LUND UNIVERSITY

Utilizing Monte Carlo methods and shifting to preserve Causality in particle simulation

Aryan Singh Chauhan

Project duration: 6 months
Supervised by Korinna Zapp and
Co-supervised by Chiara Le Roux

Department of Physics
Division of Particle and Nuclear Physics

November 2025

Abstract

This thesis explores the use of Monte Carlo methods and a parton shifting technique to preserve causality in particle simulations, specifically within the framework of parton cascade models. These models are crucial for understanding the dynamics of high-energy nuclear collisions, such as those producing quark-gluon plasma. The primary challenge addressed is the violation of causality due to relativistic effects, where simultaneous events in one reference frame may not be simultaneous in another, leading to inconsistencies in the simulation. The proposed parton shifting method ensures that the collision sequences remain frame-independent, thereby maintaining the integrity of causality. The thesis provides a comprehensive analysis of the results, and discusses the potential implications and future applications of the parton shifting technique in high-energy particle physics.

Popular Science Summary

This thesis presents an innovative method for solving the Boltzmann equation in the context of parton evolution during relativistic heavy-ion collisions. In these extremely high-energy collisions, researchers study how small particles, called partons (quarks and gluons), move and interact. A significant challenge in this field of research is ensuring that events occur in the correct chronological order, a principle known as causality.

In relativistic systems, where speeds can approach the speed of light, the sequence of events can appear differently depending on the frame of reference. This can lead to unrealistic simulations and unreliable results. To address this problem, we proposed a method called "parton shifting."

By repositioning colliding partons to a common space-time point within their center-of-momentum frame (COM frame), we ensure that their interactions occur in the correct chronological order, regardless of the frame of reference used. This means that our simulations maintain their physical accuracy and reliability even when viewed from different moving reference frames.

This method has broad applications and can be used in various fields of research, such as high-energy physics, astrophysics, and plasma science. By improving the accuracy of our simulations, we can gain a deeper understanding of the fundamental processes that govern the universe at its smallest and most energetic levels.

In addition to describing the benefits of the parton shifting method, the thesis also discusses its limitations and proposes enhancements to further increase the method's accuracy and reliability. This research represents a significant step forward in understanding the complex interactions between partons at extremely high energies and contributes to our overall knowledge in the field of physics.

Contents

1 Introduction

Relativistic heavy-ion collisions provide an opportunity to study matter under extreme conditions, such as the formation of the quark–gluon plasma (QGP), a deconfined phase of quarks and gluons believed to have existed shortly after the Big Bang [1, 2]. These collisions, investigated at facilities like RHIC and the LHC, produce highly energetic partons (quarks and gluons) that undergo stochastic scatterings and can be effectively modeled using parton cascade simulations based on the Boltzmann transport equation [3].In such models, particle distributions evolve through a sequence of two-body interactions, determined by cross-sections and relativistic kinematics. While successful in many respects, a major theoretical challenge arises in maintaining causality and Lorentz invariance [4, 5]. Due to the relativity of simultaneity, the chronological ordering of collisions can vary between inertial frames, leading to frame-dependent outcomes that violate the consistency expected from a relativistic theory.
To address this issue, we investigate a method known as *parton shifting*, introduced to restore frame-independent chronology by adjusting the space-time coordinates of colliding partons to a shared interaction point, typically the midpoint in their center-of-momentum frame [6]. This approach aims to eliminate causality violations by enforcing a consistent global ordering of events. The central objective of this thesis is to implement, test, and evaluate the parton shifting procedure within a relativistic cascade simulation and assess its impact on frame invariance and physical consistency. The structure of this thesis is as follows:

- Chapter 1 introduces key principles and mathematical frameworks pertinent to this thesis, including the Boltzmann equation and Monte Carlo methods used for sampling and collision modeling.
- Chapter 2 describes the simulation setup, detailing particle dynamics, collisions in both inertial and boosted frames, and the methodology of parton shifting to analyze and compare collision events.
- Chapter 3 presents and interprets the results, emphasizing the preservation of causality through parton shifting, ensuring invariance across reference frames in parton cascade models.
- In Chapter 4, we discuss the broader implications of the findings, highlighting the enhancements in accuracy and stability of simulations, as well as the method's limitations and potential extensions for future applications.
- Chapter 5 concludes our thesis with a summary of the contributions of this thesis and proposes directions for further research.

1.1 Boltzmann Transport Equation

Originally formulated by Ludwig Boltzmann in 1872, the Boltzmann Transport Equation (BTE) describes the evolution of a distribution function in systems out of thermodynamic equilibrium. It is a nonlinear integro-differential equation involving the phase-space distribution function $f(\mathbf{r}, \mathbf{p}, t)$, which gives the expected number density of particles at

position $\mathbf{r}$, momentum $\mathbf{p}$, and time t. This distribution operates in six-dimensional phase space and accounts for the effects of both particle collisions and external forces [7]. The number of particles in an infinitesimal phase-space volume $d^3\mathbf{r}\, d^3\mathbf{p}$ is:

$$dN(t) = \frac{g}{(2\pi\hbar)^3} f(\mathbf{r}, \mathbf{p}, t)\, d^3r\, d^3p, \tag{1.1}$$

where g is the degeneracy factor and $(2\pi\hbar)^3$ is the quantum phase-space volume [8]. In systems of identical particles undergoing short-range, momentum-conserving interactions, the BTE provides a statistical, rather than deterministic description where instead of resolving individual particle trajectories, the BTE governs the evolution of $f(\mathbf{r}, \mathbf{p}, t)$, which encapsulates the ensemble-averaged behavior of many microscopic particles. It can be expressed as

$$\frac{\partial f}{\partial t} + \mathbf{v} \cdot \nabla_{\mathbf{r}} f + \mathbf{F} \cdot \nabla_{\mathbf{p}} f = \left(\frac{\partial f}{\partial t}\right)_{\text{coll}}, \tag{1.2}$$

The term on the right side of the equation accounts for particle collisions, which for binary collisions is given by a complex integral expression [7]. In contrast, the left-hand side of the BTE accounts for the local changes in the PDF due to temporal evolution, spatial transport, and forces acting on the particles. The term $\mathbf{v} \cdot \nabla_r f$ represents the spatial transport of particles and involves the velocity $\mathbf{v}$ and the spatial gradient of the PDF, thus capturing how the particle distribution shifts in space. The term $\mathbf{F} \cdot \nabla_p f$ represents the force term with $\mathbf{F}$ being an external force acting on particles. $\nabla_p f$ is the gradient of the PDF in momentum space and describes how forces change the momentum distribution of particles [9].

1.2 Monte Carlo Method

The Monte Carlo method is a stochastic computational technique for solving complex problems that are otherwise deterministic in nature. Instead of analytically solving a system, it approximates solutions through random sampling and statistical analysis. This class of methods is widely applied in optimization, numerical integration, and sampling from probability distributions [10, 11].The central idea is to evaluate a function repeatedly using randomly sampled inputs and estimate the result statistically, with convergence guaranteed under the law of large numbers [12]. Though variations exist, Monte Carlo methods generally follow this sequence:

1. Establish the range of valid inputs for the computation.[12].
2. Randomly sample inputs from the domain using an appropriate probability distribution, as this impacts the quality of results [13].
3. Performing deterministic computations on the sampled inputs (it will always produce the same output given the same input)[14].
4. Aggregating results from multiple runs (e.g., mean or median) to improve accuracy and assess performance.[15].

1.3 Causality Violation in Parton Cascade Models

Causality, a foundational principle of relativity and field theory, asserts that no effect can precede its cause and that no influence can propagate faster than light [16, 17]. In relativistic transport models such as parton cascade simulations—which describe the space–time evolution of quarks and gluons in high-energy nuclear collisions—preserving causality is essential for maintaining physical reliability [18, 19].However, conventional formulations can violate causality both geometrically and through frame dependence. In many simulations, two partons are considered to collide when the invariant separation, d_{ij} , between their worldlines satisfies the interaction criterion $d_{ij} \leq d_{\text{crit}}$, even though their trajectories do not intersect at the same spacetime point [20]. This means that the interaction can occur across a spacelike interval, where the particles are not mutually within each other's light cones, leading to an acausal exchange of information [21]. Additionally, because the sequence of collisions is typically ordered using a single simulation frame (often the laboratory or global time frame), the relativistic non-absoluteness of simultaneity causes the same set of events to appear reordered or overlapping in another inertial frame [20, 5]. As a result, a collision that occurs before another in one frame may appear later in a boosted frame, reflecting the loss of consistent causal ordering in standard cascade formulations [18, 19].

Frame-Dependent Collision Sequences: This relativity of simultaneity introduces ambiguity in how collisions are temporally ordered across frames. If a simulation schedules interactions solely based on global lab time, it risks producing sequences where a parton is seen to scatter with two others at overlapping or conflicting times in different frames. Such inconsistencies can result in non-causal behavior—for example, where the outcome of one scattering influences a prior interaction in some frame [22]. These violations have tangible consequences: they compromise the predictive power of cascade simulations and can lead to results that contradict relativistic causality [20, 23]. If left unaddressed, such violations render comparisons with experimental observables unreliable and physically unsound. Thus, causality preservation is not merely a theoretical constraint—it is essential for constructing frame-invariant and consistent simulation algorithms.

1.4 Monte Carlo Integration

Monte Carlo integration is a probabilistic technique for estimating definite integrals using random sampling. It relies on the law of large numbers and the central limit theorem to ensure convergence and accuracy as the number of samples increases. The method is especially useful for high-dimensional or analytically intractable integrals. A one-dimensional integral,

$$I = \int_a^b f(x)\,dx, \tag{1.3}$$

can be approximated by sampling points x_i uniformly from $[a, b]$ and computing:

$$I \approx \frac{(b-a)}{n} \sum_{i=1}^{n} f(x_i), \tag{1.4}$$

where each $f(x_i)$ is an unbiased estimator of the integrand. The estimate converges to the true value as $n \to \infty$, with the error decreasing as $\sim 1/\sqrt{n}$.

