

H24/3679

MONASH UNIVERSITY
THESIS ACCEPTED IN SATISFACTION OF THE
REQUIREMENTS FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

ON..... 11 November 2003

.....
Sec. Research Graduate School Committee

Under the Copyright Act 1968, this thesis must be used only under the normal conditions of scholarly fair dealing for the purposes of research, criticism or review. In particular no results or conclusions should be extracted from it, nor should it be copied or closely paraphrased in whole or in part without the written consent of the author. Proper written acknowledgement should be made for any assistance obtained from this thesis.

ERRATA

- p. 1, 5th line: "nuclear" for "Nuclear"
- p. 2, last line: "distinction" for "dinstinction"
- p. 6, 2nd last line: "and metrics (Section 2.2)" for "and metrics Section 2.2"
- p. 6, last line: "diffusion" for "diffussion"
- p. 9, 4th line: "do" for "does"
- p. 11, last paragraph: " β^i " for " β_i "
- p. 12, 4th line: "coordinate" for "coodinate"
- p. 22, last paragraph: "close" for "enclose"
- p. 23, 2nd last line: "solvers" for "Solvers"
- p. 23, 3rd paragraph: "equation" for "Equation"
- p. 41, second sentence: Remove superfluous "is"
- p. 43, 4th last line: "This" for "THis"
- p. 45, caption to Figure 3.1: space between "sound" and "speed"
- p. 46, 2nd para: "describe" for "descibe"
- p. 54, 7th line: Remove superfluous comma
- p. 54, 2nd para: "relativistic" for "Relativistic"
- p. 57, 9th line: "noticeable" for "noticable"
- p. 64, 2nd para: "Complementing" for "Complimenting"
- p. 70, 2nd last line: space between "under" and "study"
- p. 83, last paragraph: "coordinate" for "coodinate"
- p. 92, section 6.2.3, 1st para: Replace "Taub [1978]" with "[Taub, 1978]"
- p. 93, 5th sentence: Insert missing "a"
- p. 105, 6th last line: Insert "is a" to read "..the other is a three.."
- p. 119, 4th sentence: "open" for "opens"
- p. 119, 5th sentence: "look" for "lok"
- p. 119, 4th sentence: "it has" for "it as"
- p. 121, 1st para: "neutron stars" for "Neutron Stars"
- p. 121, 9th line: "principal" for "principle"
- p. 126, last para: "thermodynamic" for "thermo-dynamic"
- p. 129, 7th line: "condemning" for "comdemning"
- p. 129, 13th line: "predictive" for "predicative"
- p. 130, 1st para: space between "ability" and "of"
- p. 130, 3rd line: "quantitative" for "qualitative"
- p. 155, 10th line: "proper" for "proer"
- p. 159, section 8.6.1, 5th line: space between "of" and "large"

ADDENDUM

- p. 26, Add at start of para 1: "These transformations are of convenience, and involve the kinetic energy, E . The m used is the rest mass value."
- p. 28, Add to the end of first sentence: "Here Γ is the adiabatic index, given by the ratio of specific heats."
- p. 40: Comment: For reference, Siegler and Riffert's equations are as follows:-

$$\begin{aligned} \frac{dq_a}{dt} = & -\sqrt{-g_a} \sum_b \left(\frac{P_a + q_a}{D_a^{*2}} + \frac{P_b + q_b}{D_b^{*2}} \right) \nabla_a W_{ab} \\ & - \frac{\sqrt{-g_a}}{D_a^{*2}} ((p_a + q_a) \nabla_a (\ln \sqrt{-g_a})_a - \frac{1}{2} T_a^{\alpha\beta} \nabla_a (g_{\alpha\beta})_a) \end{aligned}$$

and

$$\begin{aligned} \frac{d\alpha E - \beta^i q_i}{dt} = & -\frac{\sqrt{-g_a}}{2} \sum_b \left(\frac{P_a + q_a}{D_a^{*2}} + \frac{P_b + q_b}{D_b^{*2}} \right) (\mathbf{v}_a + \mathbf{v}_b) \cdot \nabla_a W_{ab} \\ & - \frac{\sqrt{-g_a}}{D_a^{*2}} ((p_a + q_a) \mathbf{v}_a \cdot \nabla_a (\ln \sqrt{-g_a})_a + \frac{1}{2} T_a^{\alpha\beta} (g_{\alpha\beta,t})_a) \end{aligned}$$

- p. 46: Comment: The \mathbf{v} refers to the velocity of the moving frame K' , and is appended to γ as a subscript to specify that this factor corresponds to the frame's motion, not that of a particle.
- p. 55, first para: Comment: It is erroneous to say that no new errors are being introduced. However, the RNC analysis shows the the order of the error introduced by this modification, is of the same order as the errors introduced by the SPH interpolation algorithm and the leap-frog integration algorithm. Therefore the *order* of the errors remains unchanged.
- p. 79, 2nd para: Delete " , and how important it is" and read "It is important.."
- p. 139, 2nd para: Add to the end of the first sentence: " , where the impact parameter is defined to be the ratio between the separation of the nuclei centres and the sum of their radii.

A Relativistic, 3-Dimensional Smoothed Particle Hydrodynamics (SPH) Algorithm and its Applications

A Thesis submitted for Degree of
Doctor of Philosophy

by

Stuart Muir
B.Sc. (Hons)

School of Mathematical Sciences,
Monash University
May 2003

Contents

Quotes	iv
Summary	v
Statement	vii
Acknowledgements	viii
1 Introduction	1
2 Relativistic Hydrodynamical Theory	6
2.1 Describing Motions and the Concept of Frames	7
2.2 Metrics and Curved Space-Time	8
2.2.1 3+1 Formalism	11
2.3 Riemann Normal Coordinates	12
2.4 Conservation Laws	16
2.5 Comparison to Riemann Solvers	23
2.6 The Equation of State	25
2.6.1 RHIC EoS	32
3 Discretizations and SPH Methodology	35
3.1 Interpolation Theory	36
3.2 Discrete, Non-Dissipative, Lagrangian Equations of Motion . .	38
3.2.1 Conservation Properties	40
3.3 Dissipation terms and their derivation	42
3.4 Choice and Derivation of Signal Speeds	44
4 Numerics and Code Specifics	49
4.1 SPH: A Brief Overview	49
4.2 Choice of Splines	50
4.3 Smoothing Length Modification	58
4.4 Associated Errors and Stability of the Algorithm	64
4.5 Effects of Omitting ∇h Terms	66

