

The following content is provided under a Creative Commons license. Your support will help MIT OpenCourseWare continue to offer high quality educational resources for free. To make a donation or view additional materials from hundreds of MIT courses, visit MIT OpenCourseWare at ocw.mit.edu.

PROFESSOR: OK, so let's start to discuss ion traps. Now, in many situations, ions are just charged neutral atoms. And if the charge doesn't matter, you can do very similar physics with ions and neutral atoms.

However, the charge makes major differences when it comes to confining and manipulating particles. So when I compare neutrals to ions, the forces which we exert on neutral atoms are magnetic forces, optical forces. And typically, the traps we can create have a depth of about 1 kelvin.

Actually, you just need superconducting magnets. The magnetic traps we build here at MIT have millikelvin depths. So the trap depth is often on the order of microkelvin, millikelvin, and kelvin at the high scale.

Well, this is very different in ions. In ions, you have the Coulomb force. If you take electric fields of 100 volts per millimeter, you can create confinement forces for 10,000 kelvin. So you have no problem in confining room temperature ions.

When it comes to-- so this is trapping and confinement. If you talk about cooling, well, our workhorses are Doppler cooling and evaporative cooling. You can actually cool ions by evaporation. You can also cool ions by evaporative cooling.

But what has been very common and special to ions is the ability to resolve sidebands, sidebands of the vibration and motion in the harmonic oscillator trap. And in problem set number 10, you have looked through the physics of sideband cooling. However, sideband cooling is now also coming back to neutral atoms. For instance, if you have neutral atoms in a very tightly confining dipole trap, you can resolve sidebands and use those for cooling.

OK, what is the science? Well, a driving force for both ions and neutral atoms has been clocks. If atoms are cold, you can use them to have very accurate ways of keeping time. And I think currently, there is a race between the based neutral atom atomic clocks, neutral [INAUDIBLE] optical lattices, and ion traps in particular, the aluminum ion at Boulder.

So we have clocks. In ions-- and this is what I want to talk a little bit more about today-- a major driving force was quantum information. You have two ions and get quantum gates between them, or you have eight ions and can have now eight qubits in a prototypical quantum computer. So quantum information science and quantum algorithms was quite important. But I would say some of the physics is at least discussed and partially pursued with atoms in optical lattices.

We have talked a lot about, for neutral atoms, quantum simulations, studying super fluidity, studying phase transitions, magnetic properties of many body systems. But there are more efforts to do that with ions. Chris Monroe's group in Maryland is studying antiferromagnetic ordering of a spin chain. And the spin chain consists of a string of ions in an ion trap. So see actually on the science side that all three-- clocks, precision measurements, quantum information, qubits, quantum algorithms, and quantum simulation of many body systems-- are done with both neutral atoms and ions. So these are now two very different platforms.

So I could spend a long time teaching about special aspects of ions. How do you do a trap? How do you cool ions? How do you do quantum state control? Because this is our last lecture, I decided to focus on two aspects.

One is, I want to talk about a special way of trapping ions. Because there is one difference for ions. Ions, until recently, have not been trapped in stationary traps. They've only been trapped in dynamic traps. So you cannot trap ions by simply putting an electric potential around it.

Because the divergence of the electric field is 0. And that means there's always one direction in space where the ions are dragged out. You cannot have an inward force from all directions.

So therefore, when you want to use electric fields, you can only use dynamic traps. Well, for the experts I should say, this does not apply for the optical dipole trap. We know how we can create trapping potentials for neutral atoms using the optical dipole trap. And very recently, a group was able to trap ions not with electric and magnetic fields provided by electrodes or radio frequency, but by inner focused laser beams. So it's a very new development. But it uses exactly the mechanism which has been used for neutral atoms.

So when it comes to dynamic traps, the standard two traps which have been used for ion trapping-- one is the Penning trap, which I do not want to describe in detail. The Penning trap uses a quadrupole configuration of two positive charges. And then you have a ring with negative charge. That means that an ion is confined here but would see anti-confinement in the outward direction.

Now, what you do with the anti-confinement? Well, the solution is that you add a magnetic field. And what happens now is the particle is well-confined axially. But radially, when it wants to move out, it's turned into a cyclotron orbit, and therefore stays trapped. So this is the principal of Penning traps where the confinement in the axial direction is electrostatic, and the confinement in x, y is through cyclotron motion.