1.5 Direct Sampling

In statistics, the probability density function (PDF) $f(x)$ of a continuous random variable X specifies the likelihood of finding X within an infinitesimal interval $[x, x + dx]$, given by $f(x)\,dx$, with normalization $\int_{-\infty}^{\infty} f(x)\,dx = 1$ [15, 11]. The cumulative distribution function (CDF) naturally follows as the integral of the PDF

$$F(x) = P(X \leq x) = \int_{-\infty}^{x} f(x')\,dx', \tag{1.5}$$

representing the total probability of the variable taking a value less than or equal to x. Direct sampling is a Monte Carlo technique used to draw samples from a known probability density function (PDF) by inverting its cumulative distribution function (CDF). A uniformly distributed random number $\xi \in [0, 1]$ is mapped to a sample $\hat{x}$ such that $F(\hat{x}) = \xi$, where $F(x)$ is the CDF of the variable X [11] for a finite domain $[x_{\min}, x_{\max}]$. The normalized sampling condition becomes:

$$F(\hat{x}) - F(x_{\min}) = \xi\,[F(x_{\max}) - F(x_{\min})] \text{ [15].} \tag{1.6}$$

Solving for $\hat{x}$ gives:

$$\hat{x} = F^{-1}\left[F(x_{\min}) + \xi\,(F(x_{\max}) - F(x_{\min}))\right], \tag{1.7}$$

which maps a uniform variable ξ into the distribution defined by $f(x)$. This method is efficient but requires the CDF to be analytically invertible.

1.6 Hit-or-Miss Sampling

Hit-or-Miss sampling is an integration method used when the primitive function of $f(x)$ is non-invertible or the function exhibits complex behavior that makes traditional integration infeasible. This technique is especially effective for stochastic simulations over a bounded interval $A = [a, b]$. The method involves the following steps:

1. Identify an overestimate M such that $M \geq f(x)$ for all $x \in A$. This M should closely approximate the maximum value of $f(x)$ within A to reduce inefficiency [12].

2. Generate uniform samples of x from $[a, b]$ and y from $[0, M]$, forming a grid of points:

$$\{(x_i, y_i)\}_{i=1}^{n} \quad \text{where } x_i \sim \text{Uniform}(a, b) \text{ and } y_i \sim \text{Uniform}(0, M).$$

3. Retain the sample point (x_i, y_i) if $y_i \leq f(x_i)$ [13].

The method's efficiency, η, is determined by the ratio of the area under $f(x)$ to the area of the bounding rectangle. Efficiency decreases significantly when $f(x)$ varies widely within A, as more points are rejected. To address this, techniques like importance sampling can be used to adjust the sampling distribution to better match $f(x)$. Figure 1 demonstrates the concept of the Hit or Miss method. The overestimate M bounds the PDF $f(x)$, with

green dots representing accepted samples (hits) and red dots representing rejected ones (misses).

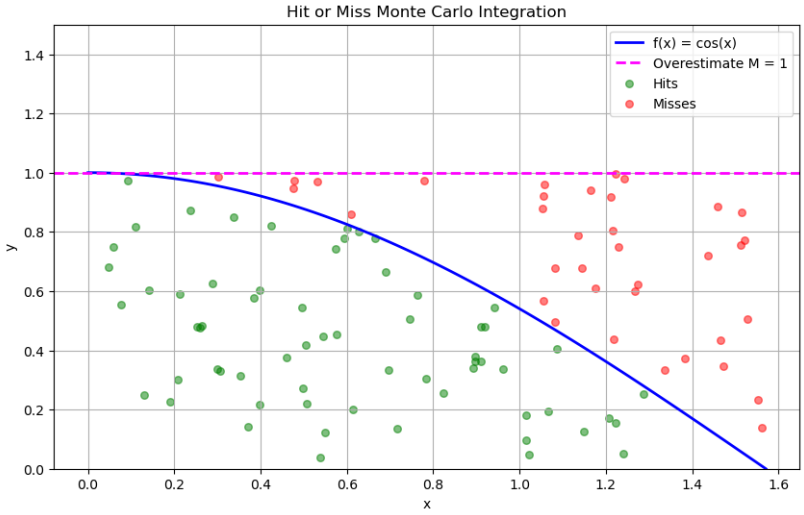

Figure 1: Illustration of the *Hit or Miss* method performed on $f(x) = \cos(x)$ with an overestimate of $M = 1$.

2 Method and Implementation

2.1 Initial Setup

To model the initial conditions for a relativistic simulation, the system is initialized as an ensemble of N partons in thermal equilibrium at temperature T. This initialization defines the reference phase-space configuration from which the partons evolve during the simulation. The phase-space distribution of these partons follows the Maxwell–Boltzmann distribution, which describes the momentum distribution of classical, distinguishable particles in thermal equilibrium. It is valid in the limit where quantum effects are negligible. The distribution is given by:

$$f(\mathbf{r}, \mathbf{p}) = e^{-\frac{E}{k_B T}} = e^{-\frac{\sqrt{\mathbf{p}^2 + m^2}}{k_B T}}, \tag{2.8}$$

with $E^2 = \mathbf{p}^2 + m^2$, and $m, T, k_B, \hbar$ denoting the particle mass, temperature, Boltzmann constant, and reduced Planck constant. In spherical coordinates, the probability of finding a particle in momentum range $[p, p + dp]$ and solid angle (θ, ϕ) is:

$$dP = f(\mathbf{r}, \mathbf{p})\, p^2 dp\, \sin\theta d\theta d\phi\, d^3r. \tag{2.9}$$

Using equations (2.8) and (1.1), assuming natural units ($k_B = \hbar = g = 1$) and the ultrarelativistic limit $E \approx p$, the total number of particles becomes [24]:

$$N = \left(\frac{1}{2\pi}\right)^3 \int e^{-p/T}\, d^3r\, d^3p. \tag{2.10}$$

For a uniform spatial volume $V = L^3$, and converting to spherical momentum coordinates:

$$N = V\left(\frac{1}{2\pi}\right)^3 \int e^{-p/T} p^2 \sin\theta dp\, d\theta\, d\phi = \frac{V}{2\pi^2} \int e^{-p/T} p^2 dp. \tag{2.11}$$

Evaluating the integral gives:

$$N = \frac{VT^3}{2\pi^2}, \quad \text{so} \quad L = \left(\frac{2\pi^2 N}{T^3}\right)^{1/3}. \tag{2.12}$$

This relation determines the spatial extent L for a fixed number of particles N. Initial particle positions are sampled uniformly in $[0, L]^3$. The system is initialized at $T = 1\,\text{eV}$, and the momentum cutoff is set to $P_{\text{max}} = 20T$ to capture the dominant tail of the distribution.

2.2 θ,ϕ,P values

The generation of angular variables θ, ϕ, and the momentum magnitude P is fundamental for defining particle momenta in simulations. These are sampled from the isotropic distribution in three dimensions. Table 1 provides an overview of their respective PDFs, cumulative distribution functions (CDFs), generation methods, and valid ranges.

Table 1: Summary of θ, ϕ, and P value generation methods.

Variable	PDF	CDF	Generation Method	Range
θ	$\sin\theta$	$F(\theta) = -\cos\theta$	Inverse Transform Sampling	$[0, \pi]$
ϕ	Constant	$F(\phi) = \frac{\phi}{2\pi}$	Uniform Sampling	$[0, 2\pi]$
P	$p^2 e^{-p/T}$	Non-invertible	Stratified Hit-or-Miss	$[0, P_{\text{max}}]$

The PDF for θ accounts for the geometric probability associated with angular directions in spherical coordinates. Since the surface area element in spherical coordinates is $dA \propto \sin\theta\, d\theta\, d\phi$, the probability density must be proportional to $\sin\theta$, ensuring uniform coverage of the sphere. To generate values for θ, the direct sampling method is used. Given the cumulative distribution function $F(\theta) = -\cos\theta$, the corresponding inverse function is applied to a uniform random variable $\xi \in [0, 1]$ to obtain:

$$\theta = \arccos(1 - 2\xi).$$

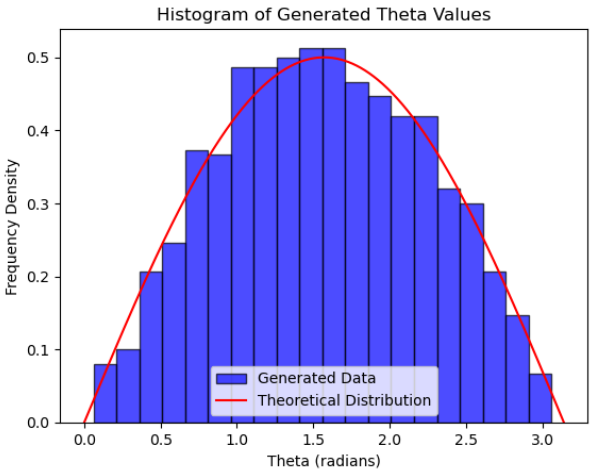

Figure 2: Illustration of Direct sampling for 1000 particles

This ensures that θ values follow the expected distribution [25], as confirmed by the histograms in Figure 2. The azimuthal angle ϕ follows a uniform distribution over

the range $[0, 2\pi]$ since all orientations around the z-axis are equally probable. Since its PDF is constant, ϕ values are sampled directly from a uniform distribution without the need for transformation [26].

Momentum Magnitude: The momentum magnitude P is governed by the PDF $p^2 e^{-p/T}$, which lacks a closed-form inverse function, necessitating the use of the hit-or-miss sampling technique. For efficiency [27, 11], the momentum range is divided into equal-width bins, each assigned an overestimated maximum PDF value, facilitating stratified sampling. The process involves:

1. The momentum range $[0, P_{\max}]$ is divided into n equal-width bins, as demostrated in Fig 3, to localize sampling efforts.
2. The maximum value of the PDF, M_i, within each bin is determined, serving as an upper bound for the sampling process.
3. Each bin is assigned a weight proportional to the area under the curve within it, calculated by numerically integrating the PDF over the bin's range.
4. A bin is selected with probability proportional to its weight, implemented using a weighted random selection command (`numpy.random.choice`) that internally performs cumulative probability sampling, ensuring higher-probability bins are sampled more frequently.
5. A random momentum value p_i is uniformly sampled from the selected bin's range.
6. For each sampled value p_i, a random number y is drawn uniformly from the interval $[0, M_i]$. The sample is accepted if:
$$y \leq f(p_i),$$
otherwise, the value is rejected, and the process is repeated with a new random sample p_i within the same bin. This ensures that the true probability used for bin generation is maintained.
7. The procedure is repeated until the required number of momentum values is obtained, ensuring statistical consistency with the target distribution.