5	Application Specifics	67
5.1	Boundary Conditions	67
5.2	Initial Conditions	73
5.3	Timestepping Conditions	74
5.4	Solution for Pressure using a Non-Linear Equation	76
5.5	Neighbour Searching and Rank Space	78
5.5.1	Ranklisting	80
6	Preliminary (Shocktube) Testing	84
6.1	The Shock Tube as a Hydrodynamical System	84
6.2	Exact Solution	87
6.2.1	The Rarefaction	88
6.2.2	The Intermediate States	90
6.2.3	The Shock Front	92
6.3	Initial Conditions	93
6.4	SPH Performance and Results	95
6.4.1	Diffusion Parameter Choices	96
6.4.2	Sod's Shock	98
6.5	Relativistic Shock Wave	104
6.5.1	3-D Shocks	105
6.5.2	Effects of the Number of Neighbours	113
7	Relativistic Heavy Ion Collisions (RHIC)	119
7.1	An Historic Overview	120
7.2	Phases of a Nuclear Collision	122
7.3	Modelling a Nuclear Collision	124
7.4	The Equation of State as an Hydrodynamical Input	127
7.5	The Hydrodynamic Phenomena of RHIC	128
7.6	Advantages of Using a Particle Method for RHIC Modelling	132
7.7	Initial Conditions	134
7.8	Data Visualisation	138
7.9	Performance and Results	139
7.9.1	Neon : Uranium Collisions	139
7.9.2	Au: Au collisions	143
8	Performance on a Curved Metric	146
8.1	Motivation	147
8.2	The Schwarzschild and Kerr Spacetimes	148
8.3	Exact Solution	152
8.4	Generating the Exact Solution	155

8.5	The Physics of the Schwarzschild Infall Problem	156
8.6	Initial Conditions	157
8.6.1	A Note on Time-Stepping	159
8.7	Performance	159
9	Conclusions and Further Applications	166
	Bibliography	170
	Quotes	177

'Science is the topography of ignorance,' Holmes

'If we knew what we were doing, it wouldn't be called research, would it?,' Einstein

'Basic research is what I am doing when I don't know what I am doing,' Wernher von Braun

Summary

Einstein's Theory of Relativity predicts the effects of high speed motion or being in the vicinity of a massive object. Such effects create difficulties for computer models. This thesis explores the design and application of an Smoothed Particle Hydrodynamics (SPH) algorithm suitable for studying relativistic gas dynamics.

In this thesis, we construct an algorithm capable of solving the equations of gas motions within the framework of Special Relativity. We produce an algorithm which specifically avoids reliance on exact solutions and 'tuning' or parameters. The basic code is designed to be modular, allowing for extra physics and routines to be added easily. This allows us to have subroutines which change the initial conditions, boundary conditions and equations of state, which we insert into the main loop and calculate a variety of simulations using the one general algorithm. We generate artificial viscosity terms and signal velocities which prevent acausal information travel.

As the algorithm works fundamentally in three spatial dimensions, we also look at ways of increasing the efficiency of memory storage, and locating nearest neighbours. This is done through carrying out much of the calculations in integer rankspace. This modification will also lend itself to future parallelisation of the algorithm.


These concepts are then tested within the framework of 1D shocktubes to

ensure the hydrodynamics is correct. The algorithm is then applied to the quantitative study of Heavy Ion collisions with the use of a nuclear equation of state and differing boundary conditions.

We finally apply the code to perform calculations in the vicinity of a black hole, taking the effects of curved space-time into account

Statement

This thesis contains no material which has been accepted for the award of any other degree or diploma in any University or other Institution. To the best of my knowledge, this Thesis contains no material previously published or written by another person except where due reference is made in the text.



Stuart Muir

Acknowledgements

First I wish to thank my supervisor, Joe Monaghan for his guidance and support during my candidature. Thanks are also due to John Lattanzio, Leo Brewin and Brian Schmidt for many conversations, thoughts and advice through the tough bits.

Thanks also go to Libby and Ray, for their careful proof reading, comments, suggestions and general forbearance.

For their support over the years, thanks go to my family, and especially to my brother for enduring animated discussions with barely a grimace. In this role, thanks also to Susan. Without a doubt, no-one has supported and understood more than Libby, so 'Ta, Bubs.'

For the contribution which they never knew they made, thanks to Jeff, Fi, Bob and Rosco, for constantly reminding me of the other world that exists out there somewhere.

And finally to Sydney, for consistantly pointing out that, 'Particles, whether they be computational, nuclear or compact object, are only ever merely approximations to the most important particles in the Universe:- Tennis Balls.'

Chapter 1

Introduction

The purpose of this thesis is to develop a hydrodynamical code capable of resolving wave phenomena and other fluid motions within the framework of Relativistic Theory. The applications of such an algorithm encompass fluid dynamical studies of environments far beyond those of human experience, and range from Nuclear physics through to astrophysical studies of the regions surrounding black holes. With such a diverse range of applications, we desire that the algorithm be modular, allowing different physics and equations to be exchanged readily, and simple, avoiding the use of problem specific solution methods. This last point also encompasses the art of parameter tuning to resolve a specific feature of the solution.

The eventual aim of such a project is to develop the ability to evolve a given initial state forward in time. The development of such a program is envisaged as following the basic historical development of gravitational N-body programs. The algorithm here will be able to evolve itself by responding to the curvature of the space-time it finds itself in, just as early N-body codes responded to a background gravitational potential. Eventually, it is hoped that there will exist codes which, given a matter distribution (such as that output

from this algorithm), will be able to deduce the curvature of the space-time. This will release us from the current restriction of analytic, vacuum space-time solutions. With this in mind, where possible, we have incorporated notations used by numerical relativists, hoping to ease the coupling of the two codes at some future time.

Hydrodynamical modelling itself is the process of approximating a coherent block of matter by a fluid which responds to known equations of motion. Whilst sounding initially limiting, it is a very powerful technique, allowing simulations of scenarios as diverse as ship wakes and aeronautics through to planetary mantle convection and rock fracture.

The numerical solution of the equations of fluid motion can be done in either a Eulerian, or a Lagrangian framework. Eulerian schemes typically observe the fluid moving from a distance. Placing a computational grid over the fluid motion means that fluid flow between cells can be monitored by fluxes across the cell's faces. Eulerian schemes therefore have a clear advantage in their ability to model difficult boundaries and single fluid motions. They are however limited in their resolution of effects smaller than the grid scale, and so respond poorly to free boundaries or interfaces, particularly where there is reconnection (cresting wave) or where there are important small-scale effects such as turbulent motions.

Lagrangian schemes on the other hand observe the fluid from within, using co-moving grids or particles from which to deduce the motions and hydrodynamic variables. Lagrangian schemes then handle free boundaries much better than their Eulerian counterparts, but fixed boundaries are a much more difficult proposition. Computational nodes free to move with a fluid are also able to resolve smaller details without the expense of refining the grid scale, and can re-connect the fluid flow easily. Perhaps the clearest distinction

between these two schools of thought however, is the readiness of Lagrangian particle methods to move to higher dimensions.

