

A Boris-like integration scheme with spatial stepping

P. H. Stoltz* and G. Penn and J. Wurtele†

Dept. of Physics

Univ. of California

Berkeley, CA, 94720

J. R. Cary‡

Tech-X Corporation

Boulder, CO, 80301

(July 24, 2002)

Abstract

The Boris integration scheme, familiar in plasma simulation, is modified to use a spatial variable as the independent variable. This is important for beam simulations, where using the distance along the beam axis as the independent variable simplifies particle tracking through a beamline. The modified Boris scheme is second-order accurate, requires only one force calculation per particle per step, and preserves phase space structure more accurately than a Runge-Kutta integration scheme. The new scheme is implemented in the ICOOL code for simulation of muon cooling. Two ICOOL simulations demonstrate the advantages of the Boris integrator: (i) for a single particle

*also: Tech-X Corporation, Boulder, CO 80301, pstoltz@txcorp.com

†also: Lawrence Berkeley Laboratory, Berkeley, CA 80309

‡also: Univeristy of Colorado, Boulder, CO 80309

moving through a uniform solenoid field, the Boris scheme improved conservation of perpendicular momentum more than two orders of magnitude over the Runge-Kutta, and (ii) for a simulation with 2500 particles and a beamline based on the Brookhaven Study 2, the overall code speed with the Boris integrator improved a factor of three over the Runge-Kutta.

A popular integration scheme in plasma physics codes is the Boris integration scheme [1]. The Boris integrator is second-order accurate, requires only one force evaluation per step, compared to four for a standard, fourth-order Runge-Kutta (RK) scheme [2], and it preserves phase space structure more accurately than an RK integrator. Although it is only second-order accurate, because of the decrease in the number of required force evaluations and the increase in phase space accuracy, the Boris scheme can be much faster for many applications. However, beam physics simulations have not taken advantage of the Boris scheme because it uses time as the independent variable, while in beam physics simulation, the spatial coordinate along the beam axis is more often the independent variable since that makes matching spatial boundary conditions in the magnetic lattice easier. The aim of this paper is to describe a modified Boris scheme that uses a spatial independent variable, making it suitable for beam physics codes. We first describe the theory of the modified Boris scheme, and then discuss the improvement resulting from implementing this scheme in the beam simulation code, ICOOL [3]. In ICOOL, momentum conservation was improved two orders of magnitude for integration of a single particle through a uniform solenoid field and overall code speed was improved by a factor of three for a simulation with 2500 particles and a beamline based on Brookhaven Study 2 [4].

Accurate simulation of muon cooling is important to the Muon Collider Collaboration to help decide between many different design proposals for a muon collider [5]. ICOOL is one of the main tools for such simulation, and speeding up the code allows more rapid progress in collider design and optimization. ICOOL has many features that make it a uniquely important tool to the Muon Collider Collaboration, including sophisticated modeling of muons interacting with materials, which allows users to study ionization beam cooling (the most widely-proposed method of creating muon beams of high enough luminosity for a collider). Further, ICOOL has many analytic field options, making it easy to create a field of any specification. Decreasing the runtime of an ICOOL simulation will let users complete their analysis more quickly and study a wider range of design options. Using the Runge-Kutta integrator, ICOOL spends a large fraction of run time in the particle

integration subroutines, so implementing a more efficient integrator is one way to improve code performance.

The motivation for developing the spatial Boris scheme comes from relativity and formal Hamiltonian theory. Both of these fields give formal methods for exchanging spatial and temporal dimensions. We begin by writing the momentum evolution equations, but substituting the energy evolution equation for the variable we want to become the independent variable (z in this paper):

$$\frac{dp_x}{dt} = q(E_x + v_y B_z - v_z B_y) \quad (1)$$

$$\frac{dp_y}{dt} = q(E_y + v_z B_x - v_x B_z) \quad (2)$$

$$\frac{d(U/c)}{dt} = q(E_x v_x + E_y v_y + E_z v_z) \quad (3)$$

where $p_z^2 = (U/c)^2 - p_x^2 - p_y^2 - m^2 c^2$ is now the energy-like quantity (MKS units). Interactions with material in the lattice, including stochastic effects, are not included here. In the ICOOL code, the algorithm for modeling such processes is independent of the particle integration scheme. One exchanges z for t on the left-hand side by multiplying through by $1/v_z$:

