z}/M\mathbb{Z}$, as

$$|j\rangle \mapsto \frac{1}{\sqrt{M}} \sum_{k=0}^{M-1} \omega^{jk} |k\rangle$$

where $\omega = e^{\frac{2\pi i}{M}}$. Also recall: the DFT is a unitary change of basis such that in the new basis, multiplication in $\mathbb{Z}/M\mathbb{Z}$ is given by a diagonal matrix, and the classical FFT writes the DFT as a product of $O(\log(M))$ sparse matrices (each with $M \ll M^2$ nonzero entries), for a total cost of $O(\log(M)M) < O(M^2)$ arithmetic operations to execute.

Write $M = 2^m$. We show that the DFT can be written as a product of $O(m^3) = O(\log(M)^3)$ controlled local unitary operators. We will thus be able to produce an approximation of the output vector by a sequence of $\text{poly}(m)$ unitary operators from our gate set.

It will be convenient to express j in binary and view $\mathbb{C}^M = (\mathbb{C}^2)^{\otimes m}$, i.e., write

$$|j\rangle = |j_1\rangle \otimes \cdots \otimes |j_m\rangle$$

where $j = j_1 2^{m-1} + j_2 2^{m-2} + \cdots + j_m 2^0$ and $j_i \in \{0, 1\}$. Write the DFT as

$$|j_1\rangle \otimes \cdots \otimes |j_m\rangle$$

$$\mapsto \frac{1}{\sqrt{M}} \sum_{k=0}^{M-1} \omega^{jk} |k\rangle$$

$$= \frac{1}{\sqrt{M}} \sum_{k_i \in \{0,1\}} \omega^{j(\sum_{l=1}^m k_l 2^{m-l})} |k_1\rangle \otimes \cdots \otimes |k_m\rangle$$

$$= \frac{1}{\sqrt{M}} \sum_{k_i \in \{0,1\}} \bigotimes_{l=1}^m [\omega^{j k_l 2^{m-l}} |k_l\rangle]$$

$$= \frac{1}{\sqrt{M}} \sum_{k_i \in \{0,1\}} \bigotimes_{l=1}^m [\omega^{(j_1 2^{2m-1-l} + \cdots + j_m 2^{m-l}) k_l} |k_l\rangle]$$

(3.5.3)

$$= \frac{1}{2^{\frac{m}{2}}} (|0\rangle + \omega^{j_m 2^{-1}} |1\rangle) \otimes (|0\rangle + \omega^{j_{m-1} 2^{-1} + j_m 2^{-2}} |1\rangle) \otimes (|0\rangle + \omega^{j_{m-2} 2^{-1} + j_{m-1} 2^{-2} + j_m 2^{-3}} |1\rangle) \\ \otimes \cdots \otimes (|0\rangle + \omega^{j_1 2^{-1} + j_2 2^{-2} + \cdots + j_m 2^{-m}} |1\rangle)$$

where for the last line if $2m-s-l > m$, i.e., $s+l < m$, there is no contribution with j_s because $\omega^{2^m} = 1$, and we multiplied all terms by $1 = \omega^{-2^m}$ to have negative exponents.

It will be notationally more convenient to write the quantum circuit for this vector with the order of factors reversed. I.e., we describe a quantum circuit that produces

(3.5.4)

$$\frac{1}{\sqrt{2}} ((|0\rangle + \omega^{j_1 2^{-1}} \omega^{j_2 2^{-2}} \cdots \omega^{j_m 2^{-m}} |1\rangle)) \otimes \cdots \otimes \frac{1}{\sqrt{2}} (|0\rangle + \omega^{j_{m-2} 2^{-1} + j_{m-1} 2^{-2} + j_m 2^{-3}} |1\rangle) \\ \otimes \frac{1}{\sqrt{2}} (|0\rangle + \omega^{j_{m-1} 2^{-1} + j_m 2^{-2}} |1\rangle) \otimes \frac{1}{\sqrt{2}} (|0\rangle + \omega^{j_m 2^{-1}} |1\rangle).$$

Set

$$(3.5.5) \quad R_k = \begin{pmatrix} 1 & 0 \\ 0 & \omega^{2^{-k}} \end{pmatrix},$$

then (3.5.4) is obtained from $|j_1\rangle \otimes \cdots \otimes |j_m\rangle$ as follows: first apply H to $(\mathbb{C}^2)_1$. Note that $\omega^{2^{-1}} = \omega^{2^{m-1}} = -1$, so at this point the first factor becomes $\frac{1}{\sqrt{2}} ((|0\rangle + \omega^{j_1 2^{-1}} |1\rangle))$. To get the next term in the product of ω 's in the first term, apply $\Lambda^1 R_2$ to $(\mathbb{C}^2)_2 \otimes (\mathbb{C}^2)_1$. Continue, applying $\Lambda^1 R_j$ to

$(\mathbb{C}^2)_j \otimes (\mathbb{C}^2)_1$ for $j = 3, \dots, m$. Note that at this point only the $(\mathbb{C}^2)_1$ -term has been altered, as all the other factors only acted as controlling qubits.

Exercise 3.5.12: Verify that at this point we have the correct term in the $(\mathbb{C}^2)_1$ -slot in (3.5.4).

From now on we leave the $(\mathbb{C}^2)_1$ -slot alone. Next apply H to $(\mathbb{C}^2)_2$ then $\Lambda^1 R_{j-1}$ to $(\mathbb{C}^2)_j \otimes (\mathbb{C}^2)_2$ for $j = 3, \dots, m$. Then apply H to $(\mathbb{C}^2)_3$ then $\Lambda^1 R_{j-2}$ to $(\mathbb{C}^2)_j \otimes (\mathbb{C}^2)_3$ for $j = 4, \dots, m$. Continue, until finally one just applies H to $(\mathbb{C}^2)_m$. In all these cases, the $(\mathbb{C}^2)_j$'s act as the control. Finally to obtain the DFT, reverse the orders of the factors (a classical operation).

Exercise 3.5.13: Verify that the above sequence of maps produces (3.5.3).

In practice, one has to fix a quantum gate set in advance, so in general we will have to approximate the transformations R_k from elements of our gate set, so we will only approximate the DFT.

3.5.5. The order finding algorithm. We are given a and N with $\gcd(a, N) = 1$ and want to find the order of a in $(\mathbb{Z}/N\mathbb{Z})^*$. Set $m = \lceil 3 \log(N) \rceil$, $M = 2^m$ and $n = \lceil \log(N) \rceil$. Work in $(\mathbb{C}^2)^{\otimes(m+n)}$. (Here and in what follows, I ignore the additional bits needed to execute the classical reversible computations.) Initialize the state to $|0^{m+n}\rangle$. Apply the quantum Fourier transform to the first m qubits to obtain

$$\frac{1}{\sqrt{M}} \sum_{x \in \mathbb{Z}/M\mathbb{Z}} |x\rangle \otimes |0^k\rangle.$$

Use our (polynomial in n) reversible powering mod N routine $|x\rangle \otimes |y\rangle \mapsto |x\rangle \otimes |y \oplus a^x \bmod N\rangle$ to obtain

$$\frac{1}{\sqrt{M}} \sum_{x \in \mathbb{Z}/M\mathbb{Z}} |x\rangle \otimes |a^x \bmod N\rangle.$$

Here and in what follows, $a^x \bmod N$ is to be considered as an element of $\{0, \dots, N-1\}$ expressed in binary. Now perform a measurement on the last n qubits (which recall involves re-normalizing the resulting vector to have length one) to obtain some $y_0 \in \{0, 1\}^n$ that appears in the image of the map $x \mapsto a^x \bmod N$, which will appear with all x 's such that $a^x \bmod N = y_0$. Let x_0 be the smallest natural number such that $a^{x_0} \equiv y_0 \bmod N$. Then y_0 will appear paired with the values $x_0 + \ell r \in \{0, \dots, M-1\}$ where r is the period. Thus the output vector is

$$(3.5.6) \quad \frac{1}{\sqrt{K}} \sum_{\ell=0}^{K-1} |x_0 + \ell r\rangle \otimes |y_0\rangle,$$

where

$$K = \lfloor \frac{M - (1 + x_0)}{r} \rfloor.$$

Compare (3.5.6) with the step in Simons' algorithm giving rise to (3.3.1). As with Simons' algorithm, we now have a vector that "sees" r that we will manipulate to get information about r .

Apply the quantum Fourier transform to the first m qubits again to obtain

$$\frac{1}{\sqrt{M}\sqrt{K}} \sum_{x \in \mathbb{Z}/M\mathbb{Z}} \sum_{\ell=0}^{K-1} \omega^{(x_0 + \ell r)x} |x\rangle \otimes |y_0\rangle.$$

Finally measure the first m qubits to obtain some $x \in \mathbb{Z}/M\mathbb{Z}$.

How will this be useful? We claim that the algorithm will tend to produce x such that $\frac{x}{M}$ will be close to a fraction $\frac{b}{r}$. Then if we take the partial fraction decomposition of $\frac{x}{M}$, the term $\frac{b}{r}$ will appear. In fact we can just test the denominators and q_j and just take the first one such that $a^{q_j} \equiv 1 \pmod{N}$.

We know if we draw x such that $|\frac{b}{r} - \frac{x}{M}| < \frac{1}{2r^2}$, we obtain r .

We need to show there are "enough" "good" x each with sufficiently high probability of being chosen that we succeed. Since we may repeat the experiment $\text{poly}(\log N)$ times while still being in polynomial time, if G is the number of good x and P the probability of drawing a good x , we need $GP = \Omega(\frac{1}{\log N})$. We will show $G = \Omega(\frac{r}{\log r})$ and $P = \Omega(\frac{1}{r})$, and since $r < N$, this will suffice.

The probability of drawing any given $|x\rangle$ is

$$\begin{aligned} & \frac{1}{MK} |\omega^{x_0 x} + \omega^{(x_0+r)x} + \omega^{(x_0+2r)x} + \dots + \omega^{(x_0+(K-1)r)x}|^2 \\ &= \frac{1}{MK} |1 + \omega^{rx} + \omega^{2rx} + \dots + \omega^{(K-1)rx}|^2 \end{aligned}$$

First consider the highly improbable but illustrative case $M = rc$ for some natural number c . Write $\eta = \omega^{rx}$. Note that $\eta^c = (\omega^{rc})^x = 1$. If $c \nmid x$, then $1 + \eta + \eta^2 + \dots + \eta^{c-1} = 0$ and the probability of drawing such x is zero, but if $c \mid x$, each power equals 1 as $\omega^{rjx} = \omega^{\frac{M}{c}jx}$ and $\frac{x}{c} \in \mathbb{Z}$, and thus we will draw $x = cb$ for a random $b \in \{0, \dots, r-1\}$ so $\frac{x}{M} = \frac{cb}{rc} = \frac{b}{r}$ and all such x 's are equally likely.

From this we can recover r with high probability: there are $\Omega(\frac{r}{\log(r)})$ numbers relatively prime to r , when we run the continued fraction algorithm for $\frac{x}{M}$, the denominator it produces will be r with high probability.
clean*

Now for the general case: the idea is similar- we will prove that the values of x that are most likely to be drawn will be such that xr is nearly

divisible by M . For the algorithm to work, we also need $\gcd(\lfloor \frac{xr}{M} \rfloor, r) = 1$. We first show there are “many” such good x .

Lemma 3.5.14. *There exist $\Omega(\frac{r}{\log(r)})$ elements $x \in \mathbb{Z}/M\mathbb{Z}$ satisfying:*

- i) $0 \leq xr \bmod M < \frac{r}{10}$, and
- ii) $\gcd(\lfloor \frac{xr}{M} \rfloor, r) = 1$.

Proof. Consider the case $\gcd(r, M) = 1$, i.e., r is odd. Then the map $x \mapsto rx \bmod M$ is a permutation on $(\mathbb{Z}/M\mathbb{Z})^*$. By Corollary 3.5.7 there are $\Omega(\frac{r}{\log(r)})$ elements x such that $xr \bmod M < \frac{r}{10}$ and $\gcd(x, r) = 1$.

Since $xr \bmod M = rx - \lfloor \frac{rx}{M} \rfloor M$, we have $\gcd(\lfloor \frac{rx}{M} \rfloor, r) = 1$ as otherwise $rx \bmod M$ would also have a factor in common with r .

For the general case, write $d = \gcd(r, M)$, set $r' = \frac{r}{d}$ and $M' = \frac{M}{d}$.

Exercise 3.5.15: Apply the same argument to show there exists $\Omega(\frac{r}{d \log(r)})$ x 's in $\mathbb{Z}/M'\mathbb{Z}$ satisfying the condition. Finally show that for all $c \in \mathbb{N}$, $x + cM$ also satisfies the condition. □

To conclude, we need to show the probability of drawing one of the good x 's is sufficiently large.

Lemma 3.5.16. *If x is such that $0 < xr \bmod M < \frac{r}{10}$, then the probability of drawing x in Shor's algorithm is $\Omega(\frac{1}{r})$.*

Proof. The probability of drawing x is

$$\frac{1}{KM} \left| \sum_{\ell=0}^{K-1} \omega^{\ell rx} \right|^2.$$

Recalling the sum of a geometric series, this is

$$\frac{1}{KM} \left| \frac{1 - e^{\frac{2\pi i rx K}{M}}}{1 - e^{\frac{2\pi i rx}{M}}} \right|^2.$$

Note that $\frac{M}{2r} < K < \frac{M}{r}$. We need to show the quantity in absolute value is $\Omega(\frac{M}{r})$. Setting $\theta = \frac{rx \bmod M}{M}$, it is at least

$$\frac{\sin(\theta \lceil \frac{M}{r} \rceil / 2)}{\sin(\theta / 2)}.$$

Now use that for small θ , $\sin(\theta) \sim \theta$ to conclude. □

The first assertion of Lemma 3.5.14 implies

$$|xr - cM| < \frac{r}{10}$$

for $c = \lceil \frac{xr}{M} \rceil$, i.e.,

$$\left| \frac{x}{M} - \frac{c}{r} \right| < \frac{1}{10M}.$$

Thus the rational number $\frac{c}{r}$ approximates $\frac{x}{M}$ to within $\frac{1}{10M}$.

Exercise 3.5.17: Show that for every $0 < \alpha < 1$ and $N > 0$, there is at most a single rational number $\frac{a}{b}$ such that $b < N$, and $|\alpha - \frac{a}{b}| < \frac{1}{2N^2}$.

By Exercise 3.5.17 our approximation is good enough to determine $\frac{c}{r}$, and thus r .

3.6. A unified perspective on quantum algorithms: the hidden subgroup problem

Given G : a discrete group with a specific representation of its elements in binary, an explicit function $f : G \rightarrow \mathbb{F}_2^n$, and the knowledge that there exists a subgroup $G' \subset G$ such that $f(x) = f(y)$ if and only if $xy^{-1} \in G'$, find G' .

For abelian groups, it is sufficient to solve the problem for $G = \mathbb{Z}^{\oplus k}$ as all abelian groups are quotients of some $\mathbb{Z}^{\oplus 2k}$.

Simons algorithm is the case $G = \mathbb{Z}_2^{\oplus m}$. The DFT_2 matrix is

$$H = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & -1 \\ -1 & 1 \end{pmatrix}$$

and G' is the subgroup generated by $a \in \mathbb{Z}_2^{\oplus m}$.

Shor's algorithm is the case $G = \mathbb{Z}$ and F is the function $x \mapsto a^x \bmod N$. Note that with Shor, as we did with multiplying polynomials, we restricted to $\mathbb{Z}/M\mathbb{Z}$ for M sufficiently large.

Both are solved via the DFT for the finite group G .

Another example, closely related to order finding, is as follows: the *discrete logarithm* of a number a at base $\zeta = e^{\frac{2\pi i}{N}}$ is the smallest positive integer s such that $\zeta^s = a$. Consider the function $f : \mathbb{Z}^2 \rightarrow \mathbb{Z}_N$ given by $f(x_1, x_2) = \zeta^{x_1} a^{x_2} \bmod N$. Take $G = \mathbb{Z}^2$, $G' = \{(x_1, x_2) \in \mathbb{Z}^2 \mid \zeta^{x_1} a^{x_2} \equiv 1 \bmod N\}$. Given G' , just find an element of the form $(s, -1) \in G'$, and s is the discrete logarithm of a at base ζ .

3.7. What is a quantum computer?

Currently there are three approaches towards building a quantum computer: adiabatic quantum optimization (D Wave), digital (IBM), topological (currently science fiction).

The main problem of quantum computing is on the one hand, one wants an isolated system to be resistant to outside noise, on the other hand one needs to be able to manipulate it.

3.7.1. D Wave's computers. Claims of 2,000 qubits. Machine only designed for quadratic optimization. Yet to have an advantage over a classical computer. Internal workings of machine kept private. Principle quantum cousin of MRI. If they can scale up significantly, despite these drawbacks, they will beat classical computers.

3.7.2. IBM's digital quantum computer. It really exists and is quantum, and the world is free to examine it. Problem: only 5 qubits.

3.7.3. Topological quantum computing. technology not yet there...

3.8. Appendix: review of basic information on groups and rings

Let S be a set, a *binary operation* on S is a map $f : S \times S \rightarrow S$. One often writes $f(x, y) = x * y$.

A *group* is a set G with a binary operation such that 1) for all $x, y, z \in G$, $x * (y * z) = (x * y) * z$ (associativity), 2) there exists and identity element $e \in G$ such that $e * x = x * e = x$ for all $x \in G$ and 3) for all $x \in G$, there exists an inverse $x^{-1} \in G$ such that $x * x^{-1} = x^{-1} * x = e$.

Examples: $(G, *) = (\mathbb{Z}, +)$, $\mathbf{U}(n)$ with operation matrix multiplication. Let $\mathbb{Z}/N\mathbb{Z}$ denote the set of equivalence classes in the integers defined by remainder under division by N . Then $(\mathbb{Z}/N\mathbb{Z})^*$, the set of elements of $\mathbb{Z}/N\mathbb{Z}$ with multiplicative inverses is a group, where the binary operation is multiplication inherited from multiplication of the integers: $[x] * [y] = [x * y]$ (one must verify this is well defined).

Non-examples: the set of stochastic matrices with operation matrix multiplication, the nonzero integers with operation multiplication.

A group G is *abelian* if $x * y = y * x$ for all $x, y \in G$. $(\mathbb{Z}, +)$, $(\mathbb{Z}/N\mathbb{Z})^*$ are abelian groups, $\mathbf{U}(n)$ is not, when $n > 1$.

An abelian group is *cyclic* if it is generated by a single element.

Exercise 3.8.1: Show that when p is prime $(\mathbb{Z}/p\mathbb{Z})^*$ is cyclic of order $p - 1$.

A *ring* is a set R with two binary operations, often denoted $+$ and $*$, such that 1) $(R, +)$ is an abelian group, 2) $*$ is associative and has an identity, often denoted 1 or 1_R , 3) (compatibility) for all $x, y, z \in R$, $(x + y) * z = x * z + y * z$ and $z * (x + y) = z * x + z * y$.

Examples $(\mathbb{Z}, +, *)$, $(\mathbb{Z}/N\mathbb{Z}, +, *)$ (where both operations are inherited from the operations on \mathbb{Z}), $\mathbb{Z}[x]$: the polynomials in one variable with integer coefficients.

Let G, H be groups. A map $f : G \rightarrow H$ is a *group homomorphism* if $f(x * y) = f(x) * f(y)$ (where the first $*$ is in G and the second in H) for all $x, y \in G$. Let R, S be rings. A map $f : R \rightarrow S$ is a *ring homomorphism* if $f(1_R) = 1_S$, $f(x + y) = f(x) + f(y)$ and $f(x * y) = f(x) * f(y)$ for all $x, y \in R$.

Classical information theory

We have been referring to the classical unit of information as a *bit*, a copy of $\{0, 1\}$. The discovery/invention of the bit by Tukey and its development by Shannon [Sha48] was one of the great scientific achievements of the twentieth century, as it changed the way we view information, giving it an abstract formalism that is discussed in this chapter. Instead of reading this chapter, we suggest just reading Shannon's classic article, as it is extremely well written, with carefully chosen examples.

The basic question is: given a physical channel (e.g., telegraph wire), what is the maximal rate of transmission, tolerating a small amount of error. I begin with toy examples, leading up to Shannon's two fundamental theorems on channel capacity.

4.1. Data compression: noiseless channels

4.1.1. A toy problem. (Following [BCHW16]) A source emits symbols x from an alphabet \mathcal{X} that we want to store efficiently so we try to encode x in a small number of bits, to say $y \in \mathcal{Y}$ in a way that we can decode it later to recover x .




Figure 4.1.1. Message from source encoded into bits then decoded

The symbols from \mathcal{X} are not necessarily emitted with the same frequency. Let $p = P_{\mathcal{X}}$ denote the associated probability distribution. We want to determine the minimum possible size for \mathcal{Y} . Since we are dealing in bits, it will be convenient to use the logarithms of cardinalities, so define $\text{Cap}(P_{\mathcal{X}}) := \min \log |\mathcal{Y}|$.

Consider the case $\mathcal{X} = \{a, b, c, d\}$ where $\Pr(a) = 0.1$, $\Pr(b) = 0$, $\Pr(c) = 0.4$ and $\Pr(d) = 0.5$. We can clearly get away with \mathcal{Y} having cardinality 3, e.g., for the encoder, send a, b to 1, c to 2 and d to 3, then for the decoder, send 1 to a , 2 to c and 3 to d . In general, we can always throw away symbols with probability zero. On the other hand, we cannot map two distinct symbols that do occur to the same symbol, as there would be no way to distinguish them when decoding. Thus $\text{Cap}(p) = \log \text{supp}(p)$, where $\text{supp}(p) = \#\{x \in \mathcal{X} \mid \Pr(x) > 0\}$.

Now say we are willing to tolerate a small error. First rephrase what we did probabilistically: Let $p^{enc}(y|x)$ denote the conditional probability distribution of the encoder \mathcal{E} and $p^{dec}(x'|y)$ that of the decoder \mathcal{D} . Our requirement was for all x ,

$$\Pr[\forall x \in \mathcal{X}, x = \mathcal{D} \circ \mathcal{E}(x)] = \sum_{x,y,x'} \Pr^{enc}(y|x) \Pr^{dec}(x'|y) \delta_{x,x'} = 1.$$

and we now relax it to

$$\sum_{x,y,x'} \Pr(x) p^{enc}(y|x) p^{dec}(x'|y) \delta_{x,x'} \geq 1 - \epsilon.$$

for some error ϵ that we are willing to tolerate. In addition to throwing out the symbols that do not appear, we may also discard the largest set of symbols whose total probability is smaller than ϵ . Call the corresponding quantity $\text{Cap}^{\epsilon}(p)$.

In the example above, if we take $\epsilon > 0.1$, we can lower storage cost, taking $|\mathcal{Y}| = 2$.

Recall that a probability distribution p on \mathcal{X} must satisfy $\sum_{x \in \mathcal{X}} \Pr(x) = 1$. We relax this to *non-normalized* probability distributions, q , where $q(x) \geq 0$ for all $x \in \mathcal{X}$ and $\sum_{x \in \mathcal{X}} q(x) \leq 1$. We obtain: $\text{Cap}^{\epsilon}(p) = \min \log \text{supp}(q)$, where the min is taken over all non-normalized probability distributions q satisfying $q(x) \leq p(x)$ and $\sum_{x \in \mathcal{X}} q(x) \geq 1 - \epsilon$.

4.1.2. The case of interest. Now say the transmission is not a single symbol, but a string of n symbols, so we seek an encoder $\mathcal{E} : \mathcal{X}^n \rightarrow \mathcal{Y}(n)$, where $\mathcal{Y}(n)$ is a set that varies with n , and decoder $\mathcal{D} : \mathcal{Y}(n) \rightarrow \mathcal{X}^n$, and we want to minimize $|\mathcal{Y}(n)|$, with a tolerance of error that goes to zero for n going to infinity. In practice one wants to send information through a communication channel (e.g. telegraph wire). The channel can only send a limited number

of bits per second, and we want to maximize the amount of information we can send per second. Define $\text{Rate}(p) := \lim_{\epsilon \rightarrow 0} \lim_{n \rightarrow \infty} \frac{1}{n} \text{Cap}^\epsilon(p^n)$. We would like to simplify the right hand side of this expression.

4.1.3. Estimating $|\mathcal{Y}(n)|$. Define a map $wt : \mathcal{X}^n \rightarrow \mathbb{R}^d$ by $\bar{x}^n \mapsto (c_1, \dots, c_d)$, where c_j is the number of times j occurs in the string. Then $E[wt(\bar{x}^n)] = (np_1, \dots, np_d)$. The weak law of large numbers (1.3.2) states that for any $\epsilon > 0$,

$$\lim_{n \rightarrow \infty} \Pr\left[\left\|\frac{1}{n}(wt(\bar{x}^n) - E[wt(\bar{x}^n)])\right\|_1 > \epsilon\right] = 0$$

where for $f : \mathcal{Z} \rightarrow \mathbb{R}^d$, define $\|f\|_1 = \sum_{z \in \mathcal{Z}} |f(z)|$. In our case, $\mathcal{Z} = \mathcal{X}^n$ and $f = wt$.

Thus we now simply throw out all strings \bar{x}^n with $\left\|\frac{1}{n}(wt(\bar{x}^n) - E[wt(\bar{x}^n)])\right\|_1 > \epsilon$, and we can take $\mathcal{Y}(n)$ of size

$$\begin{aligned} |\mathcal{Y}(n)| &= \#\{\bar{x}^n \mid \left\|\frac{1}{n}(wt(\bar{x}^n) - E[wt(\bar{x}^n)])\right\|_1 < \epsilon\} \\ &= \sum_{\substack{\bar{x}^n \\ \left\|\frac{1}{n}(wt(\bar{x}^n) - E[wt(\bar{x}^n)])\right\|_1 < \epsilon}} \binom{n}{wt(\bar{x}^n)}. \end{aligned}$$

If ϵ is small, the multinomial coefficients appearing will all be very close to

$$\binom{n}{np_1, \dots, np_d}$$

and the number of \bar{x}^n of a given weight grows like a polynomial in n , so for what follows, we can take the crude approximation (which will be justified later)

$$(4.1.1) \quad |\mathcal{Y}(n)| \leq \text{poly}(n) \binom{n}{np_1, \dots, np_d}$$

(recall that d is fixed).

