One and Many Particle Paths in a Penning Trap Simulation

Viktor Braaten Gulbrandsen, Katchi Senthil Ganesh
(Dated: November 19, 2026)

The Penning trap is an important part of high-precision measurements of the mass and magnetic moment of single charges, and holds a charged particle by adding a homogeneous magnetic field along the axis of a static electric saddle, yet only has an analytic solution for a single particle in an ideal field. We first cover the analytic solution for a single particle, then simulate both single ions and multiple interacting ions using C++ by way of the numerical methods forward Euler and fourth-order Runge–Kutta (RK4). Against this baseline, a convergence rate of 4.00 for RK4 and the order approximations of error derived from the Taylor expansions of those numerical methods are confirmed, with forward Euler at 1.40 nearing its first order as the step shrinks. Further, we add a second ion, finding that the Coulomb force hands the motion along the axis back and forth between the two without loosening the trap. We then fill the trap with 100 ions to explore under which configurations the trap can lose its ions. This we do by stepping the angular frequency of a time-dependent perturbation to the applied potential from 0.2 to 2.48 MHz at three amplitudes.

Code is available at https://github.uio.no/viktorgu/fys3150, in the folder `project3/`.

I. INTRODUCTION

A major aim for particle physics has been developing experimental techniques to better understand the properties of matter at the level of its smallest constituents. A fundamental component of this is the development of techniques which measure to high precision the mass and magnetic moment of single charges. For a long time, such measurements were made on beams of particles passing once through the apparatus, as in Thomson's measurement of the charge-to-mass ratio of the electron [1]. In work culminating in Hans Dehmelt and Wolfgang Paul being awarded the Nobel Prize in Physics in the year 1989, however, a technique known as ion trapping made it possible to hold a single charged particle in place for long times [2]—with large collaborations at CERN funded by its member states, such as ALPHA [3], AEgIS [4] and BASE [5], standing in testament to the field this work opened.

The primary goal of ion trapping is holding a charged particle in place without disturbing its dynamics by it touching the walls of a container or colliding with other matter. Efforts before the ion trap were primarily hindered by two fundamental inconveniences of electrostatics. First, by Laplace's equation (1), a particle in empty space feels a potential whose curvatures along our three spatial axes sum to zero. Thus a static field can at best offer a saddle point. This observation, that any static field which confines a charge in one direction pushes it outwards in another, is commonly referred to as Earnshaw's theorem [6]. Secondly, charges of the same sign repel each other, so the more particles a trap holds, the stronger the confinement must be kept. This repulsion makes experiments where many particles are held at once inconvenient, as every particle disturbs the motion of all the others.

The Penning trap—named by Hans Dehmelt [2] after Frans Penning, who in the 1930s used a magnetic field to lengthen the paths of electrons in a gas discharge [7]—overcomes these inconveniences by using both the electric and magnetic components of the electromagnetic field. Imposing a static voltage on its electrodes, it maintains the saddle point as given by the electric potential but adds a homogeneous magnetic field along the axis of the saddle. Because the magnetic force acts at right angles to the velocity of particles within a magnetic field, it in effect bends the motion that would escape the trap into circles. Charges held in the trap then move in three superposed motions (one oscillatory along the axis and two circular in the plane perpendicular to it) [8].

All the while an idealized Penning trap, where a single particle moves in perfectly shaped fields, has an analytic solution, experimental actual traps do not. And the details of exactly how to numerically simulate a trap with many particles and fields varying in time are still of importance today. For instance, the group at NIST runs a Penning trap with hundreds of ions as a quantum simulator [9], and the BASE collaboration at CERN uses one to compare the antiproton with the proton [10].

The topic of this paper is therefore investigating numerical simulations of the Penning trap. The natural starting point in this investigation is the analytical expressions, which we derive in Section II. We then show the results of our simulations in C++, first the path of a single particle with the forward Euler and fourth-order Runge–Kutta (RK4) methods, measured against the analytical solution. From that, we introduce two particles interacting by the Coulomb force. Finally, we fill the trap with 100 particles and vary the voltage in time, to find the frequencies at which particles are lost. This scan runs without the Coulomb force between the particles, and we then repeat it around one of the frequencies found, with the force switched on.

II. METHODS

The Penning Trap

The Penning trap consists of three electrodes and a magnet in order to set the electric potential V and the magnetic field $\mathbf{B}$ in a way keeping a particle in an enclosed radius [11]. The components of a Penning trap are

- two metal electrodes above and below the centre, most commonly called caps, which repel the particle back towards the centre whenever it moves up or down the axis of the trap, the axial confinement,
- a metal electrode encapsulating the centre, called the ring, held at a voltage V_0 below the caps, whose surface together with the caps is shaped as a surface of constant V, so that the potential between them is the one derived below,
- a superconducting coil around the electrodes, the magnet, which sets the field $\mathbf{B}$ along the axis and holds the particle in the plane across it, the radial confinement.

Geometrically, the caps are two dishes facing each other across the centre, each a distance z_0 from it, and the ring is a hollow cylinder between them, its wall a distance r_0 from the centre—Figure 1 shows the three in cross-section, with the magnet as a larger cylinder around them.

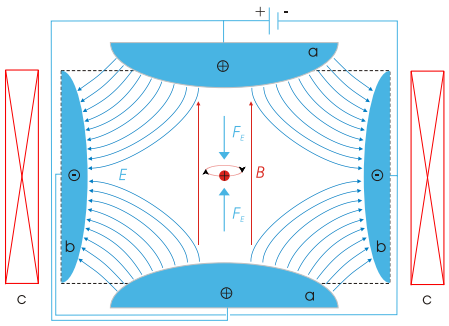

FIG. 1. Schematic of a Penning trap. The two end caps (a) are held at a voltage V_0 above the ring (b), the cylinder magnet (c) sets the magnetic field along the axis, and the red dot is a trapped positive particle. Illustration by Arian Kriesch, Wikimedia Commons, as used in the course text [12].

In order to understand the way the trap is able to hold a charge, let us consider first the electric field, and then later the magnetic.

The Electric Field of a Penning Trap

A good starting point is considering Gauss's law in the empty space between the electrodes,

$$\nabla \cdot \mathbf{E} = 0.$$

This is the field in the gap of the Penning trap resulting from sending electricity through the electrodes. From electrostatics, we know that this relates to the gradient of the potential, because a charge is pushed from higher to lower potential energy,

$$\mathbf{E} = -\nabla V.$$

From this, Laplace's equation is derived by inserting the gradient into Gauss's law, as the divergence of a gradient is the sum of the second derivatives,

$$\nabla^2 V = \frac{\partial^2 V}{\partial x^2} + \frac{\partial^2 V}{\partial y^2} + \frac{\partial^2 V}{\partial z^2} = 0. \tag{1}$$

This shows us more directly what we previously referenced as Earnshaw's theorem [6]—stating that a charge cannot be held at rest by static electric forces alone. Ideally, one would be able to get a stable set point where the potential curves upwards along all three axes (so that the field, $\mathbf{E} = -\nabla V$, points back towards the centre from every side). That is, a bowl where a charge "rolls" into the centre from every side.

Instead, by Laplace's equation the three curvatures sum to zero, so along at least one axis the field points away from the centre. This could also be read from the standpoint of Gauss's law, with the gap of the trap being a vacuum so that there is no sink at the centre for a field pointing in from every side to end on [13]. This creates the geometry of a saddle—a basic mathematical object studied throughout various fields—where the lowest point along one axis (in our trap the z-axis, between the caps) is the highest point along another. This creates a stationary point that is not a minimum in optimization, or in the current case a situation where a charge rolling back towards the centre along the axis of the caps rolls outwards towards the ring in the plane.