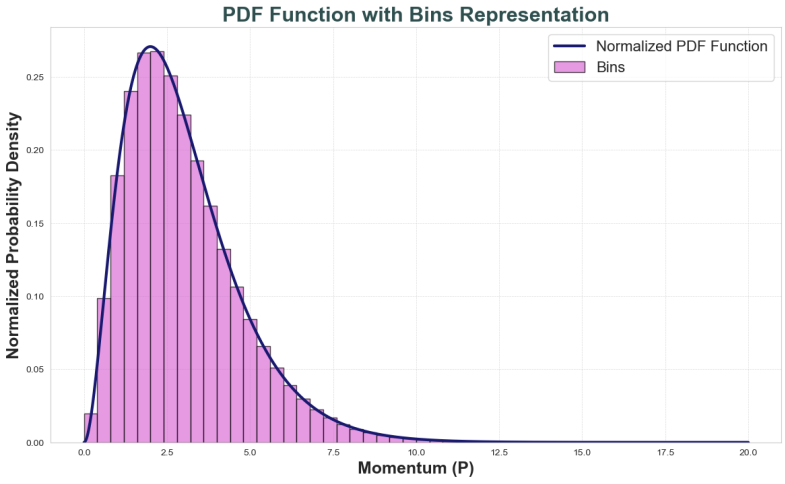

Figure 3: Illustration of creating bins to perform hit or miss in individual bins using weighted probability

Momentum Components:Once the values of θ, ϕ, and P are obtained, the Cartesian components of the momentum vector are computed as follows:

$$p_x = P\sin\theta\cos\phi, \quad p_y = P\sin\theta\sin\phi, \quad p_z = P\cos\theta.$$

Assuming massless particles, the total energy E is given by $E = P$, ensuring compliance with relativistic kinematics [28].

2.3 Collision Detection, Scattering, and Parton Shifting Procedure

Accurately modeling collisions in a relativistic particle system requires a framework that respects Lorentz invariance, conserves energy and momentum, and preserves causality across different frames [29]. This section describes two primary procedures: *Collision Detection* and *Parton Shifting*. The former establishes how collisions are identified and processed, while the latter addresses frame-dependent discrepancies arising from the relativity of simultaneity [16].

Collision Detection Methodology: Collision detection is based on the calculation of a Lorentz-invariant distance measure. The mathematical formulation of this invariant distance is given by:

$$d_{ij}^2 = -\left[(\Delta X^\mu \Delta X_\mu) - \frac{(\Delta X_\mu \Pi^\mu)^2}{\Pi^\alpha \Pi_\alpha}\right], \tag{2.13}$$

where $\Delta X^\mu = x_i^\mu - x_j^\mu$ denotes the four-position difference between particles i and j, and $\Pi^\mu = p_i^\mu + p_j^\mu$ represents their combined four-momentum [29]. Here, four-position vectors X^μ include time and spatial coordinates, $X^\mu = (ct, x, y, z)$, and four-momentum vectors p^μ include energy and momentum, $p^\mu = (E/c, p_x, p_y, p_z)$. The Lorentz-invariant metric $g_{\mu\nu}$ used here follows the convention: $g_{\mu\nu} = \text{diag}(+1, -1, -1, -1)$. Using the invariant distance as a frame-independent measure of proximity, the simulation applies the following procedure to evaluate and schedule binary collisions:

1. **Enumerate Particle Pairs**: Iterate over all unique particle pairs (i, j) without repetition.

2. **Boost to the Center-of-Momentum (COM) Frame**: Each particle pair (i, j) is Lorentz-transformed into its unique center-of-momentum frame to evaluate their interaction in a frame where their total spatial momentum vanishes [30]. The boost velocity is computed from the sum of their momenta:

$$\vec{\beta} = \frac{\vec{p}_i + \vec{p}_j}{E_i + E_j} \tag{2.14}$$

   A Lorentz boost matrix $\Lambda^\mu_{\ \nu}(\vec{\beta})$ is constructed using this velocity [28]. The four-momenta p_i^μ, p_j^μ and the four-positions x_i^μ, x_j^μ are then transformed to the COM frame:

$$p_{i,\text{COM}}^\mu = \Lambda^\mu_{\ \nu} p_i^\nu, \quad x_{i,\text{COM}}^\mu = \Lambda^\mu_{\ \nu} x_i^\nu \tag{2.15}$$

This transformation ensures that in the COM frame, the total momentum $\vec{P}_{\text{tot}}$ satisfies $\vec{P}_{\text{tot}} = 0$, allowing for a symmetric and frame-independent evaluation of the collision.

3. **Determine Closest Approach**: Within the COM frame, the time $t_{\min}$ at which the two particles are spatially closest is computed under the assumption of free-streaming motion [29]. Each particle's trajectory is extrapolated according to:

$$x^{\mu}(t) = x_0^{\mu} + \frac{p^{\mu}}{p^0}(t - t_0) \tag{2.16}$$

This formulation maintains a consistent relativistic worldline for each particle. The value of $t_{\min}$ that minimizes their invariant separation is obtained numerically by minimizing the Lorentz-invariant distance measure (Eq.2.13), typically using scalar optimization routines such as `minimize_scalar` from the `SciPy` library.

4. **Check Collision Condition**:If the computed minimum distance at $t_{\min}$ satisfies $d_{ij} \leq \sigma$, where $\sigma = \sqrt{\sigma_{\text{tot}}/\pi}$ represents the effective interaction range derived from the predefined total cross section σ_{tot}, the event is recorded as a collision.

5. **Translate and Unboost Scattering Pairs**: For valid collision candidates in the COM frame, trajectories are advanced to $t_{\min}$. The particles are then Lorentz-unboosted to the lab frame. This yields two lab–frame time coordinates, t_i and t_j, associated with the event and stored. The effective collision time is set as $\min(t_i, t_j)$, the earliest of the two. This is used to order and compare with other collision candidates.

6. **Global Sorting** :It compiles collision candidates, sorts them by ascending $t_{\min}$, and processes the earliest one first. The code then re-checks collisions, and continues until no more events are left or a max iteration is reached.

Lorentz Invariance in the Parton Cascade Model: This methodology is referred to as the *parton cascade model*, and it is not fully Lorentz invariant [20]. Conventional parton cascade models break Lorentz invariance in two primary ways: through the geometric treatment of collisions and through frame-dependent time ordering.
First, collisions are detected using a geometric approximation based on the minimum invariant distance (or impact parameter) between two parton worldlines, often evaluated at a fixed global simulation time. This procedure is not manifestly covariant because when the criterion is satisfied, particles generally occupy different space–time points. As a result, interactions can occur across spacelike intervals, outside each other's light cones, violating causality as the two particles are not in causal contact. Second, even when each individual collision is treated consistently in its own frame, a second source of Lorentz violation arises from *frame-dependent time ordering*. Standard cascade models enforce a global time evolution—typically defined in the simulation or laboratory frame—to sequence binary collisions. However, due to the relativity of simultaneity, different inertial observers may disagree on the order of such events [16].As a result, the model's outcomes can vary across frames, leading to causality violations and frame-dependent behavior.
To address these issues, our implementation of *parton shifting method*, grounded in Lorentz-covariant principles [22], should theoretically restore locality and simultaneity of collisions in a manifestly covariant way. The same interaction occurs at the same point

in space-time across all frames. In practical terms, this significantly mitigates the frame dependence seen in naive cascade models. Additionally, causal consistency is preserved: no particle is seen to scatter "before" another in one frame and "after" in their world lines. Nevertheless, the method does not completely eliminate all sources of Lorentz violation. The act of shifting particles introduces a *controlled nonlocal modification* to their world-lines, and the reliance on numerical Lorentz transformations introduces finite-precision errors.

Random Boosts: To test the Lorentz invariance of the simulation, particles are boosted to a random inertial frame, and collision results are compared with those in the unboosted (lab) frame [27]. This is to show that collision sequence depends on the observer's frame of reference [1].

- A random subluminal boost velocity $\vec{\beta}$, with $|\vec{\beta}| < 1$, is sampled.
- The corresponding Lorentz transformation matrix $\Lambda^{\mu}_{\ \nu}(\vec{\beta})$ is constructed.
- Each particle's four-position and four-momentum are transformed to $x'^{\mu} = \Lambda^{\mu}_{\ \nu} x^{\nu}$, and $p'^{\mu} = \Lambda^{\mu}_{\ \nu} p^{\nu}$ respectively

Collision detection is then repeated in the boosted frame. Any observed mismatch confirms the need for a covariant correction. The implementation proceeds through the parton shifting procedure, outlined as follows:

1. **Initialization:** A copy of all particle four–positions in the chosen simulation frame is initialized, together with empty storage variables for the next collision candidate. The following quantities are initialized:

$$t_{\min} \leftarrow +\infty, \qquad z \leftarrow (0,\, 0), \qquad x^{\mu}_{\text{commit}} \leftarrow 0.$$

   These variables are updated during the scan over all pairs and determine which pair is selected in the current iteration, and are then carried forward for the comparison with subsequent iteration that identifies the next accepted collision.

2. **Pairwise Evaluation:** Iterate over all unique particle pairs (i, j) in the system.

3. **Boost to COM Frame**: The total four-momentum of the pair is used to construct a Lorentz boost that brings both particles into their center-of-momentum frame, as defined in Eq. 2.14. The corresponding transformation is applied to their four-momenta and four-positions to yield COM-frame quantities, setting the total spatial momentum to zero.

4. **Check Collision Condition**: After calculating the time $t_{\min}$ for the pair the same way as in non-shifted method, the computed minimum distance at $t_{\min}$ is checked if $d_{ij} \leq \sigma$, the pair (i, j) is treated as a *candidate* collision for this iteration.

5. **Align Interaction Point:** For each colliding candidate pair (i, j), the code obtains $x_i^{\mu}(t_{\min}),\ x_j^{\mu}(t_{\min}),\ p_i^{\mu},\ p_j^{\mu}$. and translates the pair via the free-streaming assumption to $(t_{\min})$ in COM frame. The instantaneous positions $x^{\mu}_{i,\text{COM}}(t_{\min})$ and $x^{\mu}_{j,\text{COM}}(t_{\min})$ are generally distinct because their worldlines do not intersect exactly. Therefore

a common interaction event is enforced and both particles are shifted to the mean spacetime position

$$x^{\mu}_{\text{shift}} = \frac{x^{\mu}_{i,\text{COM}}(t_{\min}) + x^{\mu}_{j,\text{COM}}(t_{\min})}{2}. \tag{2.17}$$

This point serves as the effective collision vertex in the COM frame, ensuring that both partons interact at a single, frame-independent spacetime event within their mutual light cone.