When we take logarithms, the right hand side of (4.1.1) becomes $nH(\bar{p}) + O(\log(n))$. Thus

$$\frac{1}{n} \log |\mathcal{Y}(n)| \leq H(\bar{p}) + o(1)$$

and $\text{Rate}(\bar{p}) \leq H(\bar{p})$.

Theorem 4.1.1. [Sha48] $\text{Rate}(\bar{p}) = H(\bar{p})$

We give a proof in §4.3. Since $H(\bar{p})$ is a fundamental quantity, we first discuss some of its properties.

4.2. Entropy, i.e., uncertainty

The entropy is a measure of the uncertainty of an outcome. For example, consider the case $d = 2$, with probabilities $p, 1 - p$. The graph of $H(p)$ is

graph here

Note that it has a unique maximum when $p = \frac{1}{2}$ (situation of maximal uncertainty) and is 0 in the two certain cases.

We arrived at H while approximating logs of multinomial coefficients. Say we had not yet discovered it but were looking for a function h on probability distributions with the following properties:

- (1) h is continuous in the p_i .
- (2) If $p_i = \frac{1}{d}$ for all i and we increase d , h is monotonically increasing.
- (3) h is additive with respect to breaking an event into a sequence of conditional events: Write $\mathcal{X} = A_1 \sqcup \dots \sqcup A_k$ and assume $p(A_j) > 0$ for all j . Write $\bar{p}_A = (p(A_1), \dots, p(A_k))$. Then $h(\bar{p}_{\mathcal{X}}) = h(\bar{p}_A) + \sum_{i=1}^k p(A_i)h(\bar{p}_{\mathcal{X}|A_i})$.

For example, if $\mathcal{X} = \{1, 2, 3\}$, with $p_1 = \frac{1}{2}$, $p_2 = \frac{1}{3}$, $p_3 = \frac{1}{6}$, we can choose one of the three at the outset, so we have $H(\frac{1}{2}, \frac{1}{3}, \frac{1}{6})$, or we can first decide between the sets $\{1\}$ and $\{2, 3\}$, both of which have probability $\frac{1}{2}$ and then if we choose the second, decide between 2 and 3, the first with probability $\frac{2}{3}$, the second with probability $\frac{1}{3}$. Thus we require the equality $H(\frac{1}{2}, \frac{1}{3}, \frac{1}{6}) = H(\frac{1}{2}, \frac{1}{2}) + \frac{1}{2}(0) + \frac{1}{2}H(\frac{2}{3}, \frac{1}{3})$.




Figure 4.2.1. a choice of three outcomes viewed as a choice of two followed by a second choice of two

Theorem 4.2.1. [Sha48] *The only function h satisfying 1,2,3 is, up to a constant, the entropy.*

Proof. Set $A(d) = h(\frac{1}{d}, \dots, \frac{1}{d})$. By 3, $A(s^m) = mA(s)$ (recall that $\log(s^m) = m \log(s)$).

To see this, consider the example $s = 2, m = 3$:

Exercise 4.2.2: Show that $A(s) = C \log(s)$ for some constant C .

Now say we have a choice of D equally likely outcomes, which we break up as $D = \sum_{i=1}^d d_i$. Write $p_i = \frac{d_i}{D}$, and assume the d_i are natural numbers.




Figure 4.2.2. repeated application of (3) gives $A((\frac{1}{2})^3) = A(\frac{1}{2}) + \frac{1}{2}A(\frac{1}{4}) + \frac{1}{2}A(\frac{1}{4}) = 3A(\frac{1}{2})$

By property (3), $C \log(D) = H(p_1, \dots, p_d) + C \sum_i p_i \log d_i$, i.e.,

$$\begin{aligned} H(p_1, \dots, p_d) &= -C \left[\sum_i p_i \log(d_i) - \log(D) \right] \\ &= -C \left[\sum_i p_i \log(d_i) - \sum_i p_i \log(D) \right] \\ &= -C \sum_i p_i \log\left(\frac{d_i}{D}\right) \\ &= -C \sum_i p_i \log(p_i). \end{aligned}$$

Finally use property 1 to extend by continuity to all probability distributions. \square

Here are some further properties of entropy that are easily verified:

- (1) $0 \leq H(\bar{p}) \leq \log(d)$ with $H(\bar{p}) = 0$ if and only if the outcome is certain, i.e., $\bar{p} = (0, \dots, 0, 1, 0, \dots, 0)$, and with $H(\bar{p}) = \log(d)$ if and only if $\bar{p} = (\frac{1}{d}, \dots, \frac{1}{d})$.
- (2) Say our space is $\mathcal{X} \times \mathcal{Y}$, so $H(\bar{p}_{\mathcal{X} \times \mathcal{Y}}) = - \sum_{i,j} p_{\mathcal{X} \times \mathcal{Y}}(i, j) \log p_{\mathcal{X} \times \mathcal{Y}}(i, j)$. Then we can recover $H(\bar{p}_{\mathcal{X}})$ as:

$$H(\bar{p}_{\mathcal{X}}) = - \sum_{i,j} p_{\mathcal{X} \times \mathcal{Y}}(i, j) \log\left(\sum_k p_{\mathcal{X} \times \mathcal{Y}}(i, k)\right).$$

In particular, $H(\bar{p}_{\mathcal{X} \times \mathcal{Y}}) \leq H(\bar{p}_{\mathcal{X}}) + H(\bar{p}_{\mathcal{Y}})$, with equality if and only if \mathcal{X}, \mathcal{Y} are independent, i.e., $p_{\mathcal{X} \times \mathcal{Y}}(i, j) = p_{\mathcal{X}}(i)p_{\mathcal{Y}}(j)$. (The uncertainty of $p_{\mathcal{X} \times \mathcal{Y}}$ is at most the sum of the uncertainties of the marginal distributions $p_{\mathcal{X}}$ and $p_{\mathcal{Y}}$.)

- (3) A modification of the p_i towards equalization increases H . More precisely:

Exercise 4.2.3: Let A be a doubly stochastic matrix, i.e., the entries of A are non-negative and the column and row sums of A

are one. Show that $H(A\bar{p}) \geq H(\bar{p})$ and unless A is a permutation matrix, that there exists \bar{p} where the inequality is strict.

4.3. Shannon's noiseless channel theorem

Theorem 4.3.1. [Sha48] *Given $\epsilon > 0, \delta > 0$, there exists n_0 such that for all $n \geq n_0$, there is a decomposition $\mathcal{X}^{\times n} = \mathcal{X}_{\epsilon\text{-typ}}^n \sqcup \mathcal{X}_{\delta\text{-small}}^n$ where*

- (1) $\Pr(\mathcal{X}_{\delta\text{-small}}^n) < \delta$, and
- (2) $\forall \bar{x} \in \mathcal{X}_{\epsilon\text{-typ}}^n$,

$$\left| \frac{1}{n} \log(\Pr(\bar{x})) - H(\bar{p}) \right| < \epsilon,$$

(3)

$$(1 - \delta)2^{n(H(\bar{p}) - \epsilon)} \leq |\mathcal{X}_{\epsilon\text{-typ}}^n| \leq 2^{n(H(\bar{p}) - \epsilon)}.$$

The set $\mathcal{X}_{\epsilon\text{-typ}}^n$ will play the role of $\mathcal{Y}(n)$ from §4.1. Informally, the probability of not being ϵ -typical is small, if ϵ -typical, the probability is close to the expectation, and if the entropy is large, most sequences are ϵ -typical, and if it is small, there are few such.

Proof. The strong law of large numbers (1.3.3) means that given iid random variables X_j , for all $\epsilon, \delta > 0$, there exists n_0 such that for all $n \geq n_0$ such that

$$P\left(\left|\frac{X_1 + \cdots + X_n}{n} - E[X]\right| > \epsilon\right) < \delta$$

Let $\mathcal{X}_{\delta\text{-small}}^n$ be all events $\bar{x} = (x_1, \dots, x_n)$ where $|\frac{1}{n}(x_1 + \cdots + x_n) - E[X]| \geq \epsilon$, and $\mathcal{X}_{\epsilon\text{-typ}}^n$ the events \bar{x} with $|\frac{1}{n}(x_1 + \cdots + x_n) - E[X]| < \epsilon$. In our case $X = -\log(p(\bar{x}))$ and $E[-\log(p(\bar{x}))] = H(\bar{p})$. Note that (3) is a quantitative version of (4.1.1). To prove it, note that the second condition implies that for all $\bar{x} \in \mathcal{X}_{\epsilon\text{-typ}}^n$, and $n > \frac{1}{\epsilon}$,

$$(4.3.1) \quad \Pr(\bar{x}) \leq 2^{-n(H(\bar{p}) - \epsilon)}.$$

On the other hand

$$1 - \delta \leq \Pr(\mathcal{X}_{\epsilon\text{-typ}}^n) \leq 1$$

i.e.,

$$1 - \delta \leq \sum_{\bar{x} \in \mathcal{X}_{\epsilon\text{-typ}}^n} \Pr(\bar{x}) \leq 1$$

so plugging in (4.3.1) we conclude. □

It will be useful to give two variants of Theorem 4.3.1.

Theorem 4.3.2. [Sha48] Fix $q \in (0, 1)$. Let $\text{num}(q)$ denote the minimum cardinality of a subset S of $\mathcal{X}^{\times n}$ such that $\Pr(S) \geq q$. Then

$$\lim_{n \rightarrow \infty} \frac{\log(\text{num}(q))}{n} = H(\bar{p}).$$

Note that the right hand side is independent of q , thus for large n , the ratio is nearly independent of q and n . Theorem 4.3.2 will be a consequence of:

Theorem 4.3.3. [Sha48] Let $R < H(\bar{p})$ and let $S(n) \subset \mathcal{X}^{\times n}$ be a sequence of subsets with $|S(n)| \leq 2^{nR}$, i.e., $\frac{1}{n} \log |S(n)| \leq R$. Then for any $\eta > 0$, there exists n_0 such that for all $n > n_0$, $\Pr(S(n)) < \eta$.

In other words, *any* subset of size less than capacity can only accumulate a small amount of the probability.

Proof of Theorem 4.3.3. Write $S(n) = S(n)_{\delta\text{-small}} \sqcup S(n)_{\epsilon\text{-typ}}$, where $S(n)_{\delta\text{-small}} = S(n) \cap \mathcal{X}_{\delta\text{-small}}^n$ and $S(n)_{\epsilon\text{-typ}} = S(n) \cap \mathcal{X}_{\epsilon\text{-typ}}^n$. Then $|S(n)_{\epsilon\text{-typ}}| < |S(n)| < 2^{nR}$. By (4.3.1) each element of $S(n)_{\epsilon\text{-typ}}$ has probability at most $2^{-n(H(\bar{p})-\epsilon)}$, so $\Pr(S(n)_{\epsilon\text{-typ}}) \leq 2^{-n(H(\bar{p})-\epsilon)} 2^{nR}$. So just take, $\epsilon < H(\bar{p}) - R$, and n_0, δ such that $2^{-n_0(H(\bar{p})-R-\epsilon)} + \delta < \eta$. \square

Theorem 4.3.4. [Sha48] Let a source \mathcal{X} have entropy $H(\bar{p})$ (bits per symbol) and let a channel have capacity C (bits per second). Then for any $\epsilon > 0$, it is possible to encode the output of the source to transmit at the rate $\frac{C}{H(\bar{p})} - \epsilon$ symbols/sec., and it is not possible to reliably transmit at an average rate greater than $\frac{C}{H(\bar{p})}$.

The idea of the proof is clear: just transmit all the ϵ -typical sequences and discard the others.

4.4. Transmission over noisy channels

Say we transmit symbols x and receive symbols y over a channel subject to noise, so we may or may not have $y = x$. Intuitively, if the noise is small, with some redundancy we should be able to communicate accurate messages most of the time. Let Rate denote the maximal possible rate of transmission. In a noiseless channel this is just $H(p_{\mathcal{X}})$, but now we must subtract off something to account for the uncertainty that, upon receiving y , that it was the signal sent. This something will be the *conditional entropy*, $H(\bar{p}_{\mathcal{X}}|\bar{p}_{\mathcal{Y}})$ defined below. The punch line will be:

Given a channel with noise and symbols sent according to $p_{\mathcal{X}}$, and noise given by $p_{\mathcal{Y}}$, the maximum rate of transmission is $H(\bar{p}_{\mathcal{X}}) - H(\bar{p}_{\mathcal{X}}|\bar{p}_{\mathcal{Y}})$.

Given a channel, with noise its maximal capacity for transmission is $\max_{q_{\mathcal{X}}} (H(q_{\mathcal{X}}) - H(q_{\mathcal{X}}|p_{\mathcal{Y}}))$.

Note that the second result is trivial in the noiseless case - one just takes a uniform distribution for the symbols. What is going on here is that if some symbols are more susceptible to corruption than others, the uniform distribution will no longer be optimal.

4.4.1. Conditional entropy. Recall the conditional probability of i occurring given knowledge that j occurs (assuming $\Pr(j) > 0$): $\Pr_{\mathcal{X}|\mathcal{Y}}(i|j) = \frac{\Pr_{\mathcal{X},\mathcal{Y}}(i,j)}{\Pr_{\mathcal{Y}}(j)}$ (also recall $\Pr_{\mathcal{Y}}(j) = \sum_i \Pr_{\mathcal{X},\mathcal{Y}}(i,j)$). Define the conditional entropy

$$H(\bar{p}_{\mathcal{Y}}|\bar{p}_{\mathcal{X}}) := - \sum_{i,j} \Pr_{\mathcal{X},\mathcal{Y}}(i,j) \log \Pr_{\mathcal{Y}|\mathcal{X}}(j|i).$$

Note that

$$(4.4.1) \quad H(\bar{p}_{\mathcal{Y}}|\bar{p}_{\mathcal{X}}) = H(\bar{p}_{\mathcal{X},\mathcal{Y}}) - H(\bar{p}_{\mathcal{X}})$$

or equivalently $H(\bar{p}_{\mathcal{X},\mathcal{Y}}) = H(\bar{p}_{\mathcal{X}}) + H(\bar{p}_{\mathcal{Y}}|\bar{p}_{\mathcal{X}})$, the uncertainty of $p_{\mathcal{X},\mathcal{Y}}$ is the uncertainty of $p_{\mathcal{X}}$ plus the uncertainty of $p_{\mathcal{Y}}$ given $p_{\mathcal{X}}$.

In particular, we have $H(\bar{p}_{\mathcal{X}}) + H(\bar{p}_{\mathcal{Y}}) \geq H(\bar{p}_{\mathcal{X},\mathcal{Y}}) = H(\bar{p}_{\mathcal{X}}) + H(\bar{p}_{\mathcal{Y}}|\bar{p}_{\mathcal{X}})$, i.e.,

$$H(\bar{p}_{\mathcal{Y}}) \geq H(\bar{p}_{\mathcal{Y}}|\bar{p}_{\mathcal{X}}),$$

i.e., with extra knowledge, our uncertainty about $p_{\mathcal{Y}}$ cannot increase, and decreases unless $p_{\mathcal{X}}$ and $p_{\mathcal{Y}}$ are independent).

4.4.2. Examples. We first give an example where $p_{\mathcal{X}}$ is fixed. Say 0's and 1's are transmitted with each having a probability of error .01, so $H(\bar{p}_{\mathcal{X}}|\bar{p}_{\mathcal{Y}}) = -[.99 \log(.99) + .01 \log(.01)] \sim .81$ bits/symbol, and we can transmit 1000 bits/second.

Exercise 4.4.1: Verify the above assertion about $H(\bar{p}_{\mathcal{X}}|\bar{p}_{\mathcal{Y}})$.

The above discussion predicts a transmission rate of $1000 - 81 = 919$ bits/second. If the probability of error goes up to $\frac{1}{2}$, then $H(\bar{p}_{\mathcal{X}}|\bar{p}_{\mathcal{Y}}) = 1$ and our discussion predicts a rate of $1000 - 1000 = 0$ bits/second, which agrees with our intuition that what is received is just noise.

Here is an example where we choose $p_{\mathcal{X}}$ to optimize capacity. Now say $\mathcal{X} = \{a, b, c\}$ where the symbol a is never effected by noise, b has probability p being transmitted correctly and $1 - p$ of being flipped to c , and c has probability p being transmitted correctly and $1 - p$ of being flipped to b .

Let p_a be the probability that a is transmitted, let p_b be the probability that b is transmitted and p_c the probability that c is transmitted: we get to choose these. Given the symmetry of the situation, we should set $p_b = p_c =: p_{b,c}$.

Thus

$$\begin{aligned} H(\bar{p}_{\mathcal{X}}) &= -p_a \log p_a - 2p_{b,c} \log p_{b,c} \\ H(\bar{p}_{\mathcal{X}}|\bar{p}_{\mathcal{Y}}) &= 2p_{b,c}H(p, 1-p) \end{aligned}$$

We want to maximize $H(\bar{p}_{\mathcal{X}}) - H(\bar{p}_{\mathcal{X}}|\bar{p}_{\mathcal{Y}})$ by a good choice of $p_a, p_{b,c}$, subject to the constraint that $p_a + 2p_{b,c} = 1$. We have

$$H(\bar{p}_{\mathcal{X}}) - H(\bar{p}_{\mathcal{X}}|\bar{p}_{\mathcal{Y}}) = -p_a \log p_a - 2p_{b,c} \log p_{b,c} - 2p_{b,c}H(p, 1-p).$$

Viewed as a function of $p_a, p_{b,c}$, we have a standard maximization problem from calculus.

Exercise 4.4.2: Differentiating and imposing the constraint, show that it is optimal to take

$$(4.4.2) \quad p_a = \frac{e^{H(p,1-p)}}{e^{H(p,1-p)} + 2}, \quad p_{b,c} = \frac{1}{e^{H(p,1-p)} + 2}.$$

Using (4.4.2), we obtain

$$H(\bar{p}_{\mathcal{X}}) - H(\bar{p}_{\mathcal{X}}|\bar{p}_{\mathcal{Y}}) = \log \frac{e^{H(p,1-p)} + 2}{e^{H(p,1-p)}}.$$

For example, when $p = 1$, i.e., no errors, we obtain $\log(3)$, as we have a noiseless channel. When $p = \frac{1}{2}$, we obtain $\log(2)$, as the second and third symbols are indistinguishable, so we essentially have two channels. In general, the capacity is between these two, and the first channel used somewhere between the same amount and twice as often as the other two channels.

4.4.3. Warm up: Communication with the help of a correction channel. Before giving the proof of Shannon's theorem, here is a warm-up problem giving intuition into the conditional entropy $H(\bar{p}_{\mathcal{X}}|\bar{p}_{\mathcal{Y}})$, which Shannon calls the "equivocation", as the amount of extra information that must be supplied to correct errors on a noisy channel.

Consider the following picture ***

In the transmission from the source to the receiver, an observer is allowed to see what is transmitted and what the receiver gets. The observer is then allowed to send correction data to allow the receiver to correct the message. The question is how much information must the observer send to the receiver to enable correction, i.e., how much capacity does the correction channel need to correct errors (assume the correction channel is not subject to noise).

Theorem 4.4.3. [Sha48] *If $\text{Cap}(\text{correction channel}) \geq H(\bar{p}_{\mathcal{X}}|\bar{p}_{\mathcal{Y}})$, then in the scheme above, all but an arbitrarily small fraction of the errors can be corrected.*

If $\text{Cap}(\text{correction channel}) < H(\bar{p}_X|\bar{p}_Y)$, then in the scheme above, reliable error correction is not possible.

Proof. Assume we are in the first case, say \bar{y}^n is received when \bar{x}^n is sent over t seconds. By the same argument as in Theorem 4.3.1(3) applied to $\#\{\bar{x}^n \mid \Pr_{\mathcal{X}|Y}(\bar{x}^n, \bar{y}^n) \geq 1 - \delta\}$, there exist on the order of $2^{tH(\bar{p}_X|\bar{p}_Y)}$ possible \bar{x}^n 's that could have reasonably produced \bar{y}^n . Thus we need to send $tH(\bar{p}_X|\bar{p}_Y)$ bits each t seconds, which is possible with an ϵ -tolerance of error on a channel of capacity $H(\bar{p}_X|\bar{p}_Y)$.

To prepare for the second case:

Exercise 4.4.4: For any random variables x, y, z , determining probability distributions p_X, p_Y, p_Z , show that

$$H(\bar{p}_X|\bar{p}_Y, \bar{p}_Z) \geq H(\bar{p}_X|\bar{p}_Y) - H(\bar{p}_Z|\bar{p}_Y) \geq H(\bar{p}_X|\bar{p}_Y) - H(\bar{p}_Z).$$

Now let x be the output of the source, y the received signal, and z the signal sent over the correction channel. We see $H(\bar{p}_X|\bar{p}_Y, \bar{p}_Z) > 0$ so we cannot recover x reliably by the noiseless theorem. \square

Note that we have three interpretations of the rate:

$$\begin{aligned} \text{Rate} &= H(\bar{p}_X) - H(\bar{p}_X|\bar{p}_Y) \text{ the information sent minus the uncertainty of what sent} \\ &= H(\bar{p}_Y) - H(\bar{p}_Y|\bar{p}_X) \text{ the information received minus the noise} \\ &= H(\bar{p}_X) + H(\bar{p}_Y) - H(\bar{p}_{X \times Y}) \text{ the sum of the information sent and received minus the} \\ &\quad \text{joint entropy, essentially the bits/sec. common to } x, y \end{aligned}$$

The Rate is also called the *mutual information*.

4.4.4. Capacity of a noisy channel. Define the *capacity* of a noisy channel to be the maximum rate over all possible probability distributions on the source:

$$\text{Cap} := \max_{q_X} (H(q_X) - H(q_X|p_Y)).$$

Theorem 4.4.5. [Sha48] Let a discrete channel have capacity Cap and entropy per second H . If $H < \text{Cap}$, then there exists an encoding \bar{p} of the source such that information can be transmitted over the channel with an arbitrarily small frequency of errors $H(\bar{p}_X|\bar{p}_Y)$ ***check***. If $H > \text{Cap}$, then there exists an encoding \bar{p} such that the equivocation is less than $H - \text{Cap} + \epsilon$ for any $\epsilon > 0$, and there does not exist any \bar{p} with equivocation less than $H - \text{Cap}$.

The basic idea is the same as the noiseless case, however there is a novel feature that now occurs frequently in complexity theory arguments - that instead of producing an algorithm to find the efficient encoding, Shannon

showed that a random choice of encoding will work. More on this after the proof.

Proof. Split the transmitter and receiver $\mathcal{X}^{\times n}$ and $\mathcal{Y}^{\times n}$ into the union of ϵ -typical δ -small subsets. For a high probability message \bar{y}^n in $\mathcal{Y}_{\epsilon\text{-typ}}^n$, there are roughly $2^{H(\bar{p}_X|\bar{p}_Y)t}$ “reasonable” \bar{x}^n 's, i.e., elements of $\mathcal{X}_{\epsilon\text{-typ}}^n$, that could have been sent for \bar{y}^n to be received, i.e., roughly $2^{H(\bar{p}_X|\bar{p}_Y)t}$ elements of $\mathcal{X}_{\epsilon\text{-typ}}^n$. On the other hand, if some $\bar{x}^n \in \mathcal{X}_{\epsilon\text{-typ}}^n$ is sent, there are about $2^{H(\bar{p}_Y|\bar{p}_X)t}$ elements of $\mathcal{Y}_{\epsilon\text{-typ}}^n$ that could be received.

pic*

Every t seconds we have 2^{tR} high probability messages. Say \bar{y}^n is observed, we want to know the probability that more than one message in $\mathcal{X}_{\epsilon\text{-typ}}^n$ could arrive as \bar{y}^n , based on our choice of distribution.

Take a *random* encoding of $\mathcal{X}_{\epsilon\text{-typ}}^n$, so we have 2^{tR} messages distributed at random among $2^{tH(p_X)}$ points. The probability of a particular message being received as \bar{y}^n is $\frac{2^{tR}}{2^{tH(p_X)}} = 2^{t(R-H(p_X))}$. The probability that no $\bar{x}^n \in \mathcal{X}_{\epsilon\text{-typ}}^n$ (other than \bar{y}^n) is sent to \bar{y}^n is

$$(4.4.3) \quad [1 - 2^{t(R-H(p_X))}]^{2^{tH(\bar{p}_X|\bar{p}_Y)}}.$$

Now say $R < \text{Cap}$, and write $R - H(p_X) = -H(\bar{p}_X|\bar{p}_Y) - \eta$ for some $\eta > 0$, so (4.4.3) becomes

$$[1 - 2^{-tH(\bar{p}_X|\bar{p}_Y) - t\eta}]^{2^{tH(\bar{p}_X|\bar{p}_Y)}}.$$

This limits to 1 as $t \rightarrow \infty$.

To prove the second assertion, just send Cap bits/sec. of x 's generated and throw away the rest. This gives $H(\bar{p}_X|\bar{p}_Y)$ equal to $H(p_X) - \text{Cap}$ plus the ϵ from the first case. □

Exercise 4.4.6: Verify that, for $H \geq 0$, $\eta > 0$, $\lim_{t \rightarrow \infty} [1 - 2^{-tH - t\eta}]^{2^{tH}} = 1$.

⊙

Note that the phrase “Take a random encoding” is not constructive, as it gives no recipe how to do so.

After presenting the proof, Shannon remarks: “An attempt to obtain a good approximation to ideal coding by following the method of the proof is generally impractical. ... Probably this is no accident but is related to the difficulty of giving an explicit construction for a good approximation to a random sequence”. To our knowledge, this is the first time that the difficulty of “finding hay in a haystack” (phrase due to Howard Karloff) is mentioned in print. This problem is central to complexity: for example,

Valiant's algebraic version of $\mathbf{P} \neq \mathbf{NP}$ can be phrased as the problem of finding a sequence of explicit polynomials that are difficult to compute, while it is known that a random sequence is indeed difficult to compute. (According to A. Wigderson, the difficulty of writing down random objects was explicitly discussed by Erdős, in the context of random graphs, at least as early as 1947, in relation to his seminar paper [Erd47]. This paper, along with [Sha48] gave rise to the now ubiquitous probabilistic method in complexity theory.

