GPU-ACCELERATED SPIN DYNAMICS AND ANALYSIS FOR RHIC*

Dan T. Abell†, Dominic Meiser, Tech-X Corp., Boulder, CO, USA
Desmond P. Barber, DESY, Hamburg, Germany
Mei Bai, Vahid H. Ranjbar, Brookhaven National Laboratory, Upton, NY, USA

Abstract

The elucidation of nucleon spin structure benefits from highly polarized beams in storage rings. Accurate and fast spin dynamics simulations are a valuable tool for optimizing beam polarization. We describe the integrators and the performance of a new spin-orbit tracking code. These integrators use quaternions and Romberg quadratures to accelerate both the computation and the convergence of spin rotations. We exploit the inherently data-parallel nature of spin tracking to accelerate our algorithms on graphics processing units.

INTRODUCTION

Elucidating the origin of nucleon spin [1] is the principal focus of polarized beam experiments at the Relativistic Heavy Ion Collider (RHIC) at Brookhaven National Lab [2]. Because statistical uncertainties scale inversely with the polarization squared [3], optimizing the polarization is essential for efficient use of experimental resources.

Computer simulations serve an important rôle in understanding and improving beam polarization. For example, the invariant spin field (ISF) places an important upper bound on the maximum attainable polarization of a stored beam [4], and finding the ISF requires fast and accurate spin tracking. Other motivations for fast and accurate spin tracking include the computationally challenging assessment of proposed storage-ring-based searches for a permanent electric dipole moment (EDM) in protons and deuterons [5]. Assessing the sensitivity of such experiments will require long-term spin-orbit simulations of unprecedented accuracy [6].

We have developed a very accurate and efficient spin-orbit tracking code, TEASPINK, now built on top of the UAL framework [7]. Here we present a brief description of that code and our error analysis, with a particular focus on the spin integration. As numerical efficiency is an important consideration for spin tracking codes, we have implemented all our integrators on a graphics processing unit (GPU). The embarrassingly parallel nature of spin-orbit tracking (in the absence of space-charge) makes this type of computation an ideal fit for the highly parallel architecture of GPUs. This aspect of our work, along with more in-depth information about the integrators, will be presented elsewhere [8].

* This work is supported by the US Department of Energy, Office of Science, Office of Nuclear Physics, including grant No. DE-SC0004432.
† dabell@txcorp.com

SPIN INTEGRATION

In this study, we ignore the effects of synchrotron radiation and space-charge forces. We therefore model the orbital dynamics in an accelerator using a single-particle Hamiltonian appropriate to the externally applied magnetic and electric fields of a particle accelerator. For the spin dynamics, we treat spin-½ particles with the spin orientation described by a unit three vector that obeys the Thomas-BMT equation [9]:

$$\mathrm{d}\vec{\mathcal{S}}/\mathrm{d}s = \vec{\Omega} \times \vec{\mathcal{S}}. \tag{1}$$

Here the precession vector $\vec{\Omega}$ is a function of the particle velocity and the local magnetic and electric fields.

Because we have found that the accuracy of the orbital data has a significant impact on the accuracy of the spin tracking, our code is based on first performing very accurate symplectic integration for the orbital motion [10]. With orbital data in hand, we then use the Thomas-BMT equation to integrate the spin motion. The most significant aspect of this work is that we have found a means of accelerating the convergence of spin integration.

Piece-wise Constant Spin Precession

A commonly used integration strategy for spin treats the fields and velocity vectors in $\vec{\Omega}$ as constant throughout a slice of length Δs. One then integrates Eq. 1 to the form

$$\vec{\mathcal{S}}(s + \Delta s) = \mathbf{R}(\vec{\omega}) \cdot \vec{\mathcal{S}}(s), \tag{2a}$$

where $\mathbf{R}(\vec{\omega})$ denotes the 3×3 matrix that describes rotation about axis $\vec{\omega}$ by angle $|\vec{\omega}|$. One approximates this rotation vector as Δs times $\vec{\Omega}$. Then to compute the spin rotation across a whole element, one simply multiplies the contributions from each slice. For four slices, say, one thus transports an initial spin $\vec{\mathcal{S}}^i$ to a final spin $\vec{\mathcal{S}}^f$ according to

$$\vec{\mathcal{S}}^f = \mathbf{R}(\vec{\omega}_4) \cdot \mathbf{R}(\vec{\omega}_3) \cdot \mathbf{R}(\vec{\omega}_2) \cdot \mathbf{R}(\vec{\omega}_1) \cdot \vec{\mathcal{S}}^i \tag{2b}$$

