UNIVERSITY OF CAPE TOWN

MASTERS DISSERTATION

Full 3+1 dimensional simulation of the relativistic Boltzmann equation

Author:
William GRUNOW

Supervisor:
Prof. Andre PESHIER

A thesis submitted in fulfilment of the requirements
for the degree of Master of Science

in the

Department of Physics

June 2020

Declaration of Authorship

I, William GRUNOW, declare that this thesis titled, 'Full 3+1 dimensional simulation of the relativistic Boltzmann equation' and the work presented in it are my own. I confirm that:

- This work was done wholly or mainly while in candidature for a research degree at this University.
- Where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated.
- Where I have consulted the published work of others, this is always clearly attributed.
- Where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work.
- I have acknowledged all main sources of help.
- Where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself.

Signed by candidate

Signed:

Date: 10 July 2019

Abstract

Faculty of Science
Department of Physics

Master of Science

Full 3+1 dimensional simulation of the relativistic Boltzmann equation

by William GRUNOW

Relativistic hydrodynamics has been the tool of choice to simulate the dynamics of the quark-gluon plasma produced in heavy-ion collisions. Despite the success of hydrodynamics, it has several shortcomings stemming from the fact that it assumes a system close to equilibrium. An alternative to hydrodynamics is solving the Boltzmann equation, which describes the evolution of the full distribution function of the system without the close to equilibrium requirement.Large scale simulations using the Boltzmann equation, however, has hitherto proved computationally intractable due to their computational expense. By using a novel algorithm, and leveraging the computational power of graphical processor units, we numerically integrate the Boltzmann equation in the relaxation time approximation.

Acknowledgements

I owe a great debt of gratitute to my supervisor, Prof. Andre Peshier, who with deep insight, and boundless patience guided me, and regularly challenged my ideas with his own unique way of "thinking computationally" .

I would like to thank the National Institute of Theoretical Physics (NITheP) for their generous funding during this thesis.

I would like to thank the community in the UCT Physics department, it has been a true pleasure working in the department, and the interactions I've had in the deparment have taught me stuff that no curriculum could.

Finally, I would like to thank my family. Words cannot express how supportive they have been.

Contents

Chapter 1

Introduction

1.1 The Strong Nuclear force and Heavy-ion collisions

Our current understanding of nature is described by four fundamental forces, gravity, electromagnetism, the weak nuclear force and strong nuclear force all of which are described by two theoretical models, the Standard Model and General Relativity. While these theories have proven very successful in their respective domains, we know they must be incomplete as they are not yet compatible with one another.

Our focus in this thesis will be to study the properties of the strong nuclear force in particular. The strong force is described in the Standard Model by Quantum Chromodynamics (QCD), which is a non-Abelian gauge theory of $SU(3)$.

The strong force is the force that binds neutrons and protons together to form atomic nuclei. In order to do this it has to overcome the repulsive electromagnetic force that protons will generate between one another. Protons and neutrons can be further subdivided into constituent particles called, quarks and gluons. Quarks and gluons, quarks are massive Fermions, while gluons are massless Bosons and mediate the interactions between quarks. Analogously to how electromagnetism generates interactions between objects that

are electromagnetically charged, the strong nuclear force generates interactions between objects that have colour charge. There are three distinct colour charges (red, green and blue). Two of the most important phenomena of the strong force are *colour confinement* and *asymptotic freedom*. Colour confinement is the phenomenon where colour charged particles (such as quarks and gluons) cannot be isolated, and thus under normal conditions, we can only observe colour neutral particles. Asymptotic freedom is the phenomenon where the strong interaction becomes asymptotically weaker at higher energy and shorter length scales

The primary experimental efforts for studying the physics of the strong nuclear force comes from particle collider experiments. Currently, the Relativistic Heavy Ion Collider (RHIC) and the Large Hadron Collider (LHC) run a variety of collision experiments including proton-proton collisions, proton-nucleus collision and heavy-ion collisions.

In heavy-ion collisions, which we will focus on in this thesis, the energy densities generated are so large that colour neutral hadrons dissociate into their constituent quarks and gluons creating a plasma of quarks and gluons, the Quark-Gluon Plasma (QGP). This dissociation of quarks and gluons is due to asymptotic freedom. The QGP being a deconfined state of quarks and gluons provides an excellent opportunity to study the phenomenology of the strong nuclear force and as a test-bed for the theoretical predictions of many-body QCD. Experimentally, this has historically been challenging to do due to colour confinement.

Furthermore, these high-energy collisions produce energy densities so large that they have not occurred in nature since microseconds after the Big Bang. Thus an understanding of the dynamics of matter at these densities provides us with insight into the dynamics of the early universe beyond the veil of the Cosmic Microwave Background.

A heavy ion collision event is often broken down into four distinct stages where different physics processes dominate the dynamics.

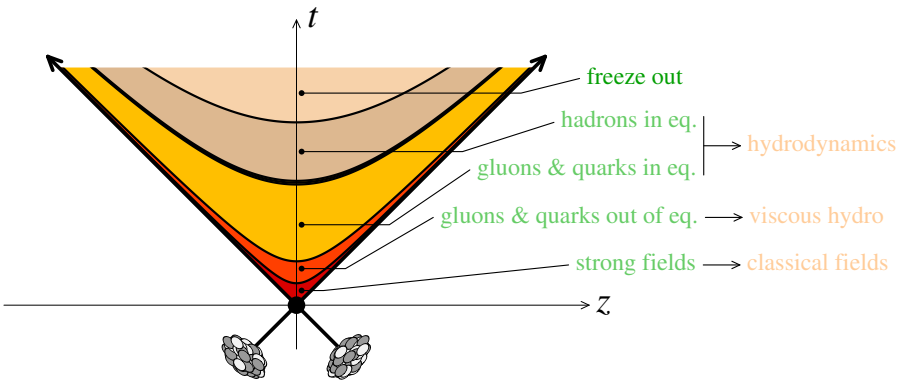

Figure 1.1: A spacetime illustration of the various stages if a heavy ion collision. Source [1]

In chronological order, these stages, as depicted in Fig. (1.1), are

- Pre-collision dynamics: At the ultra-relativistic speeds, the nuclei are strongly Lorentz contracted along the beam axis. The dynamics of this pre-collision stage is described by the Color Glass Condensate (CGC) model [2, 3] which predicts that at high collision energies the number density of gluons in each nucleus becomes very large, up to some characteristic momentum, Q_s, known as the saturation scale. The dynamics of these gluons can be approximated by the solutions to the transverse classical Yang-Mills equations.

- Pre-equilibrium dynamics: In the initial phase of the heavy ion collision, the system enters an intermediate state between the CGC and the QGP known as the Glasma. Similarly to the CGC, this stage is well described by solutions to the classical Yang-Mills equations; however, the dynamics are now predominantly longitudinal. Some of the strongest electric and magnetic fields ever observed in nature occur during this stage. There are various phenomenological Glasma models that are currently used to provide the initial conditions for hydrodynamic evolution which is the next stage in the collision process [4–6].

- Hydrodynamic evolution: From ~ 1 fm/c the system is well described by relativistic hydrodynamics. The use of a hydrodynamical model assumes that the system is at least close to thermal equilibrium, and this leads to one of the big open questions in the study of the QGP, how is it

possible that thermalises on such a short timescale? Previous theoretical attempts at solving this question have thus far been unsatisfactory. Furthermore, hydrodynamic modelling has determined the shear viscosity to entropy density ratio to be very small [8] (almost saturating the universal lower bound predicted by AdS/CFT calculations [9, 10]) earning the QGP the moniker of a "near perfect fluid" [11]. There are new experimental results, one recent development in the filed is the observation of hydrodynamic flow in very central pA collisions at the LHC; a system previously thought to be too small to undergo hydrodynamic flow [12]. Relativistic hydrodynamics is also continuing to develop, with the sophistication of relativistic hydrodynamics codes ever increasing by considering viscous corrections and considering local momentum anisotropies [13].

- Freezeout and hadronisation: From ~ 10 fm/c the system has expanded and cooled down sufficiently a QGP state can no longer be sustained, and the constituent quarks and gluons recombine into colourless hadronic particles. The usual procedure for determining the hadronic yields at freezeout is the Cooper-Frye prescription [14]. At this stage, the system is still a strongly interacting hot hadronic gas, but its constituent particle ratios are now fixed. Modelling of the hadron interactions post chemical freezeout is done by more kinetic theory oriented approaches, such as *hadronic afterburners* [15, 16]. At some point at ~ 20 fm/c the system becomes so dilute that the hadrons no longer interact with one another, and will free-stream towards the detectors.

Our goal in this thesis will be to develop tools for simulating, in particular, the viscous-hydrodynamical stage of the QGP expansion. Our precise motivation for doing so will be elaborated on in the next parts of this chapter. We will begin with an exploration of how we describe the macroscopic properties of the QGP. In particular, we will focus on how to track the evolution of the distribution function, which will require numerical simulation. We will develop the numerical tools, including leveraging the computational power of Graphical Processor Units (GPUs) in order to do so. We then compare to some known analytical limits and see the effects of relaxing those assumptions.

1.2 Notation and conventions

Throughout this thesis, it will be assumed that we are working with natural units, i.e. $\hbar = c = k_B = 1$.

It will also be assumed that we will always be using the mostly-minus convention for the Minkowski metric,

$$g^{\mu\nu} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & -1 & 0 & 0 \\ 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & -1 \end{pmatrix}. \tag{1.1}$$

We define the shorthand for the usual momentum space component of the Lorentz invariant phase space integration measure,

$$\int d\Gamma = \frac{d}{(2\pi)^3} \int \frac{d^3\boldsymbol{p}}{p^0} \tag{1.2}$$

where d is the degeneracy factor of the system we are integrating over, $p^0 = \sqrt{\boldsymbol{p}^2 + m^2}$, where m is the mass of the particle, such that the normalisation of p^μ is $p^\mu p_\mu = m^2$. For massless particles the momentum $\boldsymbol{p}$ is equal to the wavevector of the particle $\boldsymbol{k}$. If we are assuming that particles are massless (as will be the case for the majority of the thesis), we will explicitly use k instead of p.

We will orient our coordinate system such that the coordinates x, y will always describe directions transverse to the beam axis, while the coordinate z will be longitudinal to the beam axis.

The notation for transverse coordinates will often be condensed as follows,

$$\frac{\partial}{\partial \boldsymbol{x}_\perp} = \left(\frac{\partial}{\partial x}, \frac{\partial}{\partial y} \right) \tag{1.3}$$

and

$$d^2\boldsymbol{x}_\perp = dx\, dy. \tag{1.4}$$

1.3 Description of the gluon plasma system

We will model the gluon plasma system by simulating the evolution of the *one-particle distribution function* f (just the distribution function henceforth). The distribution function describes the number of particles in a particular phase-space volume element $d^3x\, d^3p$ around a phase space point $(\boldsymbol{x}, \boldsymbol{p})$ at time t.

$$f(\boldsymbol{x}, \boldsymbol{p}, t) = \frac{dN(t)}{d^3x d^3p}. \tag{1.5}$$

f is a Lorentz scalar (i.e. it is independent of reference frame) as the phase-space volume element is a Lorentz invariant.

When referring to the distribution function f is always assumed to be a function of postition, momentum and time, but for notational economy, some, or all of the arguments may be omitted (i.e. $f(\boldsymbol{x}, \boldsymbol{p}, t) = f = f(\boldsymbol{p})$ etc.).

Relevant observables are obtained from taking various moments of the distribution function.

In particular the first moment given by,

$$J^u = \int d\Gamma\, k^\mu f, \tag{1.6}$$

is knwon as the *particle current* four-vector.

From J^μ we can extract the number density n,

$$n = J^0 \tag{1.7}$$

If the local conservation equation,

$$\partial_\mu J^u = 0, \tag{1.8}$$

is satisfied then the total particle number of the system is conserved.

The second moment of the distribution function,

$$T^{\mu\nu} = \int d\Gamma \, k^\mu k^\nu f, \tag{1.9}$$

is also known as the *energy-momentum* tensor.

The local energy density ϵ is given by the T^{00} component.

Conservation of energy and momentum can be easily expressed in terms of the energy-momentum tensor in the following way,

$$\partial_\mu T^{\mu\nu} = 0. \tag{1.10}$$

We introduce the relativistic generalization of entropy and its associated flow via the *entropy four-current* defined as,

$$S^\mu = -\int d\Gamma \, k^\mu \left(f \ln f + \frac{1}{\alpha}(1-\alpha f)\ln(1-\alpha f) \right) \tag{1.11}$$

where α specifies the type of particle we are considering, $\alpha = 1$ for Fermions, $\alpha = -1$ for Bosons ($\alpha \to 0$ for Maxwellian particles).

Since our goal is to study gluon dynamics, we will henceforth set $\alpha = -1$.

The relativistic generalisation of the second law of thermodynamics reads,

$$\partial_\mu S^\mu \geq 0, \tag{1.12}$$

which implies that the total entropy of the system

$$S = \int d^3x S^0, \tag{1.13}$$

will eventually be maximized.

The equilibrium distributions maximise the local entropy density of the system under the particular local constraints of that system, and take the form

$$f^{eq}(k;\beta,\mu,u_\alpha) = \frac{1}{\exp(\beta(k_\alpha u^\alpha - \mu)) + \alpha} \tag{1.14}$$

The parameters are, β the inverse temperature, μ the chemical potential, and u^α the relativistic fluid flow velocity. These parameters are Lagrange multiplier that ensure energy, momentum and particle number are conserved.

There are two competing definitions of the fluid flow velocity found in the literature, the *Eckart definition* and the *Landau definition*.

The Eckart definition defines the flow velocity in terms of the local particle flux as follows,

$$u^\mu_{\text{Eckhart}} = \frac{J^\mu}{\sqrt{J^\mu J_\mu}}. \tag{1.15}$$

In contrast the Landau definition defines the flow velocity in terms of the local energy flux,

$$T^\mu_{\ \nu} u^\nu_{\text{Landau}} = \epsilon_{\text{LRF}} u^\mu_{\text{Landau}}. \tag{1.16}$$

This defines the flow velocity as the unique normalised timelike eigenvector of the energy-momentum tensor. In both cases, the local rest frame (LRF) is defined as the frame where the flow velocity is given by $u^\mu_{\text{LRF}} = (1, 0, 0, 0)$.

As, in general, the particle number of gluons is not conserved we will henceforth exclusively use the Landau definition of flow velocity.

The energy-momentum tensor of an equilibrium distribution can be decomposed in the following way

$$T^{\mu\nu}_{eq} = (\epsilon_{\text{LRF}} + P) u^\mu u^\nu - P g^{\mu\nu}, \tag{1.17}$$

where ϵ_{LRF} is the local energy density in the Landau rest frame, and u^μ is the local flow velocity of the system.

Similarly, the particle current density of an equilibrium distribution can be decomposed as,

$$J^\mu_{eq} = n_{\text{LRF}} u^\mu, \tag{1.18}$$

where n_{LRF} is the local number density in the Landau rest frame.

The local energy and particle densities of the Bose-Einstein distribution function are given by,

$$\epsilon^{eq}_{\text{LRF}} = d_g \frac{3}{\pi^2} \text{Li}_4 \left(\exp\left(\frac{\mu}{T}\right) \right) T^4, \tag{1.19}$$

$$n^{eq}_{\text{LRF}} = d_g \frac{1}{\pi^2} \text{Li}_3 \left(\exp\left(\frac{\mu}{T}\right) \right) T^3, \tag{1.20}$$

where $\text{Li}_n(x)$ is the polylogarithm of order n and d_g is the degeneracy factor associated with gluons ($d_g = 16$).

When we consider physical scenarios that do not impose particle number conservation we have that $\mu = 0$, and we can write down an analytical expression for the temperature,

$$T = \left(\frac{30}{d_g \pi^2} \epsilon^{eq}_{\text{LRF}} \right)^{1/4}. \tag{1.21}$$

It is evident from the definition the energy-momentum tensor is symmetric under exchange of indices, and for massless particles, we have the condition that $T^{\mu}{}_{\mu} = 0$. This implies that there are 9 degrees of freedom associated with the energy-momentum tensor.

If the distribution function is not an equilibrium distribution, then the energy-momentum tensor can be decomposed into equilibrium ($T^{\mu\nu}_{eq}$) and non-equilibrium ($T^{\mu\nu}_{\delta f}$) terms,

$$T^{\mu\nu} = T^{\mu\nu}_{eq} + T^{\mu\nu}_{\delta f} = \epsilon u^{\mu} u^{\nu} + P \Delta^{\mu\nu} + \pi^{\mu\nu}. \tag{1.22}$$

P is the bulk viscosity of the system, $\Delta^{\mu\nu}$ is the so called *projection tensor* defined by,

$$\Delta^{\mu\nu} = g^{\mu\nu} - u^{\mu} u^{\nu}. \tag{1.23}$$

As the name suggests, the projection tensor is a projector,

$$\Delta^{\mu\alpha} \Delta_{\alpha}{}^{\nu} = \Delta^{\mu\nu}. \tag{1.24}$$

The projection tensor is transverse to the flow velocity,

$$u_\mu \Delta^{\mu\nu} = 0, \tag{1.25}$$

and has a non-vanishing trace

$$\Delta^\mu{}_\mu = 3. \tag{1.26}$$

The shear stress tensor $\pi^{\mu\nu}$ is symmetric ($\pi^{\mu\nu} = \pi^{\nu\mu}$), traceless ($g_{\mu\nu}\pi^{\mu\nu} = 0$) and transverse to the flow velocity ($u_\mu \pi^{\mu\nu} = 0$)

Let us perform a quick sanity check on how we have defined the components of the energy-momentum tensor.

If we consider,

$$\begin{aligned} u_\mu u_\nu T^{\mu\nu} &= \int d\Gamma \, u_\mu k^\mu u_\nu k^\nu f \\ &= \int d\Gamma \, (u_\mu k^\mu)^2 f \end{aligned} \tag{1.27}$$

since the integrand is strictly non-zero we have that

$$u_\mu u_\nu T^{\mu\nu} \geq 0. \tag{1.28}$$

By Eq. (1.22) we also have that,

$$u_\mu u_\nu T^{\mu\nu} = \epsilon u_\mu u^\mu \tag{1.29}$$

and hence

$$\epsilon u_\mu u^\mu \geq 0. \tag{1.30}$$

This implies that positive eigenvalues are associated with timelike eigenvectors and negative eigenvalues with spacelike eigenvectors.

Hence we are guaranteed that physically realistic energy densities (i.e. non-negative) have physically realistic flow velocities (i.e. timelike), and vice-versa.

1.4 Bose-Einstein condensation

In order to find the equilibrium distribution for some given energy and number density, one needs to find the equilibrium parameters that satisfy Eqs. (1.19, 1.20). While it is always true that a root to the system of equations exists, it is not guaranteed that those solutions are consistent with the constraints imposed by physics, for instance, that the distribution is strictly non-negative. This constraint would be violated if the solution for chemical potential μ were non-negative.

To proceed further, we introduce the dimensionless population parameter ξ defined by,

$$\xi = \frac{n^4}{\epsilon^3}, \tag{1.31}$$

which for the equilibrium distribution can be parameterised by the ratio $\frac{\mu}{T}$ as follows,

$$\xi = d_g \frac{1}{27\pi^2} \frac{\mathrm{Li}_3\left(\exp\left(\frac{\mu}{T}\right)\right)^4}{\mathrm{Li}_4\left(\exp\left(\frac{\mu}{T}\right)\right)^3}. \tag{1.32}$$

Since any strictly positive ξ is physically relevant, we should be able to extract physically appropriate solutions to the equilibrium parameters for any positive value of ξ.

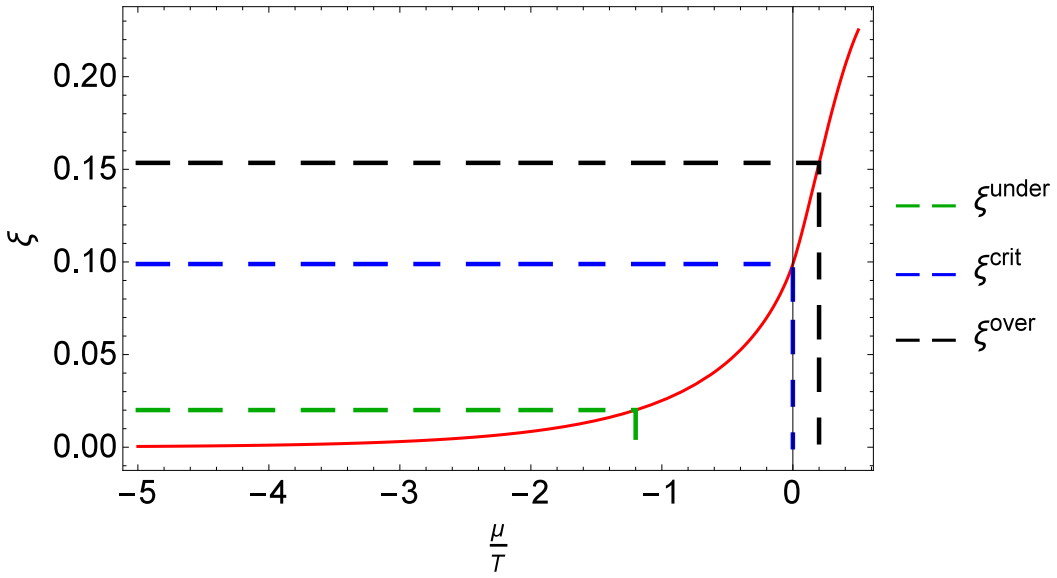

Figure 1.2: Plot of how $\xi = \frac{n^4}{e^3}$ varies with $\frac{\mu}{T}$. (Using $d_g = 16$)

In Fig. (1.2) we see that ξ is a monotonically increasing function of $\frac{\mu}{T}$, but reaches a maximal permissible value (ξ_{crit}) at $\frac{\mu}{T} = 0$.

If $\xi < \xi_{crit}$, the system is said to be *underpopulated*, and physically admissible solutions to the Eqs. (1.19, 1.20) exist.

Conversely, if $\xi > \xi_{crit}$, the system is said to be *overpopulated*, and no physically admissible solution to Eqs. (1.19, 1.20) can be found, as they would predict a complex temperature.

Since overpopulated initial conditions are physically relevant as the number densities in heavy-ion collisions are large, we need to modify the equilibrium distribution in order to accommodate the additional particle number density; this is achieved by the introduction of a *condensate*, a state of zero-momentum but finite particle density, i.e

$$f^{eq}(k;\beta,\mu,u_\alpha) = \frac{1}{\exp(\beta k_\alpha u^\alpha) - 1} + n_c\delta(\boldsymbol{k}), \tag{1.33}$$

where n_c is the particle number density present in the condensate.

Evidence of Bose-Einstein condensates has been measured in several low-temperature systems [17–19].

In the context of heavy-ion collisions, the Color Glass Condensate (CGC) model predicts that in very high energy collisions, the number density of gluons becomes large. Recent papers [20, 21] have shown that this initial condition of gluons could form an overpopulated initial condition and that at least in the early stages of the evolution of the quark-gluon plasma, number changing processes are suppressed.

A Bose-Einstein condensate being part of the equilibrium distribution is only possible if particle number is conserved. Gluon number, however, is not conserved; thus, if a Bose-Einstein condensate were to form, it would be transient. The presence of a Bose-Einstein condensate of gluons could significantly affect the dynamics of the quark-gluon plasma, for example, it has been shown that the presence of BEC of gluons would reduce the shear viscosity of the quark-gluon plasma [22].

It is thus important to check whether or not the conditions suitable for the formation of a Bose-Einstein condensate survive for long enough for a significant Bose-Einstein condensate to form. There have been several previous attempts to study the dynamical formation of a gluon BEC [23–26] these approaches have only considered the momentum-space dynamics to establish the time-scale over which the Bose-Einstein condensate will form, however, in order to get a complete picture one needs to consider the full phase-space dynamics, including configuration space dynamics. Thus our goal will be to implement a simulation package that implements configuration space dynamics that can be used in conjunction with different models of momentum-space gluon dynamics. For proof of concept, we will assume the momentum-space model to be the so-called relaxation time approximation, which we will discuss in the next chapter.

Chapter 2

Kinetic theory of a gluon plasma

In the previous chapter, we have developed several tools for describing the properties of the QGP; all of these tools require us to know the distribution function. So we will dedicate this chapter to developing the techniques for determining the evolution of the distribution.

In general, the evolution of the distribution function is given by an infinite set of coupled differential equations known as the BBGKY hierarchy. This hierarchy couples the evolution of the one-particle distribution function to the evolution of higher-order distribution functions (a distribution function of order n describes the position and momentum space correlations between n particles).

The BBGKY hierarchy can be truncated under the assumption of *Molecular chaos* (also known as the *Stosszahlansatz*). This is the assumption that the position and momenta of the particles are sufficiently decorrelated that the pairwise distribution function $f_2(\boldsymbol{x_1}, \boldsymbol{x_2}, \boldsymbol{k_1}, \boldsymbol{k_2})$ can be decomposed as a product of one particle distribution functions,

$$f_2(\boldsymbol{x_1}, \boldsymbol{x_2}, \boldsymbol{k_1}, \boldsymbol{k_2}; t) \approx f(\boldsymbol{x_1}, \boldsymbol{k_1}; t) f(\boldsymbol{x_2}, \boldsymbol{k_2}; t). \tag{2.1}$$

This assumption is usually satisfied when the mean free path between collisions λ_{mfp} is much larger than the thermal wavelength $\lambda_{\text{th}} \sim 1/T$.

Once this truncation is made, we are left with the Boltzmann equation. As the Boltzmann equation explicitly does not depend on the pairwise (or higher-order) distribution functions, the complexity of finding a solution is significantly reduced compared to the BBGKY hierarchy.

2.1 The Boltzmann equation

The Boltzmann equation is a non-linear integro-differential equation that describes the evolution of the distribution function. The relativistically covariant form is given by,

$$k^\mu \partial_\mu f(\boldsymbol{x}, \boldsymbol{k}, t) = \mathcal{C}[f(\boldsymbol{x}, \boldsymbol{k}, t)] \tag{2.2}$$

where there LHS corresponds to the so-called *convective derivative* which describes the free propagation of a given momentum through configuration space. While the RHS corresponds to the so-called *collision term* which describes how the particles interact (e.g. exchange momenta).

The specific physics (e.g. QCD, QED, etc.) of the system is contained within the collision term.

The general format of a collision term for interaction models that permit binary elastic collisions that correctly reproduces quantum statistical effects is given by,

$$\begin{aligned}\mathcal{C}[f(\boldsymbol{x}, \boldsymbol{k}, t)] = \frac{(2\pi)^4}{2E_k} \int \frac{d^3k_2}{(2\pi)^3 2E_2} \frac{1}{\nu} \int \frac{d^3k_3}{(2\pi)^3 2E_3} \int \frac{d^3k_4}{(2\pi)^3 2E_4} |\mathcal{M}_{kk_2 \to k_3k_4}|^2 \\ \times \left[f(k_3) f(k_4) \bar{f}(k) \bar{f}(k_2) - f(k) f(k_2) \bar{f}(k_3) \bar{f}(k_4) \right] \delta^4(k + k_2 - k_3 - k_4).\end{aligned} \tag{2.3}$$

Where we have introduced $\bar{f} = (1 \pm f)$, the sign depending on whether f is a distribution of Bosons$(+)$ or Fermions$(-)$. This term implements either Bose-enhancement or Fermi blocking.

The $\mathcal{M}_{kk_2\to k_3k_4}$ term is the transition amplitude of two particles with momenta k and k_2 respectively scattering into two outgoing particles with momenta k_3 and k_4. This transition amplitude is supplied by the underlying physical model.

If the output particles are identical $\nu = 2$ to account for double counting, else $\nu = 1$.

The delta function enforces energy-momentum conservation.

A key feature of the Boltzmann equation is the fact that it will drive the distribution function into a higher entropy state in accordance with the Second law of thermodynamics (this property is called the Boltzmann H-theorem).

We would expect that the Boltzmann equation eventually drives the system to an equilibrium distribution, and once in equilibrium, the expectation is that the distribution function would stop evolving i.e. $\mathcal{C}[f^{eq}] = 0$.