6. **Inverse Boost and Candidate Time Extraction:** The shifted COM–frame point x^{μ}_{shift} is Lorentz–unboosted to obtain the corresponding lab–frame four–position $\tilde{x}^{\mu} = \left(\Lambda^{-1}\right)^{\mu}_{\ \nu}(\vec{\beta})\, x^{\nu}_{\text{shift}}$ for the pair (i, j) . The lab–frame time component $\tilde{t} = \tilde{x}^0$, taken as the candidate interaction time, and the unboosted midpoint $\tilde{x}^{\mu}$ together with the pair label (i, j) is temporarily stored for possible selection.

7. **Earliest–Event Selection and Update:** Whenever a colliding pair produces a lab–frame time $\tilde{t} < t_{\min}$, the algorithm updates

$$t_{\min} \leftarrow \tilde{t}, \qquad z \leftarrow (i, j), \qquad x^{\mu}_{\text{commit}} \leftarrow \tilde{x}^{\mu}.$$

After the full scan, the pair z is taken as the next interaction: the global four–positions of that pair are updated to x^{μ}_{commit}, the scattering is performed, and the procedure restarts until no further collisions are found.

The entire logic is summarized in Figure 4.

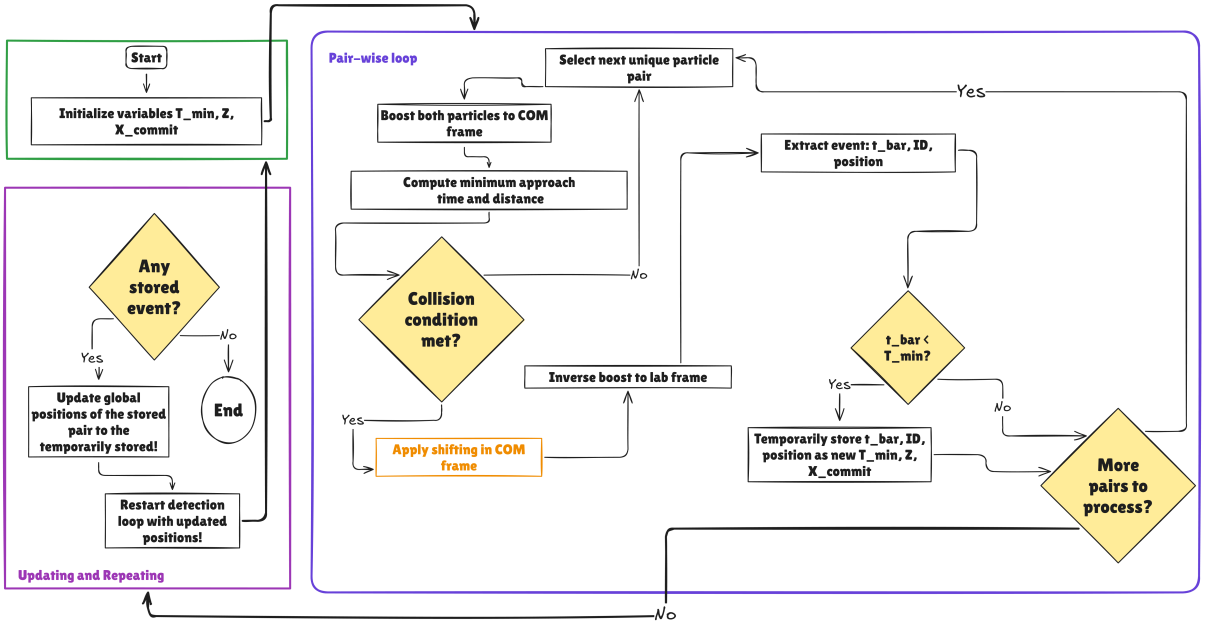

Figure 4: Flowchart of the pair-wise collision detection and updating algorithm.

3 Results

Multiple simulations were conducted to study the frame dependence of relativistic collisions under varying conditions. While the broader analysis spans several configurations

and random seeds, we focus here on one representative simulation to present the methodology and results in detail. A system consisting of 30 partons is initialized according to a thermal momentum distribution at temperature $T = 1\,\text{eV}$, for massless particles ($M = 0$) and with a momentum cut-off set to $P = 20T$. All spatial positions are sampled randomly within a cubic volume of side length $L \approx 8.39750\,\text{eV}^{-1}$ as derived from Eq.2.12. A cross-section of $\sigma_{\text{tot}} = 6\pi\ \text{eV}^{-2}$ is considered[1] [22]. To investigate frame dependence, each simulation is performed in two reference frames:

- **Lab (Unboosted) Frame:** This serves as the baseline frame, where the initial conditions are defined and the total system momentum is zero.
- **Boosted Frame:** All particle momenta and positions are Lorentz-boosted to a frame moving at random velocity, although in this case it is $0.2c$ in the $+x$ direction relative to the lab frame. This transformation enables analysis of how the sequence and nature of collisions vary in a moving reference frame [28].

3.1 Unshifted Simulation Results and Frame-Dependent Effects

Table 2 and 3 summarizes all two-body scatterings that occur in a single run of the simulation in the lab frame (unboosted).In a typical run,40+ scattering events were observed in the lab frame. 42 collisions were recorded in this representative run and these collisions are listed chronologically in ascending time order, which enumerates each pair of parton that scatters, the distance of closest approach, and the time interval of the interaction in the lab frame. All times are measured from the start of the simulation ($t = 0$) in that frame. Most scatterings occur sequentially between distinct pairs; however, several apparent overlaps arise where a single parton participates in multiple interactions during overlapping time intervals. Such behavior is not physically permissible, as a parton cannot undergo two independent interactions simultaneously. The cause of this artifact is the absence of a corrective mechanism; collisions were scheduled strictly by ascending global time. The simulation algorithm, upon finding a new pair satisfying the geometric collision criterion, simply queues that collision even if one of the particles is already mid-collision, since nothing in the naive scheme prevents a parton from colliding again as long as the distance criterion is met. This is a direct manifestation of the causality problem in the unshifted approach [20].

Frame-Dependent Collision Ordering: When the same unshifted simulation is examined from the boosted frame , it is found that the set of collisions that occur is physically the same – the same pairs of particles collide with the same closest-approach distances as expected from Lorentz invariance. This is confirmed by Table 3 , which lists the collision events as observed in the boosted frame. However, the chronological sequence of these events is markedly different between the two frames. The order in which scatterings happen is not invariant but depends on the observer's frame [21].

For instance, in the lab frame particle 7 scatters with particle 3 before later scattering with particle 9, whereas in the boosted frame the same two collisions involving particle 7 occur in the opposite order (7 collides with 9 before colliding with 3). Similarly, the detailed history of particle 10 illustrates this frame-dependent reordering: in the lab

[1] All quantities are expressed in natural units, with $\hbar = c = 1$. Consequently, distances and times are expressed in units of eV^{-1}, and momenta in eV.

Scattering pairs	Distance (eV^{-1})	Beginning (eV^{-1})	End (eV^{-1})
(22, 25)	0.355619	0.044580	0.562214
(16, 20)	2.309322	0.106935	1.043895
(17, 29)	1.609667	0.117785	0.305766
(10, 27)	0.799490	0.212572	0.563519
(10, 14)	1.070405	0.221038	1.053253
(2, 8)	2.322179	0.306968	1.401207
(0, 19)	2.435821	0.328988	0.578387
(11, 18)	1.395971	0.342135	0.787096
(3, 5)	1.663963	0.373641	0.832931
(20, 24)	2.296371	0.397762	1.090396
(16, 24)	0.144602	0.439285	0.447811
(19, 28)	2.636005	0.566329	5.490561
(3, 18)	2.081545	0.582184	1.354017
(7, 29)	1.627032	0.600991	3.291518
(7, 11)	1.482734	0.656477	0.742877
(10, 24)	2.852115	0.709291	1.401311
(11, 14)	2.118114	0.772551	2.558642
(3, 29)	1.115696	0.929747	1.041605
(1, 26)	2.710069	1.044034	1.437879
(0, 26)	1.942200	1.151640	1.407239
(16, 25)	2.411579	1.175127	1.770080
(3, 17)	1.072386	1.185125	1.899511
(9, 18)	1.501315	1.335202	1.462876
(14, 24)	2.258690	1.469145	1.495232
(3, 7)	1.139080	1.550657	1.801182
(7, 9)	2.720390	1.560399	1.665693
(2, 24)	2.974300	1.645971	2.165072
(12, 28)	1.836776	1.699012	2.054147
(11, 26)	1.783412	1.708904	2.455267
(9, 11)	1.629252	1.769253	2.207582
(2, 27)	1.368031	1.864735	2.257384
(0, 15)	1.997415	1.894959	3.367980
(24, 27)	1.719067	2.232279	2.915340
(20, 27)	1.901464	2.404426	2.591628
(14, 26)	1.443520	2.468967	3.009747
(4, 26)	2.378178	2.501401	2.703982
(4, 16)	2.266863	2.541592	3.205128
(2, 21)	2.880384	2.862448	6.568773
(20, 21)	2.396763	3.115861	4.167936
(2, 20)	0.290817	3.467934	3.696688
(3, 24)	2.988777	3.620471	4.775240
(3, 9)	2.603975	3.913020	4.048765

Table 2: Non-shifted scattering in unboosted frame (30 particles, $\sigma_{\text{tot}} = 6\pi$ eV^{-2}).