A single spatial dimension Eulerian code must look at the way in which fluid moves between cells in 1D. The analogous particle code looks at how two particles interact upon radial orbits. Moving up to two and three dimensions, Eulerian codes need to 'understand' the fluid's behaviour as it crosses cell's faces (which can reduce to a 1D application as above), the more difficult cell edge interactions, and the most difficult, the fluid's behaviour across cell vertices. Particle codes on the other hand can always reduce to single particle-particle interactions, with no additional complexity other than freedom of motion being restricted to radial, planar or 3-dimensional orbits. Because of this simplification, 3-dimensional particle codes are relatively simple to write. They do however, suffer from a drawback.

Particle methods represent the fluid as a series of discrete, computational nodes. Physically, two colliding regions of fluid will heat, and not penetrate. If however, these two packets of fluid are represented by two particles, free to move in three-space, then the initial colliding orbits are easily deflected, resulting in the particles passing each other. This would represent large fluid interpenetration, something not seen physically, and not desirable computationally. To combat this, Lagrangian particle methods rely on artificial viscosities. One needs to be careful in the application of such numerical diffusion, as it is easy to introduce unphysical results and communications. We will be looking into this problem in the thesis.

Despite this difficulty, it has been decided that a Lagrangian particle method gives the most general algorithm, capable of resolving unknown hydrodynamical effects. The methodology used will be that of Smoothed Particle Hydrodynamics, a numerical method devised independantly by Lucy

[1977] and Gingold and Monaghan [1977], and presented in detail within the work.

It is argued that once one knows what the basic gas behaviour is for a given system, one can either choose a more appropriate or fine-tuned formulation, or even a different numerical method, chosen specifically for its ability to resolve or capture the noteworthy phenomena.

As the background spacetime upon which we will derive our equations of motion is presumed to be static, the mass of the actual hydrodynamical system must remain much less than that of the dominant, curvature inducing mass energy. This is not a great restriction in itself, but it does mean that there is no particle self gravity. To add this in a meaningful relativistic calculation is far and above beyond the scope of such a work as this. But in abiding by the assumption of keeping the hydrodynamic mass-energy small (relative to the source mass-energy) a Newtonian self-gravity for the particles can be added at a later date, and would be a worthwhile addition.

A condition is also placed upon the types of background space-time within which we can perform hydrodynamical calculations. The current form of our derived equations assumes that the shift vector used to describe subsequent hypersurface translations is zero. Developing the equations further to incorporate a non-zero shift vector would open the work again to more physical space-times.

It should also be noted that this is not the first work attempting to couple the SPH algorithm to a curved space metric. Laguna et al. [1993] introduced a similar method, and used it to model 1D shocks, gas infall and Bondi collapse problems. Siegler and Riffert [1999] have also presented similar equations. Both these works however integrate the thermal energy equation, and have awkward terms such as time derivatives of the relativistic contraction factor,

γ , or non-symmetries in their equations. An excellent work by Oechslin et al. [2001] solves both the relativistic field equations and the hydrodynamics on a curved background. It is however limited to conformal metrics, and loses the grid-free nature of SPH by applying a gravity grid.

Throughout this paper we will use the convention of superscripts and subscripts where Greek indices, μ and ν denote full four-vector components $(0, \dots, 3)$ and Latin indices (i, j, k) denote spatial components only $(1, 2, 3)$. The covariant derivative with respect to the coordinate x^μ is shown by the subscript $;\mu$, and the subscripts a and b are reserved to denote quantities associated with particle a and b respectively. A 0 subscript is used to denote the proper rest-frame quantity. The gas calculations are done in a Cartesian coordinate frame specified by $x^\mu = [t, x, y, z]$, however we attempt to generate the equations in a general form, applicable to an arbitrary choice of coordinates.

Chapter 2

Relativistic Hydrodynamical Theory

In modelling any motions where the speeds approach significant fractions of the speed of light we must rely upon relativistic theory, which tells us how space and time are intrinsically coupled. As a consequence, the well known and understood equations of hydrodynamics become increasingly complex, even in the simplest case of flat spacetime. To calculate motions near an object massive enough to dominate the curvature, it becomes necessary to include local curvature effects, accomplished through the metric terms. If, however, one wants to calculate the evolution of a changing spacetime induced by the redistribution of the matter, then the full Einstein equations must be solved, a problem of significant complexity which will not be attempted here.

We begin in Section 2.1 by establishing the relativistic concepts of space and time such as coordinate frames (Section 2.1 and Section 2.3) and metrics Section 2.2, with which we can develop the diffusionless equations of motion. Section 2.4 deals with some attempts to handle diffusion, and what one

expects the terms to look like. We also introduce the method of Riemann Solvers to solve for the relativistic equations of fluid motion. Examination of the theory behind these methods allows us to derive the diffusive term, Λ . (Section 2.5) The following sections (Section 2.6) examine the equations of state to be used.

2.1 Describing Motions and the Concept of Frames

Newtonian theories are built around the concept of instantaneous information travel, a perfectly valid approximation if the notional velocities are much less than that of light and distances much less than cosmological. When dealing with relativistic speeds, one needs to be more careful with regards to terms involving time and simultaneity, which can be rendered ambiguous, if not meaningless, in this context.

The notion of frames is used to distinguish where a measurement was taken. The laboratory frame is an inertial frame, within which a fluid is moving. The other important frame required is a co-moving, or proper frame. This inertial frame is instantaneously co-moving with the fluid element of interest, and leads to the concept of proper time, τ , or the time the fluid element measures.

With proper time as an obvious parameter, we can define the particle's path to be $x(\tau) = \{t = t(\tau), x = x(\tau), y = y(\tau), z = z(\tau)\}$. This leads to the definition of the fluid element's 4-velocity

$$U^\mu = \frac{dx^\mu(\tau)}{d\tau} \quad (2.1)$$

which is constrained by the normalisation condition (in geometric coordinates

where velocities are scaled to c)

$$U^\mu U_\mu = -1 \quad (2.2)$$

In the laboratory frame, in which the fluid element is moving, we have the relation $d\tau = \frac{1}{\gamma} dt$, where γ is the usual 'Lorentz contraction factor' of Special Relativity, given by

$$\gamma = \frac{1}{\sqrt{1 - v^i v_i}} \quad (2.3)$$

Here we see the use of v^μ which is the velocity of the co-moving frame with respect to the laboratory, and is referred to as the 'transport velocity,' or 'fluid velocity', and is related to U^μ via

$$U^\mu = \gamma v^\mu \quad (2.4)$$

A relativistic fluid consisting entirely of baryons (a perfect one-fluid) is fully described by specifying its baryon number density, ρ , its momentum flux density, Q , its energy density, E , internal or thermal energy, ϵ , and the isotropic pressure, P , which together give us a natural choice for computational variables.

