SAND2022-14260C

DSMC Collisions under Special Relativity

William J McDoniel,* Christopher H Moore, and Keith L Cartwright
Sandia National Laboratories, Albuquerque, NM, USA
(Dated: October 14, 2022)

This paper details a modification to the no time counter collision scheme of the Direct Simulation Monte Carlo (DSMC) method that extends its validity to relativistic particle velocities. Several challenges arise when extending non-relativistic DSMC to the relativistic regime that will be discussed. Some implications of special relativity are non-obvious in their impact on kinetic theory and can easily be overlooked. Numerical issues can introduce significant error (and non-conservation of momentum or energy) for a naive implementation of the physics into code. A non-trivial amount of algebra is required to take equations from the form one might see in a physics textbook to something suitable for numerical implementation in DSMC. Our goal is to offer a "cookbook" that lays out the relativistic algorithm as clearly as possible so that a non-expert in special relativity can take an existing program for computing DSMC collisions and extend it to relativistic speeds without much effort. We detail the changes required, offer some explanation of why they are necessary, and references are included that offer more detail on the physics for those interested. We implement this in Sandia's plasma simulation code EMPIRE and briefly show some results.

I. INTRODUCTION

Sandia National Laboratory's EMPIRE code [1] is used for a variety of particle-in-cell (PIC) simulations with Direct Simulation Monte Carlo (DSMC) collisions [2], often involving high-energy electron beams which are interacting (ionizing and scattering) with a neutral gas. Especially when electric and magnetic fields are considered, electrons in this and similar applications achieve relativistic speeds. Beyond high-energy plasmas, relativistic conditions might arise at the enormous scales present in astrophysical flows and especially astrophysical plasmas, to which DSMC has also been applied [3][4] and to laser-driven wakefields [5]. These applications also have in common that a specialized relativistic-only algorithm that breaks down at low energies is inappropriate, since non-relativistic collisions remain frequent. Many simulations appear to still use non-relativistic physics regardless, and there does not appear to be a single, easy-to-follow reference that an expert in DSMC but not necessarily in special relativity can take and apply to their simulation code. When we set out to extend our existing implemention of no time counter [6] DSMC collisions with variable-weight particles [7] to handle these energy regimes in a physically accurate way, we found that, while the actual physics involved is well-understood, its application to DSMC is non-obvious. Therefore we present here the important modifications that must be made to textbook DSMC as described in, e.g., [8] to extend it to relativistic energies without loss of capability at low energies, hopefully saving others from having to deal with the same stumbling blocks we encountered. This paper is intended for an audience familiar with DSMC who is aware of but not expert in special relativity. This was the position we found ourselves in, and we apologize to physicists for any imprecision or error in our discussions of special relativity.

It is important to note that we are only concerned here with the collisions themselves. The movement and acceleration of particles is beyond the scope of this paper, but fortunately is described well elsewhere. In EMPIRE we use the Boris push [9].

In Section II, we first discuss some general issues in computing collisions across such huge energy ranges. Then we break down textbook DSMC and describe what small modifications need to be made to each part of the algorithm. In Section III, we demonstrate that our algorithm can correctly produce a relativistic equilibrium and explain why this is actually a little harder to do than one might think, and also show that we can correctly simulate a transient relaxation even when the distribution's bulk velocity is very relativistic. Finally, we show some results from simulations demonstrating our use of the algorithm in real applications.

*Corresponding author: wmcdoni@sandia.gov

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.

II. ALGORITHM

Now we proceed through the steps of the textbook DSMC algorithm and show how they need to be modified. At a medium level of abstraction, these steps are: finding the number of pairs to sample, sampling a random pair, testing it for a collision, and then performing a collision. Special relativity impacts each of these differently. There is no modification necessary to sampling random pairs, and so we will not address it. Finding the number of pairs to sample and testing each pair require many of the same underlying modifications to calculating the product of the relative velocity and cross-section. Actually performing the collision is a fairly involved step.

Here is a list of modifications to be made to the algorithm. This is our "cookbook", a simple checklist of what needs to be changed, with pointers to the later equations that give the details of the implementation.

- Keep track of particles' proper velocities $\boldsymbol{u} = \gamma \boldsymbol{v}$ (instead of tracking $\boldsymbol{v}$, Section II A).
- Determining a number of pairs to sample and testing pairs for collision:
  - Use the Møller flux (Section II B 1) instead of the relative velocity (Equation 4).
  - Ensure the cross-section is a function of the collision energy rather than relative speed. (Section II B 2)
  - Use the relativistic collision energy (Equation 10).
- Performing a collision:
  - Find the center of momentum reference frame $\boldsymbol{u}_{CM}$ (Equation 12).
  - Shift the particles' proper velocities to and from that frame with a Lorentz transform (Equation 13).
  - For inelastic collisions, find the post-collision speeds of the particles from the total energy (Equation 16).
- Use a relativistic temperature calculation where temperature is desired (Equation 20).