Quantum information

In this chapter we give the rudiments of quantum information theory. We first, in §5.1, give a convenient reformulation of the postulates of quantum mechanics in terms of density operators. In order to discuss when one density operator is “close” to another, we need to discuss distance functions on the space of density operators, which is done in §5.2. We present a quantum version of Shannon’s noiseless channel theorem in §5.3. This requires the introduction of von Neumann entropy. Properties of von Neumann entropy are discussed in §5.4. We conclude in §5.5, with a discussion of measures of entanglement.

5.1. Reformulation of quantum mechanics

We discuss two inconveniences about our formulation of the postulates of quantum mechanics, leading to a formulation of the postulates in terms of density operators.

5.1.1. Partial measurements. Before, we defined a measurement of a state $|\psi\rangle = \sum z_I |I\rangle$ as a procedure that gives us $I = (i_1, \dots, i_n) \in \{0, 1\}^n$ with probability $|z_I|^2$. But in our algorithms, this is not what we did: we were working not in $(\mathbb{C}^2)^{\otimes n}$, but $(\mathbb{C}^2)^{\otimes n+m}$ where there were m “workspace” qubits we were not interested in measuring. So our measurements were the projections onto the *spaces* $|I\rangle \otimes (\mathbb{C}^2)^{\otimes m}$. We define a generalized notion of measurement that allows for projection onto spaces.

To make the transition, first rewrite

$$\begin{aligned} |z_I|^2 &= |\langle \psi | I \rangle|^2 \\ &= \langle \psi | I \rangle \langle I | \psi \rangle \\ &= \langle \psi | \text{Proj}_I | \psi \rangle, \end{aligned}$$

where $\text{Proj}_I : (\mathbb{C}^2)^{\otimes n} \rightarrow \mathbb{C}|I\rangle$ is the orthogonal projection onto the line spanned by $|I\rangle$.

Now say we are only interested in the first n qubits of a system of $n+m$ qubits, and we want to know the probability a measurement gives rise to some I represented by a vector $|I\rangle \in (\mathbb{C}^2)^{\otimes n}$, but we have $|\psi\rangle \in (\mathbb{C}^2)^{\otimes n+m}$. Then the probability of obtaining $|I\rangle$ given $|\psi\rangle$ is

$$\begin{aligned} \Pr(|I\rangle \mid |\psi\rangle) &= \sum_{J \in \{0,1\}^m} \Pr(|\psi\rangle, |IJ\rangle) \\ &= \sum_J \langle \psi | IJ \rangle \langle IJ | \psi \rangle \\ &= \langle \psi | (|I\rangle \langle I| \otimes \text{Id}_{(\mathbb{C}^2)^{\otimes m}}) | \psi \rangle \\ &= \langle \psi | \text{Proj}_{\mathcal{M}} | \psi \rangle \end{aligned}$$

where $\text{Proj}_{\mathcal{M}} : (\mathbb{C}^2)^{\otimes n+m} \rightarrow |I\rangle \otimes (\mathbb{C}^2)^{\otimes m} =: \mathcal{M}$ is the orthogonal projection operator. Then $\Pr(|I\rangle \mid |\psi\rangle) = \langle \psi | \text{Proj}_{\mathcal{M}} | \psi \rangle$. With this definition, we can allow $\mathcal{M} \subset \mathcal{H}$ to be *any* linear subspace, which will simplify our measurements. (Earlier, if we wanted to measure the probability of a non-basis state, we had to change bases before measuring.)

One may think of projection operators as representing outside interference of a quantum system, like adding a filter to beams being sent that destroy states not in \mathcal{M} .

The following chart from [KSV02] illustrates the comparison between the classical (ℓ_1) and quantum (ℓ_2) situations so far:

	Classical (ℓ_1) probability	Quantum (ℓ_2) probability
Event	subset $M \subset \mathcal{X}$	subspace $\mathcal{M} \subset \mathcal{H}$
Probability distribution	$p : \mathcal{X} \rightarrow \mathbb{R}_{\geq 0}$ with $ p _1 = 1$	$ \psi\rangle \in \mathcal{H}$ with $ \psi _2 = 1$
Probability	$\Pr(M) = \sum_{j \in M} p_j$	$\Pr(\mathcal{M}) = \langle \psi \text{Proj}_{\mathcal{M}} \psi \rangle$

Recall that in classical probability, one has the identity:

$$(5.1.1) \quad \Pr(M_1 \cup M_2) = \Pr(M_1) + \Pr(M_2) - \Pr(M_1 \cap M_2).$$

The quantum analog is *false* in general: Let $\mathcal{H} = \mathbb{C}^2$, $\mathcal{M}_1 = \mathbb{C}|0\rangle$ and $\mathcal{M}_2 = \mathbb{C}(|0\rangle + |1\rangle)$. Let $|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$ with $|\alpha|^2 + |\beta|^2 = 1$.

Exercise 5.1.1: Show that $\Pr(\text{span}\{\mathcal{M}_1, \mathcal{M}_2\}) \neq \Pr(\mathcal{M}_1) + \Pr(\mathcal{M}_2) - \Pr(\mathcal{M}_1 \cap \mathcal{M}_2)$. \odot

However, we can recover (5.1.1) if the projection operators commute:

Proposition 5.1.2. *If $\text{Proj}_{\mathcal{M}_1} \text{Proj}_{\mathcal{M}_2} = \text{Proj}_{\mathcal{M}_2} \text{Proj}_{\mathcal{M}_1}$ then $\Pr(\text{span}\{\mathcal{M}_1, \mathcal{M}_2\}) = \Pr(\mathcal{M}_1) + \Pr(\mathcal{M}_2) - \Pr(\mathcal{M}_1 \cap \mathcal{M}_2)$.*

In particular, if $\mathcal{M}_1 \perp \mathcal{M}_2$, so $\text{Proj}_{\mathcal{M}_1} \text{Proj}_{\mathcal{M}_2} = \text{Proj}_{\mathcal{M}_2} \text{Proj}_{\mathcal{M}_1} = 0$, then $\Pr(\text{span}\{\mathcal{M}_1, \mathcal{M}_2\}) = \Pr(\mathcal{M}_1) + \Pr(\mathcal{M}_2)$.

Exercise 5.1.3: Prove Proposition 5.1.2.

5.1.2. Mixing classical and quantum probability. A typical situation in probability is as follows: you want a cookie, but can't make up your mind which kind, so you decide to take one at random from the cookie jar to eat. However when you open the cupboard, you find there are two different cookie jars H and T , each with a different distribution of cookies, say P_H and P_T . You decide to flip a coin to decide which jar and say your coin is biased with probability p for heads (choice H). The resulting probability distribution is

$$pP_H + (1 - p)P_T.$$

Let's encode this scenario with vectors. Classically, if vectors corresponding to P_H, P_T are respectively v_H, v_T , the new vector is $pv_H + (1 - p)v_T$. The probability of drawing a chocolate chip (CC) cookie is $pP_H(CC) + (1 - p)P_T(CC) = pv_{H,CC} + (1 - p)v_{T,CC}$.

But what should we take in quantum probability, where we use the ℓ_2 norm instead of the ℓ_1 norm? Given $|\psi_A\rangle = \sum z_I |I\rangle, |\psi_B\rangle = \sum w_J |J\rangle$, we want to make a measurement that gives us $p|z_{CC}|^2 + (1 - p)|w_{CC}|^2$. Unfortunately $|pz_{CC} + (1 - p)w_{CC}|^2 \neq p|z_{CC}|^2 + (1 - p)|w_{CC}|^2$ in general. To fix this problem we will enlarge the notion of state and further modify our notion of measurement.

Our problem comes from having a mixture of ℓ_1 and ℓ_2 norms. Our fix will be to rewrite $|\psi\rangle$ in a way that the ℓ_2 norm becomes an ℓ_1 norm. That is, we construct an object that naturally contains the squares of the norms of the coefficients of $|\psi_A\rangle$. Consider the endomorphism $|\psi_A\rangle\langle\psi_A| = \sum_{I,J} z_I \bar{z}_J |I\rangle\langle J|$. It is rank one, and its diagonal entries are the quantities we want.

To measure them, let Proj_J denote the projection onto the J -th coordinate. Then

$$\text{trace}(\text{Proj}_J |\psi_A\rangle\langle\psi_A|) = |z_{A,J}|^2$$

is the desired quantity.

Now back to our cookie jars, set

$$\rho = p|\psi_A\rangle\langle\psi_A| + (1-p)|\psi_B\rangle\langle\psi_B|$$

and observe that

$$\text{trace}(\text{Proj}_J \rho) = p|z_{A,J}|^2 + (1-p)|z_{B,J}|^2$$

as desired.

Given a finite set of states $\{|\psi_1\rangle, \dots, |\psi_s\rangle\}$, with $\Pr(|\psi_i\rangle) = p_i$, and $\sum_i p_i = 1$, set $\rho = \sum_k p_k |\psi_k\rangle\langle\psi_k| \in \text{End}(\mathcal{H})$. Then our probability of measuring our state in \mathcal{M} is $\text{trace}(\text{Proj}_{\mathcal{M}} \rho)$. Note that ρ has the properties

- (1) $\rho = \rho^\dagger$, i.e., ρ is Hermitian,
- (2) $\forall |\eta\rangle, \langle\eta|\rho|\eta\rangle \geq 0$, i.e., ρ is *positive*,
- (3) $\text{trace}(\rho) = 1$.

This motivates the following definition:

Definition 5.1.4. An operator $\rho \in \text{End}(\mathcal{H})$ satisfying 1,2,3 above is called a *density operator*.

Exercise 5.1.5: Show that property (2) implies property (1). \odot

Density operators will replace states in our quantum notion of probability. Note that a density operator that is diagonal in the standard basis of \mathbb{C}^d corresponds to a probability distribution on $\{1, \dots, d\}$, so the definition includes classical probability as well as our old notion of state (which are the rank one density operators).

Exercise 5.1.6: Show that the set of density operators is invariant under the induced action of $\mathbf{U}(\mathcal{H})$ on $\text{End}(\mathcal{H})$.

Different scenarios can lead to the same density operator. However, two states with the same density operator are physically indistinguishable.

5.1.3. Reformulation of the postulates of quantum mechanics.

Postulate 1. Associated to any isolated physical system is a Hilbert space \mathcal{H} , called the *state space*. The system is described by a density operator $\rho \in \text{End}(\mathcal{H})$.

Remark 5.1.7. This postulate explains the $d = n^2$ in the Hardy description from §2.1.2, as the (real) dimension of the space of Hermitian operators is n^2 .

Postulate 2. The evolution of an isolated system is described by the action of unitary operators on ρ .

It will be convenient to have two formulations of the third postulate (even more appear in the literature).

Postulate 3. (Projective) Measurements correspond to positive trace one operators $X \in \text{End } \mathcal{H}$. Let $X = \sum_j \lambda_j \text{Proj}_{\mathcal{M}_j}$ be its eigenspace decomposition. The probability that ρ is in measured in state \mathcal{M}_j is $\lambda_j \text{trace}(\text{Proj}_{\mathcal{M}_j} \rho)$.

Note that X is the same mathematical object as a density operator, but its role here is completely different. **Above is different from NC, because what they have does not make sense mathematically, at least to me – need to check** Such X are called *observables*.

Other than for the uncertainty principle, we will generally only need the following types of measurements:

Postulate 3'. (POVM) Measurements correspond to a collection of projection operators $\text{Proj}_{\mathcal{M}_j}$ such that $\sum_k \text{Proj}_{\mathcal{M}_k} = \text{Id}_{\mathcal{H}}$. The probability that ρ is in measured in state \mathcal{M}_j is $\text{trace}(\text{Proj}_{\mathcal{M}_j} \rho)$.

(Some texts use even more general measurement operators, but they again give rise to an equivalent theory.) POVM stands for *positive operator valued measure*, but since it is the primary measurement we will deal with, we generally omit the POVM.

Postulate 4 is unchanged:

Postulate 4. The state of a composite physical system, that is a physical system arising from physical systems, is the tensor product of the Hilbert spaces: $\mathcal{H}_{AB} = \mathcal{H}_A \otimes \mathcal{H}_B$.

Exercise 5.1.8: **Add exercises on teleportation and super-dense coding from this perspective following Christandl

5.1.4. Remarks on composite systems. Just as in classical probability one is interested in marginal distributions, in the quantum setting we have already seen several instances where one takes “marginals”. We now formalize this.

We first give an interpretation of the map $\text{trace} : \text{End}(\mathcal{H}) \rightarrow \mathbb{C}$. Since $\text{End}(\mathcal{H}) = \mathcal{H} \otimes \mathcal{H}^*$, we may view trace as an element of $(\mathcal{H} \otimes \mathcal{H}^*)^* = \mathcal{H}^* \otimes \mathcal{H} = \text{End}(\mathcal{H}^*) \simeq \text{End}(\mathcal{H})$. It is some map $\mathcal{H} \rightarrow \mathcal{H}$ that is invariant under the action of $GL(\mathcal{H})$. There is a unique such up to scale, namely the identity map, and checking the trace of any nonzero linear map, we see as an endomorphism $\text{trace} = \text{Id}_{\mathcal{H}}$. To give yet another perspective, for $X \in \text{End } \mathcal{H} = \mathcal{H}^* \otimes \mathcal{H}$, $\text{trace}(X)$ is the image of X under the contraction map $\mathcal{H}^* \otimes \mathcal{H} \rightarrow \mathbb{C}$, $\langle v | \otimes | w \rangle \mapsto \langle v | w \rangle$.

For $Y \in \text{End}(\mathcal{H}_1 \otimes \mathcal{H}_2) = (\mathcal{H}_1^* \otimes \mathcal{H}_2^*) \otimes (\mathcal{H}_1 \otimes \mathcal{H}_2)$, define the partial trace $\text{trace}_{\mathcal{H}_1}(Y)$ to be the image of Y under the contraction $\mathcal{H}_1^* \otimes \mathcal{H}_2^* \otimes \mathcal{H}_1 \otimes \mathcal{H}_2 \rightarrow \mathcal{H}_2^* \otimes \mathcal{H}_2$ given by $\langle \phi | \otimes \langle \psi | \otimes | v \rangle \otimes | w \rangle \mapsto \langle \phi | v \rangle \langle \psi | \otimes | w \rangle = \langle \phi | v \rangle | w \rangle \langle \psi |$.

Compare the classical and quantum situations:

Classical probability	Quantum probability
for $p = p_1 \times p_2$ on $\mathcal{X}_1 \times \mathcal{X}_2$, and $M_j \subset \mathcal{X}_j$ $\Pr(M_1 \times M_2) = \Pr(M_1) \Pr(M_2)$	for $\rho = \rho_1 \otimes \rho_2$ on $\mathcal{H}_1 \otimes \mathcal{H}_2$, and $\mathcal{M}_j \subset \mathcal{H}_j$ $\text{trace}(\rho_1 \otimes \rho_2 \text{Proj}_{\mathcal{M}_1 \otimes \mathcal{M}_2})$ $= \text{trace}(\rho_1 \text{Proj}_{\mathcal{M}_1}) \text{trace}(\rho_2 \text{Proj}_{\mathcal{M}_2})$
for p on $\mathcal{X}_1 \times \mathcal{X}_2$ general and $M = M_1 \times \mathcal{X}_2$ $\Pr(M) = \sum_{j \in \mathcal{X}_2} \Pr(M_1, j)$	$\mathcal{M} = \mathcal{M}_1 \otimes \mathcal{H}_2$ $\text{trace}_{\mathcal{H}_1 \otimes \mathcal{H}_2}(\rho \text{Proj}_{\mathcal{M}})$ $= \text{trace}_{\mathcal{H}_1}(\text{trace}_{\mathcal{H}_2}(\rho) \text{Proj}_{\mathcal{M}_1})$

Exercise 5.1.9: Verify the above properties of quantum probability.

5.1.5. Expectation and the uncertainty principle. Allowing for non-commuting measurements has dramatic consequences.

Let $A \in \text{End}(\mathcal{H})$ be a Hermitian operator with eigenvalues $\lambda_1, \dots, \lambda_k$ and eigenspaces \mathcal{M}_j . If our system is in state ρ , we can consider A as a random variable that takes the value λ_j with probability $\text{trace}(\text{Proj}_{\mathcal{M}_j} \rho)$.

Recall the expectation of a random variable $X : \mathcal{X} \rightarrow \mathbb{R}$ is $E[X] := \sum_{j \in \mathcal{X}} X(j) \Pr(j)$.

Proposition 5.1.10. *If a system is in state ρ , the expectation of a Hermitian operator $A \in \text{End}(\mathcal{H})$ is $\text{trace}(A\rho)$.*

Proof.

$$\begin{aligned}
 E[A] &= \sum_j \lambda_j \text{trace}(\text{Proj}_{\mathcal{M}_j} \rho) \\
 &= \text{trace}\left(\left(\sum_j \lambda_j \text{Proj}_{\mathcal{M}_j}\right)\rho\right) \\
 &= \text{trace}(A\rho).
 \end{aligned}$$

□

One way mathematicians describe the famous Heisenberg uncertainty principle is that it is impossible to localize both a function and its Fourier transform. Another interpretation comes from probability:

First note that given a random variable, or Hermitian operator X (and a system in state ρ), we can replace it with an operator of mean zero $\hat{X} := X - E[X] \text{Id}$. For notational convenience, we state the uncertainty principle for such shifted operators.

Recall from (1.3.4) that the standard deviation $\sigma(X) = \sqrt{\text{var}(X)}$ of a random variable X is a measure of the failure of the corresponding probability distribution to be concentrated at a point, i.e., failure of the induced probability distribution to have a certain outcome.

Proposition 5.1.11. [Heisenberg Uncertainty Principle] Let X, Y be Hermitian operators of mean zero, corresponding to observables on a system in state ρ . Then

$$\sigma(X)\sigma(Y) \geq \frac{|\text{trace}([X, Y]\rho)|}{2}.$$

The uncertainty principle says that the failure of two Hermitian operators to commute lower bounds the product of their uncertainties. In particular, if they do not commute, neither can give rise to a classical (certain) measurement.

For $X, Y \in \text{End}(\mathcal{H})$, introduce the notation $\{X, Y\} := XY + YX$.

Exercise 5.1.12: For any Hermitian operators $A, B \in \text{End} \mathcal{H}$, and a density operator ρ , show that

$$|\text{trace}([A, B]\rho)|^2 + |\text{trace}(\{A, B\}\rho)|^2 = 4|\text{trace}(AB\rho)|^2.$$

Exercise 5.1.13: Prove Proposition 5.1.11. \odot

5.1.6. Pure and mixed states.

Definition 5.1.14. Let $\rho \in \text{End}(\mathcal{H})$ be a density operator. If $\text{rank}(\rho) = 1$, i.e. $\rho = |\xi\rangle\langle\xi| =: |\xi\rangle\langle\xi|$, ρ is called a *pure state*, and otherwise it is called a *mixed state*.

Exercise 5.1.15: Show that indeed, a rank one density operator is of the form $|\xi\rangle\langle\xi|$ (as opposed to just $|\xi\rangle\langle\psi|$).

Exercise 5.1.16: Show that ρ is pure if and only if $\rho \circ \rho = \rho$.

The partial trace of a pure state can be a mixed state. For example, if $\rho = |\psi\rangle\langle\psi|$ with $\psi = \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle) \in \mathcal{H}_1 \otimes \mathcal{H}_2$, then $\text{trace}_{\mathcal{H}_2}(\rho) = \frac{1}{2}(|0\rangle\langle 0| + |1\rangle\langle 1|)$.

The following proposition shows that we could avoid density operators altogether by working on a larger space:

Proposition 5.1.17. An arbitrary mixed state $\rho \in \text{End}(\mathcal{H})$ may be represented as the partial trace $\text{trace}_{\mathcal{H}'} |\psi\rangle\langle\psi|$ of a pure state in $\text{End}(\mathcal{H} \otimes \mathcal{H}')$ for some Hilbert space \mathcal{H}' . In fact, one can always take $\mathcal{H}' = \mathcal{H}^*$.

Exercise 5.1.18: Show that given a density operator $\rho \in \text{End}(\mathcal{H})$, there is a well defined operator $\sqrt{\rho} \in \text{End}(\mathcal{H})$ whose eigenspaces are the same as for ρ , and whose eigenvalues are the positive square roots of the eigenvalues of ρ .

Proof. We are given $\rho \in \mathcal{H} \otimes \mathcal{H}^*$. Consider $|\sqrt{\rho}\rangle\langle\sqrt{\rho}| \in \text{End}(\mathcal{H} \otimes \mathcal{H}^*)$. Then $\rho = \text{trace}_{\mathcal{H}^*}(|\sqrt{\rho}\rangle\langle\sqrt{\rho}|)$. \square

Exercise 5.1.19: Verify $\rho = \text{trace}_{\mathcal{H}^*}(|\sqrt{\rho}\rangle\langle\sqrt{\rho}|)$.

A pure state whose partial trace is ρ is called a *purification* of ρ . The purification $|\sqrt{\rho}\rangle\langle\sqrt{\rho}|$ is called the *standard purification*.

Exercise 5.1.20: Show that if $|\psi\rangle\langle\psi|, |\xi\rangle\langle\xi| \in \text{End}(\mathcal{H} \otimes \mathcal{H}')$ are purifications of ρ , then $|\psi\rangle = U|\xi\rangle$ for some $U \in \mathbf{U}(\mathcal{H}')$.

Exercise 5.1.21: Show that in a purification, one may take $\dim \mathcal{H}' = \text{rank } \rho$.

Now that we have our generalizations of probability distributions, we are nearly ready for quantum information theory. It remains to discuss when a sequence of density operators converges to a given density operator. For this we need a measure of distance. We will give two after a detour on distances between classical probabilities.

5.2. Distances between classical and quantum probability distributions

We first consider when two classical probability distributions are close. It is convenient to define two notions of closeness because the quantum analog of the first does not behave well under purification ****check****.

5.2.1. Classical distances. Say we have a source $\mathcal{X} = \{1, \dots, d\}$ generating elements according to a probability distribution p , and then we send the symbol generated through a channel where some corruption occurs, altering the symbol with probability at most ϵ . Let the resulting distribution of p followed by corruption be q . We would like a measure of how much p has been corrupted by noise. We have already seen the ℓ_1 -norm:

$$\|p - q\|_1 := \frac{1}{2} \sum_j |p_j - q_j|.$$

We included the $\frac{1}{2}$ because another natural definition of the ℓ_1 -norm is then given by the following exercise:

Exercise 5.2.1: Show that $\|p - q\|_1 = \max_{M \subset [d]} |p(M) - q(M)|$.

We will also use a second measure of how distributions p, q are close. Given p, q define corresponding vectors $\bar{p}, \bar{q} \in \mathbb{R}^d$, where $\bar{p} = (p_1, \dots, p_d)$. Consider the vectors $\sqrt{\bar{p}} := (\sqrt{p_1}, \dots, \sqrt{p_d}), \sqrt{\bar{q}} \in S^{d-1} \subset \mathbb{R}^d$, where S^{d-1} denotes the unit sphere.

Define the *fidelity* of p and q to be

$$F(p, q) := \langle \sqrt{\bar{p}} | \sqrt{\bar{q}} \rangle = \sum_j \sqrt{p_j q_j}.$$

Note that $F(p, p) = 1$. If one prefers a measure that is zero when $p = q$, one can define the *fidelity distance*

$$d_F(p, q) := \sqrt{2(1 - F(p, q))}.$$

Given this language, we may rephrase Shannon's noiseless theorem as saying that if one transmits below capacity, i.e., $|\mathcal{Y}(n)| > 2^{nH(\bar{p})}$, there exist $\mathcal{E}_n, \mathcal{D}_n$ such that the encoding and decoding

$$\mathcal{X}^{\times n} \xrightarrow{\mathcal{E}_n} \mathcal{Y}(n) \xrightarrow{\mathcal{D}_n} \mathcal{X}^{\times n}$$

satisfies $F(\mathcal{D}_n \circ \mathcal{E}_n(p), p) \rightarrow 1$ as $n \rightarrow \infty$ or $\|\mathcal{D}_n \circ \mathcal{E}_n(p) - p\|_1 \rightarrow 0$ as $n \rightarrow \infty$.

5.2.2. Distances between density operators. Recall that if $\rho \in \text{End}(\mathbb{C}^d)$ is diagonal, it corresponds to a classical probability distribution, so whatever distances we define on density operators, we could ask them to specialize to the ℓ_1 distance and fidelity distance when the density operators measured are classical. The quantum cousin of the ℓ_1 norm is the trace norm: The trace norm of $A \in \text{End}(\mathcal{H})$ is $\|A\|_{tr} := \text{trace}(\sqrt{A^\dagger A})$. Recall that $A^\dagger A$ is Hermitian and non-negative so the square root makes sense.

Note that if A is Hermitian, $\|A\|_{tr}$ is the sum of the absolute values of the eigenvalues and thus for density operators ρ , $\|\rho\|_{tr}$ is the sum of the eigenvalues. In particular, if ρ, σ are diagonal with diagonals p, q , then $\|\rho - \sigma\|_{tr} = \|p - q\|_1$.

Exercise 5.2.2: Show that $\|A\|_{tr}$ is indeed a norm and that

$$\|A\|_{tr} = \max_{U \in \mathbf{U}(\mathcal{H})} |\text{trace } AU|.$$

⊙

Exercise 5.2.3: Show that for $A \in \text{End}(\mathcal{H}_1)$ and $B \in \text{End}(\mathcal{H}_2)$, for $A \otimes B \in \text{End}(\mathcal{H}_1 \otimes \mathcal{H}_2)$, one has $\|A \otimes B\|_{tr} = \|A\|_{tr} \|B\|_{tr}$.

Now for the quantum version of fidelity and fidelity distance: For density operators ρ, σ the *fidelity* is

$$\begin{aligned} F(\rho, \sigma) &:= \max \left\{ |\langle \xi | \eta \rangle| \mid \xi, \eta \in \mathcal{H} \otimes \mathcal{H}' \text{ and} \right. \\ &\quad \left. \rho = \text{trace}_{\mathcal{H}'} |\xi\rangle\langle\xi|, \sigma = \text{trace}_{\mathcal{H}'} |\eta\rangle\langle\eta| \right\} \\ &= \text{trace}(\sqrt{\rho}\sqrt{\sigma}) \\ &= \text{trace} \sqrt{\sqrt{\rho}\sigma\sqrt{\rho}} \end{aligned}$$

where the maximum is over all possible purifications.

