Entangling Ions for Fun and Profit

Cirac–Zoller Entanglement through the Shared Motion of Trapped Ions

Jonathan Feasby
Pathfinder Quantum Technologies

July 2026

Abstract

We start with a somewhat impossible-sounding claim: we are going to flip two random coins in a way that guarantees they always agree—both heads, or both tails, but never one of each. Run it a thousand times and we never once see a mismatch.

This counterintuitive bit of physics was proposed by Ignacio Cirac and Peter Zoller in 1995 [1], and is pulled off by shining lasers at ions held still in an electromagnetic trap. Their proposal established the central idea underlying many later trapped-ion entangling gates: using shared motion to entangle otherwise separate qubits.

The article runs in three parts. First we frame the problem: create a single-ion superposition, then convince ourselves that the obvious two-ion extension fails. Next we assemble the toolkit—shared motion, phonons, and sideband coupling. Finally we derive a stripped-down Bell-state preparation that uses the same central mechanism as Cirac and Zoller: write into a shared motional mode with one ion, and read it out with another.

1. Framing the problem

1.1. Creating superposition: one ion, one laser

We start with a single ion held in a trap. Out of two of its internal energy levels we define a qubit: a lower-energy state we call $|g\rangle$ (the ground state) and a higher-energy state we call $|e\rangle$ (the excited state). In quantum computing notation, we map $|g\rangle$ to $|0\rangle$ and $|e\rangle$ to $|1\rangle$.

In order to perform computations, we need the ability to control this two-level system. At the very minimum we need to be able to take an ion prepared in $|g\rangle$ and drive it into $|e\rangle$. How, physically, does one flip $|g\rangle$ to $|e\rangle$?

Here we can copy the homework of previous physicists. In 1913 Niels Bohr showed that atomic energy levels are discrete; an atom cannot absorb just any photon; it absorbs light efficiently when the photon energy matches the energy gap between two allowed states. That is to say, *the ground-to-excited state transition has a specific resonant frequency.*

This is not quite the same as saying the ion simply absorbs one photon, sits in the excited state, and later emits it again. For qubit operations, the useful point is that the laser field drives a coherent oscillation between the two states.

If we leave the resonant laser on, the ion oscillates between $|g\rangle$ and $|e\rangle$. The rate of this oscillation is the *Rabi frequency*, denoted Ω, and is proportional to the strength of the light. If the laser strength is constant, the speed of the $|g\rangle \leftrightarrow |e\rangle$ oscillation is constant as well.

This answers the question of how to flip between $|g\rangle \leftrightarrow |e\rangle$. A full flip from $|g\rangle$ to $|e\rangle$ is a π-pulse; a flip from $|g\rangle$ to $|e\rangle$ and back to $|g\rangle$ is a 2π-pulse.

Generating an equal superposition of the two states is now as simple as turning the laser off halfway through the ground-to-excited cycle. This is a $\pi/2$-pulse; the atom is then in a superposition of ground and excited states:

$$|g\rangle \;\longrightarrow\; \frac{1}{\sqrt{2}}\big(|g\rangle + |e\rangle\big). \tag{1}$$

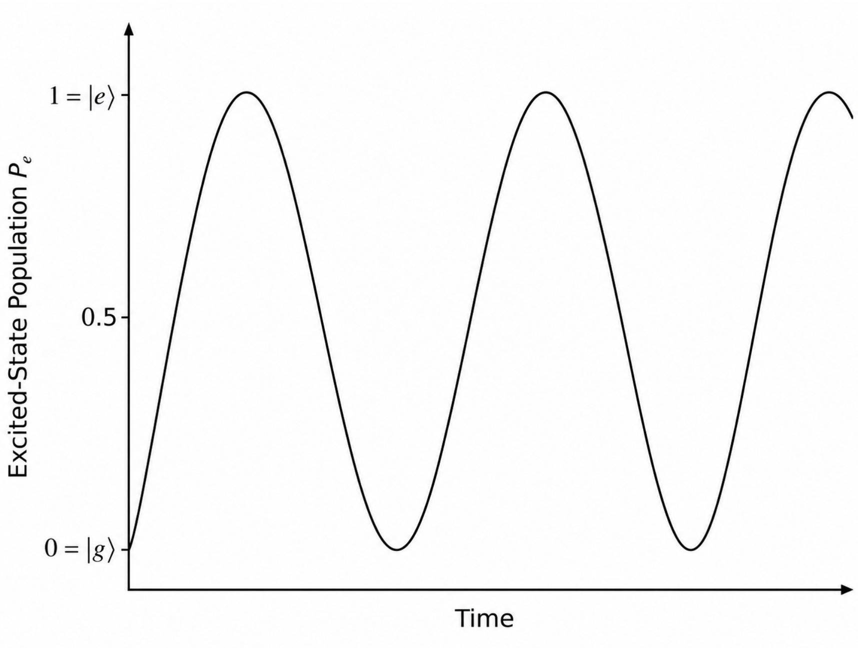

Figure 1. The ground-to-excited-state transition is discrete, but the population in each state changes continuously during a coherent laser pulse. At the top of the curve, an excited-state population of 1 means a 100% probability of finding the ion in the excited state at that time.