A. Numerical Cancellation and Proper Velocity

A fundamental problem commonly encountered when doing numerical calculations across a huge range of energies is that computers do finite-precision math. According to special relativity, no particle in any reference frame can go faster than the speed of light, even if accelerated by an arbitrary force or shifted to another reference frame itself moving in the other direction at high speed. Ultimately, close to c, tiny absolute differences in velocity can imply enormous differences in energy (or in velocity in another reference frame). A standard way to get around this for relativistic calculations is to track the *proper velocity* of particles rather than their velocity $\boldsymbol{v}$, defined as $\boldsymbol{u} = \gamma(v)\boldsymbol{v}$, where $\gamma = \frac{1}{\sqrt{1-v^2/c^2}}$ is the Lorentz factor. Proper velocity therefore behaves exactly like velocity at low speeds where $\gamma \approx 1$ and goes to infinity as $v \to c$. In EMPIRE we now only store $\boldsymbol{u}$ on the particles and convert back to $\boldsymbol{v}$ as needed (which is rare). Note that physicists often use w or sometimes η for proper velocity, but we already use w for particle weight. As well, the Boris push that we use for moving the particles also operates on $\boldsymbol{u}$, so using it was desirable anyway.

Beyond the storage of velocity, some calculations end up very imprecise at certain energies. Often, as with the Lorentz factor at $v \approx c$, the fundamental problem is in taking the difference between two numbers which are very close in value. For example, in special relativity the kinetic energy of a particle is given by $E_K = (\gamma - 1)mc^2$ (this is the total energy minus the familiar rest mass energy mc^2). At moderate speeds there is no issue here. At high speeds, lack of precision in v will be a problem for computing $\gamma(v)$. However, using u, we can rewrite and get $\gamma(u) = \sqrt{1+u^2/c^2}$, which is well-behaved. There remains a problem at low speeds even when tracking u, because $\gamma \approx 1$ and we care about $\gamma - 1$. With some algebra we can rewrite and obtain:

$$E_K = \frac{mu^2}{1+\sqrt{1+u^2/c^2}} \tag{1}$$

Equation 1 gives a precise value for a particle's kinetic energy at any proper speed u, including at low speeds where it goes to the familiar $\frac{1}{2}mv^2$. For comparison, for a slow object with $m = 1$ kg and $v = 10$ m/s, attempting to directly compute $(\gamma(v) - 1)mc^2$ using python gives a kinetic energy of 59.952 J rather than the non-relativistic value of 50. Using Equation 1 instead we obtain 50.00000000000004. This problem of numerical cancellation arises frequently, but in almost all cases the operations required for DSMC can be recast to be suitable for all energies.

B. Sampling and Testing Pairs

Textbook DSMC begins by computing the number of pairs to be sampled:

$$N_{pairs,A\neq B} = \frac{N_A N_B w \Delta t}{V}(v_r\sigma)_{max} \tag{2}$$

$$N_{pairs,A=B} = \frac{N_A(N_A-1) w \Delta t}{2V}(v_r\sigma)_{max} \tag{3}$$

Where N is the number of particles of some species in an element, w is the particles' weight, Δt is the timestep, V is the element volume, v_r is a relative speed, and σ is the cross-section (generally a function of v_r). The quantity $(v_r\sigma)_{max}$ denotes the maximum value of the product of relative speed and cross-section over all of the possible pairs of particles in the element.

1. Møller Flux

One might think that much of this will change at relativistic speeds. We know that "length contraction" and "time dilation" occur, and perhaps that should affect what we use for the timestep or element volume. However, we can do most of the calculation in the lab frame of reference, where the timestep size and element volume are just what we expect them to be. In fact the collision rate is a *Lorentz-invariant* – every observer should agree on the number of collisions that a particle is expected to undergo as it crosses from one wall of a system to the other. The trick is defining $(v_r\sigma)_{max}$ in such a way that the whole expression is Lorentz-invariant. This is not simple because the usual "moving particle sweeping out a cylinder against a stationary background" argument for the collision rate depends on shifting to a particular reference frame, and that becomes more complicated. This is fortunately a solved problem in relativistic particle physics. The cross-section can be used as-is, although the cross-section itself should of course use a model appropriate for relativistic energies and developed with special relativity in mind if high-energy collisions are expected. The idea is that the cross-section *in the frame of the collision* is Lorentz-invariant, even if it might change in other reference frames. However, for the relative speed we have to introduce the Møller flux factor [10]. The physical meaning of this is a little contentious even among physicists (e.g., [11]), but for our purposes it is sufficient to say that instead of computing $v_r = |\boldsymbol{v_2} - \boldsymbol{v_1}|$, we take:

