1 Answers: quantum computation

8. Principles of quantum computing

1. The SWAP gate was explored in Part I and can be implemented using the network

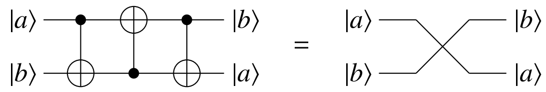

which works for both classical and quantum inputs. The classical CLONE network is just a single controlled-NOT gate

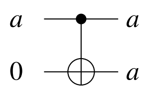

(if $a = 1$ then the second qubit is flipped from 0 to 1) but CLONE does not work on a quantum computer unless the qubits are in eigenstates: the no cloning theorem.

2. To build NOT and controlled-NOT gates from Toffoli gates just set both inputs (NOT) or one of the two inputs (controlled-NOT) to one. To build an OR gate use De Morgan's laws, a OR b = NOT(NOT a AND NOT b) and implement AND using a Toffoli gate.
3. Both NOT and AND gates can be built from Fredkin gates with appropriate patterns of inputs, though it takes a bit of thought to see why these gates work.

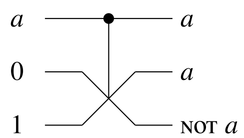

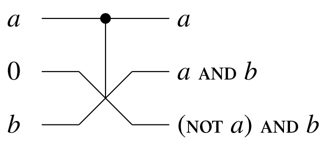

The circuit for a NOT gate also copies the input a at the same time, and so implements CLONE on a classical computer.

4. Since the Fedkin gate is a controlled-SWAP gate it can be built from the standard SWAP network by adding an additional control to each gate, turning each controlled-NOT gate into a Toffoli gate.

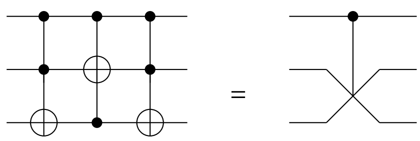

It is not possible to build a Toffoli gate using only Fredkin gates without using ancilla bits; this is most easily seen by noting that the Fredkin gate only swaps bits or leaves them alone, so the number of 0s and 1s in the output must be the same as in the input. However since the Fredkin gate is *universal* there must be some construction of a Toffoli gate using multiple Fredkin gates and ancilla bits.

5. If the control bit is 0 then the central controlled-gate is not applied to the target qubit, which just experiences $VV^\dagger = \mathbb{1}$; if the control bit is 1 the target qubit experiences $VUV^\dagger$ as desired. This idea obviously generalises to controlled-controlled-gates, and since controlled-NOT is self-inverse we can simplify the construction above by replacing the outer Toffoli gates by simple controlled-NOT gates.
6. As previously noted the controlled-NOT gate implements the bitwise sum, that is the sum without carry, while the carry bit is 1 if and only if both a and b are 1. There is no need to explicitly preserve the second input as all gates applied to it are reversible.

9. Elementary quantum algorithms

1. The oracle will take the form of an f-controlled-NOT gate, and its parity can be determined in two calls with just two bits

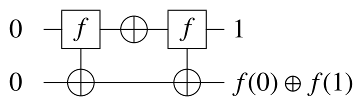

2. The circuit

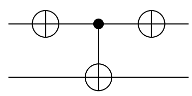

will achieve the desired result.

3. We have already proved f_{11} and f_{01}, and f_{00} is trivial, so the only interesting case is f_{10}. Using $H^2 = \mathbb{1}$ the circuit can be written as

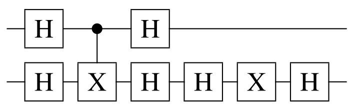

and the remaining steps follow easily by combining results for f_{01} and f_{11}.