Resonant laser pulses are how all single-qubit rotations are performed, and they are relatively easy to do well.

1.2. Two ions, and why the obvious idea fails

Given how straightforward creating a single-ion superposition is, one may be tempted to assume the two-ion case is just more of the same. Take the same resonant laser, shine it at both target ions for the right amount of time, and prepare two superpositions. Surely two superpositions, side by side in the same chain, will give us something entangled?

Alas, it does not.

This is analogous to flipping two coins at once: each lands heads or tails with a 50/50 chance, but the result of one coin tells you nothing about the other. We could never infer the result of the second coin purely by looking at the first. Yet it is exactly this correlated outcome that we need. We want to flip two coins and *know* that heads on one guarantees heads on the other, and tails on one guarantees tails on the other.

The straightforward "shine the superposition-creating light at both" approach gives the uncorrelated, coin-like state. The truth table for the coin flips looks like this:

	Heads	Tails
Heads	HH	HT
Tails	TH	TT

Each outcome has a 25% chance of occurring. The corresponding ion state is

$$\frac{1}{\sqrt{2}}(|g\rangle + |e\rangle) \otimes \frac{1}{\sqrt{2}}(|g\rangle + |e\rangle) = \frac{1}{2}(|gg\rangle + |ge\rangle + |eg\rangle + |ee\rangle). \tag{2}$$

This is a separable state. It factorises cleanly into "ion one is doing its thing" and "ion two is doing its thing", and a state that factorises like this is, by definition, not entangled. Measuring one ion tells you nothing about the other.

This makes sense physically: each laser only talks to the ion it is pointed at. Ion one's evolution never references ion two, so the operation we applied is just a product of two independent single-ion operations—two unrelated coin flips.

The quantum version is not merely that the outcomes are correlated after measurement; before measurement, the system is in a coherent superposition of the two joint possibilities. The state we want, following the coin analogy, is

	Heads	Tails
Heads	HH	
Tails		TT

where each result has a 50% chance of occurring. The corresponding ion state is

$$\frac{1}{\sqrt{2}}(|gg\rangle + |ee\rangle). \tag{3}$$

This tells us that, in order to entangle the two ions, the physics acting on one must somehow depend on the state of the other.

The whole apparatus of the operation we are circling in on is an answer to the question: how *exactly* do we make one ion feel the state of the other? Cirac and Zoller's idea was to use something both ions already share: the collective motion of the ion chain.

The ions may have separate internal states, but they are not mechanically independent. As we explain in the next section, if one ion is pushed, the chain's shared vibrational mode can carry that disturbance to the other ion. The motion becomes a temporary "bus" between them. That is the basic trick: do not try to make the ions talk directly; make them talk *through the motion of the chain.*

2. Defining the toolkit

2.1. Shared motion

In a linear trap, multiple ions can settle into a one-dimensional structure known as a linear Coulomb crystal (Fig. 2).

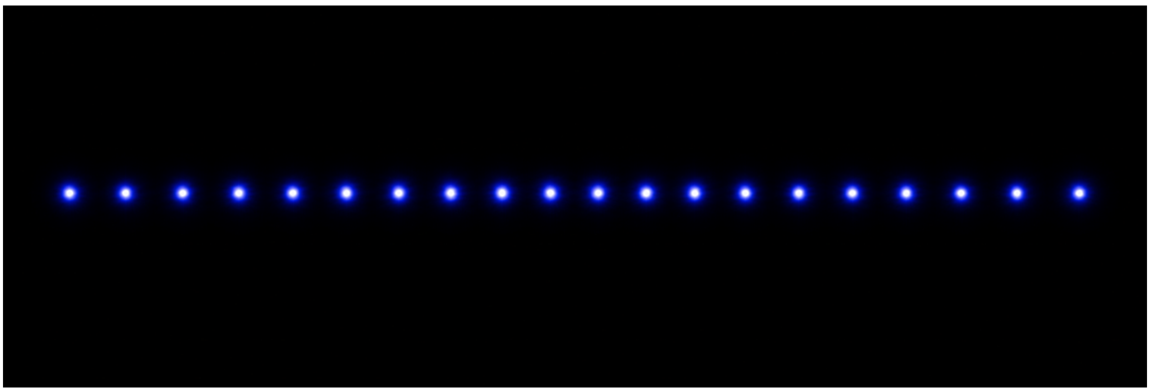

Figure 2. A linear Coulomb crystal of trapped ions. Like charges repel; the trap confinement balances that repulsion so that a push on one ion excites a shared vibration of the whole chain.