We can ensure $C[f] = 0$ if we impose the condition,

$$f(k_3)f(k_4)\bar{f}(k)\bar{f}(k_2) = f(k)f(k_2)\bar{f}(k_3)\bar{f}(k_4), \tag{2.4}$$

or equivalently,

$$\log\frac{f(k)}{\bar{f}(k)} + \log\frac{f(k_2)}{\bar{f}(k_2)} = \log\frac{f(k_3)}{\bar{f}(k_3)} + \log\frac{f(k_4)}{\bar{f}(k_4)}. \tag{2.5}$$

This condition is met if each term is parameterised by terms that are conserved in the collision (i.e. energy and momentum)

$$\log\frac{f(k)}{\bar{f}(k)} = \beta(\mu - E - \boldsymbol{u}\cdot\boldsymbol{k}), \tag{2.6}$$

which gives us,

$$f(\boldsymbol{k}) = \frac{1}{e^{\beta(E(\boldsymbol{k})-\mu+\boldsymbol{u}\cdot\boldsymbol{k})} \pm 1}, \tag{2.7}$$

the expected Bose-Einstein/Fermi-Dirac distribution functions.

2.1.1 The relaxation time approximation

The collision term is in general non-trivial to calculate. We, however, know, by Boltzmann's H-theorem, that regardless of the details of the collision term, the system will be driven towards the equilibrium distribution. The relaxation time approximation takes advantage of this fact in order to drastically simplify the collision term, replacing it with a term that just monotonically drives the system towards its equilibrium distribution with a characteristic timescale.

The relativistic version of the Boltzmann equation under the relaxation time approximation was introduced by Anderson and Witting in [27] reads,

$$k^\alpha \partial_\alpha f(t, \boldsymbol{x}, \boldsymbol{k}) = k^\alpha u_\alpha \frac{f^{eq}(k, \boldsymbol{u}; \beta, \mu) - f(t, \boldsymbol{x}, \boldsymbol{k})}{\tau_{rel}}. \tag{2.8}$$

The type of physical system being modelled is characterised by choice of the equilibrium distribution, f^{eq}, used (typically Bose-Einstein distribution for Bosons and Fermi-Dirac for Fermions), and the relaxation time τ_{rel}. The relaxation time τ_{rel} is a free parameter in this model and controls the rate at which the distribution function will evolve towards its *local* equilibrium distribution.

The condition for energy-momentum conservation is

$$\partial_\alpha T^{\alpha\beta} = 0, \tag{2.9}$$

which implies

$$\int d\Gamma k^\alpha k^\beta \partial_\alpha f = 0, \tag{2.10}$$

using Eq. (2.8) we can rewrite the above as

$$\int d\Gamma k^\alpha k^\beta u_\alpha \frac{f^{eq}(k; T, \mu, \boldsymbol{u}) - f(t, \boldsymbol{x}, \boldsymbol{k})}{\tau_{rel}} = 0, \tag{2.11}$$

and by Eq. (1.9) this simplifies into,

$$u_\alpha \left(T^{\alpha\beta}_{[f]} - T^{\alpha\beta}_{eq} \right) = 0, \tag{2.12}$$

finally, using Eq. (1.22) we get an eigenvalue equation,

$$T^{\beta}_{[f]\ \alpha} u^{\alpha} = \epsilon_{eq} u^{\beta}, \tag{2.13}$$

where the local energy density and flow velocity are the eigenvalue and eigenvector, respectively.

Ifwe choose to impose particle number conservation, the additional constraint is

$$\partial_{\mu} J^{\mu} = 0. \tag{2.14}$$

Once more we substitute in Eq. (2.8) so that we get

$$u_{\alpha} J^{\alpha}_{[f]} = n_{eq}. \tag{2.15}$$

Since we have been able to extract the local energy density and particle number density from the energy-momentum tensor and particle-current four-vector respectively, we need to invert Eqs. (1.19,1.20) in order to extract the equilibrium parameters. In general, inverting these equations to extract the equilibrium parameters is in general not analytically tractable, and we resort to using a numerical root finding method to extract the parameters, once we have the energy and number densities.

Henceforth, when referring to the temperature, chemical potential and flow velocity of a system (including non-equilibrium systems), these values correspond to the unique values of these parameters that solve Eqs. (1.19,1.20).

2.1.2 Running equilibrium

If we consider a entirely spatially homogeneous distribution function the spatial derivatives in Eq. (2.2) vanish, and so there are only dynamics in the momentum domain. Furthermore, since energy and momentum density are locally constant in this scenario, and the equilibrium distribution is fully specified given the energy and momentum density, and so one is able to *apriori* determine the equilibrium parameters.

Naively, the solution to Eq. (2.8) seems to be trivial, the momentum distribution approaches the equilibrium distribution exponentially in time. This is only true if the flow velocity is zero. If there is a non-zero flow velocity $\boldsymbol{u}$ then it is an empirical observation that it is a requirement that the "equilibrium" parameters used in the RHS of Eq. (2.8) are time-dependent (this is not to say that the energy density is time-dependent; it is constant).

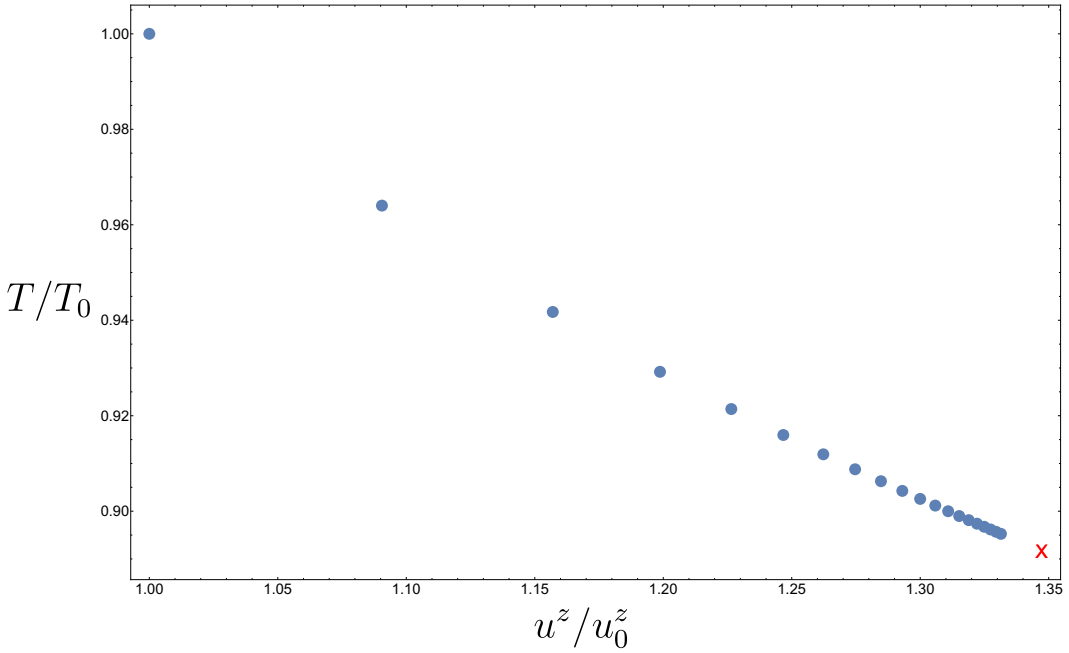

Figure 2.1: The evolution of the temperature and longitudinal flow velocity (blue dots) towards the asymptotic equilibrium (red cross) for a specific initial condition. This is an illustrative example of how the equilibrium parameters evolve.

It is true that as $t \to \infty$ the parameters that characterise f^{eq} in Eq. (2.8) become equal to those that we predicted apriori for the equilibrium distribution (as must be the case for energy conservation). This has a rather significant implication for us, in order to conserve energy in our model we require a timestep smaller than the relaxation time in order to smoothly transition to the final equilibrium temperature. Thus it is a strict requirement that we have a non-zero relaxation time; which will prevent us from considering the strict ideal hydrodynamic limit ($\tau_{\text{rel}} = 0$).

Chapter 3

Solving the Boltzmann equation

In this chapter, we will develop the tools needed to solve the Boltzmann equation under the relaxation time approximation. This will include a discussion on how we recast the Boltzmann equation in a form amenable to a numerical solution in the specific case of strong longitudinal expansion as relevant for heavy-ion collisions, as well as some of the details on the algorithmic implementation of the numerical solver. The fact that the distribution function we are solving for exists in a 6-dimensional phase space means that the computational complexity of the problem grows rapidly with an increasing system size, we will, therefore, emphasize techniques for reducing the size of the system we need to simulate.

Apart from developing efficient numerical tools, we also leverage the computational power of modern graphics processor units. Some detail of how this this is done is discussed in Appendix(A)

3.1 Choice of coordinate system

An important consideration that we must make is the choice of coordinate system we use. By choosing an appropriate coordinate system, as dictated by the physical scenario we are considering, we can significantly reduce the complexity of the problem.

The use of hydrodynamical techniques in heavy ion collisions was initially suggested by Landau [28]. In Landau's model he imagined that in the collision, the two nuclei come to a complete stop and all energy is deposited at the point of the collision, and further expansion is generated by hydrodynamical evolution.

In high energy collisions, the assumption that the nuclei are completely stopped breaks down as the majority of particles that in are produced high-flow velocity cells are later and further away from the collision point. This suggests it would be natural that the initial condition is not specified at a particular laboratory time (as it would be in the Landau model), but rather, as suggested by Bjorken, at a particular proper time. Bjorken also makes the assumption, that around the collision point, the system is invariant under Lorentz boost (i.e. the system is *boost-invariant*).

As a matter of convenience, as flow velocities are, at least initially, longitudinally dominated, we will use the longitudinal proper time (or *eigentime*) $\tau = \sqrt{t^2 - z^2}$.

The Cartesian coordinate system is not convenient for specifying an initial condition at a specified longitudinal proper time.

The system is rapidly expanding in the longitudinal direction; thus if we were to use a Cartesian grid, we would require an ever larger grid to accommodate the simulation at a given resolution. Ideally, we will find a coordinate system that mitigates this effect so that we minimise the number of sampled points.

Thus our goal will be to recast the Boltzmann equation in a more appropriate coordinate system.

We first consider the transformation of the temporal and longitudinal space, coordinates, in terms of the eigentime τ and space-time rapidity η.

This hyperbolic coordinate transformation is given by,

$$\begin{aligned} t &= \tau \cosh \eta \\ z &= \tau \sinh \eta. \end{aligned} \tag{3.1}$$

This transformation can be inverted as follows,

$$\begin{aligned} \tau &= \sqrt{t^2 - z^2} \\ \eta &= \frac{1}{2}\log\left(\frac{t+z}{t-z}\right). \end{aligned} \tag{3.2}$$

We want to recast the (RTA) Boltzmann equation in these coordinates; thus we need to express the derivatives ∂_t and ∂_z in terms of ∂_η and ∂_τ.

In particular, we have that,

$$\begin{aligned} \partial_t &= \frac{\partial\eta}{\partial t}\frac{\partial}{\partial\eta} + \frac{\partial\tau}{\partial t}\frac{\partial}{\partial\tau} \\ &= \cosh\eta\frac{\partial}{\partial\tau} - \frac{\sinh\eta}{\tau}\frac{\partial}{\partial\eta} \end{aligned} \tag{3.3}$$

and,

$$\begin{aligned} \partial_z &= \frac{\partial\eta}{\partial z}\frac{\partial}{\partial\eta} + \frac{\partial\tau}{\partial z}\frac{\partial}{\partial\tau} \\ &= -\sinh\eta\frac{\partial}{\partial\tau} + \frac{\cosh\eta}{\tau}\frac{\partial}{\partial\eta}. \end{aligned} \tag{3.4}$$

The RTA Boltzmann equation in Cartesian space is given by

$$\left(\frac{\partial}{\partial t} + \frac{k_z}{k_0}\frac{\partial}{\partial z} + \frac{\boldsymbol{k}_\perp}{k_0}\frac{\partial}{\partial\boldsymbol{x}_\perp}\right) f = \frac{k^\mu u_\mu}{k^0}\frac{f^{eq} - f}{\tau_{\text{rel}}} \tag{3.5}$$

which, under the hyperbolic transformation, becomes

$$\left(\frac{\partial}{\partial\tau} + \frac{1}{\tau}v_\eta\frac{\partial}{\partial\eta} + \boldsymbol{v}_\perp\frac{\partial}{\partial\boldsymbol{x}_\perp}\right) f = \frac{k^\mu u_\mu}{k_0\cosh\eta - k_z\sinh\eta}\frac{f^{eq} - f}{\tau_{\text{rel}}} \tag{3.6}$$

where we have defined the longitudinal velocity,

$$v_\eta = \frac{k_z\cosh\eta - k_0\sinh\eta}{k_0\cosh\eta - k_z\sinh\eta} \tag{3.7}$$

and transverse velocity,

$$\boldsymbol{v}_\perp = \frac{\boldsymbol{k}_\perp}{k_0 \cosh\eta - k_z \sinh\eta}. \tag{3.8}$$

Another consideration is that we will have a large variation in longitudinal momentum scales with typical longitudinal momentum scale of the system becoming exponentially larger as we consider points of increasing spacetime-rapidity. Thus for an accurate simulation, we require an appropriate resolution across all the relevant momentum scales, while to be numerically efficient, we need to minimise the number of points we need to sample to achieve that resolution.

It will also be convenient to perform a similar procedure for the momentum components by introducing the momentum-rapidity variable, Y, defined by

$$Y = \frac{1}{2} \ln \frac{k_0 + k_z}{k_0 - k_z}. \tag{3.9}$$

The inverse of this transformation is given by,

$$k_0 = k_\perp \cosh Y \tag{3.10}$$

$$k_z = k_\perp \sinh Y \tag{3.11}$$

which allows us to rewrite the Boltzmann equation fully in hyperbolic coordinates,

$$\left(\frac{\partial}{\partial \tau} + \frac{1}{\tau} v_\eta \frac{\partial}{\partial \eta} + \boldsymbol{v}_\perp \frac{\partial}{\partial \boldsymbol{x}_\perp} \right) f = \frac{k_\mu u^\mu}{k_\perp \cosh(Y - \eta)} \frac{f^{eq} - f}{\tau_{\text{rel}}} \tag{3.12}$$

where we have simplified Eq. (3.7) to,

$$v_\eta = \tanh(Y - \eta). \tag{3.13}$$

We can relate the momentum rapidity back to longitudinal velocity by,

$$v_z = \frac{k_z}{k_0} = \tanh(Y). \tag{3.14}$$

3.1.1 Conserved quantities in hyperbolic coordinates

In Cartesian coordinates, energy-momentum conservation is expressed by the following equation

$$\partial_\mu T^{\mu\nu} = 0 \tag{3.15}$$

which generates the conserved quantities,

$$\frac{d}{dt}\int d^3x\, T^{0\nu} = 0. \tag{3.16}$$

These equations specify that the total energy and momentum of the system are constant between discrete equal time hypersurfaces. However, in hyperbolic coordinates, we generate solutions along discrete equal eigentime hypersurfaces. Each eigentime hypersurface will cross multiple equal time hypersurfaces rendering the integral in Eq. (3.16) impractical.

We nonetheless want to ensure that energy remains conserved, In order to get the equivalent conserved quantities in hyperbolic coordinates, we replace the Cartesian derivative in Eq. (3.15) with the hyperbolic equivalents Eqs. (3.3, 3.4).

Upon making this substitution, we get

$$\cosh\eta\frac{\partial}{\partial\tau}T^{0\nu} - \sinh\eta\frac{\partial}{\partial\tau}T^{3\nu} + \frac{\cosh\eta}{\tau}\frac{\partial}{\partial\eta}T^{3\nu} - \frac{\sinh\eta}{\tau}\frac{\partial}{\partial\eta}T^{0\nu} + \frac{\partial}{\partial\boldsymbol{x}_\perp}T^{\boldsymbol{x}_\perp\nu} = 0. \tag{3.17}$$

It is readily checked that is equivalent to

$$\frac{\partial}{\partial\tau}\left(\tau\cosh\eta T^{0\nu} - \tau\sinh\eta T^{3\nu}\right) + \frac{\partial}{\partial\eta}\left(\cosh\eta T^{3\nu} - \sinh\eta T^{0\nu}\right) + \tau\frac{\partial}{\partial\boldsymbol{x}_\perp}T^{\boldsymbol{x}_\perp\nu} = 0 \tag{3.18}$$

which gives us the conserved quantities [29],

$$\frac{d}{d\tau}\int d^2\boldsymbol{x}_\perp d\eta\left(\tau\cosh\eta T^{0\nu} - \tau\sinh\eta T^{3\nu}\right) = 0 \tag{3.19}$$

In the cylindrically symmetric case, the only relevant cases are where $\nu = 0, 3$, which produces the following quantities,

$$\frac{d}{d\tau}\int d\eta\, d^2\boldsymbol{x}_\perp\, d^2\boldsymbol{k}_\perp\, dY\ \tau\, k_\perp^2 \cosh(Y)\cosh(Y-\eta) f = 0, \tag{3.20}$$

which correspond to energy conservation and,

$$\frac{d}{d\tau}\int d\eta\, d^2\boldsymbol{x}_\perp\, d^2\boldsymbol{k}_\perp\, dY\ \tau\, k_\perp^2 \sinh(Y)\cosh(Y-\eta) f = 0, \tag{3.21}$$

which corresponds to momentum conservation.

In a similar vein we can derive a conserved quantity that corresponds to particle number conservation, in particular,

$$\frac{d}{d\tau}\int d\eta\, d^2\boldsymbol{x}_\perp\ \left(\tau\cosh\eta J^0 - \tau\sinh\eta J^3\right) = 0 \tag{3.22}$$

These conserved quantities provide a valuable metric to test the accuracy of our simulations.

3.2 Discretisation of the Boltzmann equation

In order to generate a numerical solution to the Boltzmann equation, it first needs to be recast into a discrete form (both in time and space).

The first step is to separate the Boltzmann equation into two independent components, the streaming(S) and relaxation(R) steps, each of which can be solved separately.

So given a first order which consists of a linear combination of operators, as follows,

$$\frac{\partial}{\partial t} f(t) = \left(\hat{S} + \hat{R}\right) f(t), \tag{3.23}$$

we are able to write down a formal solution

$$f(t) = e^{\left(\hat{S}+\hat{R}\right)t} f(t=0), \tag{3.24}$$

$$f\left(t+\Delta t\right)=e^{\left(\hat{S}+\hat{R}\right)\Delta t}f(t) \tag{3.25}$$

using the Baker-Campbell-Hausdorff (BCH) formula.

Given

$$e^{\hat{X}}e^{\hat{Y}}=e^{\hat{Z}}, \tag{3.26}$$

the BCH formula tell us how to calculate $\hat{Z}$

$$\hat{Z}=\hat{X}+\hat{Y}+\frac{1}{2}\left[\hat{X},\hat{Y}\right]+\frac{1}{12}\left[\hat{X},\left[\hat{X},\hat{Y}\right]\right]+\frac{1}{12}\left[\hat{Y},\left[\hat{Y},\hat{X}\right]\right]+... \tag{3.27}$$

Using the BCH formula we can evaluate the truncation error of approximations to $e^{\left(\hat{S}+\hat{R}\right)\Delta t}$.

Expanding the exponential via the *Lie Splitting* [30] approximation, given by

$$e^{\left(\hat{S}+\hat{R}\right)\Delta t}\approx e^{\hat{S}\Delta t}e^{\hat{R}\Delta t} \tag{3.28}$$

shows a truncation error $\mathcal{O}(\Delta t^2)$.

Similarly, expanding the exponential via *Strang Splitting* [31] approximation, given by

$$e^{\left(\hat{S}+\hat{R}\right)\Delta t}\approx e^{\hat{S}\frac{\Delta t}{2}}e^{\hat{R}\Delta t}e^{\hat{S}\frac{\Delta t}{2}} \tag{3.29}$$

shows a truncation error $\mathcal{O}(\Delta t^3)$.

Our next step will be to decompose the Boltzmann equation into these linear operators and solve equations generated by these operators.

3.3 The Streaming Operator

The first linear operator in the Strang splittling expansion we will consider is the *streaming* operator.

In the streaming step we solve the equation

$$\left(\frac{\partial}{\partial\tau}+\frac{1}{\tau}v_\eta\frac{\partial}{\partial\eta}+\boldsymbol{v}_\perp\frac{\partial}{\partial\boldsymbol{x}_\perp}\right)f=0 \tag{3.30}$$

Where we identify the streaming operator $\hat{S}$, as

$$\hat{S}=\frac{1}{\tau}v_\eta\frac{\partial}{\partial\eta}+\boldsymbol{v}_\perp\frac{\partial}{\partial\boldsymbol{x}_\perp} \tag{3.31}$$

This equation describes the propagation of particles in straight lines (in Cartesian coordinates) with some constant velocity. We subdivided the problem further into independent propagation along each axis, turning the partial differential equation in 3 spatial dimensions into 3 1-dimensional differential equations.

In the hyperbolic form of the equation, the 2 transverse directions have the same solution while the longitudinal direction needs to be solved independently.

Since the homogeneous Boltzmann equation admits an analytic solution by the method of characteristics,

The method of characteristics is a method for solving first-order hyperbolic partial differential equations by reducing the PDE into a set of ordinary differential equations which can be integrated given some initial condition.

In the transverse plane, we have a simple advection equation,

$$\frac{d}{d\tau}f=-\boldsymbol{v}_\perp\frac{d}{d\boldsymbol{x}_\perp}f \tag{3.32}$$

for which we can make a simple ansatz solves the equation,

$$f(\tau,\boldsymbol{x}_\perp)=f_0(\boldsymbol{x}_\perp-\boldsymbol{v}_\perp(\tau-\tau_0)) \tag{3.33}$$

for some initial condition f_0 specified at some τ_0.

The longitudinal component of the streaming step is given by,

$$\frac{d}{d\tau}f = -\frac{v_\eta}{\tau}\frac{d}{d\eta}f. \tag{3.34}$$

This is once again a first-order linear differential equation, and so, admits a solution by the method of characteristics, if given some initial condition $f_0(\eta)$ at an initial $\tau = \tau_0$.

This case is less straightforward than the transverse components, as the velocity term has both an η and a τ dependence.

The pair of differential equations generated by the method of characteristics are,

$$\begin{aligned} \frac{d\tau}{ds} &= 1, \\ \frac{d\eta}{ds} &= \frac{1}{\tau}\tanh(Y-\eta). \end{aligned} \tag{3.35}$$

The solution of these equations generate lines along which the function is constant. The solutions of the above equations are given by,

$$\eta(\tau) = Y - \sinh^{-1}\left(\frac{\tau_0 \sinh(Y-\eta_0)}{\tau}\right) \tag{3.36}$$

inverting this expression to extract η_0, we get solution to Eq. (3.34) as,

$$f(\tau,\eta,\boldsymbol{k}_\perp,Y) = f_0(Y - \sinh^{-1}\left(\frac{\tau \sinh(Y-\eta)}{\tau_0}\right), \boldsymbol{k}_\perp, Y). \tag{3.37}$$

Visualising how the advection equation evolves a typical initial condition in these coordinates makes it clear as to why this hyperbolic transformation is useful in contrast to a Cartesian description. By comparing how an initial condition Fig. (3.1) to Fig. (3.2) we see, that instead of an increase in the required size of the spatial grid as the system expands that we see in Cartesian, the spatial rapidity domain required to represent the distirbution function does not increase in size.

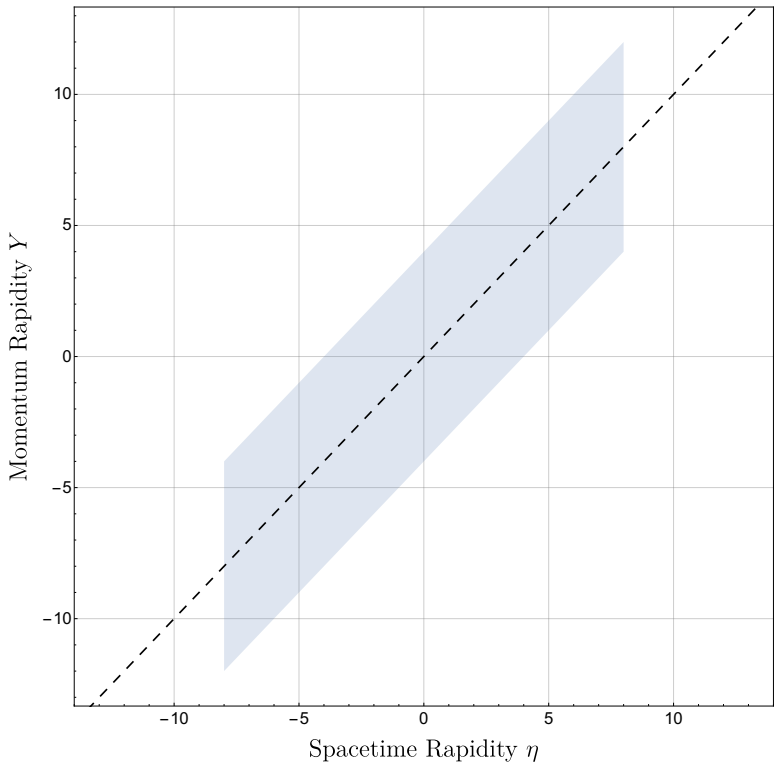

Figure 3.1: Region of support(i.e where the distribution function is non-zero) of a typical boost-invariant initial condition in rapidity space.

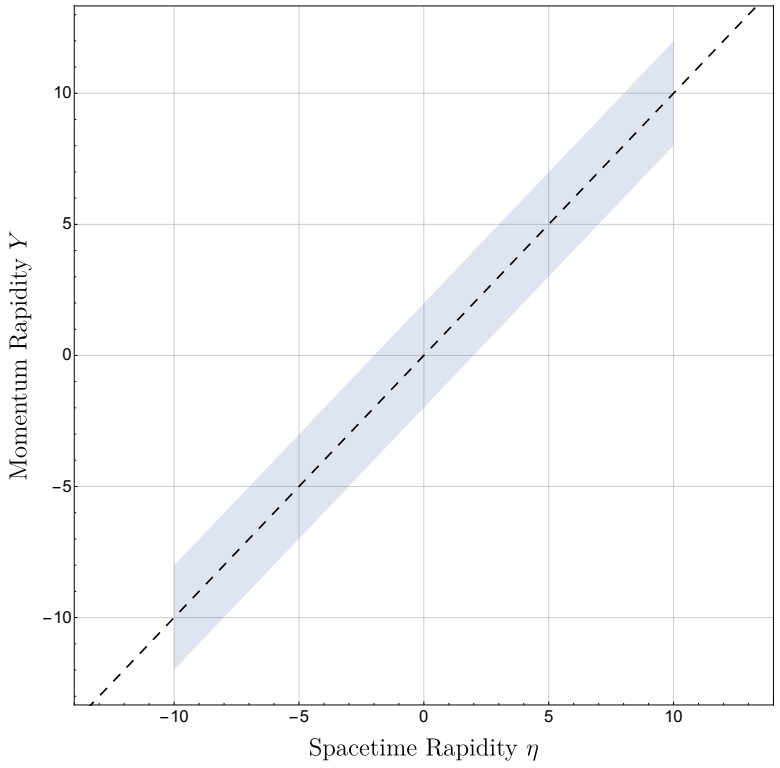

Figure 3.2: Region of support (i.e where the distribution function is non-zero) of a typical boost-invariant initial condition after free-streaming.

3.3.1 Transverse discretisation

If we knew the solution of the equation on a hypersurface of constant eigentime, we could evolve the solution in time because we know the exact solution via the method of characteristics. When we do numerical simulations, we do not, in general, have access to the solution across the entire equal-eigentime hypersurface but rather just points sampled at discrete intervals. There are various methods of sampling and updating these discrete points. As the solutions to the transverse and longitudinal streaming components are different, the discretisation schemes employed we choose for each case will be different.

