

An arbitrary order time-stepping algorithm for tracking particles in inhomogeneous magnetic fields

Krasymyr Tretiak, Daniel Ruprecht*

School of Mechanical Engineering, University of Leeds, United Kingdom



ARTICLE INFO

Article history:

Received 20 December 2018
Received in revised form 14 June 2019
Accepted 19 July 2019
Available online 23 August 2019

Keywords:

Boris integrator
Particle tracking
High-order time integration
Spectral deferred corrections
Fusion reactor

ABSTRACT

The Lorentz equations describe the motion of electrically charged particles in electric and magnetic fields and are used widely in plasma physics. The most popular numerical algorithm for solving them is the Boris method, a variant of the Störmer-Verlet algorithm. Boris method is phase space volume conserving and simulated particles typically remain near the correct trajectory. However, it is only second order accurate. Therefore, in scenarios where it is not enough to know that a particle stays on the right trajectory but one needs to know where on the trajectory the particle is at a given time, Boris method requires very small time steps to deliver accurate phase information, making it computationally expensive. We derive an improved version of the high-order Boris spectral deferred correction algorithm (Boris-SDC) by adopting a convergence acceleration strategy for second order problems based on the Generalised Minimum Residual (GMRES) method. Our new algorithm is easy to implement as it still relies on the standard Boris method. Like Boris-SDC it can deliver arbitrary order of accuracy through simple changes of runtime parameter but possesses better long-term energy stability. We demonstrate for two examples, a magnetic mirror trap and the Solev'ev equilibrium, that the new method can deliver better accuracy at lower computational cost compared to the standard Boris method. While our examples are motivated by tracking ions in the magnetic field of a nuclear fusion reactor, the introduced algorithm can potentially deliver similar improvements in efficiency for other applications.

© 2019 The Author(s). Published by Elsevier Inc. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

1. Introduction

The Lorentz equations

$$\dot{\mathbf{x}} = \mathbf{v}, \quad (1a)$$

$$\dot{\mathbf{v}} = \alpha [\mathbf{E}(\mathbf{x}, t) + \mathbf{v} \times \mathbf{B}(\mathbf{x}, t)] =: \mathbf{f}(\mathbf{x}, \mathbf{v}) \quad (1b)$$

model movement of charged particles in electro-magnetic fields. Here, $\mathbf{x}(t)$ is a vector containing all particle position at some time t , $\mathbf{v}(t)$ contains all particle velocities, α is the charge-to-mass ratio, \mathbf{E} the electric field (both externally applied and internally generated from particle interaction) and \mathbf{B} the magnetic field. The Lorentz equations are used in many ap-

* Corresponding author.

E-mail addresses: k.tretiak@leeds.ac.uk (K. Tretiak), d.ruprecht@leeds.ac.uk (D. Ruprecht).

plications in computational plasma physics, for example laser-plasma interactions [1], particle accelerators [2] or nuclear fusion reactors [3].

The Boris method, introduced by Boris in 1970 [4], is the most popular numerical scheme used for solving (1) although other numerical time stepping methods like Runge-Kutta-4 are used as well. It is based on the Leapfrog algorithm but uses a special trick to resolve the seemingly implicit dependence that arises from the fact that the Lorentz force depends on \mathbf{v} . Its popularity is because it is computationally cheap, second order accurate and phase space conserving [5]. While it was recently shown that for general magnetic fields this does not guarantee energy conservation and that the Boris method can exhibit energy drift [6], it is nevertheless a surprisingly good algorithm. In most cases, particles will remain close to their correct trajectory because of its conservation properties. However, Boris' method can introduce substantial phase errors and, for long time simulations, it only ensures that particles are near the right trajectory – it does not provide information about where on the trajectory they are at a given time.

For some applications this is not an issue because the only required information is whether a particle passes through some region but not when it does so. In these cases, phase errors are of no concern and the Boris algorithm is highly competitive, combining low computational cost with high quality results. There are other applications, however, where accurate phase information is crucial. One example are particle-wave interactions triggering Alfvén instabilities due to resonances between orbit frequencies and wave velocities [7]. Because it is only second order accurate, the Boris method requires very small time steps, creating substantial computational cost. In these cases, methods of order higher than two can be more efficient.

For separable Hamiltonians, the development of explicit symmetric integrators has been studied for decades [8]. However, the Lorentz equations (1) give rise to a non-separable Hamiltonian, making development of higher order methods challenging, see the overview by He et al. [9]. Quandt et al. suggest a high order integrator based on a Taylor series expansion and demonstrate high convergence order for relativistic and non-relativistic test cases [10]. The method needs derivatives of the electric and magnetic field, though, which may be difficult to obtain. A recently introduced new class of methods are so-called explicit symplectic shadowed Runge-Kutta methods or ESSRK for short [11]. They are symplectic and therefore have bounded long-term energy error. ESSRK has been shown to be more accurate than Runge-Kutta-4 with respect to both energy and phase error but also require substantially more sub-steps. No comparison with respect to computational efficiency seems to exist. He et al. introduce a high-order volume preserving method based on splitting and composition of low order methods [12]. A class of symmetric multi-step methods is derived by Hairer and Lubich but not analysed with respect to computational efficiency [13]. Instead of building higher order methods, Umeda [14] constructs a three-step version of the Boris method that can be about a factor of two faster.

Spectral deferred correction (SDC), introduced by Dutt et al. in 2000 [15], are iterative time stepping methods based on collocation. In each time step, they perform multiple sweeps with a low order integrator (often a form of Euler method) in order to generate a higher order approximation. This paper presents a new high order algorithm for solving the Lorentz equations (1) called Boris-GMRES-SDC or BGSDC for short. Its key advantages are that it is straightforward to implement since it heavily relies on the classical Boris scheme which will be available in almost any plasma modelling code. Furthermore, it allows to flexibly tune the order of accuracy by simply changing runtime parameters without the need to solve equations for order conditions. SDC also provides dense output and allows to generate a high order solution anywhere within a time step. We use this feature to accurately compute the turning points of particles in a magnetic mirror. The codes used to generate the numerical examples are available for download from GitHub [16,17].

BGSDC is an extension of Boris spectral deferred corrections (Boris-SDC), introduced and tested for homogeneous electric and magnetic fields by Winkel et al. [18]. The present paper expands their results in multiple ways. First, it provides a slightly simplified version of the method with almost identical performance. Second, it integrates a GMRES-based convergence accelerator, originally introduced by Huang et al. [19] for first order problems, with Boris-SDC. We show that this leads to a substantial improvement in the long-term energy error. Third, it studies the performance of BGSDC for inhomogeneous magnetic fields, in contrast to Winkel et al. who only studied the homogeneous case.

While BGSDC can be applied to problems where an electric field is present, we focus here on tracking fast particles in the core region of the plasma in a nuclear fusion reactor. There, the effect of the electric field generated from particle-interactions is small, although not totally negligible, and often ignored [3,20]. However, to include the effect of $\mathbf{E} \times \mathbf{B}$ drift on the numerical accuracy of BGSDC, we add a weak, external electric field. Also, to be able to quantitatively compare the accuracy of BGSDC and the classical Boris algorithm for single trajectories, we make two simplifying assumptions. First, we only consider test cases where the magnetic field is given by a mathematical formula (a magnetic mirror and a Solov'ev equilibrium), in contrast to a real reactor where the field is given by a numerical solution to the Grad-Shafranov equation fitted to experimental data. Second, we neglect the stochastic models used to capture the effect of interactions of fast ions with the plasma. An implementation of BGSDC into the LOCUST-GPU simulation software [21] and experiments for realistic use cases for the DIII-D, JET and ITER experimental fusion reactors are ongoing work.

Verlet-based versus Leapfrog-based Boris integrator. Boris-SDC relies on the classical velocity-Verlet scheme applied to (1), which reads

$$\mathbf{x}_{n+1} = \mathbf{x}_n + \Delta t \left(\mathbf{v}_n + \frac{\Delta t}{2} \mathbf{f}(\mathbf{x}_n, \mathbf{v}_n) \right), \quad (2a)$$

$$\mathbf{v}_{n+1} = \mathbf{v}_n + \frac{\Delta t}{2} (\mathbf{f}(\mathbf{x}_n, \mathbf{v}_n) + \mathbf{f}(\mathbf{x}_{n+1}, \mathbf{v}_{n+1})), \quad (2b)$$

with $\mathbf{x}_{n+1} \approx \mathbf{x}(t_{n+1})$, $\mathbf{v}_{n+1} \approx \mathbf{v}(t_{n+1})$ being numerical approximations of the analytical solution at some time step t_{n+1} . The seemingly implicit dependence in \mathbf{v}_{n+1} is resolved using the trick sketched in Algorithm 1 introduced by Boris in 1970 [4]. What is typically referred to as ‘‘Boris algorithm’’ is the staggered Leapfrog method

$$\mathbf{v}_{n+1/2} = \mathbf{v}_{n-1/2} + \Delta t \mathbf{f}(\mathbf{x}_n, \mathbf{v}_n) \quad (3a)$$

$$\mathbf{x}_{n+1} = \mathbf{x}_n + \Delta t \mathbf{v}_{n+1/2} \quad (3b)$$

which can be rewritten in ‘‘kick-drift-kick’’ form

$$\mathbf{v}_{n+1/2} = \mathbf{v}_n + \frac{\Delta t}{2} \mathbf{f}(\mathbf{x}_n, \mathbf{v}_n) \quad (4a)$$

$$\mathbf{x}_{n+1} = \mathbf{x}_n + \Delta t \mathbf{v}_{n+1/2} \quad (4b)$$

$$\mathbf{v}_{n+1} = \mathbf{v}_{n+1/2} + \frac{\Delta t}{2} \mathbf{f}(\mathbf{x}_{n+1}, \mathbf{v}_{n+1}) \quad (4c)$$

where the Boris-trick is used in (4c). While Velocity-Verlet (2) and Leapfrog (4) are similar they are not equivalent, see for example the analysis by Mazur [22]. In particular, in the absence of an electric field, the staggered version conserves kinetic energy exactly. Below, we refer to (2) plus the Boris trick as ‘‘unstaggered Boris’’ and to (4) with Boris trick as ‘‘staggered Boris’’ method.