OK, it's wonderful physics to describe the motion of ions in this trap. It's the combination of an axial motion, cyclotron motion, magnetron motion. It's very interesting physics. And for many years, Dave Pritchard here at MIT used such Penning traps to measure the masses, determine the masses, of ions with the highest precision ever possible. MIT still has the world record of the most precise mass measurements at the 10 to the minus 11 level. And this was done in Penning traps.

But I want to talk more today about RF traps, or Paul traps, because they are used in the majority of experiments. And they're used in pretty much all experiments on quantum information processing with ions. So the neat thing-- and I try to teach you concepts in physics-- the interesting thing about Paul traps or RF traps is that you

alternate at a radio frequency between confinement and anti-confinement.

You can say you have a settle point potential, which is confining in x and anti-confining in y . But then, you switch x and y . So therefore, the situation is just imagine you're an ion, and you feel half of the time a confining harmonic oscillator and half of the time an anti-confining harmonic oscillator. The question I have for you, and this is a general physics question, well, is the sum of a confining and an anti-confining potential of the same strength give 0, or does it give something which in the end is confining?

So if you were in a potential, and the potential switches permanently from parabolic upward to parabolic downward, at any given position, the average force is 0. Because in one phase, you have a force to the left. In the other case, you have the opposite. You have the same force but to the right with the opposite sign. But what's going on here? Yes.

AUDIENCE: So I can't answer this mathematically or even physically. But I just remember this very cool example that someone showed me once where they took a saddle shape in space, put a ball in the middle, and then you spin the saddle. And as soon as the ball tries to fall to one side, the saddle will spin over and pick it up. So I know the answer has to be, it is confining. But I can't explain why.

PROFESSOR: OK, Colin?

AUDIENCE: Could it possibly be like a threshold? Because if you consider spinning it sort of slowly, then the particle can just sort of roll off. Maybe there's a threshold.

PROFESSOR: Yeah, there is a threshold. I mean, if there's a confining potential for a long, long time, and then there's an anti-confining potential, and I just fall down the anti-confining potential, and I disappear to infinity, the confining part doesn't help me. So there will definitely be a threshold in time. We have to switch between confining and anti-confining sufficiently rapidly. Yes.

AUDIENCE: Does it have something to do with fictitious force [INAUDIBLE]?

PROFESSOR: No, but this would have been a possibility. Maybe there is a formalism. But the standard formalism does not use the rotating frame. Let me--

AUDIENCE: You could think of the classical analog where you have a bead on a wire, and you have a stable position at sort of a finite angle only at certain frequencies.

PROFESSOR: Yes, I'm not sure. Anyway, you all-- I think everybody here works with optics or knows what optics is. So why don't I discuss the following. If you take a positive lens and a negative lens with equal strength, what happens now to your laser beam? You send a laser beam through, it gets an inward bend by the convex lens, an outward bend by the concave lens. Is the net effect nothing, or is it focusing?

Well, it depends. What happens is the following. If the lenses are put in close proximity, the two bending actions on the rays just completely cancel. But now we can do the following. We can put the two lenses further apart.

And now the following happens. We have a laser beam which is bent inward. And you probably all know the lens maker's equation, that the further out the rays are, the more they have to be bent inward to come to a focus.

So what happens is those beams are now bent inward. But when they reach the other lens, they are now bent outward somewhat. I'm not very accurate in my drawing. But the angle by which they're bent inward is proportional to the separation here. The angle by which they're bent outward is proportional to the separation there. So the inward bend is stronger than the outward bend.

So therefore what I'm telling you is convex and concave, with a certain distance between the two, results in focusing. Now maybe you can think about it. What happens if I reverse convex and concave?

AUDIENCE: [INAUDIBLE]

PROFESSOR: Is it now still focusing, or is it the opposite of the example above? Now the first lens bends the beam outward, yes? And the second lens bends it inward. But since we reach the second lens further out, the inward kick is stronger than the outward bend

initially. So therefore, when you have positive and negative lenses of equal strength, no matter in what order you apply it, you always get an inward force. So this should actually explain--

AUDIENCE: As long as they're within the focal length. Could you imagine if I had positive and negative that was larger than the focal length away?

PROFESSOR: Yeah, everything within reason here-- there will be a stability diagram. But I want to first make it very plausible what happens. And just to address the question which I raised before, if I'm in a potential, and I get an equal amount of inward and outward force, shouldn't that average to 0? The question is how you average.