2.2 Metrics and Curved Space-Time

Just as important as knowing where a given measurement was taken and the details of the observer who took it, is understanding the spacetime of that event. This information is contained within the metric $g_{\mu\nu}$.

As previously mentioned, if one wishes to calculate how the changes in

mass distribution affect the curvature, then the components of the metric become evolving variables, whose evolution is dictated by Einstein's field equations. If, however, one assumes that the hydrodynamics and particles' distribution does not affect the metric, that is the motion or hydrodynamics occur on a background spacetime, then the fluid equations of evolution will reduce to familiar forms, analogous with Newtonian theory, albeit with the addition of extra source functions due to the curvature.

For this assumption to retain an acceptable level of accuracy, i.e. for the curvature to remain static and stationary, the total mass-energy of the hydrodynamical system under study must be sufficiently small relative to the mass-energy producing the background curvature.

For early test problems, and then in Heavy Ion collisions, the effects of gravity are minimal compared with the body forces involved in the reaction. Therefore we can neglect curvature terms completely and use the flat-space, Cartesian metric of Minkowski's space-time, given by

$$g_{\mu\nu} = \text{diag}\{-1, 1, 1, 1\} \quad (2.5)$$

The second situation where the change in $g_{\mu\nu}$ can be neglected is regions of space where the influence of the mass of the fluid body is much less than that of some dominant, background mass. This could be as simple as a small body under the influence of the Earth's field, or more interestingly, the effects on stellar mass objects trapped in the gravitational well of a supermassive black hole, such as those believed to be in the centre of most galaxies.

In this work, we will use two curved space-time metrics other than that prescribed by Minkowski. The first is the well known Schwarzschild metric

of an isolated, static singularity, given by

$$ds^2 = -(1 - \frac{2M}{r})dt^2 + (1 + \frac{2M}{r})dr^2 + r^2 d\theta^2 + r^2 \sin^2(\theta) d\phi^2 \quad (2.6)$$

The second form is that of the isotropic Schwarzschild space-time, where the line element can be described by

$$ds^2 = -(\frac{1 - \frac{M}{2\bar{r}}}{1 + \frac{M}{2\bar{r}}})^2 dt^2 + (1 + \frac{M}{2\bar{r}})^4 \{d\Omega^2\} \quad (2.7)$$

and using $d\Omega^2$ as an appropriate flat 3-space metric and

$$r = \bar{r}(1 + \frac{M}{2\bar{r}})^2 \quad (2.8)$$

A third possible description of this space-time is provided by Hawking and Ellis [1973] through their description of a stationary, rotating singularity. It is given in Cartesian, Kerr-Schild coordinates.

$$ds^2 = dx^2 + dy^2 + dz^2 - dt^2 + \frac{2mr^3}{r^4 + a^2z^2} \left\{ \frac{r(xdx + ydy) - a(xdy - ydx)}{r^2 + a^2} + \frac{zdz}{r} + dt \right\}^2 \quad (2.9)$$

where a is the specific angular momentum per unit mass of the central object (black hole) and r can be determined (to a sign) from the relation

$$r^4 + (x^2 + y^2 + z^2 - a^2)r^2 - a^2z^2 = 0 \quad (2.10)$$

This metric can be seen to reduce to a Schwarzschild spacetime (although in Kerr-Schild coordinates) when $a = 0$. This metric is different to the others

shown as it has a shift vector (see Section 2.2.1) and so is unsuitable currently, but is of interest for further investigation and application.

2.2.1 3+1 Formalism

In an effort to reduce the computational labour of 4-dimensional calculations, it is possible to split a given spacetime into a series of 3-manifolds, described by their associated spatial metrics, layered along the time axis. This is the premise behind Arnowitt, Deser & Misner's 3+1 formulation [Witten, 1962], of which a full and detailed explanation is above and beyond this work. Although much of the application is above the realm of this work, it is important to be aware of its implications and notation. The 3+1 formalism, or some variation of it, is used almost exclusively by numerical relativists. Not only can we align this study with metric studies, but using this notation introduces a number of terms which help clarify the equations. These are the lapse function, referred to as α , and the shift vector, β_i . In this notation, an arbitrary 3+1 spacetime is given by the line element:

$$ds^2 = -(\alpha + \beta^i \beta_i) dt^2 + \beta_i \beta_j dx^i dx^j + \eta_{ij} dx^i dx^j, \quad (2.11)$$

where η_{ij} is the spatial metric.

These (as shown in Figure 2.1) are effectively the distance between the hypersurfaces along the time axis (αn^μ), and how the surfaces are stacked relative to each other (β_i). They are necessary for an evolution in time of the hypersurfaces, as they tell exactly how the time derivative, and hence time-integral, is formulated. Whilst not purely necessary at this level, eventually, hydrodynamical codes will be required to function on numerically derived metrics which most likely will be 3+1 notations.




FIGURE 2.1: 3 + 1 space-time, two spacelike 3-manifolds shown

2.3 Riemann Normal Coordinates

In later sections the local region of a particle in curved space will become an important issue, particularly when we look at how the numerical method will interpolate thermodynamic variables. In a curved space-time, it is impossible to define a set of coordinates where the connections, $\Gamma_{\beta\alpha}^\alpha$, are zero everywhere for all time. That is, for an arbitrary space-time, it is not possible to find a coordinate transformation to recover Minkowski (flat) space everywhere. It is not impossible though to construct a local inertial reference frame, centred on a particle O , where local, free particles move along straight lines. These are the local Riemann Normal Coordinates. For a full description of their derivation and uses, the reader is advised to turn to Misner et al. [1973a] or Brewin [1997], available from

http://newton.maths.monash.edu.au:8000/papers/rnc_notes.ps.gz The basic principles relating to how we will use these frames are outlined in more detail below.

If we take a particle at a given location in our space-time, designated O as before, and envision all the geodesics of the space-time that go through that point, we will have O at the focus of a radial explosion of vectors (Figure 2.2). If we then look at a nearby particle, P , and provided it is not far away enough that geodesics can intercept, then we can uniquely specify P 's location as

$$x_P^\mu = \lambda a^\mu, \quad (2.12)$$

where λ is the geodesic distance, and a^μ are the components of the tangent vector uniquely specifying the geodesic.