While a variant of Boris-SDC can be derived based on the staggered Leapfrog method, it requires additional storage of solutions at half-points and, in tests not documented here, was not found to improve performance over the velocity-Verlet based Boris-SDC. Substantial differences between Verlet and Leapfrog seem only to arise in simulations with very large time steps with nearly no significant digits left (phase errors well above 10^{-1}), where staggered Boris shows better stability. In such regimes, BGSDC is not going to be competitive anyway so that we focus here on the simpler Verlet-based variant.

Algorithm 1 : Boris’ trick for unstaggered Velocity-Verlet (2). See Birdsall and Langdon [23, Section 4–4] for the geometric derivation.

input : $\mathbf{x}_{m-1}, \mathbf{x}_m, \mathbf{v}_{m-1}, \Delta t$
output : \mathbf{v}_m solving $\mathbf{v}_m = \mathbf{v}_{m-1} + \Delta t \mathbf{E}_{m-1/2} + \Delta t \frac{\mathbf{v}_{m-1} + \mathbf{v}_m}{2} \times \mathbf{B}(\mathbf{x}_m)$

- 1.1 $\mathbf{E}_{m-1/2} = \frac{1}{2} (\mathbf{E}(\mathbf{x}_{m-1}) + \mathbf{E}(\mathbf{x}_m))$
- 1.2 $\mathbf{t} = \frac{\Delta t}{2} \mathbf{B}(\mathbf{x}_m)$
- 1.3 $\mathbf{s} = 2\mathbf{t} / (1 + \mathbf{t} \cdot \mathbf{t})$
- 1.4 $\mathbf{v}^- = \mathbf{v}_{m-1} + \frac{\Delta t}{2} \mathbf{E}_{m-1/2}$
- 1.5 $\mathbf{v}^* = \mathbf{v}^- + \mathbf{v}^- \times \mathbf{t}$
- 1.6 $\mathbf{v}^+ = \mathbf{v}^- + \mathbf{v}^* \times \mathbf{s}$
- 1.7 $\mathbf{v}_m = \mathbf{v}^+ + \frac{\Delta t}{2} \mathbf{E}_{m-1/2}$

2. Spectral deferred corrections

Spectral deferred corrections [15] are based on collocation. Therefore, we first summarise the collocation formulation of the Lorentz equations (1) before deriving the BGSDC algorithm.

2.1. Collocation

Consider a single time step $[t_n, t_{n+1}]$. Integrating (1) from t_n to some $t_n \leq t \leq t_{n+1}$ turns them into the integral equations

$$\mathbf{x}(t) = \mathbf{x}_0 + \int_{t_n}^t \mathbf{v}(s) ds \quad (5a)$$

$$\mathbf{v}(t) = \mathbf{v}_0 + \int_{t_n}^t \mathbf{f}(\mathbf{x}(s), \mathbf{v}(s)) ds \quad (5b)$$

denoting $\mathbf{x}_0 = \mathbf{x}(t_n)$ and $\mathbf{v}_0 = \mathbf{v}(t_n)$. The exact solution at the end of the time step can theoretically be found by inserting $t = t_{n+1}$. In the original paper introducing Boris-SDC for homogeneous fields [18], the second equation is substituted into the first, resulting in double integrals over \mathbf{f} in the equation for the position \mathbf{x} . For the test cases studied in this paper we could not see any meaningful improvement in performance and, since the substitution leads to more complicated notation, we omit it and work directly with equations (5).

To discretise the integral equations (5) we introduce a set of quadrature nodes $t_n \leq \tau_1 < \dots < \tau_M \leq t_{n+1}$, set $t = t_{n+1}$ and approximate

$$\int_{t_n}^{t_{n+1}} \mathbf{v}(s) ds \approx \sum_{m=1}^M q_m \mathbf{v}_m \quad (6)$$

and

$$\int_{t_n}^{t_{n+1}} \mathbf{f}(\mathbf{x}(s), \mathbf{v}(s)) ds \approx \sum_{m=1}^M q_m \mathbf{f}(\mathbf{x}_m, \mathbf{v}_m) \quad (7)$$

with $\mathbf{x}_j, \mathbf{v}_j$ being approximations of $\mathbf{x}(\tau_j), \mathbf{v}(\tau_j)$, that is of the analytical solution at the quadrature nodes. Then, approximations at t_{n+1} can be found from

$$\mathbf{x}_{\text{new}} = \mathbf{x}_0 + \sum_{m=1}^M q_m \mathbf{v}_m \quad (8a)$$

$$\mathbf{v}_{\text{new}} = \mathbf{v}_0 + \sum_{m=1}^M q_m \mathbf{f}(\mathbf{x}_m, \mathbf{v}_m). \quad (8b)$$

To turn this into a usable numerical method, we require equations for the $\mathbf{x}_m, \mathbf{v}_m$. Those can be obtained from discrete counterparts of (5) when setting $t = \tau_m$, for $m = 1, \dots, M$, resulting in

$$\mathbf{x}_m = \mathbf{x}_0 + \sum_{j=1}^M q_{m,j} \mathbf{v}_j \approx \mathbf{x}_0 + \int_{t_n}^{\tau_m} \mathbf{v}(s) ds \quad (9a)$$

$$\mathbf{v}_m = \mathbf{v}_0 + \sum_{j=1}^M q_{m,j} \mathbf{f}(\mathbf{x}_j, \mathbf{v}_j) \approx \mathbf{v}_0 + \int_{t_n}^{\tau_m} \mathbf{f}(\mathbf{x}(s), \mathbf{v}(s)) ds. \quad (9b)$$

The quadrature weights $q_{m,j}$ are given by

$$q_{m,j} = \int_{t_n}^{\tau_m} l_j(s) ds \quad (10)$$

with l_j being Lagrange polynomials with respect to the τ_m .

Solving (9) directly using Newton's method gives rise to a collocation method. Collocation methods are a special type of fully implicit Runge-Kutta methods with a full Butcher tableau. Depending on the type of quadrature nodes, they have favourable properties like symmetry (Gauss-Lobatto or Gauss-Legendre nodes) [24, Theorem 8.9] or symplecticity (Gauss-Legendre nodes) [25, Theorem 16.5] and A- and B-stability [26, Theorem 12.9]. However, note that even for formally symplectic implicit methods, accumulation of round-off error from the nonlinear solver can still lead to energy drift [27].

2.2. Boris-SDC

By packing the solutions $\mathbf{x}_m, \mathbf{v}_m$ at the quadrature nodes into a single vector

$$\mathbf{U} = (\mathbf{x}_1, \dots, \mathbf{x}_M, \mathbf{v}_1, \dots, \mathbf{v}_M)^T, \quad (11)$$

the discrete collocation problem (9) can be written as

$$\mathbf{U} - \mathbf{QF}(\mathbf{U}) = \mathbf{U}_0, \quad (12)$$

with $\mathbf{U}_0 = (\mathbf{x}_0, \dots, \mathbf{x}_0, \mathbf{v}_0, \dots, \mathbf{v}_0)$ and

$$\mathbf{Q} = \begin{pmatrix} \tilde{\mathbf{Q}} \otimes \mathbf{I}_{3M} & \mathbf{0} \\ \mathbf{0} & \tilde{\mathbf{Q}} \otimes \mathbf{I}_{3M} \end{pmatrix} \text{ with } \tilde{\mathbf{Q}} = \begin{pmatrix} q_{1,1} & \dots & q_{1,M} \\ \vdots & & \vdots \\ q_{M,1} & \dots & q_{M,M} \end{pmatrix}, \quad (13)$$

see the Appendix in Winkel et al. for details [18]. First, consider the case where \mathbf{f} is linear. For the Lorentz equations, this would be the case, for example, if \mathbf{B} is homogeneous and $\mathbf{E} = 0$. In a slight abuse of notation we write \mathbf{F} for the matrix denoting the operator

$$\mathbf{F}(\mathbf{U}) = \begin{pmatrix} \mathbf{v}_1 \\ \vdots \\ \mathbf{v}_M \\ \mathbf{f}(\mathbf{x}_1, \mathbf{v}_1) \\ \vdots \\ \mathbf{f}(\mathbf{x}_M, \mathbf{v}_M) \end{pmatrix}, \quad (14)$$

so that the nonlinear collocation problem (12) reduces to the linear system

$$(\mathbf{I} - \mathbf{QF}) \mathbf{U} = \mathbf{U}_0. \quad (15)$$

One sweep of Boris-SDC can be written as

$$\mathbf{U}^{k+1} = (\mathbf{I} - \mathbf{Q}_\Delta \mathbf{F})^{-1} \mathbf{U}_0 + [\mathbf{I} - (\mathbf{I} - \mathbf{Q}_\Delta \mathbf{F})^{-1} (\mathbf{I} - \mathbf{QF})] \mathbf{U}^k, \quad (16)$$

with

$$\mathbf{Q}_\Delta = \begin{pmatrix} \mathbf{Q}_{\Delta,E} \otimes \mathbf{I}_{3M} & \frac{1}{2} \mathbf{Q}_{\Delta,E}^{(2)} \otimes \mathbf{I}_{3M} \\ \mathbf{0} & \mathbf{Q}_{\Delta,T} \otimes \mathbf{I}_{3M} \end{pmatrix} \quad (17)$$

where

$$\mathbf{Q}_{\Delta,E} = \begin{pmatrix} 0 & \dots & 0 \\ \Delta\tau_2 & 0 & \dots \\ \Delta\tau_2 & \Delta\tau_3 & 0 \dots \\ \dots & \dots & \dots \\ \Delta\tau_2 & \Delta\tau_3 & \dots \Delta\tau_M & 0 \end{pmatrix} \quad (18)$$

and

$$\mathbf{Q}_{\Delta,I} = \begin{pmatrix} \Delta\tau_1 & 0 & \dots \\ \Delta\tau_1 & \Delta\tau_2 & 0 \dots \\ \dots & \dots & \dots \\ \Delta\tau_1 & \Delta\tau_2 & \dots \Delta\tau_M \end{pmatrix} \quad (19)$$

and $\mathbf{Q}_{\Delta,T} = \frac{1}{2} (\mathbf{Q}_{\Delta,E} + \mathbf{Q}_{\Delta,I})$ and $\mathbf{Q}_{\Delta,E}^{(2)} := \mathbf{Q}_{\Delta,E} \circ \mathbf{Q}_{\Delta,E}$, see again Winkel et al. for details [18].¹ Iteration (16) can be understood as a Picard iteration applied to the preconditioned system

$$(\mathbf{I} - \mathbf{Q}_\Delta \mathbf{F})^{-1} (\mathbf{I} - \mathbf{QF}) \mathbf{U} = (\mathbf{I} - \mathbf{Q}_\Delta \mathbf{F})^{-1} \mathbf{U}_0. \quad (20)$$

2.3. Boris-GMRES-SDC (BGSDC)

For linear first order differential equations, Huang et al. showed that performing k iterations of the Generalized Minimum Residual (GMRES) algorithm on (20) often gives better results than performing k standard SDC iterations [19,28]. Here, we adopt their strategy to the second order Lorentz equations for cases where the magnetic field varies only weakly over a single time step. Note that while we rely on a self-written GMRES implementation in the accompanying code, we verified that it gives identical results to the GMRES implementation in the SciPy library [29].