The ions in a trap are all positively charged, and like charges repel. Squeeze two positively charged ions towards each other and Coulomb repulsion shoves them back apart. The trap apparatus is tuned such that the Coulomb repulsion is perfectly balanced by the trap's confinement strength. This is a pain to realise physically, but the upshot is worth it: the ions no longer move as completely independent particles. A small push on one ion excites the shared vibration of the whole chain.

We wanted a way for one ion to feel another, and the shared motion of the chain is exactly that. The shared motion is a wire, or "bus", as it is commonly called, connecting the two ions, a channel through which whatever we do to ion one can be felt by ion two. The rest of the gate is an exercise in sending information into and out of this motional bus.

2.2. Phonons

We have our bus. Now, what travels along it?

Whatever the signal is, it has to be made of motion, since motion is the only thing the shared mode can carry. So picture giving the chain a push... how big a push can you give it? One might reasonably guess any size one likes, as if it were a smooth dial ranging from "barely moving" to "shaking hard", but that turns out to be incorrect.

Motion in a trap is quantised. More precisely, each collective vibrational mode behaves like a quantum harmonic oscillator. It does not behave like a smooth dial; rather it comes in discrete, equal steps like rungs on a ladder, and one cannot add half a step. Each step is one *phonon*.

A phonon is not a little ball moving down the chain. It is one quantum of excitation in a chosen collective motional mode. If the shared mode has zero phonons, we write $|n=0\rangle$. If it has one phonon, we write $|n=1\rangle$. In the general case, it is simply $|n\rangle$.

So when we write $|g,n\rangle$, we mean: the ion is in its internal ground state $|g\rangle$, and the shared motional mode contains n phonons.

2.3. Sending a signal through the bus

We have a bus, which is the shared motion. We have a language for it, which is phonons. How, then, do we actually send a signal?

The way we interact with our trapped ions is through lasers. If the laser is tuned to the ordinary qubit transition, it drives the *carrier* $|g,n\rangle \leftrightarrow |e,n\rangle$. The ion flips, but the phonon count stays the same.

But there are also neighbouring transitions where the ion flips and the motion changes at the same time. These are the *motional sidebands*.

- The *blue sideband* sits at $\omega_0+\omega_m$, where ω_0 is the qubit transition frequency and ω_m is the frequency of the shared motion we are using as the bus. Driving the blue sideband flips the ion and *adds* a phonon: $|g,n\rangle \leftrightarrow |e,n+1\rangle$.
- The *red sideband* sits at $\omega_0-\omega_m$. Driving the red sideband flips the ion and *removes* a phonon: $|g,n\rangle \leftrightarrow |e,n-1\rangle$.

This is the key mechanism: a sideband pulse lets us tie the ion's internal state to the motion of the chain. Repeat that, because it is really important. *A sideband pulse lets us entangle the ion's internal energy state with the motional state of the chain.* For example, suppose the shared mode starts in $|n=0\rangle$ and the ion starts in $|g\rangle$. If we apply a blue-sideband $\pi/2$-pulse, we coherently split the state into two branches:

$$|g,0\rangle \longrightarrow \frac{1}{\sqrt{2}}\big(|g,0\rangle+|e,1\rangle\big). \tag{4}$$

In one branch, the ion remains in $|g\rangle$ and the motion remains in $|n=0\rangle$. In the other branch, the ion is excited to $|e\rangle$ and the motion has gained one phonon. Because the two alternatives remain in a coherent superposition, and because the joint state cannot be factorised into an independent ion state and motional state, the ion and motion are entangled.

It is not exactly the kind of entanglement we want—the ion's energy level is entangled with the motion, not yet with the other ion—but we can work with it (Fig. 3).