$$v_r = \sqrt{(\boldsymbol{v_2} - \boldsymbol{v_1})^2 - \frac{(\boldsymbol{v_1} \times \boldsymbol{v_2})^2}{c^2}} \tag{4}$$

Where $\times$ indicates a vector cross product and we recover $\boldsymbol{v}$ from our stored $\boldsymbol{u}$ by dividing by $\gamma(u) = \sqrt{u^2/c^2+1}$. This reduces to the ordinary relative speed when the particles have parallel velocities and can attain a value of nearly $2c$ when the particles are moving towards each other in the lab frame, each at nearly c. We can determine a number of pairs to sample using Equations 2 and 3 in the ordinary way, except that we use Equation 4 for the relative speed.

2. Cross-Sections

While it is permissible to use the cross-section in the frame of the collision, it is vital that the cross-section used be a function of the collision energy. This is true in the non-relativistic case as well, but when we talk about cross-section models we often talk about them as being functions of the relative speed. This works out in the non-relativistic case because the collision energy is a function of the relative speed, but this is no longer true at relativistic speeds. Instead we need to compute the kinetic energy between the particles in the frame of the collision and use that to find a value for σ. In general it is pretty easy to convert functions of non-relativistic speed to energy so that low-energy capability can be maintained while also extending the cross-section to relativistic speeds. For example, for Maxwell molecules where $\sigma = \sigma_{ref}/v_r$, write $v_r = \sqrt{2E_{CM}/m_r}$, which is just the non-relativistic kinetic energy formula where m_r is the reduced mass. Substitute this into σ so that it is a function of energy, but then evaluate E_{CM} relativistically (Equation 10):

$$\sigma_{Maxwell} = \frac{\sigma_{ref}}{\sqrt{2E_{CM}/m_r}} \tag{5}$$

When relativistic effects can be neglected, this gives the expected result, but continues to be a function of energy at high speeds too.

3. Collision Energy

Computing the relativistic kinetic energy of a pair of particles in their center-of-momentum (CM) frame is conceptually simple: we find the CM frame, shift the particle velocities to it, and then find their kinetic energies in that frame (Equation 1). But this is a lot of unnecessary work for a collision that we might not even accept, when all we want is the kinetic energy. With a lot of algebra, we can make use of the energy-momentum relation:

$$E^2 = (\boldsymbol{p}c)^2 + (mc^2)^2 \tag{6}$$

Where $E = \gamma mc^2$ and $\boldsymbol{p} = \gamma m\boldsymbol{v} = m\boldsymbol{u}$ is the relativistic momentum. Because the quantity mc^2 is Lorentz-invariant, so is $E^2 - (\boldsymbol{p}c)^2$ (this is the square of the Minkowski norm of the four-momentum). In the lab frame, we can evaluate this as the sum of contributions from the individual particles. In the CM frame, where $\boldsymbol{p} = 0$ for the two-particle system and $\gamma = 1$, we have only $E^2 = (Mc^2)^2$. M is the rest mass of the two-particle system, and is not equal to $m_1 + m_2$ because it also captures the system's internal energy, which is the kinetic energy of the individual particles in the CM frame. In principle a particle's internal energy, such as that stored in vibration or excitation, should be included in its individual rest mass m_1 or m_2, but this adjustment should be very small compared to its ground state rest mass. We set these invariant quantities equal to each other:

$$(E_1 + E_2)^2 - (\boldsymbol{p}_1 c + \boldsymbol{p}_2 c)^2 = M^2 c^4 \tag{7}$$

