Accounting for backflow in hydrodynamic-Boltzmann interfaces

Scott Pratt

Department of Physics and Astronomy and National Superconducting Cyclotron Laboratory, Michigan State University,

East Lansing, Michigan 48824, USA

(Received 1 January 2014; revised manuscript received 10 February 2014; published 24 February 2014)

Methods for building a consistent interface between hydrodynamic and Boltzmann modules is presented. These methods account for the backflow across the hydrodynamic-Boltzmann hypersurface. The algorithms are efficient, relatively straight-forward to implement, and exactly satisfy conservation laws across the hypersurface. The methods also account for the interactions between particles in the backflow and other particles by following the subsequent impact of such particles. Since the number of altered trajectories grows exponentially in time, a cutoff is built into the procedure so that the effects of the backflow are ignored beyond a certain number of collisions.

DOI: 10.1103/PhysRevC.89.024910

PACS number(s): 25.75.-q, 24.10.Jv

I. INTRODUCTION AND MOTIVATION

Relativistic viscous hydrodynamics is a popular choice for modeling high-energy heavy-ion collisions. Hydrodynamics is appropriate when collisions are sufficiently rapid to keep the various species moving with a single collective velocity and local kinetic temperature, and to keep the stress-energy tensor sufficiently isotropic to warrant a viscous treatment. However, these conditions are lost near the end of the reaction, when the various hadrons begin to cool separately and move with different collective velocities [1,2]. The final stage of the reaction and the decoupling are then best modeled with a microscopic simulation, which in the limit of many particles, or with a high oversampling, becomes equivalent to a Boltzmann description. Often these simulations are referred to as hadronic cascades.

Particles are emitted into the hadronic cascade through a hypersurface separating the hydrodynamic region from the cascade. If the phase-space density at the boundary is

$$f(\mathbf{p}, \mathbf{r}, t) = \frac{1}{(2\pi)^3} \frac{dN}{d^3 p d^3 r},$$
 (1)

the number of particles emitted into the simulation side of the boundary through a hypersurface element is given by the Cooper–Frye formula [3],

$$dN = \frac{d^3p}{(2\pi)^3} \frac{p \cdot \Omega}{E_p} f(\mathbf{p}).$$
(2)

Here, Ω_{μ} is a small hypersurface element located at position **r** and time *t*. Variations of the Cooper–Frye formula have been applied to numerous hybrid models [4–7]. In each of these approaches the Boltzmann equation is solved by sampling techniques, i.e., rather than storing information for each phase-space element $d^3 p d^3 r / (2\pi)^3$, one follows the evolution of sample particles chosen consistently with the phase-space density. For sampling ratios of unity, this becomes a one-to-one simulation of the hadronic stage, and the simulations are often labeled as a cascade. For higher sampling ratios, the models approach the limit of a Boltzmann equation. Since hadronic cascades are modeling the low-density stage at the end of the collision where velocity gradients are reduced, sampling ratios

of unity are nearly indistinguishable from the Boltzmann limit. This is in contrast to the case of simulating the early partonic stage, where sampling ratios need to be of order 10 or more to approach the Boltzmann limit [8,9]. Whether the sampling ratios are unity or not, one needs to generate particles into the cascade code consistently with the hydrodynamic description at the hypersurface.

For a timelike element, $\Omega^2 > 0$, one can consider the emission from a frame where the emission is simultaneous across the element and Ω_0 represents a volume element undergoing sudden emission. For a spacelike element, one can choose a frame where the surface is stationary. In this frame, Ω_i represents the area of the element multiplied by the time of the emission. Depending on the hypersurface element, Ω_{μ} , and the particle's momentum p^{μ} , the number dNcan be either positive or negative. The positive contribution describes particles being fed into the cascade, whereas the negative contribution represents the backflow, i.e., those particles which leave the cascade region, cross the interface, and enter the hydrodynamic domain. Both the positive and negative contributions are necessary if energy, momentum, and charge are to be conserved across the interface. Similar issues to those being discussed here in the context of a hydrodynamics-cascade interface have also been studied with regards to coupling hydrodynamics directly to the vacuum. This involves ignoring the subsequent interactions of those particles as they decouple from the hydrodynamic medium; see, for example, Ref. [10]. When coupling directly to the vacuum the stress-energy tensor is discontinuous, while with coupling to a Boltzmann code, one can aim for a continuous description if the interface occurs at a sufficiently high density that collision rates justify the hydrodynamic description.

Most hadronic cascades ignore backflow. Such codes apply a step function, $\Theta(p \cdot \Omega)$, to Eq. (2), and do not erase particles from the simulation that flow backward across the boundary. For modeling of particles at midrapidity, it has been shown that the backflow is on the order of one half of one percent [11,12]. However, it has been reported that the error grows to the level of several percent away from midrapidity [11]. The purpose of this short paper is to introduce a method for correcting for backflow. The approach accounts for the backflow through the production and propagation of negative-weight particles. Local conservation laws are exactly conserved in the limit of high sampling. The theoretical underpinnings of this approach are described in the next two sections. A hadronic cascade was altered to incorporate these changes, and a brief evaluation of the behavior is presented in the subsequent section. The Appendix provides a description of the algorithm used to numerically sample the particle flow across the hydrodynamiccascade interface.