3.3.1.1 Finite volume methods

The particular method we will employ to discretise the transverse plane belongs to a class of methods called finite volume methods [32], the general idea behind these methods is to, instead of storing and evolving discrete values of the function value, store the average value of the function within some cell and evolve that average.

When numerically solving the advection equation, pointwise methods frequently suffer significant problems with numerical stability, with solutions experiencing exponential growth or decay when the solution is known to be conserved. With finite-volume methods, we will calculate fluxes between adjacent cells so that if the function value in a particular cell increases/decreases, there was precisely the same decrease/increase in the adjacent cells, this implies that the numerical solution generated is guaranteed to be conserved.

Now that we have introduced the concept of a finite-volume method, let us delve into the details of how they work and are implemented.

Starting with the advection equation,

$$\frac{d}{dt}f(x,\eta,\tau) + v\frac{d}{dx}f(x,t) = 0 \tag{3.38}$$

we perform a spatial integration across one of the cells to get

$$\frac{d}{dt}\int_{x_{i-1/2}}^{x_{i+1/2}} dx\, f(x,t) = -v\left(f(x_{i+1/2},t) - f(x_{i-1/2},t)\right), \tag{3.39}$$

where we have discretised our spatial dimension into N cells $I_i = [x_{i-1/2}, x_{i+1/2}]$ with width $\Delta x = |x_{i+1/2} - x_{i-1/2}|$ with $i = 1, .., N$ indexing the cells.

Additionally, performing a temporal integration from t_n to t_{n+1} (where n indexes discrete times with timestep $t_{n+1} - t_n = \Delta t$) we get,

$$\begin{aligned}\int_{x_{i-1/2}}^{x_{i+1/2}} dx\, f(x,t_{n+1}) = &\int_{x_{i-1/2}}^{x_{i+1/2}} dx\, f(x,t_n) \\ &- v\left(\int_{t_n}^{t_{n+1}} dt\, f(x_{i+1/2},t) - \int_{t_n}^{t_{n+1}} dt\, f(x_{i-1/2},t)\right).\end{aligned} \tag{3.40}$$

Eq. (3.40) is an expression that states that the change of particle number within one cell in a given time is equal to the time-integrated flux of particles through the boundaries of the cell.

Using the solution given by the method of characteristics in Sec. (3.3) we can write

$$f(x_j,t) = f(x_j - v(t-t_n), t_n) \tag{3.41}$$

upon substituting this result into Eq. (3.40) we get,

$$\begin{aligned}&\int_{x_{i-1/2}}^{x_{i+1/2}} dx\, f(x,t_{n+1}) = \int_{x_{i-1/2}}^{x_{i+1/2}} dx\, f(x,t_n) \\ &- v\left(\int_{t_n}^{t_{n+1}} dt\, f(x_{i+1/2} - v(t-t_n), t_n) - \int_{t_n}^{t_{n+1}} dt\, f(x_{i-1/2} - v(t-t_n), t_n)\right)\end{aligned} \tag{3.42}$$

Performing the change of variables $x = x_{i+1/2} - v\Delta t$ and $x = x_{i-1/2} - v\Delta t$ on the latter two integrals respectively we get

$$\int_{x_{i-1/2}}^{x_{i+1/2}} dx\, f(x,t_{n+1}) = \int_{x_{i-1/2}}^{x_{i+1/2}} dx\, f(x,t_n) - \int_{x_{i+1/2}-v\Delta t}^{x_{i+1/2}} dx\, f(x,t_n) + \int_{x_{i-1/2-v\Delta t}}^{x_{i-1/2}} dx\, f(x,t_n) \tag{3.43}$$

Defining,

$$F_i^n = \frac{1}{\Delta x}\int_{x_{i-1/2}}^{x_{i+1/2}} dx\, f(x,t_n) \tag{3.44}$$

as the average over the cell, which will be the value that we store in our algorithm.

We further define,

$$\Phi_i^{\pm} = \int_{x_{i\pm1/2-v\Delta t}}^{x_{i\pm1/2}} dx\, f(x,t_n) \tag{3.45}$$

which represent the particle flux through the two faces of the cell.

Combing Eqs. (3.43, 3.44, 3.45) we are left with an update equation,

$$f_i^{n+1} = f_i^n + \Delta x\,(\Phi_i^- - \Phi_i^+). \tag{3.46}$$

This is still formally an exact equation; however, since we only store the values f_i^n, we have to approximate the flux integrals $\Phi_i^{\pm}$. To approximate these integrals, we perform an interpolation of f given the f_i^n values and integrate the interpolation which is covered in the next section.

In this streaming algorithm we choose a set of transverse coordinates that will be convenient to work with (e.g. for numerical integration), these set of coordinates remain fixed and the cell-aveaged function values are updated.

3.3.1.2 Positive and flux conservative interpolation

The interpolation required by Eq. (3.45) should satisfy several constraints, firstly f is strictly non-negative, and hence the interpolation should also be strictly non-negative.

Secondly, in order to maintain particle number, the average needs to be conserved

$$\int_{x_{i-1/2}}^{x_{i+1/2}} f_h(x, t_n) = \Delta x f_i^n. \tag{3.47}$$

In order to ensure numerical stability it should also satisfy the maximum principle, in particular $\forall x \in (x_{min}, x_{max})$

$$f_h(x) \leq f_\infty, \tag{3.48}$$

with $f_\infty = \max_{j \in I} f_j^n$.

We use the interpolation suggested in [33],

$$\begin{aligned} f_h(x,t) =& f_i^n + \frac{\epsilon_i^+}{6\Delta x^2}\left[2(x-x_i)(x-x_{i-3/2}) + (x-x_{i-1/2})(x-x_{i+1/2})\right](f_{i+1}^n - f_i^n) \\ &- \frac{\epsilon_i^-}{6\Delta x^2}\left[2(x-x_i)(x-x_{i+3/2}) + (x-x_{i-1/2})(x-x_{i+1/2})\right](f_i^n - f_{i-1}^n) \end{aligned} \tag{3.49}$$

with the slope limiters ϵ_+, ϵ_- defined as

$$\epsilon_i^\pm = \begin{cases} \min\left(1; f_i^n/(f_{i\pm1}^n - f_i^n)\right) & \text{if } f_{i\pm1} - f_i^n > 0 \\ \min\left(1; -2(f_\infty - f_i^n)/(f_{i\pm1}^n - f_i^n)\right)\right) & \text{if } f_{i\pm1}^n - f_i^n < 0 \end{cases} \tag{3.50}$$

which are non-linear functions of f_i^n introduced in order satisfy the maximum principle. This interpolation scheme satisfies the above constraints.

This allows us to calculate the fluxes in Eq. (3.45), for example,

$$\Phi_i^+ = \frac{1}{\Delta x}\int_{x_{i+1/2}-\alpha\Delta x}^{x_{i+1/2}} dx\, f_h^n = \alpha\Delta x\Big(f_i^n + \frac{\epsilon_i^+}{6}(2-3\alpha+\alpha^2)(f_{i+1}^n - f_i^n) - \frac{\epsilon_i^-}{6}(\alpha^2-1)(f_i - f_{i-m})\Big) \tag{3.51}$$

where $\alpha = \frac{v\Delta t}{\Delta x}$.

3.3.2 Longtitudinal streaming

The longitudinal streaming solution Eq. (3.37) is not amenable to calculating the flux between cells.To get around this limitation we recognise that in the streaming case any function value at a given position and (eigen)time can be obtained directly from the initial condition as follows,

$$f(\tau,\eta,\boldsymbol{k}_\perp,Y) = f_0(Y - \sinh^{-1}\left(\frac{\tau\sinh(Y-\eta)}{\tau_0}\right),\boldsymbol{k}_\perp,Y). \tag{3.52}$$

If we require access to the function values at some fixed set of points (e.g. for integration) it is highly unlikely that these points will map back to points we sampled in our initial condition. Thus once we've mapped those point back we use an inverse distance weighted interpolation to estimate the initial condition at that point,

$$f_{initial}^n(x) = \sum_{<i>} w_i^n(x) f_{initial}^n(x_i), \tag{3.53}$$

where $\sum_{<i>}$ indicates a sum over the indices of the nearest neighbours to the point x, and

$$w_i^n = \frac{\sum_{<j>}|x - x_j^n|}{|x - x_i^n|}, \tag{3.54}$$

is the weight applied to the neighbouring sample point.

In our algorithm, we use 9 points to calculated this weighted average. First we find the nearest grid point to our candidate point and then the 4 nearest neighbours and the 4 near-to-nearest neighbours to that grid point.

3.4 The Relaxation Operator

The second linear operator in the Strang splitting expansion is the "relaxation" step. The relaxation operator in this case is a non-linear functional defined as

$$\hat{R}f = \frac{k_\mu u^\mu}{k_\perp \cosh(Y-\eta)} \frac{f^{eq} - f}{\tau_{\text{rel}}}. \tag{3.55}$$

The equation we solve in the relaxation step is given by

$$\frac{d}{dt} f = \frac{k_\mu u^\mu}{k_\perp \cosh(Y-\eta)} \frac{f^{eq} - f}{\tau_{\text{rel}}}. \tag{3.56}$$

If we know f^{eq} (and hence u^μ as well) we can use a numerical differential update algorithm (such as an Euler update) to advance the distribution,

$$f(\tau + \Delta\tau) = f(\tau) + \frac{k_\mu u^\mu}{k_\perp \cosh(Y-\eta)} \frac{\Delta\tau}{\tau_{rel}} (f^{eq} - f(\tau)). \tag{3.57}$$

In this step it is necessary to determine f^{eq}. In order to do so, we need to extract the local energy and number density and solve Eqs. (1.19,1.20). This requires us to calculate the energy-momentum tensor and particle-current four-vector, which requires us to calculate various integrals of the distribution function. We can efficiently numerically evaluate these integrals using a Gauss-Legendre integration procedure outlined in Appendix(C).

Once we have the energy-momentum tensor we need to determine its eigenvalues and eigenvectors, to do this we use the GNU Scientific Library(GSL) [34]. The root finding to extract the temperature and chemical potential is also handled by GSL.

3.4.1 The numerics of eigenvalue extraction

One feature of the system that became apparent when attempting to implement this system is that, at large spacetime-rapidities, the extraction of the energy eigenvalue becomes numerically inaccurate. This inaccuracy is present even if one increases the order of integration. To see why this is the case, we first investigate a cylindrically symmetric distribution function.

In the case of a tranversely homogenous and isotropic system the energy-momentum tensor simplifies dramatically (the number of degrees of freedom are reduced down to three).

$$T^{\mu}_{\ \nu} = \begin{pmatrix} T^{00} & 0 & 0 & -T^{03} \\ 0 & -T^{\perp\perp} & 0 & 0 \\ 0 & 0 & -T^{\perp\perp} & 0 \\ T^{03} & 0 & 0 & 2T^{\perp\perp} + T^{00} \end{pmatrix} \tag{3.58}$$

The simple structure of the energy-momentum tensor under this symmetry allows us to analytically extract the eigenvalues, and in particular the positive eigenvalue which corresponds to the local energy density

$$\epsilon = T^{\perp\perp} + \sqrt{(T^{00} - T^{\perp\perp})^2 - (T^{03})^2}. \tag{3.59}$$

At large rapidities the components T^{00} and T^{03} become very large relative to the perpendicular component $T^{\perp\perp}$, and it is this property that leads to the numerical inaccuracies. So our goal will be to recast such that we avoid the subtraction of numbers that vary greatly in scale.

To do this we begin with

$$T^{00} = \frac{1}{(2\pi)^3} \int d^2\boldsymbol{k}_\perp dY |k_\perp|^2 \cosh(Y)^2 f, \tag{3.60}$$

$$T^{03} = \frac{1}{(2\pi)^3} \int d^2\boldsymbol{k}_\perp dY |k_\perp|^2 \cosh(Y) \sinh(Y) f, \tag{3.61}$$

$$T^{\perp\perp} = \frac{1}{(2\pi)^3} \int d^2\boldsymbol{k}_\perp dY \frac{1}{2} |k_\perp|^2 f. \tag{3.62}$$

We decompose the hypertrigonometric functions into their exponential representation,

$$\cosh(Y)^2 = \frac{1}{4}\left(e^{2Y} + e^{-2Y}\right) + \frac{1}{2} \tag{3.63}$$

$$\cosh(Y)\sinh(Y) = \frac{1}{4}\left(e^{2Y} - e^{-2Y}\right). \tag{3.64}$$

With this decomposition in hand, we can rewrite the components of the energy-momentum tensor as linear combinations of the following quantities,

$$A = \frac{1}{(2\pi)^3}\int d^2\boldsymbol{k}_\perp dY |k_\perp|^2 e^{2Y} f, \tag{3.65}$$

$$B = \frac{1}{(2\pi)^3}\int d^2\boldsymbol{k}_\perp dY |k_\perp|^2 e^{-2Y} f, \tag{3.66}$$

$$C = \frac{1}{(2\pi)^3}\int d^2\boldsymbol{k}_\perp dY |k_\perp|^2 f. \tag{3.67}$$

The behaviour of these integrals at large spacetime-rapidities is apparent. For large (positive) spacetime-rapidities the quantity A is exponentially enhanced while the quantity B is exponentially suppressed (and vice-versa for large negative rapidities). The quantity C has no explicit spacetime-rapidity dependence.

The components of the energy-momentum tensor rewritten in terms of these quantities look as follows,

$$T^{00} = \frac{1}{4}(A + B) + \frac{1}{2}C, \tag{3.68}$$

$$T^{03} = \frac{1}{4}(A - B), \tag{3.69}$$

$$T^{\perp\perp} = \frac{1}{2}C. \tag{3.70}$$

The benefit of this exercise becomes clear once we rewrite the energy eigenvalue in terms of A, B and C,

$$\epsilon = \frac{1}{2}\left(\sqrt{AB} + C\right). \tag{3.71}$$

As the two relatively large and small values no longer but add or subtract from one another, but instead multiply, the result is numerically stable.

The flow velocity in the longitudinal direction (the only non-trivial component of the flow velocity in the transversely isotropic case) is given,

$$u^z = \frac{A - B}{A + B + 2\sqrt{AB}}. \tag{3.72}$$

This solves the problem of numerical stability in the cylindrically symmetric case, but not yet in the general case.

To expand this procedure to the general case we can consider boosting the energy-momentum tensor into a new frame with a relative velocity of $-u^z$, since the eigenvalues of the energy-momentum tensor are invariant under boosts.

The resulting energy-momentum tensor has no net flow in the longitudinal direction; hence the problematic large values generated by a large flow velocity are no longer present.

The resulting matrix, unfortunately, does not permit a simple expression for the energy eigenvalue as in the transversely isotropic case; empirically, however, numerical eigenvalue solvers can reliably extract the corresponding eigenvectors and eigenvalues once this boost procedure has been performed.

Chapter 4

The transversely homogenous and isotropic case

By assuming that nuclei are large (effectively infinite) and homegenous in the transverse plane, we eliminate any possibility of transverse plane dynamics, reducing the problem to 2 spatial dimensions.This assumption is a handy test case, as we can easily compare to known analytical and numerical results to establish the validity of our results.

4.1 Existing solutions

4.1.1 Bjorken longitudinal hydrodynamic model

The Bjorken hydrodynamic model [35] assumes that the distribution function is invariant under Lorentz boosts along the beam axis. While a fully boost invariant distribution function would have infinite energy, it is plausible, given the experimentally observed distribution of charged particles $dN_{ch}/d\eta$ around mid-rapidity, that the assumption of boost-invariance is at least approximately valid in that region (see Fig. (4.1)).

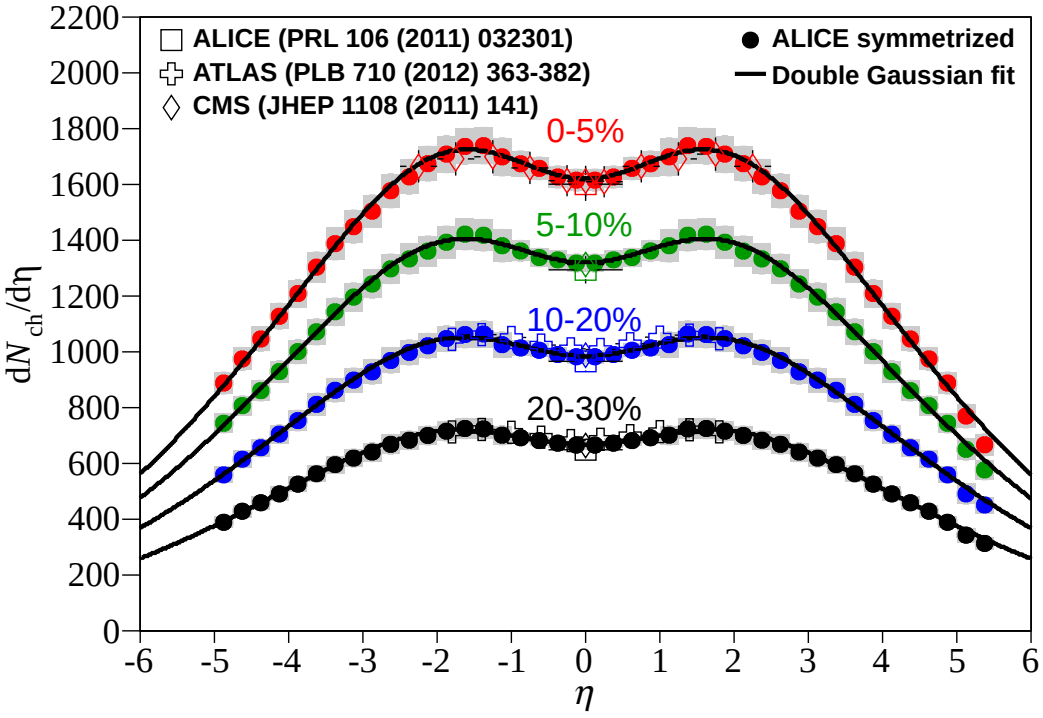

Figure 4.1: The number of charged particles per pseudorapidity bin $dN_{ch}/d\eta$ for various centralities measured by experiments at the LHC taken at energy $\sqrt{s_{\rm NN}} = 2.76$ TeV. The relatively flat peaks in the mid-rapidity region are indicative of approximate boost invariance [36].

For a cylindrically symmetric distribution function the Boltzmann equation becomes

$$\partial_\tau f + \frac{v_\eta}{\tau}\partial_\eta f = \frac{k^\mu u_\mu}{k_0 \cosh\eta - k_z \sinh\eta}\frac{f^{eq} - f}{\tau_{\rm rel}}. \tag{4.1}$$

Under the assumption of full boost-invariance, solving for the distribution function at mid-rapidity is equivalent to solving for the distribution function at all rapidities, and at mid-rapidity we have that $\partial_\eta f = -\partial_Y f$, $v_\eta = k_z/k_0$ and $k_\mu u^\mu = k_0$ giving us,

$$\partial_\tau f - \frac{1}{\tau}\frac{k_z}{k_0}\partial_Y f = \frac{f^{eq} - f}{\tau_{rel}}. \tag{4.2}$$

Integrating the above to find the energy density we get

$$\partial_\tau \epsilon + \frac{1}{\tau}(\epsilon + T^{33}) = 0, \tag{4.3}$$

the RHS side vanishing by the condition set by Eq. (2.12).

Under the hydrodynamic assumption of $\tau_{rel} = 0$ the distribution function is always in local equilibrium, since the distribution function is isotropic and the energy moementum tensor is traceless we have that $T^{33} = \frac{1}{3}\epsilon$, leaving us with,

$$\partial_\tau \epsilon + \frac{4}{3}\frac{\epsilon}{\tau} = 0 \tag{4.4}$$

which is readily integrated to give us,

$$\epsilon(\tau) = \epsilon_0 \left(\frac{\tau}{\tau_0}\right)^{-4/3}. \tag{4.5}$$

If the system does not conserve particle number, we can explicitly write down the eigentime dependence of the temperature using Eq. (1.21), giving us

$$\frac{T(\tau)}{T_0} = \left(\frac{\tau}{\tau_0}\right)^{-1/3}, \tag{4.6}$$

where T_0 is the initial temperature of the system.

4.1.2 Non-boost invariant longitudinal ideal hydrodynamics

Another model we will compare against is a non-boost invariant longitudinal ideal hydrodynamic model.

As the name suggests, it assumes ideal hydrodynamics (i.e. $\tau_{\text{rel}} = 0$), but without making the assumption made by Bjorken and Baym models of boost-invariance.

The set of coupled ordinary differential equations generated by the model are given by [37],

$$\partial_\tau e = \frac{1}{\tau}\tanh(\eta - Y)\partial_\eta e + (e+p)\left[\tanh(\eta - Y)\partial_\tau Y - \frac{1}{\tau}\partial_\eta Y\right], \tag{4.7}$$

$$\partial_\tau Y = \frac{1}{\tau}\tanh(\eta - Y)\partial_\eta Y + (e+p)^{-1}\left[\tanh(\eta - Y)\partial_\tau p - \frac{1}{\tau}\partial_\eta p\right], \tag{4.8}$$

which can be numerically integrated.

If particle number is conserved then the evolution of the particle number density is given by,

$$\partial_\tau n = \frac{1}{\tau}\tanh(\eta - Y)\partial_\eta n + n\left[\tanh(\eta - Y)\partial_\tau Y - \frac{1}{\tau}\partial_\eta Y\right]. \tag{4.9}$$

We can extract the time evolution of the equilibrium paramters via Eqs. (1.19,1.20).

4.1.3 Baym model of the boost-invariant RTA

We now consider a semi-analytic solution to the Boltzmann equation under the relaxation time approximation first introduced by Baym [38].

It is similar to the Bjorken hydrodynamical solution in that it assumes full boost-invariance, but it also assumes that the initial momentum distribution at mid-rapidity is entirely transverse. Furthermore, it only provides the evolution of the energy-density rather than the full distribution function. The significant advantage provided by this model is that one has a free choice of relaxation times.

An outline of the solution is provided in Appendix(C).

We can consider the free-streaming limit of the assumed initial condition. Under this initial condition $T^{33} = 0$ so Eq. (4.3) becomes,

$$\partial_\tau \epsilon + \frac{\epsilon}{\tau} = 0 \tag{4.10}$$

which is readily integrated to give us,

$$\epsilon(\tau) = \epsilon_0 \left(\frac{\tau}{\tau_0}\right)^{-1}. \tag{4.11}$$

If the system does not conserve particle number, we can explicitly write down the eigentime dependence of the temperature using Eq. (1.21), giving us

$$\frac{T(\tau)}{T_0} = \left(\frac{\tau}{\tau_0}\right)^{-1/4}, \tag{4.12}$$

where T_0 is the initial temperature of the system.

4.2 Initial conditions

While several models [4–6] provide initial conditions for the energy-momentum tensor which is required for relativistic hydrodynamics simulations heavy-ion collisions, the full initial distribution function for heavy-ion collisions is not known.

Our choice of initial distribution function should satisfy several criteria,

- We should be able to control the total amount of energy present in the system to match the energy provided by the experiment.
- The distribution function should be approximately boost invariant at mid-rapidities, but break boost-invariance at large rapidities, as particle colliders have a finite rapidity range.
- We should be able to control whether or not the system is under or overpopulated.
- It should be able to replicate the assumptions made by the Baym model, but not be restricted to those assumptions.

Our proposal of distribution function that satisfies the above criteria is the following,

$$f(Y,\eta,|\boldsymbol{k}_\perp|) = \epsilon_0 \frac{\sqrt{2\pi^3}\,(1-\tanh(\sigma_Y^2))}{\sigma_Y \sigma_\perp^4} e^{-(Y-\eta)^2/\left(2\sigma_Y^2\right)} e^{-|\boldsymbol{k}_\perp|^2/\left(2\sigma_\perp^2\right)} (e^{s(|\eta|-\eta_{\text{width}})}+1)^{-1}. \tag{4.13}$$

.

The $e^{-(Y-\eta)^2/\left(2\sigma_Y^2\right)}$ term controls the longitudinal momentum distribution. The $(Y-\eta)$ dependence ensures that this term is explicitly boost-invariant. σ_Y is a parameter that controls the whether the momentum distribution at mid-rapidity is longitudinally or transversely biased

The $e^{-|\boldsymbol{k}_\perp|^2/\left(2\sigma_\perp^2\right)}$ term controls the transverse momentum distribution, with $\sigma_\perp$ setting the transverse momentum scale, and since the transverse momentum scale also sets the longitudinal momentum scale, it is the only momentum scale that we set. Since we will we will measure all momenta according this scale so we set $\sigma_T = 1$.

We will consider three different choices for σ_Y, such that we get initial conditions that are anisotropic biased transversely $(\sigma_Y = 0.1)$, approximately isotropic $(\sigma_Y = 0.58)$ and anisotropic biased longitudinally $(\sigma_Y = 1)$.

The assumption made of made in the Baym model [38] is satisfied in the limit $\sigma_Y \to 0$, thus for small values of σ_Y we expect to reproduce the behaviour predicted by the Baym model.

Hydrodynamic models assume that the system is close to local thermal equilibrium which would imply an isotropic local momentum distribution. Thus we would expect $\sigma_Y = 0.58$ combined with a small relaxation time would most closely reproduce the hydrodynamic results.

The $(\exp(|\eta| - \eta_{\text{width}}) + 1)^{-1}$ term breaks boost invariance by damping out contributions at rapidities greater than η_{width}, whilst maintaining approximate boost invariance at mid-rapidities. The steepness as which rapidities aboce η_{width} are damped out is controlled by the parameter s.

The distribution function is normalised so that the energy density at mid-rapidity is $d_g\,\epsilon_0$ for all choices of σ (formally, only in the limit $s \to \infty$).

The number density (at mid-rapidity) is,

$$n = d_g \sqrt{\frac{\pi}{2}} \epsilon_0 e^{\sigma_Y^2/2} \left(\sigma_\perp (1 + e^{\sigma_Y^2/2}) \right)^{-1}. \tag{4.14}$$

This allows us to determine the overpopulation parameter (ξ) as discussed in Sec. (1.4)

$$\xi = d_g \frac{\pi^2}{4} \epsilon_0 \, e^{2\sigma_Y^2} \left(\sigma_\perp (1 + e^{\sigma_Y^2/2}) \right)^{-4} \tag{4.15}$$

The specific combination of ϵ_0, σ_Y and $\sigma_\perp$ allow us to control whether or not the initial distribution is under or over-populated.

We set our rapidity width to align with the maximum rapidity achieved by the LHC of ~ 8.

Visualising the initial temperature profile when we do not enforce particle number conservation in Fig. (4.2) we have a simple relation between temperature and energy density $T \propto \epsilon_{\text{LRF}}^{1/4}$, with T_0 being the local temperature at mid-rapidity.