$$\frac{dp_x}{dz} = \frac{1}{v_z} \frac{dp_x}{dt} = q \left(\frac{E_x}{v_z} + \frac{v_y B_z}{v_z} - B_y \right) \quad (4)$$

$$\frac{dp_y}{dz} = \frac{1}{v_z} \frac{dp_y}{dt} = q \left(\frac{E_y}{v_z} - \frac{v_x B_z}{v_z} + B_x \right) \quad (5)$$

$$\frac{dU/c}{dz} = \frac{1}{v_z} \frac{dU/c}{dt} = q \left(\frac{v_x E_x}{v_z c} + \frac{v_y E_y}{v_z c} + \frac{E_z}{c} \right). \quad (6)$$

Exchanging v_z for p_z on the right-hand side, one can write this as a matrix equation, splitting into terms that involve p_z and those that do not:

$$\frac{d}{dz} \begin{pmatrix} p_x \\ p_y \\ U/c \end{pmatrix} = \frac{q}{p_z} \begin{pmatrix} 0 & B_z & E_x/c \\ -B_z & 0 & E_y/c \\ E_x/c & E_y/c & 0 \end{pmatrix} \begin{pmatrix} p_x \\ p_y \\ U/c \end{pmatrix} + q \begin{pmatrix} -B_y \\ B_x \\ E_z/c \end{pmatrix} \quad (7)$$

We define

$$w = \begin{pmatrix} p_x \\ p_y \\ U/c \end{pmatrix}, \quad (8)$$

$$b = q \begin{pmatrix} -B_y \\ B_x \\ E_z/c \end{pmatrix}, \quad (9)$$

and

$$M = \frac{q}{p_z} \begin{pmatrix} 0 & B_z & E_x/c \\ -B_z & 0 & E_y/c \\ E_x/c & E_y/c & 0 \end{pmatrix} \quad (10)$$

so that Eq. 7 can be written

$$\frac{dw}{dz} = Mw + b. \quad (11)$$

The equation for the evolution of the generalized particle position, s , can be written:

$$\frac{ds}{dz} = \frac{w}{p_z}, \quad (12)$$

where $s = [x, y, ct]$. Equations 11 and 12 are the equations one needs to advance in the numerical integration scheme.

The goal is to find a scheme that will advance these two sets of equations with second-order accuracy while requiring only one force evaluation per step. One can leapfrog [6] the advance of the generalized positions (Eq. 12) with the generalized momenta (Eq. 11). This means advancing the positions one-half a step, advancing the momenta a full step, and advancing the positions half a step. In this scheme, one assumes that the momenta are constant when advancing the positions and that the positions are constant when advancing the momenta. This will be second-order accurate so long as each piece is at least second-order accurate. The integration of the positions is exact assuming constant momenta, so one only must find a way to integrate the momentum equation to second-order accuracy. As with the temporal Boris scheme, the approach here will be to further split the advance of Eq. 11 in a leapfrog way: (i) advance w first by only the vector term, b , for one-half a step, (ii) advance by only the matrix term, M , a full step, then (iii) advance by the vector b a

final one-half step. Because the positions are assumed constant in the momentum advance, all the terms in b are constant, and so steps (i) and (iii) are exact. All that is left, then, is to show step (ii) can be done to second-order accuracy. So long as the elements of M are all constant during this step, a space-centered advance (i.e. using the average of w on the right-hand side and solving the resulting implicit equation) is second-order accurate.