II. COOPER-FRYE FORMALISM

In order to apply the Cooper–Frye formalism of Eq. (2) one needs to first generate a list of hypersurface elements Ω_{μ}^{I} , where the index *I* denotes each small element. Formally, Ω_{μ}^{I} can be described as the portion of the hypersurface that falls within a small four-volume element $\int_{I} d^{4}x$. The surface is described by a criteria based on local properties, such as the temperature or density, and can be written parametrically as

$$C(x) = C(T(x), \rho(x)) = 0,$$
 (3)

where the function ρ could be any quantity propagated through the hydrodynamic evolution, such as the particle or energy density that depends on the spacetime coordinate x. For instance, if the interface is chosen to occur at fixed temperature T_0 , the location of the surface would be described by $C(x) = T(x) - T_0 = 0$. The function should be defined such that one is inside the hydrodynamic region if C(x) > 0and outside if C(x) < 0. The hypersurface element is then

$$\Omega^{I}_{\mu} = -\int_{I} d^{4}x \partial_{\mu} \Theta(C(x)) = -\int_{I} d^{4}x \delta(C(x)) \partial_{\mu} C(x), \quad (4)$$

where Θ is a step function. For example, if *C* is defined by a surface of constant temperature, Ω^{I}_{μ} is parallel to $-\partial_{\mu}T$.

For timelike elements, $\Omega^2 > 0$, one can always consider the element in a frame where Ω points purely in the time direction. In this frame the emission is simultaneous across the hypersurface, and Ω_0 represents a small volume element. If $\Omega_0 > 0$ the criteria *C* is falling with time, e.g., the temperature is falling, and particles are being emitted into the cascade region. The majority of the emission in heavy-ion collisions comes from such elements. For $\Omega_0 < 0$, the criteria is rising, and the hydrodynamic region is absorbing the volume element. This can be understood by considering the sign of the $p \cdot \Omega$ term in the Cooper–Frye formula. Having $\Omega_0 < 0$ is unusual in a timelike element, because it is difficult to find an area where the density or temperatures are rising. Even with lumpy initial conditions, the longitudinal expansion is driving the densities downward in regions where two lumps are expanding into one another.

For spacelike elements, $\Omega^2 < 0$, one can always find a boost and rotation which point Ω_{μ} along the positive *x* axis. In this frame the emission surface is temporarily stationary, and the emission *dN* is positive for particles with $p_x > 0$ and negative for $p_x < 0$. If the flow velocity in this frame, u_x , is greater than zero, there will more positive contributions than negative contributions. For explosive collisions this is usually the case. The negative contributions represent those particles flowing into the hydrodynamic region, i.e., the backflow. Both the positive and negative contributions are required to conserve charge. For instance, the charge that travels through an element, from the hydrodynamic to the cascade region between times t_1 and t_2 , is

$$\Delta Q = \sum_{I} \Omega^{I}_{\mu} j^{\mu} = \int_{t_{1} < x_{0} < t_{2}} d^{4}x j^{\mu} \partial_{\mu} \Theta(C(x))$$

$$= -\int_{t_{1} < x_{0} < t_{2}} d^{4}x \Theta(C(x)) \partial \cdot j + \int_{x_{0} = t_{1}} d^{3}x j_{0}(x) \Theta(C(x))$$

$$-\int_{x_{0} = t_{2}} d^{3}x j_{0}(x) \Theta(C(x)).$$
(5)

The first term vanishes from current conservation and the last two terms describe the difference between the net charges in the C(x) > 0 region at the two times. The net current density for all the hadron species *a* can be written in terms of the phase-space density as

$$j^{\mu}(x) = \sum_{a} Q_{a} \int \frac{d^{3}p}{(2\pi)^{3}} \frac{p^{\mu}}{E_{p}} f_{a}(p,x).$$
(6)

Since the contributions from all momenta are required to construct the conserved current density at a position x, one cannot throw away the contributions to the Cooper–Frye formula from those momenta with $p \cdot \Omega < 0$ without violating current conservation. Similar expressions can be derived for the energy and momentum, thus showing that the negative contributions to the Coope—Frye expressions are essential if one wishes to satisfy any of the conservation laws.

There exist numerous algorithms for finding the hypersurface, depending on whether the initial conditions are lumpy or smooth, whether the calculation is one dimensional, two dimensional, or three dimensional, and depending on what sort of mesh is used to model the hydrodynamic expansion. A particularly robust algorithm that works for three-dimensional systems of arbitrary topology is described in Ref. [11].