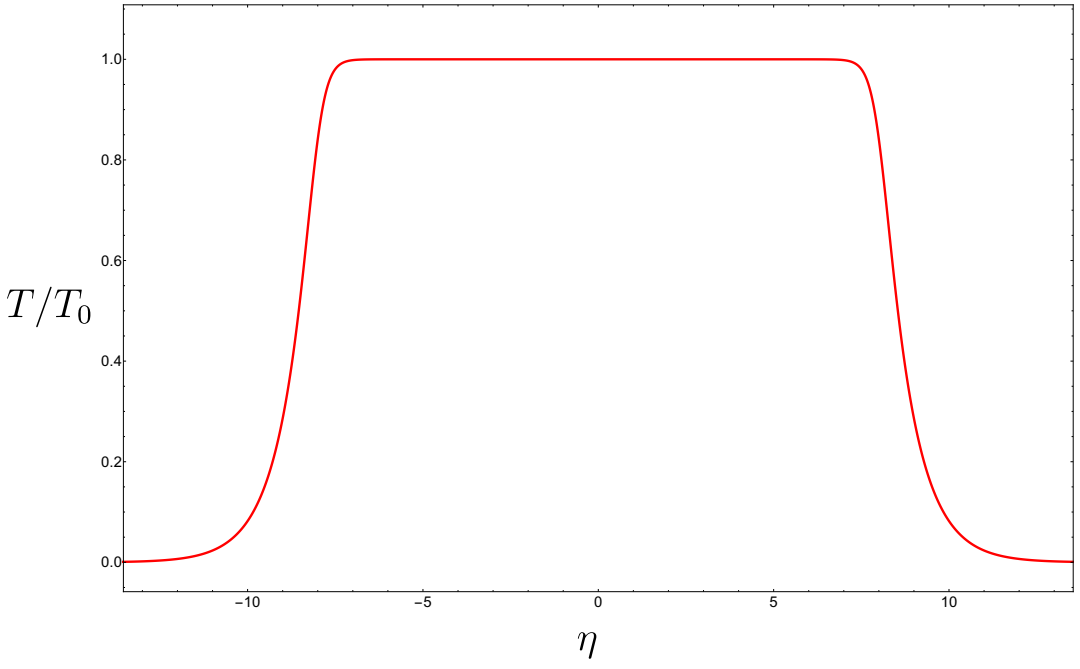

Figure 4.2: Initial temperature profile of a system without particle number conservation

In contrast, when particle number is conserved, we have both an initial temperature profile Fig. (4.3) and chemical potential profile Fig. (4.4) determined by Eqs. (1.19,1.20)

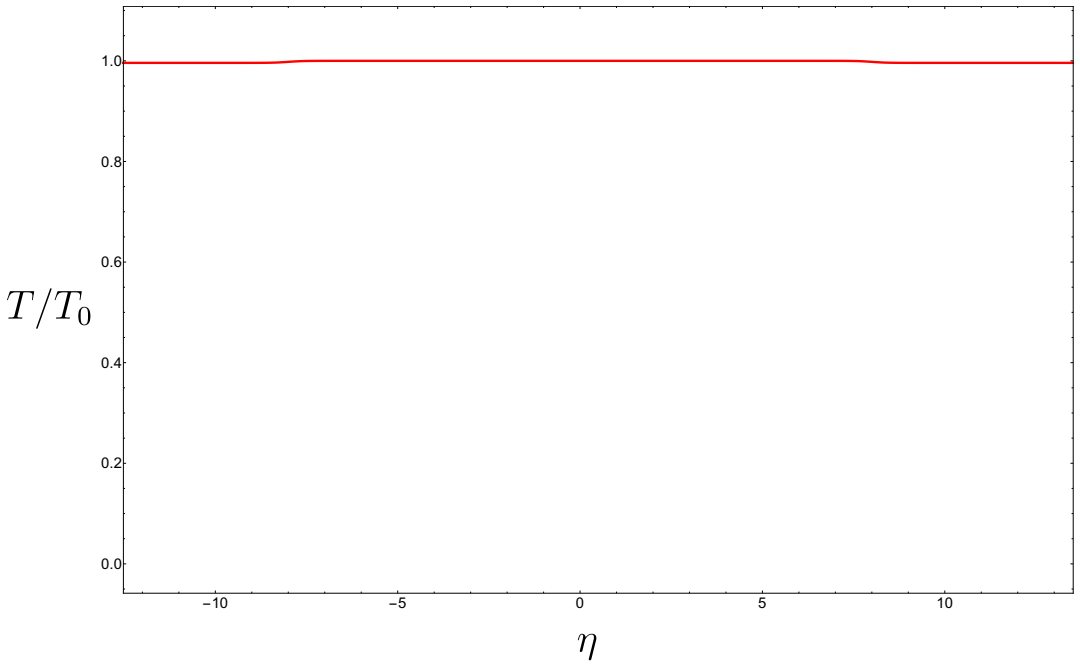

Figure 4.3: Initial temperature profile of a system with particle number conservation

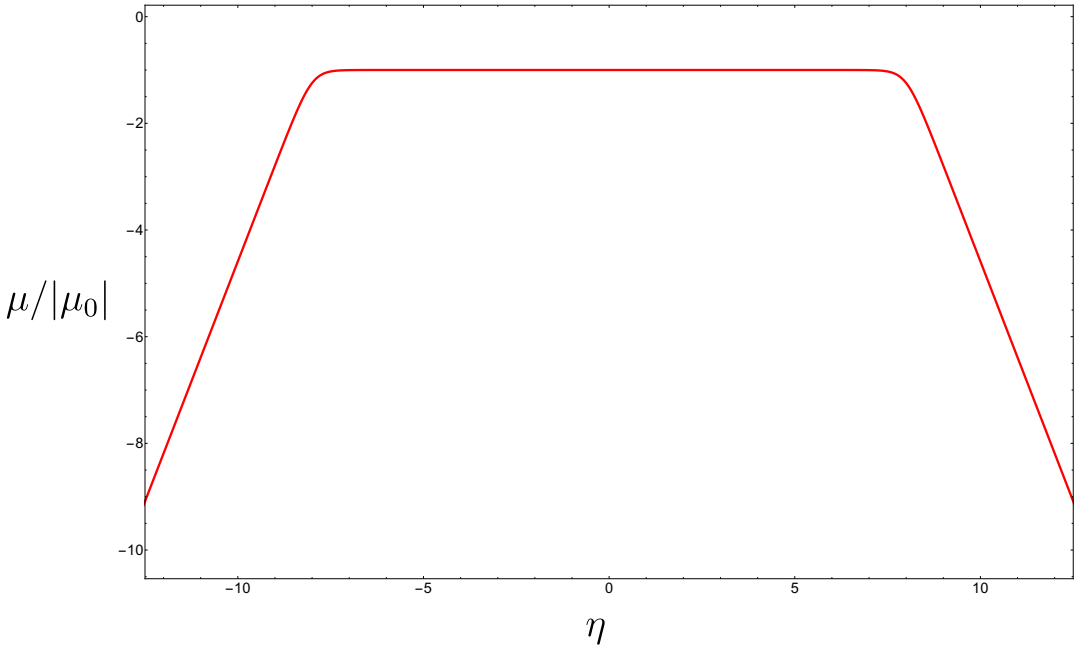

Figure 4.4: Initial chemical potential profile of a system with particle number conservation

We make an interesting observation in models where particle current is conserved. At large rapidities, the initial condition has a vanishingly small energy density, which is accounted for by having a large negative chemical potential (i.e. the system being very dilute). At these extreme rapidities, the system is too dilute for a quark-gluon plasma to realistically form, yet a naive freezeout condition(i.e when to the quark-gluon plasma rehadronizes), $(T_{eq} < T_{\text{freeze}})$ would predict a quark-gluon plasma does form. Thus we need an alternative condition. We will henceforth adopt a critical energy density as our freezeout criterion in the case where particle number is conserved.

We choose this energy density to be the same as the energy density of the Boltzmann distribution at T_{freeze} and $\mu = 0$.

$$\epsilon_{\text{freeze}} = d_g \frac{\pi^2}{30} (T_{\text{freeze}})^4 \tag{4.16}$$

For a purely gluonic system, the freezeout temperature is on the order of $T_{\text{freeze}} \approx 260$ MeV [39], compared to a system that includes quarks which $T_{\text{freeze}} \approx 150$ MeV [8].

4.3 Verification of numerics by conservation laws

We want to ensure that that our scheme is at least self-consistent for the various initial conditions and relaxation times we will consider. To this end we will make sure that our scheme conserves energy according to the conservation law Eq. (3.20)

To this end we measure the evolution of the relative error with respect to the energy,

$$R_E = \left| \frac{E(\tau)}{E(1)} \right|, \tag{4.17}$$

which measures the fraction of deviation of the energy from the initial condition.

We will investigate this relative error for various relaxation times, as well as both assumptions about the conservation of particle number.Since the momentum grid is asymmetric in the longitudinal and transverse directions,

we want to evaluate whether or not this causes a significant variation in accuracy depending on the momentum distribution of the initial condition as well.

The grid size and spacing in both momentum an position space was optimized by trial and error to strike a balance between simulation accuracy and size (smaller grid spacing increase accuracy but decrease the maximum size of the simulation).

The following grid parameters were used in the following experiments

The transverse momentum grid was constructed out of 676 points generated by the Gauss-Legendre method. Similarly, the longitudinal momentum grid was had 44 points generated by the Gauss-Legendre method.

In position space we had 13 points sampled along the transverse radius, and in the longitudinal direction in which the flow is more rapid, 260 points.

The boundary conditions set were sharp cutoffs (i.e all values clamped to zero past the edge of the grid) thus the grid needs to be large enough so that by the end of the evolution no significant flow has reached the edge of the grid.

The intial eigientime at which we will start these experiments will be a constant $\tau_0 = 1fm$ and as such provides a temporal scale.

Hydrodynamic codes typically run till $\tau \approx 12fm$ so we will halt our code there as a benchmark.

4.3.1 $\mu = 0$

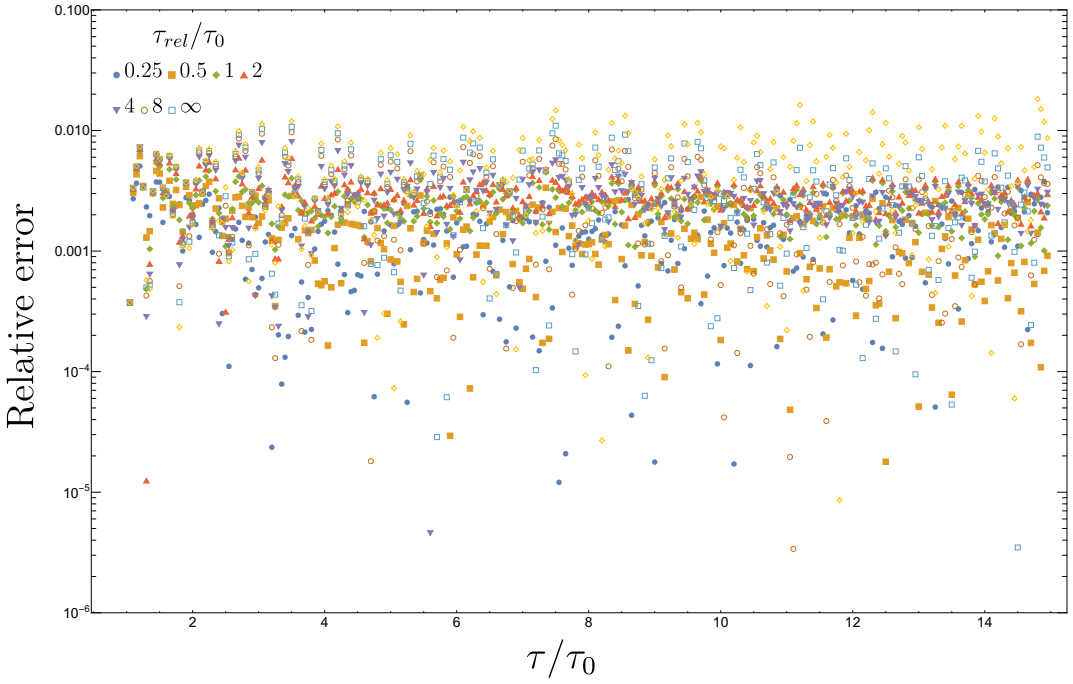

Figure 4.5: Relative error in the hyperbolic energy for a variety of relaxation times, particle number not conserved, with anisotropic initial condition ($\sigma_Y = 0.1$)

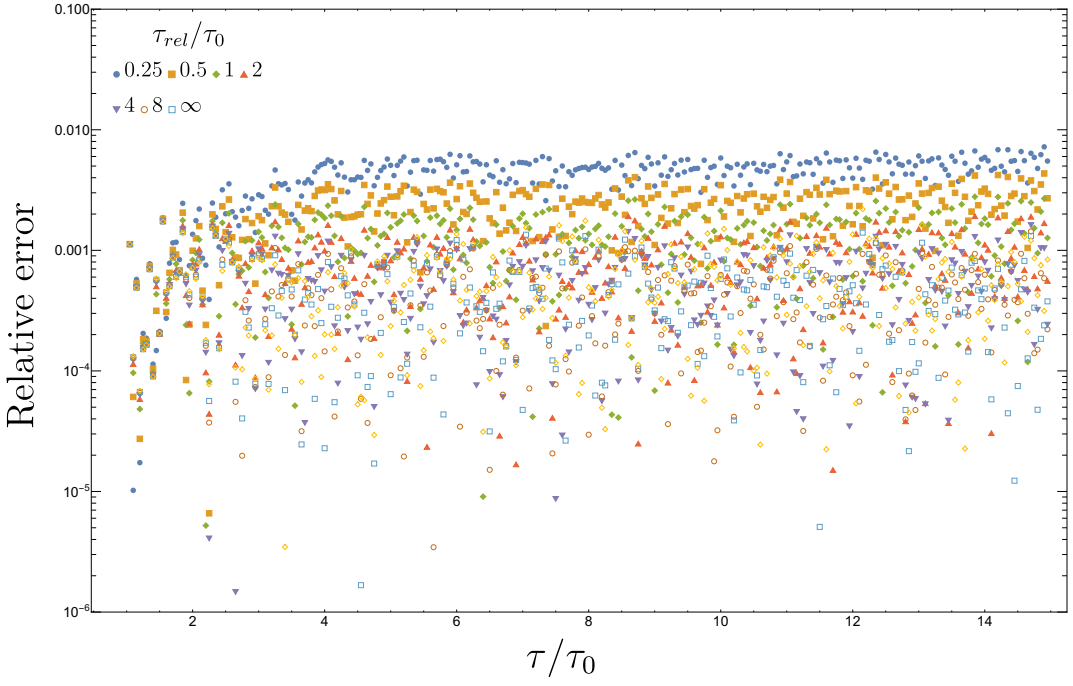

Figure 4.6: Relative error in the hyperbolic energy for a variety of relaxation times, particle number not conserved, with an isotropic initial condition ($\sigma_Y \approx 0.5$)

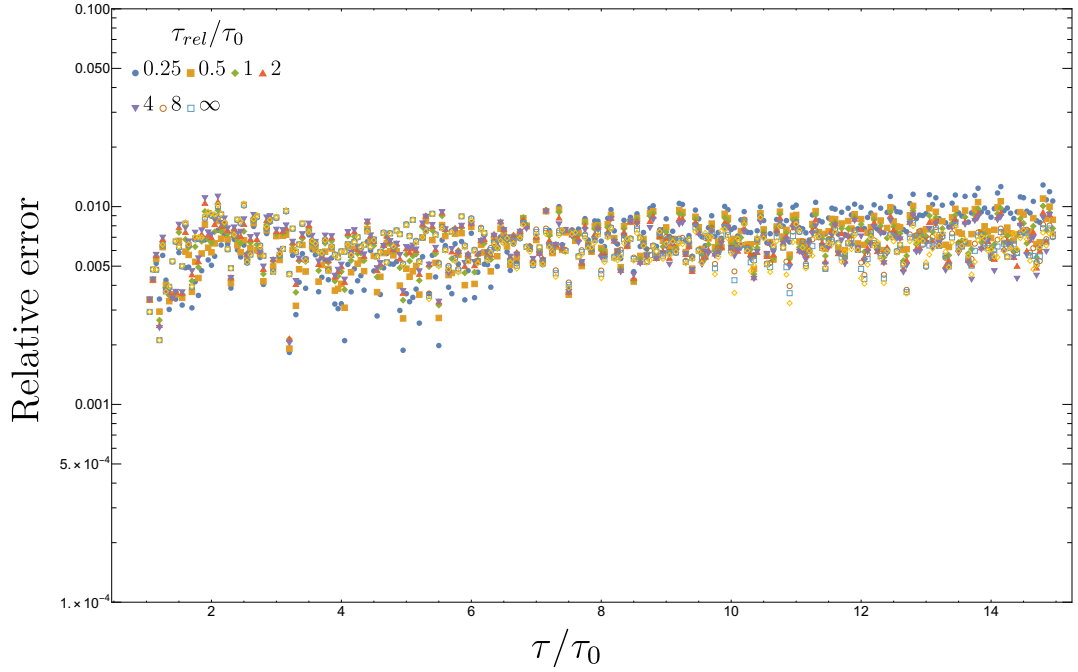

Figure 4.7: Relative error in the hyperbolic energy for a variety of relaxation times, particle number not conserved, with anisotropic initial condition ($\sigma_Y = 1$)

In Figs. (4.5, 4.6, 4.7) we see that the variation in the conserved quantity is, at worst, on the order of $\sim 1\%$. As we see larger variations in the error varying the the momentum distribution of the initial condition, the most significant source of error when there is no particle number conservation, likely stems from the streaming steps (in particular longitudinal streaming) rather than the relaxation steps.

4.3.2 $\mu \neq 0$

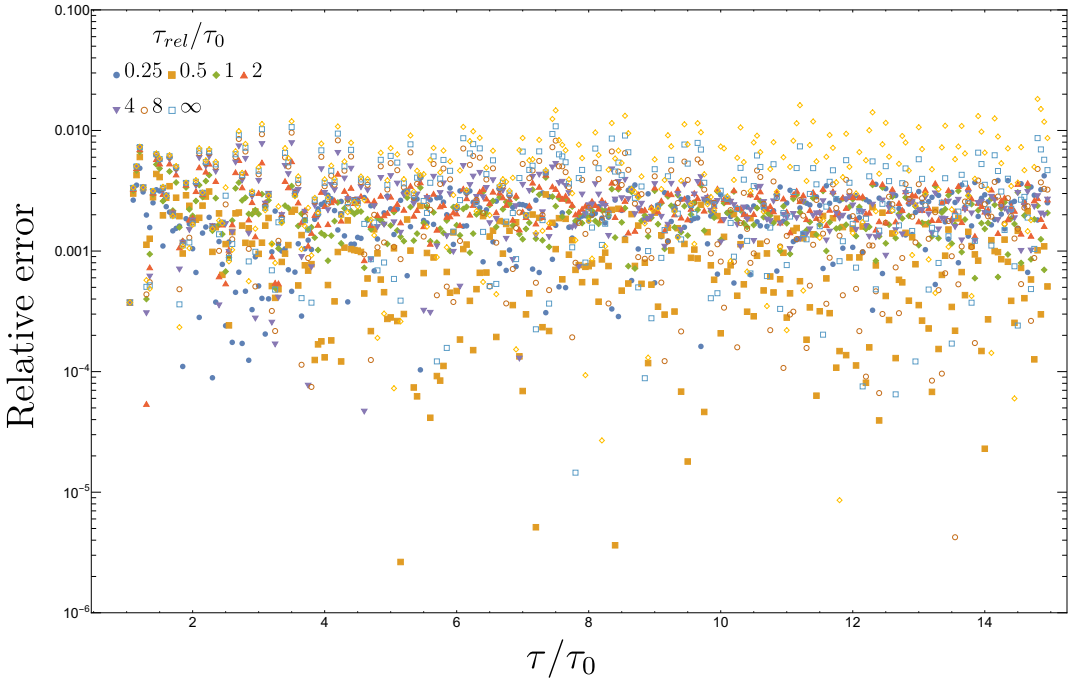

Figure 4.8: Relative error in the hyperbolic energy for a variety of relaxation times, particle number conserved, with a transversely-biased initial condition ($\sigma_Y = 0.1$)

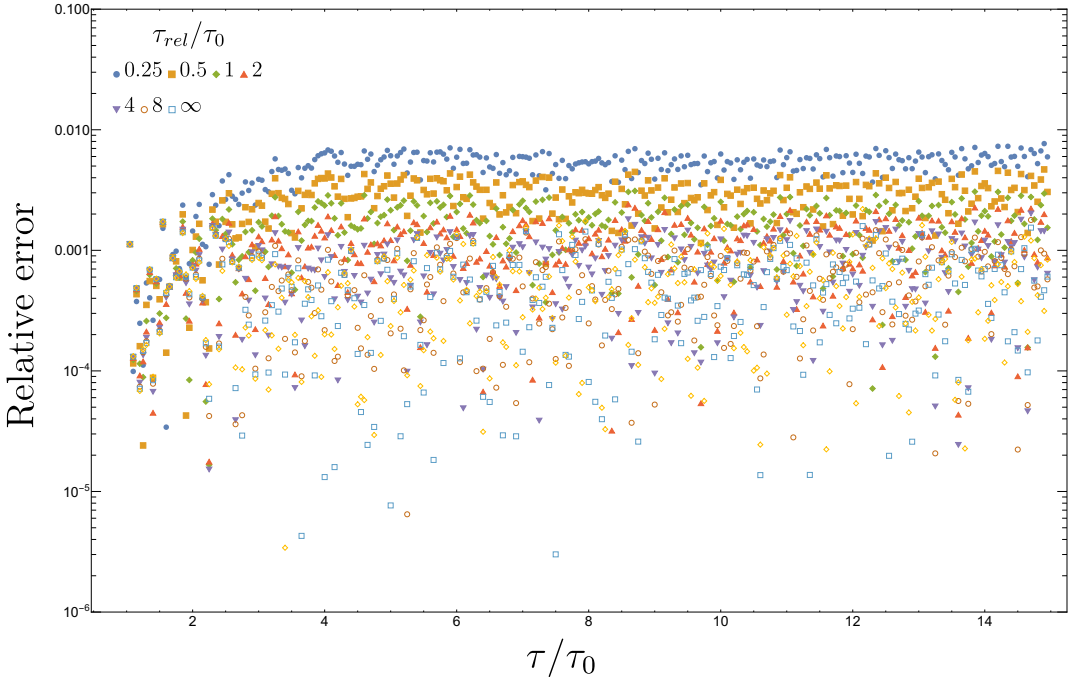

Figure 4.9: Relative error in the hyperbolic energy for a variety of relaxation times, particle number conserved, with an isotropic initial condition ($\sigma_Y \approx 0.5$)

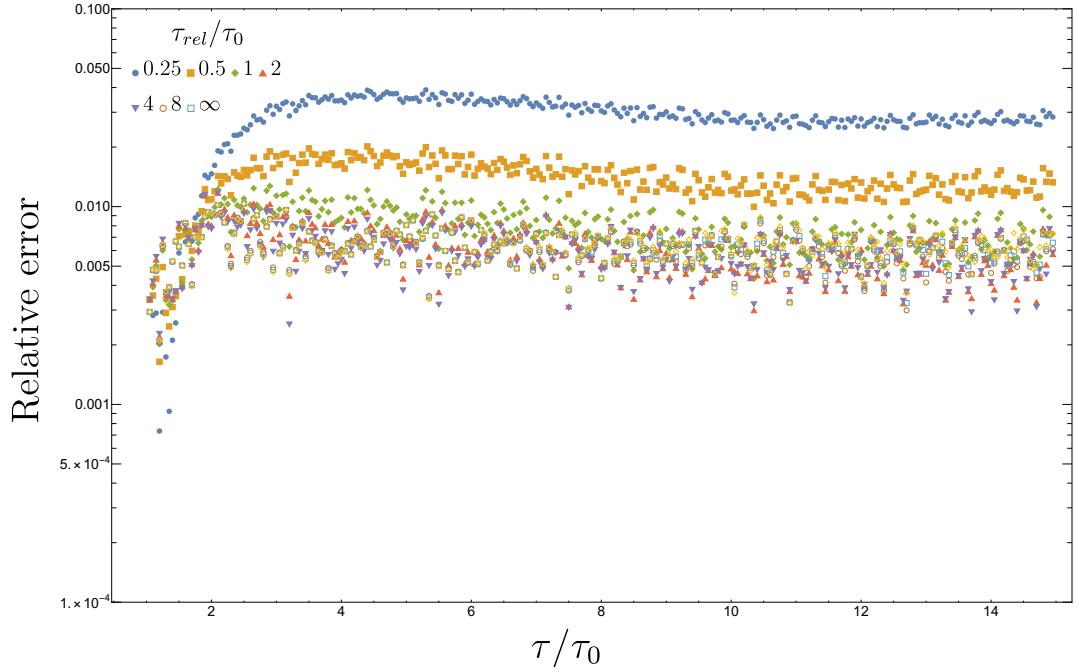

Figure 4.10: Relative error in the hyperbolic energy for a variety of relaxation times, particle number conserved, with a longitudinally-biased initial condition ($\sigma_Y = 1$)

From Figs. (4.8, 4.9 4.10) it seems that turning on particle number conservation increases the sensitivity to relaxation time when the system becomes more isotropic, possibly requiring smaller timesteps for smaller relaxation times. Nevertheless, even in the worst case, the error seems to be bounded to be on the order of $\sim 3\%$.

It is clear that the the relative errors are significantly larger when there is particle number conservation, this is due to the fact that one needs to numerically root find the equilibrium parameters; whilst in the case where we do not require particle number conservation, we have analytic expressions for the equilibrium parameters.

4.4 Comparison to existing solutions

4.4.1 Effect of relaxation time on evolution at mid rapidity

Our first comparison will be the energy density at mid-rapidity, we will compare our simulations against those predicted by the Baym model (Sec. (4.1.3)) and these Bjorken hydrodynamic model (Sec. (4.1.3)).

4.4.1.1 $\mu = 0$

We first consider the comparison when in our simulations, we are not enforcing particle number conservation.

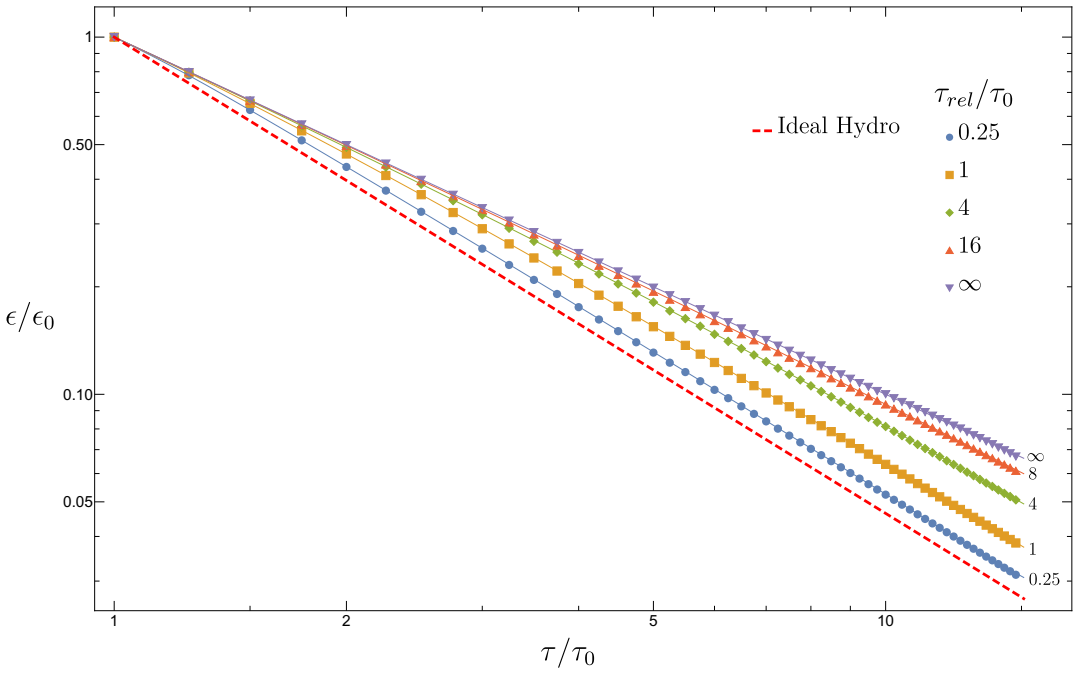

Figure 4.11: Mid-rapidity energy density evolution for a transversely-biased initial condition ($\sigma_Y = 0.1$) for a variety of relaxation times, particle number not conserved. The discrete data points represent results from our simulations while solid lines represent the Baym model solutions.

In Fig. (4.11) we see, as expected, when $\sigma_Y << 1$ our solution coincides with the Baym solutions across a variety of relaxation times including the free-streaming limit ($\tau_{rel} \to \infty$). We can also observe that starting from the transversely biased anisotropic momentum distribution the effect of a smaller relaxation time will tend to isotropise the momentum distribution, so a transverse bias decreases the rate at which the energy density decreases.