To show that M is constant, one needs to show that the coefficient p_z is constant during step (ii). Because M involves only the field components B_z , E_x , and E_y that do not directly modify p_z , one expects that p_z will be constant. Formally, one can show that p_z does not change magnitude when operated on by M by considering:

$$\frac{dp_z}{dz} = \frac{1}{p_z} \left(U/c \frac{dU/c}{dz} - p_x \frac{dp_x}{dz} - p_y \frac{dp_y}{dz} \right). \quad (13)$$

Using the equations of motion from Eq. 7 but including only the matrix term gives:

$$\begin{aligned} U/c \frac{dU/c}{dz} &= U/c \left[\frac{q}{p_z} \left(\frac{E_x p_x}{c} + \frac{E_y p_y}{c} \right) \right] \\ p_x \frac{dp_x}{dz} &= p_x \left[\frac{q}{p_z} \left(B_z p_y + \frac{E_x U}{c^2} \right) \right] \\ p_y \frac{dp_y}{dz} &= p_y \left[\frac{q}{p_z} \left(-B_z p_x + \frac{E_y U}{c^2} \right) \right]. \end{aligned} \quad (14)$$

Substituting into Eq. 13 yields (for the matrix term alone):

$$\frac{dp_z}{dz} = 0. \quad (15)$$

Thus all change in the magnitude of p_z is due to the vector b . This means that for a step involving only the matrix term, the elements of the matrix are constant, and a space-centered advance will be second-order accurate.

Using the space-centered advance scheme for step (ii) means that one must solve an implicit equation. One uses the average w on the right-hand side of the equation:

$$w^+ - w^- = M \left(\frac{w^+ + w^-}{2} \right) \Delta z, \quad (16)$$

where w^- is the vector before the matrix operation, and w^+ is the vector after. Solving for w^+ gives

$$w^+ = \left(I - M \frac{\Delta z}{2}\right)^{-1} \left(I + M \frac{\Delta z}{2}\right) w^-. \quad (17)$$

To calculate $(I - M \Delta z/2)^{-1}$, one needs the eigenvalues of M . The eigenvalues of M are

$$\lambda = 0, \pm i \frac{q}{p_z} \sqrt{B_z^2 - E_x^2 - E_y^2} = 0, \pm i \hat{\lambda}. \quad (18)$$

In the corresponding basis of eigenvectors, this leads to

$$\left(I - M \frac{\Delta z}{2}\right)^{-1} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & \frac{1+i\hat{\lambda}\Delta z/2}{1+\hat{\lambda}^2\Delta z^2/4} & 0 \\ 0 & 0 & \frac{1-i\hat{\lambda}\Delta z/2}{1+\hat{\lambda}^2\Delta z^2/4} \end{pmatrix}. \quad (19)$$

One can then write the matrix from Eq. 17 in a basis-independent way:

$$\left(I - M \frac{\Delta z}{2}\right)^{-1} \left(I + M \frac{\Delta z}{2}\right) = I + \frac{\Delta z}{1 + \hat{\lambda}^2 \Delta z^2 / 4} M + \frac{\Delta z^2 / 2}{1 + \hat{\lambda}^2 \Delta z^2 / 4} M^2. \quad (20)$$

So one can advance from w^- to w^+ using Eq. 17 rewritten as

$$w^+ = w^- + \frac{\Delta z}{1 + \hat{\lambda}^2 \Delta z^2 / 4} M w^- + \frac{\Delta z^2 / 2}{1 + \hat{\lambda}^2 \Delta z^2 / 4} M^2 w^-. \quad (21)$$

The full operator for advancing w^- to w^+ can be written out as:

$$w^+ = w^- + R w^-, \quad (22)$$

where from Eq. 21:

$$R = \frac{\Delta z M + \frac{\Delta z^2 M^2}{2}}{1 + \frac{\hat{\lambda}^2 \Delta z^2}{4}} \quad (23)$$

$$= \frac{q \Delta z / p_z}{1 + \frac{q^2 \Delta z^2}{4 p_z^2} \left(B_z^2 - \frac{E_x^2 + E_y^2}{c^2}\right)} \begin{pmatrix} \frac{q \Delta z}{2 p_z} \left(\frac{E_x^2}{c^2} - B_z^2\right) & B_z + \frac{q \Delta z}{2 p_z} \frac{E_x E_y}{c^2} & \frac{E_x}{c} + \frac{q \Delta z}{2 p_z} \frac{B_z E_y}{c} \\ -B_z + \frac{q \Delta z}{2 p_z} \frac{E_x E_y}{c^2} & \frac{q \Delta z}{2 p_z} \left(\frac{E_y^2}{c^2} - B_z^2\right) & \frac{E_y}{c} - \frac{q \Delta z}{2 p_z} \frac{B_z E_x}{c} \\ \frac{E_x}{c} - \frac{q \Delta z}{2 p_z} \frac{B_z E_y}{c} & \frac{E_y}{c} + \frac{q \Delta z}{2 p_z} \frac{B_z E_x}{c} & \frac{q \Delta z}{2 p_z} \left(\frac{E_x^2}{c^2} + \frac{E_y^2}{c^2}\right) \end{pmatrix}.$$

The final steps for the spatial Boris push to move from position z^n to z^{n+1} are as follows:

- i. Push the generalized positions one-half step ($\Delta z/2$) using the velocities at z^n .
- ii. Evaluate the fields at this mid-point time and position.