III. METHODS FOR ACCOUNTING FOR BACKFLOW

The output of hydrodynamic codes is a list of the hypersurface elements Ω^I , their spacetime coordinates, and any additional information required to reconstruct the phase-space density, e.g., the collective flow, temperature, densities, and anisotropies of the stress-energy tensor. For each $d^3 p$ within each Ω^I , one can calculate the probability of creating a particle, dN. One then decides to create a particle with probability dN. An efficient algorithm for doing this is presented in the Appendix. The issue of backflow occurs when dN is negative. In the field of relativistic heavy ions, two approaches have been applied thus far. The first approach is simply to neglect the negative contribution. As mentioned in the introduction, this violates energy, momentum and charge conservation below the one percent level at midrapidity, and higher away from midrapidity.

A second method for handling the backflow has been to erase those particles from the cascade that cross back into the hydrodynamic region. This requires storing the location of the hypersurface from the hydrodynamic module for consideration by the cascade. This method would satisfy the conservation laws as long as the phase-space density is continuous across the boundary. A strong discontinuity would suggest that the hydrodynamic treatment was not justified for the densities or temperatures that determined the interface, or that the viscous corrections to the stress-energy tensor in the hydrodynamic treatment is not consistent with the dynamics of a hadron gas. Choosing a higher interface density or temperature, or better accounting for viscous effects, might solve the problem. For the purpose of this paper, we will assume that the phase-space density is continuous, although with the recommendation that the continuity be better studied in the future. A more daunting problem with this approach comes from the book keeping required to decide whether particles crossed the boundary. Hydrodynamic approaches tend to use very small cells, on the order of 0.1 fm. Particles would typically cross hundreds of such cell boundaries, with a very small fraction of such cell boundaries representing an interface with the hydrodynamic treatment. Although there are a variety of algorithms used in cascades, it would seem that the majority of the numerical cost in performing the cascades might then be applied to tracking the backflow. Given the field's interest in fluctuating (lumpy) initial conditions, the topology of the breakup surface could be complex and could vary event-to-event. This makes it difficult to find a robust algorithm that avoids constantly checking for interface crossings. This approach was applied in Ref. [12] but that was for a smooth, azimuthally symmetric, boost-invariant hydrodynamic treatment, where the hypersurface could be represented by giving the radius of the transition hypersurface as a function of time.

A third scheme is presented here. Instead of tracking particles relative to the interface, this scheme simulates the evolution of particles emitted with the negative weights, $p \cdot \Omega < 0$, from Eq. (2). A weight of negative unity is originally assigned to such particles, whereas a weight of positive unity is assigned to particles emitted with $p \cdot \Omega > 0$. Since the negative-weight particles are a small fraction of the overall particles, the additional numerical cost should be on the order of a few percent. The particles are then evolved through the cascade, but with the products of each collision being assigned weights as described below. When analyzing the final products for their impact on specific observables, each particle would contribute proportional to its weight. For instance, when incrementing a bin used to calculate spectra, a particle with weight of -1 would reduce the bin count by one.

If the phase-space density is continuous across the interface, the particles being scattered from the cascade region back into the hydrodynamic region should exactly cancel the negatively weighted particles emitted according to the Cooper–Frye formula. Even if the phase-space density turns out to be discontinuous, this method will still satisfy all the conservation laws.

It is propagating the weights through the subsequent collisions where the weighting becomes complicated. If a pair of incoming particles each has a weight of positive unity, both incoming particles are removed from the particle list after the collision (a weight of zero is assigned to the incoming particles) and a weight of positive unity is assigned to each of the collision products. However, if the *N* incoming particles have arbitrary weights, $w_{i,1}, \ldots, w_{i,N}$, both the new

outgoing particles and incoming particles need to be assigned nonzero postcollision weights if the conservation laws are to be satisfied. If the new weights assigned to the incoming particles are nonzero, those particles cannot be deleted from the particle list after the collision.

The weights for the M outgoing particles (those that are newly created or are scattered into new momentum states) are labeled $w_{f,b}$, and are all set to the product of the weights of the incoming particles,

$$w_{f,b} = W, \quad b = 1, M,$$
 (7)

$$W \equiv \prod_{a=1}^{N} w_{i,a}.$$
 (8)

The weights for the scattered or newly created particles are all equal to the product of the incoming weights. This ensures that if any of the incoming weights are changed by a factor F, that the weight of each of the outgoing products is changed by the same factor. Instead of being erased, the incoming particles are assigned postcollision weights $w'_{f,a}$,

$$w'_{f,a} = w_{i,a} - W, \quad a = 1, N.$$
 (9)

If all the incoming weights are unity, $w'_{f,a} = 0$ for all a, and the incoming particles can be erased. Assigning nonzero postcollision weights to the incoming particles as prescribed in Eq. (9) ensures that conservation laws are enforced in each collision,

$$\sum_{a=1,N} w_{i,a} Q_a = \sum_{b=1,M} w_{f,b} Q_b + \sum_{a=1,N} w'_{f,a} Q_a, \quad (10)$$

given that charge is conserved in the reaction,

$$\sum_{a=1,N} Q_a = \sum_{b=1,M} Q_b.$$
 (11)

The conserved charge Q could just as easily refer to a continuous variable like momentum or energy.