There are two sources of errors in the above approach to spin integration. First, errors in the orbital data feed into the spin integration via errors in the fields and velocities needed in $\vec{\Omega}$. Second, treating the rotation axis as piece-wise constant introduces errors that arise from the non-commutativity of spin rotations around non-parallel axes.

When using drift-kick integrators for orbital motion, the first source of error usually dominates. Then increasing the number of orbital slices to improve the orbital accuracy automatically diminishes the magnitude of the second type

ISBN 978-3-95450-122-9

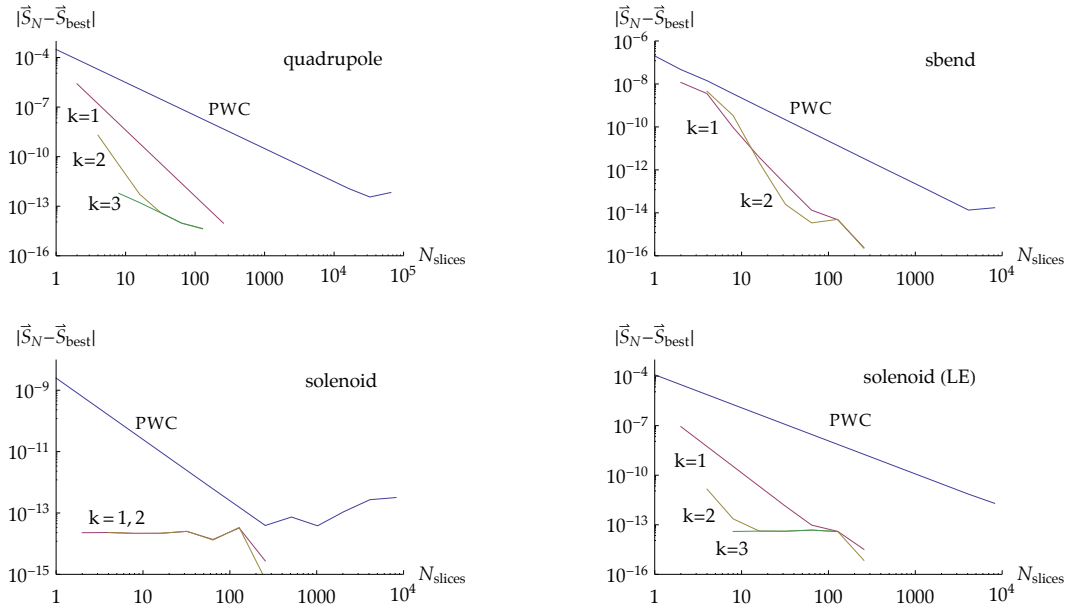

Figure 1: Absolute error in the computed spin versus number of slices for a proton crossing several different beamline elements. These graphics show results for a 200 GeV p^+ crossing a RHIC IR quad (upper left), a RHIC arc dipole (upper right), and a 2.1 m, 1.3 T solenoid (lower left). The lower right graphic show the result for a 25 GeV p^+ traversing the same solenoid. The different curves correspond to different methods of integrating the spin; see text for details.

of error. The situation changes when we use more accurate bend-kick and matrix-kick integrators [10]. These integrators allow us to take such large steps through both dipoles and quadrupoles that the lack of commutativity between consecutive spin rotations can now become an issue.

To speed the accumulation of spin rotations across a set of orbital slices, we use quaternions to represent rotations [11]. Compared to matrix multiplication, this saves a factor of about two in the required arithmetic operations.