4. The amplitude amplification operator is given by

$$U_{AA} = \frac{1}{2}\begin{pmatrix}1 & 1 & 1 & 1\\ 1 & -1 & 1 & -1\\ 1 & 1 & -1 & -1\\ 1 & -1 & -1 & 1\end{pmatrix}\begin{pmatrix}-1 & 0 & 0 & 0\\ 0 & 1 & 0 & 0\\ 0 & 0 & 1 & 0\\ 0 & 0 & 0 & 1\end{pmatrix}\frac{1}{2}\begin{pmatrix}1 & 1 & 1 & 1\\ 1 & -1 & 1 & -1\\ 1 & 1 & -1 & -1\\ 1 & -1 & -1 & 1\end{pmatrix}$$
$$= \frac{1}{2}\begin{pmatrix}1 & -1 & -1 & -1\\ -1 & 1 & -1 & -1\\ -1 & -1 & 1 & -1\\ -1 & -1 & -1 & 1\end{pmatrix}.$$

Suppose the satisfying function is f_{00}; the state after the function evaluation will be

$$\psi_{00} = \frac{1}{2}\begin{pmatrix}-1\\ 1\\ 1\\ 1\end{pmatrix}$$

and the final state can then be evaluated by multiplication to get

$$U_{AA}\psi_{00} = \frac{1}{4}\begin{pmatrix}1 & -1 & -1 & -1\\ -1 & 1 & -1 & -1\\ -1 & -1 & 1 & -1\\ -1 & -1 & -1 & 1\end{pmatrix}\begin{pmatrix}-1\\ 1\\ 1\\ 1\end{pmatrix} = \begin{pmatrix}-1\\ 0\\ 0\\ 0\end{pmatrix}$$

which is $-|00\rangle$. The other three possibilities can be evaluated in exactly the same way, and the answers are obvious by symmetry.

5. For the case of two satisfying inputs, it is simplest to choose a concrete case again, such as f_{00} and f_{01} matching. Then explicit matrix calculations show that no amplitude amplification occurs: a measurement is equally likely to give any of the four possible results. As before this argument applies whatever the two matches are. With three matches the situation is slightly more interesting, and amplitude amplification results in the final state being the single non-matching input, which now is the state marked with a unique phase.

6. Easily shown by direct calculation. As controlled-NOT gates are self inverse the decoding networks can be obtained by applying the same gates in reverse order.

7. The state of the five qubits as they enter the network is

$$(\alpha|000\rangle + \beta|111\rangle) \otimes |0\rangle \otimes |0\rangle = \alpha|00000\rangle + \beta|11100\rangle$$

and in general we have four possible states

$$|\psi_0\rangle \otimes |00\rangle = \alpha|00000\rangle + \beta|11100\rangle$$
$$|\psi_1\rangle \otimes |00\rangle = \alpha|10000\rangle + \beta|01100\rangle$$
$$|\psi_2\rangle \otimes |00\rangle = \alpha|01000\rangle + \beta|10100\rangle$$
$$|\psi_3\rangle \otimes |00\rangle = \alpha|00100\rangle + \beta|11000\rangle$$

where the subscript identifies the bit which has experienced a spin-flip error (0 indicating no error). Now run through the network of controlled-NOT gates.

$$|\psi_0\rangle \otimes |00\rangle \xrightarrow{\mathrm{CN}_{14}} \alpha|00000\rangle + \beta|11110\rangle \xrightarrow{\mathrm{CN}_{24}} \alpha|00000\rangle + \beta|11100\rangle$$
$$\xrightarrow{\mathrm{CN}_{15}} \alpha|00000\rangle + \beta|11101\rangle \xrightarrow{\mathrm{CN}_{35}} \alpha|00000\rangle + \beta|11100\rangle = |\psi_0\rangle \otimes |00\rangle$$

$$|\psi_1\rangle \otimes |00\rangle \xrightarrow{\mathrm{CN}_{14}} \alpha|10010\rangle + \beta|01100\rangle \xrightarrow{\mathrm{CN}_{24}} \alpha|10010\rangle + \beta|01110\rangle$$
$$\xrightarrow{\mathrm{CN}_{15}} \alpha|10011\rangle + \beta|01110\rangle \xrightarrow{\mathrm{CN}_{35}} \alpha|10011\rangle + \beta|01111\rangle = |\psi_1\rangle \otimes |11\rangle$$

$$|\psi_2\rangle \otimes |00\rangle \xrightarrow{\mathrm{CN}_{14}} \alpha|01000\rangle + \beta|10110\rangle \xrightarrow{\mathrm{CN}_{24}} \alpha|01010\rangle + \beta|10110\rangle$$
$$\xrightarrow{\mathrm{CN}_{15}} \alpha|01010\rangle + \beta|10111\rangle \xrightarrow{\mathrm{CN}_{35}} \alpha|01010\rangle + \beta|10110\rangle = |\psi_2\rangle \otimes |10\rangle$$