The exact shape of this saddle, where the low point along the axis is the high point in the plane, is set by the electrodes. Thus one can order them in a way so the saddle is symmetric around the z axis. We can present this as a somewhat coarse symmetry condition: V stays the same when any one coordinate is mirrored,

$$V(x,y,z) = V(-x,y,z) = V(x,-y,z) = V(x,y,-z),$$

so that the trap looks the same from above and below, and similarly so that when x and y are swapped,

$$V(x,y,z) = V(y,x,z),$$

it looks the same from every side. The expression

$$V(x,y,z) = az^2 + b\left(x^2 + y^2\right)$$

is a second-order polynomial which maintains these conditions. Check:

$$\begin{aligned} V(-x,y,z) &= az^2 + b\left((-x)^2 + y^2\right), && \text{flip } x,\\ V(x,-y,z) &= az^2 + b\left(x^2 + (-y)^2\right), && \text{flip } y,\\ V(x,y,-z) &= a(-z)^2 + b\left(x^2 + y^2\right), && \text{flip } z,\\ V(y,x,z) &= az^2 + b\left(y^2 + x^2\right), && \text{swap } x, y. \end{aligned}$$

By imposing Laplace's equation (setting the sum of the second partial derivatives to zero), we get:

$$\frac{\partial^2 V}{\partial x^2} + \frac{\partial^2 V}{\partial y^2} + \frac{\partial^2 V}{\partial z^2} = 2b + 2b + 2a = 0,$$

and hence $b = -a/2$. Following the course book in writing $a = V_0/d^2$ [12], this yields the expression

$$V(x,y,z) = \frac{V_0}{2d^2}\left(2z^2 - x^2 - y^2\right), \tag{2}$$

where V_0 is the voltage between caps and ring, and d is the characteristic dimension of the trap, $d^2 = z_0^2 + r_0^2/2$.

Eq. (2) allows a simple view of how the electric field functions in a Penning trap. The z axis is the only direction where V has a positive coefficient (meaning that the electric field, being minus the slope of V, points back towards the centre there), while x and y carry a negative one (so that the field points outwards in the plane). This means that a charge is held along z and lost in the plane. Let us now consider how the magnetic field in the Penning trap relates to this.

The Magnetic Field of a Penning Trap

The magnetic component of the Penning trap is induced by the superconducting coil around the electrodes, the magnet (c) in Figure 1, which sets a homogeneous field along the z axis. This can be described as

$$\mathbf{B} = B_0\hat{\mathbf{e}}_z = (0,0,B_0), \qquad B_0 > 0,$$

where B_0 is the strength of the field and $\hat{\mathbf{e}}_z$ is the unit vector along the z axis.

This is enough to begin describing the dynamics of a point charge q with mass m in the trap. Recall by Newton's second law of motion

$$m\ddot{\mathbf{r}} = \mathbf{F}, \tag{3}$$

that the total force on a body of mass m is m times its acceleration, the second time derivative of its position $\mathbf{r}(t)$. It is this position $\mathbf{r}(t)$ we focus on through the rest of the report. Let us first, however, see that the force on such a point charge in the trap itself is made up of two parts, one from the electric field and one from the magnetic, which can be described by the Lorentz force

$$\mathbf{F} = q\mathbf{E} + q\dot{\mathbf{r}} \times \mathbf{B}.$$

To make this expression precise for our system, let us first consider the first term of the sum (i.e. $q\mathbf{E}$). With $\mathbf{E} = -\nabla V$ and V as given in Eq. (2), this is

$$q\mathbf{E} = -q\nabla V = -\frac{qV_0}{2d^2}(-2x,-2y,4z) = \frac{qV_0}{d^2}(x,y,-2z),$$

and the second term, the cross product, is

$$q\dot{\mathbf{r}} \times \mathbf{B} = q(\dot{x},\dot{y},\dot{z}) \times (0,0,B_0) = qB_0(\dot{y},\ -\dot{x},\ 0).$$

Adding the two, the total force is

$$\mathbf{F} = \frac{qV_0}{d^2}(x,\ y,\ -2z) + qB_0(\dot{y},\ -\dot{x},\ 0).$$

Placing this back into Newton's second law, Eq. (3), yields

$$m\ddot{\mathbf{r}} = \frac{qV_0}{d^2}(x,\ y,\ -2z) + qB_0(\dot{y},\ -\dot{x},\ 0).$$

This equation combines a lot of the conditions of the Penning trap, both the electric field as constrained by the electrodes and Laplace's equation, and the magnetic field as set by the coil, all relating it back to the position $\mathbf{r}(t)$ of the charge by Newton's second law. It also allows us to write the motion as a set of linear differential equations, by dividing by m and taking one component at a time,

$$\begin{aligned} \ddot{x} &= \frac{qV_0}{md^2}x + \frac{qB_0}{m}\dot{y},\\ \ddot{y} &= \frac{qV_0}{md^2}y - \frac{qB_0}{m}\dot{x},\\ \ddot{z} &= -\frac{2qV_0}{md^2}z. \end{aligned}$$

The three lines share two combinations of the constants, which we follow the course text in naming [12]

$$\omega_0 \equiv \frac{qB_0}{m}, \qquad \omega_z^2 \equiv \frac{2qV_0}{md^2}.$$

The first is the rate at which the magnetic field turns the velocity of the charge, known as the cyclotron frequency, and the second is the frequency of the oscillation along the axis, the axial frequency [11]. With these, the equations of motion read

$$\ddot{x} - \omega_0\dot{y} - \tfrac{1}{2}\omega_z^2 x = 0, \tag{4}$$

$$\ddot{y} + \omega_0\dot{x} - \tfrac{1}{2}\omega_z^2 y = 0, \tag{5}$$

$$\ddot{z} + \omega_z^2 z = 0. \tag{6}$$

Throughout, we take $q > 0$ and $V_0 > 0$, so that both frequencies are real.

Solving the Single-Particle Trap

We wish to solve the set of linear differential equations defined above. Out of these three, the z equation,

Eq. (6), is the simplest because it depends only on itself (i.e. is decoupled). Because the second derivative of z is minus a constant times z itself, it has the general form of a harmonic oscillator—recall both $\cos(\omega_z t)$ and $\sin(\omega_z t)$ return themselves times $-\omega_z^2$ when differentiated twice. A general solution for that dynamics is

$$z(t) = C_1 \cos(\omega_z t) + C_2 \sin(\omega_z t), \tag{7}$$

where the constants C_1 and C_2 are fixed by the initial position and velocity of the particle. The particle hence oscillates along the z-axis with angular frequency ω_z.

The equations for the x and y axes are coupled. This is reasonable because we have already seen the magnetic force turn the velocity by 90°, so that $\dot{y}$ enters the equation for x and, symmetrically, $\dot{x}$ the equation for y. The flip side of this is nonetheless that they allow themselves to be expressed as one equation. The complex plane allows us to write this turn in a single operation, as multiplying by i is equivalent to turning 90°. For example, in a general complex function

$$f(t) = a(t) + i\, b(t),$$

multiplying by i gives $if = -b + ia$, which is the point (a, b) turned by 90° to $(-b, a)$. We can thus pick

$$f(t) = x(t) + iy(t),$$

with $\dot{f} = \dot{x} + i\dot{y}$ and $\ddot{f} = \ddot{x} + i\ddot{y}$. Adding Eq. (4) to i times Eq. (5) gives

$$(\ddot{x} + i\ddot{y}) - \omega_0 \dot{y} + i\omega_0 \dot{x} - \tfrac{1}{2}\omega_z^2 (x + iy) = 0,$$

and since

$$-\omega_0 \dot{y} + i\omega_0 \dot{x} = i\omega_0 (\dot{x} + i\dot{y}),$$

this collects into the single equation

$$\ddot{f} + i\omega_0 \dot{f} - \tfrac{1}{2}\omega_z^2 f = 0. \tag{8}$$

As a complex number vanishes only when both its real and imaginary parts vanish, the real part of Eq. (8) returns Eq. (4) and the imaginary part Eq. (5), so the two descriptions are equivalent.