GMRES does not require the matrix representing the linear system or the preconditioner to be assembled explicitly. It only requires functions that compute $(\mathbf{I} - \mathbf{QF}) \mathbf{U}$ given some \mathbf{U} and solve

$$(\mathbf{I} - \mathbf{Q}_\Delta \mathbf{F}) \mathbf{U} = \mathbf{b} \quad (21)$$

given some \mathbf{b} [30]. Applying \mathbf{QF} amounts to computing the sums in (9) for $m = 1, \dots, M$ as in the original Boris-SDC. Systems of the form $(\mathbf{I} - \mathbf{Q}_\Delta \mathbf{F}) \mathbf{U} = \mathbf{B}$ can be solved by elimination in a sweep-like fashion. For $M = 3$ nodes, Eq. (21) becomes

¹ As pointed out by one of the reviewers, it is also possible to use an implicit midpoint rule instead of trapezoidal rule to update the velocity. In tests not documented here, this variant of Boris-SDC showed improved long-term energy errors compared to the variant using trapezoidal rule. However, it cannot directly interpreted in the form of a preconditioned iteration given by Eq. (20), so that it is not clear how to apply GMRES acceleration to this variant. Nevertheless, it would certainly warrant further study, in particular for problems with strong nonlinearities.

$$\begin{bmatrix} \mathbf{x}_1 \\ \mathbf{x}_2 \\ \mathbf{x}_3 \end{bmatrix} - \begin{bmatrix} 0 & 0 & 0 \\ \Delta\tau_2 \mathbf{I} & 0 & 0 \\ \Delta\tau_2 \mathbf{I} & \Delta\tau_3 \mathbf{I} & 0 \end{bmatrix} \begin{bmatrix} \mathbf{v}_1 \\ \mathbf{v}_2 \\ \mathbf{v}_3 \end{bmatrix} - \frac{1}{2} \begin{bmatrix} 0 & 0 & 0 \\ \Delta\tau_2^2 \mathbf{I} & 0 & 0 \\ \Delta\tau_2^2 \mathbf{I} & \Delta\tau_3^2 \mathbf{I} & 0 \end{bmatrix} \begin{bmatrix} \mathbf{F}(\mathbf{x}_1, \mathbf{v}_1) \\ \mathbf{F}(\mathbf{x}_2, \mathbf{v}_2) \\ \mathbf{F}(\mathbf{x}_3, \mathbf{v}_3) \end{bmatrix} = \begin{bmatrix} \mathbf{b}_1 \\ \mathbf{b}_2 \\ \mathbf{b}_3 \end{bmatrix}$$

and

$$\begin{bmatrix} \mathbf{v}_1 \\ \mathbf{v}_2 \\ \mathbf{v}_3 \end{bmatrix} - \frac{1}{2} \begin{bmatrix} \Delta\tau_1 \mathbf{I} & 0 & 0 \\ (\Delta\tau_1 + \Delta\tau_2) \mathbf{I} & \Delta\tau_2 \mathbf{I} & 0 \\ (\Delta\tau_1 + \Delta\tau_2) \mathbf{I} & (\Delta\tau_2 + \Delta\tau_3) \mathbf{I} & \Delta\tau_3 \mathbf{I} \end{bmatrix} \begin{bmatrix} \mathbf{F}(\mathbf{x}_1, \mathbf{v}_1) \\ \mathbf{F}(\mathbf{x}_2, \mathbf{v}_2) \\ \mathbf{F}(\mathbf{x}_3, \mathbf{v}_3) \end{bmatrix} = \begin{bmatrix} \mathbf{b}_4 \\ \mathbf{b}_5 \\ \mathbf{b}_6 \end{bmatrix}.$$

This system can be solved for $\mathbf{x}_i, \mathbf{v}_i$ by computing

$$\begin{aligned} \mathbf{x}_1 &= \mathbf{b}_1 \\ \mathbf{v}_1 &= \mathbf{b}_4 + \frac{1}{2} \Delta\tau_1 \mathbf{F}(\mathbf{x}_1, \mathbf{v}_1) \\ \mathbf{x}_2 &= \mathbf{b}_2 + \Delta\tau_2 \mathbf{v}_1 + \frac{1}{2} \Delta\tau_2^2 \mathbf{F}(\mathbf{x}_1, \mathbf{v}_1) \\ \mathbf{v}_2 &= \mathbf{b}_5 + \frac{1}{2} (\Delta\tau_1 + \Delta\tau_2) \mathbf{F}(\mathbf{x}_1, \mathbf{v}_1) + \frac{1}{2} \Delta\tau_2 \mathbf{F}(\mathbf{x}_2, \mathbf{v}_2) \\ \mathbf{x}_3 &= \mathbf{b}_3 + \Delta\tau_2 \mathbf{v}_1 + \Delta\tau_3 \mathbf{v}_2 + \frac{1}{2} \Delta\tau_2^2 \mathbf{F}(\mathbf{x}_1, \mathbf{v}_1) + \frac{1}{2} \Delta\tau_3^2 \mathbf{F}(\mathbf{x}_2, \mathbf{v}_2) \\ \mathbf{v}_3 &= \mathbf{b}_6 + \frac{1}{2} (\Delta\tau_1 + \Delta\tau_2) \mathbf{F}(\mathbf{x}_1, \mathbf{v}_1) + \frac{1}{2} (\Delta\tau_2 + \Delta\tau_3) \mathbf{F}(\mathbf{x}_2, \mathbf{v}_2) + \frac{1}{2} \Delta\tau_3 \mathbf{F}(\mathbf{x}_3, \mathbf{v}_3) \end{aligned}$$

using Boris' trick to compute the velocities. The generalisation to other values of M is straightforward.

BGSDC for inhomogeneous magnetic fields

GMRES is a solver for linear systems and will not work if \mathbf{f} and thus \mathbf{F} are nonlinear. In their original work, Huang et al. suggest to adopt GMRES-SDC to nonlinear problems by employing an outer Newton iteration and using GMRES-SDC as inner iteration to solve the arising linear problems. In tests not documented in this paper we found that this approach requires too many sweeps and was not competitive for the problems studied here. Instead, we propose a different strategy for scenarios where \mathbf{B} is changing slowly over the course of a time step and the nonlinearity is therefore weak.

It starts with a single sweep with standard non-staggered Boris to generate approximate values $\mathbf{x}_m^0, \mathbf{v}_m^0$ at all nodes. In the notation above this is equivalent to solving

$$\mathbf{U}^0 - \mathbf{Q}_\Delta \mathbf{F}(\mathbf{U}^0) = \mathbf{U}_0 \quad (22)$$

by block-wise elimination. Then, we linearize the function \mathbf{F} by setting

$$\mathbf{F}_{\text{lin}}(\mathbf{X}_0)(\mathbf{U}) = \begin{pmatrix} \mathbf{v}_1 \\ \vdots \\ \mathbf{v}_M \\ \mathbf{f}(\mathbf{x}_1^0, \mathbf{v}_1) \\ \vdots \\ \mathbf{f}(\mathbf{x}_M^0, \mathbf{v}_1) \end{pmatrix}. \quad (23)$$

That is, the magnetic field applied to the velocity \mathbf{v}_m is not $\mathbf{B}(\mathbf{x}_m)$ but $\mathbf{B}(\mathbf{x}_m^0)$ and remains fixed during the GMRES iteration. We then apply a small number of GMRES iterations to the preconditioned linearised collocation equation

$$(\mathbf{I} - \mathbf{Q}_\Delta \mathbf{F}_{\text{lin}}(\mathbf{X}_0))^{-1} (\mathbf{I} - \mathbf{Q} \mathbf{F}_{\text{lin}}(\mathbf{X}_0)) \mathbf{U} = (\mathbf{I} - \mathbf{Q}_\Delta \mathbf{F}_{\text{lin}}(\mathbf{X}_0))^{-1} \mathbf{U}_0. \quad (24)$$

Algorithm 2: Single time step of BGSDC($k_{\text{gmres}}, k_{\text{picard}}$) for weakly nonlinear problems.

input : $\mathbf{x}_0, \mathbf{v}_0$

output : $\mathbf{x}_{\text{new}}, \mathbf{v}_{\text{new}}$

2.1 Set $\mathbf{U}_0 = (\mathbf{x}_0, \dots, \mathbf{x}_0, \mathbf{v}_0, \dots, \mathbf{v}_0)$

2.2 Perform a single nonlinear Boris-SDC sweep to solve $\mathbf{U}^0 - \mathbf{Q}_\Delta \mathbf{F}(\mathbf{U}^0) = \mathbf{U}_0$

2.3 $\mathbf{U}_{\text{lin}} \leftarrow$ Perform k_{gmres} iterations of GMRES-SDC on the linearised collocation equation (24) using \mathbf{U}^0 as starting value

2.4 Perform k_{picard} Picard iterations $\mathbf{U}^k = \mathbf{U}_0 + \mathbf{Q} \mathbf{F}(\mathbf{U}^{k-1})$ with $\mathbf{U}^0 = \mathbf{U}_{\text{lin}}$