- iii. Push the generalized momenta vector, w , from w^n to an intermediate state w^- with a half-step using only the vector b :

$$w^- = w^n + \frac{\Delta z}{2}b. \quad (24)$$

- iv. Evaluate p_z at this point, plug into the matrix R , and advance w^- to w^+ with a full spatial step of the matrix part of Eq 11:

$$w^+ = w^- + R w^-. \quad (25)$$

- v. Advance w^+ to the final state w^{n+1} with a half-step using only the vector b :

$$w^{n+1} = w^+ + \frac{\Delta z}{2}b. \quad (26)$$

- vi. Push the generalized positions one-half step using the velocities at z^{n+1} .

The above steps require only one evaluation of the fields. In a traditional leap-frog scheme, steps i) and vi) are combined, and the positions are then known one-half step off from the momenta. By keeping steps i) and vi) separate, one knows the generalized positions and momenta at the same spatial location at the end of a step.

Next we discuss the application of this scheme to two simulations using ICOOL. The first simulation is of single-particle motion through a uniform solenoid and is meant to show the improvement in preserving phase space structure using the Boris integrator. Because phase space quantities like beam temperature and energy spread are important for determining how well a neutrino factory or muon collider will perform, accurately modeling phase space is important for simulations of muon beams. Runge-Kutta integration schemes are known to do a poor job preserving phase space structures [7]. Because the modified Boris scheme is based on operator splitting like a symplectic scheme, we expect it to do a much better job at accurately preserving phase space.

To test this, we used the ICOOL code with a beamline consisting of a uniform, $7T$ solenoid field and a beam with a z-momentum of $0.2 GeV/c$. We chose roughly 30 steps

per period as the stepsize (for these parameters, the period is approximately $0.6 m$, and so we choose a stepsize of $0.02 m$). Because typical ICOOL simulations run for a few thousand steps, we chose to run the simulation $600m$. In a uniform solenoid field total energy, total momentum, and perpendicular momentum ($p_{perp}^2 = p_x^2 + p_y^2$) are conserved. For the parameters relevant to a muon beam (like those given above), the total energy and momentum are dominated by motion parallel to the magnetic field. This motion is force-free and so is not sensitive to numerical integrator properties. Thus, we chose to examine the perpendicular momentum. Figure 1 shows the comparison of the perpendicular momentum for simulations using the modified Boris and the RK schemes. The simulation using the modified Boris scheme conserves perpendicular momentum to better than 0.02%, while the simulation using the RK scheme lost approximately 2.0% of the perpendicular momentum. This represents an improvement of more than two orders of magnitude. This does not imply ICOOL simulations using the RK integrator are incorrect; for a smaller stepsize (64 steps per period), the RK integrator performs equally well. Rather, this example illustrates the Boris scheme is more robust than the RK scheme with regard to properties like phase space conservation.

For the second simulation, we wanted to study to overall speed improvement of the code for a more realistic ICOOL simulation. We chose to use solenoid focusing lattices and a beamline geometry based on a simplified version of Brookhaven Feasibility Study 2 [4]. The objectives of Study 2 were to examine the possibility of upgrading the Brookhaven AGS facility (a high-energy proton source) as the first element of a muon-based neutrino factory, and to study the feasibility and cost of a high-performance version of such a machine. Because of the significant cost of such a channel, a great deal of effort has been put into attempts to minimize the length and required magnetic structure of the cooling channel, while at the same time maintaining the cooling channel in order to provide the large neutrino flux is required for the physics studies which would be done at such a facility. Therefore, a rapid, accurate prediction of the performance of the cooling channel is a critical aspect of the neutrino factory design process.