The first description illustrates a convenience of fidelity.

Exercise 5.2.4: Verify the three quantities are equal.

Note that if ρ, σ are classical (i.e., diagonal in the standard basis), this fidelity agrees with the classical fidelity.

As before, the fidelity distance is $d_F(\rho, \sigma) = \sqrt{2(1 - \sqrt{F(\rho, \sigma)})}$.

Exercise 5.2.5: Show that

$$d_F(\rho, \sigma) = \min\{\| |\xi\rangle - |\eta\rangle \| \mid \rho = \text{trace}_{\mathcal{H}'}(|\xi\rangle\langle\xi|), \sigma = \text{trace}_{\mathcal{H}'}(|\sigma\rangle\langle\sigma|)\}.$$

Here again we see a utility of fidelity when taking purifications.

5.3. The quantum noiseless channel theorem

5.3.1. What is a quantum channel? A quantum channel should be a linear map sending $\rho \in \text{End}(\mathcal{H}_A)$ to some $\Phi(\rho) \in \text{End}(\mathcal{H}_B)$.

First consider the special case $\mathcal{H}_A = \mathcal{H}_B$. We should allow coupling with an auxiliary system, i.e.,

$$(5.3.1) \quad \rho \mapsto \rho \otimes \sigma \in \text{End}(\mathcal{H}_A \otimes \mathcal{H}_C).$$

We should allow the state $\rho \otimes \sigma$ to evolve in $\text{End}(\mathcal{H}_A \otimes \mathcal{H}_C)$, i.e., be acted upon by an arbitrary $U \in \mathbf{U}(\mathcal{H}_A \otimes \mathcal{H}_C)$.

Finally we should allow measurements, i.e., tracing out the \mathcal{H}_C part.

In summary, $\rho \mapsto \text{trace}_{\mathcal{H}_C}(U(\rho \otimes \sigma)U^{-1})$. More generally to go from \mathcal{H}_A to \mathcal{H}_B , one needs to allow isometries as well. That is, we should allow maps of the form:

$$(5.3.2) \quad \rho \mapsto \Phi_{A,B} \text{trace}_{\mathcal{H}_C}(U(\rho \otimes \sigma)U^{-1}),$$

where $\Phi_{A,B} : \mathcal{H}_A \rightarrow \mathcal{H}_B$ is an isometry.

Exercise 5.3.1: Show that any isometry $\mathbb{C}^n \rightarrow \mathbb{C}^m$ may be viewed as the restriction of an element of $\mathbf{U}(n+m)$ to \mathbb{C}^n .

Maps of the form (5.3.2) are *completely positive trace preserving maps* (CPTP), where a map Λ is *completely positive* if $\Lambda \otimes \text{Id}_{\mathcal{H}_E}$ is positive for all \mathcal{H}_E . These maps should be viewed as the discrete time evolution of a quantum system.

Exercise 5.3.2: Show that the transpose map $X \mapsto X^T$ is positive, but not completely positive. \odot

Moreover, all CPTP maps $\mathcal{H}_A \rightarrow \mathcal{H}_A$ are of the form (5.3.2):

Theorem 5.3.3 (Stinespring dilation [Sti55]). *Let $\Lambda : \text{End}(\mathcal{H}_A) \rightarrow \text{End}(\mathcal{H}_B)$ be CPTP. Then there exists an isometric embedding $Z : \mathcal{H}_A \rightarrow \mathcal{H}_B \otimes \mathcal{H}_E$ for some space \mathcal{H}_E , inducing $\tilde{Z} : \text{End}(\mathcal{H}_A) \rightarrow \text{End}(\mathcal{H}_B \otimes \mathcal{H}_E)$ (i.e., $\tilde{Z}(X) = ZXZ^\dagger$), such that $\Lambda = \text{trace}_{\mathcal{H}_E} \circ \tilde{Z}$, i.e., $\Lambda(X) = \text{trace}_{\mathcal{H}_E}(ZXZ^\dagger)$.*

Proof. First note that we have a canonical isomorphism:

$$\begin{aligned} \text{Hom}(\text{End}(\mathcal{H}_A), \text{End}(\mathcal{H}_B)) &= (\mathcal{H}_A^* \otimes \mathcal{H}_A)^* \otimes (\mathcal{H}_B^* \otimes \mathcal{H}_B) \\ &= (\mathcal{H}_A^* \otimes \mathcal{H}_B)^* \otimes (\mathcal{H}_A^* \otimes \mathcal{H}_B) \\ &= \text{End}(\mathcal{H}_A^* \otimes \mathcal{H}_B). \end{aligned}$$

Given $\Lambda \in \text{Hom}(\text{End}(\mathcal{H}_A), \text{End}(\mathcal{H}_B))$, write $\rho^\Lambda \in \text{End}(\mathcal{H}_A^* \otimes \mathcal{H}_B)$ for the induced operator.

Exercise 5.3.4: Show that if Λ is CPTP, then ρ^Λ is a density operator.

Exercise 5.3.5: Show that if Λ is CPTP, then $\text{trace}_{\mathcal{H}_B}(\rho^\Lambda) = \frac{1}{\dim \mathcal{H}_A} \text{Id}_{\mathcal{H}_A^*}$.

Exercise 5.3.6: Show that the canonical isomorphism restricted to CPTP operators surjects onto the set of density operators whose partial trace on \mathcal{H}_B is the rescaled identity map, and thus identifies the two sets.

add hints

The above isomorphism is called the *Choi-Jamilkowski isomorphism*, and we will denote it CJ .

To prove the theorem, given a CPTP Λ , apply CJ to get a density operator, then take a purification, to get some $|\Psi^\Lambda\rangle \in \mathcal{H}_A^* \otimes \mathcal{H}_B \otimes \mathcal{H}_C$, where $\mathcal{H}_C \simeq (\mathcal{H}_A^* \otimes \mathcal{H}_B)^*$ is the purifying space. Then apply CP^{-1} . The upshot is we have realized Λ as stated in the theorem starting with ρ^Λ . \square

The map Z is called the *Stinespring dilation* of Λ .

Call a CPTP map an *instrument* if $\Lambda = \sum_i |i\rangle\langle i| \otimes \Lambda_i$ where each Λ_i is completely positive. From this perspective, we could view a measurement ****say which kind**** as the application of an instrument to a state.

5.3.2. Set up. Now instead of having a source $\mathcal{X}^{\times n}$ our “source” is $\mathcal{H}^{\otimes n}$, where we think of $\mathcal{H}^{\otimes n} = \mathcal{H}_A^{\otimes n}$, and Alice will transmit a state to Bob via a CPTP map, and instead of a probability distribution p we have a density operator ρ .

We seek an encoder \mathcal{E} and decoder \mathcal{D} and a compression space \mathcal{H}_{0n} :

$$\mathcal{H}_A^{\otimes n} \xrightarrow{\mathcal{E}} \mathcal{H}_{0n} = (\mathbb{C}^2)^{\otimes nR} \xrightarrow{\mathcal{D}} \mathcal{H}_B^{\otimes n}$$

with R as small as possible such that $F(\rho^{\otimes n}, \mathcal{E} \circ \mathcal{D}(\rho^{\otimes n})) \rightarrow 1$ as $n \rightarrow \infty$. To determine R , we need a quantum version of entropy.

5.3.3. von Neumann entropy.

Definition 5.3.7. The *von Neumann entropy* of a density operator ρ is $H(\rho) = -\text{trace}(\rho \log(\rho))$.

Here $\log(\rho)$ is defined as follows: write ρ in terms of its eigenvectors and eigenvalues, $\rho = \sum_j \lambda_j |\psi_j\rangle\langle\psi_j|$, then $\log(\rho) = \sum_j \log(\lambda_j) |\psi_j\rangle\langle\psi_j|$.

Note that if $\rho = \sum_j \lambda_j |\psi_j\rangle\langle\psi_j|$, then $H(\rho) = -\sum_j \lambda_j \log(\lambda_j)$ so if ρ is classical (i.e., diagonal), one obtains the Shannon entropy.

Proposition 5.3.8. *The von Neumann entropy has the following properties:*

- (1) $H(\rho) \geq 0$ with equality if and only if ρ is pure.
- (2) Let $\dim \mathcal{H} = d$. Then $H(\rho) \leq \log(d)$ with equality if and only if $\rho = \frac{1}{d} \text{Id}_{\mathcal{H}}$.
- (3) If $\rho = |\psi\rangle\langle\psi| \in \text{End}(\mathcal{H}_A \otimes \mathcal{H}_B)$, then $H(\rho_A) = H(\rho_B)$, where $\rho_A = \text{trace}_{\mathcal{H}_B}(\rho) \in \text{End}(\mathcal{H}_A)$.

Exercise 5.3.9: Prove the first two assertions.

Proof of (3). Write $|\psi\rangle = \sum_j z_j |\phi_j\rangle \otimes |\eta_j\rangle$ where the $|\phi_j\rangle \in \mathcal{H}_A$ are a basis as are the $|\eta_j\rangle \in \mathcal{H}_B$ and $\sum |z_j|^2 = 1$. Then $|\psi\rangle\langle\psi| = \sum_{i,j} z_i \bar{z}_j |\phi_i\rangle\langle\phi_j| \otimes |\eta_i\rangle\langle\eta_j|$, so $\text{trace}_{\mathcal{H}_B}(|\psi\rangle\langle\psi|) = \sum_j |z_j|^2 |\phi_j\rangle\langle\phi_j|$ and $\text{trace}_{\mathcal{H}_A}(|\psi\rangle\langle\psi|) = \sum_j |z_j|^2 |\eta_j\rangle\langle\eta_j|$. \square

5.3.4. Quantum typical subspace theorem. Recall that in the classical case, given a probability distribution on \mathcal{X} we had a splitting $\mathcal{X}^{\times n} = \mathcal{X}_{\epsilon\text{-typ}}^n \sqcup \mathcal{X}_{\delta\text{-small}}^n$. Here is a quantum analog:

Theorem 5.3.10. *Given (\mathcal{H}, ρ) and $\epsilon, \delta > 0$, there exists n_0 such that for all $n \geq n_0$, there exists an orthogonal direct sum decomposition $\mathcal{H}^{\otimes n} = \mathcal{H}_{\epsilon\text{-typ}}^n \oplus \mathcal{H}_{\delta\text{-small}}^n$ satisfying:*

- (1) $\text{trace}(\text{Proj}_{\mathcal{H}_{\delta\text{-small}}^n} \rho^{\otimes n}) \leq \delta$, i.e., events in $\mathcal{H}_{\delta\text{-small}}^n$ are improbable,
- (2) $\mathcal{H}_{\epsilon\text{-typ}}^n$ is spanned by eigenvectors with eigenvalues close to the expected value. More precisely, it is spanned by vectors of the form $|\bar{x}\rangle = |x_1\rangle \otimes \cdots \otimes |x_n\rangle$ where $|x_j\rangle$ is an eigenvector for ρ with eigenvalue $\lambda(x_j)$, so $|\bar{x}\rangle$ is an eigenvector for $\rho^{\otimes n}$ with eigenvalue $\lambda(\bar{x}) := \lambda(x_1) \cdots \lambda(x_n)$. The for all $|\bar{x}\rangle \in \mathcal{H}_{\epsilon\text{-typ}}^n$,

$$\left| \frac{1}{n} \log\left(\frac{1}{\lambda(\bar{x})}\right) - H(\rho) \right| < \epsilon,$$

- (3) $(1 - \delta)2^{nH(\rho) - \epsilon} \leq \dim \mathcal{H}_{\epsilon\text{-typ}}^n \leq 2^{nH(\rho) + \epsilon}$.

The proof is identical to the classical case Theorem 4.3.1. Here is a quantum analog of Theorem 4.3.3.

Theorem 5.3.11. *Given (\mathcal{H}, ρ) , let $R < H(\rho)$ and $\eta > 0$. Then there exists \tilde{n} such that for all $n \geq \tilde{n}$, and all sequences $\mathcal{M}_n \subset \mathcal{H}^{\otimes n}$ with $\dim \mathcal{M}_n \leq 2^{nR}$, $\text{trace}(\text{Proj}_{\mathcal{M}_n} \rho^{\otimes n}) \leq \eta$.*

Proof. The proof is the same as the classical case except for one step. Write

$$\text{trace}(\text{Proj}_{\mathcal{M}} \rho^{\otimes n}) = \text{trace}(\text{Proj}_{\mathcal{M}} \rho^{\otimes n} \text{Proj}_{\mathcal{H}_{\epsilon\text{-typ}}^n}) + \text{trace}(\text{Proj}_{\mathcal{M}} \rho^{\otimes n} \text{Proj}_{\mathcal{H}_{\delta\text{-small}}^n}).$$

Since the decomposition $\mathcal{H}^{\otimes n} = \mathcal{H}_{\epsilon\text{-typ}}^n \oplus \mathcal{H}_{\delta\text{-small}}^n$ is orthogonal and the first space is spanned by eigenvectors of $\rho^{\otimes n}$, whose eigenvectors are orthogonal, the second is as well, so it commutes with $\rho^{\otimes n}$. Moreover, $\text{Proj}(\mathcal{H}_{\delta\text{-small}}^n)^2 = \text{Proj}_{\mathcal{H}_{\delta\text{-small}}^n}$, so we may rewrite the second term as

$$\text{trace}(\text{Proj}_{\mathcal{M}} \text{Proj}_{\mathcal{H}_{\delta\text{-small}}^n} \rho^{\otimes n} \text{Proj}_{\mathcal{H}_{\delta\text{-small}}^n}).$$

Thus even without the projection to \mathcal{M} , by choosing n sufficiently large, the second term is at most some δ of our choosing. The operators $\rho^{\otimes n}$ and $\text{Proj}_{\mathcal{H}_{\epsilon\text{-typ}}^n}$ also commute, and since $\text{Proj}_{\mathcal{H}_{\epsilon\text{-typ}}^n}^2 = \text{Proj}_{\mathcal{H}_{\epsilon\text{-typ}}^n}$, we have $\rho^{\otimes n} \text{Proj}_{\mathcal{H}_{\epsilon\text{-typ}}^n} = \text{Proj}_{\mathcal{H}_{\epsilon\text{-typ}}^n} \rho^{\otimes n} \text{Proj}_{\mathcal{H}_{\epsilon\text{-typ}}^n}$, the projection of $\rho^{\otimes n}$ onto $\mathcal{H}_{\epsilon\text{-typ}}^n$. Thus

$$\text{trace}(\text{Proj}_{\mathcal{M}} \rho^{\otimes n} \text{Proj}_{\mathcal{H}_{\epsilon\text{-typ}}^n}) \leq 2^{nR} 2^{-n(H(\rho)-\epsilon)},$$

the first factor because we are projecting to a space of dimension at most 2^{nR} and the second because the eigenvalues in the typical subspace are all at most $2^{-n(H(\rho)-\epsilon)}$. Now simply choose ϵ, δ and \tilde{n} such that $2^{\tilde{n}\epsilon} 2^{-n(H(\rho)-R)} + \delta < \eta$. \square

5.3.5. The quantum noiseless channel theorem.

Theorem 5.3.12. [Sch95] *Let (\mathcal{H}, ρ) be an i.i.d. quantum source. If $R > H(\rho)$, then there exists a reliable compression scheme of rate R . That is, there exists a compression space \mathcal{H}_{0n} , of dimension 2^{nR} , and encoder $\mathcal{E} : \mathcal{H}^{\otimes n} \rightarrow \mathcal{H}_{0n}$ and a decoder $\mathcal{D} : \mathcal{H}_{0n} \rightarrow \mathcal{H}^{\otimes n}$ such that $\lim_{n \rightarrow \infty} F(\rho^{\otimes n}, \mathcal{D} \circ \mathcal{E}(\rho^{\otimes n})) = 1$. If $R < H(\rho)$, then any compression scheme is unreliable.*

The idea of the proof is the same as the classical case. Here are details.

Proof. Say $R > H(\rho)$, write $R = H(\rho) + \epsilon$. The quantum typical subspace theorem implies that for all $\delta > 0$ and n sufficiently large, $\text{trace}(\rho^{\otimes n} \text{Proj}_{\mathcal{H}_{\epsilon\text{-typ}}^n}) \geq 1 - \delta$ and $\dim(\mathcal{H}_{\epsilon\text{-typ}}^n) \leq 2^{nR}$. Fix $\mathcal{H}_0^n = \mathbb{C}^{2^{nR}} \supseteq \mathcal{H}_{\epsilon\text{-typ}}^n$. Let $P_{0,n} : \mathcal{H}^{\otimes n} \rightarrow \mathbb{C}|\psi_n\rangle$ be an operator sending $\mathcal{H}_{\epsilon\text{-typ}}^n$ to zero and $\mathcal{H}_{\delta\text{-small}}^n$ to $\mathbb{C}|\psi_n\rangle$. Define the encoder $\mathcal{E} : \mathcal{H}^{\otimes n} \rightarrow \mathcal{H}_0^n$ to be $\mathcal{E}(\rho^{\otimes n}) = \text{Proj}_{\mathcal{H}_{\epsilon\text{-typ}}^n} \rho^{\otimes n} \text{Proj}_{\mathcal{H}_{\epsilon\text{-typ}}^n} + P_{0,n} \rho^{\otimes n} P_{0,n}^\dagger$, and the decoder $\mathcal{D} : \mathcal{H}_0^n \rightarrow \mathcal{H}^{\otimes n}$ to simply be the inclusion. We estimate

the fidelity:

$$\begin{aligned}
F(\rho^{\otimes n}, \mathcal{D} \circ \mathcal{E}(\rho^{\otimes n})) &= [\text{trace} \left(\sqrt{\rho^{\otimes n}} \mathcal{D} \circ \mathcal{E} \rho^{\otimes n} \sqrt{\rho^{\otimes n}} \right)^{\frac{1}{2}}]^2 \\
&= [\text{trace} \left(\sqrt{\rho^{\otimes n}} \text{Proj}_{\mathcal{H}_{\epsilon-ty}^n} \rho^{\otimes n} \text{Proj}_{\mathcal{H}_{\epsilon-ty}^n} \sqrt{\rho^{\otimes n}} \right) + \sqrt{\rho^{\otimes n}} P_{0,n} \rho^{\otimes n} P_{0,n}^\dagger \sqrt{\rho^{\otimes n}}]^{\frac{1}{2}}]^2 \\
&\geq [\text{trace} \left(\sqrt{\rho^{\otimes n}} \text{Proj}_{\mathcal{H}_{\epsilon-ty}^n} \rho^{\otimes n} \text{Proj}_{\mathcal{H}_{\epsilon-ty}^n} \sqrt{\rho^{\otimes n}} \right)^{\frac{1}{2}}]^2 \\
&= [\text{trace}(\rho^{\otimes n} \text{Proj}_{\mathcal{H}_{\epsilon-ty}^n})]^2 \text{ because } \text{Proj}_{\mathcal{H}_{\epsilon-ty}^n}, \sqrt{\rho^{\otimes n}} \text{ commute} \\
&\geq (1 - \delta)^2
\end{aligned}$$

Now just let $\delta \rightarrow 0$.

The proof of the other direction is similar to the classical case, see Exercise 5.3.13 below. \square

Exercise 5.3.13: Prove that for any \mathcal{D}, \mathcal{E} used when $\dim \mathcal{H}_{0n} < 2^{nH(\rho)}$, that $F(\rho^{\otimes n}, \mathcal{D} \circ \mathcal{E}(\rho^{\otimes n}))$ becomes arbitrarily small as $n \rightarrow \infty$. \odot

5.4. Properties of von Neumann entropy

5.4.1. Von Neumann entropy of classical/quantum systems. We introduced density matrices to facilitate the mixing of classical and quantum probability. Here is another example:

Proposition 5.4.1. *If $\rho = \bigoplus_i p_i \rho_i$ where $\rho_i \in \text{End}(\mathcal{H}_i)$, $\mathcal{H} = \bigoplus_i \mathcal{H}_i$ is an orthogonal decomposition, and \bar{p} is a probability distribution, then $H(\rho) = H(\bar{p}) + \sum_i p_i H(\rho_i)$.*

Proof. Let λ_i^j be the eigenvalues of ρ_i .

$$\begin{aligned}
H\left(\sum_i p_i \rho_i\right) &= - \sum_{i,j} (p_i \lambda_i^j) \log(p_i \lambda_i^j) \\
&= - \sum_{i,j} p_i \lambda_i^j \log(p_i) - \sum_{i,j} p_i \lambda_i^j \log(\lambda_i^j) \\
&= H(\bar{p}) + \sum_i p_i H(\rho_i)
\end{aligned}$$

where the last line holds because for each fixed i , $\sum_j \lambda_i^j = 1$. \square

Exercise 5.4.2: Show that if $\rho = \sum_i p_i |i\rangle\langle i| \otimes \rho_i \in \text{End}(\mathbb{C}^d \otimes \mathcal{H})$, with terms as above, then $H(\rho) = H(\bar{p}) + \sum_i p_i H(\rho_i)$.

5.4.2. Relative entropy. We introduce yet another quantity that will facilitate the comparisons of density operators. First for the classical case, define the relative entropy $H(\bar{p}||\bar{q}) := -\sum p_i \log \frac{q_i}{p_i} = -H(\bar{p}) - \sum_i p_i \log(q_i)$. Note that it is zero when $\bar{p} = \bar{q}$. It is otherwise positive:

Proposition 5.4.3. $H(\bar{p}||\bar{q}) \geq 0$ with equality if and only if $\bar{p} = \bar{q}$.

Proof. Recall that $\log(x) \ln(2) = \ln x \leq x - 1$ for all $x > 0$ with equality if and only if $x = 1$, i.e., $\log(x) \leq \frac{1}{\ln(2)}(x - 1)$. We have

$$\begin{aligned} H(\bar{p}||\bar{q}) &= \sum p_i \log \frac{q_i}{p_i} \\ &\geq -\sum p_i \frac{1}{\ln(2)} \left(\frac{q_i}{p_i} - 1 \right) \\ &= \frac{1}{\ln(2)} \sum_j (p_j - q_j) \\ &= \frac{1}{\ln(2)} (0 - 0). \end{aligned}$$

□

Define the *relative von Neumann entropy* $H(\rho||\sigma) := \text{trace}(\rho \log(\rho)) - \text{trace}(\rho \log(\sigma))$. It shares the positivity property of its classical cousin:

Proposition 5.4.4. [Kle] $H(\rho||\sigma) \geq 0$ with equality if and only if $\rho = \sigma$.

Proof. Write the eigenvalue decompositions $\rho = \sum_i p_i |v_i\rangle\langle v_i|$, $\sigma = \sum_j q_j |w_j\rangle\langle w_j|$. Then

$$\begin{aligned} H(\rho||\sigma) &= \sum_i p_i \log(p_i) - \text{trace}\left(\left(\sum_i p_i |v_i\rangle\langle v_i|\right)\left(\sum_j \log(q_j) |w_j\rangle\langle w_j|\right)\right) \\ &= \sum_i p_i \log(p_i) - \text{trace}\left(\left(\sum_{i,j} p_i \log(q_j) \langle v_i|w_j\rangle\right) |v_i\rangle\langle w_j|\right) \\ &= \sum_i p_i \log(p_i) - \sum_{i,j} p_i \log(q_j) \langle v_i|w_j\rangle\langle w_j|v_i\rangle. \end{aligned}$$

Write $P_{ij} = \langle v_i|w_j\rangle\langle w_j|v_i\rangle$. Note that $P_{ij} \geq 0$ for all i, j , and for all j , $\sum_i P_{ij} = 1$ and for all i , $\sum_j P_{ij} = 1$. This is the definition of the matrix (P_{ij}) being doubly stochastic (cf. Exercise 4.2.3).

Recall that a function $f(x)$ is *concave* if for all $0 \leq \lambda_1, \lambda_2$ with $\lambda_1 + \lambda_2 = 1$, $f(\lambda_1 x_1 + \lambda_2 x_2) \geq \lambda_1 f(x_1) + \lambda_2 f(x_2)$.

Now $\log(x)$ is a concave function, so

$$\sum_j P_{ij} \log(q_j) \leq \log\left(\sum_j P_{ij} q_j\right).$$

Set $\tilde{q}_i = \sum_j P_{ij} q_j$ and recall from Exercise 4.2.3 that \tilde{q}_i is a probability distribution. Thus $H(\rho||\sigma) \geq \sum_i p_i \log(\frac{\tilde{q}_i}{p_i}) = H(\bar{p}||\bar{q})$ and we conclude by the classical case. \square

Proposition 5.4.4 shows that relative entropy is a measure of distance between two density operators

Corollary 5.4.5 (von Neumann Entropy is non-decreasing under projective measurements). *Let Proj_i be a complete set of orthogonal projectors, set $\rho' = \sum_i \text{Proj}_i \rho \text{Proj}_i$. Then $H(\rho') \geq H(\rho)$ with equality if and only if $\rho' = \rho$.*

Proof. We have $0 \leq H(\rho||\rho') = -H(\rho) - \text{trace}(\rho \log(\rho'))$. Now

$$\begin{aligned} \text{trace}(\rho \log(\rho')) &= \text{trace} \left(\sum_i \text{Proj}_i \rho \log(\rho') \right) \\ &= \text{trace} \left(\sum_i \text{Proj}_i \rho \log(\rho') \text{Proj}_i \right) \end{aligned}$$

because $\text{Proj}_i^2 = \text{Proj}_i$ and $\text{trace}(ABC) = \text{trace}(BCA)$. Now Proj_i commutes with ρ' and $\log(\rho')$ because $\text{Proj}_i \text{Proj}_j = 0$ if $i \neq j$, so

$$\begin{aligned} \text{trace}(\rho \log(\rho')) &= \text{trace} \left(\sum_i \text{Proj}_i \rho \text{Proj}_i \log(\rho') \right) \\ &= \text{trace}(\rho' \log(\rho')) \\ &= -H(\rho') \end{aligned}$$

Putting it all together, we obtain the result. \square

5.4.3. von Neumann entropy and composite systems. In what follows ρ_{AB} is a density operator on $\mathcal{H}_A \otimes \mathcal{H}_B$ and $\rho_A = \text{trace}_{\mathcal{H}_B}(\rho_{AB})$, $\rho_B = \text{trace}_{\mathcal{H}_A}(\rho_{AB})$ are respectively the induced density operators on \mathcal{H}_A , \mathcal{H}_B .