$$|\psi_3\rangle \otimes |00\rangle \xrightarrow{\mathrm{CN}_{14}} \alpha|00100\rangle + \beta|11010\rangle \xrightarrow{\mathrm{CN}_{24}} \alpha|00100\rangle + \beta|11000\rangle$$
$$\xrightarrow{\mathrm{CN}_{15}} \alpha|00100\rangle + \beta|11001\rangle \xrightarrow{\mathrm{CN}_{35}} \alpha|00101\rangle + \beta|11001\rangle = |\psi_3\rangle \otimes |01\rangle$$

The first three qubits (which are always control bits) are not changed by any of the controlled-NOT gates. Furthermore the states of the ancilla qubits 4 and 5 are the same in both components of the superposition, and so can be factored out as indicated. If a quantum state is separable then measuring one part has no effect on the other, and so the ancillas can be measured without affecting the logical qubit. Finally note that the four different ancilla states are all orthonormal, and so can be perfectly distinguished.

8. From the results of the previous question it is easy to write down the error correcting steps, as measuring the ancillas in the computational basis gives four distinct results with corresponding actions. For example if the ancillas are in $|01\rangle$ then the encoded qubits are in state $|\psi_3\rangle$, which can be fixed by applying a NOT gate to qubit 3; similar results apply in the other cases. For quantum control, note that these actions can all be implemented using generalised Toffoli gates, but implementing all these Toffoli gates is a lot of work. Another problem is that the ancilla qubits need to be reinitialized to $|0\rangle$ at the end; this is easy if the ancillas have been measured, as any ancillas in state $|1\rangle$ can be rest with NOT gates.
9. In a classical code, if two errors occur on different bits then two bits have the wrong value, and the majority vote approach "corrects" the third bit to the wrong value. (If the same bit is flipped both times then the situation is indistinguishable from the error free case). For the quantum code the state $|\psi_L\rangle$ is "corrected" to $\text{NOT}_L|\psi_L\rangle$.
10. For an initial state $|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$ the relevant density matrices are

$$\rho_a = \tfrac{1}{2}\left(\phi_{+z}|\psi\rangle\langle\psi|\phi_{+z} + \phi_{-z}|\psi\rangle\langle\psi|\phi_{-z}\right)$$

and

$$\rho_b = (1-p) \times |\psi\rangle\langle\psi| + p \times \mathrm{Z}|\psi\rangle\langle\psi|\mathrm{Z}.$$

Now $\phi_{\pm z} = \cos(\phi/2)\mathbb{1} \mp i\sin(\phi/2)\mathrm{Z}$ and so

$$\rho_a = \cos^2(\phi/2)|\psi\rangle\langle\psi| + \sin^2(\phi/2)\mathrm{Z}|\psi\rangle\langle\psi|\mathrm{Z}$$

with the other two terms cancelling. Clearly ρ_a and ρ_b have the same form, and they are identical if $p = \sin^2(\phi/2)$, which rearranges to various forms such as $p = (1 - \cos\phi)/2$ or $\phi = \arccos(1 - 2p)$.

10. More advanced quantum algorithms

1. There are four possible inputs, each of which has two possible outputs, giving a total of $2^4 = 16$ possible functions of which two are constant and six are balanced, with the last eight functions being neither constant or balanced (four give mostly 0 and four give mostly 1). For the rest of the question we only consider the constant and balanced cases. A single value of $f(x)$ tells us nothing while two values that disagree with each other indicates a balanced function. After three queries we know the result with certainty (either we have a disagreement, or three values are the same, and the function is constant). Hence the minimum number of queries is two and the maximum is three.

   For the average case, suppose the function is balanced and that $f(x_1) = 1$: then $f(x_2)$ will be 1 with probability $1/3$ and 0 with probability $2/3$. In the latter case we can stop; otherwise we will need one more query. So for a balanced function the average number of queries required is $2/3 \times 2 + 1/3 \times 3 = 7/3$, while for a constant function it is always necessary to use 3 queries. If the function is chosen to be constant or balanced with 50% probability, then the average number of queries is $(7/3 + 3)/2 = 8/3$. (If the function was chosen from amongst the 8 possible functions at random then the average query count would be $(6 \times 7/3 + 2 \times 3)/8 = 5/2$, but this is not what was asked!) On a quantum computer the minimum, maximum, and average query counts are all 1.