Working in direct analogy with Brewin's work, we can define x_P^μ vector as a function of the coordinates of O , x_O^μ and a new coordinate system y^μ to give

$$x_P^\mu = x_O^\mu + \delta y^\mu \quad (2.13)$$

Here δ is the length scale of the coordinate patch where our local RNC will be defined. The line element will then be

$$\begin{aligned} ds^2 &= g_{\mu\nu}(x) dx^\mu dx^\nu \\ &= \delta^2 g_{\mu\nu}(x_O + \delta y) dy^\mu dy^\nu \end{aligned} \quad (2.14)$$

This is most naturally remapped to a conformal metric, denoted by a \sim ,




FIGURE 2.2: Geodesics radiating outwards from the point O to all locations on the patch

$$\begin{aligned}\tilde{ds}^2 &= g_{\mu\nu}(x_O + \delta y) dy^\mu dy^\nu \\ &= \tilde{g}_{\mu\nu}(y, \delta) dy^\mu dy^\nu\end{aligned}\tag{2.15}$$

Knowing that $\Gamma_{\beta\alpha}^\alpha$ is zero at the origin of this new coordinate system, and by simple examination of the above Equations 2.14 and 2.15, the following deductions can be made, noting that the conformal, tilde metric is in terms of the y^μ coordinates, and the normal metric is in terms of x^μ




FIGURE 2.3: The Riemann Normal Coordinate patch, depicted as a plane, centred on the location O on a curved space manifold (2D analogy)

$$\tilde{g}_{\mu\nu} = g_{\mu\nu} \quad (2.16)$$

$$\tilde{g}_{\mu\nu,\alpha} = \delta g_{\mu\nu,\alpha} \quad (2.17)$$

$$\tilde{g}_{\mu\nu,\alpha\beta} = \delta^2 g_{\mu\nu,\alpha\beta} \quad (2.18)$$

$$\tilde{R}_{\alpha\beta\mu\nu} = \delta^2 R_{\alpha\beta\mu\nu} \quad (2.19)$$

It should be noted that $R_{\alpha\beta\mu\nu}$ is independent of the proportions of the coordinate patch, leaving

$$\tilde{R}_{\alpha\beta\mu\nu} = O(\delta^2) \quad (2.20)$$

Therefore, to order δ^2 , the Riemann Normal Coordinates represent flat space (Figure 2.3).

2.4 Conservation Laws

The relativistic analogues to the familiar conservation laws from Newtonian hydrodynamics can be generated by enforcing that the (4-)divergence of the stress-energy tensor, $T^{\mu\nu}$, is null. This gives us the four equations

$$T^{\mu\nu}{}_{;\nu} = 0 \quad (2.21)$$

The choice made for $T^{\mu\nu}$ should be dependent only upon the fluid being modelled. The most complete stress-energy tensor available is probably that of the so called 'first-order' theories of Eckart [1940] and Landau and Lifshitz [1959], which are neatly expressed in the notation of Misner et al. [1973b] as

$$T^{\mu\nu} = (\rho_0 + \rho_0\epsilon + P - \zeta\theta)U^\mu U^\nu + (P - \zeta\theta)g^{\mu\nu} - 2\xi\sigma^{\mu\nu} + q^\mu U^\nu + q^\nu U^\mu \quad (2.22)$$

where ζ and ξ are the bulk and shear viscosity coefficients, q^μ is the heat flux, θ is the expansion of the fluid, given by the divergence of the 4-velocity, $U^\mu{}_{;\mu}$, and $\sigma^{\mu\nu} = \frac{1}{2}(U^\mu{}_{;\alpha}(g^\alpha{}_\nu + U_\alpha U^\nu) + U^\nu{}_{;\alpha}(g^\alpha{}_\mu + U_\alpha U^\mu)) - \frac{1}{3}\theta(g^{\mu\nu} + U^\mu U^\nu)$ is the shear tensor. The remaining terms (P , ρ_0 , and ϵ) are the hydrodynamic variables of isotropic pressure, rest baryon number density and internal energy respectively.

As with their Newtonian counterparts, these dissipation forms give rise to parabolic equations of heat flow, with their associated infinite speeds of propagation. This is unsettling in non-relativistic scenarios, but its inherent violation of the principles of causality make it completely unacceptable for a relativistic application.

Hiscock and Lindblom [1985] generalised these theories, and showed that not only are they unstable to minor perturbations, but the forms of Eckart

and Landau and Lifshitz do not even permit stable equilibria. These findings are in agreement with the works of Israel [1976], who attributes these effects to the arguably arbitrary fashion in which quadratic and higher order terms (of deviations from the equilibrium values) of the heat flux and viscosity are omitted from the entropy definition. Hence the title 'first-order' theories which is attributed to these formulations. Israel goes on to formulate a second order theory, where the quadratic terms are restored [Israel and Stewart, 1979] and which permits stable equilibria, and causal propagation of linear perturbations.

Unfortunately for this work, the Israel-Stewart formalism produces many terms of second order in the evolution equations, greatly complicating the system. The coefficients of these terms are unrecoverable from the equation of state [Olson and Hiscock, 1990], requiring experimental deduction, and so are unsuitable for these applications.

In light of these complications, we remind ourselves of the heavy reliance in numerical hydrodynamics of artificial viscosities to stabilise calculations, and focus our attentions on fluids where many of these viscous terms can be ignored.

Accordingly, in the interests of simplicity and clarity, the conductive terms and dissipation terms are removed from the tensor. Artificial dissipation will be added at a later stage in the guise of an additional pressure term to help resolve shocks, and its form is guided by the terms above and by comparisons with analytical Riemann solvers.

The stress-energy tensor then becomes

$$\begin{aligned} T^{\mu\nu} &= (\rho_0 + \rho_0\epsilon + P)U^\mu U^\nu + Pg^{\mu\nu} \\ &= \rho_0 h U^\mu U^\nu + Pg^{\mu\nu} \end{aligned} \tag{2.23}$$

where $g^{\mu\nu}$ are the contravariant metric components and

$$h = 1 + \epsilon + \frac{P}{\rho_0} \quad (2.24)$$

is the relativistic enthalpy, which is the same as the Newtonian enthalpy, with an additional term taking into account the rest mass-energy. Natural, or geometric, units ($c = 1$) have been used.

It is interesting to consider these components in the laboratory frame (Minkowski flat space), where the relation between the rest mass density ρ_0 and that measured is given by

$$\begin{aligned} \rho &= \frac{\rho_0}{\sqrt{1-v^2}} \\ &= \gamma \rho_0, \end{aligned} \quad (2.25)$$

and γ is the Lorentz contraction factor.