Exercise 5.4.6: Show that $\text{trace}(\rho_{AB} \log(\rho_A \otimes \rho_B)) = -H(\rho_A) - H(\rho_B)$. \odot

Proposition 5.4.7 (Sub-additivity of von Neumann entropy). $H(\rho_{AB}) \leq H(\rho_A) + H(\rho_B)$ with equality if and only if $\rho_{AB} = \rho_A \otimes \rho_B$.

Proof. We have $0 \leq H(\rho_{AB}||\rho_A \otimes \rho_B) = H(\rho_{AB}) - \text{trace}(\rho_{AB} \log(\rho_A \otimes \rho_B))$, and we conclude by Exercise 5.4.6. \square

Proposition 5.4.8 (Triangle inequality for von Neumann entropy). $H(\rho_{AB}) \geq |H(\rho_A) - H(\rho_B)|$

Proof. Let $\mathcal{H}_R \simeq (\mathcal{H}_A \otimes \mathcal{H}_B)^*$ be a space purifying ρ_{AB} , i.e., ρ_{ABR} is a pure state density operator on $\mathcal{H}_A \otimes \mathcal{H}_B \otimes \mathcal{H}_R$ such that $\text{trace}_{\mathcal{H}_R}(\rho_{ABR}) = \rho_{AB}$. Sub-additivity implies $H(\rho_R) + H(\rho_A) \geq H(\rho_{AR})$, and since ρ_{ABR} is pure

we have $H(\rho_R) = H(\rho_{AB})$ and $H(\rho_{AR}) = H(\rho_B)$. Putting these together gives $H(\rho_{AB}) \geq H(\rho_B) - H(\rho_A)$ and combining this with the analogous calculation with the roles of A and B reversed gives the result. \square

5.4.4. von Neumann entropy and classical-quantum density. Here is another example how the reformulation in terms of density operators facilitates the combining of classical and quantum probability.

Proposition 5.4.9 (Concavity of von Neumann entropy). *Let ρ_i , $1 \leq i \leq d$ be density operators on \mathcal{H} , let \bar{p} be a probability distribution on $[d]$, and consider the density operator $\sum_i p_i \rho_i$. Then*

$$\sum_i p_i H(\rho_i) \leq H\left(\sum_i p_i \rho_i\right).$$

Proof. Let $\mathcal{H}' = \mathbb{C}^d$ and let $\rho_{\mathcal{H}\mathcal{H}'} := \sum_i p_i \rho_i \otimes |i\rangle\langle i|$. By Exercise 5.4.2, $H(\rho_{\mathcal{H}\mathcal{H}'}) = H(\bar{p}) + \sum_i p_i H(\rho_i)$. On the other hand $H(\rho_{\mathcal{H}'}) = H(\bar{p})$ (because each ρ_i has trace one) and $H(\rho_{\mathcal{H}}) = H(\sum_i p_i \rho_i)$, so we conclude. \square

Remark 5.4.10. It is also true that $H(\sum_i p_i \rho_i) \leq \sum p_i H(\rho_i) + H(\bar{p})$, see, e.g., [NC00, §11.3.6].

5.4.5. Conditional von Neumann entropy. Recall the conditional Shannon entropy is defined to be $H(\bar{p}_{\mathcal{X}}|\bar{p}_{\mathcal{Y}}) = -\sum_{i,j} p_{\mathcal{X}\times\mathcal{Y}}(i,j) \log p_{\mathcal{X}|\mathcal{Y}}(i|j)$, the entropy of $p_{\mathcal{X}}$ conditioned on $y = j$, averaged over \mathcal{Y} . It is not clear how to “condition” one density matrix on another, so we need to find a different definition. Recall that Shannon entropy satisfies $H(\bar{p}_{\mathcal{X}}|\bar{p}_{\mathcal{Y}}) = H(\bar{p}_{\mathcal{X}\times\mathcal{Y}}) - H(\bar{p}_{\mathcal{Y}})$, and the right hand side of this expression makes sense for density operators, so define, for ρ_{AB} a density operator on $\mathcal{H}_A \otimes \mathcal{H}_B$,

$$(5.4.1) \quad H(\rho_A|\rho_B) := H(\rho_{AB}) - H(\rho_B).$$

WARNING: it is possible that the conditional von Neumann entropy is *negative*. Consider the following example: Let $|\psi\rangle = \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle) \in \mathcal{H}_A \otimes \mathcal{H}_B$. Then $\rho_A = \frac{1}{2} \text{Id}_{\mathcal{H}_A} = \frac{1}{2}(|0\rangle\langle 0| + |1\rangle\langle 1|)$ so $H(\rho_A) = 1$, but $H(|\psi\rangle\langle\psi|) = 0$ because $|\psi\rangle\langle\psi|$ is pure.

Despite this, we will see that for $\rho_{ABC} \in \mathcal{H}_A \otimes \mathcal{H}_B \otimes \mathcal{H}_C$, one always has

$$H(\rho_{AC}) - H(\rho_A) + H(\rho_{BC}) - H(\rho_B) \geq 0.$$

This result is called *strong sub-additivity*, and we will have to work a little to prove it. Before doing that, note that as with the classical case, we have

three expressions for the mutual information, each of which has a quantum-informational interpretation:

$$\begin{aligned} \text{Minfo}(\rho_A : \rho_B) &:= H(\rho_A) + H(\rho_B) - H(\rho_{AB}) \\ &= H(\rho_A) - H(\rho_A|\rho_B) \\ &= H(\rho_B) - H(\rho_B|\rho_A) \end{aligned}$$

5.4.6. Lieb's lemma. We prove a concavity Lemma that will be the key to proving strong sub-additivity.

Definition 5.4.11. A function $f : \mathbb{R}^N \times \mathbb{R}^N \rightarrow \mathbb{R}$ is *jointly concave* in its arguments if $f(\lambda_1 x_1 + \lambda_2 x_2, \lambda_1 y_1 + \lambda_2 y_2) \geq \lambda_1 f(x_1, y_1) + \lambda_2 f(x_2, y_2)$.

Lemma 5.4.12 (Lieb [Lie73]). *Let $X \in \text{End}(\mathcal{H})$. For all $0 \leq p \leq 1$,*

$$f(A, B) := \text{trace}(A^p X^\dagger B^{1-p} X),$$

defined on pairs of positive operators, is jointly concave.

Note that A^p, B^{1-p} make sense as A, B are positive.

Proof. We follow the proof in [Rus04]. First note that it is sufficient to prove the case $A = B$ because

$$\text{trace}(A^p X^\dagger B^{1-p} X) = \text{trace} \begin{pmatrix} A & 0 \\ 0 & B \end{pmatrix}^p \begin{pmatrix} 0 & X^\dagger \\ 0 & 0 \end{pmatrix} \begin{pmatrix} A & 0 \\ 0 & B \end{pmatrix}^{1-p} \begin{pmatrix} 0 & 0 \\ X & 0 \end{pmatrix}.$$

We need to show, for $\lambda_j \geq 0$ and $\lambda_1 + \lambda_2 = 1$, that

(5.4.2)

$$\lambda_1 \text{trace}(A_1^p X^\dagger A_1^{1-p} X) + \lambda_2 \text{trace}(A_2^p X^\dagger A_2^{1-p} X) \leq \text{trace}(\lambda_1 A_1 + \lambda_2 A_2)^p X^\dagger (\lambda_1 A_1 + \lambda_2 A_2)^{1-p} X.$$

We may assume $C := \lambda_1 A_1 + \lambda_2 A_2$ is invertible (the general case will follow by continuity). Write

$$M = C^{\frac{1-p}{2}} X C^{\frac{p}{2}}.$$

Set

$$f_k(p) = \text{trace}(A_k^p C^{-\frac{p}{2}} M^\dagger C^{-\frac{1-p}{2}} A_k^{1-p} C^{-\frac{1-p}{2}} M C^{-\frac{p}{2}}),$$

then (5.4.2) becomes

$$(5.4.3) \quad \lambda_1 f_1(p) + \lambda_2 f_2(p) \leq \text{trace}(M^\dagger M).$$

We need to prove this for $0 \leq p \leq 1$. Now we use a standard trick from complex analysis: Extend $f_j(p)$ to be defined for complex numbers. We can do so on the strip $0 \leq \text{Re}(z) \leq 1$. Why would we want to do this? First recall that if Z is hermitian, then Z^{iy} is unitary, in particular $\|Z^{iy}\|_{op} = 1$. The maximum principle (see, e.g., [Ahl78, §4.3.4]) implies that if $f(z)$ is uniformly bounded on a region, then the maximum of $f(z)$ is achieved on the boundary of the region. So we will trade proving (5.4.3) on the interval $(0, 1)$ to proving $\lambda_1 f_1(p) + \lambda_2 f_2(p)$ is bounded on a strip, and proving the

inequality (5.4.3) on the lines $0 + iy$ and $1 + iy$ where we will have more control over the operators.

Using that $|\text{trace}(XY)| \leq \text{trace}|XY| \leq \|X\|_{op} \text{trace}(Y)$, we see

$$\begin{aligned} f_k(z) &\leq \|A_k\| \|C^{-1}\| \text{trace}(M^\dagger M) \\ &\leq \|A_k\| \|C^{-1}\| \|C\| \text{trace}(X^\dagger X), \end{aligned}$$

so $f_k(z)$ is indeed uniformly bounded.

Consider

$$\begin{aligned} f_k(0 + iy) &= \text{trace}(A_k^{iy} C^{-\frac{iy}{2}} M^\dagger C^{\frac{iy}{2}} C^{-\frac{1}{2}} A_k^{1-iy} C^{-\frac{iy}{2}} C^{-\frac{1}{2}} M C^{-\frac{iy}{2}}) \\ &= \text{trace} \left((A_k^{\frac{iy}{2}} C^{-\frac{iy}{2}} M^\dagger C^{\frac{iy}{2}} C^{-\frac{1}{2}} A_k^{\frac{1}{2}}) (A_k^{\frac{1}{2}} A_k^{-iy} A_k^{\frac{1}{2}} C^{-\frac{iy}{2}} C^{-\frac{1}{2}} M C^{-\frac{iy}{2}} A_k^{\frac{iy}{2}}) \right) \end{aligned}$$

where we used $A_k^{iy} = A_k^{\frac{iy}{2}} A_k^{\frac{iy}{2}}$, $\text{trace}(ZW) = \text{trace}(WZ)$, and split $A_k^{1-iy} = A_k^{\frac{1}{2}} A_k^{\frac{1}{2}} A_k^{-iy}$.

Note that $|\text{trace}(ZW)| \leq [\text{trace}(ZZ^\dagger) \text{trace}(WW^\dagger)]^{\frac{1}{2}}$ by the Cauchy-Schwartz inequality. We have

$$\begin{aligned} |f_k(0 + iy)| &\leq \left[\text{trace}(A_k^{\frac{iy}{2}} C^{-\frac{iy}{2}} M^\dagger C^{\frac{iy}{2}} C^{-\frac{1}{2}} A_k^{\frac{1}{2}}) \text{trace}(A_k^{\frac{1}{2}} A_k^{-iy} A_k^{\frac{1}{2}} C^{-\frac{iy}{2}} C^{-\frac{1}{2}} M C^{-\frac{iy}{2}} A_k^{\frac{iy}{2}}) \right]^{\frac{1}{2}} \\ &= \left[\text{trace}(M^\dagger C^{\frac{iy}{2}} C^{-\frac{1}{2}} A C^{-\frac{1}{2}} C^{-\frac{iy}{2}} M) \text{trace}(A_k^{\frac{1}{2}} C^{-\frac{1}{2}} C^{-\frac{iy}{2}} M M^\dagger C^{-\frac{iy}{2}} C^{-\frac{1}{2}} A_k^{\frac{1}{2}}) \right]^{\frac{1}{2}} \end{aligned}$$

Noting that these two traces are the same, we obtain

$$|f_k(0 + iy)| \leq \text{trace}(M^\dagger C^{\frac{iy}{2}} C^{-\frac{1}{2}} A C^{-\frac{1}{2}} C^{-\frac{iy}{2}} M).$$

Putting it all together,

$$\begin{aligned} |\lambda_1 f_1(0 + iy) + \lambda_2 f_2(0 + iy)| &\leq \lambda_1 |f_1(0 + iy)| + \lambda_2 |f_2(0 + iy)| \\ &\leq \text{trace}(M^\dagger C^{\frac{iy}{2}} C^{-\frac{1}{2}} (\lambda_1 A_1 + \lambda_2 A_2) C^{-\frac{1}{2}} C^{-\frac{iy}{2}} M) \\ &= \text{trace}(M^\dagger M). \end{aligned}$$

The case of $f(1 + iy)$ is similar. \square

Theorem 5.4.13. *The relative entropy $H(\rho||\sigma)$ is jointly convex in ρ, σ .*

Proof. Let $A, X \in \text{End}(\mathcal{H})$ with A positive. Define

$$I_t(A, X) := \text{trace}(X^\dagger A^t X A^{1-t}) - \text{trace}(X^\dagger X A),$$

where the second term is just to normalize $I_0(X, A) = 0$. Lieb's lemma implies the first term is concave in A , and since the second is linear in A , $I_t(A, X)$ is concave in A . Write $I'_0(A, X) = \frac{d}{dt}|_{t=0} I_t(A, X)$. Recall that

$\frac{d}{dt}|_{t=0}a^t = \ln(a)$, and the same holds for positive operators, so $I'_0(A, X) = \text{trace}(X^\dagger \ln(A)XA) - \text{trace}(X^\dagger X \ln(A)A)$. Consider

$$\begin{aligned} I'_0(\lambda_1 A_1 + \lambda_2 A_2, X) &= \lim_{\epsilon \rightarrow 0} \frac{I_\epsilon(\lambda_1 A_1 + \lambda_2 A_2, X) - 0}{\epsilon} \\ &\geq \lim_{\epsilon \rightarrow 0} \lambda_1 \frac{I_\epsilon(A_1, X)}{\epsilon} + \lambda_2 \frac{I_\epsilon(A_2, X)}{\epsilon} \quad \text{by concavity} \\ &= \lambda_1 I'_0(A_1, X) + \lambda_2 I'_0(A_2, X), \end{aligned}$$

so $I'_0(A, X)$ is a concave function of A . Now let $A = \begin{pmatrix} \rho & 0 \\ 0 & \sigma \end{pmatrix}$ and $X = \begin{pmatrix} 0 & 0 \\ \text{Id} & 0 \end{pmatrix}$, we obtain

$$I'_0(A, X) = -H(\rho||\sigma)$$

and we conclude $H(\rho||\sigma)$ is jointly convex in ρ, σ . \square

Corollary 5.4.14. *The conditional entropy $H(\rho_A|\rho_B)$ is concave in ρ_{AB} .*

Proof. Let $d = \dim \mathcal{H}_A$, consider

$$H(\rho_{AB}||\frac{1}{d} \text{Id}_{\mathcal{H}_A} \otimes \rho_B) = -H(\rho_{AB}) - \text{trace}_{\mathcal{H}_A \otimes \mathcal{H}_B}(\rho_{AB} \log(\frac{1}{d} \text{Id}_{\mathcal{H}_A} \otimes \rho_B)).$$

Exercise 5.4.6 implies that

$$\begin{aligned} \text{trace}_{\mathcal{H}_A \otimes \mathcal{H}_B}(\rho_{AB} \log(\frac{1}{d} \text{Id}_{\mathcal{H}_A} \otimes \rho_B)) &= -\text{trace}(\rho_B \log(\rho_B) + \log(d)) \\ &= -H(\rho_B) + \log(d) \end{aligned}$$

so

$$H(\rho_A|\rho_B) = \log(d) - H(\rho_{AB}||\frac{1}{d} \text{Id}_{\mathcal{H}_A} \otimes \rho_B)$$

and the second term is jointly concave in both its arguments, so we conclude. \square

Finally we have

Theorem 5.4.15 (Strong sub-additivity). *Let ρ_{ABC} be a density operator on $\mathcal{H}_A \otimes \mathcal{H}_B \otimes \mathcal{H}_C$. Then*

$$(5.4.4) \quad H(\rho_C|\rho_A) + H(\rho_C|\rho_B) \geq 0$$

i.e.,

$$(5.4.5) \quad H(\rho_{AC}) - H(\rho_A) + H(\rho_{CB}) - H(\rho_B) \geq 0.$$

Proof. Let $f(\rho_{ABC}) = H(\rho_C|\rho_A) + H(\rho_C|\rho_B)$, which is a concave function of ρ_{ABC} because it is a sum of two concave functions of ρ_{ABC} . Write

$$\rho_{ABC} = \sum \lambda_i |v_i\rangle\langle v_i|.$$

The concavity of f implies $f(\rho_{ABC}) \geq \sum \lambda_i f(|v_i\rangle\langle v_i|) = 0$ because $|v_i\rangle\langle v_i|$ is pure, and we conclude $f(\rho_{ABC}) \geq 0$. \square

Corollary 5.4.16. *Let ρ_{ABC} be a density operator on $\mathcal{H}_A \otimes \mathcal{H}_B \otimes \mathcal{H}_C$. Then*

$$(5.4.6) \quad H(\rho_{ABC}) - [H(\rho_{AB}) + H(\rho_{BC})] + H(\rho_B) \geq 0.$$

Proof. Introduce a purification ρ_{ABCR} of ρ_{ABC} , $\rho_{ABC} = \text{trace}_{\mathcal{H}_R}(\rho_{ABCR})$. Then by (5.4.5) $H(\rho_R) + H(\rho_B) \leq H(\rho_{CR}) + H(\rho_{BC})$, but since ρ_{ABCR} is pure $H(\rho_R) = H(\rho_{ABC})$ and $H(\rho_{CR}) = H(\rho_{AB})$ and we conclude. \square

Applications:

Corollary 5.4.17. *[Conditional Entropy is non-increasing under further conditioning] Let ρ_{ABC} be a density operator on $\mathcal{H}_A \otimes \mathcal{H}_B \otimes \mathcal{H}_C$, then $H(\rho_A|\rho_{BC}) \leq H(\rho_A|\rho_B)$.*

Exercise 5.4.18: Prove Corollary 5.4.17.

Corollary 5.4.19. *[Mutual information is non-increasing when quantum systems are discarded] Let ρ_{ABC} be a density operator on $\mathcal{H}_A \otimes \mathcal{H}_B \otimes \mathcal{H}_C$, then $\text{Muinfo}(\rho_A : \rho_B) \leq \text{Muinfo}(\rho_A : \rho_{BC})$.*

Exercise 5.4.20: Prove Corollary 5.4.19.

Theorem 5.4.21 (Sub-additivity of conditional entropy).

(1) *Let ρ_{ABCD} be a density operator on $\mathcal{H}_A \otimes \mathcal{H}_B \otimes \mathcal{H}_C \otimes \mathcal{H}_D$, then*

$$H(\rho_{AB}|\rho_{CD}) \leq H(\rho_A|\rho_C) + H(\rho_B|\rho_D).$$

(2) *Let ρ_{ABC} be a density operator on $\mathcal{H}_A \otimes \mathcal{H}_B \otimes \mathcal{H}_C$, then*

$$H(\rho_{AB}|\rho_C) \leq H(\rho_A|\rho_C) + H(\rho_B|\rho_C)$$

$$H(\rho_A|\rho_{BC}) \leq H(\rho_A|\rho_B) + H(\rho_A|\rho_C)$$

Proof. Proof of (1): Considering $\mathcal{H}_B \otimes \mathcal{H}_C$ as a single space, (5.4.6) implies

$$H(\rho_{ABCD}) + H(\rho_C) \leq H(\rho_{AC}) + H(\rho_{BCD}).$$

Now add $H(\rho_D)$ to both sides and observe $H(\rho_{BCD}) + H(\rho_D) \leq H(\rho_{AC}) + H(\rho_{BD}) + H(\rho_{BC})$, substitute this in and rearrange to obtain the result.

The first assertion in (2) is just (5.4.6). The second is equivalent to

$$H(\rho_{ABC}) + H(\rho_B) + H(\rho_C) \leq H(\rho_{AB}) + H(\rho_{BC}) + H(\rho_{AC}).$$

By strong sub-additivity, at least one of the inequalities $H(\rho_C) \leq H(\rho_{AC})$, $H(\rho_B) \leq H(\rho_{AB})$ must hold. Say the first does, adding it to the equation gives the result. \square

We also obtain more evidence that the relative entropy is a useful distance measure:

Corollary 5.4.22. [Monotonicity of the relative entropy] Let ρ_{AB}, σ_{AB} be density operators on $\mathcal{H}_A \otimes \mathcal{H}_B$. Then

$$H(\rho_A || \sigma_A) \leq H(\rho_{AB} || \sigma_{AB}).$$

Corollary 5.4.22 says that ignoring part of a physical system makes it harder to distinguish the state of the system.

5.5. Entanglement and LOCC

We have seen several ways that *entanglement* is a resource already for the space $\mathcal{H}_A \otimes \mathcal{H}_B = \mathbb{C}^2 \otimes \mathbb{C}^2$: given a shared qubit $|epr\rangle = \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle)$, one can transport two bits of classical information using only one qubit (“super dense coding”) and one can also transmit one qubit of quantum information from Alice to Bob by sending two classical bits (“teleportation”).

In this section we study entanglement as a resource. Unlike quantities such as quantum channel capacity and von Neumann entropy, there is no obvious classical cousin of entanglement. For pure states $|\psi\rangle\langle\psi| \in \text{End}(\mathcal{H}_1 \otimes \cdots \otimes \mathcal{H}_n)$, ψ was defined to be entangled if it is not a decomposable (i.e. rank one) tensor. We have not yet defined what it means for a mixed state to be entangled. This will be rectified by Definition 5.5.18.

In this section we will assume several different laboratories can communicate classically, have prepared some shared states in advance, and can perform unitary and projection operations on their parts of the states, as was the situation for quantum teleportation. More precisely, we make the following assumptions:

- $\mathcal{H} = \mathcal{H}_1 \otimes \cdots \otimes \mathcal{H}_n$, and the \mathcal{H}_j share an entangled state ρ . Often we will just have $\mathcal{H} = \mathcal{H}_A \otimes \mathcal{H}_B$ and $\rho = \alpha|00\rangle + \beta|11\rangle$.
- The laboratories can communicate classically.
- Each laboratory is allowed to perform unitary and measurement operations on their own spaces.

The above assumptions are called *LOCC* for “local operations and classical communication”. It generalizes the set-up for teleportation §2.4.

Restrict to the case $\mathcal{H} = \mathcal{H}_A \otimes \mathcal{H}_B$, each of dimension two. We will use $|epr\rangle$ as a benchmark for measuring the quality of entanglement.

We will not be concerned with a single state $|\psi\rangle$, but the tensor product of many copies of it, $|\psi\rangle^{\otimes n} \in (\mathcal{H}_A \otimes \mathcal{H}_B)^{\otimes n}$. We ask “how much” entanglement does $|\psi\rangle^{\otimes n}$ have? More precisely, how many copies of $|epr\rangle$ (or any $|\phi\rangle$) can we construct from it via LOCC? We will be content with output that is “close” to many copies of $|\psi\rangle$, i.e., has high fidelity with $|\psi\rangle^{\otimes n}$.

To gain insight as to which states can be produced via LOCC from a given density operator, we return to the classical case. For the classical cousin of LOCC, by considering diagonal density operators, we see we should allow alteration of a probability distribution by permuting the p_j (permutation matrices are unitary), and more generally averaging our probability measure under some probability measure on elements of \mathfrak{S}_d (the classical cousin of a projective measurement), i.e., we should allow

$$(5.5.1) \quad \bar{p} \mapsto \sum_{\sigma \in \mathfrak{S}_d} q_\sigma \mu(\sigma) \bar{p}$$

where $\mu : \mathfrak{S}_d \rightarrow GL_d$ is the representation, and q is a probability distribution on \mathfrak{S}_d .

This is because the unitary and projection local operators allowed amount to

$$\rho \mapsto \sum_{j=1}^k p_j U_j \rho U_j^{-1}$$

where the U_j are unitary and p is a probability distribution on $\{1, \dots, k\}$ for some finite k .

Recall from Exercise 4.2.3 that Shannon entropy is non-increasing under an action of the form (5.5.1). We want to understand the partial order on probability distributions determined by (5.5.1).

5.5.1. A partial order on probability distributions compatible with entropy. Given probability distributions p, q on $\{1, \dots, d\}$, consider the corresponding vectors $\bar{p}, \bar{q} \in \mathbb{R}^d$. Recall that the entropy $H(\bar{p})$ gives a measure of the uncertainty of a probability distribution, telling us how much information we can transfer along a channel if it sends out signals according to the distribution p . We would like a partial order on distributions that is compatible with the degenerations (5.5.1), and thus with entropy. The *dominance order* satisfies these properties:

Definition 5.5.1. Let $x, y \in \mathbb{R}^d$, write x^\downarrow for x re-ordered such that $x_1 \geq x_2 \geq \dots \geq x_d$. We say $x \prec y$ if for all $k \leq d$, $\sum_{j=1}^k x_j^\downarrow \leq \sum_{j=1}^k y_j^\downarrow$. The partial order determined by \prec is called the *dominance order*.

Note that if p is a probability distribution concentrated at a point, then $\bar{q} \prec \bar{p}$ for all probability distributions q , and if p is such that $p_j = \frac{1}{d}$ for all j , then $\bar{p} \prec \bar{q}$ for all q , and more generally the dominance order is compatible with the entropy in the sense that $\bar{p} \prec \bar{q}$ implies $H(\bar{p}) \geq H(\bar{q})$.

We will show that p can degenerate to q in the sense of (5.5.1) if and only if $\bar{p} \prec \bar{q}$.

Exercise 5.5.2: Fix $y \in \mathbb{R}^d$. Show that the set $\{x \mid x \prec y\}$ is convex.

Now observe that the set of matrices of the form $\sum_{\sigma \in \mathfrak{S}_d} p_\sigma \mu(\sigma)$ is the convex hull of $\mu(\mathfrak{S}_d)$, which we will denote $\text{conv}(\mu(\mathfrak{S}_d))$, where in general, $z \in \text{conv}\{w_1, \dots, w_k\}$ means $z = \sum_{j=1}^k p_j w_j$, for some probability distribution on $\{1, \dots, k\}$.