Our ultimate goal is to find the kinetic energy in the CM frame as $E_{CM} = (M - m_1 - m_2)c^2$, which is that portion of the system's rest energy not attributable to the rest masses of the individual particles. After a lot of manipulation to get there and then to deal with the numerical cancellation issues discussed above, we obtain:

$$Y = \frac{u_1^2 u_2^2}{c^2} + u_1^2 + u_2^2 \tag{8}$$

$$Z = \frac{2m_1 m_2}{(m_1 + m_2)^2 c^2} \left(\frac{Y}{1 + \sqrt{Y/c^2 + 1}} - \boldsymbol{u}_1 \cdot \boldsymbol{u}_2 \right) \tag{9}$$

$$E_{CM} = \frac{Z}{1 + \sqrt{Z + 1}} (m_1 + m_2)c^2 \tag{10}$$

Which gives the kinetic energy in the CM frame as a function of the particles' rest masses and proper velocities in the lab frame (or in any other frame). This kinetic energy can then be used to determine a cross-section. Note that the result for Z has the potential for some numerical cancellation still, when the two particles have nearly identical large velocities – if both particles have speed $\approx c$ in the same direction then the parenthetical is $c^2 - c^2 \approx 0$. If this is very important then doing the full reference frame shift is likely a better option, but for our applications it is rare to have an important cross-section which is very sensitive to small changes in kinetic energy at low energies between particles which are moving extremely quickly (this would be a cold but fast beam where it is also important to accurately capture its internal collisions).

C. Performing a Collision

Once a collision is accepted, the process again looks conceptually similar to textbook DSMC. We want to rotate a relative velocity vector according to a scattering distribution and apply it back to the particles. The main complication is that it is now much harder to move between reference frames.

The collision occurs in the CM frame of the two particles. In the non-relativistic case, this doesn't mean much, because rotating a vector is the same regardless of where that vector is in velocity space. For a relativistic collision, however, this is not so. The set of points in (proper) velocity space where the two particles conserve momentum and energy undergoes a complicated transformation between the lab frame and the CM frame, and the scattering distribution deforms too. Rather than attempting to grapple with this, it is conceptually and mathematically simpler to transform the particle velocities to the CM frame, perform the collision in that frame, and then transform their velocities back to the lab frame.

1. Reference Frame Transformations

First we need to find the CM frame. The CM velocity looks a lot like the non-relativistic result, substituting γm for m, and this and the expressions for shifting reference frames below are largely taken from [12].

$$\boldsymbol{v}_{CM} = \frac{m_1 u_1 + m_2 u_2}{\gamma_1 m_1 + \gamma_2 m_2} \tag{11}$$

This is going to become imprecise near c, so instead we seek a proper velocity corresponding to the CM frame. We take $\boldsymbol{u}_{CM} = \boldsymbol{v}_{CM}/\gamma_{CM}$, and observe that we can again use the energy-momentum relation to simplify the result to:

$$\boldsymbol{u}_{CM} = \frac{m_1 u_1 + m_2 u_2}{M} = \frac{m_1 u_1 + m_2 u_2}{E_{CM}/c^2 + m_1 + m_2} \tag{12}$$

Where M is again the rest mass of the two-particle system, which we can compute directly from Equation 7 (this will be more precise) or else use our previous result for E_{CM} (which can have trouble with cold, fast pairs).

We can transform a particle's proper velocity in the lab frame, $\boldsymbol{u}$ to its velocity in the CM frame, $\boldsymbol{u'}$, by taking:

$$\boldsymbol{u'} = \boldsymbol{u} + \boldsymbol{u}_{CM}\left(\frac{\boldsymbol{u}_{CM}\cdot\boldsymbol{u}}{c^2(1+\gamma_{CM})} - \gamma\right) \tag{13}$$

Where γ corresponds to the particle's velocity and γ_{CM} to $\boldsymbol{u}_{CM}$. This same formula can also be used to transform a particle's velocity from the CM frame back to the lab frame: using $-\boldsymbol{u}_{CM}$ to move back gives the same result except we add γ instead of subtracting it.

2. Post-Collision Velocities

Now that we have the particles' proper velocities in the CM frame, scattering is simple. The direction of the relative velocity vector is the same as the direction of the relative proper velocity vector $\boldsymbol{u}'_2 - \boldsymbol{u}'_1$, since in this frame the two velocities just point in opposite directions. We scatter this direction exactly as we would in the non-relativistic case.