To find the path in the plane, we use the general solution of Eq. (8) given in the course book [12],

$$f(t) = A_+ e^{-i(\omega_+ t + \phi_+)} + A_- e^{-i(\omega_- t + \phi_-)}, \tag{9}$$
$$\omega_\pm = \frac{\omega_0 \pm \sqrt{\omega_0^2 - 2\omega_z^2}}{2},$$

where the amplitudes $A_\pm > 0$ and the phases $\phi_\pm$ are constants set by the initial conditions, and the path in the plane is $x(t) = \mathrm{Re}\, f(t)$ and $y(t) = \mathrm{Im}\, f(t)$. The motion in the plane is thus two circular motions added together, a fast one with frequency ω_+ and a slow one with frequency ω_-.

Whether this path stays bounded is decided by the two frequencies. Writing a frequency as $\omega = \alpha + i\beta$, each term splits into

$$e^{-i\omega t} = e^{-i\alpha t}\, e^{\beta t},$$

where the real part α sets the direction of the term and the imaginary part β its length. A bounded path thus requires both frequencies to be real, that is, the square root to be real, and the two frequencies distinct, since a double root adds a term growing linearly in time,

$$\omega_0^2 > 2\omega_z^2.$$

Inserting the definitions of the two frequencies and dividing by $q/m > 0$, this reads

$$\frac{q}{m} > \frac{4V_0}{B_0^2 d^2}, \tag{10}$$

so a given trap holds only particles whose charge-to-mass ratio is large enough.

When this holds, the distance of the particle from the z-axis is

$$\sqrt{x^2 + y^2} = |f|, \quad |f|^2 = A_+^2 + A_-^2 + 2A_+ A_- \cos\theta(t),$$

with $\theta(t) = (\omega_+ - \omega_-)t + (\phi_+ - \phi_-)$. As $\theta(t)$ grows linearly in time, the cosine takes every value between -1 and 1, and the largest and smallest distances from the z-axis are

$$R_+ = A_+ + A_-, \qquad R_- = |A_+ - A_-|,$$

reached when the cosine equals 1 and -1, respectively. The particle thus oscillates along the z-axis with frequency ω_z and circles in the plane between the radii R_- and R_+.

To compare with the numerical solution later, we need this path for one set of initial conditions. We use those of the course book [12], $x(0) = x_0$, $y(0) = 0$, $z(0) = z_0$ and $\dot{x}(0) = 0$, $\dot{y}(0) = v_0$, $\dot{z}(0) = 0$, with $x_0, v_0 > 0$. The z motion then has $C_1 = z_0$ and $C_2 = 0$,

$$z(t) = z_0 \cos(\omega_z t), \tag{11}$$

and the plane motion has

$$A_+ = \frac{v_0 + \omega_- x_0}{\omega_+ - \omega_-}, \qquad A_- = \frac{v_0 + \omega_+ x_0}{\omega_+ - \omega_-}, \tag{12}$$
$$\phi_+ = \pi, \qquad \phi_- = 0.$$

This is the exact solution the numerical methods are checked against in the Results.

Solving the Many-Particle Trap

Many traps hold more than one particle—meaning that any expression claiming to describe a real trap needs to consider the linear dynamics of a single particle in the trap all the while folding those in with the interactions

between the particles. Thus describing the dynamics of one particular particle relies on the other particles in the system. This fundamentally alters the problem class the many-particle trap belongs to, as it is now nonlinear [13]. Because a closed form is in general not feasible to find, we progress by introducing numerical methods.

Consider such a trap $P = \{1, \ldots, n\}$ with n particles of charges q_j, masses m_j and positions $\mathbf{r}_j$, where j enumerates the particles pushing and i the particle being pushed. The Coulomb force describes their push on each other, through the electric field at particle i from all the others [12],

$$\mathbf{E}_{\text{int}}(\mathbf{r}_i) = k_e \sum_{j \in P,\, j \neq i} q_j \frac{\mathbf{r}_i - \mathbf{r}_j}{|\mathbf{r}_i - \mathbf{r}_j|^3}, \tag{13}$$

where k_e is Coulomb's constant, with Newton's second law for particle $i \in P$ reading

$$m_i \ddot{\mathbf{r}}_i = q_i \mathbf{E}(\mathbf{r}_i) + q_i \dot{\mathbf{r}}_i \times \mathbf{B} + q_i \mathbf{E}_{\text{int}}(\mathbf{r}_i), \tag{14}$$

and the equations of motion are coupled across all particles through the last term. Each one is of second order in time. We split it into two equations of first order, one for the position and one for the velocity,

$$\dot{\mathbf{r}}_i = \mathbf{v}_i, \qquad \dot{\mathbf{v}}_i = \frac{\mathbf{F}_i}{m_i}, \tag{15}$$

with $\mathbf{F}_i$ the total force on particle i, which we integrate numerically in the next section.

Before considering how that integration is done, let us note that a consequence of our system being an instance of the more general problem class of nonlinear dynamical systems is bifurcation behaviour. A famous example of bifurcation generally is the van der Pol oscillator [14], where a small change in a parameter turns an oscillation that dies out into one that keeps itself going, and similarly the Lorenz system [15], where a steady state turns into a chaotic one. In our case, the trap will under a periodically varying voltage show resonances. These show up by altering the voltage V_0 between the caps and the ring—the parameter regulating the electric field as given in Eq. (2)—as

$$V_0 \to V_0 \left(1 + f \cos(\omega_V t)\right), \tag{16}$$

where f is the amplitude of the variation, as a fraction of V_0, and ω_V its angular frequency [12]. In the single particle case, the particle moves at three frequencies of its own, the axial frequency ω_z and the two plane frequencies $\omega_\pm$ of Eq. (9), and a voltage that swings in step with one of these adds to that motion on every cycle until the particle is lost. With more than one particle, each moves in the field of the trap plus the field of the rest, so these frequencies are potentially shifted, the size of which we can only find in the results of the numerical simulation, as we will now come to.

Numerical Methods

The equations of the last section give, for any position and velocity, the force on each particle and with it how fast its position and velocity change. In general, we can say that a numerical method for solving such a differential equation is a method which progresses by solving for a short time step h at a time, from a state held in memory, the positions and velocities of all particles at the current moment, to the state one step later. The two methods we use differ in how much of the step they look at, where forward Euler uses only the rate of change at its start, at a lower level of granularity, the RK4 method samples the rate of change four times across the step. We write $\mathbf{y} = (\mathbf{r}, \mathbf{v})$ for the state and $\mathbf{g}(t, \mathbf{y}) = (\mathbf{v}, \mathbf{F}/m)$ for its rate of change, Eq. (15), so that each method is a rule for going from $\mathbf{y}_n$ at time t_n to $\mathbf{y}_{n+1}$ at $t_n + h$.

Forward Euler

The simplest rule is to follow the slope at the start of the step for the whole step,

$$\mathbf{y}_{n+1} = \mathbf{y}_n + h\, \mathbf{g}(t_n, \mathbf{y}_n),$$

which for the position and velocity reads

$$\mathbf{r}_{n+1} = \mathbf{r}_n + h\, \mathbf{v}_n, \qquad \mathbf{v}_{n+1} = \mathbf{v}_n + h\, \frac{\mathbf{F}(\mathbf{r}_n, \mathbf{v}_n, t_n)}{m}. \tag{17}$$

How wrong this is per step follows from the Taylor expansion of the true path,

$$\mathbf{y}(t_n + h) = \mathbf{y}(t_n) + h\, \mathbf{g} + \frac{h^2}{2}\, \dot{\mathbf{g}} + \ldots,$$

of which forward Euler keeps the first two terms. The error of one step, the local error, is thus the first term dropped, and over a run of T/h steps these add up to the global error,

$$\text{local error} = \mathcal{O}(h^2),$$
$$\text{global error} = \frac{T}{h}\, \mathcal{O}(h^2) = \mathcal{O}(h),$$

so halving h halves the error of the run [16].

