Entanglement and entangling gates

Artur Ekert and Alastair Kay

I. ENTANGLEMENT

Single qubit gates can be used to transform the input state $|0\rangle|0\rangle...|0\rangle$ of the n qubit register into any state of the type $|\Psi_1\rangle\ |\Psi_2\rangle...\ |\Psi_n\rangle$, where $|\Psi_i\rangle$ is an arbitrary superposition of $|0\rangle$ and $|1\rangle$. These are rather special n-qubit states, called the product states or the separable states. In general, a quantum register of size $n > 1$ can be prepared in states which are not separable – they are known as entangled states. For example, for two qubits ($n = 2$), the state

$$\alpha\ |00\rangle + \beta\ |01\rangle = |0\rangle \otimes (\alpha\ |0\rangle + \beta\ |1\rangle)$$

is separable, $|\Psi_1\rangle = |0\rangle$ and $|\Psi_2\rangle = \alpha\ |0\rangle + \beta\ |1\rangle$, whilst the state

$$\alpha\ |00\rangle + \beta\ |11\rangle \neq |\Psi_1\rangle \otimes |\Psi_2\rangle$$

is entangled ($\alpha, \beta \neq 0$), because it cannot be written as a tensor product. That is, subsystems do not have states of their own. We will usually drop the tensor product symbol and write product states as $|\Psi_1\rangle|\Psi_2\rangle$. In order to entangle two (or more qubits) we have to extend our repertoire of quantum gates to two-qubit gates.

II. CONTROLLED-NOT

The most popular two-qubit gate is the controlled-NOT (C-NOT). It flips the second (target) qubit if the first (control) qubit is $|1\rangle$ and does nothing if the control qubit is $|0\rangle$. In the computational basis $\{|00\rangle, |01\rangle, |10\rangle, |11\rangle\}$ the gate is represented by the unitary matrix

$$C = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix}$$

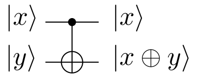

where $x, y = 0$ or 1 and $\oplus$ denotes XOR or addition modulo 2. If we apply the C-NOT to Boolean data in which the target qubit is $|0\rangle$ and the control is either $|0\rangle$ or $|1\rangle$ then the effect is to leave the control unchanged while the target becomes a copy of the control, *i.e.*

$$|x\rangle|0\rangle \mapsto |x\rangle|x\rangle \qquad x = 0, 1.$$

One might suppose that this gate could also be used to copy superpositions such as $|\Psi\rangle = \alpha\ |0\rangle + \beta\ |1\rangle$, so that

$$|\Psi\rangle|0\rangle \mapsto |\Psi\rangle|\Psi\rangle$$

for any $|\Psi\rangle$. This is not so! The unitarity of the C-NOT requires that the gate turns superpositions in the control qubit into *entanglement* of the control and the target. If the control qubit is in a superposition state $|\Psi\rangle = \alpha|0\rangle + \beta|1\rangle$, ($\alpha, \beta \neq 0$), and the target in $|0\rangle$ then the C-NOT generates the entangled state

$$(\alpha|0\rangle + \beta|1\rangle)\,|0\rangle \mapsto \alpha|00\rangle + \beta|11\rangle.$$

For example, the following simple network

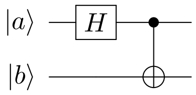

evolves the four inputs, $|ab\rangle$, $a, b = 0, 1$, into the four entangled states of two qubits known as the Bell states,

$$|00\rangle \mapsto \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle), \quad |01\rangle \mapsto \frac{1}{\sqrt{2}}(|01\rangle + |10\rangle)$$
$$|10\rangle \mapsto \frac{1}{\sqrt{2}}(|00\rangle - |11\rangle), \quad |11\rangle \mapsto \frac{1}{\sqrt{2}}(|01\rangle - |10\rangle)$$

The Bell states form another convenient basis in the four dimensional tensor product space of two qubits. Measurements in the Bell basis, also referred to as the Bell measurement, can be implemented by running the network backwards.

III. THOU SHALT NOT CLONE

Let us notice in passing that it is impossible to construct a universal quantum cloning machine effecting the transformation

$$|\Psi\rangle|0\rangle|W\rangle \mapsto |\Psi\rangle|\Psi\rangle|W'\rangle$$

where $|W\rangle$ refers to the state of the rest of the world and $|\Psi\rangle$ is *any* quantum state. To see this take any two normalised states $|\Psi\rangle$ and $|\Phi\rangle$ which are non-identical ($|\langle\Phi|\Psi\rangle| \neq 1$) and non-orthogonal ($\langle\Phi|\Psi\rangle \neq 0$), and run the cloning machine,

$$|\Psi\rangle|0\rangle|W\rangle \ \mapsto \ |\Psi\rangle|\Psi\rangle|W'\rangle \qquad (1)$$
$$|\Phi\rangle|0\rangle|W\rangle \ \mapsto \ |\Phi\rangle|\Phi\rangle|W^{''}\rangle \qquad (2)$$

As this must be a unitary transformation which preserves the inner product hence we must require

$$\langle\Phi|\Psi\rangle = \langle\Phi|\Psi\rangle^2\langle W^{'}|W^{''}\rangle \qquad (3)$$

and this can only be satisfied when $|\langle\Phi|\Psi\rangle| = 0$ or 1, which contradicts our assumptions. Thus states of qubits, unlike states of classical bits, cannot be faithfully cloned. This leads to interesting applications, quantum cryptography being one such.