Recall that a matrix $D \in \text{Mat}_{d \times d}$ is doubly stochastic if $D_{ij} \geq 0$ and all column and row sums equal one. Let $\mathcal{DS}_d \subset \text{Mat}_{d \times d}$ denote the set of doubly stochastic matrices.

Theorem 5.5.3. (*G. Birkoff 1946 [Bir46]*) $\mathcal{DS}_d = \text{conv}(\mu(\mathfrak{S}_d))$.

Note that it is clear $\text{conv}(\mu(\mathfrak{S}_d)) \subseteq \mathcal{DS}_d$.

Theorem 5.5.4. (*Hardy-Littlewood-Polya [HLP52]*) $\{x \mid x \prec y\} = \mathcal{DS}_d \cdot y$.

Note that it is clear $\mathcal{DS}_d \cdot y \subseteq \{x \mid x \prec y\}$.

We prove both theorems following [Mir58].

Lemma 5.5.5. *Let $D \in \mathcal{DS}_d \setminus \text{Id}$. Then there exists $\sigma \in \mathfrak{S}_d \setminus \text{Id}$ such that the diagonal elements of $\mu(\sigma)A$ are all nonzero.*

Proof. Say not, so for all $\sigma \in \mathfrak{S}_d \setminus \text{Id}$ there exists some k such that $A_{j, \sigma(j)} = 0$. Then $\det(A) = \sum_{\sigma \in \mathfrak{S}_d} \text{sgn}(\sigma) A_{1\sigma(1)} \cdots A_{d\sigma(d)} = A_{11} \cdots A_{dd}$. Similarly, $\det(A + t \text{Id}) = \text{Proj}_j(A_{jj} + t)$ so the eigenvalues of A are the diagonal elements. But at least one eigenvalue is 1 because of the eigenvector $(1, \dots, 1)^T$. Say this is A_{jj} . Then strike out the j -th row and column of A and apply the argument again, continuing, we see $A = \text{Id}$, a contradiction. \square

Exercise 5.5.6: Show that $x \in \text{conv}\{y_1, \dots, y_m\}$ if and only if for all $a \in \mathbb{R}^d$, $\langle a \mid x \rangle \leq \max_{\alpha \in \{1, \dots, m\}} \langle a \mid y_\alpha \rangle$. **give hint**

Exercise 5.5.7: Prove Theorem 5.5.4.

Exercise 5.5.8: Show that for all $A \in \text{Mat}_{d \times d}$, $\sup_{D \in \mathcal{DS}_d} \text{trace}(DA) = \max_{\sigma \in \mathfrak{S}_d} (\text{trace}(\mu(\sigma)A))$. Hint: assume wlog the RHS is $\text{trace} A$ as both sides are invariant under permutation matrices.

Exercise 5.5.9: Using Exercise 5.5.6 with $\{y_1, \dots, y_m\}$ the set of permutation matrices, prove Theorem 5.5.3.

5.5.2. Dominance order for Hermitian operators. Let X, Y be Hermitian operators and write $\text{spec}(X)$ for the set of eigenvalues of X (the spectrum of X). We will say $X \prec Y$ if $\text{spec}(X) \prec \text{spec}(Y)$.

Theorem 5.5.10. *Let X, Y be Hermitian operators on $\mathcal{H} = \mathbb{C}^d$. Then $Y \prec X$ if and only if there exists a probability distribution on \mathfrak{S}_d and unitary matrices $U_\sigma \in \mathbf{U}(\mathcal{H})$ such that $X = \sum_{\sigma \in \mathfrak{S}_d} p_\sigma U_\sigma Y U_\sigma^{-1}$.*

Proof. Say $Y \prec X$ so that $\text{spec}(X) = \sum_{\sigma \in \mathfrak{S}_d} p_\sigma \mu(\sigma) \text{spec}(Y)$. Write $X = U \Lambda(X) U^{-1}$ where U is unitary and $\Lambda(X)$ is a diagonal matrix with

the eigenvalues of X on the diagonal and similarly $Y = V\Lambda(Y)V^{-1}$. By hypothesis $\Lambda(X) = \sum_{\sigma \in \mathfrak{S}_d} p_\sigma \mu(\sigma) \Lambda(Y) \mu(\sigma)^{-1}$, so

$$X = U^{-1} \left(\sum_{\sigma} p_\sigma \mu(\sigma) V^{-1} Y V \mu(\sigma)^{-1} \right) U$$

so just set $U_\sigma = U^{-1} \mu(\sigma) V^{-1}$.

For the other direction, we have

$$U\Lambda(X)U^{-1} = \sum_{\sigma} p_\sigma U_\sigma V \Lambda(Y) V^{-1} U_\sigma^{-1}$$

i.e.,

$$\Lambda(X) = \sum_{\sigma} p_\sigma (UU_\sigma V) \Lambda(Y) (UU_\sigma V)^{-1}.$$

Write $W_\sigma = UU_\sigma V$. Since $\Lambda(X), \Lambda(Y)$ are diagonal, conjugation by W_σ must take diagonal matrices to diagonal matrices. If the eigenvalues of X are distinct, then W_σ must be a permutation matrix, and in general, without loss of generality, we may assume it to be so. \square

Exercise 5.5.11: Let $X \in \text{End}(\mathbb{C}^d)$ be Hermitian. Show that there exists a probability distribution p_j and $U_j \in \mathbf{U}_d$ such that $\sum_i p_i U_i X U_i^{-1} = \frac{1}{d} \text{trace}(X) \text{Id}_{\mathbb{C}^d}$. Hint: find U that diagonalizes X then apply permutations to average.

5.5.3. A reduction theorem. The study of LOCC is potentially unwieldy because there can be numerous rounds of local operations and classical communication, making it hard to model. The following result eliminates this problem:

Proposition 5.5.12. *If $|\psi\rangle \in \mathcal{H}_A \otimes \mathcal{H}_B$ can be transformed into $|\phi\rangle$ by LOCC, then it can be transformed to $|\phi\rangle$ by the following sequence of operations:*

- (1) Alice performs a single measurement with operators Proj_{M_j} .
- (2) She sends the result of her measurement (some j) to Bob classically.
- (3) Bob performs a unitary operation on his system.

Proof. The key point is that for any vector spaces V, W , an element $f \in V \otimes W$, may be considered as a linear map $W^* \rightarrow V$. In our case, $\mathcal{H}_B^* \simeq \mathcal{H}_B$ so $|\psi\rangle$ induces a linear map $\mathcal{H}_B \rightarrow \mathcal{H}_A$ which gives us the mechanism to transfer Bob's measurements to Alice.

Now for the details. Write $\mathcal{H}_A = \text{Lker } |\psi\rangle \oplus \mathcal{H}'_A$ and $\mathcal{H}_B = \text{Rker } |\psi\rangle \oplus \mathcal{H}'_B$, where $\text{Lker } |\psi\rangle \subset \mathcal{H}_A$, $\text{Rker } |\psi\rangle \subset \mathcal{H}_B$ are the kernels of the induced linear maps. Thus $|\psi\rangle$ induces an isomorphism $iso_\psi : \mathcal{H}'_B \rightarrow \mathcal{H}'_A$, which in turn induces an isomorphism $endis_\psi : \text{End}(\mathcal{H}'_B) \rightarrow \text{End}(\mathcal{H}'_A)$, which we may extend by zero to a map $end_\psi : \text{End}(\mathcal{H}_B) \rightarrow \text{End}(\mathcal{H}_A)$. Now say $M \subset \mathcal{H}_B$.

Note that $\text{Id}_{\mathcal{H}_A} \otimes \text{Proj}_M |\psi\rangle = \text{Id}_{\mathcal{H}_A} \otimes \text{Proj}_M |\mathcal{H}'_B\rangle |\psi\rangle$ (the right hand side makes sense because $|\psi\rangle \in \mathcal{H}'_B$).

Write the singular value decomposition $|\psi\rangle = \sum_{\mu} \sqrt{\lambda_{\mu}} |v_{\mu}\rangle \otimes |w_{\mu}\rangle$ with $|v_{\mu}\rangle \in \mathcal{H}'_A$, $|w_{\mu}\rangle \in \mathcal{H}'_B$. Let $M_{\mu,\nu}$ denote the matrix for π_M in the basis $\{|w_{\mu}\rangle\}$. On Bob's side,

$$\begin{aligned} \text{Id}_{\mathcal{H}_A} \otimes \text{Proj}_M |\psi\rangle &= \sum_{\mu,\nu} \sqrt{\lambda_{\mu}} |v_{\mu}\rangle \otimes M_{\mu,\nu} |w_{\nu}\rangle \\ (5.5.2) \qquad \qquad \qquad &= \sum_{\mu,\nu} \sqrt{\lambda_{\mu}} M_{\mu,\nu} |v_{\mu}\rangle \otimes |w_{\nu}\rangle. \end{aligned}$$

On the other hand

$$(5.5.3) \qquad \text{Proj}_{\text{iso}_{\psi}M} \otimes \text{Id}_{\mathcal{H}_B} |\psi\rangle = \sum_{\mu,\nu} \sqrt{\lambda_{\mu}} (M_{\mu,\nu} |\nu_A\rangle) \otimes |\mu_B\rangle.$$

The expressions (5.5.2),(5.5.3) are elements of $\mathcal{H}_A \otimes \mathcal{H}_B$ with the same singular values, so there exist $U_A \in \mathbf{U}(\mathcal{H}_A)$, $U_B \in \mathbf{U}(\mathcal{H}_B)$ such that $\text{Id}_{\mathcal{H}_A} \otimes \text{Proj}_M |\psi\rangle = U_A \otimes U_B \cdot \text{Proj}_{\text{iso}_{\psi}M} \otimes \text{Id}_{\mathcal{H}_B} |\psi\rangle$. So we may effect Bob's Proj_M by Alice's $\text{Proj}_{\text{iso}_{\psi}M}$, followed by isometries U_A and U_B . So we may get rid of all of Bob's measurements, and Bob's communication with Alice (since she has the result!). Finally, Alice's local actions commute with the local actions of Bob, so we can just combine all of Alice's local actions into one, then she sends the results all at once to Bob, who combines all his unitary actions to a single unitary operator. \square

Now we can state the main theorem on LOCC:

Theorem 5.5.13. [Nie99] $|\psi\rangle \rightsquigarrow |\phi\rangle$ by LOCC if and only if $\text{singvals}(|\psi\rangle) \prec \text{singvals}(|\phi\rangle)$.

Recall that $\text{singvals}(|\psi\rangle) = \text{spec}(\rho_{\psi,A}) = \text{spec}(\rho_{\psi,B})$.

Exercise 5.5.14: Given $A \in \text{End}(\mathcal{H})$, show that there exists $U \in \mathbf{U}(\mathcal{H})$ such that $A = U\sqrt{A^{\dagger}A} = \sqrt{AA^{\dagger}}U$, and that if A is invertible, U is unique. Hint: $A^{\dagger}A$ is Hermitian so it has a spectral decomposition.

Proof. ψ can be transformed to ϕ by LOCC means that there exist $M_j \subset \mathcal{H}_A$ giving an orthogonal decomposition, and a probability distribution p_j , such that

$$\text{Proj}_{M_j} \rho_{\psi,A} \text{Proj}_{M_j}^{\dagger} = p_j \rho_{\phi,A}.$$

By Exercise 5.5.14, there exist $U_j \in \mathbf{U}(\mathcal{H}_A)$ such that

$$\begin{aligned} \text{Proj}_{M_j} \sqrt{\rho_{\psi,A}} &= \sqrt{\text{Proj}_{M_j} \sqrt{\rho_{\psi,A}} (\text{Proj}_{M_j} \sqrt{\rho_{\psi,A}})^{\dagger}} U_j \\ &= \sqrt{p_j \rho_{\phi,A}} U_j \end{aligned}$$

and similarly with $\sqrt{\rho_{\psi,A}} \text{Proj}_{M_j}^\dagger$, so

$$\sqrt{\rho_{\psi,A}}^\dagger \text{Proj}_{M_j} \sqrt{\rho_{\psi,A}} = p_j U_j^{-1} \rho_{\phi,A} U_j.$$

Now sum on j , the projections sum to the identity and we conclude

$$\rho_{\psi,A} = \sum_j p_j U_j^{-1} \rho_{\phi,A} U_j$$

which means $\text{spec}(\rho_{\psi,A}) \prec \text{spec}(\rho_{\phi,A})$. If ρ_{ψ} is invertible, the argument can be run in the other direction to get the reverse conclusion, if not, one splits $\mathcal{H}_A, \mathcal{H}_B$ as in the proof of Proposition 5.5.12 and the argument still goes through. \square

Exercise 5.5.15: (Entanglement catalysis) Say $\mathcal{H}_A, \mathcal{H}_B$ are four dimensional and Alice and Bob share $|\psi\rangle = \sqrt{.4}(|00\rangle + |11\rangle) + \sqrt{.1}(|22\rangle + |33\rangle)$. Show that $|\psi\rangle$ cannot be degenerated to $|\phi\rangle = \sqrt{.5}|00\rangle + \sqrt{.25}|11\rangle + \sqrt{.25}|22\rangle$. But now say a bank is willing to loan them $|c\rangle = \sqrt{.6}|00\rangle + \sqrt{.4}|11\rangle$. Show that $|\psi\rangle \otimes |c\rangle$ can be degenerated by LOCC to $|\phi\rangle \otimes |c\rangle$, so they can obtain $|\phi\rangle$ and return $|c\rangle$ to the bank. In this context, $|c\rangle$ is called a ‘‘catalyst’’.

5.5.4. Entanglement distillation (concentration) and dilution. To compare the entanglement resources of two states $|\phi\rangle$ and $|\psi\rangle$, we will consider $|\phi\rangle^{\otimes m}$ for large m , and determine the largest $n = n(m)$ such that $|\phi\rangle^{\otimes m}$ may be degenerated to $|\psi\rangle^{\otimes n}$ via LOCC. Due to the approximate and probabilistic nature of quantum computing, we will be content to degenerate $|\phi\rangle^{\otimes m}$ to a state that has high fidelity with $|\psi\rangle^{\otimes n}$.

There is a subtlety for this question worth pointing out. Teleportation was defined in such a way that Alice did not need to know the state she was teleporting, but for distillation and dilution, she will need to know its right singular vectors are standard basis vectors. More precisely, if she is in possession of $|\psi\rangle = \sqrt{p_1}|v_1\rangle \otimes |1\rangle + \sqrt{p_2}|v_2\rangle \otimes |2\rangle$, she can teleport the second half of it to Bob if they share $|epr\rangle \in \mathcal{H}_A \otimes \mathcal{H}_B$. More generally, we will see that if she is in possession of $|\psi\rangle = \sum_{j=1}^d \sqrt{p_j}|v_j\rangle \otimes |j\rangle \in \mathcal{H}_{A'} \otimes \mathcal{H}_{A''}$, she can teleport it to Bob if they share enough EPR states. In most textbooks, Alice is assumed to possess states whose singular vectors are $|jj\rangle$'s and we will follow that convention here. Similarly, if $|\psi\rangle = \sum_{j=1}^d \sqrt{p_j}|jj\rangle \in \mathcal{H}_A \otimes \mathcal{H}_B$, we will discuss how many shared EPR states they can construct from a shared $|\psi\rangle^{\otimes m}$.

We define the *entanglement cost* $E_C(\psi)$ to be $\inf_m \frac{n(m)}{m}$ where $n(m)$ copies of ψ can be constructed from $|epr\rangle^{\otimes m}$ by LOCC with fidelity going to 1 as $m \rightarrow \infty$. Similarly, define the *entanglement value*, or *distillable entanglement* $E_V(\psi)$ to be $\sup_m \frac{n(m)}{m}$ where $n(m)$ copies of $|epr\rangle$ can be constructed from $|\psi\rangle^{\otimes m}$ by LOCC.

Since entanglement cannot be created by LOCC, $E_V(\psi) \leq E_C(\psi)$. Otherwise by going through rounds of LOCC, one could construct an arbitrary number of EPR states. ****better mathematical justification?***** We will show that (asymptotically) cost equals value.

Say $|\psi\rangle = \sqrt{p_1}|11\rangle + \dots + \sqrt{p_d}|dd\rangle \in \mathcal{H}_{A'} \otimes \mathcal{H}_{A''}$. Consider

$$|\psi\rangle^{\otimes n} = \sum \sqrt{p_{i_1} \cdots p_{i_n}} |i_1 \cdots i_n\rangle \otimes |i_1 \cdots i_n\rangle.$$

Project $|\psi\rangle^{\otimes n}$ to the ϵ -typical subspace for some small ϵ . Recall that this subspace has dimension at most $2^{n(H(|\psi\rangle\langle\psi|) + \epsilon)}$.

Now $|\psi\rangle_{\epsilon\text{-typ}}^{\otimes n}$ can be teleported using $2^{n(H(|\psi\rangle\langle\psi|) + \epsilon)}$ classical bits and a pre-shared $|epr\rangle^{\otimes n(H(|\psi\rangle\langle\psi|) + \epsilon)}$ with probability of error at most δ . However we can make ϵ, δ as small as we want, so we conclude $E_C(\psi) \leq H(|\psi\rangle\langle\psi|)$.

Now say Alice and Bob share $|\psi\rangle^{\otimes m} \in \mathcal{H}_A^{\otimes m} \otimes \mathcal{H}_B^{\otimes m}$ and they want to construct $|epr\rangle^{\otimes n}$ for some $n = n(m)$. For simplicity, assume $|\psi\rangle = \sqrt{p}|00\rangle + \sqrt{1-p}|11\rangle$. Project $|\psi\rangle^{\otimes m}$ onto the ϵ -typical subspace. The largest coefficient is $2^{-m(H(|\psi\rangle\langle\psi|) - \epsilon)}$ and after renormalization to have a vector of length one, this coefficient grows at most by a factor of $\frac{1}{\sqrt{1-\delta}}$. Take any n such that

$$(5.5.4) \quad 2^{-n} \geq \frac{2^{-m(H(|\psi\rangle\langle\psi|) - \epsilon)}}{1 - \delta}.$$

Then

$$\text{spec}(|\psi\rangle\langle\psi|_{\epsilon\text{-typ}}^{\otimes m}) \prec (2^{-n}, \dots, 2^{-n}),$$

so we can create $|epr\rangle^{\otimes n}$ by LOCC from $|\psi\rangle\langle\psi|_{\epsilon\text{-typ}}^{\otimes m}$. Further note that if $n \sim mH(|\psi\rangle\langle\psi|)$, then (5.5.4) will hold. We conclude $E_V(\psi) \geq H(|\psi\rangle\langle\psi|)$.

Putting together the inequalities $H(|\psi\rangle\langle\psi|) \leq E_V(\psi) \leq E_C(\psi) \leq H(|\psi\rangle\langle\psi|)$, we see they are all equalities.

Remark 5.5.16. In classical computation one can reproduce information, but this cannot be done with quantum information in general. This is because the map $|\psi\rangle \mapsto |\psi\rangle \otimes |\psi\rangle$, called the *Veronese map* in algebraic geometry, is not a linear map. This observation is called the *no cloning theorem* in the quantum literature. However, one can define a linear map, e.g., $\mathbb{C}^2 \rightarrow \mathbb{C}^2 \otimes \mathbb{C}^2$ that duplicates basis vectors, i.e., $|0\rangle \mapsto |0\rangle \otimes |0\rangle$ and $|1\rangle \mapsto |1\rangle \otimes |1\rangle$. But then of course $\alpha|0\rangle + \beta|1\rangle \mapsto \alpha|0\rangle \otimes |0\rangle + \beta|1\rangle \otimes |1\rangle \neq (\alpha|0\rangle + \beta|1\rangle)^{\otimes 2}$.

Remark 5.5.17. The level of noise in a classical channel can support before it becomes useless is higher than the level of noise a quantum channel can support before it becomes useless ****ref****. However, via LOCC, one can raise the admissible level of noise of a quantum channel and still have it useful... ****more detail here****

5.5.5. Cost and Value of mixed states. For mixed states ρ on $\mathcal{H}_A \otimes \mathcal{H}_B$, we can still define $E_C(\rho)$ and $E_V(\rho)$, but there exist examples where they differ, so there is not a canonical measure of entanglement. ****give example**** In fact, at this point we still don't even have a definition of what it means for a mixed state to be entangled.

Let's make a wish list of what we might want from an entanglement measure E

- Non-increasing under LOCC.
- If ρ is a product state, i.e., $\rho = |\phi_A\rangle\langle\phi_A| \otimes |\psi_B\rangle\langle\psi_B|$, then $E(\rho) = 0$.

The two conditions together imply any state constructible from such a ρ by LOCC should also have zero entanglement. Hence the following definition:

Definition 5.5.18. A density operator $\rho \in \text{End}(\mathcal{H}_1 \otimes \cdots \otimes \mathcal{H}_n)$ is *separable* if $\rho = \sum_i p_i \rho_{i,1} \otimes \cdots \otimes \rho_{i,n}$, where $\rho_{i,\alpha} \in \text{End}(\mathcal{H}_\alpha)$ are density operators, $p_i \geq 0$, and $\sum_i p_i = 1$. If ρ is not separable, we say ρ is *entangled*.

So we replace our second condition by requiring that $E(\rho) = 0$ for any separable ρ .

Finally, we would like any new entanglement measure to agree with E_C, E_V on pure states.

Definition 5.5.19. An *entanglement monotone* E is a function on density operators on $\mathcal{H}_A \otimes \mathcal{H}_B$ that is non-increasing under LOCC.

This implies E is zero on separable states. One would also like E to agree with E_C, E_V on pure states.

Recall the mutual information $\text{Muinfo}(\rho_A : \rho_B) := H(\rho_A) + H(\rho_B) - H(\rho_{AB})$. This vanishes on product states.

If ρ is pure, then $\text{Muinfo}(\rho_A : \rho_B) = 2H(\rho_A) = 2H(\rho_B)$, so a first idea would be to take half the mutual information.

However, the mutual information fails to be zero on separable states. To fix this, define the *squashed entanglement* [CW04]

$$E_{sq}(\rho_{AB}) := \inf_C \left\{ \frac{1}{2} \text{Muinfo}(A : B|C) \mid \rho_{AB} = \text{trace}_{\mathcal{H}_C}(\rho_{ABC}) \right\}.$$

Exercise 5.5.20: Show that E_{sq} vanishes on separable states.

The squashed entanglement is (at least as of this writing) hard to compute in general, but it does have the desirable property that it is additive under tensor products:

$$E_{sq}(\rho_{AB} \otimes \rho_{CD}) = E_{sq}(\rho_{AB}) + E_{sq}(\rho_{CD}).$$

add more and/or references

Representation theory and Quantum information

In this chapter we show how many results in quantum information theory can be understood in terms of representation theory. We also discuss the quantum marginal problem: what are the conditions on density operators $\rho_A \in \text{End}(\mathcal{H}_A)$, $\rho_B \in \text{End}(\mathcal{H}_B)$, $\sigma \in \text{End}(\mathcal{H}_A \otimes \mathcal{H}_B)$, such that $\rho_A = \text{trace}_{\mathcal{H}_B}(\sigma)$ and $\rho_B = \text{trace}_{\mathcal{H}_A}(\sigma)$?

We begin with a crash course in representation theory.

6.1. Representation theory

6.1.1. Basic definitions. We will be primarily concerned with the representation theory of the permutation group \mathfrak{S}_d and the general linear group $GL(V)$. Informally, a representation of a group G is a realization of G as a subgroup of the group of $n \times n$ matrices for some n .

Definition 6.1.1. Let G be a group. A *representation* of G is a group homomorphism $\mu : G \rightarrow GL(V)$ for some vector space V . One says G *acts* on V and that V is a *G -module*.

For example \mathfrak{S}_d acts on \mathbb{C}^d by permuting basis vectors and extending the action linearly.

Definition 6.1.2. Let V be a G -module, and let $W \subset V$ be a proper subspace. We say W is a *submodule* if for all $g \in G$ and $w \in W$, $\mu(g)w \in W$. The representation V is said to be *irreducible* if it has no proper submodules.

For example the action of \mathfrak{S}_d on \mathbb{C}^d is not irreducible.

Exercise 6.1.3: Write $\mathbb{C}^d = W_1 \oplus W_2$ where W_1, W_2 are \mathfrak{S}_d -submodules.

A G -module V is *trivial* if $\mu(g)v = v$ for all $g \in G$ and $v \in V$.

If V, W are G -modules, a linear map $f : V \rightarrow W$ commuting with the actions of G is called a *G -module map*. The modules V, W are said to be *isomorphic* if there exists a G -module map between them that is a linear isomorphism.

Lemma 6.1.4 (Schur's Lemma). *Let V, W be irreducible G -modules and $f : V \rightarrow W$ a G -module map. Then either $f = 0$ or f is an isomorphism. If furthermore $V = W$, then $f = \lambda \text{Id}$ for some constant λ .*

Exercise 6.1.5: Prove Schur's Lemma. Hint: show that the kernel and image of a G -module map are submodules.

If G acts on V by μ_V , it acts on V^* by μ_{V^*} , where $[\mu_{V^*}(g)(\alpha)](v) = \alpha(\mu_V(g)v)$ for all $\alpha \in V^*$, $v \in V$, $g \in G$. If G acts on V_1, V_2 , by μ_1, μ_2 , then it acts on $V_1 \otimes V_2$ by $\mu(g)(v_1 \otimes v_2) = \mu_1(g)v_1 \otimes \mu_2(g)v_2$. In particular, if G acts on V , it acts on all tensor powers of V and V^* . These actions are called *induced actions*.

Exercise 6.1.6: Show that if G acts on V , then the induced action on $\text{End}(V) = V \otimes V^*$ contains a trivial submodule. Show that moreover if V is irreducible, the trivial submodule is unique.

A basic problem is: given G , determine the irreducible G -modules up to isomorphism.

6.1.2. Representations of the permutation group and Schur-Weyl duality. We describe the irreducible modules for the permutation group. For a proof, see e.g., [Lan17, §8.6.8], [FH91, §I.4], or [Mac95, I.7].

Proposition 6.1.7. *The irreducible representations of \mathfrak{S}_d are in one to one correspondence with the partitions of d .*

To a partition $\pi = (p_1, \dots, p_d)$, let $[\pi]$ denote the corresponding irreducible \mathfrak{S}_d -module.