For the several simple processes involving $2 \rightarrow M$ scatterings or $1 \rightarrow M$ decays, the weights of the outgoing particles are shown in Eqs. (12). Each incoming particle is labeled (p, w), where p refers to all information about the particle and its trajectory, while w is its weight. The outgoing particles are similarly referenced, with q referring to the particle's information. If all weights were unity, the weights could be ignored and the reactions would be described as $p_1, p_2 \rightarrow q_1, \ldots, q_M$. One can understand the weightings for each line of Eq. (12)from Eq. (7). For example, one can consider a case where a positive-weight particle collides with a negative-weight particle, $w_{i,1} = 1, w_{i,2} = -1$. The positive-weight particle should not have scattered since the negative-weight particle should not be there, plus the negative-weight particle is in principle canceling the effect of a spurious positive-weight particle that should have been erased from cascade for having re-entered the hydrodynamic region. Thus, not only should particle 1 not scatter, but one needs to correct for the fact that such particles sometimes scatter spuriously, which means the forward-going track weight needs to go from 1 to 2. The scattered tracks are given negative weight to cancel the effect of particle 1 scattering off a spurious track. Similar arguments follow for each row in the list. If the phase-space density is continuous across the border, the original negative-weight tracks should exactly cancel the positive-weight tracks that would have been erased when re-entering the hydrodynamic region if one were applying the second method described earlier:

$$W \equiv \prod_{a=1}^{N} w_{a},$$

$$(p_{1},w_{1}), \dots, (p_{N},w_{N}) \rightarrow (p_{1},w_{1}-W), \dots, (p_{N},w_{N}-W), (q_{1},W), \dots, (q_{M},W),$$

$$(p_{1},+1) \rightarrow (q_{1},+1), \dots, (q_{M},+1),$$

$$(p_{1},-1) \rightarrow (q_{1},-1), \dots, (q_{M},-1),$$

$$(p_{1},+1), (p_{2},+1) \rightarrow (q_{1},+1), \dots, (q_{M},+1),$$

$$(p_{1},+1), (p_{2},-1) \rightarrow (p_{1},+2), (q_{1},-1), \dots, (q_{M},-1),$$

$$(p_{1},-1), (p_{2},-1) \rightarrow (p_{1},-2), (p_{2},-2), (q_{1},+1), \dots, (q_{M},+1).$$

$$(12)$$

Implementing the procedure described in Eq. (12) turns out to be nontenable due to the growth in the weights of the outgoing products. For example, if two particles with weight 2 collide, the scattered particles would have weights of 4. In practice, this leads to exponentially growing weights as a function of the number of collisions. In addition to the growing weights, the number of trajectories being considered increases since the incoming particles are often reweighted, rather than deleted. For a central Au + Au collision at the highest energy of the Relativistic Heavy Ion Collider (RHIC), final-state weights would often exceed the numerical range of the computer, and the number of tracks would exceed the available memory. Even if the cascade were allowed to finish, the high-weighted tracks would overwhelm the answer with noise. Thus, the procedure needs to be modified so that the effects of scattering with nonunity-weighted particles is regulated.

To limit the growth and associated noise of heavily weighted tracks, the procedure described above was modified. Particles were first divided into two sets: base particles and backflow tracer particles. The base particles all have weights of positive unity. They are created and evolve exactly as if the backflow particles had never been included in the cascade. The set of tracer particles begins as those negative-weighted particles representing the backflow. They also include all the trajectories required to trace the influence of the backflow. These trajectories can have weights of either positive or negative unity. Tracer particles are allowed to interact with base particles, but are not allowed to collide with one another. For this reason the effect of the backflow is handled correctly to linear order in the backflow. Since the fraction of backflow particles is unlikely to exceed a few percent, the error associated with this approximation should be on the order of a tenth of a percent or less. Charge, energy and momentum is conserved in each collision as well as in the generation of the particles through the hydrodynamiccascade boundary.

All of the base particles are assigned weights of positive unity, whereas the tracer particles are allowed to have weights of ± 1 . The weights for scattering processes need to be assigned consistently with those given in Eq. (12). However, because collisions between two tracer particles are neglected, the list of interactions is abbreviated and given in Eq. (13). For a collision between incoming particles whose momenta, position, and charges are labeled by p_1 and p_2 while the outgoing particles are denoted by q_1, \ldots, q_M , the cascade must consider three cases. In the first case both particles are base particles and have weights of positive unity. Such particles are described by the notation (p_i ,+1,base). In the second case, one of the particles is a tracer particle with positive weight, (p_i ,+1,tracer), and in the third case one has a tracer particle of negative weight:

$$(p_{1},+1,\text{base}), (p_{2},+1,\text{base}) \to (q_{1},+1,\text{base}), \dots, (q_{m},+1,\text{base}), (p_{1},+1,\text{base}), (p_{2},+1,\text{tracer}) \to (p_{1},+1,\text{base}), (p_{1},-1,\text{tracer}), (q_{1},+1,\text{tracer}), \dots, (q_{m},+1,\text{tracer}), (13)$$

$$(p_{1},+1,\text{base}), (p_{2},-1,\text{tracer}) \to (p_{1},+1,\text{base}), (p_{1},+1,\text{tracer}), (q_{1},-1,\text{tracer}), \dots, (q_{m},-1,\text{tracer}).$$