If I would keep the distance from the x as constant, I would say here as a function of-- let me call it x . The bending angle as a function of x is positive. The bending angle at the same x is negative. So if I would do the average at constant x , I would get 0.

But what I do is I do the average following the trajectory of the beam. So in other words, if I'm sitting on this potential, which switches from attractive to repulsive, if I would be here at my position, the average force is 0.

But when the potential is repulsive, in $1/2$ cycle I'm pushed outwards. And further out, the attractive force is now stronger than it was before. And therefore, if you average the force at my position, you will always find an inward force. But if you average the force at a fixed location, it averages out to 0.

So this is exactly how RF ion traps work. And well, the rest are just equations to describe that. So I think we have to make a choice. I could derive for you the equations, how you average over the rapid oscillations of the potential, and you get a harmonic oscillator potential, and this is how you trap ions. But it's pretty much showing to you how the concept I've just explained to you is reflected in equations. Or the second option we have that takes about 20 minutes is to explain to you how quantum computation is done with two ions, in particular how the famous Zoller-Cirac gate works. Do you have a clear preference?

AUDIENCE: Second one.

PROFESSOR: OK, fine, but there is one piece of mathematics you should understand. Really, and this is the only essence of the mathematical derivation. If I'm standing here, and I feel the alternating force, it's a cosine ωt , and it averages to 0. But if the force drives me, it shakes me. This is micromotion.

Like the beam here, if you go through many lasers, the beam goes up and down, up and down. This is micromotion. The micromotion is at cosine ωt . And the net attractive force comes because the force which oscillates at cosine ωt , and I'm oscillating at cosine ωt , then you get cosine squared ωt . And this gives you the net force which provides the net attractive force. And it's the cosine square term averaged over the micromotion, which is one half, which is responsible for the trapping potential in ion traps.

Anyway, you know the physics very well. You understand where the math comes from in that it's really just the product of the two. It's a rectification. It's an alternating force. But since the system itself, the particle itself, is driven at the micromotion, there's rectification going on, which leads to the cosine squared term. [? Timor, ?] yes?

AUDIENCE: Two quick questions. So I think ion trappers always complain about the micromotion. But in fact, they need it for trapping to work-- is pretty much what's going on here, right?

PROFESSOR: Yes, actually, the one interesting side of the mathematics is the micromotion is kinetic energy. The harmonic trapping potential is nothing else than the micromotion. In other words, if people show in an ion trap that you have a harmonic oscillator, the harmonic oscillator potential is high here. But mathematically, this harmonic oscillator potential is the kinetic energy of the micromotion.

In other words, if you look at the micromotion, it has a kinetic energy. The kinetic energy appears as the potential energy in this slow secular motion. So in other words, the potential energy for the trap is the kinetic energy of the micromotion. So

you can't avoid it. This is responsible for the restoring force.

AUDIENCE: Is it a really conservative trap, or can I have an out of phase component to the micromotion response? Is there like a sine part to the response?

PROFESSOR: Well, the secular motion is conservative. You're now asking--

AUDIENCE: I'm trying to think if there would be a reason why my response, my dipole response, would have a different [INAUDIBLE]. Is there a reason?

PROFESSOR: Well, to the best of my knowledge, everything is wonderfully conservative if you have a separation of timescale. If you want to sort of get something strange and non-conservative out of the micromotion, you have to violate the adiabatic separation of timescales. There is a possibility that if particles move quickly through the potential, they're not following the micromotion adiabatically. And then there is a collision where micromotion is released as collisional energy. But on the similar particle level, it's a pretty good approximation.

OK, we want to talk about quantum computation with ions. [INAUDIBLE] has a wonderful write-up on the wiki on the basic concept of quantum computation and how it applies to ions. But let me immediately jump to qubit operations. You've heard many, many times that a general quantum algorithm can be constructed out of single qubit rotations and two qubit gates.

Single qubit rotations are just oscillations on the Bloch sphere. If you excite a particle, move it around on the Bloch sphere, this is how you can rotate a single qubit from ground state to excited state, any possible superposition state. So the physics of single atoms, Rabi oscillation, rotation on the Bloch sphere, this is all you have to know about single qubit operations.

The interesting thing is how do we create two qubit operations? And two qubit operation would mean you want to do something to one ion here depending what the state of the other ion is. And this is much more subtle.