For an elastic collision, we could now just rotate the particle velocities and be done. However, for an inelastic collision, the particles' energy will change. In general for many-product collisions, after we otherwise determine how much energy a given particle should have, we can find its speed in the CM frame from Equation 1. For two-product collisions, just like in the non-relativistic case, a total kinetic energy and a total momentum (which is 0 in this frame) will fix the particles' speeds. We know that the particles have equal amounts of momentum and that their kinetic energy sums to E_{CM}. Note that re-use of E_{CM} from the earlier calculation of collision energy can produce energy non-conservation because Equation 10 is imprecise for low energy pairs with a large CM velocity. E_{CM} should generally be recomputed from the particles' CM frame velocities using Equations 1 or 10, then adjusted by the heat of reaction, and once we have a good value for the post-collision energy E^*_{CM} we can write:

$$m_1 u'^*_1 = m_2 u'^*_2 \tag{14}$$

$$\frac{E^*_{CM}}{c^2} = (\gamma'^*_2 - 1)m_1 + (\gamma'^*_2 - 1)m_2 \tag{15}$$

It is convenient to define $\alpha = \gamma - 1$. We can solve these equations for α'^*_1 and α'^*_2 and obtain:

$$\alpha'^*_i = \frac{E^*_{CM}(m_j + E^*_{CM}/2)}{m_i(E^*_{CM} + m_i + m_j)} \tag{16}$$

$$u'^*_i = c\sqrt{\alpha'^*_i(\alpha'^*_i + 2)} \tag{17}$$

This gives us the post-collision speeds of the particles in the CM frame at a new total energy E^*_{CM}. We now just apply these speeds to the direction that the particles have scattered into as normal to obtain post-collision velocities in the CM frame. Finally, we transform these velocities back to the lab frame using Equation 13 (remember to change the sign on γ).

3. Temperature

While not a necessary part of the particle simulation, when initializing particles or looking at the macroscopic properties in a simulation, a relativistic temperature calculation is useful. The basic problem here is that the non-relativistic temperature formula breaks down because the Hamiltonian changes in the relativistic limit. Rather than $\overline{E}_K = \frac{3}{2}k_B T$, it becomes $\overline{E}_K = 3k_B T$. This is because the kinetic energy of a particle is proportional to the square of momentum in a non-relativistic regime but merely proportional to momentum at extreme speeds (consider what happens in the energy-momentum relation as rest mass becomes negligible). There does not seem to be an unambiguous definition of relativistic temperature, but for our purposes we can go back to the Hamiltonian and the energy-momentum relation and write:

$$E_K = \sqrt{(\boldsymbol{p}c)^2 + (mc^2)^2} - mc^2 \tag{18}$$

$$3k_B T = \left\langle p_x \frac{\partial E_K}{\partial p_x} \right\rangle + \left\langle p_y \frac{\partial E_K}{\partial p_y} \right\rangle + \left\langle p_z \frac{\partial E_K}{\partial p_z} \right\rangle \tag{19}$$

$$T = \frac{m}{3k_B} \overline{u^2/\gamma} \tag{20}$$

Where the average $\overline{u^2/\gamma}$ is computed in the CM frame. This gives the correct non-relativistic result when $u \approx v$ and $\gamma \approx 1$, and the correct result in the relativistic limit when $\gamma \approx u/c$. This works well to parameterize relativistic initial conditions or to examine relativistic flow. However, note that because the contribution that a particle makes to the temperature is not simply proportional to its kinetic energy, this temperature can change even as energy is conserved while a velocity distribution changes shape.

III. RESULTS AND DISCUSSION

After implementing all of this, we need to verify that it works. We can easily check whatever tests we were previously using to establish confidence in a non-relativistic DSMC code – nothing in this algorithm should be discernably different for typical room temperature flows. For EMPIRE, among other tests we like to use the Bobylev-Krook-Wu relaxation [13] since it features an analytic solution to a time-varying problem. Our test [14] is fairly heavy-duty, with hundreds of thousands of simulations checking convergence with increasing particle count, and as far as it was concerned this new relativistic algorithm made no difference. For typical cases it would take a truly heroic number of particles/simulations to try to find the on-average difference amidst all of the DSMC noise.

Once we're confident that we haven't broken anything in the non-relativistic regime, we need tests for the new capability. We are unaware of something like the analytic BKW relaxation for relativistic massive particles, but the relativistic equilibrium distribution is known. The Maxwell-Jüttner (MJ) distribution is given by:

$$f(\boldsymbol{p}) = \frac{1}{4\pi m^3 c^3 \theta K_2(1/\theta)} e^{-\gamma/\theta} \tag{21}$$

$$f(\gamma) = \frac{\gamma\sqrt{\gamma^2 - 1}}{\theta K_2(1/\theta)} e^{-\gamma/\theta} \tag{22}$$

Where $\theta = kT/mc^2$ and K_2 is the modified Bessel function of the second kind. f can be written either in terms of the three-dimensional momentum vector $\boldsymbol{p}$ or in terms of the Lorentz factor γ, the latter taking advantage of the distribution's spherical symmetry akin to the Maxwell-Boltzmann speed distribution.