For decays, $p \rightarrow q_1, \ldots, q_M$, the outgoing decay products all have the same weights as the incoming particle, and are all either base or all tracer particles depending on whether p is a base or tracer particle. For decays, the original track is deleted.

From inspecting the three reactions in Eq. (13) one can see that the evolution of the base particles is exactly the same as

it would be without the tracer particles, because there is no change to the base particles when a base particle collides with a tracer particles. The tracer particles for the final states are all added into the equations on the right-hand sides of the last two expressions in Eq. (13) to make the reactions consistent with the weights in Eq. (12). For the two reactions involving

tracer particles there are two outgoing particles that differ only by their weights or by whether they are tracer particles. Since having two particles with the same trajectory would cause numerical problems, the second particle is translated randomly in relative rapidity in the case of Bjorken boost invariance or, if boost invariance is not implemented, moved by a small random step in coordinate space.

This procedure correctly reproduces the evolution described by the algorithm described in Eq. (12) to first order in the backflow. However, for the $2 \rightarrow 2$ processes described in Eq. (13) that involve an incoming tracer particle, the number of tracer particles triples from one to three. Thus, if a large number of collisions occur the number or tracer particles quickly grows and the method can become noisy. For that reason, the scattering of tracer particles is curtailed once a tracer particle and its ancestors have suffered a number of collisions N_{max} . This is accomplished by storing a number n_{coll} for each tracer particle. Then, when a incoming tracer particle scatters according to the list in Eq. (13), the number is incremented by one and assigned to all the tracer products of the final state. If $n_{\text{coll}} > N_{\text{max}}$, the tracer particle is no longer allowed to collide.

IV. TESTING THE APPROACH

The method was incorporated into the hadronic cascade B3D [7] and applied to the modeling of a longitudinally boostinvariant system. The cascade evolved hadrons with spatial rapidities between 1 and -1 that emitted from a hydrodynamic code, also described in Ref. [7]. B3D incorporates boost invariance by applying cyclic boundary conditions in longitudinal rapidity. The collisions being modeled were 100A+ 100A GeV Au + Au collisions with centralities of 0%-5%. The parameters for this particular run are not important for the considerations of this study, but were chosen so that the model output was consistent with data from RHIC according to the analysis in Ref. [7]. The hydrodynamic code produced a list of hypersurface elements. Each element was listed along with a description of the element's collective velocity, position, temperature, and the viscous correction to its stress-energy tensor. Using the method described in the Appendix, particles were generated from the elements to be considered in the cascade. Less than one half of one percent of these particles had negative weights. For each proper time τ , the radial location of the hypersurface, $r(\tau, \phi)$, was estimated for each of several azimuthal angles ϕ . This produced a smoothly evolving hypersurface. Some approaches constrain the hypersurface to cell boundaries. The surfaces might then hold steady for a few time steps before suddenly absorbing a cell. These varying approaches should yield the same spectra if the backflow is accounted for correctly, even though the number of additional tracer particles could vary depending on the algorithm used to find the hypersurface.

Figure 1 shows the initial positions and radial velocities of the negative-weight particles being fed into the cascade from 1000 events. Each event typically had \approx 4.5 particles with negative weight being fed into the cascade. The positions of the emitted particles were confined to those times for which the surface moved with subluminar velocities, i.e., the hypersurface elements were spacelike. At later times, the surface collapsed so quickly that its position moved faster than the speed of light. The resulting elements were then time like and no negative-weight particles were emitted. The radial positions of the negative-weight emissions are shown in the left-side panel of Fig. 1, where one can see that such emissions



FIG. 1. (Color online) The position of negative-weight tracer particles at their time of emission from the hydro module for consideration by the cascade module is shown in the left-side panel. Results for 1000 events are displayed. Emission occurs along the breakup hypersurface during early times, but once the collapsing hypersurface falls in faster than the speed of light, emission ceases. The radial velocities of the same particles are displayed in the right-hand panel. Since the surface moves inward, the negative-weight particles all have negative velocities. As the surface falls more quickly, the velocities of those particles crossing the surface with $p \cdot \Omega < 0$ are confined to smaller values.