So what I want to explain to you now is the famous Cirac-Zoller gate, which was

really a pioneering paper which has thousands of citations now. Because it showed that quantum computation can be experimentally realized in a rather straightforward way. So it's a fundamental theorem of quantum information processing that if you have single qubit operations, and one to qubit gate, which is a CNOT gate, that this is a universal set of quantum gates, any arbitrary quantum algorithm can be constructed out of those gate operations.

So let me just remind you what the CNOT gate is. We have qubits ion one, ion two. They can be in 0 or 1. These can be two different hyperfine states. Or in the case of the calcium ion, it can be the ground state or metastable excited state-- just two different states of ion one, two different states of ion two.

The CNOT gate would now say the following. The first ion is the control bit. And whenever the control bit is 1, it flips the qubit, the second qubit. If the control bit is 0, it does nothing to the second ion. So therefore, the table which shows how the quantum states perform is 1, 1, 0, 0, 1, 1.

It just shows that I'm writing down the matrix in these spaces. It's a truth table 0, 0, 0, 1, 1, 0, 1, 1. And it shows that when you have 1, 1 at the input, you get 1, 0 at the output. If you have 1, 0 at the input, you get 1, 1 at the output.

Well, the two qubit gate I want to discuss, which can be realized in the simplest possible way with the Cirac-Zoller algorithm, is the controlled phase qubit, which has a truth table of 1, 1, 1, minus 1. And it means it simply changes the phase of the two qubit states when both qubits are 1.

Now, the CNOT gate can be constructed out of the controlled phase gate by having two single qubit rotations done with the Hadamard gate. The Hadamard gate is taking the bases 1 and 2 and creates 1 plus 2, 1 minus 2, the symmetric and anti-symmetric superposition state. So therefore, the CPHASE gate is as general as the CNOT gate. Because one can be used to construct the other.

So how do we create the Cirac-Zoller gate? So how do we get now the controlled phase gate out of trapped ions? So the effect is the following. Remember, if we

have two ions, we want to change the phase to minus 1 of the state only if the two ions are qubit up. So in other words, we want to do 180 degree rotation of the wave function, but only if both ions are up.

And the question is, how do we do that? How does one ion feel what the other ion is doing? And let me explain that now. The critical element in the Cirac-Zoller case is a center of mass motion. We want at some point-- say when this ion is spin up, we swap the quantum information from the ion to the center of mass motion.

In other words, when this ion is up, we use laser pulses. I will explain that to you. Then all the ions are moving together. And then, this ion, because they're all moving, knows that this ion was spin up before.

So the center of mass motion, the common dipole oscillation of all ions together in the trap, this is the data bus. So this is the key idea. But in order to explain it to you, I have to go in two steps. So let's just take that as a genetic level structure. We have an S state. And it has two different hyperfine states, spin up and spin down. And these are our two qubits.

There may be a P state, which we have used for laser cooling. But everything is in the ground state of motion using sideband cooling, which you've studied in your homework. And what I want to show you first is how we can create a phase gate. But I'm only focusing on ion one. So hold the thought. Ion two will come in a few moments, will come back.

So what we can do is the following. Each qubit state has vibrational structure-- ground state, excited state without center of mass motion. But then, as a function of vibrational quantum number for the center of mass motion, we have g_1 , g_2 , e_1 , e_2 , and such.

And what we can now do is, let's assume we have another state, a D state, which is an auxiliary state here. Fortunately, ions and atoms have many states. And we can drive a laser transition from S to D, which is only in resonance with the first vibrational excited state to the D state. So we use a [INAUDIBLE] sideband here,

which only drives the transition from e_1 to the non-vibrational state of the auxiliary state.

And that means automatically that this laser does nothing to e_0 , does nothing to g_0 , does nothing to g_1 . So if you have this Hilbert space of g_0, g_1, e_0, e_1 , the only transition which is effected here, or the only state which is effected here, is the e_1 state. OK, I'm talking about a single ion. I'm just talking about talking to a single ion right now. It's really simple.

So if we switch on this laser, what happens is we simply induce Rabi oscillations-- simple physics. The Rabi oscillation, however, after one Rabi period has changed the phase of the initial state by minus 1. If you carefully look what Rabi oscillations are, the wave function goes up, down, up, down. But we usually discuss what happens to the probability. And this is shown in green.

So when we are back in the original state, we have changed the phase of the wave function by minus 1. So let's just assume that we switch on this laser to this auxiliary state, and we have a-- well, it's a 2π pulse in population. But it makes a phase shift of π to the wave function, which is minus 1.