The cooling channel works by passing the muon beam through chambers filled with liquid hydrogen, leading to energy loss because the muons slow down. The muons are then re-accelerated in the forward direction using RF cavities, thus leading to overall cooling of the beam. The cooling rate is proportional to the rate of energy loss, but is counteracted by Coulomb scattering events, both in the hydrogen and in metal foils which are an intrinsic part of the cooling channel; the scattering tends to randomize the direction of the particles. The evolution of the beam emittance depends on the balance between scattering events, which are simulated using the Moliere model [8], and the energy loss in material, leading to a minimum transverse emittance that can itself be reduced by focusing the beam more tightly at the hydrogen. In addition, due to the strong focusing and large energy spread of the beam, nonlinear effects of beam propagation and in the RF cavities used for bunching and acceleration strongly affect the performance of the cooling channel.

To compare the different integration schemes, we chose to compare two quantities: beam emittance and single-particle transverse orbits. Emittance is of the most interest scientifically, while single-particle orbits give a more direct measure of the integrator's performance. For the single-particle orbit measurements, we chose to suppress stochastic processes such as Coulomb scattering and muon decay. This reduces disagreement between the particle orbits from sources other than the integrator. In this example, the accuracy of the integration is determined by the first-order-accurate solver for the energy loss, so to achieve the same accuracy, the same step size is used for both schemes.

Figure 2 shows the comparison of the emittance for the two different integration schemes. The solid line is the emittance from the simulation using the RK integrator, and the dashed line is from the Boris integrator. These simulations used 2500 particles and the same stepsize of approximately 1.0 cm to model 90 m of the cooling channel. The Boris integrator ran this problem approximately three times faster. Figure 3 shows the percent difference between the RK run and the Boris run. This plot is the dashed line. With 2500 particles, statistically one can expect agreement only to within $1/\sqrt{2500} \approx$ a few percent. We expect the percent error to grow as z increases because the emittance is decreasing throughout the calculation.

The solid line is the percent difference between the RK run and another RK run with a different random seed. Because the errors between the Boris scheme and the RK scheme are of the order we expect from statistical arguments, and they are of the same order as errors from the RK scheme run with different random seeds, we conclude that the Boris scheme provides satisfactory results while running three times faster.

Figure 4 shows the comparison of the single-particle orbits for the two different integration schemes (random processes turned off). The orbits shown are at the end of the run (approximately 10,000 steps into the simulation). Again, the solid line is simulation using the RK integrator, and the dashed line is from the Boris integrator. Both runs had an RMS difference of a few percent with a control run using the RK integrator and a stepsize 10 times smaller. Errors of this size are acceptable given that one is interested in calculating emittance, and statistical and random errors in the emittance calculation were larger than these errors.

In conclusion, we have adapted the temporal Boris integration scheme to use a spatial variable as the independent variable. This makes the Boris scheme suitable for beam physics simulations, where using a spatial step is easier for matching boundary conditions in the magnetic lattice. The Boris scheme is second-order accurate and requires only one force evaluation per step, making it more efficient for many applications than a Runge-Kutta scheme. We implemented the new scheme in the muon tracking code, ICOOL. For a simulation of single-particle motion in a uniform solenoid, the modified Boris scheme improved perpendicular momentum conservation by more than two orders of magnitude. On simulations typical in muon cooling (2500 particles), runs with the Boris integrator were three times faster than runs with the RK integrator for similar accuracy.

Acknowledgments: The authors thank W. Fawley and R. Fernow for many suggestions regarding this work.

REFERENCES

- [1] J. Boris, *Proceedings of the Fourth Conference on Numerical Simulation of Plasma* (Naval Research Lab, Washington DC, 1970).
- [2] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical Recipes in C: The Art of Scientific Computing* (Cambridge Univ. Press, Cambridge, 1993).
- [3] R. Fernow, *Proceedings of the 1999 Particle Accelerator Conference* (IEEE, Piscataway, N.J., 1999), p. 3020.
- [4] S. Ozaki, R. Palmer, M. Zisman, and J. Gallardo, 2001, bNL-52623.
- [5] C. M. Ankenbrandt *et al.*, Phys. Rev. ST Accel. Beams **2**, (1999).
- [6] S. Gudunov, Math. USSR-Sb. **47**, 271 (1959).
- [7] A. Dullweber, B. Leimkuhler, and R. McLachlan, J. Chem. Phys. **15**, 5840 (1997).
- [8] C. M. Ankenbrandt *et al.*, Phys. Rev. **89**, 1256 (1953).