FIG. 2. (Color online) The net weight of tracer pions in the final state binned by transverse momentum. Since the majority of tracer particles have negative weight, the correction is negative. Results are displayed for several cutoffs on the maximum numbers of tracer scatterings, $N_{\text{max}} = 0$ (solid black line), $N_{\text{max}} = 2$ (red circles), $N_{\text{max}} = 4$ (green squares), $N_{\text{max}} = 6$ (blue downward-pointing triangles) and $N_{\text{max}} = 8$ (cyan upward-pointing triangles). The correction mainly occurs at low relative momentum, and there is little observable difference once $N_{\text{max}} > 2$. Calculations for larger N_{max} are plagued by increasing noise.

cease by 5 fm/*c*, even though many of the positive-weight particles were yet to be emitted. The right-hand side of the panel shows the radial velocities of those same particles. Since the surface always moves inward, only particles with negative radial velocities cross from the cascade-side to the hydrodynamic side of the interface. For any given τ the only emissions were for particles with velocities below the velocity of the hypersurface. Once that velocity exceeds -c, emissions of negative-weight particles ceased.

Accounting for the negative-weight particles lowers the spectra, mainly at low transverse momentum, or p_t . Since the number of particles is so low, the change to the spectra is too small to easily visualize, so instead we show the change to the spectra in Fig. 2 by evaluating

$$\Delta N \equiv \sum_{\text{tracers}i} w_i. \tag{14}$$

Since the tracer particles have mostly negative weight, the change to the spectral distributions are negative. From 10 000 cascade events, ΔN was calculated for p_t bins of width 50 MeV/c and is displayed in Fig. 2 for several cutoffs on the maximum number of collisions a tracer particle is allowed to have, N_{max} . For $N_{\text{max}} = 0$, the negative-weight particles do not scatter. For $N_{\text{max}} = 2$, there is a modest change to the p_t distribution compared to that for $N_{\text{max}} = 0$. Distributions are shown for $N_{\text{max}} = 4$ and 7, but they vary little compared with that for $N_{\text{max}} = 2$. Although the distributions change little, it is easy to see the increase in the noise, even though a large number of events were being simulated.



FIG. 3. (Color online) From a simulation of the central two units of rapidity in a 100A GeV on 100A GeV Au + Au collision, the number of tracer particles, the additional particles required to model the effects of the backflow, is shown in panel (b) as a function of N_{max} , which limits the amount of scattering the tracer particles can undergo. For $N_{\text{max}} = 0$ backflow particles never scatter and only 4.5 particles are required. However, the effect of a scattering spreads over more and more tracer particles as N_{max} is increased, and by the time $N_{\text{max}} = 7$, the number of tracer particles is greater than the number of base particles. Panel (a) shows the net transverse energy carried by tracer particles in the final state. The effect rapidly saturates for large N_{max} . Large N_{max} is numerically more expensive and can introduce more statistical noise into the analysis.

To better analyze the effect of changing N_{max} , we consider the change in the transverse energy of the tracer particles. This is denoted by ΔE_t ,

$$\Delta E_t \equiv \sum_{\text{tracers}i} w_i \sqrt{m_i^2 + p_{i,t}^2}, \qquad (15)$$

and is shown in the upper panel of Fig. 3 as a function of N_{max} . For $N_{\text{max}} = 0$, the reduction to the total transverse energy in the rapidity window is near 2.1 GeV. This is only a few tenths of a percent of the total transverse energy. By increasing N_{max} , that number is increased to near -1.6 GeV. The changes for $N_{\text{max}} > 4$ are too small to overcome the statistical noise, and the change for $N_{\text{max}} > 2$ are also on the order of 150 MeV and are far smaller than errors associated with the approximations with the calculation and are also far smaller than experimental uncertainties in the measurement of transverse energy. Even if the number of original negative-weight particles were increased by an order of magnitude, a setting of $N_{\text{max}} = 2$ would be sufficient.

For higher values of N_{max} , the number of tracer particles increases exponentially. Even though the net weight of tracer particles, those with negative weight minus those with positive weights, is only a half dozen, the total number exceeds the number of regular particles by the time $N_{\text{max}} = 7$, as shown in the lower panel of Fig. 3. This increase is accompanied by a noticeable increase in computational cost. Considering the results from the upper panel of Fig. 3, it would make sense to cut off N_{max} at lower values. Given that most particles collide approximately 3 times, it is surprising that the number of tracer particles continues to rise for $N_{\text{max}} > 3$. The explanation comes from the fact that the negative-weight particles are those with negative radial velocities and are emitted directly into the outgoing blast wave. Many of them would collide 10 times or more.

The method was implemented into the Cooper–Frye generator and hadronic cascade used in the code B3D [7]. This code incorporates the list of resonances and decays described in the particle data book [13]. For every decay, the inverse process is included through a Breit–Wigner cross section. Additionally, a 10 mb *s*-wave cross section is included for all channels. B3D can generate particles from the hypersurface using the algorithms described in the Appendix, and perform the hadronic cascade in approximately 0.25 seconds on a single core in a typical laptop. Adding backflow corrections with $N_{max} = 3$ slowed the code down by 3% and increased the size of the final-state data files by 5%. The results of Fig. 3 show that using $N_{max} = 3$ allows one to calculate the effects of the backflow to better than 10%. Given that backflow is a small effect, this should be more than sufficient for most purposes.