Example 6.1.8. $\pi = (d) = (d, 0, \dots, 0)$ corresponds to the one-dimensional trivial representation. $\pi = (1, \dots, 1) = (1^d)$ corresponds to the one-dimensional sign representation $\mu(\sigma)v = \text{sgn}(\sigma)v$. The partition $\pi = (d-1, 1)$ corresponds to the action on $\mathbb{C}^{d-1} \subset \mathbb{C}^d$ where \mathbb{C}^{d-1} is the subspace of vectors whose entries add to zero.

We will give a recipe for constructing all the irreducible \mathfrak{S}_d -modules as submodules of $V^{\otimes d}$, where $\dim V \geq d$.

First note that for any vector space V , \mathfrak{S}_d acts on $V^{\otimes d}$ by permuting the factors: $\mu(\sigma)(v_1 \otimes \cdots \otimes v_d) = v_{\sigma^{-1}(1)} \otimes \cdots \otimes v_{\sigma^{-1}(d)}$. One takes the inverse so that $\mu(\sigma)\mu(\tau)T = \mu(\sigma\tau)T$.

To visualize a partition π , define the *Young diagram* associated to π to be a collection of left-aligned boxes with p_j boxes in the the j -th row, as in Figure 6.1.1.






Figure 6.1.1. Young diagram for $\pi = (4, 2, 1, 1)$

Label the boxes in the diagram by $\{1, \dots, d\}$, such is called a *Young tableau without repetitions*. For example



The *default Young tableau* is labeled left to right and top to bottom:



Given any Young tableau without repetitions, consider the following projection operator on $V^{\otimes d}$ constructed from it: first write $V^{\otimes d} = V_1 \otimes V_2 \otimes \cdots \otimes V_d$, where the subscript is to remember the position. Then for each row in the Young tableau, symmetrize the corresponding copies of V , e.g., for $\pi = (4, 2, 1, 1)$ after the symmetrization, one obtains an element of $S^4 V \otimes S^2 V \otimes V \otimes V \subset V^{\otimes 8}$. Next, skew symmetrize along the columns.

For example, if we take the default Young tableau for $\pi = (2, 1)$, the maps are $v_i \otimes v_j \otimes v_k \mapsto v_i \otimes v_j \otimes v_k + v_j \otimes v_i \otimes v_k$ followed by $v_i \otimes v_j \otimes v_k + v_j \otimes v_i \otimes v_k \mapsto v_i \otimes v_j \otimes v_k - v_i \otimes v_k \otimes v_j + v_j \otimes v_i \otimes v_k - v_j \otimes v_k \otimes v_i$.

For the default Young tableau, write the resulting map as

$$P_{\pi-def} : V^{\otimes d} \rightarrow V^{\otimes d}$$

and let $S_{\pi-def} V := P_{\pi-def}(V^{\otimes d})$ denote the image.

Since $P_{\pi-def}$ is a $GL(V)$ -module map, $S_{\pi-def} V$ is a $GL(V)$ -module.

Fact : $S_{\pi-def}V$ is an irreducible $GL(V)$ -module, and if $\pi \neq \nu$, then $S_{\pi-def}V$ is not isomorphic to $S_{\nu-def}V$.

Let $V = \mathbb{C}^d$ be equipped with its standard basis e_1, \dots, e_d . Set

$$v_{\pi-def} := P_{\pi-def}(e_1^{\otimes p_1} \otimes e_2^{\otimes p_2} \otimes \dots \otimes e_d^{\otimes p_d}).$$

For any group G and any G -module W , the span of the G -orbit of any $v \in W$ is a submodule (or all of W).

Fact : The span of the \mathfrak{S}_d -orbit of $v_{\pi-def}$ is an irreducible \mathfrak{S}_d -module isomorphic to $[\pi]$.

If one defines the corresponding map for a different Young tableau without repetitions associated to π , one obtains an isomorphic $GL(V)$ -module. Let $S_\pi V$ denote the isomorphism class. Similarly, the \mathfrak{S}_d -module one obtains by a similar process for a different Young tableau without repetitions is also isomorphic to $[\pi]$.

Definition 6.1.9. If W, M are G -modules with M irreducible, the *isotypic component* of M in W is the largest submodule of W isomorphic to $M^{\oplus m}$ for some m . The integer m is called the *multiplicity* of M in W .

Let $\sigma \in \mathfrak{S}_d$, the map $\mu(\sigma) : V^{\otimes d} \rightarrow V^{\otimes d}$ is a $GL(V)$ -module map, as is $\mu(\sigma)|_{S_{\pi-def}V}$. Thus by Schur's lemma, its image is either zero or a module isomorphic to $S_\pi V$. It is clearly not zero as $\mu(\sigma)$ is an isomorphism.

Fact : The span of $\mu(\mathfrak{S}_d)S_{\pi-def}V$ is the isotypic component of $S_\pi V$ in $V^{\otimes d}$.

Similarly, for $g \in GL(V)$, the image under g of the span of $\mathfrak{S}_d \cdot v_{\pi-def}$ is an \mathfrak{S}_d -module isomorphic to $[\pi]$.

All these observations are consequences of:

Theorem 6.1.10 (Schur-Weyl duality). *As a $GL(V) \times \mathfrak{S}_d$ -module,*

$$V^{\otimes d} = \bigoplus_{|\pi|=d} S_\pi V \otimes [\pi].$$

In particular $\text{mult}([\pi], V^{\otimes d}) = \dim S_\pi V$ and $\text{mult}(S_\pi V, V^{\otimes d}) = \dim[\pi]$.

The projection $P_\pi : V^{\otimes d} \rightarrow S_\pi V \otimes [\pi]$ may be obtained as the direct sum of projection operators $P_{\pi, std}$ where the Young tableaux without repetitions are labeled such that the numbers increase along the rows and columns. Such tableaux are called *standard*. For example

$$P_{(2,1)} = P_{\begin{array}{|c|c|} \hline 1 & 2 \\ \hline 3 & \\ \hline \end{array}} \oplus P_{\begin{array}{|c|c|} \hline 1 & 3 \\ \hline & 1 \\ \hline \end{array}}.$$

6.1.3. Decomposition of tensor products. One is often interested in decompositions of a module under the action of a subgroup. For example

$S^d(V \otimes W)$ is an irreducible $GL(V \otimes W)$ -module, but as a $GL(V) \times GL(W)$ -module it has the decomposition, called the *Cauchy formula*,

$$(6.1.1) \quad S^d(V \otimes W) = \bigoplus_{|\pi|=d} S_\pi V \otimes S_\pi W.$$

For the quantum marginal problem, we will be particularly interested in the decomposition of $S^d(U \otimes V \otimes W)$ as a $GL(U) \times GL(V) \times GL(W)$ -module. An explicit formula for this decomposition is *not known*. Write

$$S^d(U \otimes V \otimes W) = \bigoplus_{|\pi|, |\mu|, |\nu|=d} (S_\pi U \otimes S_\mu V \otimes S_\nu W)^{\oplus k_{\pi, \mu, \nu}}.$$

The numbers $k_{\pi, \nu, \mu}$ that record the multiplicities are called *Kronecker coefficients*. They have several additional descriptions. For example, we could try to obtain the decomposition of $S^d(U \otimes V \otimes W)$ first using the Cauchy formula to write $S^d(U \otimes V \otimes W) = \bigoplus_{|\pi|=d} S_\pi U \otimes S_\pi(V \otimes W)$ and then further decomposing $S_\pi(V \otimes W)$. Comparing the formulas, we see

$$S_\pi(V \otimes W) = \bigoplus_{|\mu|, |\nu|=d} (S_\mu V \otimes S_\nu W)^{\oplus k_{\pi, \mu, \nu}}.$$

For yet another perspective, Schur-Weyl duality allows us to define the $GL(V)$ -module $S_\pi V$ as $S_\pi V := \text{Hom}_{\mathfrak{S}_d}([\pi], V^{\otimes d})$. From this perspective

$$\begin{aligned} S_\pi(V \otimes W) &= \text{Hom}_{\mathfrak{S}_d}([\pi], (V \otimes W)^{\otimes d}) \\ &= \text{Hom}_{\mathfrak{S}_d}([\pi], V^{\otimes d} \otimes W^{\otimes d}) \\ &= \text{Hom}_{\mathfrak{S}_d}([\pi], \left(\bigoplus_{|\mu|=d} S_\mu V \otimes [\mu] \right) \otimes \left(\bigoplus_{|\nu|=d} S_\nu W \otimes [\nu] \right)) \\ &= \bigoplus_{|\mu|, |\nu|=d} \text{Hom}_{\mathfrak{S}_d}([\pi], [\mu] \otimes [\nu]) \otimes S_\mu V \otimes S_\nu W. \end{aligned}$$

We conclude $k_{\pi, \mu, \nu} = \dim \text{Hom}_{\mathfrak{S}_d}([\pi], [\mu] \otimes [\nu])$. To recover the symmetry from the permutation group perspective, we use that fact that representations of the permutation group are self-dual: $[\pi]^* \simeq [\pi]$, so

$$\begin{aligned} \text{Hom}_{\mathfrak{S}_d}([\pi], [\mu] \otimes [\nu]) &= ([\pi]^* \otimes [\mu] \otimes [\nu])^{\mathfrak{S}_d} \\ &= ([\pi] \otimes [\mu] \otimes [\nu])^{\mathfrak{S}_d}, \end{aligned}$$

i.e., $k_{\pi, \mu, \nu} = \dim([\pi] \otimes [\mu] \otimes [\nu])^{\mathfrak{S}_d}$.

In other words $k_{\pi, \mu, \nu} = \text{mult}([d], [\pi] \otimes [\mu] \otimes [\nu]) = \text{mult}([\pi], [\mu] \otimes [\nu])$.

6.2. Projections onto isotypic subspaces of $\mathcal{H}^{\otimes d}$

Above we discussed representations of the general linear group $GL(V)$ where V is a complex vector space. In quantum theory, we are interested in representations on the unitary group $\mathbf{U}(\mathcal{H})$ on a Hilbert space \mathcal{H} . A subtlety we ignored before is that the unitary group is a real Lie group, not a complex Lie group, because complex conjugation is not a complex linear map. It is a special case of a general fact about representations of a maximal compact subgroups of complex Lie groups have the same representation theory as the the original group, so in particular the decomposition of $\mathcal{H}^{\otimes d}$ as a $\mathbf{U}(\mathcal{H})$ -module coincides with its decomposition as a $GL(\mathcal{H})$ -module.

For a partition $\pi = (p_1, \dots, p_d)$ of d , introduce the notation $\bar{\pi} = (\frac{p_1}{d}, \dots, \frac{p_d}{d})$ which is a probability distribution on $\{1, \dots, d\}$. Recall the relative entropy $H(\bar{p}||\bar{q}) = -\sum_i p_i \log \frac{q_i}{p_i}$, which may be thought of as measuring how close p, q are because it is non-negative, and zero if and only if $p = q$.

6.2.1. The quantum marginal problem: statement of results. The results in the following three theorems appeared almost at the same time:

Theorem 6.2.1. [CM06] *Let ρ_{AB} be a density operator on $\mathcal{H}_A \otimes \mathcal{H}_B$. Then there exists a sequence (π_j, μ_j, ν_j) of triples of partitions such that $k_{\pi_j, \mu_j, \nu_j} \neq 0$ for all j and*

$$\begin{aligned} \lim_{j \rightarrow \infty} \bar{\pi}_j &= \text{spec}(\rho_{AB}) \\ \lim_{j \rightarrow \infty} \bar{\mu}_j &= \text{spec}(\rho_A) \\ \lim_{j \rightarrow \infty} \bar{\nu}_j &= \text{spec}(\rho_B). \end{aligned}$$

Theorem 6.2.2. [Kly04] *Let ρ_{AB} be a density operator on $\mathcal{H}_A \otimes \mathcal{H}_B$ such that $\text{spec}(\rho_{AB})$, $\text{spec}(\rho_A)$ and $\text{spec}(\rho_B)$ are all rational vectors. Then there exists an $M > 0$ such that $k_{M \text{spec}(\rho_A), M \text{spec}(\rho_B), M \text{spec}(\rho_C)} \neq 0$.*

Theorem 6.2.3. [Kly04] *Let π, μ, ν be partitions of d with $k_{\pi, \mu, \nu} \neq 0$ and satisfying $\ell(\pi) \leq mn$, $\ell(\mu) \leq m$, and $\ell(\nu) \leq n$. Then there exists a density operator ρ_{AB} on $\mathbb{C}^n \otimes \mathbb{C}^m = \mathcal{H}_A \otimes \mathcal{H}_B$ with $\text{spec}(\rho_{AB}) = \bar{\pi}$, $\text{spec}(\rho_A) = \bar{\mu}$, and $\text{spec}(\rho_B) = \bar{\nu}$.*

Klyatchko's proofs are via co-adjoint orbits and vector bundles on flag varieties, while the proof of Christandl-Mitchison is information-theoretic in flavor. In the spirit of this course, we will give the information-theoretic proof, as well as information-theoretic proofs of Klyatchko's results following [CHM07].

6.2.2. The Keyl-Werner theorem. Informally, the following theorem states that for a density operator ρ on \mathcal{H} , the projection of $\rho^{\otimes d}$ onto the

$GL(\mathcal{H}) \times \mathfrak{S}_d$ submodules $S_\pi \mathcal{H} \otimes [\pi]$ of $\mathcal{H}^{\otimes d}$ for large d is negligible unless $\bar{\pi}$ is close to $\text{spec}(\rho)$.

Theorem 6.2.4. *Let $\rho \in \text{End}(\mathcal{H})$ be a density operator, where $\dim \mathcal{H} = n$. Let $|\pi| = d$ and let $P_\pi : \mathcal{H}^{\otimes d} \rightarrow S_\pi \mathcal{H} \otimes [\pi]$ be the projection operator. Then*

$$\text{trace}(P_\pi \rho^{\otimes d}) \leq (d+1) \binom{n}{2} e^{-dH(\bar{\pi} \parallel \text{spec}(\rho))}.$$

Before giving the proof we will need a few more notions from representation theory. Let e_1, \dots, e_n be a basis of \mathcal{H} and write elements of the induced basis of $\mathcal{H}^{\otimes d}$ as $e_I = e_{i_1} \otimes \dots \otimes e_{i_d}$. Define the *weight* of e_I , $wt(e_I) := (w_1, \dots, w_n)$ where w_j is the number of i_t 's equal to j .

For a partition $\pi = (p_1, \dots, p_d)$, let $\ell(\pi)$ denote the number of nonzero p_j 's, the *length* of π .

Consider the projection P_π . If e_I contains less than $\ell(\pi)$ distinct indices, then $P_\pi(e_I) = 0$, because the projections skew-symmetrize over $\ell(\pi)$ slots. More generally, consider the Young diagram of π : its first column has ℓ boxes. Say its second column has q_2 boxes, then in addition to I containing ℓ distinct indices, taking away those ℓ indices, among the remaining indices, there must be q_2 distinct, and taking these away as well, there must be q_3 (the height of the third column) distinct indices remaining, etc...

For a partition π , let π' denote the partition whose Young diagram is the transpose of the Young diagram of π .

Exercise 6.2.5: Show that for partitions μ, ν of d , that $\mu \prec \nu$ if and only if $\nu' \prec \mu'$.

Exercise 6.2.6: Show that the projection of e_I to $S_\pi V \otimes [\pi]$ is nonzero if and only if $wt(e_I) \prec \pi$. Hint for the sufficiency: we are projecting to the entire isotypic component, and one can always construct a Young tableau tailor made to the ordering of I .

So for example, when $\pi = (d)$ no basis vector maps to zero and when $\pi = (1, \dots, 1) = (1^d)$ only basis vectors with weights having at least d positive w_i have nonzero images.

Proof of the Keyl-Werner Theorem. Choose a basis (e_1, \dots, e_n) of \mathcal{H} consisting of eigenvectors of ρ , so the eigenvectors of $\rho^{\otimes d}$ are the e_I . Write $\bar{r} = \text{spec}(\rho) = (r_1, \dots, r_n)$ with $r_j \geq r_{j+1}$, so $\rho^{\otimes d} = \sum_I r_I |e_I\rangle \langle e_I|$, where $r_I = r_{i_1} \dots r_{i_d}$. The eigenvalues that do not project to zero are $\{r_I \mid wt(e_I) \prec \pi\}$. By the definition of the dominance order, these satisfy

$$r_I \leq r_1^{p_1} \dots r_d^{p_d}.$$

We conclude

$$\text{trace}(P_\pi \rho^{\otimes d}) \leq \dim(S_\pi \mathcal{H} \otimes [\pi]) r_1^{p_1} \dots r_d^{p_d}.$$

To finish the proof we need an estimate of $\dim(S_\pi \mathcal{H} \otimes [\pi])$. The dimension of this space may be computed as follows (see, e.g., **** for proofs): identify π with its Young diagram, and write $x \in \pi$ for a box in the diagram. Define the *hook length* of x to be the number of boxes to the right of it in its row, plus the number of boxes below it in its column, plus one, and define the content $c(x)$ of x to be zero if x is on the main diagonal, j , if it is on the j -th diagonal above the main diagonal, and $-j$ if it is on the j -th diagonal below the main diagonal. For example we have the following hook lengths for $(4, 2, 1, 1)$:

7	4	2	1
4	1		
2			
1			

Then

(6.2.1)

$$\dim S_\pi \mathbb{C}^n = \prod_{x \in \pi} \frac{n + c(x)}{h(x)} = \prod_{1 \leq i < j \leq n} \frac{\pi_i - \pi_j + j - 1}{j - i}, \text{ and}$$

(6.2.2)

$$\dim[\pi] = \frac{d!}{\prod_{x \in \pi} h(x)} = \frac{d! \prod_{1 \leq s < t \leq d} \ell_s - \ell_t}{\prod_{u=1}^d \ell_u!}, \text{ where } \ell_s := p_s + d - s.$$

Exercise 6.2.7: Show that $\dim S_\pi \mathbb{C}^n \leq (d+1) \binom{n}{2}$ and $\frac{d!}{\prod_{j=1}^d (p_j + d - j)!} \leq \dim[\pi] \leq \frac{d!}{\prod_{i=1}^d p_i!}$.

We conclude

$$\text{trace}(P_\pi \rho^{\otimes d}) \leq (d+1) \binom{n}{2} \frac{d!}{\prod p_i!} r_1^{p_1} \cdots r_d^{p_d}.$$

Finally

$$\begin{aligned} e^{-dH(\bar{\pi} \parallel \text{spec}(\rho))} &= \exp\left(-d\left(-\sum (p_i \log(r_i) - p_i \log(p_i))\right)\right) \\ &= e^d \frac{\prod r_i^{p_i}}{\prod p_i^{p_i}} \\ &> \frac{d!}{\prod p_i!} r_1^{p_1} \cdots r_d^{p_d} \end{aligned}$$

where for the last line, recall that $m! > (\frac{m}{e})^m$ and $d = \sum p_i$. □

Let $SPEC_n$ denote the set of possible spectra for density operators on \mathbb{C}^n .

Corollary 6.2.8. *Let ρ be a density operator and let $S \subset \text{SPEC}_n$. Set*

$$P_S := \sum_{|\pi|=d, \bar{\pi} \in S} P_\pi.$$

Then

$$\text{trace}(P_X \rho^{\otimes d}) \leq (d+1) \binom{n}{2}^{+n} \exp(-d \min_{\bar{\pi} \in S} H(\bar{\pi} \| \text{spec}(\rho))).$$

Proof. The number of Young diagrams with d boxes in n rows is at most $(d+1)^n$, and we are taking the worst case in the exponential. \square

For $r \in \text{SPEC}_n$, let

$$\begin{aligned} B_\epsilon(r) &:= \{r' \in \text{SPEC}_n \mid \|r - r'\|_1 < \epsilon\} \\ S_{\epsilon,r} &:= \text{SPEC}_n \setminus B_\epsilon(r). \end{aligned}$$

Corollary 6.2.9. *For all $\epsilon, \delta > 0$, there exists a d_0 such that for all $d > d_0$, $\text{trace}(P_{S_{\epsilon, \text{spec}(\rho)}} \rho^{\otimes d}) < \delta$, i.e., $\text{trace}(P_{B_\epsilon(\text{spec}(\rho))} \rho^{\otimes d}) \geq 1 - \delta$.*

6.2.3. Proof of Theorem 6.2.1.

Proof of Theorem 6.2.1. It will be more convenient to view the theorem symmetrically by taking a purification $|\psi\rangle \in \mathcal{H}_A \otimes \mathcal{H}_B \otimes \mathcal{H}_C$, so $\rho_C = \rho_{AB} = \text{trace}_{\mathcal{H}_A \otimes \mathcal{H}_B}(|\psi\rangle\langle\psi|)$. By Corollary 6.2.9, for all $\epsilon, \delta > 0$, there exists d_0 such that for all $d \geq d_0$

$$\begin{aligned} \text{trace}(P_{B_\epsilon(\text{spec}(\rho_A))} \rho_A^{\otimes d}) &\geq 1 - \delta \\ \text{trace}(P_{B_\epsilon(\text{spec}(\rho_B))} \rho_B^{\otimes d}) &\geq 1 - \delta \\ \text{trace}(P_{B_\epsilon(\text{spec}(\rho_C))} \rho_C^{\otimes d}) &\geq 1 - \delta. \end{aligned}$$

Now, for all projection operators $P \in \text{End}(\mathcal{H}_A)$, $Q \in \text{End}(\mathcal{H}_B)$, $R \in \text{End}(\mathcal{H}_C)$, and density operators ρ_{ABC} on $\mathcal{H}_A \otimes \mathcal{H}_B \otimes \mathcal{H}_C$, we have

(6.2.3)

$$\text{trace}((P \otimes Q \otimes R) \rho_{ABC}) \geq \text{trace}(P \rho_A) + \text{trace}(Q \rho_B) + \text{trace}(R \rho_C) - 2.$$

Exercise 6.2.10: Verify (6.2.3). Hint: First show that for all P, Q ,

$$\text{trace}((P \otimes Q) \rho_{AB}) \geq \text{trace}(P \rho_A) + \text{trace}(Q \rho_B) - 1$$

by considering $\text{trace}((\text{Id}_{\mathcal{H}_A} - P) \otimes (\text{Id}_{\mathcal{H}_B} - Q) \rho_{AB}) \geq 0$.

We obtain

$$\text{trace} \left(P_{B_\epsilon(\text{spec}(\rho_A))} \otimes P_{B_\epsilon(\text{spec}(\rho_B))} \otimes P_{B_\epsilon(\text{spec}(\rho_C))} (|\psi\rangle\langle\psi|)^{\otimes d} \right) \geq 1 - 3\delta.$$

Assuming $\delta < \frac{1}{3}$, since each of $P_{B_\epsilon(\text{spec}(\rho_A))}$, $P_{B_\epsilon(\text{spec}(\rho_B))}$, $P_{B_\epsilon(\text{spec}(\rho_C))}$ is a sum of projection operators, there must be one triple (μ, ν, π) , with $\bar{\mu} \in B_\epsilon(\text{spec}(\rho_A))$, $\bar{\nu} \in B_\epsilon(\text{spec}(\rho_B))$, $\bar{\pi} \in B_\epsilon(\text{spec}(\rho_C))$, with $k_{\bar{\pi}, \bar{\mu}, \bar{\nu}} \neq 0$ for which the projection is nonzero. Now just take a sequence of such as $\epsilon \rightarrow 0$. \square

6.2.4. Consequences. We can recover standard facts about von Neumann entropy from Theorem 6.2.1.

Corollary 6.2.11 (subadditivity of von Neumann entropy). $H(\rho_{AB}) \leq H(\rho_A) + H(\rho_B)$.

Proof. Since $[\mu] \otimes [\nu] = \bigoplus_{\pi} [\pi]^{\oplus k_{\pi\mu\nu}}$, if $k_{\pi\mu\nu} \neq 0$, then $\dim([\mu] \otimes [\nu]) \geq \dim[\pi]$. Now take a sequence (π_j, μ_j, ν_j) as in the proof of Theorem 6.2.1, so in particular, for each j , $\dim[\pi_j] \leq \dim[\mu_j] \dim[\nu_j]$. Write $\text{spec}(\rho_{AB}) = (p_1^{AB}, \dots, p_{mn}^{AB})$, by Exercise 6.2.7,

$$\frac{d!}{\prod_x (p_x + d - x)!} \leq \dim[\pi] \leq \frac{d!}{\prod_x p_x!}$$

so $\frac{1}{d_j} \log(\dim[\pi_j]) \sim -\sum \log(p_j!)$ tends to $-\sum_i p_i^{AB} \log p_i^{AB} = H(\rho_{AB})$ and analogously for $H(\rho_A)$, $H(\rho_B)$, so we obtain the result. \square

Corollary 6.2.12 (triangle inequality for von Neumann entropy). $H(\rho_{AB}) \geq |H(\rho_A) - H(\rho_B)|$.

This is an immediate consequence of the symmetry of the Kronecker coefficients and Corollary 6.2.11.

6.2.5. Weights, Cartan products and a result on Kronecker coefficients. Before giving the proof of Theorem 6.2.2 we need a little more representation theory. We say $v \in (\mathbb{C}^n)^{\otimes d}$ is a *weight vector* if $v = \sum_s c_{I_s} e_{I_s}$ where $\text{wt}(e_{I_s}) = \text{wt}(e_{I_t})$ for all s, t in the sum.

Note that for all I, π , the projection $P_{\pi}(e_I)$ is either a weight vector of weight $\text{wt}(e_I)$ or zero.

A weight vector v is a *highest weight vector* if $g \cdot v = v$ for all $g \in N_n$, where $N_n \subset GL_n$ is the subgroup of matrices with 1's on the diagonal and zero's below the diagonal.

We extend the notion of weight vectors and highest weight vectors to $G = GL_m \times GL_n \times GL_k$ in the natural way, e.g., the weight of $e_I \otimes f_J \otimes h_K$ is a triple of weights and a weight vector is a highest weight vector if it is invariant under $N_m \times N_n \times N_k$.

Fact: If V is an irreducible G -module, where G is reductive group (a class of groups including GL_n and products of general linear groups) then there exists a unique up to scale highest vector in V .

Exercise 6.2.13: If $v \in S_{\pi} \mathbb{C}^n$ is a highest weight vector, then $\text{wt}(v) = \pi$.