2.5 Perform update step (8) to compute $\mathbf{x}_{\text{new}}, \mathbf{v}_{\text{new}}$

Algorithm 3 : Picard iteration.

```

input :  $\mathbf{U}^k = (\mathbf{x}_1^k, \mathbf{v}_1^k, \dots, \mathbf{x}_M^k, \mathbf{v}_M^k), \mathbf{x}_0, \mathbf{v}_0$ 
output :  $\mathbf{U}^{k+1} = \mathbf{U}_0 + \mathbf{QF}(\mathbf{U})$ 
3.1 for  $m = 1, \dots, M$  do
3.2  $\mathbf{x}_m^{k+1} = \mathbf{x}_0 + \sum_{j=1}^M q_{m,j} \mathbf{v}_j^k$ 
3.3  $\mathbf{v}_m^{k+1} = \mathbf{v}_0 + \sum_{j=1}^M q_{m,j} \mathbf{f}(\mathbf{x}_j^k, \mathbf{v}_j^k)$ 
3.4 end
3.5  $\mathbf{U}^{k+1} = (\mathbf{x}_1^{k+1}, \mathbf{v}_1^{k+1}, \dots, \mathbf{x}_M^{k+1}, \mathbf{v}_M^{k+1})$ 

```

For a slowly varying magnetic field this will provide an approximation \mathbf{U}_{lin} that is close to the solution of the nonlinear collocation problem (12). We then apply a small number of Picard iterations as sketched in Algorithm 3 using \mathbf{U}_{lin} as starting value. Picard iterations only require application of \mathbf{QF} and do not need Boris' trick, so they are computationally cheap. However, they only converge for starting values that are close to the collocation solution or for small time steps. Therefore, Picard iterations alone were not found to be competitive with either standard Boris-SDC or BGSDC. But for weakly nonlinear problems, the solution to the linearised collocation problem (24) is close to the nonlinear collocation solution (12) so that the output from the linearised GMRES procedure is a very accurate starting value. Using full Boris-SDC sweeps instead of Picard iterations is also possible and, in tests not documented here, resulted in smaller errors in some cases. We found the reduction in error is likely not significant enough to justify the higher complexity of full sweeps but leave a detailed comparison for future work.

It was recently observed that the entries in the \mathbf{Q}_Δ matrix can be changed without losing the sweep-like structure of SDC. For first order problems, this allows to build more efficient sweeps resembling DIRK schemes [31]. In particular, one can use optimization routines to find entries for \mathbf{Q}_Δ that provide rapid convergence. We tried to adopt this approach to second order problems but were unable to find a robust strategy that delivered improved results for a reasonably wide range of parameters.

Computational effort. We use the number of evaluations of \mathbf{f} required by each method as a proxy for computational effort W . While Boris' trick requires some additional computation, in realistic simulations with experimentally given magnetic fields, evaluation of $\mathbf{B}(\mathbf{x}_m)$ dominates the computational cost because of the required interpolation. Therefore, we count each application of Boris trick as one evaluation of \mathbf{f} , ignoring the cost of computing vector products. Non-staggered Boris (2) requires one evaluation of \mathbf{f} per time step. Thus, its total cost when computing N_{steps} many time steps is simply

$$W_{\text{boris}} = N_{\text{steps}}. \quad (25)$$

In contrast, the initial predictor step in BGSDC requires $M - 1$ Boris steps and the computation of $\mathbf{f}(\mathbf{x}_0, \mathbf{v}_0)$ for a total of M evaluations. Computing $F_{\text{lin}}(\mathbf{X}_0)$ requires $M - 1$ evaluations of \mathbf{f} for Gauss-Lobatto nodes.² Because we keep the magnetic field fixed in the GMRES iterations, there is no additional cost in terms of evaluations of \mathbf{f} . Finally, Picard iterations each require $M - 1$ evaluations of \mathbf{f} and the update step requires another $M - 1$. Therefore, the total estimated cost of BGSDC is

$$W_{\text{gmres}} = N_{\text{steps}} \left(\underbrace{M}_{\text{predictor}} + \underbrace{(M-1)}_{\text{Compute } \mathbf{F}(\mathbf{X}_0)} + \underbrace{(M-1)K_{\text{picard}}}_{\text{Picard iteration}} + \underbrace{M-1}_{\text{Update step (8)}} \right). \quad (26)$$

In principle, all of those steps except the predictor can be parallelised by using $M - 1$ threads to do the \mathbf{f} evaluations for all quadrature nodes in parallel. That would allow them to be computed in the wall clock time required for a single evaluation. Parallelisation would reduce the cost of the algorithm to

$$W_{\text{gmres}} = N_{\text{steps}} (M + 1 + K_{\text{picard}} + 1 + \tau_{\text{overhead}}). \quad (27)$$

Here, τ_{overhead} accounts for any overheads, for example from threads competing for memory bandwidth. There are approaches available to parallelise a full SDC sweep instead of only the Picard iteration [32] but those have not yet been adopted for second order problems. We leave those as well as the development of an effective parallel implementation and a detailed assessment of required wall clock times for future work.

3. Numerical results

We compare the performance of BGSDC against both the staggered and non-staggered Boris method for two fusion-related test problems. The first is a magnetic mirror where particles are confined by a magnetic field generated by two

² We experimented with Gauss-Legendre nodes but found the resulting BGSDC method not competitive for the studied examples. Therefore, we use Gauss-Lobatto nodes throughout the paper.




Fig. 1. Example article trajectory in a magnetic mirror trap. Evolution in time is indicated by the line's changing colour from blue ($t = 0$) to green ($t = t_{end}$).

coils. Our second benchmark uses a Solev'ev equilibrium which resembles the magnetic field in the Joint European Torus (JET) experimental Tokamak reactor.

3.1. Magnetic mirror trap

We use a simple mathematical model that has similar characteristics as a magnetic mirror trap. The static but non-uniform magnetic field between the coils $\mathbf{B} = (B_x, B_y, B_z)$ has components

$$B_x = -B_0 \frac{xz}{z_0^2} \quad (28a)$$

$$B_y = -B_0 \frac{yz}{z_0^2} \quad (28b)$$

$$B_z = B_0 \left(1 - \frac{z^2}{z_0^2}\right). \quad (28c)$$

Here, $B_0 = \omega_B/\alpha$ is the magnetic field at the centre of trap, ω_B is the cyclotron frequency, $\alpha = q/m$ is the particle's charge-to-mass ratio and z_0 the distance between coil and centre. Note that (28) is not a valid approximation of a mirror trap's magnetic field outside of the two coils [33].

Fig. 1 shows an example trajectory of a particle that remains vertically confined, reflecting back and forth between points at around $z = -3$ and $z = 3$. Note that the parameters, see Table 1 (right), were chosen to create a recognisable trajectory and are different than the ones for the two examples reported below, see Table 1 (left and middle).

The basic physical principle of magnetic mirroring [34,35] is that charged particles in a longitudinal axially symmetric static magnetic field bounded by coils with higher value of magnetic field on both sides will be reflected from these high field side regions when moving along magnetic field lines. This is due to the invariance of a charged particle's magnetic moment

$$\mu = \frac{1}{2} \frac{m\mathbf{v}_\perp^2}{B}, \quad B = (B_x^2 + B_y^2 + B_z^2)^{1/2}, \quad (29)$$

in the adiabatic limit where

$$\varepsilon = \frac{\rho_L}{L} \ll 1 \quad \text{or} \quad \frac{B}{|\nabla B|} \gg \rho_L, \quad (30)$$

with ρ_L being the Larmor radius of particle, L the radius of curvature of the magnetic field line and the particle's velocity $\mathbf{v} = \mathbf{v}_\perp + \mathbf{v}_\parallel$ being split into a part perpendicular to the magnetic field lines and a parallel part. As a particle moves from a low field to a high field side region, B increases and therefore, according to (29), \mathbf{v}_\perp^2 must increase in order to keep μ constant. Since the particle's kinetic energy

$$E_{\text{kin}} = \frac{m\mathbf{v}_\perp^2}{2} + \frac{m\mathbf{v}_\parallel^2}{2} \quad (31)$$

remains constant, the parallel velocity \mathbf{v}_\parallel^2 must decrease. When B becomes large enough, \mathbf{v}_\parallel approaches zero and the particle is reflected and travels back along the field line.

Table 1

Parameter for scenario 1 with $\varepsilon \sim 10^{-4}$ (left), scenario 2 with $\varepsilon \sim 10^{-2}$ (middle) and for visualization (right) of a single classical particle's trajectory in Fig. 1 for the magnetic mirror trap.

t_{end}	50	t_{end}	16	t_{end}	0.485
α	1	α	1	α	1
z_0	200	z_0	16	z_0	8
ω_B	2000	ω_B	400	ω_B	200
$\vec{x}(t=0)$	(1.0, 0.5, 0)	$\vec{x}(t=0)$	(1.0, 0, 0)	$\vec{x}(t=0)$	(5.25, 5.25, 0)
$\vec{v}(t=0)$	(100, 0, 50)	$\vec{v}(t=0)$	(100, 0, 50)	$\vec{v}(t=0)$	(100, 0, 50)
N_{steps}	variable	N_{steps}	variable	N_{steps}	1000
Δt	variable	Δt	variable	Δt	t_{end}/N_{steps}

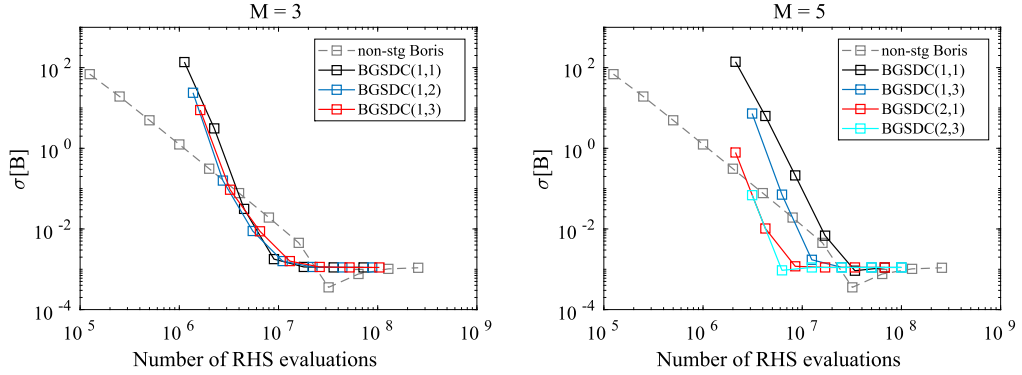


Fig. 2. Error $\sigma[B]$ measured against the analytically computed value of B at reflection points in the limit $\varepsilon \rightarrow 0$ plotted against the number of \mathbf{f} evaluations.