V. SUMMARY

The backflow contribution from a Cooper–Frye interface between hydrodynamic and Boltzmann-cascade models can indeed be accounted for with the methods described here. The method correctly accounts for the backflow in the limit that the backflow is sufficiently small so that the effects are linear in the backflow. Additionally, a second approximation was applied that limits the simulation of backflow to a finite number of scatterings. Most of the secondary scattering effects occur in the first few scatterings, which allows one to impose the cutoff for small N_{max} . This prevents the method from consuming significant additional numerical resources or from adding significant statistical noise to the output. Even with $N_{\text{max}} = 0$, charge, energy, and momentum conservation was enforced by the algorithm.

This work was mainly motivated by the desire to find a method to account for backflow without requiring disproportionate resources to describe an effect whose impact is modest at best. At midrapidity in high-energy heavy-ion collisions, the number of particles traveling back into the hydrodynamic region is only a fraction of a percent of the total number of particles. For slightly lower energies, or away from midrapidity, that fraction might increase, although is unlikely to exceed a few percent. The method described here seems ideal for such cases.

ACKNOWLEDGMENTS

This work was supported by the Department of Energy Office of Science through Grant No. DE-FG02-03ER41259.

APPENDIX: ALGORITHM FOR GENERATING PARTICLES

For a small hypersurface element Ω_{μ} , the number of particles emitted is

$$dN = \frac{p \cdot \Omega}{(2\pi)^3} \frac{d^3 p}{E_p} f(\mathbf{p}) = |w(p)| \Omega_{\text{opt}} \frac{d^3 p'}{(2\pi)^3} f'(\mathbf{p}'),$$

$$w(p) = \frac{\Omega \cdot p}{\Omega_{\text{opt}}(u \cdot p)}, \quad \Omega_{\text{opt}} = \Omega \cdot u + \sqrt{(\Omega \cdot u)^2 - \Omega^2}.$$
(A1)

Aside from the factor of w(p), the second line looks like the usual thermal emission. The primed quantities refer to momenta and phase-space densities measured in the reference frame of the fluid, i.e., the frame where $u \cdot p = E'$. This formula works for any Ω_{opt} , but for a Monte Carlo procedure it must be chosen large enough that |w(p)| never exceeds unity for any p. Efficiency is lost if it becomes larger than necessary. The choice above is the optimum value as it corresponds to the smallest acceptable value that keeps $|w(p)| \leq 1$.

The strategy is as follows:

- (i) Choose particles according to a static thermal distribution with volume Ω_{opt} . The quantity f should incorporate any viscous corrections.
- (ii) Boost the particles by the four-velocity u.
- (iii) Keep or reject the particle with probability |w(p)|. If a particle is created, a weight of ± 1 is assigned depending on whether w(p) is greater or less than zero. If rejected, continue to the next species.

This procedure should be repeated for each species. However, volume elements tend to be small and there might be many species with small probabilities. Rather than throwing random numbers for each species it is more efficient to implement the following algorithm:

- (i) Generate a series of thresholds separated by random amounts, $-\ln(\operatorname{ran}())$, where ran() generates a random number between zero and unity. This produces thresholds separated by lengths with probabilities e^{-x} . If one is at a position x and one increments by $dx \rightarrow 0$, the chance of hitting a threshold is dx, independent of any previous threshold.
- (ii) Thus one can make a cumulative sum for each species *i*,

$$S_i = \sum_{j \leqslant i} \Omega_{\text{opt}} n_j, \qquad (A2)$$

where n_j is the density in the rest frame of species *j*. If S_i crosses a threshold, one considers creating a particle as described above. If S_i crosses two thresholds, one attempts to create two particles.

For each differential probability the chance of crossing a threshold will be independent of any other differential probability, so the emissions are independent, and thus Poissonian. With this approach, the number of random numbers one generates in deciding whether to attempt creating a particle is the number of particles one would make if the weights, w(p) in Eq. (A1), were unity. This can be a significant improvement compared to throwing random numbers for each species (perhaps over 300) in each volume element (perhaps many millions). Since one knows the total density of hadrons, n_{had} , one can check to see whether $n_{had}\Omega_{opt}$ crosses a threshold. If not, one can increment the running total in one swipe and avoid considering each species independently.

This procedure assumes that you already know n_j . For calculations where the interfaces is defined by a fixed temperature, and with $\mu = 0$, one can calculate these densities once, and store them. For calculations with a variety of breakup densities, one might store an array of values.