Algorithm 1 denotes one such step for all n particles at once.

Algorithm 1 One forward Euler step for n particles

Require: positions $\mathbf{r}_i$ and velocities $\mathbf{v}_i$ for $i = 1, \ldots, n$, time t, step h
  for $i = 1, \ldots, n$ **do**
    $\mathbf{F}_i \leftarrow$ force on particle i at $(\mathbf{r}_i, \mathbf{v}_i, t)$
  for $i = 1, \ldots, n$ **do**
    $\mathbf{r}_i \leftarrow \mathbf{r}_i + h\, \mathbf{v}_i$
    $\mathbf{v}_i \leftarrow \mathbf{v}_i + h\, \mathbf{F}_i / m_i$
  return $\mathbf{r}_i$, $\mathbf{v}_i$, $t + h$

Fourth-Order Runge–Kutta

Instead of one slope at the start, RK4 evaluates the slope four times inside the step, at the start, twice at the midpoint and at the end, and takes their weighted average,

$$\begin{aligned}\mathbf{k}_1 &= \mathbf{g}(t_n, \mathbf{y}_n),\\ \mathbf{k}_2 &= \mathbf{g}(t_n + \tfrac{h}{2}, \mathbf{y}_n + \tfrac{h}{2}\mathbf{k}_1),\\ \mathbf{k}_3 &= \mathbf{g}(t_n + \tfrac{h}{2}, \mathbf{y}_n + \tfrac{h}{2}\mathbf{k}_2),\\ \mathbf{k}_4 &= \mathbf{g}(t_n + h, \mathbf{y}_n + h\mathbf{k}_3),\\ \mathbf{y}_{n+1} &= \mathbf{y}_n + \tfrac{h}{6}\left(\mathbf{k}_1 + 2\mathbf{k}_2 + 2\mathbf{k}_3 + \mathbf{k}_4\right). \end{aligned} \tag{18}$$

The weights are chosen so that this average matches the Taylor expansion of the true path through the h^4 term, so that

$$\begin{aligned}\text{local error} &= \mathcal{O}(h^5),\\ \text{global error} &= \frac{T}{h}\,\mathcal{O}(h^5) = \mathcal{O}(h^4),\end{aligned}$$

and halving h divides the error of the run by sixteen [16]. Each slope is evaluated at its own time, which matters once the force depends on t, and with several particles every particle is moved to the same intermediate stage before the next slope is computed, since the force on each depends on where all the others are. One step is Algorithm 2.

Algorithm 2 One RK4 step for n particles

Require: state $\mathbf{y} = (\mathbf{r}_i, \mathbf{v}_i)$ for $i = 1, \ldots, n$, time t, step h
  $\mathbf{k}_1 \leftarrow \mathbf{g}(t, \mathbf{y})$
  $\mathbf{k}_2 \leftarrow \mathbf{g}(t + h/2,\ \mathbf{y} + \frac{h}{2}\mathbf{k}_1)$
  $\mathbf{k}_3 \leftarrow \mathbf{g}(t + h/2,\ \mathbf{y} + \frac{h}{2}\mathbf{k}_2)$
  $\mathbf{k}_4 \leftarrow \mathbf{g}(t + h,\ \mathbf{y} + h\mathbf{k}_3)$
  $\mathbf{y} \leftarrow \mathbf{y} + \frac{h}{6}\left(\mathbf{k}_1 + 2\mathbf{k}_2 + 2\mathbf{k}_3 + \mathbf{k}_4\right)$
  return $\mathbf{y}$, $t + h$

  function G(t, $\mathbf{y}$)
    for $i = 1, \ldots, n$ **do**
      $\mathbf{F}_i \leftarrow$ force on particle i at $(\mathbf{r}_i, \mathbf{v}_i, t)$
    return $(\mathbf{v}_i,\ \mathbf{F}_i/m_i)$ for $i = 1, \ldots, n$

To assess the relation between a numerical path and the analytical path, we run a single particle for $T = 50\,\mu$s from the initial conditions of the course book and compare its computed position at every step with the exact solution, through the relative error

$$\frac{|\mathbf{r}_{i,\text{exact}} - \mathbf{r}_i|}{|\mathbf{r}_{i,\text{exact}}|}. \tag{19}$$

We do this four times for each method, with $n = 4000$, 8000, 16000 and 32000 steps, so that the step size $h = T/n$ halves from one run to the next and the order of the method can be read off how fast the error falls.

To search for resonances we fill the trap with 100 singly charged calcium ions, their positions and velocities drawn from normal distributions of width $0.1\,d$, and run for $500\,\mu$s under the drive of Eq. (16). At the end we record the fraction of ions still inside the trap, $|\mathbf{r}| < d$, both fields having been set to zero outside. The scan covers ω_V from 0.2 to 2.48 MHz in steps of 0.02 MHz at each of the amplitudes $f = 0.1$, 0.4 and 0.7. It is run without the Coulomb interaction, which makes it fast. One resonance found this way is then scanned in finer steps, with and without the interaction, three times each from different random starts.

All quantities are given in micrometres, microseconds, atomic mass units and elementary charges [12]. In these units $k_e = 1.38935333 \times 10^5$, and the trap has $B_0 = 1.00\,\text{T} = 96.5$, $V_0 = 25.0\,\text{mV} = 2.41 \times 10^6$ and $d = 500$, so that $V_0/d^2 = 9.65$.

Tools

The code is written in C++20 with the Armadillo library for vectors and matrices [17], and built with make. The figures and the notebooks, one per problem, are made in Python with numpy [18], pandas and matplotlib [19], in an environment managed by uv.

Large language model based tools are transforming workflows in both industry and academia. At the time of writing, the sentiment towards their adoption in coding is polarized. Some companies urge their engineers to rely on coding agents, others are wary of ending up with code no engineer has read, and universities are split the same way. We believe the utility of a tool, be it the compass in geometry or the calculator in arithmetic, is to be measured by the goal it serves and the outcome it is used for. We wish 1) to get the full educational takeaway, all the while 2) producing sound, good and reproducible code. We thus set up the folder structure shown in Figure 2 and use Claude Code [20].

To do our share against the flood of text no one has either read or written, however, every word of this report, every comment, and most of the code are written by us. Claude Code was used to look things up, to check derivations and code against the sources in Figure 2, and to draft what we then rewrote.

```
fys3150/
├── CLAUDE.md          encourages grep search and validation
├── notes/
│   ├── course_repo/   course material [12]
│   └── docs/
│       ├── index.md   information for grep
│       ├── standard/  C++ reference
│       └── armadillo/ docs [17], one file per function
├── project3/          the code
└── reports/           LaTeX code
```

FIG. 2. The folder structure behind our use of Claude Code. Only `project3/` is part of the graded repository.

III. RESULTS AND DISCUSSION

A Single Particle

Figure 3 shows the first test of our implementation, using a single Ca^+ ion simulation compared to the analytical solution (Eq. (11)) in the default trap, run for 50 μs with RK4 at 8000 steps, with the initial conditions from the course book (Table I). At this level, the two results cannot be told apart.