At the very least, a relativistic DSMC code should be able to simulate relaxation to this equilibrium distribution (Fig. 1). One might think that this is a pretty weak requirement – we are used to thinking that it is difficult to screw up DSMC badly enough that it does not go to the right equilibrium – but actually this tests almost every aspect of the relativistic algorithm. The particular things that needed to change are important for getting the correct equilibrium.

Like the Maxwell-Boltzmann distribution, Maxwell-Jüttner has the property that the likelihood of sampling a particle with a particular velocity depends only on the energy of the particle, and the joint probability of sampling some set of particles depends only on the sum of their energies. This is one way that the principle of detailed balance works. For an elastic collision, a given pre-collision and post-collision set of velocities are equally likely to be sampled from the distribution since they have the same total energy. And the differential cross-section that controls the probability with which the pre-collision set of velocities scatters into the post-collision set is equal to that for the reverse reaction, because the cross-section is a function of the CM energy and is symmetric.

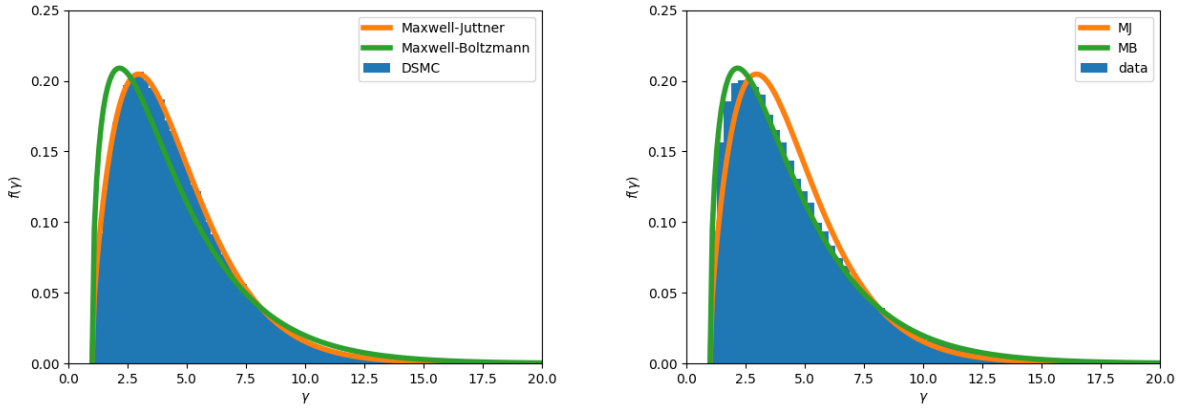

FIG. 1: Steady-state velocity distributions for relativistic relaxations. 100k particles are initially created in two delta functions at $u_x = \pm 2.99777490 \times 10^8$ m/s and are allowed to relax. The left image shows the analytic Maxwell-Jüttner and Maxwell-Boltzmann distributions at the same total energy, and the DSMC result which closely tracks the MJ distribution. The right image illustrates the types of errors that are possible if the algorithm is implemented partially or incorrectly – in this case all pairs of particles are equally likely to collide, which is an impossible cross-section under special relativity and violates detailed balance, whereas in the left image the relativistic Maxwell-like cross-section of Equation 5 is used.

An important and annoying consequence of this is that there is no relativistic cross-section model that gives every pair of particles an equal probability of colliding. We saw earlier with Equation 5 that a Maxwell model won't do this, and in fact there is no model that will because the model must be a function of the collision energy and the collision energy is not a function of the relative velocity (or more accurately the Møller flux) in the lab frame. If you force this to occur in DSMC, you violate detailed balance. We did this at first, thinking we were testing a simple Maxwell model, and we observed that we got the wrong equilibrium distribution. Only after evaluating our cross-sections as functions of relativistic energy did we get the correct result. Likewise, use of the non-relativistic relative velocity instead of the Møller flux yields the wrong equilibrium distribution comparable to the error in Fig. 1(right).

When performing collisions, errors in shifting the particles to the CM frame and back or in determining post-collision speeds from the energy should also show up in a simple relaxation to equilibrium as momentum or energy non-conservation, but this appears to be exact almost to machine precision (12 decimal places in the case of Fig. 1). And so we conclude that the ability to reproduce the correct MJ equilibrium at relativistic energies along with the ability to correctly perform the typical array of non-relativistic DSMC tests is actually a strong indication that the algorithm is correct and implemented correctly.