FIGURES




FIG. 1. The perpendicular momentum ($p_{perp} = \sqrt{p_x^2 + p_y^2}$), normalized to its initial value, as a function of distance along a uniform solenoid channel. The solid line is from a simulation using a Runge-Kutta integration scheme. The dashed line is from a simulation using the modified Boris integration scheme. The simulation using the modified Boris scheme conserved perpendicular momentum to better than 0.02%, more than two orders of magnitude better than the Runge-Kutta.

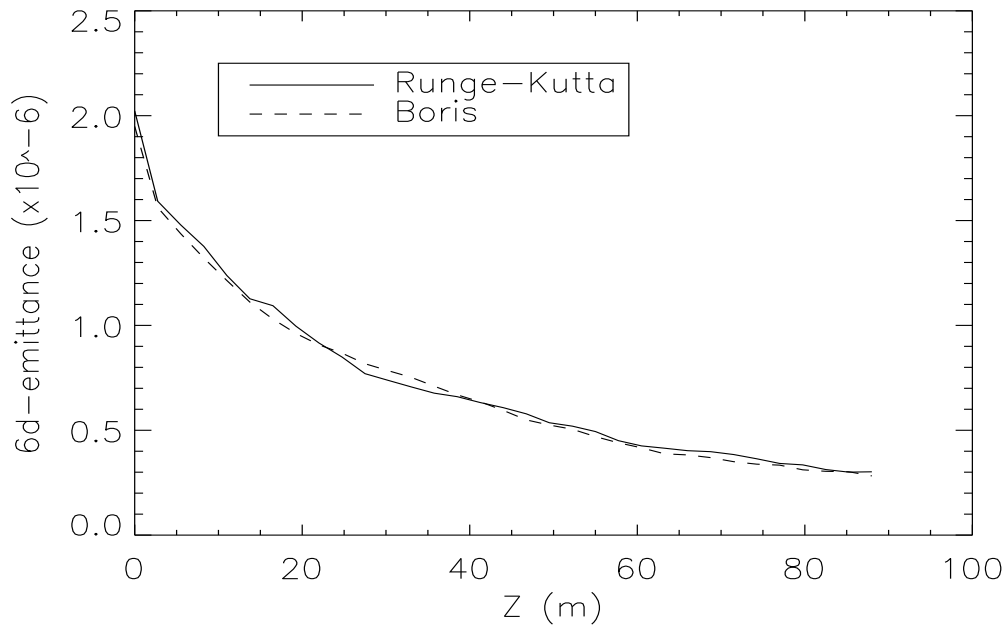


FIG. 2. The emittance as a function of distance along the cooling channel for a simulation with 2500 particles and using a beamline based on Brookhaven Study 2. The solid line is from a simulation using a Runge-Kutta integration scheme. The dashed line is from a simulation using the modified Boris integration scheme. The simulation using the Boris scheme ran this problem three times faster.