TABLE I. Initial conditions of the two ions [12], and the frequencies of the default trap for Ca^+.

	position [µm]	velocity [µm/µs]
ion 1	$(20, 0, 20)$	$(0, 25, 0)$
ion 2	$(25, 25, 0)$	$(0, 40, 5)$
frequencies [MHz]	$\omega_z = 0.694$	$\omega_0 = 2.41$
	$\omega_+ = 2.30$	$\omega_- = 0.105$

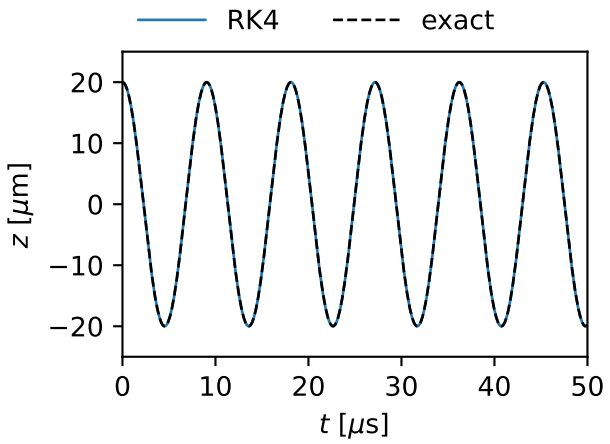

FIG. 3. RK4 for one Ca^+ ion follows the axial motion. $z(t)$ over 50 µs with 8000 steps (solid) against the analytic solution $z_0 \cos(\omega_z t)$ (dashed). The initial conditions are read from Table I. This also checks the axial frequency, because the time between every second zero crossing gives a period of 9.05 µs, equal to $2\pi/\omega_z$ for this trap and ion.

Figure 4 shows, in more detail, the discrepancy between the RK4 path and the analytic solution by computing the relative error (Eq. (19)) for every step, over four step counts. This confirms what we expected from the error orders of Section II, where each doubling of n lowers the curve by a factor of $2^4 = 16$. The figure also expresses oscillations of a period of about 2.9 μs. This is the time the fast circular motion of Section II takes to lap the slow one,

$$\frac{2\pi}{\omega_+ - \omega_-} = 2.86\,\mu\text{s}.$$

Each time the two circles point opposite ways, the ion comes closest to the axis, and the relative error, which divides by the distance to the centre, peaks there, while the error itself grows smoothly with time.

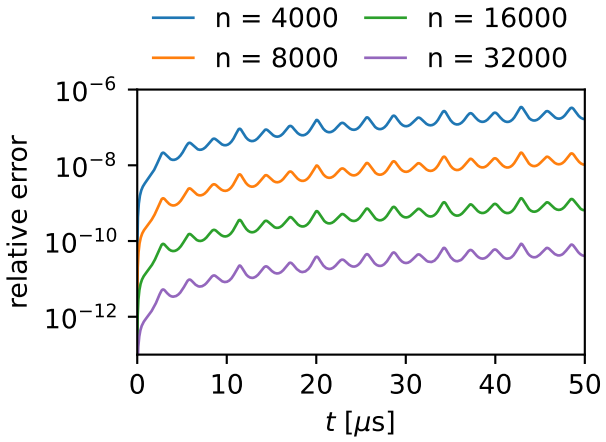

FIG. 4. Relative error on a log axis for RK4 against time for $n = 4000$, 8000, 16000 and 32000 steps over 50 µs shows RK4 falls by $2^4 = 16$ for each doubling of the step count n. On the log axis, this factor of 16 is a gap of a little over one power of ten between neighbouring curves. After a steep start, every curve rises slowly and stays below 10^{-6}. The small oscillations on each curve have a period of about 2.9 µs.

Figure 5 shows the relative error for forward Euler with the same step size, instead of RK4. Because forward Euler is first order, a doubling of n should result in the curve falling by a factor $2^1 = 2$. The ratio between the largest errors with n and $2n$ steps comes to be 3.27 for the first doubling. For the second, 2.50, and 2.23 for the third. This suggests that forward Euler only reaches its first order as the step shrinks. The error of this method is, as expected, far larger than RK4. With 4000 steps, the error passes 1 at $t = 33.9\,\mu$s, and at every peak after that the computed ion is further from its exact position than that position is from the centre of the trap. At the end of the run, the error is 52.5 μm, against 2.9 μm with 32000 steps. The oscillations of Figure 4 appear here too, with the same period, so they come from the path of the ion and not from the method.

Both the RK4 and the forward Euler errors lend themselves to analysis by an estimate of their convergence rate. This can be written as

$$r_{\text{err}} = \frac{1}{3}\sum_{k=2}^{4} \frac{\log\left(\Delta_{\text{max},k}/\Delta_{\text{max},k-1}\right)}{\log\left(h_k/h_{k-1}\right)},$$

with $\Delta_{\text{max},k}$ the largest distance between the computed and the exact position over run k, and h_k its step size, for both the Euler and RK4 methods [12]. This makes sense as each term measures how fast the largest error falls when the step is halved, and the mean over the three pairs

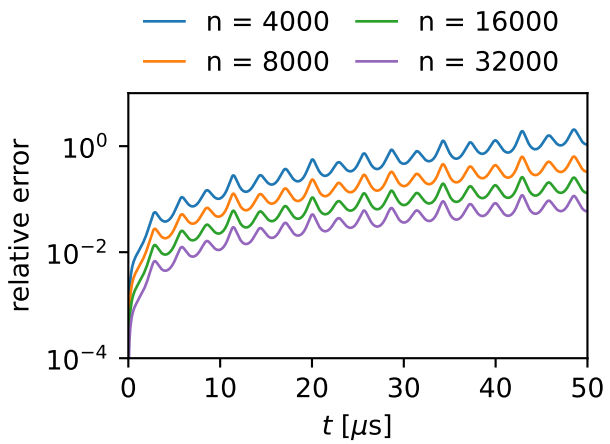

FIG. 5. Relative error on a log plot in the position of a single Ca^+ ion over time for forward Euler ($n = 4000$, 8000, 16000 and 32000 steps over 50 µs). Each doubling of n lowers the curve by more than $2^1 = 2$. With 4000 steps, the relative error passes 1 and peaks at 2.06. The small oscillations on each curve have the same period of about 2.9 µs as in Figure 4.

of neighbouring runs is then the order of the method. Table II shows the largest errors of each run and the resulting rates, 4.00 for RK4 and 1.40 for forward Euler. This is in accordance with the results shown thus far: RK4 at its order of four, and forward Euler above its order of one, nearing it only as the step shrinks.

TABLE II. Largest error for forward Euler and RK4 with $n = 4000$, 8000, 16000 and 32000 steps, $\Delta_{\max}$, in the position of a single Ca^+ ion over 50 µs, and the convergence rate r_{err} of each method.

n	forward Euler [µm]	RK4 [µm]
4000	5.25×10^{1}	8.12×10^{-6}
8000	1.60×10^{1}	5.07×10^{-7}
16000	6.40×10^{0}	3.17×10^{-8}
32000	2.88×10^{0}	1.98×10^{-9}
r_{err}	1.40	4.00

Two Particles

Moving now to the many-particle trap of Eq. (14), the introduction of the Coulomb force couples the particles. Figure 6 shows the paths in the xy plane of a simulation that introduces a second ion with initial conditions read from Table I and runs this system with and without the Coulomb force for 50 μs with RK4 at 16000 steps. We see that without the force each ion moves as a single ion would, on the two circles of Eq. (9). With the force the two push each other off these circles, and the first ion reaches 69 μm from the axis, against 45 μm without.

Figure 7 plots the phase space of each ion—that is, its position against its velocity along one direction, so that each point is the state of the ion at one moment, with x against v_x above and z against v_z below. A closed path in this picture means that the ion comes back to the same state. It thus repeats its motion. Without the force, this is what we expect along the axis. Each ion is then the harmonic oscillator of Eq. (7), and nothing else acts on its swing.

With the force, the two ions push on each other along z whenever they are at different heights. Both also swing at the same axial frequency. We read the trade in panel (c) as the behaviour of two equal oscillators joined by a weak coupling, like two pendulums of equal length joined by a spring. Such a pair hands its swing back and forth. Only a small part of the swing goes into the plane, as the combined energy of the two swings changes by only about a quarter.

In the plane, this push between the ions is harder to see. The motion there is two circles added together, and so never a single closed path to begin with. That the force only moves motion around is also reasonable physically. As Figure 6 showed, the centre of mass is the same with and without the force, and the total energy stays constant in both runs. This means that the force can move motion between the ions and between directions, but cannot add or remove any.