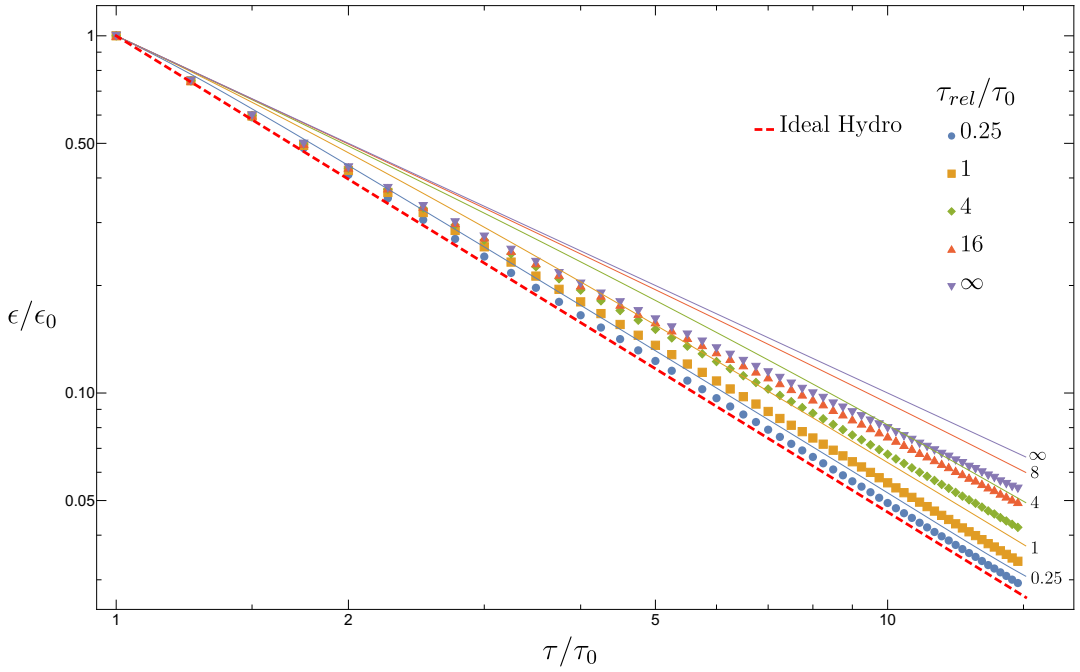

Figure 4.12: Mid-rapidity energy density evolution for an isotropic initial condition ($\sigma_Y \approx 0.58$) for a variety of relaxation times, particle number not conserved. The discrete data points represent results from our simulations while solid lines represent the Baym model solutions.

In Fig. (4.12) we consider the isotropic initial condition and observe that at small eigentimes the energy density decreases at the same rate as in the hydrodynamic limit irrespective of relaxation time. Over time, however, the larger relaxation times are unable to maintain an isotropic momentum distribution as streaming will tend to increase the transverse bias in the momentum distribution, which slows down the decrease in energy density.

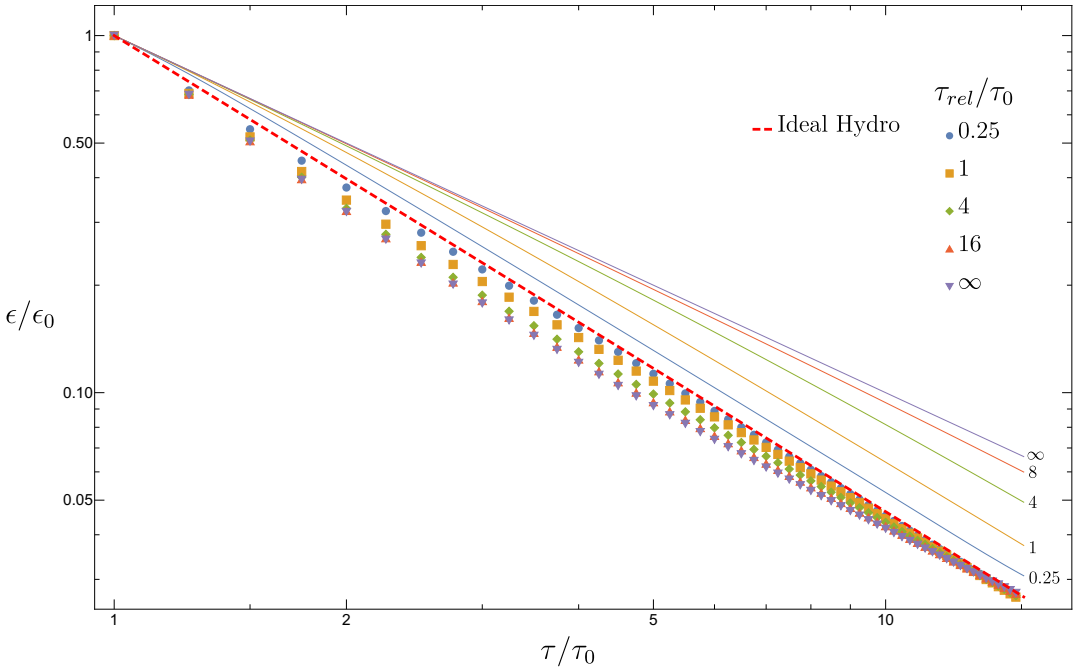

Figure 4.13: Mid-rapidity energy density evolution for a longitudinally-biased initial condition ($\sigma_Y = 1.0$) for a variety of relaxation times, particle number not conserved. The discrete data points represent results from our simulations while solid lines represent the Baym model solutions.

We see in Fig. (4.13) that the initial condition being longitudinally biased increases the rate at which the energy density decreases. By the end of the evolution, each simulation, irrespective of relaxation time, tends to the ideal hydrodynamic energy density.

4.4.1.2 $\mu \neq 0$

We repeat our simulations but for a physical system with particle number conservation. Comparing Figs. (4.11, 4.12, 4.13) to Figs. (4.14, 4.15, 4.16) respectively, we see that particle number conservation has no effect on the evolution of the energy density.

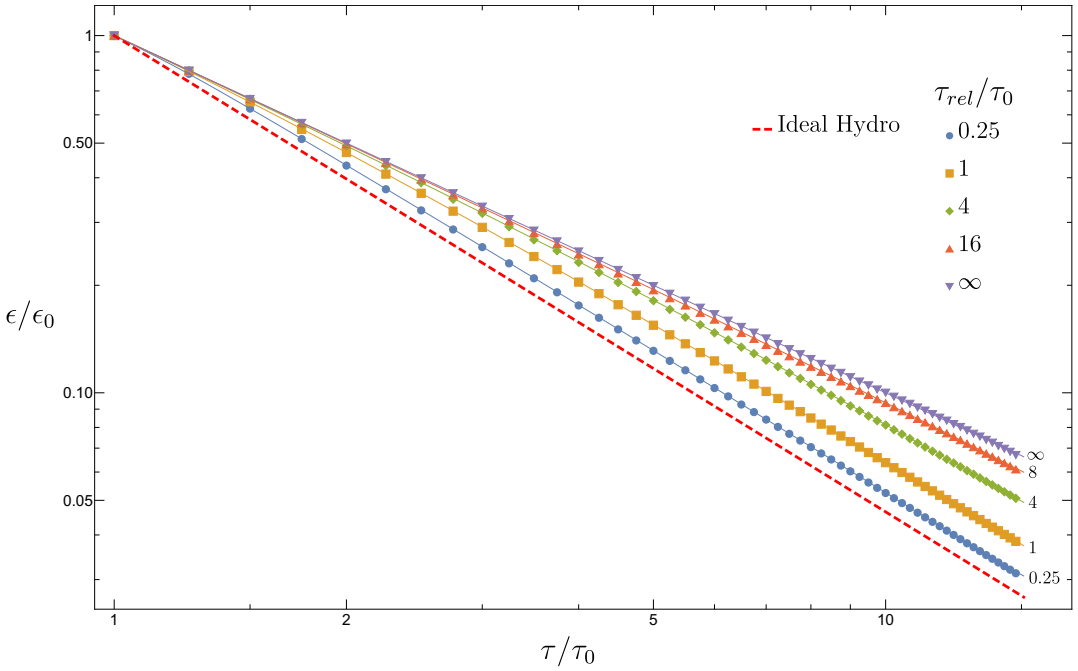

Figure 4.14: Mid-rapidity energy density evolution for an isotropic initial condition ($\sigma_Y \approx 0.58$) for a variety of relaxation times, particle number conserved. The discrete data points represent results from our simulations while solid lines represent the Baym model solutions.

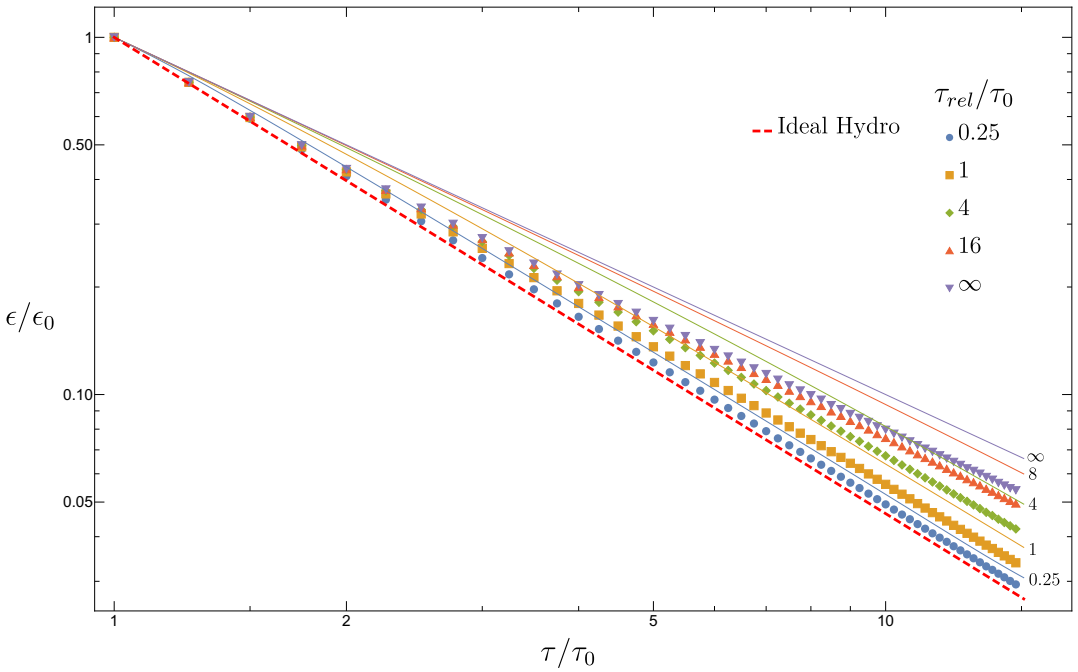

Figure 4.15: Mid-rapidity energy density evolution for an isotropic initial condition ($\sigma_Y \approx 0.58$) for a variety of relaxation times, particle number conserved. The discrete data points represent results from our simulations while solid lines represent the Baym model solutions.

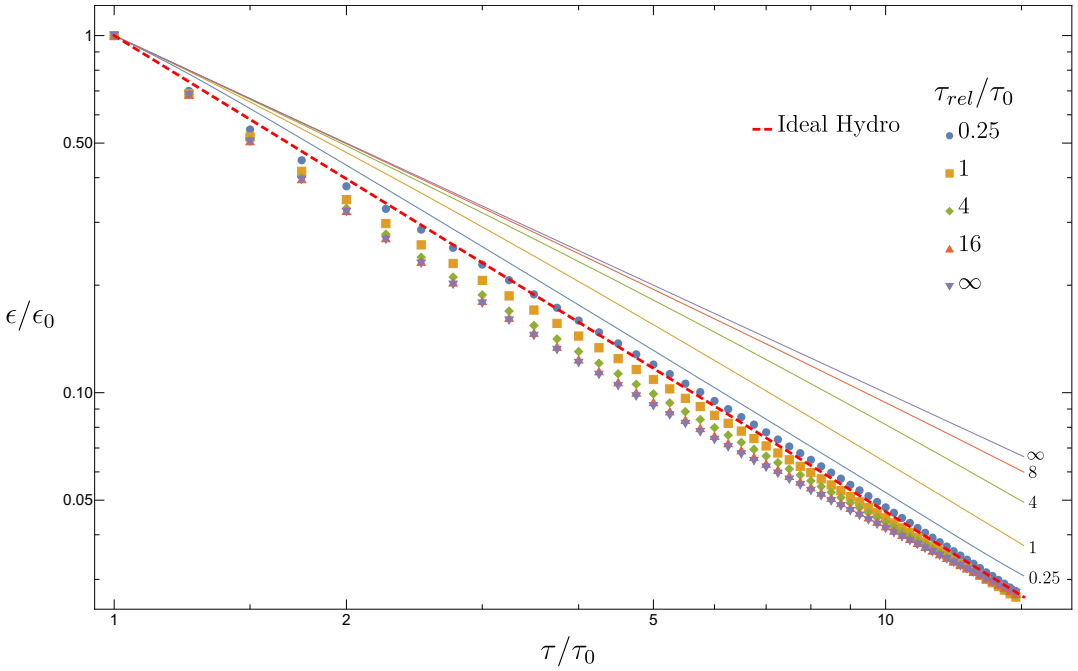

Figure 4.16: Mid-rapidity energy density evolution for a longitudinally-biased initial condition ($\sigma_Y = 1.0$) for a variety of relaxation times, particle number conserved. The discrete data points represent results from our simulations while solid lines represent the Baym model solutions.

4.4.2 Evolution of temperature and chemical potential

We will now examine how both the temperature and chemical potential evolve in time, and across spacetime rapidity.

4.4.2.1 $\mu = 0$

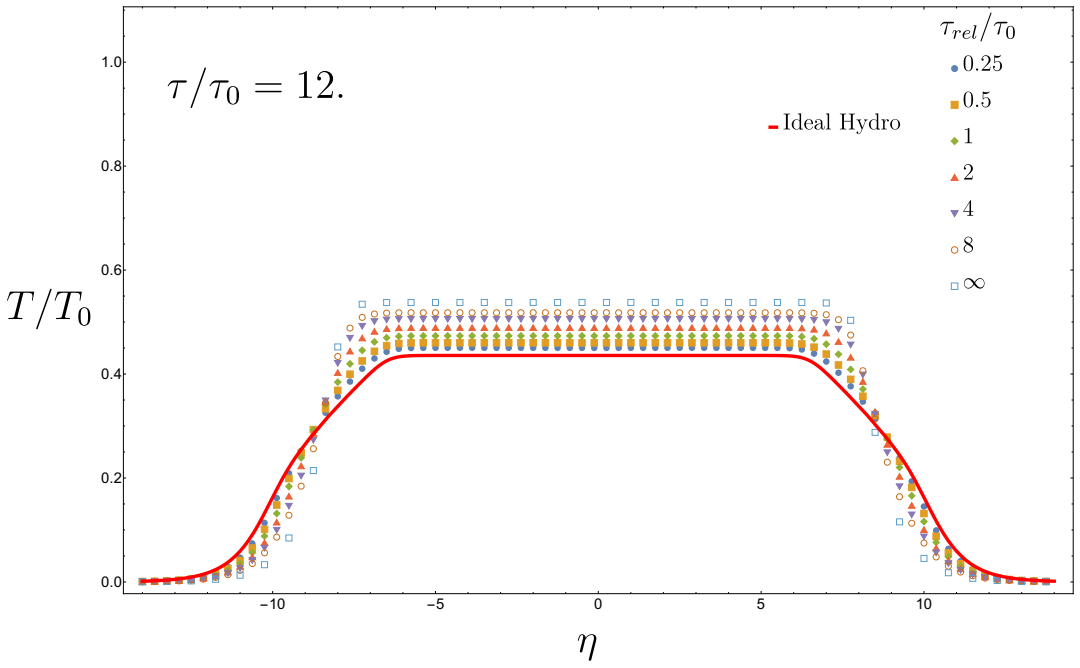

Figure 4.17: Temperature evolution of a transversely-biased initial condition ($\sigma_Y = 0.1$) for a variety of relaxation times, particle number not conserved.

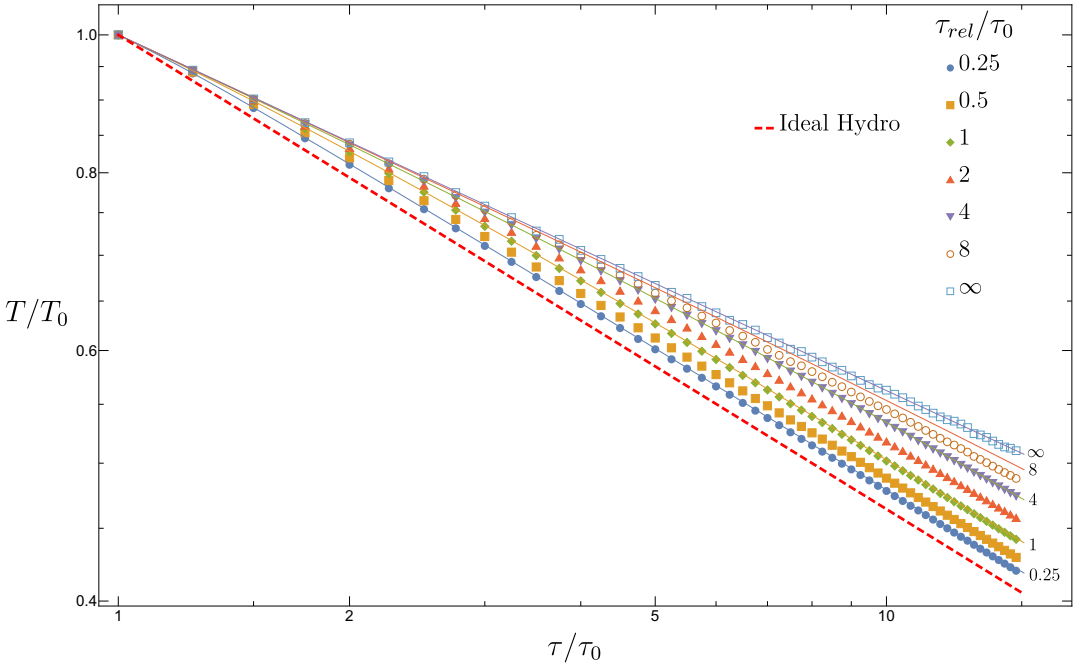

Figure 4.18: Temperature evolution at midrapidity with a transversely-biased initial condition with $\sigma_Y = 0.1$ for a variety of relaxation times, particle number not conserved.

In Fig. (4.17) we see that smaller relaxation times produce more rapid cooling along the central plateau; the broadening of the high-rapidity shoulder indicates this is due to matter flowing into higher rapidity regions as the system isotropises. Consistent with our expectations, smaller relaxation times produced solutions approach the ideal hydrodynamic limit.

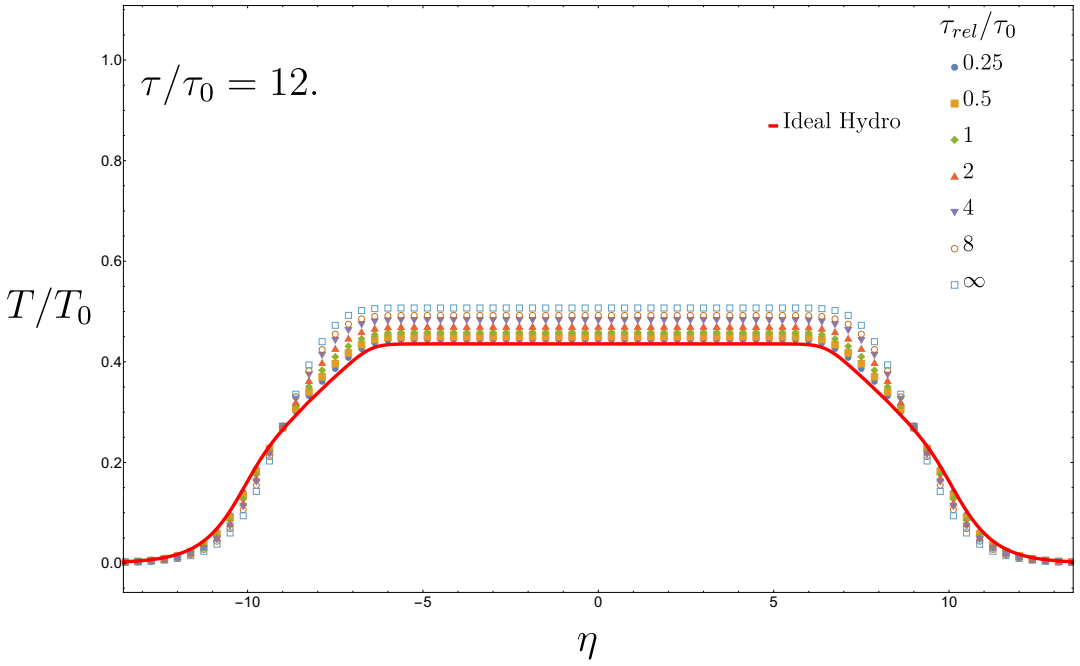

Figure 4.19: Temperature evolution of an isotropic initial condition ($\sigma_Y \approx 0.58$) for a variety of relaxation times, particle number not conserved.

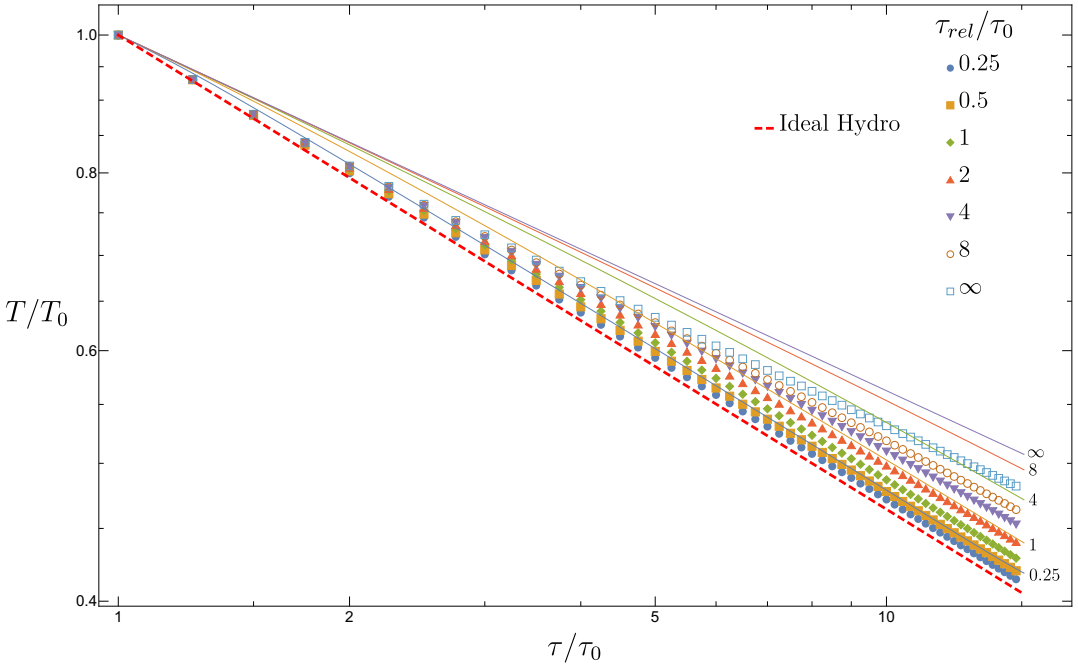

Figure 4.20: Temperature evolution at midrapidity with an isotropic initial condition ($\sigma_Y \approx 0.58$) for a variety of relaxation times, particle number not conserved.

Moving from the transversely-biased anisotropy to an isotropic initial condition in Fig. (4.19), we get the same qualitative results, albeit with a smaller effect size due to relaxation time.

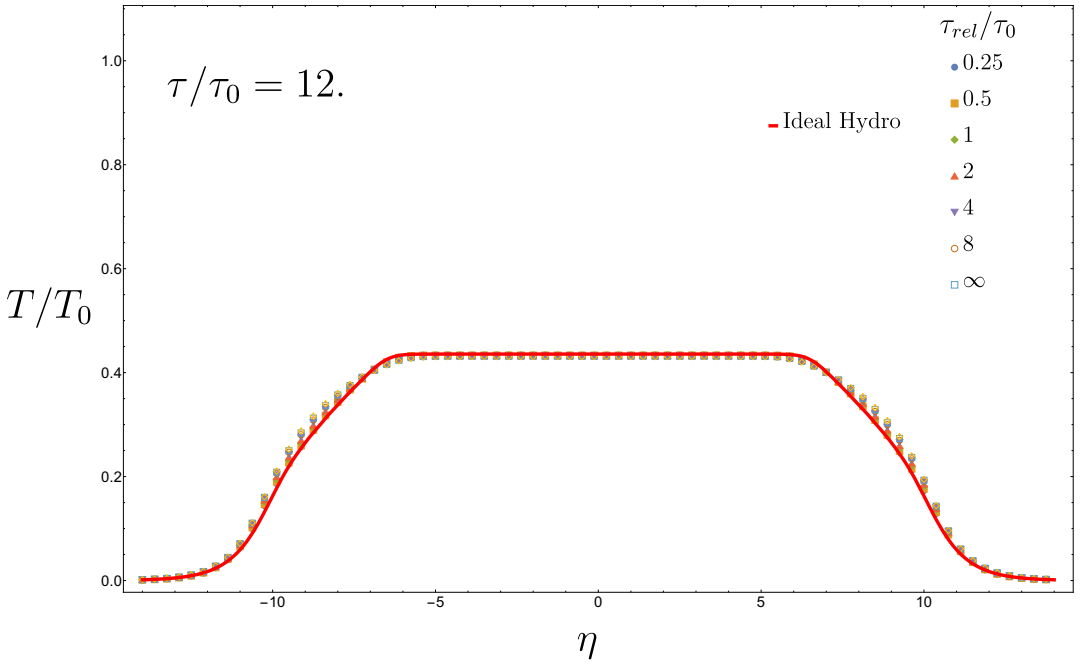

Figure 4.21: Temperature evolution of a longitudinally-biased initial condition with $\sigma_Y = 1$ for a variety of relaxation times, particle number not conserved.

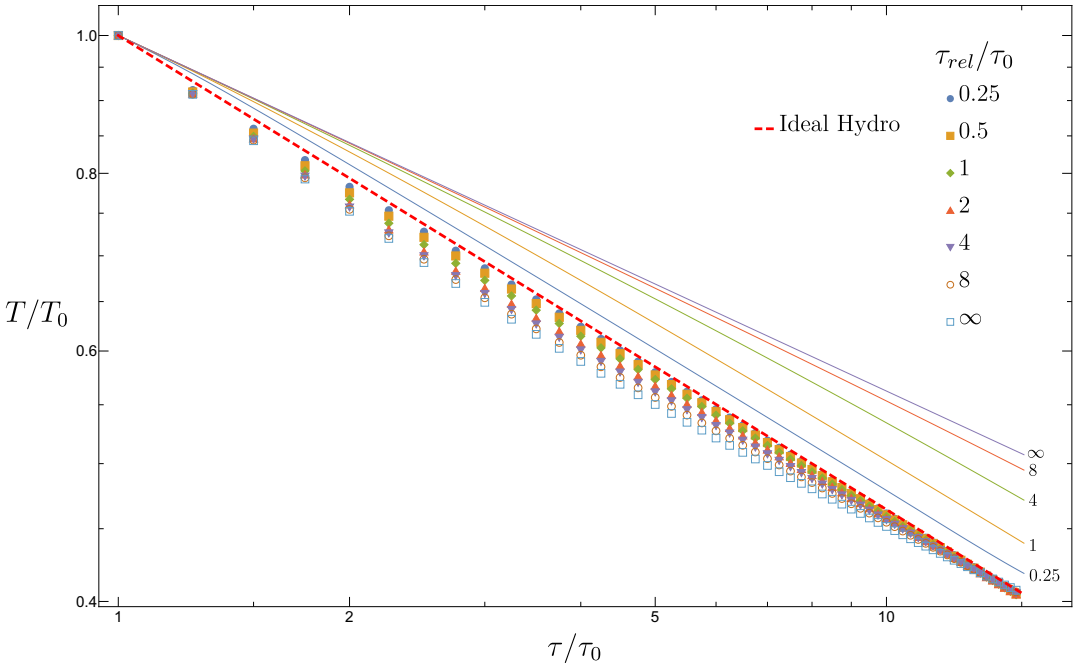

Figure 4.22: Temperature evolution at midrapidity with a longitudinally-biased initial condition with $\sigma_Y = 1$ for a variety of relaxation times, particle number not conserved.