Scattering pairs	Distance (eV^{-1})	Beginning (eV^{-1})	End (eV^{-1})
(22, 25)	0.355619	-0.766871	-0.265716
(16, 20)	2.309322	-0.493837	0.451520
(20, 24)	2.296371	-0.187836	0.610963
(16, 24)	0.144602	-0.101752	-0.095107
(17, 29)	1.609667	-0.083433	0.181888
(10, 14)	1.070405	-0.062624	0.725010
(10, 27)	0.799490	-0.045803	0.261797
(11, 18)	1.395971	-0.036972	0.524821
(0, 19)	2.435821	-0.027084	0.093235
(10, 24)	2.852115	0.192206	1.055409
(19, 28)	2.636005	0.228262	5.437387
(2, 8)	2.322179	0.241072	1.128617
(3, 5)	1.663963	0.302727	0.806781
(7, 29)	1.627032	0.328557	3.107742
(7, 11)	1.482734	0.359255	0.387501
(11, 14)	2.118114	0.468927	2.154565
(3, 18)	2.081545	0.504664	1.096120
(1, 26)	2.710069	0.569966	1.225528
(16, 25)	2.411579	0.571609	1.008454
(0, 26)	1.942200	0.678219	0.864961
(3, 29)	1.115696	0.793351	0.949532
(14, 24)	2.258690	1.027131	1.128228
(9, 18)	1.501315	1.077160	1.087915
(3, 17)	1.072386	1.088507	1.858624
(2, 24)	2.974300	1.221426	1.841740
(7, 9)	2.720390	1.300792	1.347775
(11, 26)	1.783412	1.314400	1.989675
(0, 15)	1.997415	1.319067	2.652957
(3, 7)	1.139080	1.337424	1.685049
(9, 11)	1.629252	1.374068	1.869557
(12, 28)	1.836776	1.502366	1.981212
(2, 27)	1.368031	1.550008	1.927919
(4, 26)	2.378178	1.866266	2.239884
(4, 16)	2.266863	1.904585	2.429242
(24, 27)	1.719067	1.905016	2.616205
(14, 26)	1.443520	2.003457	2.509918
(20, 27)	1.901464	2.071291	2.120495
(2, 21)	2.880384	2.492789	5.846640
(20, 21)	2.396763	2.672080	3.584736
(2, 20)	0.290817	3.058054	3.283212
(3, 24)	2.988777	3.446711	4.659856
(3, 9)	2.603975	3.659585	3.861439

Table 3: Non-shifted scattering in boosted frame (30 particles, $\sigma_{\text{tot}} = 6\pi$ eV^{-2}).

frame, particle 10's collision with particle 27 begins earlier than its subsequent collision with particle 14; but after Lorentz boosting to the moving frame, particle 10 appears to collide with particle 14 first (the boosted time t' for the 10–14 collision starts earlier) and the 10–27 collision is delayed relative to that. Such a discrepancy implies that cause and effect could be seen in reverse order in different frames – a clear violation of causality. The absence of a frame-invariant collision criterion allows two scatterings to be processed in an order that is not preserved under Lorentz transformation.

In summary, the unshifted cascade results are inconsistent across frames and exhibit unphysical behavior. Both concurrent collisions (one particle in two overlaps) and frame-dependent reordering of events is observed, indicating that the naive scheduling of parton collisions violates causality. These findings highlight the need for a mechanism to enforce a consistent, causal ordering of collisions that all observers will agree on. We address this by introducing the parton shifting technique and examining its impact on the simulation outcomes.

3.2 Shifting and Preserving Causality in Particle Simulations

With the shifting technique in place, the pathologies seen in the unshifted results disappear. Table 4 lists the sequence of scattering events after performing shifting in both the lab frame and the boosted frame. Each row corresponds to a specific scattering event (identified by the two partons involved), showing the pairing and the time of the collision in the lab frame (NB) alongside the pairing and time in the boosted frame (B). We see that the same set of scattering events occurs in both frames, but the temporal order of events is not identical when viewed globally. The shifting method guarantees that causality is preserved locally for each particle and that no event can precede its cause. In the shifted results, no parton ever undergoes two collisions at once, and each parton's individual sequence of collisions remains the same in all frames [22, 20].

Index	Scattering pairs (NB)	Scattering pairs (B)	Scattering time (NB) (eV^{-1})	Scattering time (B) (eV^{-1})	Distance (NB) (eV^{-1})	Distance (B) (eV^{-1})
0	(17, 29)	(22, 25)	0.211776	-0.516293	1.609667	0.355619
1	(22, 25)	(16, 24)	0.303397	-0.098429	0.355619	0.144602
2	(10, 27)	(16, 20)	0.388046	-0.041352	0.799490	2.336865
3	(16, 24)	(0, 19)	0.443548	0.033075	0.144602	2.435821
4	(0, 19)	(17, 29)	0.453687	0.049227	2.435821	1.609667
5	(16, 20)	(10, 27)	0.553664	0.107997	2.336865	0.799490
6	(11, 18)	(11, 18)	0.564616	0.243925	1.395971	1.395971
7	(3, 5)	(3, 5)	0.603286	0.554754	1.663963	1.663963
8	(2, 8)	(7, 11)	0.854088	0.598638	2.322179	1.378691
9	(7, 11)	(11, 29)	0.896695	0.498499	1.378691	1.406274
10	(11, 29)	(0, 26)	0.731784	0.650408	1.406274	0.894405
11	(0, 26)	(2, 8)	1.119244	0.684845	0.894405	2.322179
12	(3, 29)	(9, 18)	1.134620	0.855054	2.010963	1.478415
13	(9, 18)	(1, 26)	1.201431	0.949583	1.478415	2.834587
14	(1, 26)	(3, 29)	1.290107	0.983623	2.834587	2.010963
15	(3, 10)	(14, 24)	1.465429	1.070964	2.538659	2.255160
16	(14, 24)	(16, 25)	1.475457	1.081621	2.255160	1.635464
17	(7, 9)	(3, 10)	1.575918	1.214715	1.296471	2.538659
18	(9, 29)	(7, 9)	1.552489	1.283865	1.834906	1.296471
19	(2, 27)	(9, 29)	1.636536	1.330662	1.674199	1.834906
20	(16, 25)	(2, 27)	1.817526	1.368567	1.635464	1.674199
21	(12, 28)	(14, 20)	1.876579	1.471432	1.836776	2.053049
22	(14, 20)	(2, 24)	1.951555	1.648235	2.053049	2.225004
23	(2, 24)	(12, 28)	1.986042	1.741789	2.225004	1.836776
24	(11, 26)	(11, 26)	2.064028	1.773584	1.149004	1.149004
25	(7, 29)	(20, 27)	2.218353	1.913433	0.738851	2.280115
26	(20, 27)	(7, 29)	2.297837	1.969090	2.280115	0.738851
27	(9, 26)	(9, 26)	2.299150	2.056275	1.857238	1.857238
28	(2, 3)	(2, 3)	2.500829	2.165927	2.152576	2.152576
29	(3, 20)	(3, 20)	2.982809	2.634469	1.436255	1.436255
30	(19, 28)	(19, 28)	3.122935	2.878673	2.328718	2.328718
31	(3, 26)	(3, 26)	3.471575	3.153890	2.710856	2.710856

Table 4: Scattering pairs after shifting in un-boosted and boosted frame with 30 particles with σ_{tot}-value of 6π eV^{-2} arranged in time order.

Causal Consistency After Shifting: To illustrate, after applying shifting, particle 7 from the earlier example still collides with the same partners (say, with 29 and with 11), but those interactions are now ordered in time for particle 7 in a consistent way across all frames. If particle 7 collides with 29 before colliding with 11 in the COM frame, then in both the lab and boosted frames, particle 7 is also found to collide with 29 before 11 (even though the global timestamps of those events differ between frames). Any potential overlap is resolved because the shifting algorithm would not allow the second collision to start until the first is completed – the two events share a common point on particle 7's worldline [22]. The same holds for particle 11 and all other partons: each particle's interaction history becomes an invariant causal chain. Table 5 demonstrates this by listing the scattering events arranged by individual particles. For each parton, the table shows the collisions it participates in (with corresponding times) in the lab frame versus the boosted frame. Every entry for a given particle matches in both frames. For

example, particle 9 (which in the unshifted run had a problematic history) now has a well-defined sequence: in the lab frame, particle 9 first scatters with particle 18, then with particle 29, then with particle 26, in that order. In the boosted frame, the same three interactions occur at different global times, but in the identical physical order: $9 \to 18$, followed by $9 \to 29$, and then $9 \to 26$. In other words, the relative order of particle 9's own interactions is preserved [21]. Furthermore, there are no overlaps: particle 9's second collision (with 29) begins only after the first (with 18) has finished, and the third (with 26) begins after the second is done. This behavior holds for every parton in the simulation. Thus, local causality (the cause–effect progression for each particle) is maintained, even though observers may see the overall event list in different orders.
Lorentz Frame Independence and Causal Structure: An important consequence of the shifting procedure is that it restores Lorentz frame independence to the collision history. Since collisions are defined at invariant space-time points (by using the COM frame for each pair), no matter which inertial frame the final result is record from, the same set of scattering events and the same causal relationships is found. No interaction can "appear" in one frame and not in another, and no particle sees its cause–effect timeline scrambled by a change of frame. Every scattering lies within the future light cone of any earlier scattering involving the same particle, across all frames, thus preserving the causal order across both frames, and each particle's causal relationships are maintained within their own light cones, independent of the global sequence, or in other words, each particle's interaction history is an independent causal chain [21, 28]. By focusing on the light cone structure of each particle, the causality is preserved without needing to consider the entire global sequence as the global sequence is a result of the superposition of multiple independent local sequences.
To emphasize this, the example of particle 9 with shifting enabled is shown in Table 5. In the lab frame, particle 9's first collision (with 18) occurs at time $t_{\text{lab}} \approx 0.855$ and the next (with 29) at $t_{\text{lab}} \approx 1.331$; in the boosted frame, these same events occur at $t' \approx 1.201$ and 1.552 respectively. Despite the time coordinate differences, the second event still happens after the first in both frames, and indeed the second lies inside the future light cone emanating from the first event in space-time. The third collision for particle 9 is with 26, same in both frames.