Figure 8 puts the motion in the plane of Figure 6 and the motion along the axis of Figure 7 together in one picture. Without the force, the second ion stays close to the middle plane, as its swing along the axis is small from the start. With the force, it is lifted to the same height as the first ion. This is the trade of Figure 7(c), now seen in space. Even so, the force reshapes the paths without loosening the trap. Neither ion comes further than about 75 μm from the centre, well inside the distance $d = 500\,\mu$m at which the trap lets go. Two ions are thus still held, with their motion only moved around between them. With many ions, however, each one moves in the field of the trap plus the field of all the others, so the frequencies of Section II may shift, as Methods noted. Whether such a shift shows, and at which frequencies of the voltage the trap can be made to lose its ions at all, is the question of the next part.

Resonances

An important part of the results concerns under which configurations the trap actually can lose its ions. Let us now turn to a configuration of the trap filled with 100 randomly initialized Ca^+ particles, where we add a time-dependent perturbation to the applied potential, as in Eq. (16). To set a baseline we switch off the Coulomb interactions first, so that the particles move on their own at the frequencies previously mapped.

The result of such a perturbation is that the previous potential holding the particles no longer has any constant

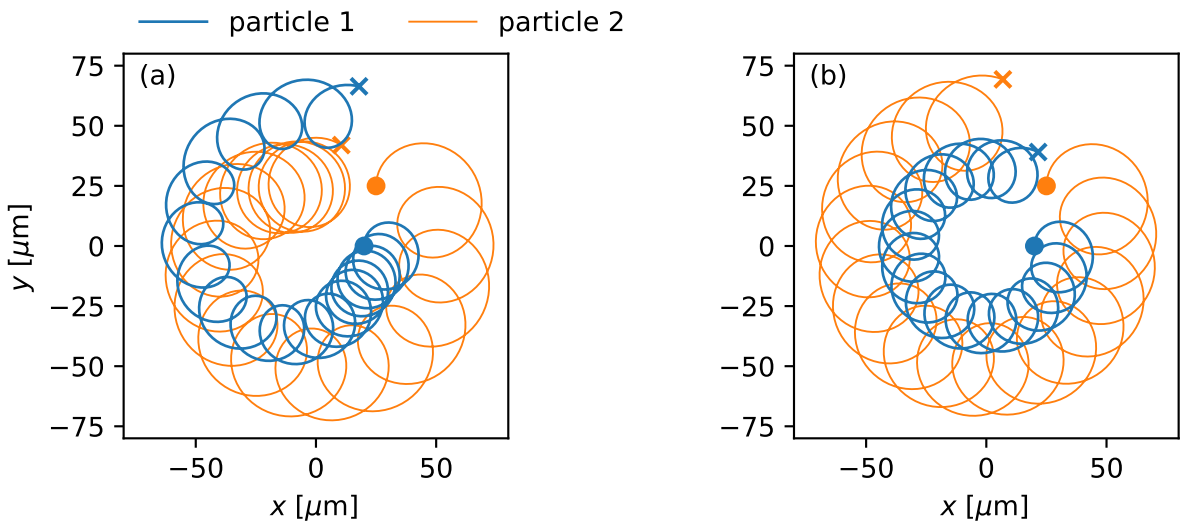

FIG. 6. Two Ca^+ ions shown over 50 µs with the Coulomb force (a) and without it (b). The dot marks where the particles begin and the cross where they end in the simulation. Because the interactions are the only difference, we see in panel (b) the two ions replicating the same two-circle motion as if they were isolated. They are each a slow circle around the axis with fast loops on top, and each ion keeps its own energy throughout. In panel (a) the ions repel each other and thus stay at a minimal distance of 23 µm, opposite to 22 µm without. The interacting ions show energy passage, where for instance the first moves out to 69 µm from the axis, different to 45 µm without. The centre of mass, however, seems stable across the conditions. This is expected since the Coulomb force is symmetric and equal in a way where the forces on the two ions cancel in their sum, leaving the centre of mass to move as a single ion would in the trap. The total energy stays constant throughout the run.

strength, and thus oscillates like a cosine around V_0 in time. Thus we have an oscillation with amplitude f, in ratio to V_0, and an angular frequency ω_V. Figure 9 shows the fraction of particles still trapped after 500 μs stepping across the applied angular frequency from 0.2 to 2.48 MHz for three amplitudes. These curves cover most of the range at a fraction of one, with every particle still trapped at no fewer than 64 of the 115 frequencies for each amplitude. The loss of particles is thus confined to a few dips, which widen as the amplitude increases. Table III lists these dips together with the frequency of motion each coincides with.

TABLE III. Applied angular frequencies at which the trap loses more than half its particles in Figure 9, in MHz, for each amplitude f of the time-dependent potential term, and the frequency of the motion each sits at. A dash means no dip at that amplitude.

	$f = 0.1$	$f = 0.4$	$f = 0.7$
$2\omega_z/3 = 0.46$	–	–	0.44
$\omega_z = 0.69$	–	0.68	0.64 to 0.70
$2\omega_z = 1.39$	1.36 to 1.42	1.26 to 1.52	1.14 to 1.62
$\omega_+ - \omega_- = 2.20$	2.20	2.16 to 2.24	2.12 to 2.26

Why only some frequencies empty the trap follows from the periodic motion of the particles. A perturbation that keeps time with this motion pushes the same way every cycle, so its pushes add up until the particle leaves the trap, while at any other frequency they cancel over a few cycles. The widest dip, present at every amplitude, sits at twice the axial frequency. We read this as the perturbation not pushing the particle but changing the stiffness of the potential it moves in. A stiffness that changes twice per swing along the axis pumps that swing on both sides, the way a child on a swing stands up and crouches twice per swing. A dip at the axial frequency itself appears from $f = 0.4$. The dip at 0.44 MHz sits close to two thirds of the axial frequency, and we do not place it further. The remaining dip sits at the difference of the two plane frequencies, and there is none at their sum, which is ω_0. We do not try to explain this pattern in the plane further. A larger amplitude pumps harder, so the dips widen, and dips absent at the smallest amplitude appear.

To see whether the Coulomb interactions have an impact on this structure, Figure 10 zooms in on the widest dip at the smallest amplitude with fine-grained frequency scans, with and without the interactions, three times from different random initializations. With the interactions the dip reaches further down in frequency, while its floor and its upper edge stay nearly where they were. We read this as the charge of the other particles partly cancelling the pull of the trap. Each particle then feels a weaker trap and swings along the axis a little slower, so the frequency that pumps it, twice its own axial frequency, moves lower. The centre of mass of the cloud is not shifted in this way, as Figure 6 showed for two par-