Once we transition over to a longitudinally-biased anisotropy in Fig. (4.21), we lose our ability to differentiate between various relaxation times based on the final temperature of the evolution, and the variation in the intermediate temperatures Fig. (4.22) is small.

In the particle conserving case we see that if there is no relaxation (i.e $\tau_{rel} = \infty$ we see that the drop in energy density is entirely captured by changes in the chemical potential, and the temperature remains constant.

4.4.2.2 $\mu \neq 0$

When we consider a system with particle number conservation, we can track both temperature and chemical potential.

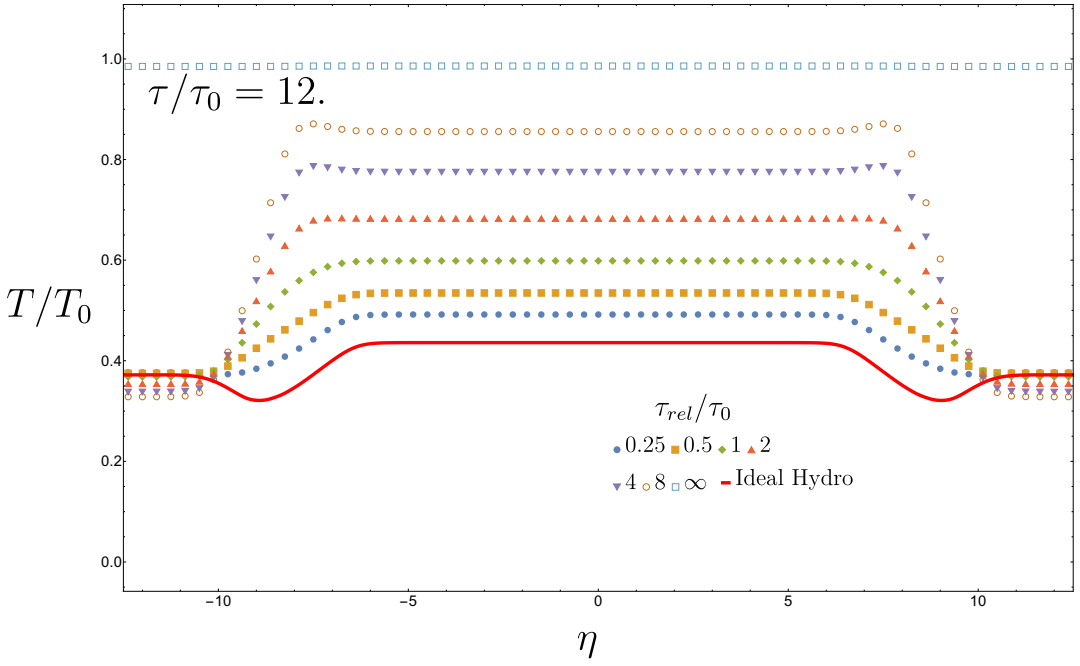

Figure 4.23: Temperature evolution of a transversely-biased initial condition ($\sigma_Y = 0.1$) for a variety of relaxation times, particle number conserved.

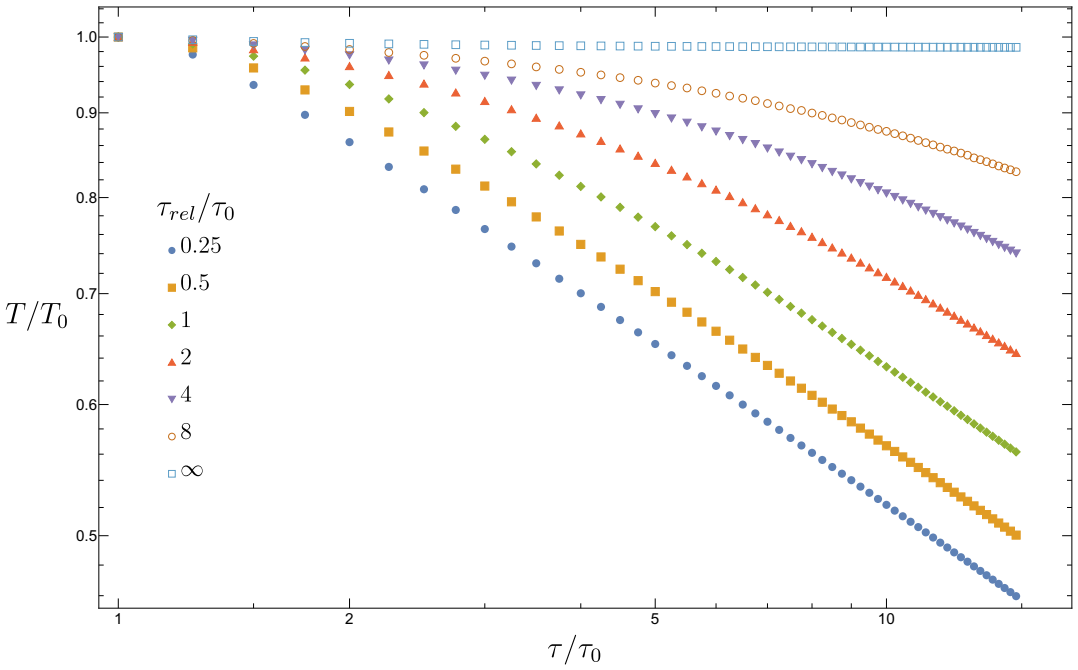

Figure 4.24: Temperature evolution at midrapidity of a transversely-biased initial condition ($\sigma_Y = 0.1$) for a variety of relaxation times, particle number conserved.

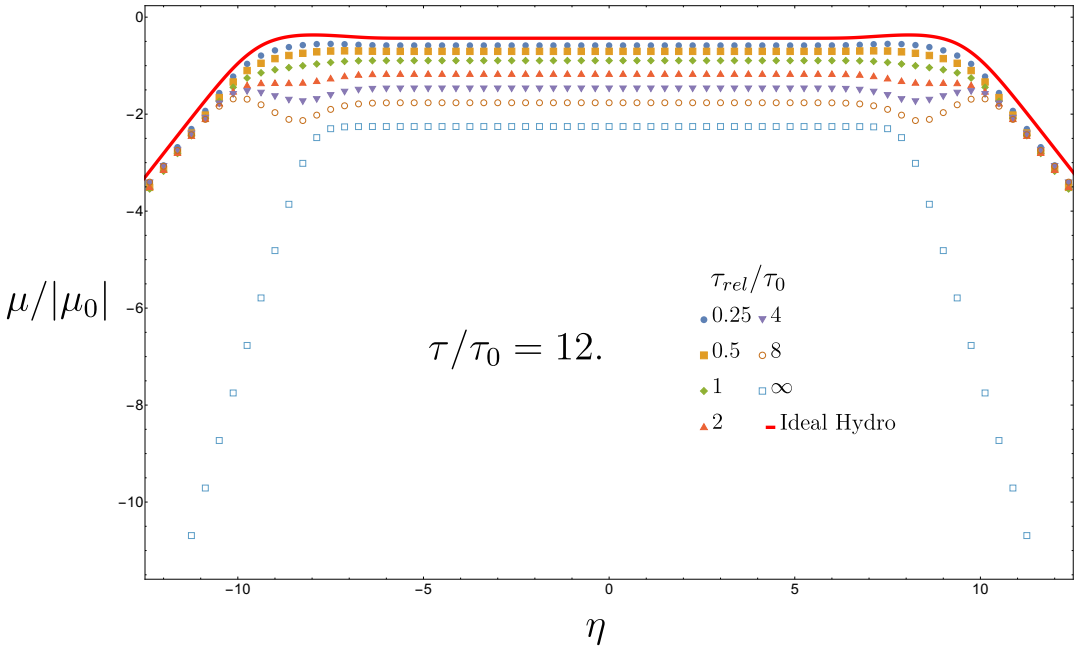

Figure 4.25: Chemical potential evolution of a transversely-biased initial condition ($\sigma_Y = 0.1$) for a variety of relaxation times.

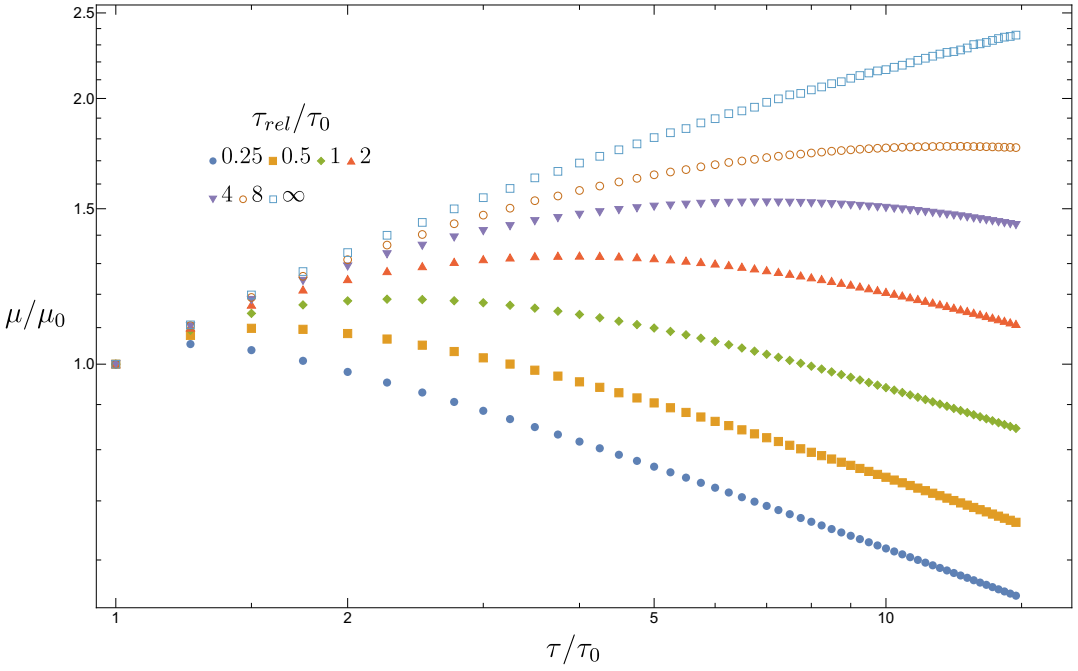

Figure 4.26: Chemical potential evolution at midrapidity of a transversely-biased initial condition ($\sigma_Y = 0.1$) for a variety of relaxation times.

In Figs. (4.23, 4.25) we now see a much stronger dependence on the relaxation time; the evolution when the relaxation time is large is dominated by a change in the chemical potential, and a much smaller relative variation in temperature. Conversely, when the relaxation time is small, the variation in the temperature dominates. We also see a qualitative difference in the evolution of the chemical potential with larger relaxation times driving the chemical potential to become more negative, and smaller relaxation times less negative, compared to the initial chemical potential. In both the temperature and chemical potential, decreasing the relaxation time brings the solution brings the solution closer to the ideal hydrodynamic limit, which is consistent with our expectations.

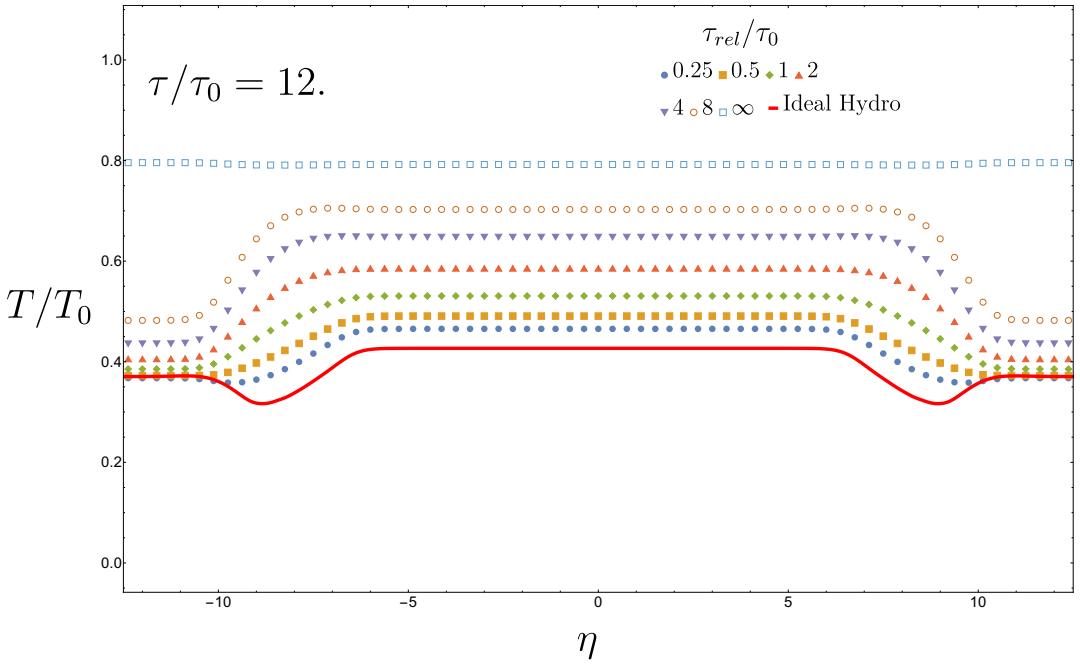

Figure 4.27: Temperature evolution of an isotropic initial condition ($\sigma_Y \approx 0.58$) for a variety of relaxation times, particle number conserved.

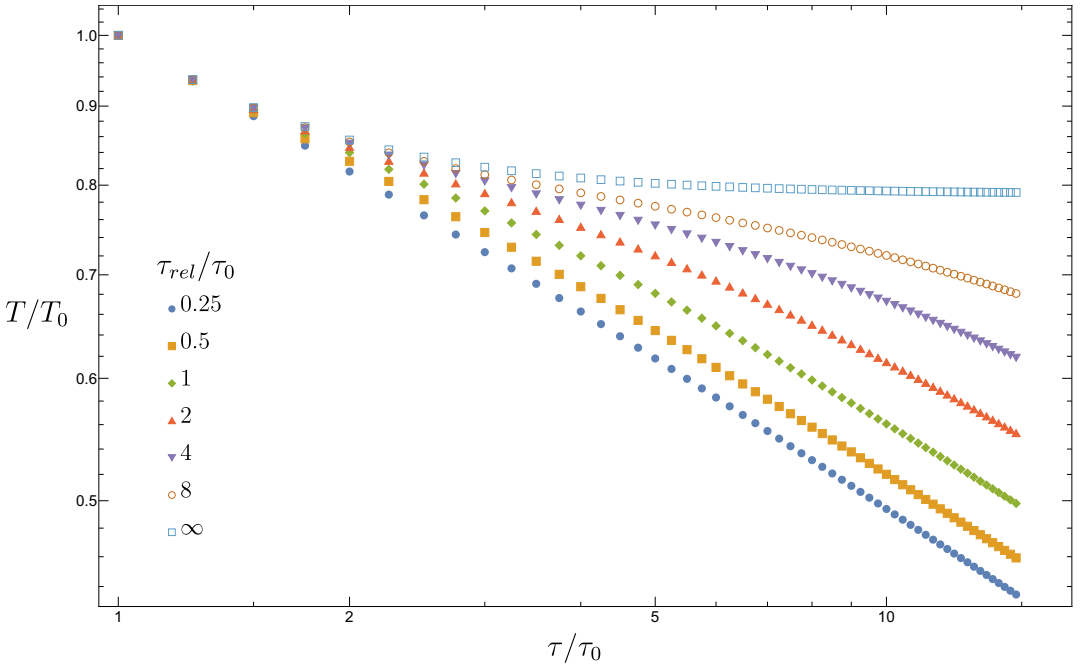

Figure 4.28: Temperature evolution of an isotropic initial condition ($\sigma_Y \approx 0.58$) for a variety of relaxation times, particle number conserved.

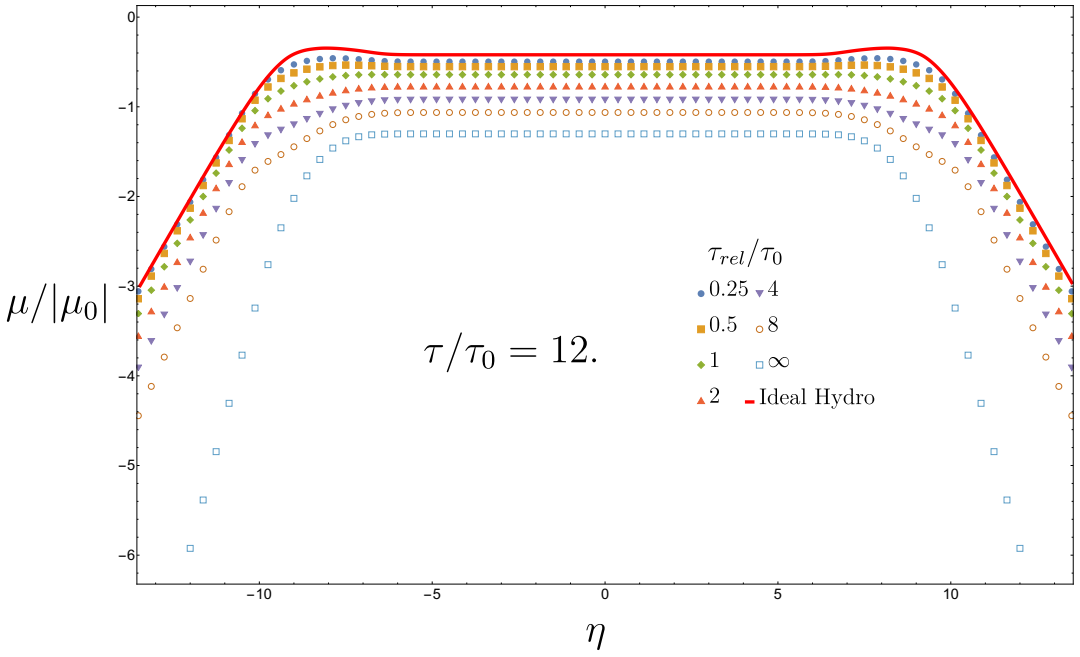

Figure 4.29: Chemical potential evolution of an isotropic initial condition ($\sigma_Y \approx 0.58$) for a variety of relaxation times.

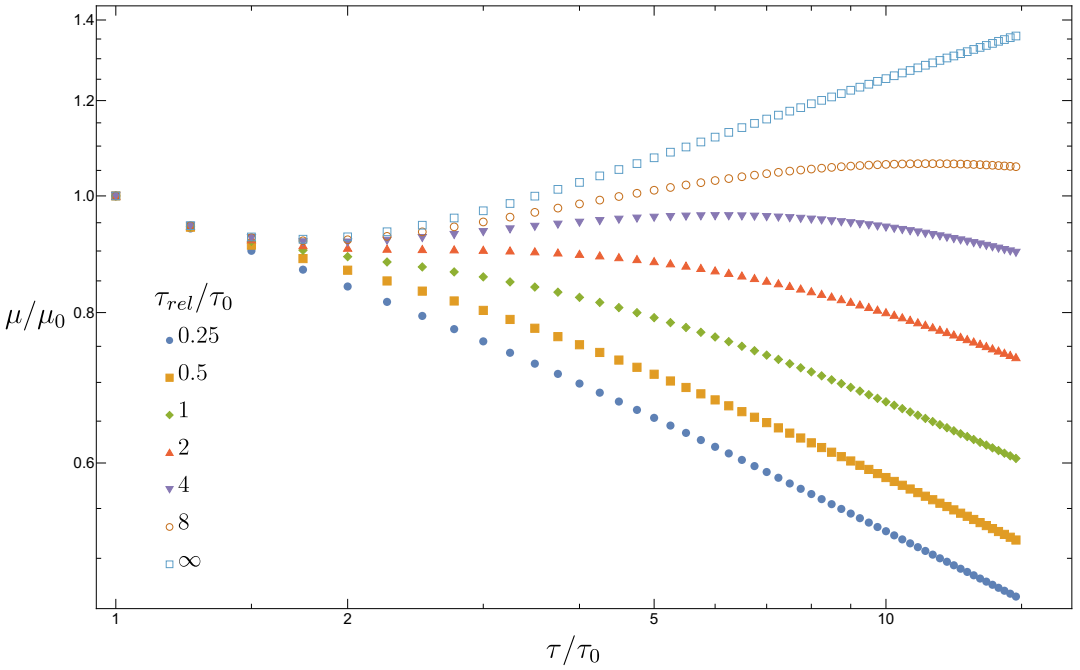

Figure 4.30: Chemical potential evolution at midrapidity of an isotropic initial condition ($\sigma_Y \approx 0.58$) for a variety of relaxation times.

In the isotropic case, we once again see the same qualitative behaviour between the various relaxation time, but the magnitude of the effect of relaxation time is smaller.

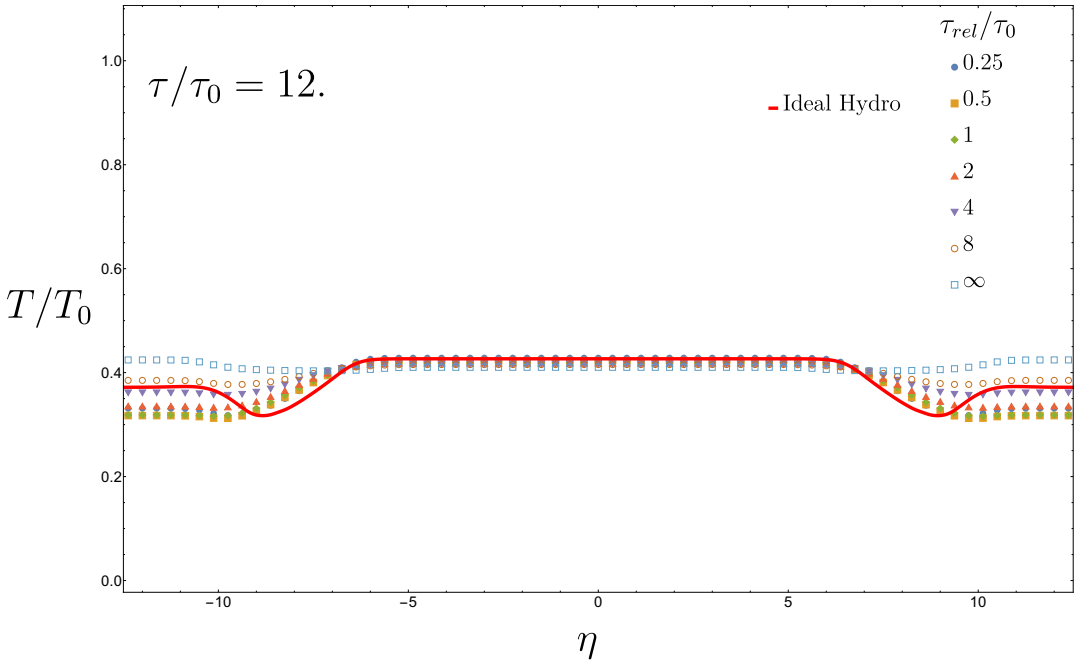

Figure 4.31: Temperature evolution of a longitudinally-biased initial condition ($\sigma_Y = 1$) for a variety of relaxation times, particle number not conserved.

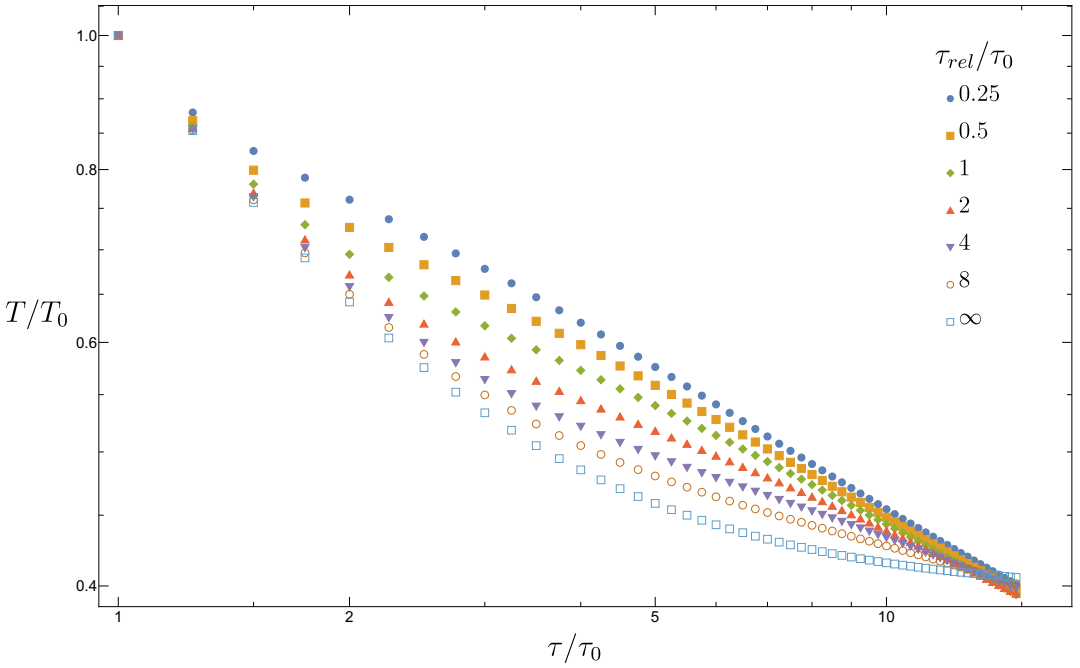

Figure 4.32: Temperature evolution at midrapidity of an longitudinally-biased initial condition ($\sigma_Y = 1$) for a variety of relaxation times, particle number not conserved.

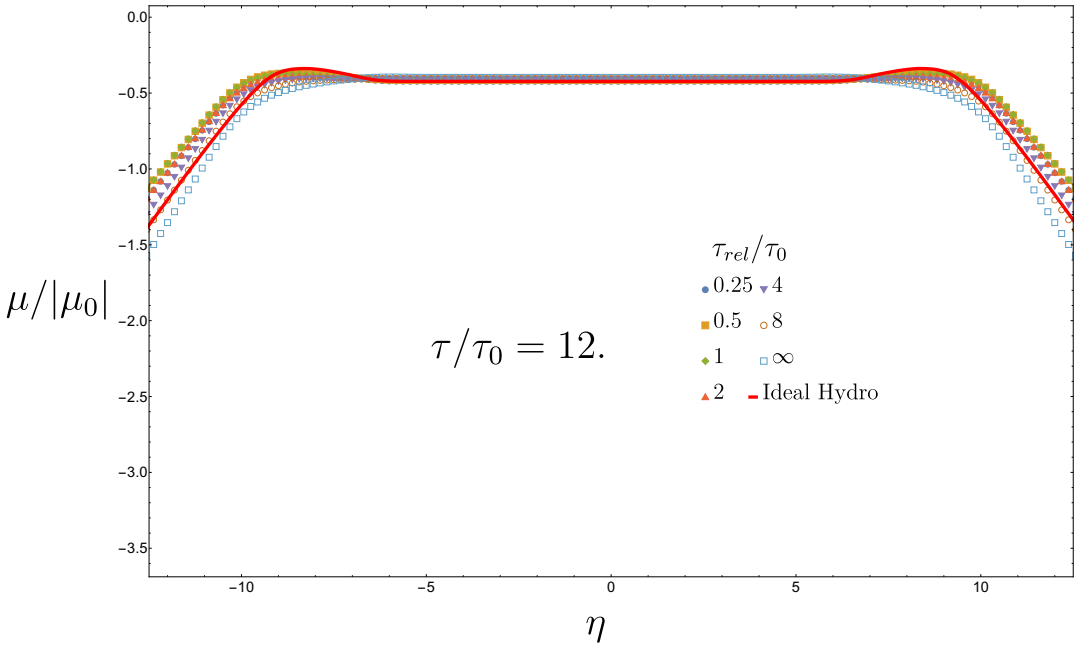

Figure 4.33: Chemical potential evolutionof a longitudinally-biased initial condition ($\sigma_Y = 1$) for a variety of relaxation times.