11. Trapped atoms and ions

1. The potential energy is given by

$$U = \frac{M}{2} \sum_{n=1}^{N} \left(\omega_r^2 r_n^2 + \omega_z^2 z_n^2\right) + \frac{e^2}{4\pi\epsilon_0} \sum_{m>n} \frac{1}{|\mathbf{r}_n - \mathbf{r}_m|}$$

   where the first group of terms is just the standard form for the potential energy of n harmonic oscillators, written in plane polar coordinates, and the second group of terms is the coulomb repulsions between the ions (the second sum goes over all pairs of ions, counting each pair only once).

2. For an ion travelling in free space the effect of the motion will be to cause Doppler shifts in the transition frequencies. The effect will depend on the velocity distribution, but the most common result is Doppler broadening. In a trap the motion is quantised as vibrations within the trap, and it is necessary to consider transitions between vibrational sub-levels of each electronic level. For a strictly harmonic trap these levels are equally spaced, with $E_n = (n + \frac{1}{2})h\nu$, and the selection rule $\Delta n = \pm 1$ results in a pair of sharp sidebands, separated from the sharp principal transition by $\pm h\nu$.

3. For Ca rearrangement gives $1/\Gamma = (3\epsilon_0\hbar\lambda^3)/(8\pi^2 e^2 a_0^2)$ and direct substitution gives a value of about $0.19\,\mu$s; for Be we get $1/\Gamma = (3\epsilon_0\hbar c^3)/(8\pi^2 e^2 a_0^2 \nu^3)$ and substitution gives a value of 6.8×10^9 s or 216 years. The $^{40}Ca^+$ ion trap avoids rapid relaxation by using

a forbidden transition, weakly allowed by electric quadrupole rules. The $^{9}Be^{+}$ ion trap achieves spatial discrimination by using optical Raman transitions.

4. The phase gate U_π negates $|01\rangle$ while leaving other states unchanged, and can be converted to the standard controlled-Z gate by applying X gates (NOT gates) to the first qubit before and after the U_π. Finally apply Hadamard gates to the second (target) qubit to get a controlled-NOT gate.
5. The first bit is just brute force:

$$|0\rangle\langle 0| \otimes \mathrm{Z} + |1\rangle\langle 1| \otimes \mathbb{1} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & -1 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix} + \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & -1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} = U_\pi.$$

Now the "massively entangled" state of 2 particles is just

$$U_\pi \mathrm{H}^{\otimes 2}|00\rangle = (|0\rangle\langle 0| \otimes \mathrm{Z} + |1\rangle\langle 1| \otimes \mathbb{1}) \times (|0\rangle|0\rangle + |0\rangle|1\rangle + |1\rangle|0\rangle + |1\rangle|1\rangle)$$

neglecting normalisation. Multiplying this out gives

$$|0\rangle\langle 0|0\rangle \otimes \mathrm{Z}|0\rangle + |0\rangle\langle 0|0\rangle \otimes \mathrm{Z}|1\rangle + |0\rangle\langle 0|1\rangle \otimes \mathrm{Z}|0\rangle + |0\rangle\langle 0|1\rangle \otimes \mathrm{Z}|1\rangle$$
$$+|1\rangle\langle 1|0\rangle \otimes \mathbb{1}|0\rangle + |1\rangle\langle 1|0\rangle \otimes \mathbb{1}|1\rangle + |1\rangle\langle 1|1\rangle \otimes \mathbb{1}|0\rangle + |1\rangle\langle 1|1\rangle \otimes \mathbb{1}|1\rangle$$

As usual all the inner products can be replaced by 0 or 1, and dropping the (pointless) $\mathbb{1}$ operators this simplifies to

$$|0\rangle \mathrm{Z}|0\rangle + |0\rangle \mathrm{Z}|1\rangle + |1\rangle|0\rangle + |1\rangle|1\rangle = (|0\rangle \mathrm{Z} + |1\rangle)(|0\rangle + |1\rangle).$$