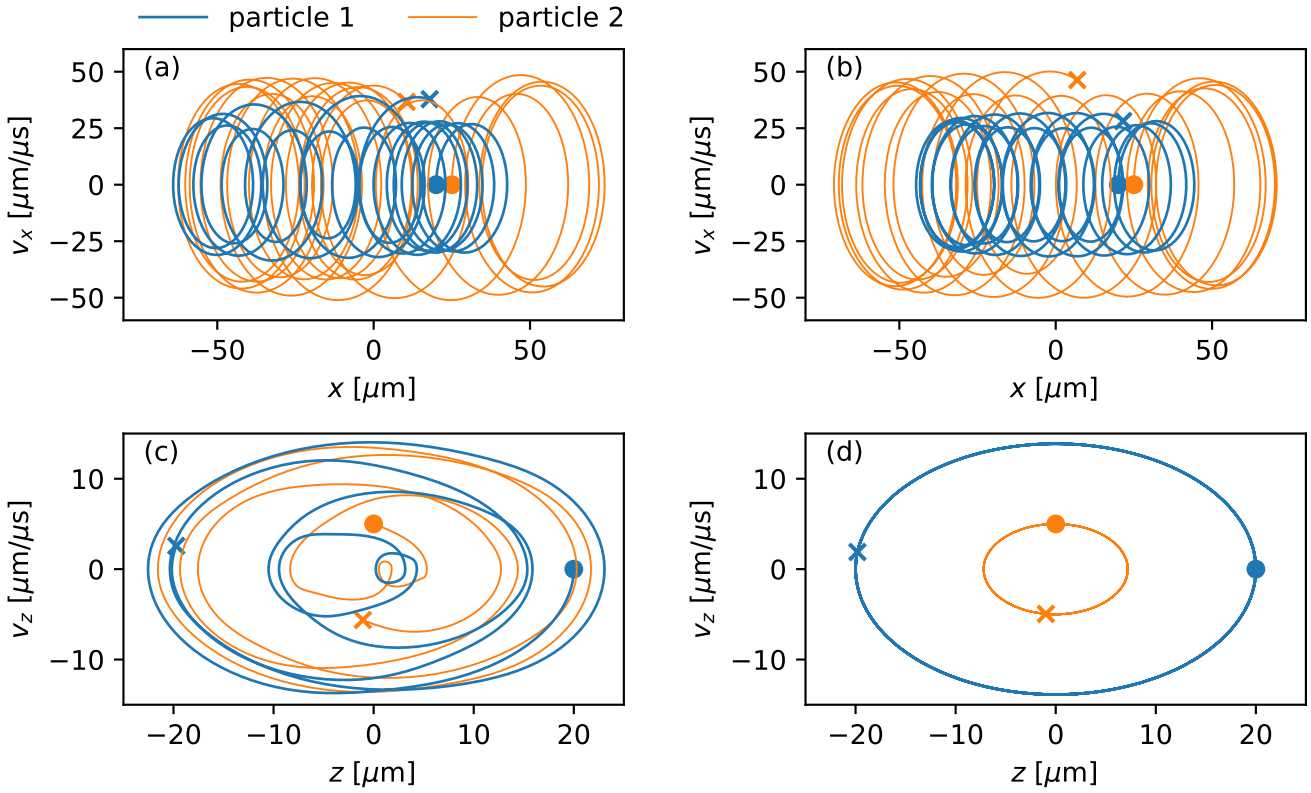

FIG. 7. Phase space, position against velocity, of the two ions of Figure 6, x against v_x above and z against v_z below, with the Coulomb force in (a) and (c) and without it in (b) and (d). A dot marks the start of each path and a cross its end, and the two rows have different axis ranges. In (d) each ion goes around one closed ellipse, the second a small one. In (c) the ellipses no longer close, as the first ion's shrinks almost to a point around 30 µs while the second ion's grows, before the two trade back. In (b) each ion draws a row of small loops whose centres slide back and forth along x. In (a) the loops keep their shape, but the first ion reaches $x = -64$ µm, against -45 µm without.

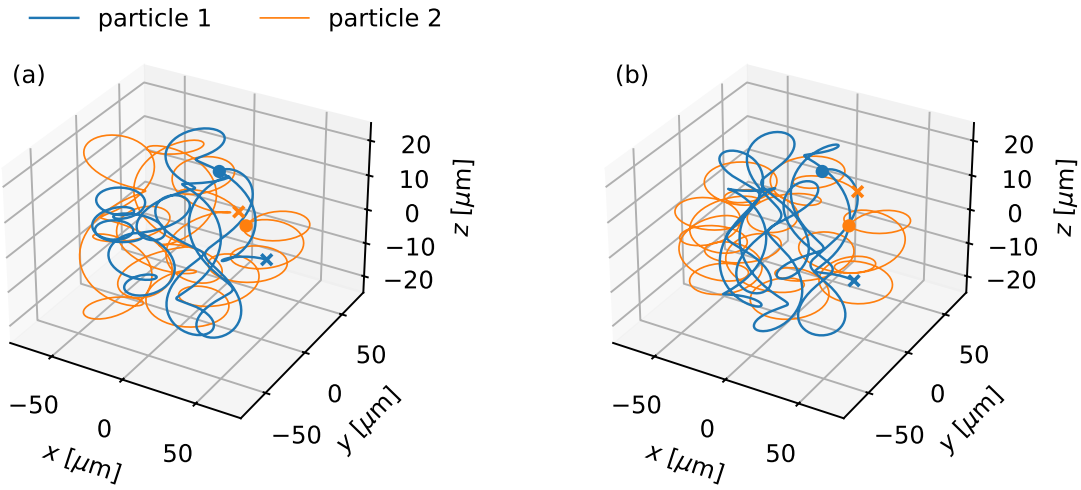

FIG. 8. The paths of the two Ca^+ ions in three dimensions over 50 µs, with the Coulomb force in (a) and without it in (b). A dot marks the start of each path and a cross its end. In (b) the first ion swings about 20 µm to either side of the middle plane, while the second stays in a thin band around it. In (a) both ions swing about 20 µm to either side. Neither ion comes further than about 75 µm from the centre of the trap in either panel.

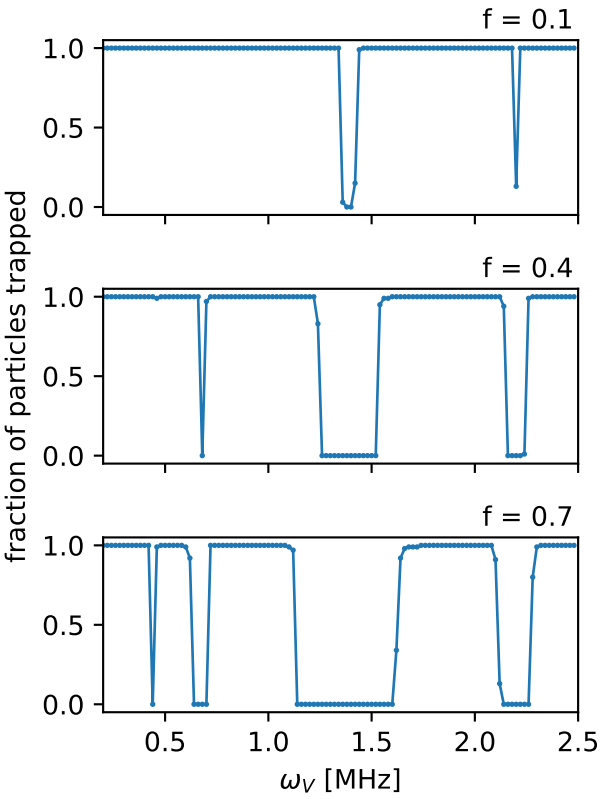

FIG. 9. Fraction of 100 Ca^+ particles still trapped after 500 µs against the applied angular frequency ω_V, for the amplitudes $f = 0.1$, 0.4 and 0.7 of the time-dependent potential term, with the Coulomb interactions switched off. Steps of 0.02 MHz, RK4 with 40000 steps. The markers are the computed points. The dip at 1.39 MHz, twice the axial frequency, is present at every amplitude and widens close to in proportion to f.

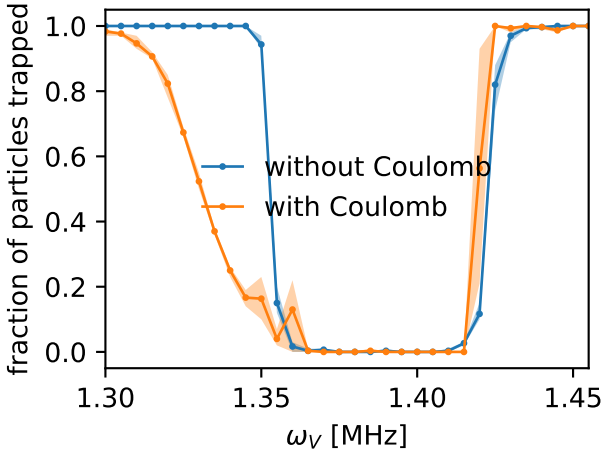

FIG. 10. Zoom on the dip near 1.4 MHz at $f = 0.1$. Fraction of 100 Ca^+ particles still trapped after 500 µs against ω_V, in steps of 0.005 MHz, with and without the Coulomb interactions. Lines are the mean of three random initializations, with a marker at each computed frequency, and the shaded bands run from the lowest to the highest of the three; where no band shows, the three agree. With the interactions fewer than half the particles stay from 1.335 MHz, against 1.355 MHz without. The upper edges lie within 0.005 MHz of each other.

ticles, which may be why the upper edge stays put. We have not tested this reading, for instance by scanning a smaller cloud.