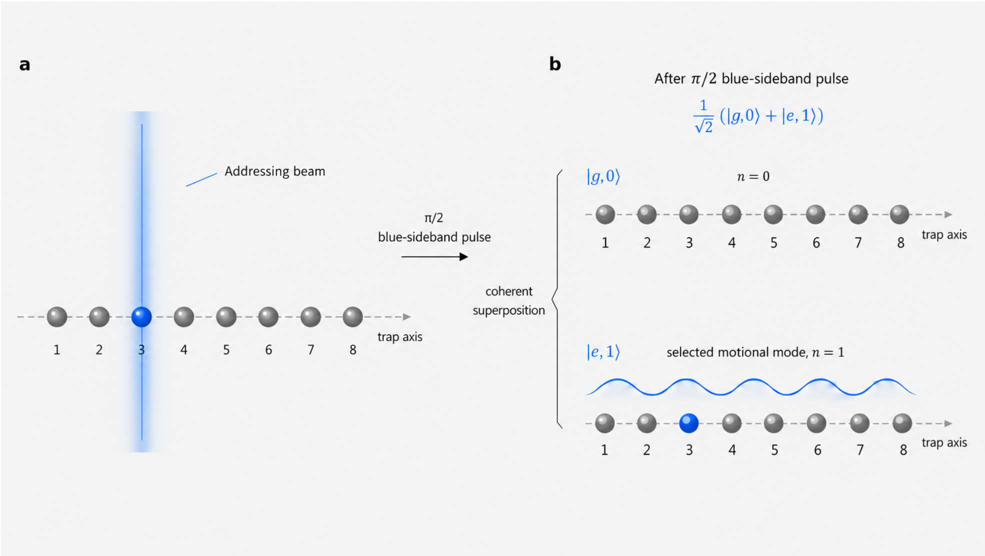

Figure 3. If the first ion can write information into the shared motion, and the second ion can read information out of the shared motion, then the motion can act as a bus between them. That is the whole trick behind sideband entanglement.

3. Deriving a solution

Note. What follows is not the complete Cirac–Zoller controlled gate. It is a stripped-down Bell-state preparation that uses the same central mechanism: writing information into a shared motional mode and reading it out through another ion.

Take a chain of ions with target ions i and j, and hit ion i with a blue sideband laser for a pulse time corresponding to $\pi/2$. This produces our familiar superposition (phase bookkeeping not included):

$$\frac{1}{\sqrt{2}}(|g,n\rangle + |e,n+1\rangle). \tag{5}$$

The motion is now carrying a record of what happened to i. Starting from the joint state $|gg,0\rangle$, the evolution is

$$|gg,0\rangle \longrightarrow \frac{1}{\sqrt{2}}(|gg,0\rangle + |eg,1\rangle). \tag{6}$$

We then hit ion j with a π-pulse (full bit-flip) of the red sideband, which excites ion j and returns the motion to its ground state in the case where ion i is in $|e\rangle$ and the shared mode contains one phonon. Crucially, the $|gg,0\rangle$ component is unaffected. A red-sideband transition must remove one phonon, but the phonon count is already $|0\rangle$. Since there is no phonon state $|-1\rangle$—there cannot be negative amounts of motion—the sideband laser cannot couple to the phonon-reducing transition:

$$|gg,0\rangle \to |gg,0\rangle, \qquad |eg,1\rangle \to |ee,0\rangle. \tag{7}$$

This is the heart of the entanglement protocol. The first pulse creates an equal coherent superposition of two joint ion–motion states: one in which ion i remains in $|g\rangle$ and the mode remains in $|0\rangle$, and one in which ion i is excited while the mode gains one phonon. Or, to follow the coin analogy: the first coin flip creates an equal superposition of heads and tails, and the second coin flip can only be heads if the first coin is heads, and can only be tails if the first coin is also tails.

3.1. Experimentally speaking

Step by step, the entanglement protocol is as follows:

1. Drive ion i with a $\pi/2$ pulse of the blue sideband, creating the superposition

$$|\psi_i\rangle = \frac{1}{\sqrt{2}}\big(|g,0\rangle + |e,1\rangle\big).$$

2. Drive ion j with a π pulse of the red sideband. This is a full population-transfer pulse that depends on the state of ion i. If i is in $|g,0\rangle$, the pulse on ion j does nothing. If i is in $|e,1\rangle$, the pulse drives j through the transition $|g,1\rangle \rightarrow |e,0\rangle$.

   This cleanly produces the desired Bell state (ignoring phase accumulation for now):

$$|\psi_{\text{final}}\rangle = \frac{1}{\sqrt{2}}\big(|gg,0\rangle + |ee,0\rangle\big). \tag{8}$$

We can now factor out the motion, since it is no longer entangled with any energy level:

$$|\psi_{\text{final}}\rangle = \frac{1}{\sqrt{2}}\big(|gg\rangle + |ee\rangle\big) \otimes |0\rangle. \tag{9}$$

This is the ion-equivalent of the heads–heads / tails–tails state from the introduction.

4. Conclusion

Cirac and Zoller's key proposal was that the quantised collective motion of a trapped-ion chain could mediate entanglement between individually addressed ions. The original scheme came with a cost: the selected motional bus mode had to be prepared near its ground state, making the gate extremely sensitive to imperfect cooling and motional heating.

Four years later, Klaus Mølmer and Anders Sørensen proposed an alternative method that could produce entanglement without requiring the bus mode to begin in its ground state [2]. That is the protocol examined in the companion article.

References

[1] J. I. Cirac and P. Zoller, *Quantum Computations with Cold Trapped Ions*, Phys. Rev. Lett. **74**, 4091 (1995).

[2] A. Sørensen and K. Mølmer, *Quantum Computation with Ions in Thermal Motion*, Phys. Rev. Lett. **82**, 1971 (1999).