Index	Scattering pairs (NB)	Scattering pairs (B)	Scattering time (NB) (eV^{-1})	Scattering time (B) (eV^{-1})	Distance (NB) (eV^{-1})	Distance (B) (eV^{-1})
0	(0, 19)	(0, 19)	0.033075	0.453687	2.435821	2.435821
1	(0, 26)	(0, 26)	0.650408	1.119244	0.894405	0.894405
2	(1, 26)	(1, 26)	0.949583	1.290107	2.834587	2.834587
3	(2, 8)	(2, 8)	0.684845	0.854088	2.322179	2.322179
4	(2, 27)	(2, 27)	1.368567	1.636536	1.674199	1.674199
5	(2, 24)	(2, 24)	1.648235	1.986042	2.225004	2.225004
6	(2, 3)	(2, 3)	2.165927	2.500829	2.152576	2.152576
7	(3, 5)	(3, 5)	0.554754	0.603286	1.663963	1.663963
8	(3, 29)	(3, 29)	0.983623	1.134620	2.010963	2.010963
9	(3, 10)	(3, 10)	1.214715	1.465429	2.538659	2.538659
10	(3, 20)	(3, 20)	2.634469	2.982809	1.436255	1.436255
11	(3, 26)	(3, 26)	3.153890	3.471575	2.710856	2.710856
12	(7, 11)	(7, 11)	0.598638	0.896695	1.378691	1.378691
13	(7, 9)	(7, 9)	1.283865	1.575918	1.296471	1.296471
14	(7, 29)	(7, 29)	1.969090	2.218353	0.738851	0.738851
15	(9, 18)	(9, 18)	0.855054	1.201431	1.478415	1.478415
16	(9, 29)	(9, 29)	1.330662	1.552489	1.834906	1.834906
17	(9, 26)	(9, 26)	2.056275	2.299150	1.857238	1.857238
18	(10, 27)	(10, 27)	0.107997	0.388046	0.799490	0.799490
19	(11, 18)	(11, 18)	0.243925	0.564616	1.395971	1.395971
20	(11, 29)	(11, 29)	0.498499	0.731784	1.406274	1.406274
21	(11, 26)	(11, 26)	1.773584	2.064028	1.149004	1.149004
22	(12, 28)	(12, 28)	1.741789	1.876579	1.836776	1.836776
23	(14, 24)	(14, 24)	1.070964	1.475457	2.255160	2.255160
24	(14, 20)	(14, 20)	1.471432	1.951555	2.053049	2.053049
25	(16, 24)	(16, 24)	-0.098429	0.443548	0.144602	0.144602
26	(16, 20)	(16, 20)	-0.041352	0.553664	2.336865	2.336865
27	(16, 25)	(16, 25)	1.081621	1.817526	1.635464	1.635464
28	(17, 29)	(17, 29)	0.049227	0.211776	1.609667	1.609667
29	(19, 28)	(19, 28)	2.878673	3.122935	2.328718	2.328718
30	(20, 27)	(20, 27)	1.913433	2.297837	2.280115	2.280115
31	(22, 25)	(22, 25)	-0.516293	0.303397	0.355619	0.355619

Table 5: Scattering arranged according to individual particles in ascending time order for individual particles.

The same result was obtained when scattering verification was carried out using an alternative method. In this procedure, each particle's position was recalculated at the exact time of every scattering event to ensure that all collisions were examined at consistent temporal coordinates. The particles were thus advanced with free-streaming assumption along their original trajectories—without changing their momenta—to confirm that the identified scattering pairs were physically consistent in time and that the results were independent of the verification method used, whether based on stored positions or recalculated ones.

3.2.1 Robustness of Causality Preservation Across Multiple Runs

To assess the robustness of the parton shifting method, the simulation was executed with varying random seeds and three different values of the cross-section parameter σ. Table 6 presents a summary of key metrics recorded across these runs, including the number of scattering events, the average minimum distance between colliding particles $\langle d_{ij} \rangle$, and the number of causality violations detected before and after shifting.

As expected, the number of scattering events increased with higher σ_{tot} values due to the greater effective interaction range. More importantly, in every single run, the simulation exhibited causality violations in the unshifted case (ranging from 2 to 7 violations per run). After applying the parton shifting procedure, all violations were eliminated, demonstrating that the method consistently preserves causality across different initial conditions and reference frames. This confirms not only the stability of the implementation but also the effectiveness of the shifting method as a frame-independent correction strategy in relativistic cascade simulations.

Table 6: Summary of 10 simulation runs showing scattering behavior and causality preservation with parton shifting.

Run	$\boldsymbol{\sigma}_{\text{tot}}$ **(eV^{-2})**	**Scatterings**	$\boldsymbol{\langle d_{ij} \rangle}$ **(eV^{-1})**	**Violations (Before)**	**Violations (After)**
1	6π	42	1.85	6	0
2	6π	40	1.91	5	0
3	5π	37	1.78	4	0
4	5π	35	1.74	5	0
5	4π	30	1.62	3	0
6	4π	28	1.59	2	0
7	6π	43	1.87	7	0
8	5π	36	1.76	4	0
9	4π	29	1.60	3	0
10	6π	41	1.89	6	0

4 Discussion

The simulation results clearly demonstrate the efficacy of the parton shifting method in preserving causality. In the unshifted cascade, multiple causality violations per run were observed, whereas with shifting applied, no violations occurred in any run. This section discusses why the shifting algorithm achieves such improvements, examining its theoretical implications (causality enforcement, Lorentz invariance, and locality), practical advantages, limitations, and possible enhancements.

Theoretical Consistency: A longstanding challenge in parton cascade models is maintaining Einstein causality, ensuring that no interaction effectively propagates information faster than light, despite many nearly simultaneous scatterings in high-density scenarios. In a naive cascade (where any two partons collide once they come within a cross-section distance), processing collisions in a fixed frame-time order can produce apparent violations: one frame's ordering may imply superluminal signals or reversed cause-effect sequences ("superluminal shockwaves" in early simulations) [18, 19].

The parton shifting solves this interaction ordering problem. Rather than using an arbitrary observer's time to sequence events, the algorithm determines each collision's space-time point in a frame-invariant way. When two partons satisfy the collision criteria, the simulation does not immediately register a collision in the current frame. Instead, it shifts the partons along their trajectories to make them collide at the space-time point of closest approach,effectively the mid-point between their worldlines in the center-of-momentum frame [22, 19]. By scheduling the collision at this common point, the interaction becomes simultaneous in the two-particle center-of-mass frame, ensuring it occurs within both particles' future light cones and thus respects locality and causality. All observers will then agree on the occurrence of that collision, yielding a Poincaré-covariant ordering of events [31].

1. **Causality Preservation** – Parton shifting imposes a Lorentz-invariant ordering of collision events, preserving causality in the simulation. Instead of relying on an observer-dependent time coordinate, each collision is scheduled at a unique space–time point (the two particle COM frame), so that all inertial observers agree on the sequence of interactions [5, 6, 31]. In practice, this method eliminates the

pathological cases seen in the unshifted cascade (a later collision overtaking an earlier one under a boost) [18, 19]. As a result, each particle's scattering history becomes independent of the global time coordinate and remains consistent under Lorentz transformations. By construction, collisions occur only when the colliding partons' future light cones overlap, thereby upholding Einstein causality and forbidding any superluminal influence [21, 17, 16]. Hence, each collision is a well-defined space–time event, resulting in a cause–effect narrative in every frame. In effect, the parton shifting approach integrates a local, relativistically consistent interaction into the cascade, making the evolution fully Lorentz-covariant and free from frame-dependent artifacts [29, 32, 33]. In combination with the covariant formulations developed in later transport frameworks such as the AMY and ALPACA approaches [32, 33], the shifting algorithm could represent a practical route towards fully causal and physically consistent parton transport in the ultrarelativistic regime.

2. **Algorithmic and Practical Benefits**: In addition to the fundamental theoretical advantages, parton shifting proved to be efficient and easily implementable. It required adding a calculational routine to determine collision points and performing Lorentz transformations, but these operations are not intensive compared to the Monte Carlo generation of collisions themselves. The overall simulation runtime increased only modestly (on the order of 10–20% for our test scenarios) after including shifting, which is a small price for the improvement in accuracy. This contrasts favorably with particle subdivision methods[5], which can increase computational costs multiplicatively by the subdivision factor. Furthermore, it did not need to rewrite the entire cascade algorithm, but only to insert the shifting step at the appropriate point in the collision handling routine. This suggests that existing transport models could adopt parton shifting without major overhaul, leveraging their current physics (cross sections, particle species, etc.) while fixing the causality problem. Also the method keeps event-by-event fluctuations intact (unlike subdivision, which artificially suppresses fluctuations by averaging over many test particles). Each collision in our approach is a real collision between two specific particles, so fluctuations in the number and pattern of collisions from event to event remain physical. This is important if one wants to study event-by-event observables like flow fluctuations or correlations. Lastly, because the parton shifting does not rely on any specific physical assumption beyond elasticity and two-body dominance, it can be applied to a wide range of scenarios (e.g. pure gluon gas, quark-gluon mixtures, etc.) as long as the basic collision logic is two-particle.

3. **Locally Interpretable Temporal Evolution**: The shifting procedure restores a *locally Lorentz-consistent* notion of time evolution for each binary interaction in the cascade. In relativistic systems, simultaneity depends on the observer's frame, so advancing the simulation by a single global time variable t_{lab} is not physically justified. Standard cascade models rely on this global clock, causing collisions to appear reordered or even overlapping in different inertial frames, thereby breaking causal correspondence between the simulation sequence and relativistic dynamics. The shifting procedure mitigates this issue locally by performing each binary collision and shifting in the pair's COM frame. This defines a unique Lorentz-covariant collision event for the pair—one that all inertial observers identify consistently—without referring to the global simulation time.

The preservation of locality, causality, and Lorentz invariance together makes the parton shifting enhanced cascade much more sound as a tool for studying relativistic many-body dynamics. It gives confidence that any collective phenomena emerging from the simulation (such as momentum anisotropies or transport coefficients) are genuine and not influenced by algorithmic artifacts. Despite its utility, the parton shifting approach – and the cascade model built around it – comes with several important limitations and assumptions. These stem from both theoretical constraints and the simplifications made in the simulation. Some limitations are as follows:

1. **Fixed Simulation Frame for Ordering**: While the method greatly reduces frame-dependent ambiguities, it is not a fully covariant solution in the multi-particle sense. It still advances the cascade using a chosen "computational frame" (in this case, the center-of-mass frame) to sequence different collision events. If two collisions are space-like separated (i.e. far apart and simultaneous in one frame), their order can be reversed by a Lorentz boost and our algorithm has no unique condition for such cases except to choose one in the simulation frame. In practice, this residual issue is minor because truly simultaneous independent collisions are rare and causally disconnected (they commute). However, this implies that the shifting method is not manifestly Lorentz-covariant for the entire N-particle system in the way of the proper-time-ordered Alpaca algorithm[6]. While the approach ensures each individual collision is covariant, the global ordering of collisions can differ across reference frames. This is a theoretical limitation: a perfect Lorentz-invariant N-particle dynamics without fields is forbidden by the no-interaction theorem[6]. This is circumvented by an approximate scheme, but an exact invariant treatment would require incorporating field-mediated interactions or a continuous time variable for all particles.

2. **Challenges at High Density (Collision Overlap)**: In extremely dense systems – such as those formed during the earliest stages of heavy-ion collisions – the average spacing between partons becomes so small that many collisions would occur nearly simultaneously. Under such conditions, many binary interactions would occur nearly simultaneously. Because the cascade algorithm processes collisions sequentially, it cannot fully resolve such concurrency, potentially leading to *collision overlap* or opacity effects, as discussed in previous studies [20, 6]. These effects may artificially suppress the effective interaction rate, since nearby scatterings cannot be handled simultaneously within a binary scheme. While not directly observed in the present study, this limitation is inherent to binary cascade frameworks, including the parton shifting where multiple scatterings overlap and interfere at high densities.