Three limits of these scans should be kept in mind. Each point of the broad scan is one run of 100 particles from one random initialization, so one particle is a hundredth of the fraction, and points on the edges of a dip change from one initialization to the next. The zoom averages three initializations, and its bands show that the spread between them is largest on the edges of the dip. The broad scan also steps by 0.02 MHz, close to the whole width of the narrowest dips it finds, so dips narrower than that are missed. Lastly, 100 interacting particles are not converged at 40000 steps inside the dip. Rerunning the low edge with 80000 steps from the same initialization, the fraction left with the interactions on at 1.36 MHz falls from 0.22 to 0.03, while the shoulder from 1.33 to 1.345 MHz moves by at most 0.03 and no point without the interactions moves. The bumps inside the dip with interactions in Figure 10 are thus the step size, and the shoulder is not. A next step would be more random initializations per point, a finer broad scan, more steps inside the dips, and a scan with a smaller cloud to test the reading of Figure 10.

IV. CONCLUSION

We have simulated calcium ions in a Penning trap in C++, from a single ion tested against the analytical solution, through two ions coupled by the Coulomb force, to a cloud of 100 ions under a time-dependent applied potential.

The single ion served as the test of the two integration methods. RK4 converged at the rate of 4.00 its Taylor expansion predicts, its largest error falling to 2×10^{-9} µm with 32000 steps. Forward Euler converged at 1.40, and with 4000 steps its relative error peaked at 2.06. RK4 was therefore the method of every later run. The error of both methods oscillated with a period of about 2.9 µs, the time the fast circular motion takes to lap the slow one, placing the oscillation in the path of the ion and not in the method.

The two ions showed the Coulomb force reshaping the paths without loosening the trap. The force pushed the first ion out to 69 µm from the axis, against 45 µm without it, yet neither ion came further than about 75 µm from the centre. Since the centre of mass was the same with and without the force, and the total energy of the pair stayed constant in both runs, the force only moved

motion between the ions and between directions.

The cloud of 100 ions was lost only at a few frequencies of the applied potential, the trap holding every ion at no fewer than 64 of the 115 frequencies scanned for each amplitude. The widest loss sat at twice the axial frequency, which we read as the potential changing the stiffness of the trap twice per swing along the axis. With the Coulomb interactions switched on, this dip reached 0.02 MHz further down in frequency, its upper edge staying in place. The dip near two thirds of the axial frequency and the missing dip at the sum of the plane frequencies remain unexplained, and our reading of the shift is untested.

A next study would use more random initializations per point, a finer broad scan and more steps inside the dips, where 100 interacting ions are not yet converged at 40000 steps. A scan of a smaller cloud would in turn test whether the weakened trap is what moves the lower edge. Many ions under a time-dependent potential are the conditions in which the analytical solution of the single ion no longer applies, and in which the traps named in the Introduction run.

[1] J. J. Thomson. XL. Cathode Rays. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 44(269):293–316, 1897.

[2] Hans G. Dehmelt. Experiments with an isolated subatomic particle at rest. *Reviews of Modern Physics*, 62(3):525–530, 1990.

[3] C. Amole et al. The ALPHA antihydrogen trapping apparatus. *Nuclear Instruments and Methods in Physics Research Section A*, 735:319–340, 2014.

[4] A. Kellerbauer et al. Proposed antimatter gravity measurement with an antihydrogen beam. *Nuclear Instruments and Methods in Physics Research Section B*, 266(3):351–356, 2008.

[5] C. Smorra et al. BASE – The Baryon Antibaryon Symmetry Experiment. *The European Physical Journal Special Topics*, 224(16):3055–3108, 2015.

[6] Samuel Earnshaw. On the nature of the molecular forces which regulate the constitution of the luminiferous ether. *Transactions of the Cambridge Philosophical Society*, 7(1):97–112, 1842.

[7] F. M. Penning. Die Glimmentladung bei niedrigem Druck zwischen koaxialen Zylindern in einem axialen Magnetfeld. *Physica*, 3(9):873–894, 1936.

[8] Lowell S. Brown and G. Gabrielse. Geonium theory: Physics of a single electron or ion in a Penning trap. *Reviews of Modern Physics*, 58(1):233–311, 1986.

[9] Justin Bohnet, Brian C. Sawyer, J. Britton, Michael L. Wall, Ana María Rey, Michael Foss-Feig, and J. J. Bollinger. Quantum spin dynamics and entanglement generation with hundreds of trapped ions. *Science*, 352(6291):1297–1301, 2016.

[10] M. J. Borchert, J. A. Devlin, Stefan Erlewein, M. Fleck, J. A. Harrington, Takashi Higuchi, B. M. Latacz, F. Voelksen, E. Wursten, F. Abbass, M. Bohman, Andreas H. Mooser, D. Popper, M. Wiesinger, C. Will, K. Blaum, Y. Matsuda, C. Ospelkaus, W. Quint, J. Walz, Y. Yamazaki, C. Smorra, and S. Ulmer. A 16-parts-per-trillion measurement of the antiproton-to-proton charge–mass ratio. *Nature*, 601(7891):53–57, 2022.

[11] K. Blaum, Yu. N. Novikov, and G. Werth. Penning traps as a versatile tool for precise experiments in fundamental physics. *Contemporary Physics*, 51(2):149–175, 2010.

[12] Anders Kvellestad and the FYS3150 teaching team. FYS3150/FYS4150 course material. https://github.com/anderkve/FYS3150, 2026. Accessed 15 October 2026.

[13] Anders Kvellestad. FYS3150/FYS4150 computational physics, lecture notes, fall 2026. https://github.com/anderkve/FYS3150/blob/master/lecture_notes/2026/lecture_notes.pdf, 2026. Version 1.25.0, pages 95–97 and 112. Accessed 21 October 2026.

[14] Balth. van der Pol. LXXXVIII. On "relaxation-oscillations". *The London Edinburgh and Dublin Philosophical Magazine and Journal of Science*, 2(11):978–992, 1926.

[15] Edward N. Lorenz. Deterministic Nonperiodic Flow. *Journal of the Atmospheric Sciences*, 20(2):130–141, 1963.

[16] Ernst Hairer, Syvert P. Nørsett, and Gerhard Wanner. *Solving Ordinary Differential Equations I.* 1993.

[17] Conrad Sanderson and Ryan Curtin. Armadillo: C++ library for linear algebra and scientific computing, API documentation, version 15.6. https://arma.sourceforge.net/docs.html, 2026. Accessed 15 October 2026.

[18] Charles R. Harris, K. Jarrod Millman, Stéfan J. van der Walt, Ralf Gommers, Pauli Virtanen, David Cournapeau, Eric Wieser, Julian Taylor, Sebastian Berg, Nathaniel J. Smith, Robert Kern, Matti Picus, Stephan Hoyer, M. H. van Kerkwijk, Matthew Brett, Allan Haldane, Jaime Fernández del Río, Mark Wiebe, Pearu Peterson, Pierre Gérard-Marchant, Kevin Sheppard, Tyler Reddy, Warren Weckesser, Hameer Abbasi, Christoph Gohlke, and Travis E. Oliphant. Array programming with NumPy. *Nature*, 2020. arXiv: 2006.10256.

[19] J. D. Hunter. Matplotlib: A 2d graphics environment. *Computing in Science & Engineering*, 9(3):90–95, 2007.

[20] Anthropic. Claude code documentation. https://code.claude.com/docs/en/overview, 2026. Accessed 15 October 2026.