Generating a particle according to a thermal phase-space density

In the absence of viscous corrections one can generate a particle in the fluid rest frame with the following algorithm. For a probability distribution $\sim x^{n-1}e^{-x}$, one can sample the distribution by taking the natural log of the product of three random numbers, i.e., $x = -\ln(r_1r_2\cdots r_n)$, where r_i are random numbers chosen uniformly between zero and unity. For a three-dimensional distribution of massless particles, the choice of coordinates is

$$p = -T \ln(r_1 r_2 r_3), \quad \cos \theta = \frac{\ln(r_1) - \ln(r_2)}{\ln(r_1) + \ln(r_2)},$$
$$\phi = \frac{2\pi [\ln(r_1 r_2)]^2}{[\ln(r_1 r_2 r_3)]^2}.$$

To check that this works one can calculate the Jacobian,

$$dN = dr_1 dr_2 dr_3 = dpd \cos\theta d\phi |J|,$$

$$J = \frac{p^2}{8\pi T^3} e^{-p/T}.$$
(A3)

For massive particles one can throw a fourth random number r_4 , and if $r_4 > e^{-(E-p)/T}$, one repeats the procedure until one satisfies the additional condition.

- [1] H. Sorge, Phys. Lett. B 373, 16 (1996).
- [2] S. Pratt and J. Murray, Phys. Rev. C 57, 1907 (1998).
- [3] F. Cooper and G. Frye, Phys. Rev. D 10, 186 (1974).
- [4] C. Nonaka and S. A. Bass, Phys. Rev. C 75, 014902 (2007).
- [5] C. Gale, S. Jeon, B. Schenke, P. Tribedy, and R. Venugopalan, Phys. Rev. Lett. **110**, 012302 (2013).
- [6] H. Song, S. A. Bass, U. Heinz, T. Hirano, and C. Shen, Phys. Rev. C 83, 054910 (2011); 86, 059903 (2012).
- [7] J. Novak, K. Novak, S. Pratt, C. Coleman-Smith, and R. Wolpert, arXiv:1303.5769.
- [8] D. Molnar and M. Gyulassy, Nucl. Phys. A 697, 495 (2002); 703, 893 (2002).

Unfortunately, this becomes inefficient for large masses m because the probability of successfully satisfying the condition for r_4 becomes small for small T/m. In that case one applies an alternative algorithm. For T/m < 0.6, it was empirically found that a more efficient method is based on the expression

$$dN \sim p^2 dp e^{-E_p/T} \sim \frac{p}{E_p} (K+m)^2 dK e^{-K/T},$$
 (A4)

where K = E - m is the kinetic energy. The strategy is to generate *K* ignoring the factor p/E_p , then do a keep-or-repeat based on the weight p/E_p . To generate a value of *K* consistent with $(K + m)^2 e^{-K/t}$, one breaks up the factor $(K + m)^2$ into three terms,

$$dN \sim (K^2 + 2mK + m^2)e^{-K/T}$$
. (A5)

One first throws a random number r_0 and chooses which of the three terms to use as a distribution based on the integrated weights for each term, $2T^3$, $2mT^2$, and m^2T . Once one has picked a given term, one can pick K as $-T \ln(r_1r_2r_3)$, $-T \ln(r_1r_2)$ and $-T \ln(r_1)$, respectively. With this value of K, one can now do a keep-or-repeat decision based on the weight p/E_p . After K is chosen, $\cos \theta$ and ϕ can be picked with new random numbers, r_4 and r_5 .

Viscous corrections can be applied according to Ref. [14]. This involves transforming the momentum according to (in the rest frame)

$$p_i = (\delta_{ij} + \gamma \tau_{ij}) p_j, \tag{A6}$$

where τ_{ij} is the shear tensor in the frame of the fluid, and γ is a constant chosen so that the generated distribution will indeed have the stress-energy tensor one wishes assuming that the viscous correction τ is much smaller than the pressure. This coefficient γ can be found analytically given the list of masses and spins of the hadrons [14]. For $\tau/P < 1/2$, the linear approximation is good to the one percent level or better in that it consistently reproduces the viscous correction τ_{ij} according to

$$\tau_{ij} = \frac{1}{\Omega} \sum_{i} \int d^3 p \frac{dN}{d^3 p} \frac{p_i p_j}{E_p}.$$
 (A7)

For modeling collisions at RHIC or at the Large Hadron Collider, the algorithm described here can generate the thousands of particles needed for a few units of rapidity in a few hundredths of a second.

- [9] S. Cheng, S. Pratt, P. Csizmadia, Y. Nara, D. Molnar, M. Gyulassy, S. E. Vance, and B. Zhang, Phys. Rev. C 65, 024901 (2002).
- [10] C. Anderlik, L. P. Csernai, F. Grassi, W. Greiner, Y. Hama, T. Kodama, Z. I. Lazar, V. K. Magas *et al.*, Phys. Rev. C 59, 3309 (1999).
- [11] P. Huovinen and H. Petersen, arXiv:1206.3371.
- [12] S. Pratt and J. Vredevoogd, Phys. Rev. C 78, 054906 (2008); 79, 069901 (2009).
- [13] J. Beringer *et al.* (Particle Data Group), Phys. Rev. D 86, 010001 (2012).
- [14] S. Pratt and G. Torrieri, Phys. Rev. C 82, 044901 (2010).