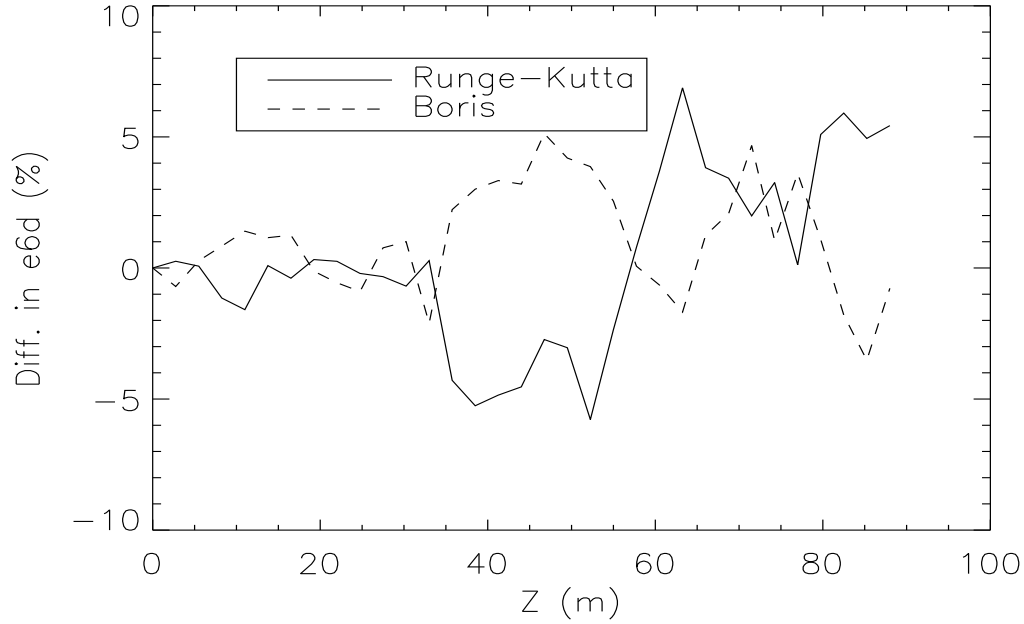


FIG. 3. The percent difference in emittance as a function of distance along the cooling channel. The solid line is the difference between the simulation shown in Fig. 2 using the Runge-Kutta scheme and another run with the Runge-Kutta scheme, but a different random seed. The dashed line is the difference between the simulation using the Runge-Kutta scheme and the simulation using the Boris scheme. These errors are of the same order, indicating that the Boris scheme is giving results of approximately the same accuracy as the Runge-Kutta scheme. We expect the percent error to grow as z increases because the emittance is decreasing throughout the calculation.

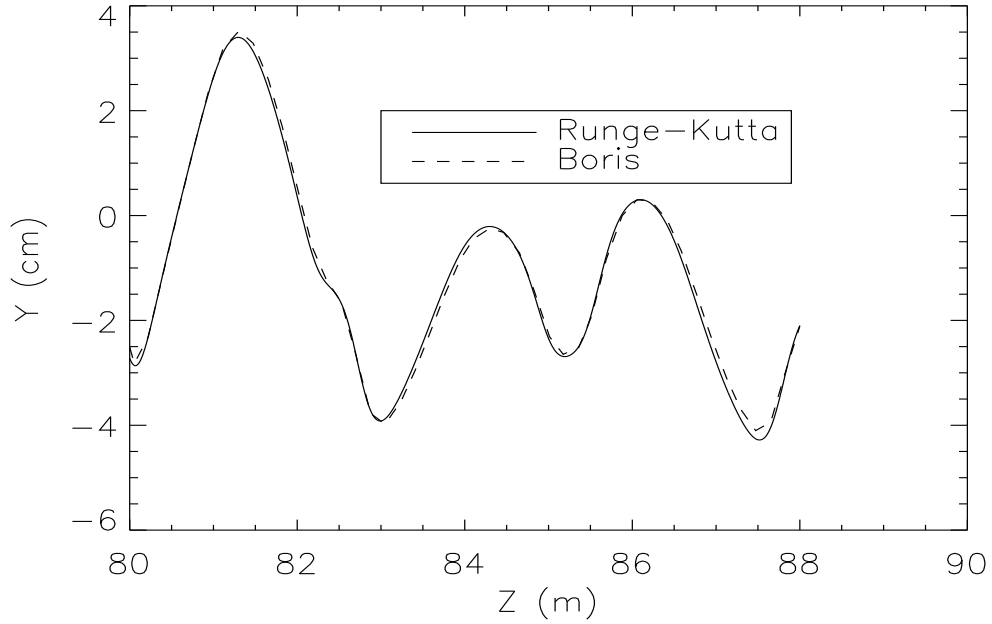


FIG. 4. The transverse position, y , as a function of the distance along the cooling channel for a single particle using the same beamline based on Study 2, but with stochastic processes turned off. The solid line is a simulation using the Runge-Kutta scheme. The dashed line is a simulation using the Boris scheme. For these simulations, random processes such as Coulomb scattering and muon decay were turned off.