So what we have accomplished now is a phase gate not between two ions. But we have a phase gate which, in this Hilbert space, between the qubit up, qubit down-- I call it ground and excited state. But the other qubit, so to speak, is the center of mass motion 0, 1, 0, 1.

So this is the truth table of the phase gate. And so technically, we have now done a local phase gate with one ion where the two qubits are ground and excited state of the ion. This is our qubit we want to work with.

But then, the second qubit to realize the phase gate was 0, 1 of the center of mass vibration. But once we have realized that, we can now immediately take this idea and get a phase gate between two ions. Because the center of mass vibration is everywhere.

So therefore, what we can do is this was ion one. We can now-- and this is the

missing step I want to show you-- go to ion two. And we can swap the qubit information-- ground, excited, spin up, spin down. We can swap that with the center of mass motion.

So therefore, what we can do is, when the second ion is spin up, we do a swapping of the information from spin up, spin down to center of mass motion being 0 or 1. So the idea is the following. Ion two should control what ion one is doing.

So we first focus on ion two. We swap the qubit of ion two into the center of mass motion. Then, we do what I just said, the controlled phase gate for ion one with the center of mass motion. And after this is done, we write the center of mass motion back into ion two.

And then we have nothing to ion two. Ion two was just putting its information into the center of mass motion. It helped ion one to do the controlled phase gate. And then, we swapped the information back. So this is the sequence.

After preparing the center of mass, the first of two lasers cool everything in the center of mass motion into the 0 state. Then we swap from ion two the information to this phonon state center of mass vibration. Then we do the controlled phase gate on ion one. And then, we swap the information back.

So what I should show you-- you may have seen it. But what I should show you explicitly-- how do we swap the information from the internal state of ion two? How do we write this information into the data bus, which is the phonon, which is the center of mass vibration in the trap? Well, we again use [INAUDIBLE] side pulse. We do a pi pulse driving this transition.

And just look how beautifully the math works out. We have an arbitrary qubit state, a superposition of ground and excited. And the center of mass motion was laser cooled to the ground state. So let me just multiply that out. So we have coefficient a times g_0 plus coefficient b times e_0 .

If we now do a pi pulse, we transfer the population from e_0 into g_1 . And this is what I've done here. But now you see I can factor out the g. And what I have now is I

have the same qubit information with coefficients e and b , but no longer in the internal state of the ion but in the phonon state in the qubit, which is the center of mass motion.

So the idea here is you can have an arbitrary number of ions. And now you want to have a quantum gate. What you do with a quantum gate, you say, I want to now take this ion here and have a quantum gate with 10 ions further down.

So you first go to your first ion, transfer the information of this qubit into the data bus, into the center of mass motion. And now with a local operation here, you can do something to this ion conditionally on the motion, the center of mass motion. But the center of mass motion reflects the ion over there, in which state it was.

And after you have done the two qubit operation between those two ions, you take the center of mass motion, put it back into this ion. And at the end of the day what you have done is you have done a two qubit operation. You've changed the state of this ion. But all the other ions have remained the same. And with that, you can now, in a chain of ion, implement arbitrary two qubit operations.

I think I said it clearer in words here. It is in a little bit more formal way the steps you will do to the wave function. You initialize. You transfer the information to the data bus, you execute the controlled phase gate, and you put the information from the data bus back into the control ion. And now your data bus is free, is back again to the initial value and ready for the next operation.

I should mention that this was pioneering paper. Right now, this quantum gate is not the most popular one. Because it has severe requirements. One is you really have to cool to the absolute ground state. The center of mass motion has to be cooled to the absolute ground state. There are other quantum gates which do not require such extreme cooling.

And secondly, the center of mass motion, because all particles move together, is one of the slowest vibrational motions in the system. And if you want to scale up the system to more and more particles, the speed of the gate operation has to be

slower than the vibrational period of the center of mass motion. So these are two disadvantages.

But there are different gates, like geometric gates, like Sorensen gates. There are other possibilities to involve the same quantum logic with ions. But the one of Cirac-Zoller really stands out by its pedagogy. It's really a wonderful way to see how we take the information, put it in a data bus, do the controlled phase operation, and then proceed to the next two qubit gate.

OK, so that's the end of ions, the end of this class. Do you have any questions about the last topic? If not, let me say thanks. You were a fabulous class.

I really enjoyed the classroom atmosphere. I enjoyed all the questions, whether it was after class or during class. And, well, I'm happy that most of you I will see around for a number of years. All right.

[APPLAUSE]