3.1.1. Scenario 1: $\varepsilon \sim 10^{-4}$

In the adiabatic limit $\varepsilon \rightarrow 0$ we can determine the strength of the magnetic field at the points where the particle is reflected. Comparing this value against the magnetic field at numerically computed reflection points allows to measure the precision of BGSDC for very small ε . Simulation parameters are summarised in Table 1 (left) and correspond to a value $\varepsilon \sim 8 \cdot 10^{-5}$.

Consider a particle with initial velocity \mathbf{v}_0 and position \mathbf{x}_0 and $B_0 = \|\mathbf{B}(\mathbf{x}_0)\|_2$ being the strength of the magnetic field at the particle's initial position. Denote as B_{ref} the strength of the magnetic field at the reflection point and as $\mathbf{v}_{\perp r}$ the perpendicular velocity. It follows from conservation of magnetic moment μ that

$$\frac{\mathbf{v}_{\perp 0}^2}{B_0} = \frac{\mathbf{v}_{\perp r}^2}{B_{ref}}. \quad (32)$$

Conservation of kinetic energy gives

$$\frac{m\mathbf{v}_{\perp 0}^2}{2} + \frac{m\mathbf{v}_{\parallel 0}^2}{2} = \frac{m\mathbf{v}_{\perp r}^2}{2} \quad (33)$$

because, by definition, $\mathbf{v}_{\parallel r} = 0$. Using (33) we can substitute $\mathbf{v}_{\perp r}$ in (32)

$$\frac{B_{ref}}{B_0} = \frac{\mathbf{v}_{\perp 0}^2 + \mathbf{v}_{\parallel 0}^2}{\mathbf{v}_{\perp 0}^2} = \frac{1}{\sin^2(\varphi)}, \quad (34)$$

allowing us to compute B_{ref} directly from the initial conditions $\mathbf{x}_0, \mathbf{v}_0$. Here, φ is the so-called pitch angle. It is assumed that the maximum magnetic field strength is on the coil $B_{max} = |\mathbf{B}(z_0)|$ and $\varphi_{min} = \arcsin(\sqrt{B_0/B_{max}})$. Note that magnetic moment is only exactly conserved in the limit $\varepsilon \rightarrow 0$. For small but finite values of ε , the actual value of B at the reflection point will be close to but not identical to B_{ref} .

Denote by B_i the strengths of the magnetic field at the numerically computed reflection points. We compute the l_2 weighted error

$$\sigma[B] = \sqrt{\frac{1}{N} \sum_{i=1}^N (B_{ref} - B_i)^2}, \quad (35)$$

where N is the number of times the particle was reflected. To compute the B_i , we exploit the fact that SDC allows to reconstruct solutions at arbitrary times in a time step with high order of accuracy. If a particle is reflected in the current time step $[t_n, t_{n+1}]$, detected by a sign change in v_{\parallel} , we construct the Lagrange polynomial

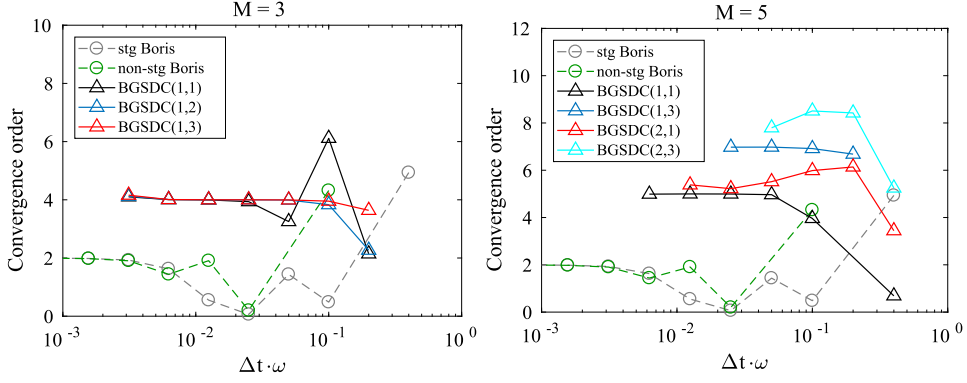


Fig. 3. Numerically observed convergence order p for the x coordinate of the particle's final position in the magnetic mirror as a function of time step size for $M=3$ (left) and $M=5$ (right) Gauss-Lobatto collocation nodes per time step and different fixed numbers of iterations per SDC sweep.

$$L(t) = \sum_{m=0}^M v_{\parallel m} l_m(t), \quad (36)$$

using values $\mathbf{x}_m, \mathbf{v}_m$ from intermediate nodes, where $v_{\parallel} = v \cos(\varphi)$ and $\cos(\varphi) = \mathbf{v} \cdot \mathbf{B} / (|\mathbf{v}| |\mathbf{B}|)$. The function $L(t)$ interpolates \mathbf{v}_{\parallel} on the interval $[t_n, t_{n+1}]$ with order M . Then, we use bisection root-finding to find the time t_{ref} at which $L(t_{\text{ref}}) = 0$. From t_{ref} we can find the position \mathbf{x}_{ref} of the reflection point using a Lagrange polynomial defined by the positions \mathbf{x}_m at the quadrature nodes and compute the value of B at that point.

Fig. 2 shows $\sigma[B]$ for BGSDC with $M=3$ quadrature nodes (left) and with $M=5$ nodes (right) against the total number of \mathbf{f} evaluations. Values from the Boris method are identical in both images. Because B_{ref} holds only in the adiabatic limit whereas we have a small but finite value of ε , errors saturate at around 10^{-3} for all numerically computed solutions. For both $M=3$ and $M=5$, BGSDC is more efficient for precisions of 10^{-1} and below, requiring fewer evaluations of \mathbf{f} than the Boris algorithm. To reach the limit error of 10^{-3} , BGSDC(2,3) with $M=5$ quadrature nodes is the most efficient choice. Boris' method requires more than ten times as many evaluations to deliver the same accuracy.

3.1.2. Scenario 2: $\varepsilon \sim 10^{-2}$

Further from the adiabatic limit we do not have an analytical solution for either the trajectory or the magnetic field at the reflection point. Therefore, we rely on a reference solution computed numerically with a very small time step. Simulation parameters are summarised in Table 1 (middle) and correspond to $\varepsilon \sim 10^{-2}$.

Convergence order. Numerically computed convergence rates for simulations with time steps from $\Delta t_0 \omega = 0.0015625$ to $\Delta t_N \omega = 0.4$ are shown in Fig. 3 for $M=3$ nodes (left) and $M=5$ nodes (right). While we only analyse the final error in the x component of a particle's final position, results for the other position components or velocities are similar and can be generated using the published code. Both variants of Boris achieve their theoretically expected order of $p=2$ for resolutions below $\Delta t \omega < 10^{-2}$, that is approximately 100 steps per gyro-period. BGSDC with $M=3$ nodes and (1,2) and (1,3) iterations achieve the fourth order accuracy of the underlying collocation solution for $\Delta t \omega_B < 10^{-1}$. BGSDC(1,1) requires a slightly smaller time step to show order $p=4$. For $M=5$, BGSDC(1,1) and BGSDC(2,1) both converge with order $p \approx 5$. This is due to having only a single Picard iteration to adjust for the nonlinearity. Using (1,3) iterations gives order $p=7$ while (2,3) delivers the theoretical convergence order of $p=8$ of the underlying collocation solution. Although the more complex interplay between GMRES and Picard iterations does not allow a simple heuristic like two orders per iteration that was found for non-accelerated Boris-SDC [18], these results show that BGSDC can deliver high orders of convergence by changing the runtime parameter M and $(K_{\text{gmres}}, K_{\text{picard}})$.

Work-precision analysis. We compare the strength of the magnetic field at the reflection point against the values delivered by a reference simulation with $\Delta t \omega = 0.005$ using standard Boris-SDC with $M=5$ and 6 iterations. Fig. 4 shows the resulting error $\sigma[B]$ against the total number of \mathbf{f} evaluations for the non-staggered Boris and BGSDC with $M=3$ (left) and $M=5$ (right) quadrature nodes with varying numbers of iterations. For $M=3$, all BGSDC variants converge with order $p=4$, in line with the order of the underlying collocation method. Increasing the number of iterations improves accuracy when keeping the time step Δt fixed, but this does not offset the additional computational work. Throughout, BGSDC(1,1) is slightly more efficient than the other BGSDC variants. To achieve errors of 10^{-1} and below, BGSDC is more efficient than Boris. It delivers a fixed accuracy with fewer \mathbf{f} evaluations or delivers a smaller error with the same amount of computational effort. For errors of 10^{-3} , the reduction in computational effort is about a factor of ten. For $M=5$, we observe higher convergence orders for BGSDC(2,1) and BGSDC(2,3), indicated by the steeper slopes. Using (2,3) iterations delivers the most efficient method for errors below 10^{-5} while both BGSDC(2,1) and BGSDC(2,3) are about equally effective for errors up to 10^{-1} . Only for errors above 0.1 does the Boris algorithm become competitive. Note that staggered Boris gives slightly smaller errors than the non-staggered Boris, but the difference is very small. Only for very large time steps

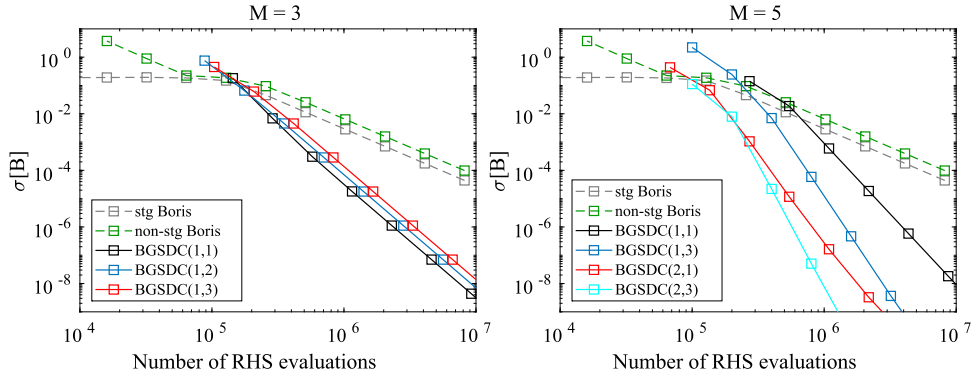


Fig. 4. Error $\sigma[B]$ of the magnetic field value at all reflection points during particle's motion in the magnetic mirror trap as a function of the number of r.h.s. evaluations performed for 3 and 5 Gauss-Lobatto collocation nodes per time step and different number of SDC iterations. The curves for the different runs result from varying the total number of time steps for fixed t_{end} . The classical Boris integrator's (both staggered and non-staggered) convergence is shown for comparison.

does a substantial difference emerge. The error for staggered Boris remains bounded at roughly 0.1 while the error for non-staggered Boris continues to increase.