Exercise 6.2.14: If $v \in V^{\otimes d}$ and $w \in V^{\otimes \delta}$ are highest weight vectors, then $v \otimes w \in V^{\otimes d+\delta}$ is a highest weight vector.

The following theorem was apparently “known to the experts” in representation theory but unknown in the quantum information theory community until 2004.

Theorem 6.2.15. *If $k_{\pi,\mu,\nu} \neq 0$ and $k_{\pi',\mu',\nu'} \neq 0$, then $k_{\pi+\pi',\mu+\mu',\nu+\nu'} \neq 0$. In particular, if $k_{\pi,\mu,\nu} \neq 0$, then $k_{M\pi,M\mu,M\nu} \neq 0$ for all $M \in \mathbb{N}$.*

Proof. Let $X_{\mu,\mu,\nu} \in S^d(\mathbb{C}^k \otimes \mathbb{C}^m \otimes \mathbb{C}^n)$ be a highest weight vector for $GL_k \times GL_m \times GL_n$ of weight (π, μ, ν) , and let $X'_{\mu',\mu',\nu'} \in S^{d'}(\mathbb{C}^k \otimes \mathbb{C}^m \otimes \mathbb{C}^n)$ be a highest weight vector for $GL_k \times GL_m \times GL_n$ of weight (π', μ', ν') , both of which exist by the non-vanishing of the Kronecker coefficients in the hypothesis. Letting $W = \mathbb{C}^k \otimes \mathbb{C}^m \otimes \mathbb{C}^n$, $X_{\mu,\mu,\nu} \otimes X'_{\mu',\mu',\nu'} \in S^d W \otimes S^{d'} W$ is a $G = GL_k \times GL_m \times GL_n$ highest weight vector. We would like to construct such a vector in $S^{d+d'} W$.

Consider the following diagram, for $g \in N_k \times N_m \times N_n$.

$$\begin{array}{ccc} S^d W \otimes S^{d'} W & \xrightarrow{g} & S^d W \otimes S^{d'} W \\ \text{mult} \downarrow & & \text{mult} \downarrow \\ S^{d+d'} W & \xrightarrow{g} & S^{d+d'} W \end{array}$$

This diagram is commutative as the vertical arrows come from an action of the permutation group and the horizontal from the action of the general linear group, and these actions commutes. Applying g first then symmetrizing is the same as just symmetrizing, so the same must be true in the other order, thus the image of $X_{\mu,\mu,\nu} \otimes X'_{\mu',\mu',\nu'}$ under the multiplication map is a highest weight vector or zero, but since the multiplication of two polynomials is nonzero, we conclude. □

Let

$$Kron_{m,n,k} := \{(\mu, \nu, \pi) \mid k_{\mu,\nu,\pi} \neq 0 \text{ and } \ell(\mu) \leq m, \ell(\nu) \leq n, \ell(\pi) \leq k\}.$$

Theorem 6.2.15 implies that $Kron_{m,n,k}$ is a semi-group.

Fact: $Kron_{m,n,k}$ is finitely generated.

The finite generation is a consequence of Hilbert’s famous basis theorem, which, in its simplest form says that ideals in the polynomial ring $\mathbb{C}[x_1, \dots, x_N]$ are finitely generated. Similarly, if G is a reductive algebraic group and \mathcal{A} is an algebra equipped with a G -action, if \mathcal{A} is finitely generated, then so is the subalgebra of G -invariants \mathcal{A}^G . The algebra of highest weight vectors for GL_n -modules in the tensor algebra of \mathbb{C}^n is finitely generated by $e_1 \wedge \dots \wedge e_u$, $1 \leq u \leq n$, the highest weight vectors for $\Lambda^u \mathbb{C}^n = S_{1^u} \mathbb{C}^n$, and similarly for $GL_m \times GL_n \times GL_m$ highest weight vectors in the tensor

algebra of $\mathbb{C}^m \otimes \mathbb{C}^n \otimes \mathbb{C}^k$. (Note that the group for which this is the algebra of invariants is $N_m \times N_n \times N_k$ which is not reductive.) We must deal with the highest weight vectors in the symmetric algebra, and the same group. An extension of Hilbert's theorem to Grosshans subgroups (see, e.g., [Dol03, Chap. 4]) gives the result.

Let $KRON_{m,n,k} := \{(\bar{\mu}, \bar{\nu}, \bar{\pi}) \mid (\mu, \nu, \pi) \in Kron_{m,n,k}\}$, the normalized Kronecker coefficients.

6.2.6. Proofs of Klyatchko's theorems.

Proof of Theorem 6.2.2. Let $(r_A, r_B, r_C) \in SPEC_{m,n,mn}$. Theorem 6.2.1 implies there exists a sequence (μ_j, ν_j, π_j) with $(\bar{\mu}_j, \bar{\nu}_j, \bar{\pi}_j)$ converging to (r_A, r_B, r_C) .

Let $(\bar{\mu}_\alpha, \bar{\nu}_\alpha, \bar{\pi}_\alpha)$ be a finite set of generators for $KRON_{m,n,mn}$. By hypothesis,

$$(6.2.4) \quad (r_A, r_B, r_C) = \sum_{\alpha} p_{\alpha} (\bar{\mu}_{\alpha}, \bar{\nu}_{\alpha}, \bar{\pi}_{\alpha})$$

for some $p_{\alpha} \in \mathbb{Q}$ with $p_{\alpha} \geq 0$ and $\sum_{\alpha} p_{\alpha} = 1$.

Say $KRON_{m,n,mn}$ has $t+1$ vertices, in fact we can get away with just $t+1$ elements on the vertices **so why start with original expression?**. Now, considering the p_{α} as unknowns, (6.2.4) is a set of $m+n+mn$ equations, which we can select a subset of t independent equations involving only vertex elements, and add the constraint that the coefficients sum to one. Since (r_A, r_B, r_C) is rational, the new solution coefficients, will be as well. Thus we can write

$$(r_A, r_B, r_C) = \sum \frac{n_{\phi}}{N} (\bar{\mu}_{\phi}, \bar{\nu}_{\phi}, \bar{\pi}_{\phi})$$

for some non-negative integers n_{ϕ}, N . Let d be the least common multiple of the $|\mu_{\phi}|$'s, so taking $M = dN$, we have that $k_{Mr_A, Mr_B, Mr_C} \neq 0$ as desired. \square

Exercise 6.2.16: Use Theorem 6.2.2 to prove Theorem 6.2.1.

To prove Theorem 6.2.3 we will need the following Lemma:

Lemma 6.2.17. *Let $(\mu, \nu, \pi) \in Kron_{m,n,mn}$ with $|\mu| = d$. Then there exists $|\psi\rangle \in \mathcal{H}_A \otimes \mathcal{H}_B \otimes \mathcal{H}_C$ such that the marginals of the corresponding density*

operator $\rho_{ABC} = |\psi\rangle\langle\psi|$ satisfy

$$\begin{aligned} \|\text{spec}(\rho_A) - \bar{\mu}\|_1 &\leq 3mn\sqrt{\frac{\log(d)}{d}} \\ \|\text{spec}(\rho_A) - \bar{\nu}\|_1 &\leq 3mn\sqrt{\frac{\log(d)}{d}} \\ \|\text{spec}(\rho_A) - \bar{\pi}\|_1 &\leq 3mn\sqrt{\frac{\log(d)}{d}} \end{aligned}$$

Theorem 6.2.3 follows, as $(j\mu, j\nu, j\pi) \in \text{Kron}_{m,n,mn}$, so we may obtain a sequence of ψ 's whose associated density operator has marginals converging to $(\bar{\mu}, \bar{\nu}, \bar{\pi})$.

To prove Lemma 6.2.17 we will need the following lemma from classical probability:

Lemma 6.2.18 (Pinsker's inequality). *Let p, q be probability distributions on $[d]$. then*

$$H(\bar{p}|\bar{q}) \geq \frac{1}{2\ln 2} \|p - q\|_1^2.$$

Proof. First consider the special case $d = 2$, so $\bar{p} = (p, 1-p)$, $\bar{q} = (q, 1-q)$. Say $p \geq q$.

Consider

$$f(p, q) := H(\bar{p}|\bar{q}) - \frac{1}{2\ln 2} \|p - q\|_1^2 = p \log \frac{p}{q} + (1-p) \log \frac{1-p}{1-q} - \frac{1}{2\ln 2} (2(p-q))^2.$$

Note that $f(p, q) = 0$ for $q = p$ and we need to show $f(p, q) \geq 0$ for $q < p$.

Exercise 6.2.19: Show that $\frac{\partial f}{\partial q} \leq 0$ to prove the special case.

Now for the general case, set $A := \{j \in [d] \mid p_j \geq q_j\}$, write $p_A = (\sum_{j \notin A} p_j, \sum_{j \in A} p_j)$ and $q_A = (\sum_{j \notin A} q_j, \sum_{j \in A} q_j)$ Now

$$\begin{aligned} \|\bar{p} - \bar{q}\|_1 &= \sum_j |p_j - q_j| \\ &= \sum_{j \in A} p_j - q_j + \sum_{j \notin A} q_j - p_j \\ &= \left| \sum_{j \in A} p_j - q_j \right| + \left| (1 - \sum_{j \notin A} p_j) - (1 - \sum_{j \notin A} q_j) \right| \\ &= \|p_A - q_A\|_1. \end{aligned}$$

Let Z be the random variable taking 1 on A and 0 on $[d] \setminus A$. Then

$$H(\bar{p}|\bar{q}) = H(p(Z)|q(Z)) + H(\bar{p}|\bar{q}|Z).$$

The first term is $H(p_A||q_A)$ and the second is non-negative. Putting it all together

$$H(\bar{p}||\bar{q}) \geq H(p_A||q_A) \geq \frac{1}{2 \ln 2} \|p_A - q_A\|_1^2 = \frac{1}{2 \ln 2} \|p - q\|_1^2.$$

□

Proof of Lemma 6.2.17. Let $|v\rangle \in S_\mu \mathbb{C}^m \otimes S_\nu \mathbb{C}^n \otimes S_\pi \mathbb{C}^{mn} \subset S^d(\mathbb{C}^m \otimes \mathbb{C}^n \otimes \mathbb{C}^{mn})$ have length one. We want to find $|\psi\rangle \in \mathbb{C}^m \otimes \mathbb{C}^n \otimes \mathbb{C}^{mn} =: W$ such that $|\psi\rangle^{\otimes d}$ is close to $|v\rangle$. Since the Veronese variety of d -th powers $\hat{v}_d(\mathbb{P}W)$ spans $\mathbb{P}S^d W$, at worst $|v\rangle$ is a sum of $\dim W$ unit vectors in $\hat{v}_d(\mathbb{P}W)$, all equidistant from $|v\rangle$, so we certainly may attain

$$|\langle \psi^{\otimes d} | v \rangle| \geq \frac{1}{\dim S^d W} \gg \frac{1}{(mn)^{2d}}.$$

Now v is just one vector in $S_\mu \mathbb{C}^m \otimes S_\nu \mathbb{C}^n \otimes S_\pi \mathbb{C}^{mn}$ so

$$\text{trace}(P_\mu \otimes P_\nu \otimes P_\pi (|\psi\rangle\langle\psi|)^{\otimes d}) \geq \text{trace}(|v\rangle\langle v| (|\psi\rangle\langle\psi|)^{\otimes d}) = |\langle \psi^{\otimes d} | v \rangle|^2 > \frac{1}{(mn)^d}.$$

In particular

$$\text{trace}(P_\mu \rho_A^{\otimes d}) = \text{trace}(P_\mu \otimes \text{Id}_{\mathcal{H}_B} \otimes \text{Id}_{\mathcal{H}_C} (|\psi\rangle\langle\psi|)^{\otimes d}) \geq \frac{1}{(mn)^d}.$$

On the other hand, the Keyl-Werner theorem 6.2.4 says

$$\text{trace}(P_\mu \rho_A^{\otimes d}) \leq (d+1) \binom{m}{2} e^{-dH(\bar{\mu}||\text{spec}(\rho_A))},$$

so

$$H(\bar{\mu}||\text{spec}(\rho_A)) \leq \frac{\binom{m}{2} \log(d+1) m^2 n^2 \log(d)}{d}.$$

Exercise 6.2.20: Finish the proof by using Pinsker's inequality.

□

Combining the theorems we also conclude

$$\text{Spec}_{m,n,mn} = \overline{KRON}_{m,n,mn}.$$

In particular, $\text{Spec}_{m,n,mn}$ is a convex polytope.

Remark 6.2.21. We can use quantum theory to deduce representation-theoretic consequences: $k_{\mu,\nu,\pi} \neq 0$ implies $H(\bar{\pi}) \leq H(\bar{\mu}) + H(\bar{\nu})$, $H(\bar{\mu}) \leq H(\bar{\pi}) + H(\bar{\nu})$, and $H(\bar{\nu}) \leq H(\bar{\mu}) + H(\bar{\pi})$.

Hints and Answers to Selected Exercises

Chapter 1.

1.2.3 see Figure 1.2.1.

1.2.4 Write

$$\begin{pmatrix} DFT_M & \Delta_M DFT_M \\ DFT_M & -\Delta_M DFT_M \end{pmatrix} = \begin{pmatrix} \text{Id}_M & \Delta_M \\ -\text{Id}_M & -\Delta_M \end{pmatrix} \begin{pmatrix} DFT_M & 0 \\ 0 & DFT_M \end{pmatrix}.$$

Also note that a $k \times k$ permutation matrix has k nonzero entries, and the product of two permutation matrices is a permutation matrix.

?? Choose the first four rows and last four columns. One obtains a 4×4 matrix M' and the associated tensor T' , so $\mathbf{R}(T_{aft,3}) \geq 8 + \mathbf{R}(T')$. Iterating the method twice yields $\mathbf{R}(T_{aft,3}) \geq 8 + 4 + 2 + 1 = 15$.

Chapter 2.

2.1.5 Consider $\langle v_i | v_j \rangle = \langle Av_i | Av_j \rangle$.

2.2.9 It is sufficient to work in bases, i.e., with matrices. First prove the case X is diagonal, then the case X is diagonalizable, then either write $X = X_s + X_n$ as the sum of a diagonalizable matrix and a nilpotent matrix or argue that the diagonalizable matrices is a dense open subset in the space of all matrices.

2.2.11 First consider the case X is diagonal, and use that the eigenvalues of a Hermitian matrix are real.

2.5.3 $\frac{7}{8}$.

Chapter 3.

3.1.9 Consider the binomial coefficients in the expansion of $(x + a)^N$.

3.1.5 First do the case $a = 2^\ell$. Then show the general case by using the binary expansion of a .

3.5.5 Write $x = \lfloor x \rfloor + \{x\}$. If $\{x\} < \frac{1}{2}$, then $2x = 2\lfloor x \rfloor + \{2x\}$ and therefore $\lfloor 2x \rfloor = 2\lfloor x \rfloor$. If $\{x\} \geq \frac{1}{2}$, then $\lfloor 2x \rfloor = \lfloor 2\lfloor x \rfloor \rfloor + 1$.

3.5.6 Use (1.3.9).

Chapter 4.

4.4.6 Prove an upper and a lower bound for the quantity.

5.1.1

$$\begin{aligned}\Pr(\text{span}\{\mathcal{M}_1, \mathcal{M}_2\}) &= 1 \\ \Pr(\mathcal{M}_1) &= |\alpha|^2 \\ \Pr(\mathcal{M}_2) &= \frac{1}{2}|\alpha + \beta|^2 \\ \Pr(\mathcal{M}_1 \cap \mathcal{M}_2) &= 0\end{aligned}$$

5.1.5 Use that any $X \in \text{End}(\mathcal{H})$ may be uniquely written as a sum of a Hermitian and an anti-Hermitian (i.e., i times a Hermitian) operator.

5.1.13 Use the Cauchy-Schwartz inequality, in the form $|\text{trace}(AB\rho)|^2 \leq \text{trace}((A\rho)^2) \text{trace}((B\rho)^2)$.

5.2.2 Use the polar or singular value decomposition.

Chapter 5.

5.3.2 Note that $\text{transpose} \otimes \text{Id}(\sum_{i,j} |ii\rangle\langle jj|) = \sum_{i,j} |ji\rangle\langle ij|$.

5.3.13 Use Theorem 5.3.11 and the Cauchy-Schwartz inequality.

5.4.6 Write $\rho_{AB} = \sum_{ij} \lambda_{i,j} |v_i\rangle\langle v_i| \otimes |w_j\rangle\langle w_j|$, the eigenbasis decomposition.

Bibliography

- [Aar13] Scott Aaronson, *Quantum computing since Democritus*, Cambridge University Press, Cambridge, 2013. MR 3058839
- [AB09] Sanjeev Arora and Boaz Barak, *Computational complexity*, Cambridge University Press, Cambridge, 2009, A modern approach. MR 2500087 (2010i:68001)
- [Ad89] V. I. Arnol' d, *Mathematical methods of classical mechanics*, Graduate Texts in Mathematics, vol. 60, Springer-Verlag, New York, [1989?], Translated from the 1974 Russian original by K. Vogtmann and A. Weinstein, Corrected reprint of the second (1989) edition. MR 1345386
- [Adl78] Leonard Adleman, *Two theorems on random polynomial time*, 19th Annual Symposium on Foundations of Computer Science (Ann Arbor, Mich., 1978), IEEE, Long Beach, Calif., 1978, pp. 75–83. MR 539832
- [Ahl78] Lars V. Ahlfors, *Complex analysis*, third ed., McGraw-Hill Book Co., New York, 1978, An introduction to the theory of analytic functions of one complex variable, International Series in Pure and Applied Mathematics. MR 510197
- [AKS04] Manindra Agrawal, Neeraj Kayal, and Nitin Saxena, *PRIMES is in P*, Ann. of Math. (2) **160** (2004), no. 2, 781–793. MR 2123939
- [BCHW16] F. G. S. L. Brandao, M. Christandl, A. W. Harrow, and M. Walter, *The Mathematics of Entanglement*, ArXiv e-prints (2016).
- [Bel64] J.S. Bell, *On the einstein-podolsky-rosen paradox*, Physics **1** (1964), 195–200.
- [Bir46] Garrett Birkhoff, *Three observations on linear algebra*, Univ. Nac. Tucumán. Revista A. **5** (1946), 147–151. MR 0020547
- [BW92] Charles H. Bennett and Stephen J. Wiesner, *Communication via one- and two-particle operators on einstein-podolsky-rosen states*, Phys. Rev. Lett. **69** (1992), 2881–2884.
- [CHM07] Matthias Christandl, Aram W. Harrow, and Graeme Mitchison, *Nonzero Kronecker coefficients and what they tell us about spectra*, Comm. Math. Phys. **270** (2007), no. 3, 575–585. MR MR2276458 (2007k:20029)

- [CHSH69] John F. Clauser, Michael A. Horne, Abner Shimony, and Richard A. Holt, *Proposed experiment to test local hidden-variable theories*, Phys. Rev. Lett. **23** (1969), 880–884.
- [Chu36] Alonzo Church, *An Unsolvable Problem of Elementary Number Theory*, Amer. J. Math. **58** (1936), no. 2, 345–363. MR 1507159
- [CM06] Matthias Christandl and Graeme Mitchison, *The spectra of quantum states and the Kronecker coefficients of the symmetric group*, Comm. Math. Phys. **261** (2006), no. 3, 789–797. MR 2197548
- [CT65] James W. Cooley and John W. Tukey, *An algorithm for the machine calculation of complex Fourier series*, Math. Comp. **19** (1965), 297–301. MR 0178586
- [CW04] Matthias Christandl and Andreas Winter, *“Squashed entanglement”: an additive entanglement measure*, J. Math. Phys. **45** (2004), no. 3, 829–840. MR 2036165
- [Dol03] Igor Dolgachev, *Lectures on invariant theory*, London Mathematical Society Lecture Note Series, vol. 296, Cambridge University Press, Cambridge, 2003. MR MR2004511 (2004g:14051)
- [EPR35] A. Einstein, B. Podolsky, and N. Rosen, *Can quantum-mechanical description of physical reality be considered complete?*, Phys. Rev. **47** (1935), 777–780.
- [Erd47] P. Erdős, *Some remarks on the theory of graphs*, Bull. Amer. Math. Soc. **53** (1947), 292–294. MR 0019911
- [FH91] William Fulton and Joe Harris, *Representation theory*, Graduate Texts in Mathematics, vol. 129, Springer-Verlag, New York, 1991, A first course, Readings in Mathematics. MR 1153249 (93a:20069)
- [Gau] C. F. Gauss, *Nachlass: Theoria interpolationis methodo nova tractata, gauss, werke, band 3*.
- [GHIL16] Fulvio Gesmundo, Jonathan D. Hauenstein, Christian Ikenmeyer, and J. M. Landsberg, *Complexity of linear circuits and geometry*, Found. Comput. Math. **16** (2016), no. 3, 599–635. MR 3494506
- [Gle11] James Gleick, *The information*, Pantheon Books, 2011.
- [Har01] L. Hardy, *Quantum Theory From Five Reasonable Axioms*, eprint arXiv:quant-ph/0101012 (2001).
- [HLP52] G. H. Hardy, J. E. Littlewood, and G. Pólya, *Inequalities*, Cambridge, at the University Press, 1952, 2d ed. MR 0046395
- [Kle] O. Klein.
- [KLPSMN09] Abhinav Kumar, Satyanarayana V. Lokam, Vijay M. Patankar, and Jayalal Sarma M. N., *Using elimination theory to construct rigid matrices*, Foundations of software technology and theoretical computer science—FSTTCS 2009, LIPIcs. Leibniz Int. Proc. Inform., vol. 4, Schloss Dagstuhl. Leibniz-Zent. Inform., Wadern, 2009, pp. 299–310. MR 2870721
- [Kly04] A. Klyachko, *Quantum marginal problem and representations of the symmetric group*, preprint arXiv:quant-ph/0409113v1 (2004).
- [KSV02] A. Yu. Kitaev, A. H. Shen, and M. N. Vyalyi, *Classical and quantum computation*, Graduate Studies in Mathematics, vol. 47, American Mathematical Society, Providence, RI, 2002, Translated from the 1999 Russian original by Lester J. Senechal. MR 1907291

- [Lan61] R. Landauer, *Irreversibility and heat generation in the computing process*, IBM Journal of Research and Development **5** (1961), 183–191.
- [Lan17] J.M. Landsberg, *Geometry and complexity theory*, Cambridge studies in advanced mathematics, vol. 169, Cambridge Univ. Press, 2017.
- [Lie73] Elliott H. Lieb, *Convex trace functions and the Wigner-Yanase-Dyson conjecture*, Advances in Math. **11** (1973), 267–288. MR 0332080
- [Lyo09] Jonathan Lyons, *The house of wisdom*, Bloomsbury Press, 2009.
- [Mac95] I. G. Macdonald, *Symmetric functions and Hall polynomials*, second ed., Oxford Mathematical Monographs, The Clarendon Press Oxford University Press, New York, 1995, With contributions by A. Zelevinsky, Oxford Science Publications. MR 1354144 (96h:05207)
- [Mir58] L. Mirsky, *Proofs of two theorems on doubly-stochastic matrices*, Proc. Amer. Math. Soc. **9** (1958), 371–374. MR 0095180
- [NC00] Michael A. Nielsen and Isaac L. Chuang, *Quantum computation and quantum information*, Cambridge University Press, Cambridge, 2000. MR MR1796805 (2003j:81038)
- [Nie99] M. A. Nielsen, *Conditions for a class of entanglement transformations*, P H Y S I C A L R E V I E W L E T T E R S **83** (1999), 436–439.
- [Rab80] Michael O. Rabin, *Probabilistic algorithm for testing primality*, J. Number Theory **12** (1980), no. 1, 128–138. MR 566880
- [Rus04] M. B. Ruskai, *Lieb’s simple proof of concavity of $\text{Tr} A^p K^* B^{(1-p)} K$ and remarks on related inequalities*, eprint arXiv:quant-ph/0404126 (2004).
- [Sch95] Benjamin Schumacher, *Quantum coding*, Phys. Rev. A (3) **51** (1995), no. 4, 2738–2747. MR 1328824
- [Sch03] Rüdiger Schack, *Quantum theory from four of Hardy’s axioms*, Found. Phys. **33** (2003), no. 10, 1461–1468, Special issue dedicated to David Mermin, Part I. MR 2039620
- [Sha48] C. E. Shannon, *A mathematical theory of communication*, Bell System Tech. J. **27** (1948), 379–423, 623–656. MR 0026286
- [Sho94] Peter W. Shor, *Algorithms for quantum computation: discrete logarithms and factoring*, 35th Annual Symposium on Foundations of Computer Science (Santa Fe, NM, 1994), IEEE Comput. Soc. Press, Los Alamitos, CA, 1994, pp. 124–134. MR 1489242
- [Sho97] ———, *Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer*, SIAM J. Comput. **26** (1997), no. 5, 1484–1509. MR 1471990
- [Sti55] W. Forrest Stinespring, *Positive functions on C^* -algebras*, Proc. Amer. Math. Soc. **6** (1955), 211–216. MR 0069403
- [Str69] Volker Strassen, *Gaussian elimination is not optimal*, Numer. Math. **13** (1969), 354–356. MR 40 #2223
- [Val77] Leslie G. Valiant, *Graph-theoretic arguments in low-level complexity*, Mathematical foundations of computer science (Proc. Sixth Sympos., Tatranská Lomnica, 1977), Springer, Berlin, 1977, pp. 162–176. Lecture Notes in Comput. Sci., Vol. 53. MR 0660702 (58 #32067)
- [VSD86] A Vergis, K Steiglitz, and B Dickinson, *The complexity of analog computation*, Math. Comput. Simul. **28** (1986), no. 2, 91–113.

Index

$E[X]$, 6
 O , 6
 $\sigma(X)$: standard deviation, 7
 $\text{var}(X)$, 7
 o , 6

big O notation, 6

convolution, 2

DFT, 3
discrete Fourier transform, 3
discrete logarithm, 46

expectation, 6

iid, 6

marginal distributions, 7

probability distribution
 discrete, 6

random variable
 discrete, 6
random variables
 identically distributed, 6
 independent, 6

standard deviation, 7
strong law of large numbers, 7

tensor product, 7

variance, 7

weak law of large numbers, 7

Young diagram, 91