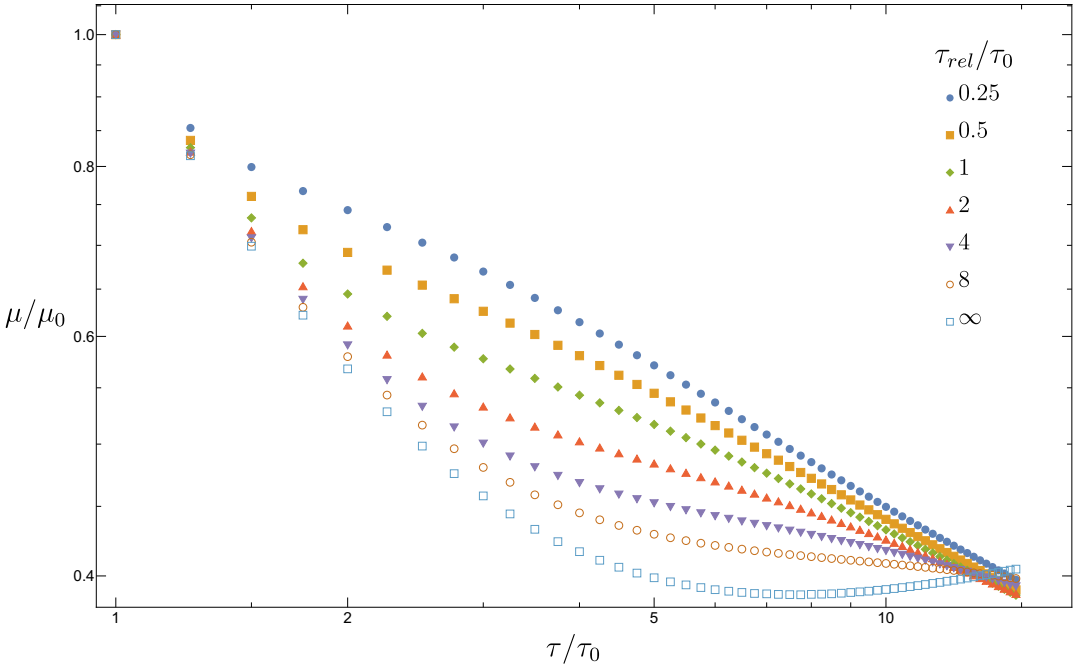

Figure 4.34: Chemical potential evolution at midrapidity of a longitudinally-biased initial condition ($\sigma_Y = 1$) for a variety of relaxation times.

The effect of the longitudinal-bias is once again to almost completely mitigate any effect of the relaxation time on the final state of the system, but a

signficant effect on the intermediate stages.

4.4.3 Effect of relaxation time on longitudinal flow acceleration

One manifestation of the boost invariance of the initial condition is the longitudinal fluid momentum rapidity is initially equal to the spacetime rapidity. Previous hydrodynamical models have shown that during the evolution, there is an acceleration of the fluid flow, in a manner that depends on the equation of state and viscosity [40, 41]. As the Baym model explicitly enforces boost-invariance for all time, no such acceleration will be present in the model. We will thus make a comparison to the ideal hydrodynamic model with various relaxation times and initial momentum anisotropies.

4.4.3.1 $\mu = 0$

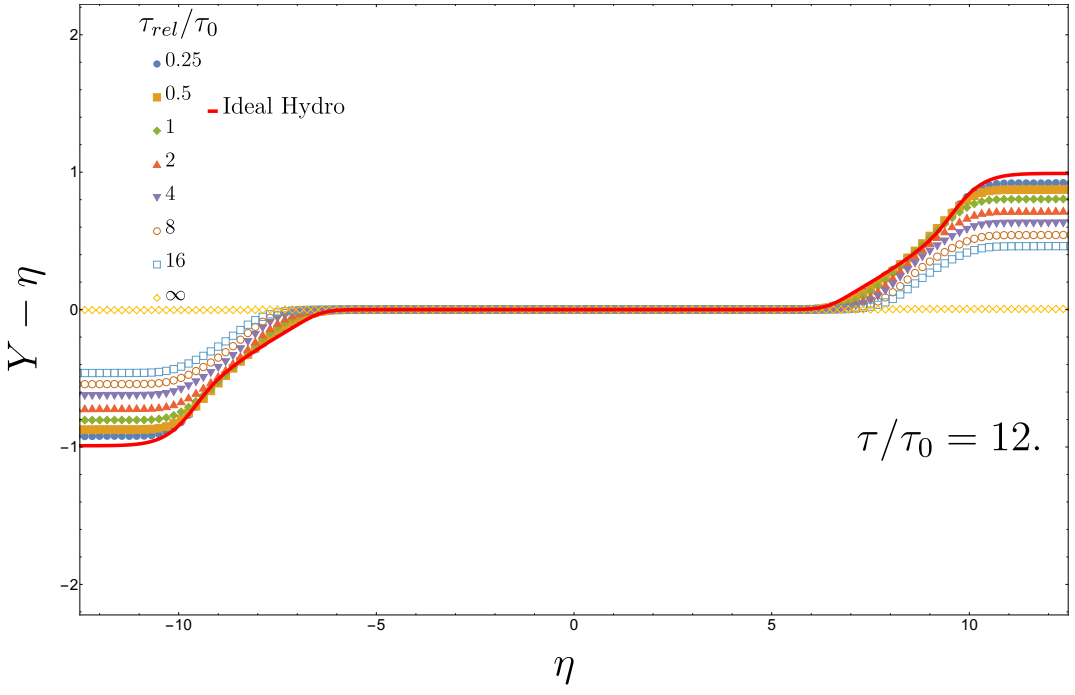

Figure 4.35: Difference between longitudinal fluid flow rapidity and spacetime-rapidity for $\sigma_Y = 0.1$ for various relaxation times, without particle number conservation.

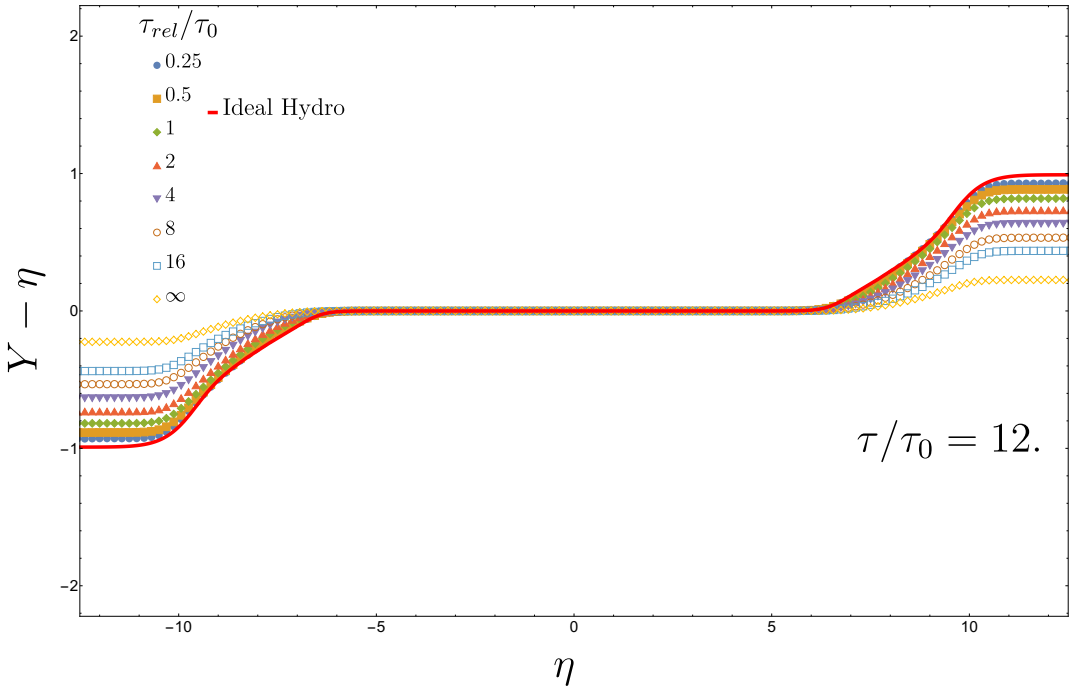

Figure 4.36: Difference between longitudinal fluid flow rapidity and spacetime-rapidity for $\sigma_Y = 0.5$ for various relaxation times, without particle number conservation.

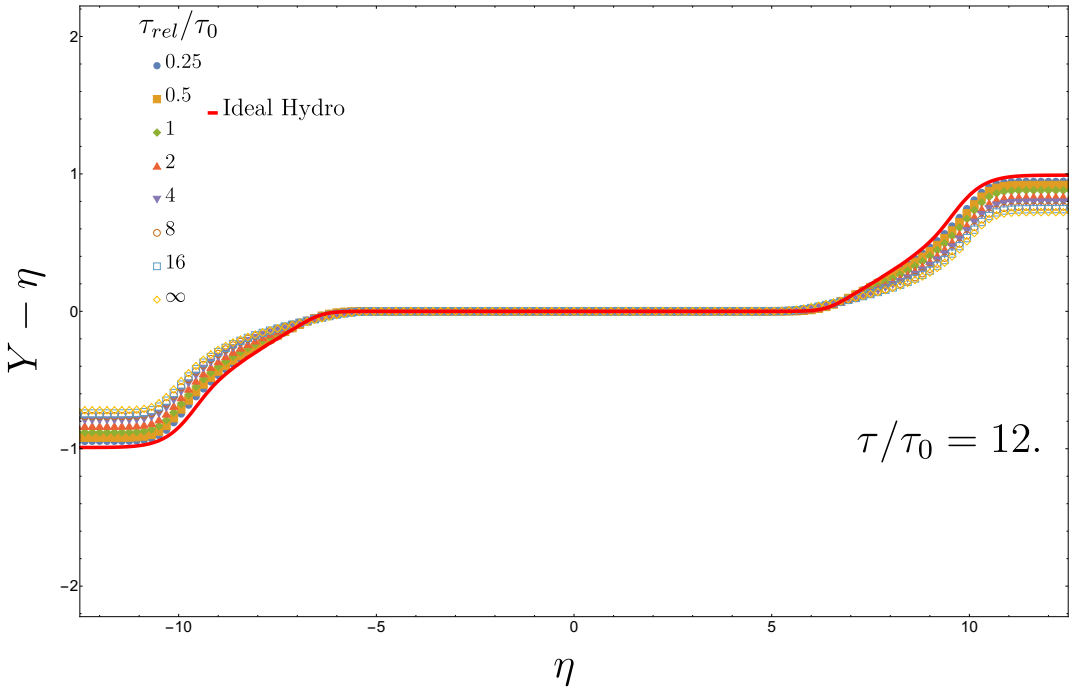

Figure 4.37: Difference between longitudinal fluid flow rapidity and spacetime-rapidity for $\sigma_Y = 1.0$ for various relaxation times, without particle number conservation

Smaller relaxation times drive a stronger acceleration at the high rapidities. The fact that we see this acceleration with an infinite relaxation time in the isotropic and longitudinally-biased scenarios, but not in the transversely-biased case indicates that this acceleration is due high longitudinal rapidity particles flowing to higher rapidities.

4.4.3.2 $\mu \neq 0$

Comparing Figs. (4.35, 4.36, 4.37) to Figs. (4.38,4.39, 4.40) respectively, we see that particle number conservation does not affect the acceleration of the flow velocity.

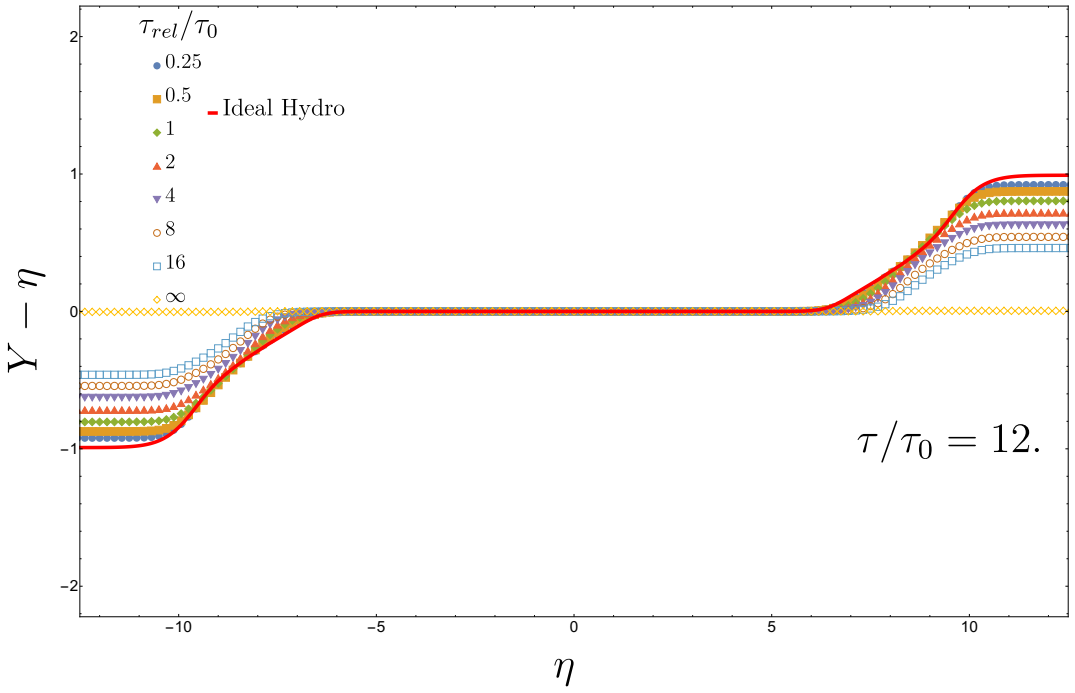

Figure 4.38: Difference between longitudinal fluid flow rapidity and spacetime-rapidity for $\sigma_Y = 0.1$ for various relaxation times, with particle number conservation.

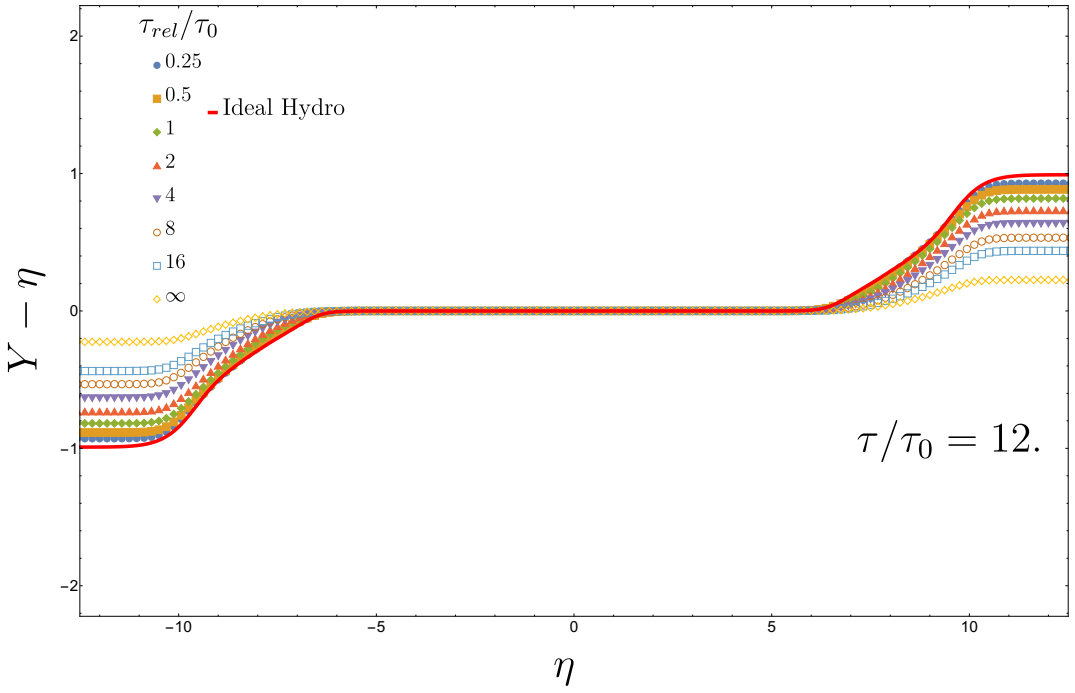

Figure 4.39: Difference between longitudinal fluid flow rapidity and spacetime-rapidity for $\sigma_Y \approx 0.58$ for various relaxation times, with particle number conservation.

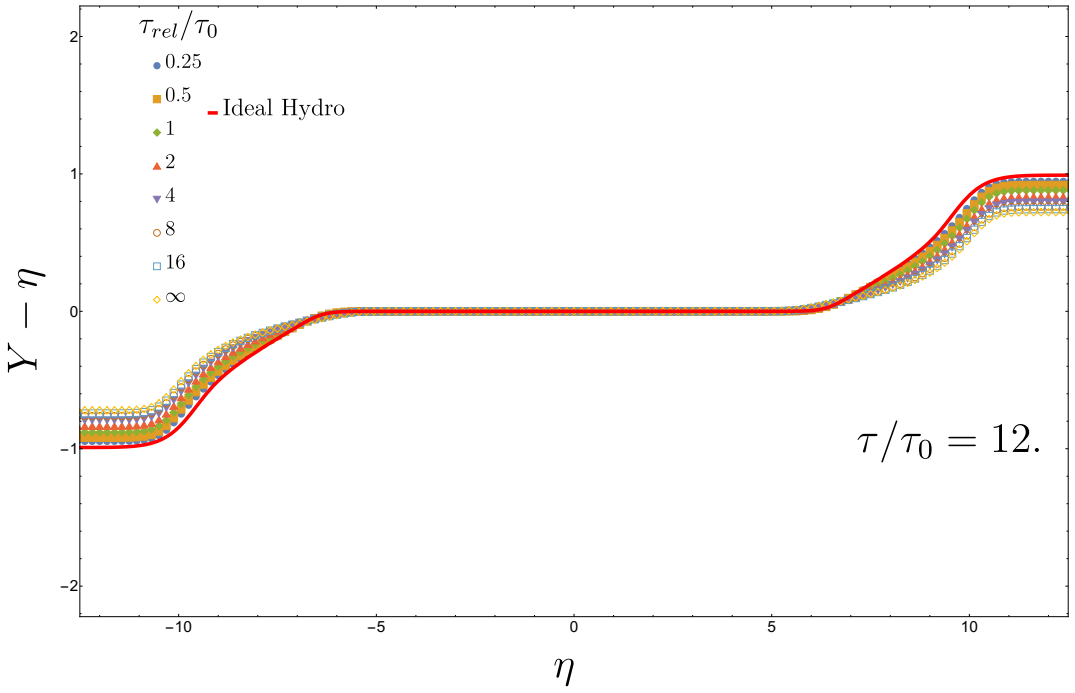

Figure 4.40: Difference between longitudinal fluid flow rapidity and spacetime-rapidity for $\sigma_Y = 1.0$ for various relaxation times, with particle number conservation.

Chapter 5

Full 3D simulation

In the previous section, we have shown that our simulation can replicate analytical results that make different assumptions in the case of tranverse isotropy and homogeniety. As proof of concept, we will now extend our simulations not to require to be homogenous and isotropic in the transverse plane.

We use the same initial as in Eq. (4.13), but we add a Gaussian term that suppresses the distribution function in the transverse plane.

For proof of concept we consider a collision of nuclei with radius 4fm, we would expect the finite nucleus size will cause more rapid cooling of the system.

5.1 $\mu = 0$

We first consider the temperature evolution for a number of relaxation time without particle number conservation.

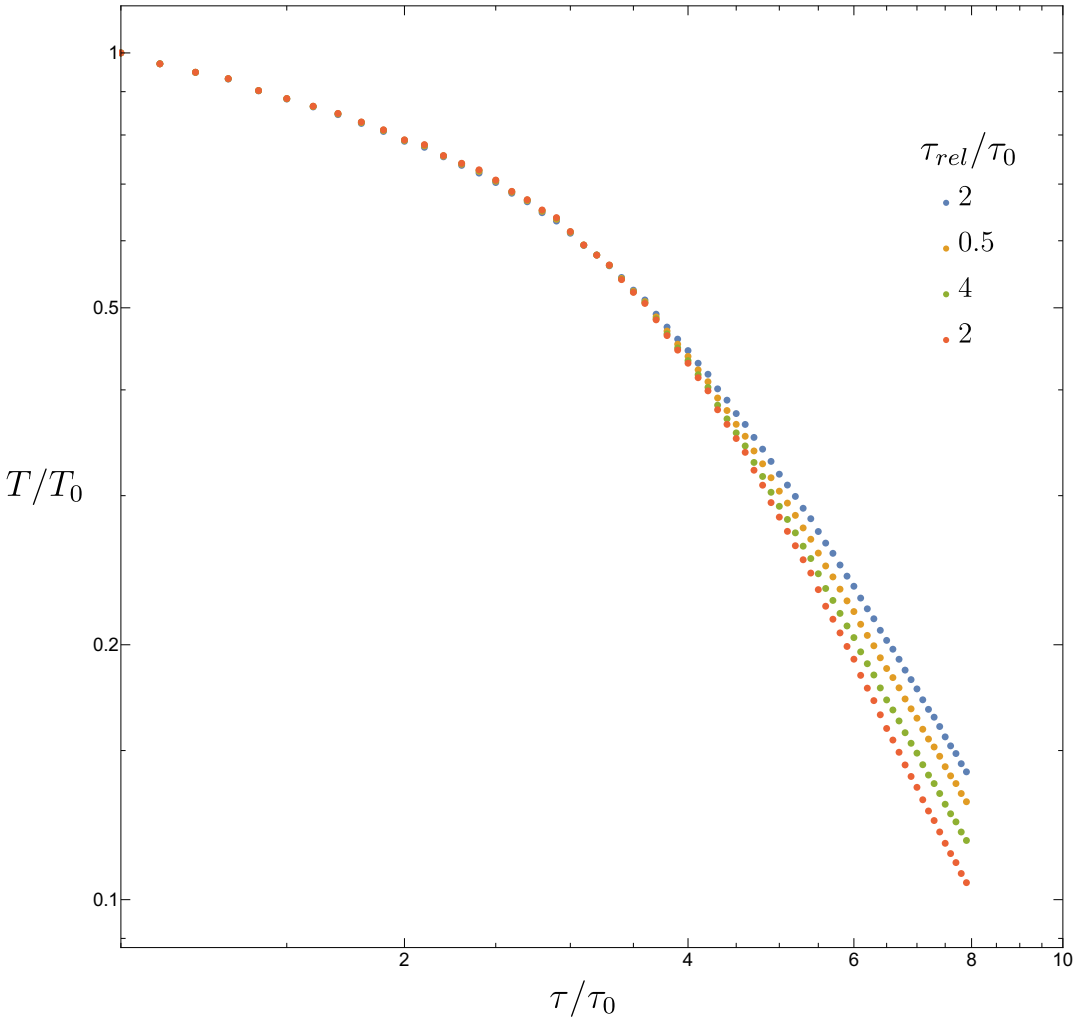

Figure 5.1: Temperature evolution at midrapidity , and centrally in the transverse plane, with a transversely-biased initial condition ($\sigma_Y = 0.3$) for a variety of relaxation times, particle number not conserved.

As expected, more rapid cooling takes place as compared to the cylindrically symmetric case.

5.2 $\mu \neq 0$

Now, if we consider the case where we do impose particle number conservation.

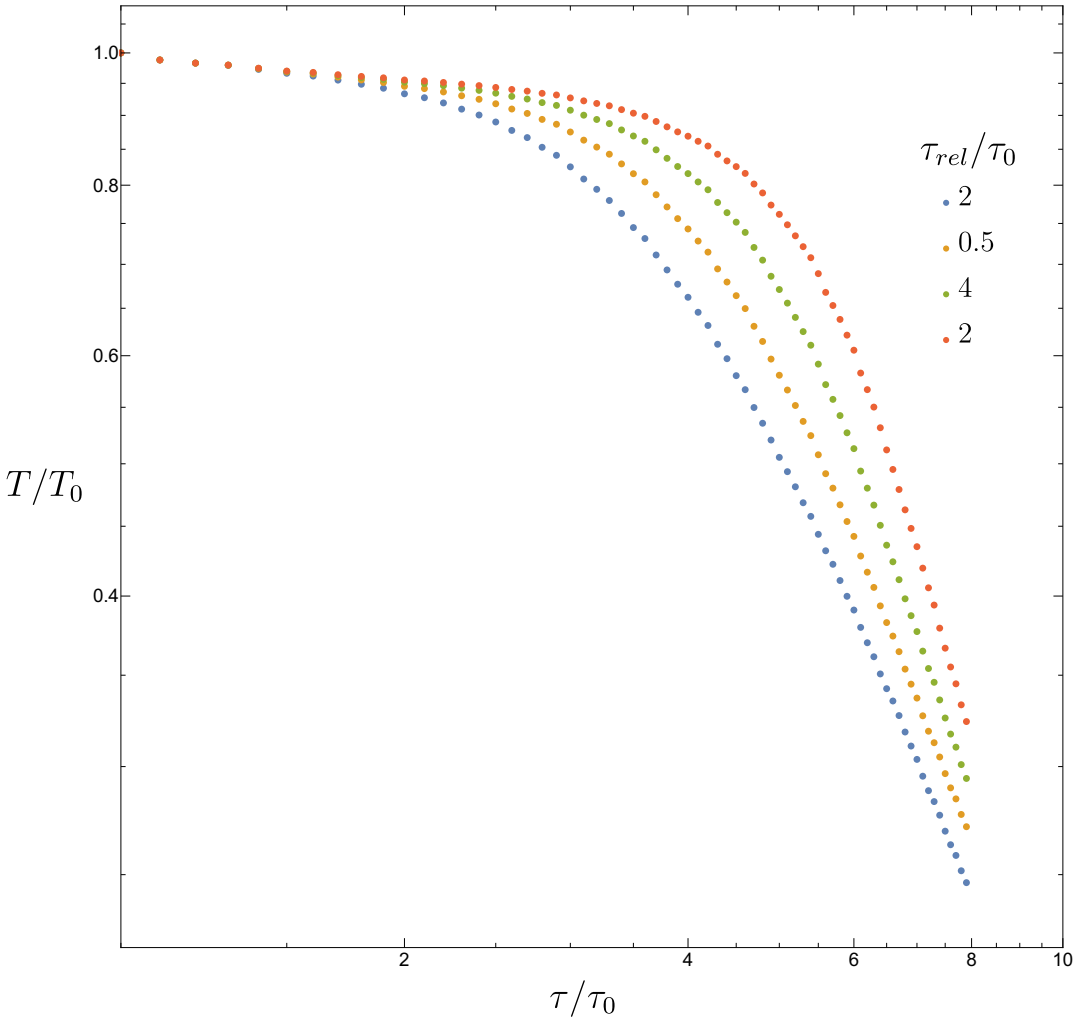

Figure 5.2: Temperature evolution at midrapidity , and centrally in the transverse plane, with a transversely-biased initial condition with $\sigma_Y = 0.3$ for a variety of relaxation times, with particle number conservation

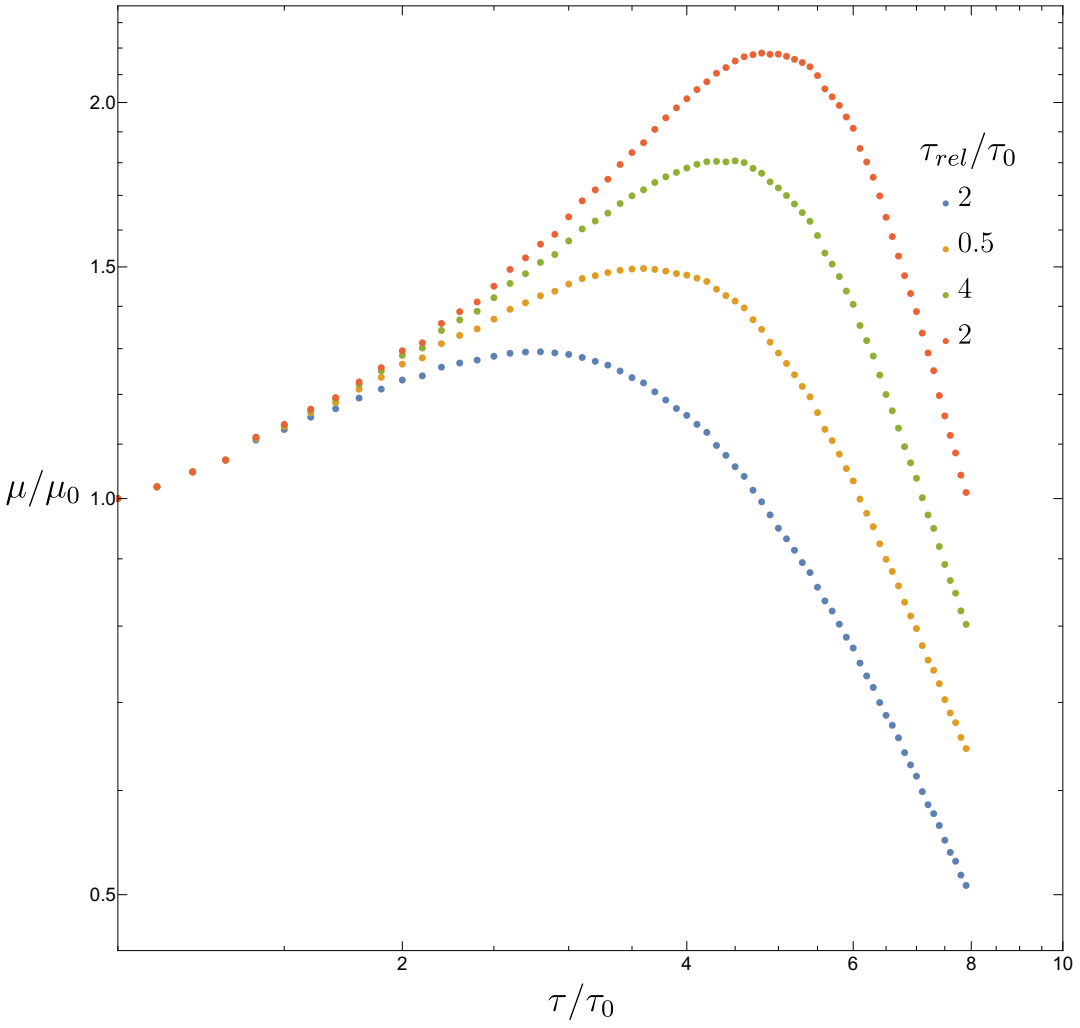

Figure 5.3: Chemical potential evolution at midrapidity, and centrally in the transverse plane, with a transversely-biased initial condition with $\sigma_Y = 0.3$ for a variety of relaxation times, with particle number conservation

Again we see more rapid cooling; however, the dynamics of the chemical potential are particularly interesting, with a strong effect as the finite size of the nucleus becoming apparent to the system.