Long-time energy error. Boris' method conserves phase-space volume [5] which typically means a bounded long-term energy error. For Boris-SDC and BGSDC, depending on the choice of quadrature nodes, the collocation solution is either symmetric (Gauss-Lobatto) or symplectic (Gauss-Legendre) and will also have bounded long-term energy error, see the discussion in Winkel et al. and references therein. However, for small numbers of iterations, both methods exhibit some energy drift.

Fig. 5 shows the relative error in the total energy over $N_{steps} = 3,840,000$ time steps (with $\Delta t\omega = 0.5$ and $t_{end} = 4800$) for $M = 3$ (left figures) and $M = 5$ (right figures) Gauss-Lobatto nodes and different iteration numbers. The two upper figures show standard Boris-SDC, the lower ones BGSDC. Except for the larger t_{end} , parameters are identical to those used for $\epsilon \sim 10^{-2}$, see Table 1 (middle). As expected, non-staggered Boris shows no drift, however its energy error is quite large at around 5×10^{-2} .

For a small number of iterations, Boris-SDC has not yet recovered the symmetry of the underlying collocation method and shows noticeable energy drift. However, for three iterations, after almost 4 million time steps, the energy error is still smaller than the one from Boris method for both $M = 3$ and $M = 5$ nodes. For five iterations and $M = 3$ nodes the method has converged and the energy error remains constant. For $M = 5$ nodes and three iterations there is drift, but the final energy error is several orders of magnitudes smaller than for Boris. Eleven iterations are required for $M = 5$ nodes for Boris-SDC to recover the bounded energy error from the collocation solution.

The lower two figures show the energy error obtained by BGSDC. Again, some energy drift is observed for small numbers of iterations and BGSDC(1,2) has a larger final energy error than Boris for both $M = 3$ and $M = 5$. BGSDC requires fewer iterations than Boris-SDC to recover the bounded energy from the collocation solution. For $M = 3$, (2,3) iterations are enough (compared to five full sweeps with standard Boris-SDC) while for $M = 5$ (3,6) iterations suffice, compared to 11 full sweeps. Although Picard iterations and Boris-SDC sweeps both need $M - 1$ evaluations of \mathbf{f} , Picard iterations don't require application of the preconditioner $(\mathbf{I} - \mathbf{Q}_\Delta \mathbf{F})$ and will thus be computationally cheaper in terms of runtime. Therefore, BGSDC delivers a smaller energy error for less computational work than Boris-SDC.

3.2. Sulev'ev equilibrium

As a second test case we consider the Sulev'ev equilibrium [36,37] with an added simple radial electric field. The magnetic field is given by

$$\begin{cases} B_R = -\left(\frac{2\tilde{y}}{\sigma^2}\right)(1 - 0.25\epsilon^2) \frac{(1 + \kappa\epsilon\tilde{x}(2 + \epsilon\tilde{x}))}{\psi R} \\ B_Z = 4(1 + \epsilon\tilde{x}) \frac{(\tilde{x} - 0.5\epsilon(1 - \tilde{x}^2) + (1 - 0.25\epsilon^2)\tilde{y}^2\kappa \frac{\epsilon}{\sigma^2})}{\psi R((r_{ma} - r_{mi})/z_0)} \\ B_\phi = \frac{B_\phi^0}{R} \end{cases} \quad (37)$$

Here, $\sigma, \epsilon, \kappa, \psi, r_{ma}, r_{mi}, z_m, B_\phi^0$ are constants given in Table 2, chosen to model an equilibrium similar to the one in the Joint European Torus (JET) fusion reactor.³ Furthermore, \tilde{x}, \tilde{y} are the intermediate coordinates

³ We thank Dr Rob Akers from Culham Centre for Fusion Energy for providing this test case and the parameter to model the JET equilibrium.

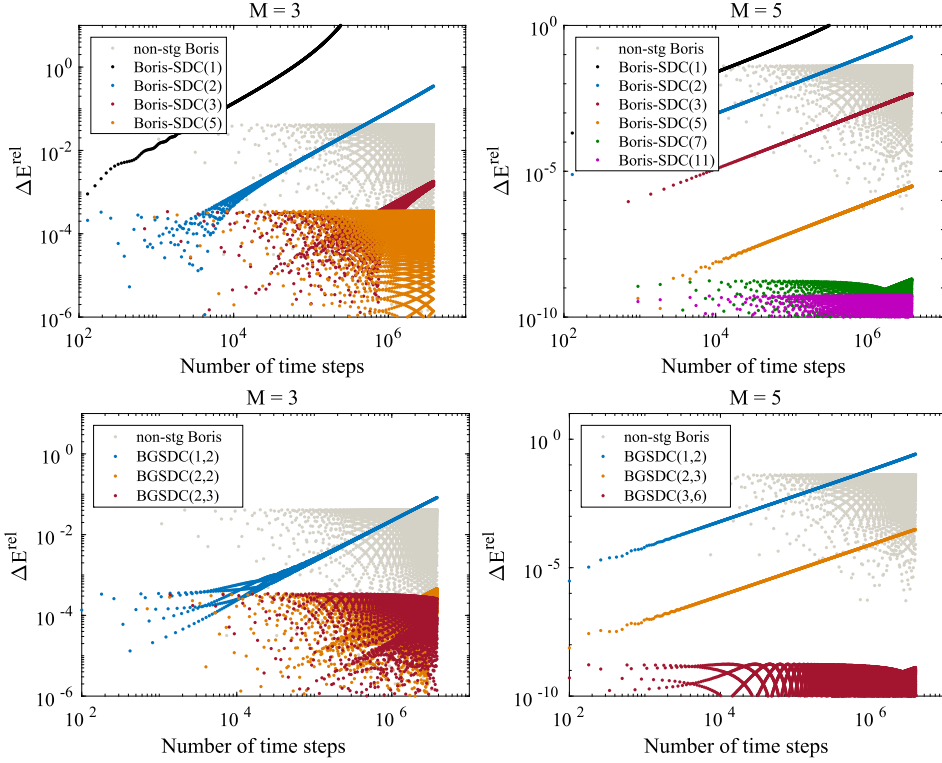


Fig. 5. Relative energy error for Boris-SDC (top) and BGSDC (bottom) over 4 million time steps with two time steps per gyro radius ($\Delta t\omega = 0.5$) for $M = 3$ (left) and $M = 5$ (right) Gauss-Lobatto collocation nodes.

Table 2

Parameters needed to reconstruct the magnetic field.

σ	1.46387369075
ϵ	0.22615668214
κ	1.43320389205
ψ	1.13333149039 [T ⁻¹ m ⁻¹]
r_{ma}	3.83120489000 [m]
r_{mi}	1.96085203000 [m]
Z_m	0.30397316800 [m]
Z_0	1.0 [m]
B_ϕ^0	-9.96056843000 [Tm]
E_0	50000 [V/m]
r_a	1.5 [m]
R_0	3.00045800 [m]
Z_0	0.30397317 [m]

$$\tilde{x} = 2 \frac{(R - r_{mi})}{(r_{ma} - r_{mi})} - 1, \quad (38)$$

$$\tilde{y} = (Z - Z_m)/Z_0.$$

The radial electric field $E_r = E_0 r^2 / r_a^2$ is given in toroidal coordinates (r, ϕ, θ) which are connected to a cylindrical system via $R = R_0 + r \cos(\theta)$, $Z = Z_0 + r \sin(\theta)$. Constants E_0, r_a, R_0, Z_0 are given in Table 2.

Fig. 6 shows the magnetic surfaces in a R - Z cross-section as well as the two example trajectories studied below. One is for a passing particle that continues to perform full revolutions in the reactor's magnetic field. The second is a trapped particle which changes direction at some point of its orbit. It thus fails to complete a full revolution and instead travels on a so-called "banana-orbit". Initial position and velocity for both the passing and the trapped particle are given in Table 3.

3.2.1. Accuracy

To assess accuracy of Boris and BGSDC, we compare their particle trajectories against a reference trajectory computed using BGSDC(2,4) $M = 5$ with a time step of $\Delta t = 0.1$ ns. We choose time steps such that the time points in every run are

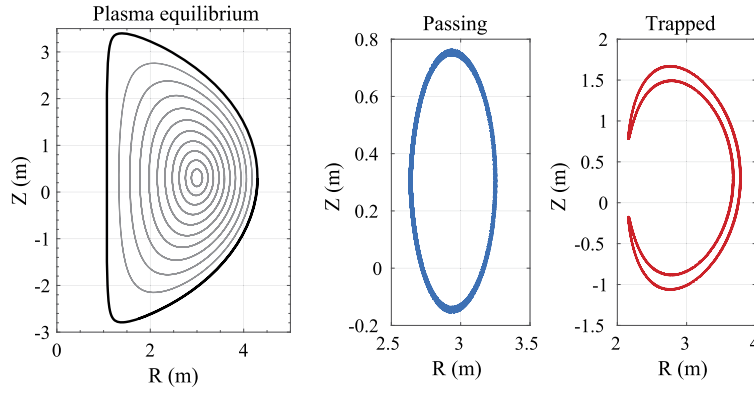


Fig. 6. Magnetic surfaces for JET-like equilibria with parameters from Table 3 (left) and example trajectories of a passing (middle) and trapped particle (right).