More sophisticated tests are possible, however. We can take a non-relativistic test like the BKW relaxation and give every particle a relativistic bulk velocity, and see if we still get the correct non-relativistic result once we post-process and transform the results back to the lab frame (Fig. 2). To readily do this requires diagnostics that most DSMC codes probably lack, like the ability to accurately compute a centered fourth moment of the particle velocities even at extreme bulk velocities, but for a demonstration we modify our code so that all particles are shifted by some extreme velocity upon entering the collision kernel and shifted back at the end. We stress again that for computing these cold collisions between very high-velocity particles, it is vital to shift colliding particles to the CM frame and compute their kinetic energy in that frame instead of using Equation 10 directly from the lab frame.

This shows that the algorithm can correctly handle even extreme cases. The $u_x = 10^3 c$ case corresponds to a bulk velocity of about 0.9999995c, while the calculation cares about differences of meters per second in the CM frame for energy conservation. There will always be some extreme where double precision variables fail to be sufficiently precise, but this is valid to quite high γ. For our applications we aimed for validity only up to $\gamma = 100$. This test demonstrates not just conservation of momentum and energy and relaxation to the correct equilibrium, but a capturing of the correct transient distribution (or at least its fourth moment). This is a powerful test because the fourth moment depends on the exact shape of the distribution at a given energy and is very sensitive to incorrect collision rates or various forms of bias.

This algorithm enables us to simulate the full range of energies present in our problems, such as in Fig. 3, where we simulate a relativistic electron beam through a background gas. The beam electrons are relativistic: they are injected at about $0.85c$, and if their energy was interpreted classically they would be superluminal. At the same time, much

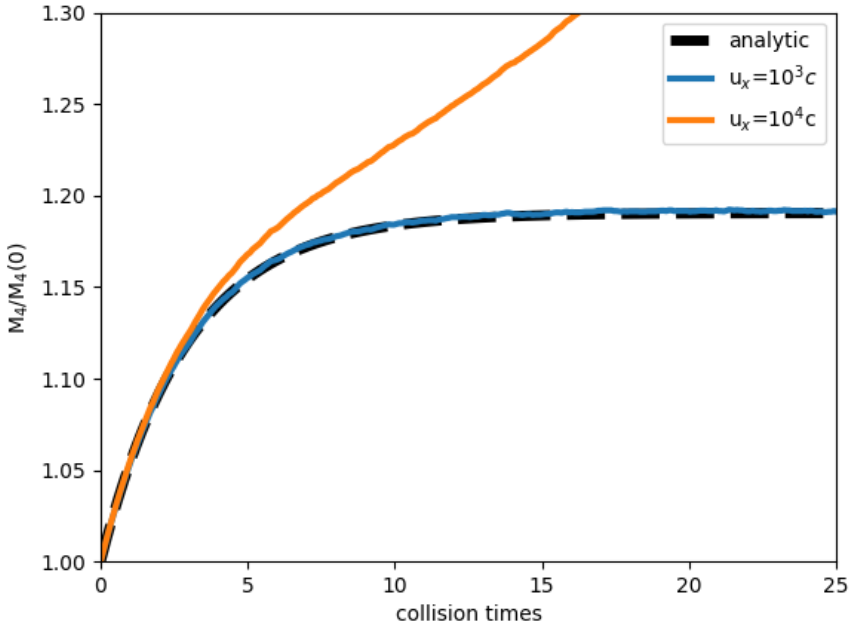

FIG. 2: Shifted BKW relaxations compared to the analytic result. One million particles are created with the BKW initial distribution at 273K. At the start of the simulation, they are shifted in u_x by 10^3c or 10^4c, and then shifted back to obtain output. We obtain excellent agreement for the 10^3c case, with significant error only apparent beyond it. The 10^4c case is simply not conservative, gaining energy over time.