These results highlight show the potential of this simulation software, having the ability to model different interaction schemes at various rapidities in 3D is unique.

Chapter 6

Summary and Outlook

In this thesis, we have developed an algorithm that efficiently simulates the Boltzmann equation under physically relevant conditions, and implemented this on GPUs. The coupling of this efficient algorithm and the computational horsepower granted by the modern GPUs allows us to numerically solve the Boltzmann equation for physically relevant initial conditions in full 3D space.

We have shown that our implementation shows agreement with analytical results, under the relevant limits in the case of homogeniety and isotropy in the transverse plane; and how these results begin to deviate once those limits are relaxed.

We then demonstrated that we can drop the requirement of cylindrical symmetry and simulate the full 3D case.

As of yet, we have not implemented a freezeout procedure for the QGP to hadron gas transition. Developing, or using an existing freezeout implementation, is a prerequisite for making direct comparisons with experimental data.

Work needs to be done on how to implement the initial conditions used in hydrodynamical simulations to make more direct comparisons.

Work has been done to build a collision term that incorporates more aspects of QCD physics, as of yet simulations that implement these new collision terms have no included streaming. The modularity of the simulation tool

we have developed means that only the relaxation needs to be modified to incorporate the new collision term. Implementing This will allow the possibility for us to determine if a Bose-Einstein condensate might form under experimental conditions and may shed light on the mystery of the rapid thermalization of the QGP.

All simulations were run on a laptop with an NVIDIA GeForce 960M with 4GB of VRAM. This VRAM limitation limited the size of systems we could simulate (especially relevant in the 3D case). At the time of writing, there exist single high-performance GPUs with up to 16GB of VRAM. It is also possible to run on multiple GPUs simultaneously, although this would likely be a non-trivial exercise. We suggest that future optimization work should focus on minimizing VRAM use by finding the optimal parameters to minimize the number of required gridpoints without sacrificing accuracy.

Appendix A

The CUDA programming model

The demand placed on hardware to render increasingly complex 3-dimensional games in real-time at ever higher resolutions has driven the development of specialised graphics processing Units (GPUs) to complement the computational resources of the central processing unit (CPU). The rendering operations that need to be performed are typically numerically intensive, and at relatively long intervals of $10^{-2}s$ compared to the time scale of a single CPU clock cycle $10^{-9}s$.

The solution GPU manufacturers have developed is a parallel, many-core processor with high memory bandwidth that prioritises total computational throughput over latency and is optimised for simple numerically intense schemes (as opposed to algorithms that have complex decision trees). The benefit is scalability, increases in CPU performance is primarily driven by increases in the clock speed. Increases in clock speeds however exponentially increase heat production. The architechture of modern GPUs however is to link hundreds or even thousands of reduced instruction set processors together on a single board which produces a linear increase in heat production. This allows a more progressive scaling of computational power.

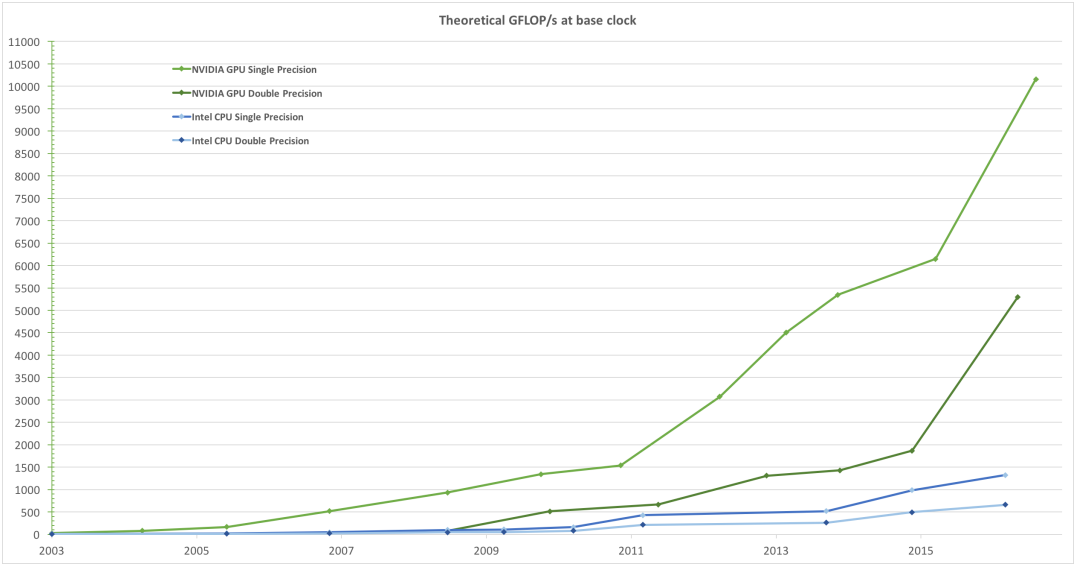

FIGURE A.1: Comparison of Floating-point operations per second theoretically possible on the CPU and GPU. Image from [42].

There is no free lunch, however, and in order to effectively harness the power of modern GPUs, it is necessary to understand the programming model and its limitations. If, however, the computational horsepower of modern GPUs can be harnessed, one gets supercomputer scale performance on a consumer grade desktop.

With the development of novel algorithms, this increase in possible computational power has recently been leveraged by many fields to accelerate their scientific computation, perhaps the most ubiquitous contemporary application is the use of GPU computation to train neural network models [43]. We however have also simulations of biological systems [44], astrophysical simulations [45, 46], hydrodynamic simulations of the quark-gluon plasma [47] as some further examples.

We discuss various considerations that go into determining if one can expect a significant performance benefit from implementing an algorithm on the GPU.

The algorithm design needs to be inherently parallel, where preferably many threads simultaneously perform completely independent calculations, as the possibly high latencies can cause threads to wait idly for a significant time if they depend on the output of another thread.

The algorithm should be numerically intensive as the reduced instruction set available on the GPU is optimised for numerical computation with many basic floating point operations implemented in hardware such that they require far fewer clock cycles to evaluate compared to the standard x86_64 architecture. Large linear algebra operations have in particular been significantly optimised.

The most significant computational burdens in our problem are numerical integration in order to determine the energy-momentum tensor and particle-current vector, and calculating the updates to the distribution function given a set of equilibrium parameters; which are inherently numerical, so this condition is met.

Memory access is expensive; each thread only requires access to a few of memory, there is less of a performance penalty of the memory access is coalesced (i.e. the memory read/writes required by a single are adjacent in memory). The fact that the Boltzmann equation is local ensures this property is met.

GPUs operate on a SIMD (single instruction, multiple data) execution model, which can suffer significant performance penalties if there is a large amount of control flow present in the algorithm. Our algorithm does not have significant control flow present.

Appendix B

Gauss-Legendre integration

Our goal is to accurately integrate distribution function at each space-time coordinate using as few momentum points as possible. We can achieve this minimisation by picking an appropriate integration procedure. The integration procedure we use is Gauss-Legendre integration.

We approximate the integral by evaluating the function at x_i (called the abscissa) and the corresponding weights w_i,

$$\int_a^b dx\, f(x) \approx \sum_{i=0}^{n} w_i f(x_i). \tag{B.1}$$

The procedure for determining these abscissae and weights is not unique, the Gauss-Legendre method is one such procedure. The benefit of using the Gauss-Legendre method for determining the abscissa and weights is that one is that the order of integration is $2n-1$ (i.e. one can integrate polynomials of order up to $2n-1$ exactly) whereas naive choices of abscissae and weights are typically to order n. Another benefit of the Gauss-Legendre procedure is that the weights generated vary only gradually instead of rapidly oscillating, which can lead to numerical instability.

The n abscissa are chosen to be the zeroes of the Legendre polynomial of order n. The zeroes of all Legendre polynomials lie in the interval $(-1, 1)$.

The Gauss-Legendre scheme is only able to integrate over the interval $(-1, 1)$ however the distribution functions that we would want to integrate (although we have an expectation of exponential fall off at sufficiently large momenta).

We need to introduce a change of variables that will map from the interval $(-\infty, \infty) \rightarrow (-1, 1)$. With some foresight we can anticipate Bose-Einstein distribution which is singular at zero; therefore our mapping function needs to exclude zero for the sake of numerical stability.

So far we have described how to perform numerical integration along a single dimension. The integrals that we need to do in practice, however, are multidimensional. Fortunately, this technique is readily extended to multiple dimensions, one first performs the Gauss-Legendre procedure for each coordinate axis independently, then the function being integrated is evaluated at every possible combination independent abscissa weighted by the product of the Gauss-Legendre weights of those points, and then summed.

$$\int d^D x \, f(x^1, .., x^D) \approx \sum_{i=0}^{n_1} ... \sum_{k=0}^{n_D} w_i^1 ... w_k^D f(x_i^1, ..., x_k^D) \tag{B.2}$$

The choice of mapping function will influence the accuracy of the integration technique. A good mapping function should be able to map typical distribution functions we expect to a function smoothly interpolated by a polynomial in the domain $(-1, 1)$.

As discussed previously, our mapping function may not map to zero. To this end, we split the integration domain into two separate halves,

$$\int_{-\infty}^{\infty} dk_i = \int_{-\infty}^{0} dk_i + \int_{0}^{\infty} dk_i \tag{B.3}$$

.

The transformation that we will use is the following,

$$k_i(x_i) = \pm\Lambda \left(\frac{1}{1 - x_i} + \frac{1}{2} \right), \tag{B.4}$$

with $x \in (0, 1)$, which has the Jacobian factor,

$$J_i(x_i) = \frac{\Lambda}{(1 - x_i)^2}, \tag{B.5}$$

associated with it, where Λ is a free parameter that sets the overall momentum scale.One trivially gets the $\int_{-\infty}^{0} dk_i$ points by just mapping $k_i \rightarrow -k_i$. The momentum scale should be proportional to the momentum scale over which the function (i.e the initial condition) being integrated varies, as the choice of Λ will generically affect the accuracy of the numerical integration.

The number of points to used will depend on the details on of the system being studied and the accuracy desired. There is, as always a tradeoff between accuracy and system size. Since there is an anistoropy between transverse and longitudinal drections, the number of number of points sampled transversely and longitudinally need not be the same.

Appendix C

Outline of the Baym solution

We provide an outline of the solution.

$$f(k_\perp, k_z, z, t) = f(k_\perp, k'_z, \tau) \tag{C.1}$$

where

$$k'_z = (k_z - ku)\gamma \tag{C.2}$$

is the boosted momentum, with boost velocity $u = z/t$ and $\gamma = 1/\sqrt{1-u^2}$.

The assumption of boost-invariance implies that it is sufficient to solve the Boltzmann equation at mid-rapidity, and then boost the solution to solutions at non-zero spacetime rapidities.

The Boltzmann equation under these assumptions is,

$$\tau_{\text{rel}} \frac{\partial f}{\partial \tau}\Big|_{k_z t} + f = f_{eq}(\tau) \tag{C.3}$$

for a constant τ_{rel}, this can be integrated to give

$$\begin{aligned} e^{\tau/\tau_{\text{rel}}} f(k_\perp, k_z, \tau) = e^{\tau_0/\tau_{\text{rel}}} f_0 \left(k_\perp, k_z \frac{\tau}{\tau_0}\right) \\ + \int_{\tau_0}^{\tau} \frac{d\tau'}{\tau_{\text{rel}}} e^{\tau'/\tau_{\text{rel}}} f_{eq} \left(\left[k_\perp^2 + \left(k_z \frac{\tau}{\tau'} \right)^2 \right]^{1/2}, \tau' \right). \end{aligned} \tag{C.4}$$

for a given initial condition $f_0(k_\perp, k_z)$ evaluated at eigentime τ_0.

Taking the second moment of the above, we get the following expression for the time dependent energy density,

$$\begin{aligned}\tau e^{\tau/\tau_{\text{rel}}}\epsilon(\tau) = \tau_0 e^{\tau_0/\tau_{\text{rel}}} \int d\Gamma\, k^2 f_0(\mathbf{k}) \left(1 - \left(1 - \frac{\tau_0^2}{\tau^2}\right)\cos^2(\theta_k)\right) \\ + \int_{\tau_0}^{\tau} \frac{d\tau'}{\tau_{\text{rel}}} \tau' e^{\tau'/\tau_{\text{rel}}} \epsilon(\tau') h\left(\frac{\tau'}{\tau}\right)\end{aligned} \tag{C.5}$$

where,

$$h(x) = \int_0^1 d\mu \sqrt{1 - \mu^2 + \mu^2 x^2} = \frac{1}{2}\left(x + \frac{\arcsin(\sqrt{1-x^2})}{\sqrt{1-x^2}}\right). \tag{C.6}$$

The assumption that the longitudinal momentum is sharply peaked about zero (i.e $f \propto \delta(k_z)$) . We get the simplification $\cos^2(\theta_k) \approx 0$, which allows us to reduce eq (C.5) to,

$$\tau e^{\tau/\tau_{\text{rel}}}\epsilon(\tau) = \tau_0 e^{\tau_0/\tau_{\text{rel}}}\epsilon(\tau_0) + \int_{\tau_0}^{\tau} \frac{d\tau'}{\tau_{\text{rel}}} \tau' e^{\tau'/\tau_{\text{rel}}} \epsilon(\tau') h\left(\frac{\tau'}{\tau}\right). \tag{C.7}$$

We introduce the a new function g defined as follows

$$g_{\tau_0}\left(\frac{\tau}{\tau_{\text{rel}}}\right) = \frac{\tau}{\tau_0}\frac{\epsilon(\tau)}{\epsilon(\tau_0)} \tag{C.8}$$

Writing eq (C.7) in terms of this new function, we get

$$g_{x_0}(x) = e^{-(x-x_0)} + \int_{x_0}^{x} dx' e^{-(x-x')} g_{x_0}(x') h\left(\frac{x'}{x}\right), \tag{C.9}$$

where $x = \frac{\tau}{\tau_{rel}}$ and $x_0 = \frac{\tau_0}{\tau_{rel}}$.

If we able to solve this integro-differential equation, we can extract the evolution of the energy density at mid-rapidity for arbitrary relaxation times.

We approximate the solution to this differential equation by using an iterative numerical solver which was developed in. [48]

Bibliography

[1] F. Gelis, "Color Glass Condensate and Glasma," *Int. J. Mod. Phys.* **A28** (2013) 1330001, arXiv:1211.3327 [hep-ph].

[2] L. McLerran and R. Venugopalan, "Computing quark and gluon distribution functions for very large nuclei," *Phys. Rev. D* **49** (Mar, 1994) 2233–2241. https://link.aps.org/doi/10.1103/PhysRevD.49.2233.

[3] F. Gelis, E. Iancu, J. Jalilian-Marian, and R. Venugopalan, "The Color Glass Condensate," *Ann. Rev. Nucl. Part. Sci.* **60** (2010) 463–489, arXiv:1002.0333 [hep-ph].

[4] B. Schenke, P. Tribedy, and R. Venugopalan, "Fluctuating Glasma initial conditions and flow in heavy ion collisions," *Phys. Rev. Lett.* **108** (2012) 252301, arXiv:1202.6646 [nucl-th].

[5] B. Schenke, P. Tribedy, and R. Venugopalan, "Event-by-event gluon multiplicity, energy density, and eccentricities in ultrarelativistic heavy-ion collisions," *Phys. Rev.* **C86** (2012) 034908, arXiv:1206.6805 [hep-ph].

[6] J. S. Moreland, J. E. Bernhard, and S. A. Bass, "Alternative ansatz to wounded nucleon and binary collision scaling in high-energy nuclear collisions," *Phys. Rev.* **C92** no. 1, (2015) 011901, arXiv:1412.4708 [nucl-th].

[7] W. Florkowski and R. Ryblewski, "Anisotropic hydrodynamics and the early-thermalization puzzle," *AIP Conf. Proc.* **1520** no. 1, (2013) 160–165, arXiv:1207.5951 [nucl-th].

[8] J. E. Bernhard, "Bayesian parameter estimation for relativistic heavy-ion collisions," *arXiv preprint arXiv:1804.06469* (2018) .

[9] G. Policastro, D. T. Son, and A. O. Starinets, "The Shear viscosity of strongly coupled N=4 supersymmetric Yang-Mills plasma," *Phys. Rev. Lett.* **87** (2001) 081601, arXiv:hep-th/0104066 [hep-th].

[10] A. Buchel and J. T. Liu, "Universality of the shear viscosity in supergravity," *Phys. Rev. Lett.* **93** (2004) 090602, arXiv:hep-th/0311175 [hep-th].

[11] H. Song, S. A. Bass, U. Heinz, T. Hirano, and C. Shen, "200 A GeV Au+Au collisions serve a nearly perfect quark-gluon liquid," *Phys. Rev. Lett.* **106** (2011) 192301, arXiv:1011.2783 [nucl-th]. [Erratum: Phys. Rev. Lett.109,139904(2012)].

[12] M. Strickland, "Small system studies: A theory overview," *Nuclear Physics A* **982** (2019) 92 – 98. http://www.sciencedirect.com/science/article/pii/S0375947418302598. The 27th International Conference on Ultrarelativistic Nucleus-Nucleus Collisions: Quark Matter 2018.

[13] M. McNelis, D. Bazow, and U. Heinz, "Viscous hydrodynamics for nonconformal anisotropic fluids," *Nuclear Physics A* **982** (2019) 915 – 918. http://www.sciencedirect.com/science/article/pii/S0375947418302422. The 27th International Conference on Ultrarelativistic Nucleus-Nucleus Collisions: Quark Matter 2018.

[14] F. Cooper and G. Frye, "Comment on the Single Particle Distribution in the Hydrodynamic and Statistical Thermodynamic Models of Multiparticle Production," *Phys. Rev.* **D10** (1974) 186.

[15] H. Petersen, D. Oliinychenko, M. Mayer, J. Staudenmaier, and S. Ryu, "SMASH – A new hadronic transport approach," *Nucl. Phys.* **A982** (2019) 399–402, arXiv:1808.06832 [nucl-th].

[16] M. Bleicher *et al.*, "Relativistic hadron hadron collisions in the ultrarelativistic quantum molecular dynamics model," *J. Phys.* **G25** (1999) 1859–1896, arXiv:hep-ph/9909407 [hep-ph].

[17] M. H. Anderson, J. R. Ensher, M. R. Matthews, C. E. Wieman, and E. A. Cornell, "Observation of Bose-Einstein Condensation in a Dilute Atomic Vapor," *Science* **269** no. 5221, (1995) 198–201. https://science.sciencemag.org/content/269/5221/198.

[18] K. B. Davis, M. O. Mewes, M. R. Andrews, N. J. van Druten, D. S. Durfee, D. M. Kurn, and W. Ketterle, "Bose-Einstein Condensation in a Gas of Sodium Atoms," *Phys. Rev. Lett.* **75** (Nov, 1995) 3969–3973. https://link.aps.org/doi/10.1103/PhysRevLett.75.3969.

[19] J. Klaers, J. Schmitt, F. Vewinger, and M. Weitz, "Bose–Einstein condensation of photons in an optical microcavity," *Nature* **468** no. 7323, (2010) 545.

[20] J.-P. Blaizot, F. Gelis, J.-F. Liao, L. McLerran, and R. Venugopalan, "Bose–Einstein Condensation and Thermalization of the Quark Gluon Plasma," *Nucl. Phys.* **A873** (2012) 68–80, arXiv:1107.5296 [hep-ph].

[21] J.-P. Blaizot, J. Liao, and L. McLerran, "Gluon transport equation in the small angle approximation and the onset of bose–einstein condensation," *Nuclear Physics A* **920** (2013) 58–77.

[22] Z. Chen, C. Greiner, Z. Xu, and P. Zhuang, "Shear viscosity of ultrarelativistic Boson systems in the presence of Bose-Einstein condensate," arXiv:1806.07594 [hep-ph].

[23] J.-P. Blaizot, F. Gelis, J. Liao, L. McLerran, and R. Venugopalan, "Thermalization and Bose-Einstein Condensation in Overpopulated Glasma," *Nuclear Physics A* **904-905** (2013) 829c – 832c. http://www.sciencedirect.com/science/article/pii/S0375947413002704. The Quark Matter 2012.

[24] B. Harrison and A. Peshier, "Bose-Einstein Condensation from the QCD Boltzmann Equation," *Particles* **2** no. 2, (2019) 231–241. http://www.mdpi.com/2571-712X/2/2/16.

[25] Z. Xu, K. Zhou, P. Zhuang, and C. Greiner, "Thermalization of Gluons with Bose-Einstein Condensation," *Phys. Rev. Lett.* **114** (May, 2015)

182301. https://link.aps.org/doi/10.1103/PhysRevLett.114.182301.

[26] F. Scardina, D. Perricone, S. Plumari, M. Ruggieri, and V. Greco, "Relativistic Boltzmann transport approach with Bose-Einstein statistics and the onset of gluon condensation," *Phys. Rev.* **C90** no. 5, (2014) 054904, arXiv:1408.1313 [nucl-th].

[27] J. Anderson and H. Witting, "A relativistic relaxation-time model for the Boltzmann equation," *Physica* **74** no. 3, (1974) 466–488.

[28] L. D. Landau, "On the multiparticle production in high-energy collisions," *Izv. Akad. Nauk Ser. Fiz.* **17** (1953) 51–64.

[29] S. Jeon and U. Heinz, "Introduction to Hydrodynamics," arXiv:1503.03931. http://arxiv.org/abs/1503.03931.

[30] H. F. Trotter, "On the Product of Semi-Groups of Operators," *Proceedings of the American Mathematical Society* **10** no. 4, (1959) 545–551. http://www.jstor.org/stable/2033649.

[31] G. Strang, "On the construction and comparison of difference schemes," *SIAM Journal on Numerical Analysis* **5** no. 3, (1968) 506–517.

[32] L. Rezzolla and O. Zanotti, *Relativistic hydrodynamics*. Oxford University Press, 2013.

[33] F. Filbet, E. Sonnendrücker, and P. Bertrand, "Conservative Numerical Schemes for the Vlasov Equation," *Journal of Computational Physics* **172** no. 1, (2001) 166–187. http://linkinghub.elsevier.com/retrieve/pii/S0021999101968184.

[34] M. Galassi et al, "GNU Scientific Library Reference Manual (3rd Ed.),". http://www.gnu.org/software/gsl/.

[35] J. D. Bjorken, "Highly Relativistic Nucleus-Nucleus Collisions: The Central Rapidity Region," *Phys. Rev.* **D27** (1983) 140–151.

[36] **ALICE** Collaboration, K. Gulbrandsen, "Charged-particle pseudorapidity density and anisotropic flow over a wide pseudorapidity range using ALICE at the LHC," *J. Phys. Conf. Ser.* **446** (2013) 012027, arXiv:1302.0894 [nucl-ex].

[37] B. Kämpfer, O. P. Pavlenko, M. I. Gorenstein, A. Peshier, and G. Soff, "Rapidity dependence of thermal dileptons resulting from hadronizing quark-gluon matter with finite baryon charge," *Zeitschrift für Physik A Hadrons and Nuclei* **353** no. 1, (Mar, 1995) 71–77. https://doi.org/10.1007/BF01297729.

[38] G. Baym, "Thermal equilibration in ultra-relativistic heavy-ion collisions," *Physics Letters B* **138** no. 1-3, (1984) 18–22.

[39] S. Borsanyi, G. Endrodi, Z. Fodor, S. D. Katz, and K. K. Szabo, "Precision SU(3) lattice thermodynamics for a large temperature range," *JHEP* **07** (2012) 056, arXiv:1204.6184 [hep-lat].

[40] L. M. Satarov, I. N. Mishustin, A. V. Merdeev, and H. Stöcker, "Longitudinal fluid dynamics for ultrarelativistic heavy-ion collisions," *Phys. Rev. C* **75** (Feb, 2007) 024903. https://link.aps.org/doi/10.1103/PhysRevC.75.024903.

[41] P. Bożek, "Viscous evolution of the rapidity distribution of matter created in relativistic heavy-ion collisions," *Phys. Rev. C* **77** (Mar, 2008) 034911. https://link.aps.org/doi/10.1103/PhysRevC.77.034911.

[42] NVIDIA Corporation, "NVIDIA CUDA C programming guide version 10.1.105," 2019. https://docs.nvidia.com/cuda/cuda-c-programming-guide/.

[43] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," in *Advances in neural information processing systems*, pp. 1097–1105. 2012.

[44] E. Elsen, V. Vishal, M. Houston, V. Pande, P. Hanrahan, and E. Darve, "N-body simulations on GPUs," *arXiv preprint arXiv:0706.3060* (2007) .

[45] E. E. Schneider and B. E. Robertson, "CHOLLA: a new massively parallel hydrodynamics code for astrophysical simulation," *The Astrophysical Journal Supplement Series* **217** no. 2, (2015) 24.

[46] P. Benítez-Llambay and F. S. Masset, "FARGO3D: a new GPU-oriented MHD code," *The Astrophysical Journal Supplement Series* **223** no. 1, (2016) 11.

[47] D. Bazow, U. Heinz, and M. Strickland, "Massively parallel simulations of relativistic fluid dynamics on graphics processing units with CUDA," *Computer Physics Communications* **225** (2018) 92–113.

[48] M. Steyn, "UCT Honours Thesis: The Boltzmann equation,".