Table 3

Initial position and velocity for a passing and trapped particle in the Solov'ev equilibrium. The charge-to-mass ratio is $\alpha = 47918787.60368$, gyrofrequency $\omega_0 = 0.159 \cdot 10^{-9}$ on the magnetic axis (R_0, Z_0) and their trajectories are simulated until $t_{\text{end}} = 10$ ms.

	Passing	Trapped
x	2.1889641172761	3.0852639552352
y	0	0
z	0.8635434778595	-0.0732997600262
v_x	2269604.3143406	814158.31065935
v_y	292264.06108651	931354.18390575
v_z	-338526.06660893	1793580.5493877

a subset of the time points of the reference to avoid the need for interpolation. The maximum defect in x at all points of the computed trajectory reads

$$d_x := \max_{n=0, \dots, N} \left| x_n - x^{(\text{ref})}(t_n) \right|, \quad (39)$$

with $t_n, n = 0, \dots, N$ being the time steps for the current resolution and $x^{(\text{ref})}(t_n)$ the reference solution at those points. Analogous expressions are used to compute d_y and d_z and we then take the maximum. Note that (39) compares positions at a specific time so that d_{max} measures not only particle drift but also errors in phase. Table 4 shows the resulting trajectory errors for the passing particle (upper two) and trapped particle (lower two) for staggered Boris and BGSDC with $M = 3$ and $M = 5$ nodes.

Passing particle. For a passing particle, if precision in the range of millimetres is required, BGSDC(2,6) with $M = 5$ nodes can deliver this with a 1 ns time step. In contrast, staggered Boris has an error of around 40 cm even with a 0.1 ns step. Assuming it is converging with its theoretical order of two, staggered Boris would require a time step of $0.1/\sqrt{0.4362/0.00159}$ or approximately 0.006 ns to be as accurate as BGSDC. Therefore, it requires $1 \text{ ns}/0.006 \text{ ns} \approx 167$ times as many steps as BGSDC to deliver the same accuracy. However, BGSDC(2,6) with $M = 5$ nodes requires $5 + 4 + 4 * 6 + 4 = 37$ f-evaluations per time step according to (26) whereas Boris needs only one. Nevertheless, BGSDC(2,6) with $\Delta t = 1$ ns will be around $167/37 \approx 4.5$ times faster than staggered Boris with $\Delta t = 0.006$ ns while delivering the same accuracy.

If only centimetre precision is required, BGSDC without parallelisation will struggle to be faster than staggered Boris. BGSDC(2,6) with $M = 3$ nodes achieves an error of about 6.6 cm for a time step of 0.5 ns. For the same accuracy, staggered Boris would require a time step of around $0.1 \text{ ns}/\sqrt{0.4362/0.0661} \approx 0.04$ ns. Therefore, staggered Boris needs about 13 times as many steps as BGSDC, but BGSDC would be around $3 + 2 + 2 * 6 + 2 = 19$ times more expensive per step, thus making it slower. The parallel BGSDC(1,3) with workload model (27) would be competitive as it is only $3 + 1 + 3 + 1 = 8$ times more expensive per step.

Trapped particle. BGSDC(2,6) with $M = 5$ nodes can deliver micrometre precision with a 1 ns time step. Staggered Boris has an error of 2.76 cm for a 0.1 ns step and would require approximately a $0.1 \text{ ns}/\sqrt{0.0276/0.0011} \approx 0.02$ ns time step to be comparable in precision to BGSDC. Note that we again assume that staggered Boris converges with its full second order accuracy, even though the reduction in error from 0.2 ns to 0.1 ns time step suggests that this is not yet the case. Staggered Boris therefore needs at least $1/0.02 = 50$ times more steps than BGSDC while every step of BGSDC(2,6) is 37 times more expensive. Therefore, we expect BGSDC to be at least $50/37 \approx 1.3$ times faster than staggered Boris.

Table 4

Maximum deviation from reference trajectory d_{\max} in m for a passing (upper two) and trapped (lower two) particle in a Solev'ev equilibrium. Note that the values for the staggered Boris method in the two upper and two lower tables are identical, because the number of nodes M does not affect it.

Trajectory error for passing particle for $M = 3$ nodes.			
Δt	Staggered Boris	BGSDC(1,3)	BGSDC(2,6)
0.1 ns	0.4362	–	–
0.2 ns	1.7210	–	–
0.5 ns	6.3286	0.2876	0.0661
1 ns	6.3089	6.3189	0.9593
2 ns	6.4776	6.4562	4.3799
Trajectory error for passing particle for $M = 5$ nodes.			
Δt	Staggered Boris	BGSDC(1,4)	BGSDC(2,6)
0.1 ns	0.4362	–	–
0.2 ns	1.7210	–	–
0.5 ns	6.3286	0.02351848	0.00000049
1 ns	6.3089	2.03153722	0.00158861
2 ns	6.4776	6.49396896	0.42631203
Trajectory error for trapped particle for $M = 3$ nodes.			
Δt	Staggered Boris	BGSDC(1,3)	BGSDC(2,6)
0.1 ns	0.0276	–	–
0.2 ns	0.0353	–	–
0.5 ns	0.1721	0.1212	0.0027
1 ns	0.6976	6.7962	0.0275
2 ns	2.7529	7.3273	5.7878
Trajectory error for trapped particle for $M = 5$ nodes.			
Δt	Staggered Boris	BGSDC(1,4)	BGSDC(2,6)
0.1 ns	0.0276	–	–
0.2 ns	0.0353	–	–
0.5 ns	0.1721	0.01407600	0.00000022
1 ns	0.6976	1.07024660	0.00109417
2 ns	2.7529	6.96382918	0.53245320

Centimetre precision can be delivered by BGSDC(2,6) with $M = 3$ nodes and a 1 ns time step or by staggered Boris with a time step of 0.1 ns. Thus, staggered Boris requires only about 10 times as many steps but BGSDC(2,6) is 19 times more expensive per step, making it slower. Even the parallel BGSDC versions is 11 times more expensive so that additional improvements are required for it to be competitive.

3.2.2. Long-time energy error

Fig. 7 shows the long-time energy error for various configurations of BGSDC for the Solev'ev test case. Although formally the collocation formulation underlying BGSDC is symmetric (because we use Gauss-Lobatto nodes), accumulation of round-off error still causes energy drift, a well documented problem of methods that rely on iterative solvers [38]. However, the growth is relatively mild and energy errors are typically small, even after millions of steps, if the number of iterations is sufficiently high. Results are similar for the passing and trapped particle. BGSDC(2,6) with $M = 3$ nodes and a time step of 1 ns delivers a final error of around 10^{-5} for the passing and 10^{-4} for the trapped case.

4. Conclusions and future work

The paper introduces Boris-GMRES-SDC (BGSDC), a new high order algorithm to numerically solve the Lorentz equations based on the widely-used Boris method. BGSDC relies on a combination of spectral deferred corrections for second order problems and a GMRES-based convergence accelerator originally devised for first order problems. Since it freezes the magnetic field over the GMRES iterations to linearise the collocation problem, its applicability is limited to cases where the magnetic field does not change substantially over the course of one time step. Parts of the introduced algorithm are amenable to parallelisation, opening up a possibility to introduce some degree of parallelism in time across the method (following the classification by Gear [39]), but this is left for future work.

The new algorithm is compared against the standard Boris method for two problems, a magnetic mirror and a Solev'ev equilibrium, the latter resembling the magnetic field of the JET experimental fusion reactor at the Culham Centre for Fusion Energy. For the Solev'ev equilibrium, our examples show that if precisions in the millimetre range are required, BGSDC can reduce computational effort by factors of up to 4 compared to the standard Boris method. Gains will be greater for even smaller accuracies and will decrease if less accuracy is needed. While the break even point from where BGSDC cannot

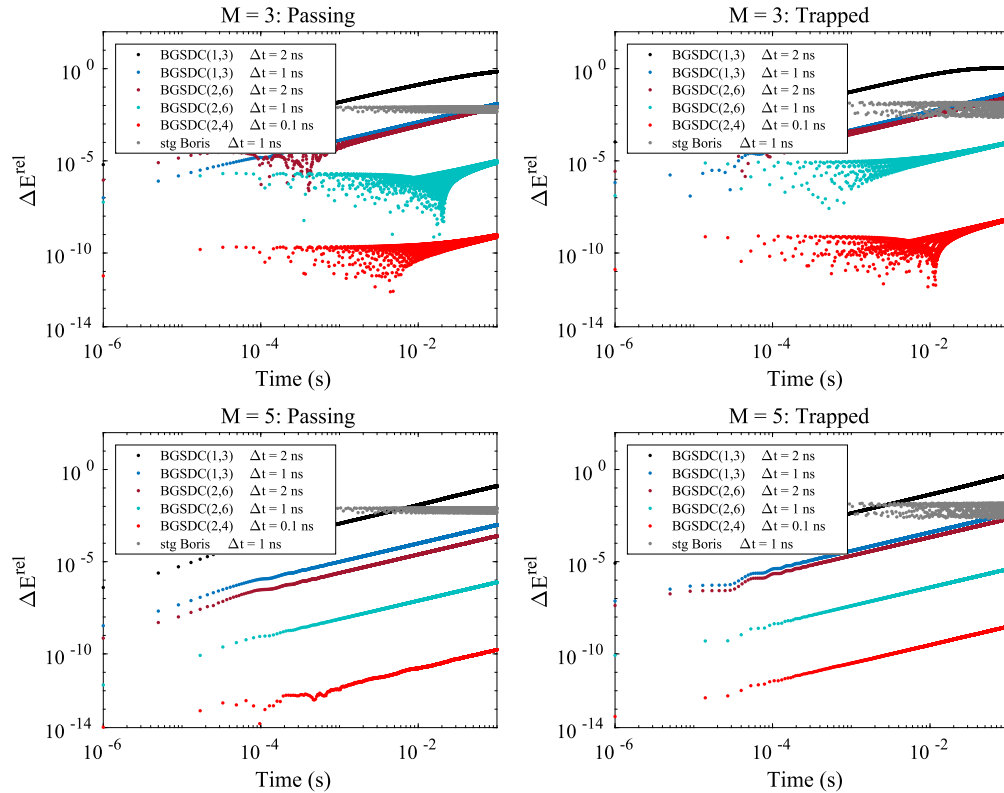


Fig. 7. Energy error for the passing (left) and trapped (right) particle for different configurations of BGSDC for simulations until $t_{\text{end}} = 100$ ms.

produce computational gains is hard to pinpoint, our results suggest it to be for precisions in the centimetre range. A properly parallelised implementation of BGSDC together with an effective adoption of the parameter optimisation strategy by Weiser [31] may still outperform the classical Boris method but is left for future work.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

This work was supported by the Engineering and Physical Sciences Research Council EPSRC under grant EP/P02372X/1 "A new algorithm to track fast ions in fusion reactors". The authors thankfully acknowledge the very helpful input from Dr Rob Akers and James Buchanan from the Culham Centre for Fusion Energy and Dr Jonathan Smith from Tech-X UK Ltd.