3. **Assumption of Massless (On-shell) Particles**: In practice, partons may be produced or propagate in an *off-shell* state, particularly during radiative stages, carrying finite virtuality and correspondingly limited lifetimes before evolving or hadronizing. By contrast, the framework here models partons as *massless and on-shell*, propagating freely between collisions with no virtuality, formation time, or dynamic mass. Virtuality measures deviation from the mass-shell condition,

$$Q^2 = p^\mu p_\mu - m^2. \tag{4.18}$$

Highly off-shell partons satisfy $|Q^2| \gg 0$ and have short lifetimes, approximately

$$\tau \sim \frac{1}{|Q|}. \tag{4.19}$$

Through successive scatterings and emissions, the parton virtuality decreases monotonically,

$$Q_{n+1}^2 < Q_n^2, \tag{4.20}$$

approaching the lightlike, massless limit,

$$p^2 \to 0 \quad \text{or} \quad E = |\mathbf{p}|. \tag{4.21}$$

This *virtuality reduction* [35, 36] accompanies thermalization, as initially off-shell excitations relax toward on-shell, long-lived quasi-particle states [32, 37]. The framework assumes this reduction has already occured, thus treating partons as instantly on-shell and classically propagating. This approximation omits off-shell propagation, parton showering, and quasi-particle mass evolution. Models incorporating finite formation times can approximate these effects, but the present formulation remains limited in describing *finite-lifetime dynamics*. [22].

Several theoretical extensions of the parton shifting framework can be formulated to further improve the causal and physical consistency of relativistic transport models. Some possible directions include:

1. **Statistical Ensemble Shifting (SES)**: While the shifting procedure checks compliance with the minimum invariant separation criterion, it is implemented through a single deterministic spacetime point $d_{ij}^{(\min)}$, thus representing one arbitrary configuration within the causal domain $d_{ij} \leq \sigma$. Such determinism can introduce subtle frame-dependent ordering biases when multiple near-simultaneous collisions occur. Rather than collapsing the overlap region to a unique midpoint, SES treats the admissible spacelike displacements Δx^μ about the nominal collision point x_{col}^μ as elements of a Lorentz-invariant probability ensemble. Using the projector

$$P^{\mu\nu} = g^{\mu\nu} - \frac{\Pi^\mu \Pi^\nu}{\Pi_\alpha \Pi^\alpha}, \tag{4.22}$$

where $\Pi^\mu = p_i^\mu + p_j^\mu$ is the total four-momentum of the pair, the invariant squared separation on the orthogonal hypersurface can be defined as

$$s = -\,\Delta x_\mu\, P^{\mu\nu}\, \Delta x_\nu \;\geq 0. \tag{4.23}$$

The ensemble kernel could then be written as

$$P(\Delta x^\mu) \;\propto\; \exp\!\left[-\,\frac{s}{2\sigma_s^2}\right], \qquad \Delta x_\mu\, \Pi^\mu = 0, \tag{4.24}$$

with σ_s taken as a small fraction of the nominal interaction range ($\sigma_s \approx 0.1$–$0.3\,\sigma$). In the pairs' COM frame, where $\Pi^\mu = (\Pi^0, \mathbf{0})$, a spacelike displacement is sampled as

$$\Delta x_{\text{COM}}^\mu = (0,\, \boldsymbol{\delta r}), \tag{4.25}$$

satisfying $\Delta x_\mu \, \Pi^\mu = 0$. The effective collision point is then

$$x^\mu_{\text{eff, COM}} = x^\mu_{\text{col, COM}} + \Delta x^\mu_{\text{COM}}, \tag{4.26}$$

which can subsequently be unboosted to the laboratory frame. SES interprets the microscopic overlap region as a causal ensemble. Each realisation corresponds to one valid spacetime ordering consistent with the invariant condition, and ensemble-averaged observables become insensitive to arbitrary microscopic shifts. This reformulates the geometric shifting step as a covariant stochastic regularisation of micro-causality approaches [37, 38].

2. **Off-Shell Propagation and Coherence (LPM Effect)**: While the shifting and SES methods may restore the spatial covariance of individual collisions, the interactions are treated as events between on-shell particles. Realistically partons propagate as off-shell quasi-particles that retain their causal memory of previous interactions over a finite formation time. During the finite coherence period, overlapping scattering amplitudes are not causally independent but interfere destructively—the **Landau–Pomeranchuk–Migdal** (LPM) effect that limits independent radiation in dense media. This way the causal consistency restored by geometric shifting is carried over to the *temporal domain*, where finite formation times govern the sequencing of interactions. Each parton can be assigned a virtuality $q^2 \neq m^2$ that evolves continuously between scatterings, reflecting its off-shell propagation. Radiative and inelastic processes such as $2 \to 3$ and $3 \to 2$ transitions are then governed by a characteristic formation time [39, 40],

$$\tau_{\text{form}} \simeq \sqrt{\frac{2\,\omega}{\hat{q}}}, \tag{4.27}$$

where ω is the energy of the radiated gluon and $\hat{q}$ denotes the local transport coefficient of the medium. If a subsequent interaction occurs before this interval elapses ($\Delta t < \tau_{\text{form}}$), the corresponding amplitudes would interfere, suppressing independent radiation. This coherence can be represented dynamically through a weighting factor [32, 37],

$$w_{\text{LPM}} = 1 - \exp\left[-\frac{\Delta t}{\tau_{\text{form}}(E, \hat{q})}\right], \tag{4.28}$$

which interpolates between the incoherent Bethe-Heitler regime and the fully coherent LPM-suppressed limit. Off-shell propagation provides the physical mechanism underlying this coherence. The parton virtuality relaxes toward its on-shell value over time [37, 41],

$$\frac{dq^2}{dt} \sim -\frac{q^2 - q^2_{\text{eq}}}{\tau_{\text{form}}}, \tag{4.29}$$

ensuring a continuous transition between interactions rather than instantaneous state changes. This framework can accommodate thermal masses and resonance states, where gluons or quarks acquire effective masses and finite spectral widths, as well as color-coherence effects such as angular ordering or color-flow correlations. These features enable the cascade to treat partons as genuine quasi-particles that propagate, radiate, and interact coherently with the surrounding medium. Thus

resulting in a fully covariant and causally consistent cascade model, bridging geometric shifting with quantum coherence and potentially restoring temporal causality. Such implementations can reproduce realistic jet-quenching and equilibration behavior observed in effective QCD transport frameworks, where LPM suppression is embedded directly into the collision kernels [32, 31].

3. **Incorporating Quantum-Statistical Effects**: An extension to the shifting framework could be to incorporate quantum-statistical effects directly into the effective local collision probability. In the current approach, scattering probabilities depend only on invariant distance and cross sections, with no account for final-state occupation. However, at high densities or near local equilibrium, quantum statistics become relevant: fermionic states may already be partially filled (*Pauli blocking*), while bosonic states can be stimulated by prior occupancy (*Bose enhancement*). To capture these effects, the collision probability can be modified by multiplicative occupation factors $(1 - f)$ for fermions and $(1 + f)$ for bosons, where $f(x, p)$ represents the local phase-space occupancy [34]. This correction ensures that the microscopic dynamics obey the correct Fermi–Dirac and Bose–Einstein limits and maintain detailed balance at the quantum level [22].

   Within the shifting method, these corrections can be integrated naturally by having the occupation factor $f(x, p)$ estimated locally around the shifted collision point after each binary pair, in their respective COM frame, satisfies the invariant collision distance criterion. A stochastic acceptance–rejection step can then be applied: a fermionic scattering that would populate an already-occupied momentum cell is rejected with probability $(1-f)$, while bosonic outcomes are accepted with enhanced probability proportional to $(1+f)$. This preserves the covariant spacetime geometry of collisions determined by shifting, while refining the statistical rate to reflect quantum occupation.

   Although a complete implementation would require dynamic evaluation of $f(x, p)$ in six-dimensional phase space, which is computationally demanding and memory-intensive, practical estimators can be used to approximate it. For example, local kernel-density reconstructions from neighboring test particles, sparse or hashed momentum binning, and adaptive sampling in dense regions have been shown to provide efficient and accurate estimates of $f(x, p)$ without requiring full grid storage [37]. These methods enable the incorporation of quantum-statistical corrections even in large-scale cascade simulations. Such occupancy-dependent corrections have already been realized in Lorentz-invariant kinetic frameworks such as Alpaca, where they enforce thermal detailed balance and ensure the correct quantum-statistical limits [33]. A quantum-corrected cascade would thus better reproduce equilibrium distributions and improve accuracy in dense regimes such as cold quark matter or the late-time evolution of the quark–gluon plasma.

5 Conclusion and Outlook

The parton shifting procedure developed and implemented in this thesis could be a significant step toward restoring Lorentz invariance and causality in relativistic cascade simulations. By assigning collision points based on pairwise center-of-momentum frames and enforcing causal order along each particle's worldline, the

method overcomes the core limitations of naive cascade algorithms, particularly their frame dependence and violation of locality [5, 20, 21]. Through extensive simulation tests, it is demonstrated that this technique eliminates causality violations, preserves the physical integrity of particle histories, and maintains consistency across inertial frames [18, 6]. These improvements were achieved without drastically increasing computational overhead or requiring a fundamental rewrite of the cascade model structure.

Nevertheless, while the shifting method corrects issues, it is not a complete solution to the broader problem of simulating interacting relativistic systems. The approach still relies on simplifications such as binary elastic collisions and sequential scheduling, which may fail in dense or quantum-dominated regimes [1, 4, 23]. Furthermore, the lack of true many-body covariance and omission of quantum effects such as Pauli blocking, Bose enhancement, and coherence limits its realism in extreme conditions [22, 6, 41]. These limitations do not invalidate the method's utility, but highlight the need for further refinements.

Looking forward, several promising directions emerge. Integrating proper-time slicing techniques (e.g., as in ALPACA [31, 33]) could extend frame invariance to full system evolution. Incorporating quantum statistical corrections and radiative processes would expand the model's fidelity in high-density environments like the early quark-gluon plasma [39, 40]. Hybridizing this approach with hydrodynamic evolution models (Knudsen-number–based hybrid framework) or quantum transport frameworks may offer a path toward simulating high-energy nuclear matter with both causal and quantum consistency [30].