The stress-energy tensor is then given by

$$\begin{aligned} T^{00} &= E \\ &= \rho_0 h \gamma^2 - P \\ &= \gamma^2 (\rho_0 + \rho_0 \epsilon + P) - P \\ &= \gamma (\rho + \rho \epsilon + \gamma P) - P \end{aligned} \quad (2.26)$$

and

$$\begin{aligned} T^{0i} &= \rho_0 h \gamma^2 v^i \\ &= \rho h \gamma v^i \\ &= Q^i \end{aligned} \quad (2.27)$$

giving

$$T^{\mu\nu} = \begin{bmatrix} E & Q^1 & Q^2 & Q^3 \\ Q^1 & Q^1 u^1 + P & Q^1 u^2 & Q^1 u^3 \\ Q^2 & Q^2 u^1 & Q^2 u^2 + P & Q^2 u^3 \\ Q^3 & Q^3 u^1 & Q^3 u^2 & Q^3 u^3 + P \end{bmatrix}$$

Accordingly, we generate the Conservation laws as follows:

$$\begin{aligned} 0 &= \partial_\beta T^{0\beta} \\ &= \partial_t[(\rho_0 + \rho_0 \epsilon + P)\gamma^2 - P] + \partial_i[(\rho_0 + \rho_0 \epsilon + P)\gamma^2 v^i] \end{aligned}$$

which leads us to the Conservation of Energy Equation

$$\partial_t E + \partial_i(E v^i) = -\partial_i(P v^i) \quad (2.28)$$

where

$$\begin{aligned} E &= \text{Relativistic Energy Density (rest mass included)} \\ &= T^{00} \\ &= (\rho_0 + \rho_0 \epsilon + P)\gamma^2 - P \end{aligned} \quad (2.29)$$

Similarly for Momentum, we get:

$$\begin{aligned} 0 &= \partial_\beta T^{i\beta} \\ &= \partial_t[(\rho_0 + \rho_0 \epsilon + P)\gamma v^i - P] + \partial_j[(\rho_0 + \rho_0 \epsilon + P)\gamma^2 v^i v^j] \end{aligned}$$

and so Conservation of Momentum is given by,

$$\partial_t Q^i + \partial_j(Q^i v^j) = -\partial_i(P) \quad (2.30)$$

where

$$\begin{aligned}
 Q^i &= \text{Momentum density in the } i^{\text{th}} \text{ direction} \\
 &= T^{0i} \\
 &= (\rho_0 + \rho_0 \epsilon + P) \gamma v^i
 \end{aligned} \tag{2.31}$$

These variables are best scaled to give the conserved variables,

$$e = \frac{E}{\rho} : \text{specific energy/unit baryon number} \tag{2.32}$$

$$q^i = \frac{Q^i}{\rho} : \text{specific Momentum} \tag{2.33}$$

The Conservation Equations (2.28) and (2.30) are further expressed in their Lagrangian form as

$$\begin{aligned}
 \frac{de}{dt} &= \partial_i e + v^i \partial_i e \\
 &= -\frac{1}{\rho} \partial_i (P v^i)
 \end{aligned} \tag{2.34}$$

$$\begin{aligned}
 \frac{dq^i}{dt} &= \partial_i q^i + v^j \partial_j q^i \\
 &= -\frac{1}{\rho} \partial_i P
 \end{aligned} \tag{2.35}$$

with

$$e = (1 + \epsilon + \frac{P}{\rho})\gamma - \frac{P}{\gamma\rho_0} \quad (2.36)$$

$$q^i = (1 + \epsilon + \frac{P}{\rho_0})\gamma v^i \quad (2.37)$$

In using these definitions we have implicitly assumed no variation in chemical potential across the fluid. This restriction precludes both any difference in constituent fluid, and the possibility of creating or destroying particles. For simulations involving molecular fluids, these particles would be the products of any chemical reactions. In the context of a baryonic fluid, no new baryons may be formed from an energy reaction. At first glance, this may appear a perfectly satisfying restriction. However, many Heavy Ion Colliders used in nuclear structure studies are capable of energies high enough to create new particles. Through these energies they are capable of exploring the nuclear equation of state in its higher energy bounds, and can probe properties of quark/gluon plasmas. Whilst there is no reason that the methods described here cannot be (easily) modified to handle these equations of state, the modifications required would demand that the general nature that we wish this algorithm to retain would be lost.

When it comes to deriving the equivalent continuity equation, one must note that while particles themselves are still conserved, their mass, (or at least their perceived mass) changes as a function of velocity. That is, if m_0 is the rest-mass of a given fluid element, then its perceived mass in the laboratory frame is given by γm_0 . So it is not true, as it is in Newtonian hydrodynamics, to say that the net mass flux is zero. (One must instead note the mass-energy flux is zero.) It is to avoid this possible misconception, that we work instead in number density ρ (as opposed to mass density), and specify the number

flux must be zero, or

$$(\rho_0 U^\mu)_{;\mu} = 0 \quad (2.38)$$

The continuity equation expressed in Eulerian form becomes

$$\partial_t \rho + \partial_i (\rho v^i) = 0 \quad (2.39)$$

or, as we will use, the Lagrangian form of

$$\frac{d\rho}{dt} = -\rho \nabla_i v^i \quad (2.40)$$

Having established the evolution of the system this way, it is natural to describe the conserved variables in a vector of unknowns,

$$\mathbf{F} = \{\rho, \mathbf{Q}, E\} \quad (2.41)$$

and its associated fluxes

$$\Phi^i = \{\rho v^i, Q^1 v^i + P \delta^{1i}, Q^2 v^i + P \delta^{2i}, Q^3 v^i + P \delta^{3i}, Q^i - \rho v^i\} \quad (2.42)$$

The system of equations (2.2), (2.21) and (2.38) require one more relation to enclose the system. This is the Equation of State, relating the pressure to the other hydrodynamical variables and generally of the form $P = P(\rho, \epsilon)$. This will be looked at in more detail in a later section (Section 2.6).

2.5 Comparison to Riemann Solvers

In recent times, there has been a number of attempts to produce High Resolution Shock Capturing methods capable of handling relativistic flows. These methods centre around conservative forms of the fluid equations, and calculation of numerical fluxes, usually through a Riemann Solver. Marti and Müller [1994] have deduced the analytical solutions for the one-dimensional problems. In higher dimensions, these exact solutions become computationally prohibitive, so many resort now to approximate solvers, or attempt to reduce the problem to one of lesser dimensions.