References

- [1] P. Gibbon, *Short Pulse Laser Interactions With Matter*, Imperial College Press, London, 2005.
- [2] J.-L. Vay, D.P. Grote, R.H. Cohen, A. Friedman, Novel methods in the particle-in-cell accelerator code-framework warp, *Comput. Sci. Discov.* 5 (1) (2012) 014019, <https://doi.org/10.1109/TNS.2015.2513752>.
- [3] R. Akers, L. Appel, P. Carolan, N. Conway, G. Counsell, M. Cox, S. Gee, M. Gryaznevich, R. Martin, A. Morris, M. Nightingale, A. Sykes, M. Mironov, M. Walsh, Neutral beam heating in the START spherical tokamak, *Nucl. Fusion* 42 (2) (2002) 122, <https://doi.org/10.1088/0029-5515/42/2/302>.
- [4] J.P. Boris, Relativistic plasma simulation-optimization of a hybrid code, in: *Proceedings of the Fourth Conference on Numerical Simulation of Plasmas*, Naval Research Laboratory, Washington, DC, 1970, pp. 3–67.
- [5] H. Qin, S. Zhang, J. Xiao, J. Liu, Y. Sun, M. Tang, Why is Boris algorithm so good? *Phys. Plasmas* 20 (8) (2013), <https://doi.org/10.1063/1.4818428>.
- [6] E. Hairer, C. Lubich, Energy behaviour of the Boris method for charged-particle dynamics, *BIT Numer. Math.* 58 (4) (2018) 969–979, <https://doi.org/10.1007/s10543-018-0713-1>.
- [7] W.W. Heidbrink, Basic physics of alfvén instabilities driven by energetic particles in toroidally confined plasmas, *Phys. Plasmas* 15 (5) (2008) 055501, <https://doi.org/10.1063/1.2838239>.
- [8] H. Yoshida, Construction of higher order symplectic integrators, *Phys. Lett. A* 150 (5) (1990) 262–268, [https://doi.org/10.1016/0375-9601\(90\)90092-3](https://doi.org/10.1016/0375-9601(90)90092-3).
- [9] Y. He, Y. Sun, J. Liu, H. Qin, Volume-preserving algorithms for charged particle dynamics, *J. Comput. Phys.* 281 (2015) 135–147, <https://doi.org/10.1016/j.jcp.2014.10.032>.

- [10] M. Quandt, C. Munz, R. Schneider, A high order relativistic particle push method for PIC simulations, in: The 30th International Electric Propulsion Conference, 2007, pp. 1–9, <http://bibliothek.fzk.de/zb/veroeff/69220.pdf>.
- [11] M. Tao, Explicit high-order symplectic integrators for charged particles in general electromagnetic fields, *J. Comput. Phys.* 327 (2016) 245–251, <https://doi.org/10.1016/j.jcp.2016.09.047>.
- [12] Y. He, Y. Sun, J. Liu, H. Qin, Higher order volume-preserving schemes for charged particle dynamics, *J. Comput. Phys.* 305 (2016) 172–184, <https://doi.org/10.1016/j.jcp.2015.10.032>.
- [13] E. Hairer, C. Lubich, Symmetric multistep methods for charged-particle dynamics, *J. Comput. Math.* 3 (2017) 205–218, <https://doi.org/10.5802/smai-jcm.25>.
- [14] T. Umeda, A three-step Boris integrator for Lorentz force equation of charged particles, *Comput. Phys. Commun.* 228 (2018) 1–4, <https://doi.org/10.1016/j.cpc.2018.03.019>.
- [15] A. Dutt, L. Greengard, V. Rokhlin, Spectral deferred correction methods for ordinary differential equations, *BIT Numer. Math.* 40 (2) (2000) 241–266, <https://doi.org/10.1023/A:1022338906936>.
- [16] K. Tretiak, D. Ruprecht, Py_Boris-GMRES-SDC: Python Boris-GMRES-SDC, <https://doi.org/10.5281/zenodo.3245381>, 2018.
- [17] K. Tretiak, D. Ruprecht, F90_GMRES-SDC: Fortran Boris-GMRES-SDC, <https://doi.org/10.5281/zenodo.3245360>, 2018.
- [18] M. Winkel, R. Speck, D. Ruprecht, A high-order Boris integrator, *J. Comput. Phys.* 295 (2015) 456–474, <https://doi.org/10.1016/j.jcp.2015.04.022>.
- [19] J. Huang, J. Jia, M. Minion, Accelerating the convergence of spectral deferred correction methods, *J. Comput. Phys.* 214 (2) (2006) 633–656, <https://doi.org/10.1016/j.jcp.2005.10.004>.
- [20] S. Kim, D. Na, J. Lee, M. You, H.-S. Kim, Y. Hwang, T. Hahm, Y.-S. Na, *Nucl. Fusion* 56 (10) (2016) 106006, <https://doi.org/10.1088/0029-5515/56/10/106006>.
- [21] R.J. Akers, S. Äkäslompolo, B. Colling, J. Hess, Y. Liu, S.D. Pinches, K. Särkimäki, M. Singh, A. Turner, J. Varje, High fidelity simulations of fast ion power flux driven by 3D field perturbations on ITER, in: 26th IAEA Fusion Energy Conference - IAEA CN-234, 2016.
- [22] A.K. Mazur, Common molecular dynamics algorithms revisited: accuracy and optimal time steps of Störmer-leapfrog integrators, *J. Comput. Phys.* 136 (2) (1997) 354–365, <https://doi.org/10.1006/jcph.1997.5740>.
- [23] C. Birdsall, A. Langdon, *Plasma Physics Via Computer Simulation*, McGraw-Hill, New York, 1985.
- [24] E. Hairer, C. Lubich, G. Wanner, *Geometric Numerical Integration: Structure-Preserving Algorithms for Ordinary Differential Equations*, Springer Verlag Berlin Heidelberg New York, 2002.
- [25] E. Hairer, S.P. Nørsett, G. Wanner, *Solving Ordinary Differential Equations I: Nonstiff Problems*, 2nd edition, Springer-Verlag Berlin Heidelberg, 1993.
- [26] E. Hairer, G. Wanner, *Solving Ordinary Differential Equations II: Stiff Problems*, Springer-Verlag, Berlin Heidelberg, 1996.
- [27] E. Hairer, R.I. McLachlan, A. Razakarivony, Achieving Brouwer's law with implicit Runge-Kutta methods, *BIT Numer. Math.* 48 (2) (2008) 231–243, <https://doi.org/10.1007/s10543-008-0170-3>.
- [28] J. Huang, J. Jia, M. Minion, Arbitrary order Krylov deferred correction methods for differential algebraic equations, *J. Comput. Phys.* 221 (2) (2007) 739–760, <https://doi.org/10.1016/j.jcp.2006.06.040>.
- [29] E. Jones, T. Oliphant, P. Peterson, et al., *SciPy: Open source scientific tools for Python*, online; accessed 4/9/2018 (2001–), <http://www.scipy.org/>.
- [30] C.T. Kelley, *Iterative Methods for Linear and Nonlinear Equations*, Society for Industrial and Applied Mathematics, 1995.
- [31] M. Weiser, Faster SDC convergence on non-equidistant grids by DIRK sweeps, *BIT Numer. Math.* 55 (4) (2014) 1219–1241, <https://doi.org/10.1007/s10543-014-0540-y>.
- [32] R. Speck, Parallelizing spectral deferred corrections across the method, *Comput. Vis. Sci.* 19 (3–4) (2018) 75–83, <https://doi.org/10.1007/s00791-018-0298-x>.
- [33] J.C. Simpson, J.E. Lane, C.D. Immer, R.C. Youngquist, Simple Analytic Expressions for the Magnetic Field of a Circular Current Loop, *Tech. rep.*, NASA, 2001.
- [34] G. Budker, Thermonuclear reactions in a system with magnetic stoppers, and the problem of direct transformation of thermonuclear energy into electrical energy, in: M.A. Leontovich (Ed.), *Plasma Physics and the Problem of Controlled Thermonuclear Reactions*, Pergamon Press, New York, 1959, pp. 1–33.
- [35] R.F. Post, Summary of UCRL Pyrotron (mirror machine) program, in: *Proc. of the Second U. N. Int. Conference on the Peaceful Uses of Atomic Energy*, vol. 32, United Nations Publication, Geneva, 1958, pp. 245–256.
- [36] L.S. Solovév, The theory of hydromagnetic stability of toroidal plasma configurations, *Sov. Phys. JETP* 26 (2) (1968) 400–407.
- [37] S.B. Zheng, A.J. Wootton, E.R. Solano, Analytical tokamak equilibrium for shaped plasmas, *Phys. Plasmas* 3 (3) (1996) 1176–1178.
- [38] E. Hairer, R.I. McLachlan, A. Razakarivony, Achieving Brouwer's law with implicit Runge-Kutta methods, *BIT Numer. Math.* 48 (2) (2008) 231–243, <https://doi.org/10.1007/s10543-008-0170-3>.
- [39] C.W. Gear, Parallel methods for ordinary differential equations, *Calcolo* 25 (1–2) (1988) 1–20, <https://doi.org/10.1007/BF02575744>.