In summary, the methodology developed in this thesis provides a robust and extensible framework for improving the physical fidelity of relativistic cascade simulations. By enforcing microscopic causality and enhancing Lorentz consistency, it contributes to the broader effort of solving the Boltzmann equation [7, 8] and simulating high-energy nuclear matter with both theoretical rigor and computational tractability [10].

References

[1] K. Geiger and B. Müller, "Dynamics of parton cascades in highly relativistic nuclear collisions," *Nucl. Phys. B*, vol. 369, pp. 600–654, 1992.

[2] J. C. Collins and M. J. Perry, "Superdense Matter: Neutrons or Asymptotically Free Quarks," *Phys. Rev. Lett.*, vol. 34, p. 1353, 1975.

[3] X.-N. Wang, "A pQCD-based approach to parton production and transport," *Phys. Rep.*, vol. 280, pp. 287–371, 1997.

[4] B. Zhang, M. Gyulassy, and Y. Pang, "Equation of state and collision rate tests of parton cascade models," *Phys. Rev. C*, vol. 58, p. 1175, 1998.

[5] D. Molnár, "On the Lorentz-invariance of parton cascade algorithms," *Eur. Phys. J. C*, vol. 49, pp. 181–186, 2007.

[6] H. Noack *et al.*, "Covariant Parton Cascade Based on Proper-Time Slicing," *Phys. Rev. C*, vol. 100, p. 064903, 2019.

[7] C. Cercignani, *The Boltzmann Equation and Its Applications*, Springer-Verlag, 1988.

[8] R. K. Pathria and P. D. Beale, *Statistical Mechanics*, 3rd ed., Elsevier, 2011.

[9] F. Reif, *Fundamentals of Statistical and Thermal Physics*, McGraw-Hill, 1965.

[10] N. Metropolis and S. Ulam, "The Monte Carlo Method," *J. Am. Stat. Assoc.*, vol. 44, no. 247, pp. 335–341, 1949.

[11] C. P. Robert and G. Casella, *Monte Carlo Statistical Methods*, 2nd ed., Springer, 2004.

[12] D. P. Kroese, T. Taimre, and Z. I. Botev, *Handbook of Monte Carlo Methods*, Wiley, 2011.

[13] R. Y. Rubinstein and D. P. Kroese, *Simulation and the Monte Carlo Method*, 2nd ed., Wiley, 2008.

[14] J. E. Gentle, *Random Number Generation and Monte Carlo Methods*, 2nd ed., Springer, 2003.

[15] R. E. Caflisch, "Monte Carlo and quasi-Monte Carlo methods," *Acta Numerica*, vol. 7, pp. 1–49, 1998.

[16] W. Rindler, *Relativity: Special, General, and Cosmological*, 2nd ed., Oxford University Press, 2006.

[17] R. M. Wald, *General Relativity*, University of Chicago Press, 1984.

[18] G. Kortemeyer, W. Bauer, K. Haglin, J. Murray, and S. Pratt, "Causality Violations in Cascade Models of Nuclear Collisions," *Phys. Rev. C*, vol. 52, pp. 2714–2718, 1995.

[19] C. C. Noack, "Causality Aspects of the Parton Cascade Approach to Ultrarelativistic Heavy Ion Reactions," arXiv:hep-ph/0312365, 2003.

[20] D. Molnár and M. Gyulassy, "New solutions to covariant nonequilibrium dynamics," *Phys. Rev. C*, vol. 62, p. 054907, 2000.

[21] A. Makhlin and E. Surdutovich, "Causality and Simultaneity in Relativistic Kinetics," *Phys. Rev. C*, vol. 53, no. 3, pp. 1316–1328, 1996.

[22] Z. Xu and C. Greiner, "Thermalization of gluons in ultrarelativistic heavy ion collisions by including three-body interactions in a parton cascade," *Phys. Rev. C*, vol. 71, p. 064901, 2005.

[23] E. V. Shuryak, "What RHIC experiments and theory tell us about properties of quark-gluon plasma?," *Nucl. Phys. A*, vol. 750, pp. 64–83, 2005.

[24] K. Huang, *Statistical Mechanics*, 2nd ed., Wiley, 1987.

[25] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical Recipes: The Art of Scientific Computing*, 3rd ed., Cambridge University Press, 2007.

[26] G. L. Ma, "Sampling in High-Energy Nuclear Simulations," *Comput. Phys. Commun.*, vol. 175, pp. 105–113, 2006.

[27] S. A. Bass *et al.*, "Microscopic Models for Ultrarelativistic Heavy Ion Collisions," *Prog. Part. Nucl. Phys.*, vol. 41, pp. 255–369, 1998.

[28] J. D. Jackson, *Classical Electrodynamics*, 3rd ed., Wiley, 1998.

[29] L. D. Landau and E. M. Lifshitz, *The Classical Theory of Fields*, 4th ed., Pergamon Press, 1987.

[30] D. H. Rischke, S. Bernard, and J. A. Maruhn, "Relativistic hydrodynamics for heavy-ion collisions. I. General aspects and expansion into vacuum," *Nucl. Phys. A*, vol. 595, pp. 346–382, 1995.

[31] A. Kurkela, R. Törnkvist, and K. Zapp, "AMY Lorentz invariant parton cascade: the thermal equilibrium case," *Eur. Phys. J. C*, vol. 84, article 74, 2024.

[32] P. B. Arnold, G. D. Moore, and L. G. Yaffe, "Effective kinetic theory for high temperature gauge theories," *JHEP*, vol. 2003, no. 01, p. 030, 2003.

[33] K. Zapp and A. Kurkela, "ALPACA: Lorentz Invariant Effective Kinetic Theory for Parton Transport," *Eur. Phys. J. C*, vol. 84, p. 74, 2024.

[34] S. R. de Groot, W. A. van Leeuwen, and C. G. van Weert, *Relativistic Kinetic Theory: Principles and Applications*, North-Holland, 1980.

[35] R. K. Ellis, W. J. Stirling, and B. R. Webber, *QCD and Collider Physics*, Cambridge University Press, 1996.

[36] T. Sjöstrand, S. Mrenna, and P. Z. Skands, "PYTHIA 6.4 Physics and Manual," *JHEP*, vol. 2006, no. 05, p. 026, 2006.

[37] W. Cassing and W. Juchem, "Semiclassical transport of particles with dynamical spectral functions," *Nucl. Phys. A*, vol. 665, p. 377, 2000.

[38] J. Zhao, C. Shen, and B. Schenke, "Collective flow in small systems from hydrodynamic transport approaches," *Phys. Rev. C*, vol. 102, p. 024903, 2020.

[39] R. Baier, Y. L. Dokshitzer, A. H. Mueller, S. Peigné, and D. Schiff, "Radiative energy loss and $p_\perp$-broadening of high energy partons in nuclei," *Nucl. Phys. B*, vol. 484, pp. 265–282, 1997.

[40] M. Gyulassy, P. Levai, and I. Vitev, "Reaction Operator Approach to Non-Abelian Energy Loss," *Nucl. Phys. B*, vol. 594, pp. 371–419, 2001.

[41] C. Greiner and S. Leupold, "Stochastic interpretation of Kadanoff–Baym equations and their relation to Langevin processes," *Ann. Phys.*, vol. 270, pp. 328–390, 1998.

6 Appendix

6.1 Special Relativity and Lorentz Transformations

In the context of special relativity, we consider two frames of reference: S and S'. The S'-frame moves at a constant velocity $\mathbf{v}$ relative to the S-frame.

In the S-frame, an event is described by coordinates (ct, x, y, z), while in the S'-frame, it's described by (ct', x', y', z'). The position vectors in these frames are $\mathbf{r}$ and $\mathbf{r}'$, respectively.

r

The position vector $\mathbf{r}$ can be decomposed into two components:

$$\mathbf{r} = \mathbf{r}_\perp + \mathbf{r}_\parallel$$

Here, $\mathbf{r}_\parallel$ is parallel to $\mathbf{v}$, and $\mathbf{r}_\perp$ is perpendicular to $\mathbf{v}$.

The perpendicular component $\mathbf{r}_\perp$ can be found by subtracting $\mathbf{r}_\parallel$ from $\mathbf{r}$:

$$\mathbf{r}_\perp = \mathbf{r} - \mathbf{r}_\parallel$$

The parallel component $\mathbf{r}_\parallel$ can be expressed as:

$$\mathbf{r}_\parallel = \left(\frac{\mathbf{r} \cdot \mathbf{v}}{|\mathbf{v}|}\right) \frac{\mathbf{v}}{|\mathbf{v}|}$$

Due to length contraction, the position vector in the S'-frame is:

$$\mathbf{r}' = \mathbf{r}_\perp + \gamma(\mathbf{r}_\parallel - \mathbf{v}t)$$

where γ is the Lorentz factor.

The time dilation equation is expanded as:

$$t' = \gamma t - \gamma \left(\frac{x v_x + y v_y + z v_z}{c^2} \right)$$

By substituting the previous equations, we obtain:

$$\mathbf{r}' = \mathbf{r} - \gamma \mathbf{v} t + (\gamma - 1) \left(\frac{\mathbf{r} \cdot \mathbf{v}}{|\mathbf{v}|} \right) \frac{\mathbf{v}}{|\mathbf{v}|}$$

The Lorentz transformation can be succinctly represented as:

$$x^\alpha = \Lambda^\alpha_\beta x^\beta$$

The Lorentz transformation matrix Λ^α_β is given by:

$$\left(\Lambda^\alpha_\beta\right) = \begin{pmatrix} \gamma & -\gamma\beta_x & -\gamma\beta_y & -\gamma\beta_z \\ -\gamma\beta_x & 1 + (\gamma - 1)\frac{\beta_x^2}{|\beta|^2} & (\gamma - 1)\frac{\beta_y \beta_x}{|\beta|^2} & (\gamma - 1)\frac{\beta_z \beta_x}{|\beta|^2} \\ -\gamma\beta_y & (\gamma - 1)\frac{\beta_x \beta_y}{|\beta|^2} & 1 + (\gamma - 1)\frac{\beta_y^2}{|\beta|^2} & (\gamma - 1)\frac{\beta_z \beta_y}{|\beta|^2} \\ -\gamma\beta_z & (\gamma - 1)\frac{\beta_x \beta_z}{|\beta|^2} & (\gamma - 1)\frac{\beta_y \beta_z}{|\beta|^2} & 1 + (\gamma - 1)\frac{\beta_z^2}{|\beta|^2} \end{pmatrix}$$