If we look at our vectors (2.42) and (2.41) the conservation laws reduce to

$$\frac{\partial \mathbf{F}}{\partial t} + \sum_{i=1}^3 \frac{\partial \Phi^i}{\partial x^i} = 0 \quad (2.43)$$

Provided that the equation of state used to close this system is causal (i.e. the sound speed is limited by c) then this system is hyperbolic [Anile, 1989]. Hyperbolic systems allow real valued eigenvalues, and a non-degenerate set of eigenvectors, for the Jacobian matrix, $\frac{\partial \Phi^i(\mathbf{F})}{\partial F^j}$. Provided the system is known at some time, these characteristics can be used to evolve the system forward in time, often using simple integration techniques.

If we take the above Equation (2.43) and naively integrate it forward in time from t^n to $t^{n+1} = t^n + \Delta t$, we get the straight forward equation

$$\mathbf{F}_j^{n+1} = \mathbf{F}_j^n - \frac{\Delta t}{\Delta x} \Phi \quad (2.44)$$

where the Φ represent some approximation to the flux. Here we can turn to Roe-type linearised Riemann Solvers, as detailed in Marti and Müller [1999]

If \mathbf{B} is the Jacobian $\frac{\partial \Phi^i(\mathbf{F})}{\partial F^j}$, then the locally linear form, $\tilde{\mathbf{B}}$, will admit d

eigenvalues $\tilde{\lambda}$, d left eigenvectors $\tilde{\mathbf{e}}_l$, and d right eigenvectors $\tilde{\mathbf{e}}_r$, where d is the number of dimensions.

Given this, the flux approximations are given by

$$\tilde{\Phi}^{roe} = \frac{1}{2}[\Phi_l + \Phi_r - \sum_d |\tilde{\lambda}^{(d)}| \tilde{\mathbf{e}}_l^{(d)} \cdot (\mathbf{F}_r - \mathbf{F}_l) \tilde{\mathbf{e}}_r^{(d)}] \quad (2.45)$$

The form of the dissipative component, with its jumps in the state vector and inclusion of eigenvalues and vectors, (which is also seen in Marti and Müller's Piecewise Parabolic Method and Falle and Komissarov's Upwind scheme), prompts us to look more closely at commonly used dissipation terms in SPH applications. It should be noted that, as with the Newtonian cases, the eigenvalues of this hyperbolic system correspond to five waves through the material, three material or fluid velocities (one for each spatial coordinate direction), and an up and down wind acoustic wave.

We can describe the up/down wind velocity as a form of signal speed, v_{sig} , and propose a numerical viscosity of the form

$$\Lambda = K v_{sig} (\mathbf{F}_r - \mathbf{F}_l), \quad (2.46)$$

where K is a parameter of order unity. This decomposes to the diffusive terms used in the evolution equations of Chow and Monaghan [1997].

It is important that the deduction of the signal speed is done in such a way to avoid the possibility of the artificial viscosity terms leading to superluminal ($v > c$) information travel. If we take the signal speed as the actual acoustic upwind speed then this cannot occur, and the code will remain causal and physical.

2.6 The Equation of State

Having deduced the above equations, the system requires an appropriate equation of state to close it and maintain hyperbolicity. As mentioned previously, this requires a causal sound speed.

This problem was tackled by Chandrasekhar [1939] and Synge [1957]. Chandrasekhar approached from the concepts of an electron gas (including the effects of degeneracy) whereas Synge looked only at a simple gas of material and photonic particles. Being more general, we will use Chandrasekhar's equations and notations. However, it is important to note that in the non-degenerate case, the two formalisms are fundamentally the same.

Taking a small volume V of relativistic gas, and defining N as the number of particles ($N = \rho V$) and U as the internal energy ($U = \rho \epsilon V$), then these values are given by

$$N = \frac{Vm^3c^3}{\pi^2\hbar^3} \int_0^\infty \frac{\sinh^2 \theta \cosh \theta d\theta}{\frac{1}{\Sigma} \exp^{\theta mc^2 \cosh \theta} + 1} \quad (2.47)$$

$$U = \frac{Vm^4c^5}{\pi^2\hbar^3} \int_0^\infty \frac{\sinh^2 \theta \cosh \theta (\cosh \theta - 1) d\theta}{\frac{1}{\Sigma} \exp^{\theta mc^2 \cosh \theta} + 1} \quad (2.48)$$

$$P = \frac{m^4c^5}{3\pi^2\hbar^3} \int_0^\infty \frac{\sinh^4 \theta d\theta}{\frac{1}{\Sigma} \exp^{\theta mc^2 \cosh \theta} + 1} \quad (2.49)$$

These are expressed using Jüttner's transformation of

$$\sinh \theta = \frac{q}{mc} \quad (2.50)$$

$$E = mc^2(\cosh \theta - 1) \quad (2.51)$$

and

$$\frac{1}{\Sigma} = \exp^{\epsilon - \theta mc^2} \quad (2.52)$$

They are best understood by considering the volume V is divided up into cells of Planck dimension ($\hbar^3 = (\frac{h}{2\pi})^3$) capable of holding one fermion in a given state. These are then integrated over the available momentum levels (q). As can be seen in (2.52), the function Σ depends on both ε which is a measure of the thermodynamic potential and ϑ which is the reciprocal temperature.

In their current forms, Chandrasekhar's equations completely describe the state of an ideal fluid. If we look at the functional $\frac{1}{\Sigma}$, and the condition of degeneracy, we see ε dominate the exponential, and $\frac{1}{\Sigma}$ becomes negligible, leaving the simpler integrals

$$N = \frac{Vm^3c^3}{\pi^2\hbar^3} \int_0^\infty \sinh^2 \theta \cosh \theta d\theta \quad (2.53)$$

$$U = \frac{Vm^4c^5}{\pi^2\hbar^3} \int_0^\infty \sinh^2 \theta \cosh \theta (\cosh \theta - 1) d\theta \quad (2.54)$$

and

$$P = \frac{m^4c^5}{3\pi^2\hbar^3} \int_0^\infty \sinh^4 \theta d\theta \quad (2.55)$$

Of more interest to our study is the non-degenerate limit, where $\frac{1}{\Sigma}$ is large relative to unity. In this case we are left with

$$N = \frac{Vm^3c^3}{\pi^2\hbar^3} \Sigma \int_0^\infty \exp^{-\vartheta mc^2 \cosh \theta} \sinh^2 \theta \cosh \theta d\theta \quad (2.56)$$

$$U = \frac{Vm^4c^5}{\pi^2\hbar^3} \Sigma \int_0^\infty \exp^{-\vartheta mc^2 \cosh \theta} \sinh^2 \theta \cosh \theta (\cosh \theta - 1) d\theta \quad (2.57)$$

and

$$PV = \frac{Vm^4c^5}{3\pi^2\hbar^3} \Sigma \int_0^\infty \exp^{-\vartheta mc^2 \cosh \theta} \sinh^4 \theta d\theta \quad (2.58)$$

$$= \frac{Vm^3c^3}{\pi^2\hbar^3} \frac{1}{\vartheta} \Sigma \int_0^\infty \exp^{-\vartheta mc^2 \cosh \theta} \sinh^2 \theta \cosh \theta d\theta \quad (2.59)$$

$$= NkT \quad (2.60)$$

These can then be solved and the solutions expressed via modified Bessel's functions of the second kind, $K_\nu(\chi)$, where $\chi = \vartheta mc^2$ giving

$$U = Nmc^2 \left\{ \frac{3K_3(\chi) + K_1(\chi)}{4K_2(\chi)} - 1 \right\} \quad (2.61)$$