Romberg Quadratures for Spin Precession

The spin precession across a beamline element is not, in reality, piece-wise constant; rather, within an element it varies smoothly with s. More significantly, the integrators based on Eq. 2 exhibit second-order convergence. This quadratic convergence suggests the use of some accelerating technique to cancel the errors. In particular, we have applied a Romberg approach [12, 13] to spin integration, and the improvement is dramatic.

Instead of computing the rotation vectors at the middle of each slice, we now, for this new approach, compute them at the edges of each slice. We then accumulate the spin precession as in Eq. 2. For the first and last edges, however, we use a half-step; *i.e.* we replace Δs by $\Delta s/2$. This is akin to using the trapezoidal rule for integration. We do this using N slices, with N a multiple of some power of two.

During the orbital integration, we record $\vec{\Omega}$ at the edge of each slice. Then by keeping every other $\vec{\Omega}$, or every fourth, *etc.*, we can approximate the net spin precession using a range of additional step-sizes related to the original by powers of two. We thus compute a set of net quaternions:

$$Q(h),\ Q(h/2),\ Q(h/2^2),\ Q(h/2^3),\ \ldots\ .$$

We then compute the *Romberg limit*: First define

$$Q_{0k} = Q(h/2^k). \tag{3a}$$

Then use the rule

$$Q_{j+1,k} = \frac{4^{j+1}Q_{j,k} - Q_{j,k-1}}{4^{j+1} - 1} \tag{3b}$$

to construct the Romberg table:

$$\begin{array}{llll} Q_{00} & & & \\ Q_{01} & Q_{11} & & \\ Q_{02} & Q_{12} & Q_{22} & \\ Q_{03} & Q_{13} & Q_{23} & Q_{33} \end{array} \tag{3c}$$

This table may have more or fewer rows than indicated here, but the number at the bottom right is the Romberg limit of the initial data given in the first column.

When integrating a well-behaved function over a finite interval, the trapezoidal rule plus Romberg limit performs remarkably well with modest computational effort. The efficiency, however, derives from the structure of the error term seen in the Euler-Maclaurin summation formula [13], and the manner in which the Romberg table cancels those errors. Here, on the other hand, we have a *product* of quaternions, and hence no *a priori* reason to suspect that the above will actually work. We tried it on a lark.

ISBN 978-3-95450-122-9

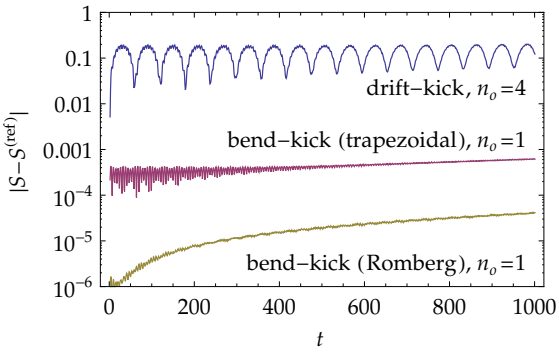

Figure 2: Accumulated spin error as a function of turn number, for multi-turn tracking in a RHIC lattice, using different methods of integration. From top to bottom: the previous standard, a TEAPOT split with $n_o = 4$ (blue), bend-kick with $n_o = 1$ and no Romberg (red), and bend-kick with $n_o = 1$ and two Romberg steps in the IR quads (tan). The simulated beam had an emittance of 45π mm mrad.

PERFORMANCE OF INTEGRATORS

Single-element Errors

The graphics in Fig. 1 show the spin integration errors that we computed when tracking a particular particle across several different beamline elements. In each of those graphics, the blue curve, labelled PWC, corresponds to piece-wise constant spin integration, where for the orbital motion we are now using the more accurate bend-kick and matrix kick integrators. In all cases the slope −2 reveals second-order convergence for PWC spin integration. The remaining curves correspond to k Romberg iterations (*i.e.* Q_{kk}) applied to the PWC data obtained for the given number of slices. Since we do not know the result of exact spin integration, we have estimated the spin integration error as the absolute difference between the result $\vec{\mathcal{S}}^N$ obtained using N slices and what we considered our "best" result, $\vec{\mathcal{S}}_{\text{best}}$.