The corresponding state for three atoms is

$$(|0\rangle \mathrm{Z} + |1\rangle)(|0\rangle \mathrm{Z} + |1\rangle)(|0\rangle + |1\rangle)$$

and multiplying this out and using $\mathrm{Z}|0\rangle = |0\rangle$ and $\mathrm{Z}|1\rangle = -|1\rangle$ gives

$$|0\rangle|0\rangle|0\rangle - |0\rangle|0\rangle|1\rangle - |0\rangle|1\rangle|0\rangle - |0\rangle|1\rangle|1\rangle + |1\rangle|0\rangle|0\rangle - |1\rangle|0\rangle|1\rangle + |1\rangle|1\rangle|0\rangle + |1\rangle|1\rangle|1\rangle$$

matching the result given (neglecting normalisation of course).

12. Nuclear magnetic resonance

1. From the exercises in Part I we know that a 12 T field gives a ^{1}H Larmor frequency of about 500 MHz, and so we need a field difference of $100 \div (500 \times 10^{6}) \times 12 = 2.4 \times 10^{-6}$ T per Å, or 24000 T/m. Generating field gradients of this size is challenging.
2. There are two reasonable approaches to this. The first is to use perturbation theory, writing the Hamiltonian as $\mathcal{H} = \mathcal{H}_0 + \mathcal{H}_1$ where

$$\mathcal{H}_0/\hbar = \omega_1 \frac{\sigma_{1z}}{2} + \omega_2 \frac{\sigma_{2z}}{2}$$

and

$$\mathcal{H}_1/\hbar = \omega_{12} \frac{\sigma_{1x}\sigma_{2x} + \sigma_{1y}\sigma_{2y} + \sigma_{1z}\sigma_{2z}}{4}$$

First order perturbation theory says that the eigenstates are unaffected by the coupling, and the eigenvalues are changed by the diagonal matrix elements of the perturbation, that is just the z terms. Thus to first order the Heisenberg coupling can be replaced by an Ising coupling. To check that this approach is valid we need the first order effect on the eigenvectors, and in general the other states get mixed in according to

$$a_k^{(1)} = \frac{\langle k|\mathcal{H}_1|m\rangle}{E_m - E_k}$$

The only off diagonal elements in the perturbation connect the two central states; these are of size $\frac{1}{2}\,\omega_{12}$ and so mixing is negligible if

$$\left|\frac{\omega_{12}}{\omega_1 - \omega_2}\right| \ll 1.$$

An alternative approach is to diagonalize the full Hamiltonian and then take appropriate limits to show that the eigenvalues can be approximated by the diagonal elements.

3. As in the previous chapter the critical step is to make a standard controlled-Z gate, as this can be converted to a controlled-NOT gate with a couple of Hadamard gates. Now the controlled-Z can be decomposed (up to an irrelevant global phase) as evolution under the Ising coupling for a time $\tau = \pi/\omega_{12}$ together with a -90_z rotation on both qubits. The Ising term can be implemented with a spin echo as usual. It might be argued that the -90_z gates are not standard, but these can be replaced by a sequence of three S gates, where $S = \sqrt{Z}$ is a standard gate. For a NOT gate that takes the same length of time start from a spin echo which refocuses everything, and put a NOT gate on the beginning or end; if you are careful this NOT gate will cancel an earlier one.
4. Assuming the couplings take the Ising form the Hamiltonian is

$$\mathcal{H}/\hbar = \tfrac{1}{2}\left(\omega_1\sigma_{1z} + \omega_2\sigma_{2z} + \omega_3\sigma_{3z}\right) + \tfrac{1}{4}\left(\omega_{12}\sigma_{1z}\sigma_{2z} + \omega_{13}\sigma_{1z}\sigma_{3z} + \omega_{23}\sigma_{2z}\sigma_{3z}\right)$$

A possible spin echo network is

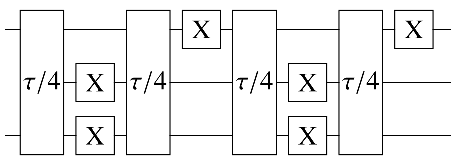

5. To reduce an apparent Larmor frequency combine a period of free precession with a period under a spin echo. To change the sign of a Larmor frequency surround a period of free precession with NOT gates. Any single component of a Hamiltonian, including couplings, can be rescaled in the range ±1, but simultaneously rescaling multiple elements gets complicated and is not always possible.