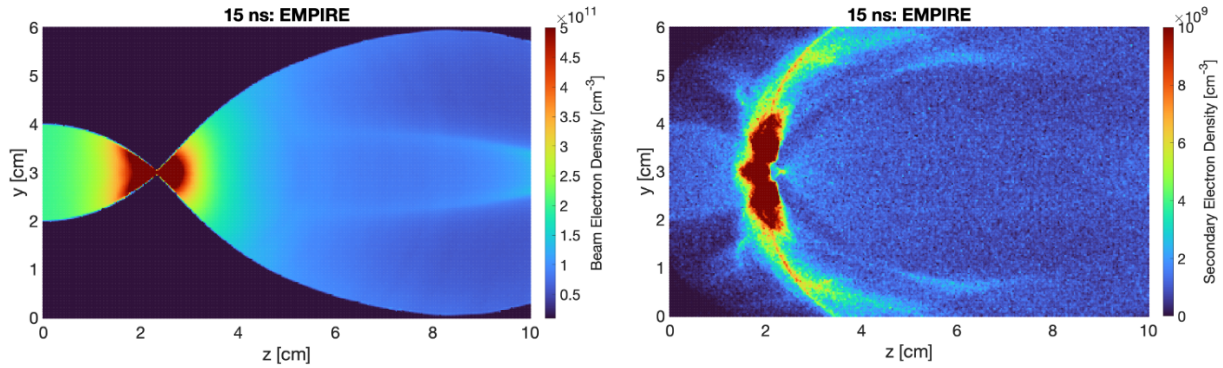

FIG. 3: EMPIRE simulation of $e-$ beam transport. Primary (left) and secondary (right) electron densities for a beam of 0.5 MeV electrons fired from the left boundary into argon at 0.1 mbar in a domain with metal boundaries and no external applied fields.

lower-energy collisions are also possible as the primaries are slowed and as low-energy secondaries are created and accelerated by the self-induced fields. The beam is initially pushed apart by its own repulsive self-force, but ionization leaves behind nearly immobile ions. This positive space charge causes the beam to actually focus and pinch down as seen in Fig. 3(left). The timescale and dynamics of this focusing is sensitive to the collisions and without the relativistic corrections shown in this work we would expect substantially different evolution of the beam.

IV. CONCLUSIONS

We have presented an algorithm for DSMC collisions which takes account of special relativity. It makes use of a single code path, minimizing maintenance and testing costs, and is not significantly more complex in concept or more costly to compute than textbook DSMC. We have attempted to do this in such a way that it is easy for a DSMC expert to modify an existing collision algorithm, without needing to be an expert in special relativity. We have attempted to point out pitfalls, such as the impossibility of a Maxwell-type cross-section, and offered some explanation of non-obvious relativistic principles. We have presented two tests that validate the algorithm and should increase confidence in our implementation, and shown some of our production runs that make use of it.

We undertook this because we had a need for a relativistic collision algorithm. When we surveyed the literature, we found that while various authors speak of using a relativistic collision capability, there was no good, single reference for how to implement one. We hope others in a similar situation will find this useful.

Acknowledgments

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.

This paper describes objective technical results and analysis. Any subjective views or opinions that might be expressed in the paper do not necessarily represent the views of the U.S. Department of Energy or the United States Government.

[1] M. T. Bettencourt, D. A. Brown, K. L. Cartwright, E. C. Cyr, C. A. Glusa, P. T. Lin, S. G. Moore, D. A. McGregor, R. P. Pawlowski, E. G. Phillips, et al., Communications in Computational Physics **30** (2021).
[2] G. Bird, *Molecular Gas Dynamics and the Direct Simulation of Gas Flows* (Claredon Press, Oxyford, 1994).
[3] T. Matsuda, K. Oka, I. Hachisu, and H. M. Boffin, in *New Developments in Computational Fluid Dynamics* (Springer, 2005), pp. 198–205.
[4] M. D. Weinberg, Monthly Notices of the Royal Astronomical Society **438**, 2995 (2014).
[5] D. L. Bruhwiler, R. E. Giacone, J. R. Cary, J. P. Verboncoeur, P. Mardahl, E. Esarey, W. Leemans, and B. Shadwick, Physical Review Special Topics-Accelerators and Beams **4**, 101302 (2001).
[6] K. Nanbu, Journal of the Physical Society of Japan **49**, 2042 (1980).
[7] D. P. Schmidt and C. Rutland, Journal of Computational Physics **164**, 62 (2000).
[8] D. Boyd, Iain and E. Schwartzentruber, Thomas, *Nonequilibrium Gas Dynamics and Molecular Simulation* (Cambridge University Press, 2017).
[9] J. P. Boris, in *Proc. Fourth Conf. Numerical Simulation Plasmas* (Washingtion, D.C., 1970), pp. 3–67.
[10] C. Møller, *General properties of the characteristic matrix in the theory of elementary particles*, vol. 23 (I kommission hos E. Munksgaard, 1945).
[11] M. Cannoni, International Journal of Modern Physics A **32**, 1730002 (2017).
[12] F. Pérez, L. Gremillet, A. Decoster, M. Drouin, and E. Lefebvre, Physics of Plasmas **19**, 083104 (2012).
[13] M. Krook and T. T. Wu, The Physics of Fluids **20**, 1589 (1977).
[14] G. A. Radtke, N. Martin, C. H. Moore, A. Huang, and K. L. Cartwright, Journal of Computational Physics **451**, 110855 (2022).