For $\vec{\mathcal{S}}_{\text{best}}$, in Fig. 1, we have used results obtained using 256 slices and three Romberg iterations. If we instead use our most finely-sliced PWC result, then small details in these graphics change, but the overall implication holds—that one or more iterations of the Romberg procedure can dramatically reduce the errors made by PWC spin integration.

Evolution of Spin Errors

We also examined how spin errors evolve over many turns in RHIC. In Fig. 2 we use the accumulated spin error to compare different methods of integration. The upper curve shows results obtained using the previous standard, a TEAPOT split with $n_o = 4$: four slices for most elements, but 16 for the strong elements in the interaction regions (IRs). We obtained the lower two curves using the new integrators with just one orbital slice for most elements, and four for those in the IRs. For the middle curve, we used just the 'trapezoidal' rule (Fig. 1's PWC). For the lower curve, we computed a Romberg limit using a maximum k of 2.

CONCLUSION

The results shown in Fig. 1 indicate that our method of applying Romberg quadratures to spin integration can yield impressive gains in accuracy and speed. When using four slices to integrate across the quadrupole, we saw that applying two Romberg iterations yields a four-decade reduction in the error—to a level that requires some four-hundred slices using just PWC spin integration. For the sector bend, we see a less dramatic absolute reduction in the error; but even there, between 8 and 64 slices, the slope −4 on the $k = 1$ curve tells us that one Romberg iteration converts second-order PWC integration to fourth-order.

Figure 2, which shows results obtained from multi-turn tracking of RHIC, indicates that the benefits of taking a Romberg limit apply over a broad range of phase space.

More details of these results, and a broader and more in-depth set of results, will be presented elsewhere [8].

REFERENCES

[1] C.A. Aidala, S.D. Bass, D. Hasch, and G.K. Mallot, *Rev. Modern Phys.*, 85(2):655, 2013.

[2] M. Harrison, S.G. Peggs, and T. Roser *Annu. Rev. Nucl. Part. Sci.*, 52:425, 2002. M. Bai, *Eur. Phys. J. Special Topics*, 162:181, 2008. H. Huang, L. Ahrens, I.G. Alekseev, *et al.*, *IPAC'2011*, 1888.

[3] W. Fischer and A. Bazilevsky, *Phys. Rev. ST Accel. Beams*, 15(4):041001, 2012.

[4] K.A. Heinemann and G.H. Hoffstätter, *Phys. Rev. E*, 54(4):4240, 1996. G.H. Hoffstaetter, *High-Energy Polarized Proton Beams: A Modern View*, Springer, New York, NY, 2006.

[5] C.J.G. Onderwater, *J. Phys.: Conf. Ser.*, 295:012008, 2011.

[6] A. Lehrach, *ICAP'2012*, 7.

[7] N. Malitsky, R. Talman, M. Blaskiewicz, *et al.*, *PAC'2003*, 272, 2003.

[8] D.T. Abell, D. Meiser, and D.P. Barber, in preparation. D.T. Abell, D. Meiser, M. Bai, V.H. Ranjbar, and D.P. Barber, in preparation.

[9] L.H. Thomas, *Philos. Mag. S 7*, 3(13):1, 1927. V. Bargmann, L. Michel, and V.L. Telegdi, *Phys. Rev. Lett.*, 2(10):435, 1959.

[10] É. Forest, *Beam Dynamics: A New Attitude and Framework*, Harwood Academic Publishers, 1998. É. Forest, *J. Phys. A: Math. Gen.*, 39(19):5321, 2006.

[11] L.C. Biedenharn and J.D. Louck, *Angular Momentum in Quantum Physics: Theory and Applications*, vol. 8 of *Encyclopedia of Mathematics and its Applications*, Addison-Wesley Publishing Co., 1981.

[12] F.S. Acton, *Real Computing Made Real: Preventing Errors in Scientific and Engineering Calculations*, Princeton University Press, 1996.

[13] W.H. Press, S.A. Teukolsky, W.T. Vetterling, and B.P. Flannery, *Numerical Recipes in C++: The Art of Scientific Computing*, Cambridge University Press, second edition, 2002.

ISBN 978-3-95450-122-9