IV. TELEPORTATION

An unknown quantum state cannot be cloned but it can be teleported. Consider the following network

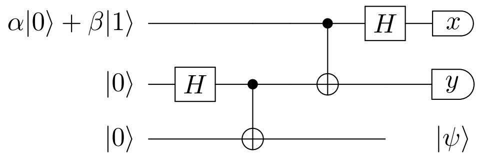

The first input qubit (counting from the top) is in some arbitrary state. The second and the third qubit are prepared in the Bell state $|00\rangle + |11\rangle$. Suppose the first and the second qubit is held by one person, say Alice, and the third qubit is given to another person, Bob, who may be miles away from Alice. In order to teleport the state of the first qubit Alice performs the Bell measurement on the first two qubits, which gives two binary digits, x and y. She then broadcasts x and y. Upon learning the values of x and y Bob chooses one of the four transformations,

$$00 \to \mathbb{1}, \quad 01 \to X, \quad 10 \to Z, \quad 11 \to ZX,$$

(e.g. if $x = 0, y = 1$ he chooses X) and applies it to his qubit. This restores the original state of the first qubit, which was destroyed when Alice performed the Bell measurement.

V. CONTROLLED UNITARIES

Another common two-qubit gate is the controlled phase shift gate $B(\phi)$ defined as

$$B(\phi) = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & e^{i\phi} \end{pmatrix}$$

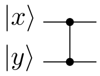

$$\left.\vphantom{\begin{matrix}a\\b\\c\\d\end{matrix}}\right\} e^{ixy\phi}|x\rangle|y\rangle.$$

Again, the matrix is written in the computational basis $\{|00\rangle, |01\rangle, |10\rangle, |11\rangle\}$. Here any of the two qubits can be viewed as the control or the target. The gate turns superpositions into entanglement, for example, the controlled sign flip gate maps

$$\begin{aligned} B(\pi) : &|00\rangle + |01\rangle + |10\rangle + |11\rangle \\ \mapsto &|00\rangle + |01\rangle + |10\rangle - |11\rangle. \end{aligned}$$

More generally, these various 2-qubit controlled gates are all of the form controlled-U, for some single-qubit unitary transformation U. The controlled-U gate applies the identity transformation to the target qubit when the control qubit is in state $|0\rangle$ and applies an arbitrary prescribed U when the control qubit is in state $|1\rangle$. The gate maps $|0\rangle|\Psi\rangle$ to $|0\rangle|\Psi\rangle$ and $|1\rangle|\Psi\rangle$ to $|1\rangle(U|\Psi\rangle)$.

Any unitary operation U can be expressed as

$$U = B^\dagger X B A^\dagger X A,$$

for some unitaries A and B. This allows to construct the controlled-U, using a couple of controlled-NOT gates, as

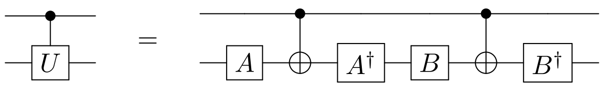

When the control qubit is $|1\rangle$, the controlled-NOT gates apply X gates on the target, and the evolution of the target qubit is U. On the other hand, if the control is $|0\rangle$, then the controlled-NOT gates do nothing and the evolution is $B^\dagger B A^\dagger A = \mathbb{1}$. Subsequently, any controlled-controlled-U (which implements U on the target qubit if all the control qubits are in the $|1\rangle$ state, and otherwise does nothing) can be applied using

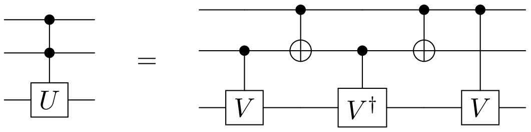

where V is any unitary matrix satisfying $V^2 = U$. The choice $V = \sqrt{\text{NOT}}$ leads to another useful gate, known as the controlled-controlled-NOT gate (c^2-NOT), or the Toffoli gate. It flips the target only if the two control qubits are both set to 1 and does nothing otherwise. This gate appears frequently in circuits which evaluate Boolean functions.

VI. UNIVERSALITY

The Hadamard gate, all phase gates, and the C-NOT, form an infinite *universal set of gates i.e.* if the C-NOT gate as well as the Hadamard and all phase gates are available then any n-qubit unitary operation can be constructed exactly with $O(4^n n)$ such gates. Here and in the following we use the asymptotic notation; given a positive function $f(n)$, the symbol $O(f(n))$ means bounded from above by $c\,f(n)$ for some constant $c > 0$ (for sufficiently large n). For example, $15n^2 + 4n + 7$ is $O(n^2)$. We should mention that there are many universal sets of gates. In fact, almost any gate which can entangle two qubits can be used as a universal gate. A *finite* universal set of gates, such as the one containing the Hadamard, $P_{\pi/4}$ and the C-NOT, can approximate any unitary operation on n qubits with arbitrary precision. The prize to pay is the number of gates — better precision requires more gates. Note that when categorising the scaling properties as exponential or polynomial, it does not matter which universal set of gates we use. For instance, if a circuit requires $3n^2$ gates drawn from the set {NOT, AND, OR}, then since each one of these can be created by no more than 3 NAND gates, the same family of circuits can be described by no more than $27n^2$ NAND gates, and hence the $O(n^2)$ scaling property in unchanged. The same argument applies to quantum circuits.