Equation 2.61 coupled with Boyle's law (Equation 2.60) and the relation between thermal energy and temperature for a material gas of


$$U = c_v T \quad (2.62)$$

reveals

$$\begin{aligned} PV &= NkT \\ &= \frac{NkU}{c_v} \\ \frac{PV}{U} &= \frac{c_p}{c_v} - 1 \\ &= \Gamma - 1 \end{aligned} \quad (2.63)$$

$$= \frac{1}{\chi} \left(\frac{3K_3(\chi) + K_1(\chi)}{4K_2(\chi)} - 1 \right)^{-1} \quad (2.64)$$

The righthand side of the relation can be plotted using Maple (see Figure

FIGURE 2.4: $\frac{U}{PV}$ against $T[K]$

2.4), revealing the asymptotic limits of $\Gamma - 1 = \frac{2}{3}$ for a non-relativistic gas, and $\frac{1}{3}$ for the relativistic limit.

Any true representation of a hydrodynamical flow where the values of $\frac{mc^2}{kT}$ traverse the values of 0-10 ($T \sim 5 \times 10^{12} K$ for a baryonic flow), need to be aware of this change in Γ . This has been highlighted by plotting the percentage difference between Chandrasekhar's equations and a constant $(\Gamma - 1)$ value. In this plot (Figure 2.5), we have assumed one changes from the non-relativistic value of $\frac{5}{3}$ to the relativistic limit $\frac{4}{3}$ at the appropriate time. It is immediately apparent that in the transition temperatures, an error of the factor 2 can be introduced by this equation of state. Note : the glitch in the curves at $\sim 4e11$ is an artefact of the plotting routine.

For smaller temperature domains it would be quite acceptable to fix the ratio of specific heats to a constant value, and use the form below.




FIGURE 2.5: Absolute Error over a constant ratio of Specific Heats model. Note the factor of 2 error induced at $T \sim 10^{13}$

$$\begin{aligned}
 P &= P(\rho, e) \\
 &= (\Gamma - 1)(e - \rho) \\
 &= (\Gamma - 1)(\rho_0 e)
 \end{aligned} \tag{2.65}$$

Synge took his own derivations of this work, and calculated the sound-speed of a material gas to be

$$\frac{c_s^2}{c^2} = \frac{1}{\chi G(\chi)} \left\{ \frac{\chi^2(1 - G(\chi)) + 5\chi G(\chi)}{\chi^2(1 - G(\chi)) + 5\chi G(\chi) - 1} \right\}, \tag{2.66}$$

where $G(\chi) = \frac{k_3(\chi)}{k_2(\chi)}$. This relationship is shown in Figure (2.6).

Due to the unrealistic computational effort required to calculate numerically the solution to Equation (2.66) for every particle at each timestep (possi-




FIGURE 2.6: c_s^2 as a function of temperature, clearly showing the relativistic limit

bly circumvented by use of a look up table), we will use the following relation and maintain Γ at a constant value to deduce the soundspeed:

$$c_s^2 = \frac{1}{h} \left(\frac{\partial P}{\partial \rho_0} + \frac{P}{\rho_0^2} \frac{\partial P}{\partial \epsilon} \right) \quad (2.67)$$

$$= \frac{\Gamma(\Gamma - 1)\epsilon}{1 + \Gamma\epsilon} \quad (2.68)$$

which can be seen to be monotonically increasing for a Γ greater than unity with a limit as the thermal energy approaches ∞ of $\Gamma - 1$. Depending on the choice of Γ , Equation 2.68 can be shown to be a reasonable approximation to Synge's formula, depicted in the following Figure 2.7.

The upper graph in Figure 2.7 depicts the approximation given by Equation 2.68 as having the correct limit, but increasingly over-estimating the value of




FIGURE 2.7: Comparison of Sygne's soundspeed (Solid Line) vs Equation 2.68 (Points)

c_s^2 , The opposite is true in the righthand view, where higher temperatures are used, and $\Gamma = \frac{4}{3}$. Here the sound speed is underestimated, but still maintains the correct limit.

The choice of the value given to Γ will range from mildly relativistic value of $\frac{5}{3}$ to the ultra-relativistic value of $\frac{4}{3}$.

2.6.1 RHIC EoS

Since the inception of nuclear accelerators, collisions between particle beams have been used to probe the nuclear equation of state. Early work by Glassgold et al. [1959] looked at the possibility of supersonic shocks passing through the nuclear fluid, and proposed that study of the distribution of matter after a collision would reveal the compressibility coefficient. Since then, study of collisions between nuclei at high relative speeds indicate large compressions within the nucleus, leading to the possibility of quasi-stable density isomers, and even exotic states of matter such as Quark-Gluon Plasmas. All of these things can be deduced from experiments such as those carried out at Fermilab's Tevatron. Theory suggests that the observed distribution of debris after one of these collisions should be directly related to the passage of a shockfront through, and over the surface of, the target nuclei. The ejection of material from the nuclear surface is assumed to be attributed entirely to the particle's Fermi motion [Amsden et al., 1977].

Hydrodynamical models have been used from as early as 1955 (Belenki and Landau) to study these collisions. Many early studies focussed on central collisions, allowing the problem to be reduced to 2 spatial dimensions under the assumption of axial symmetry. Due to the complex nature of the flows and lack of boundary conditions, particle methods would seem particularly

suited to this regime. Particle methods may also be well suited to modelling the post collision phases, where equilibration and chemical freeze out occur. These phases require that the reactivity of the fluid 'turn-off' once a certain temperature or density is reached. In SPH, this is as simple as having a flag on each particle, controlling its reaction rates. Having been constructed in three spatial dimensions, we should also be able to study the differing dynamics of off-axis collisions. These issues are explored in more detail later in Section 7.6.

Differing from Field Calculations codes, fluid dynamical modelling relies entirely upon the assumption of Local Thermal Equilibrium. There is considerable conjecture as to whether this assumption can be applicable in the extreme case of nuclear collisions [Bravina et al., 1999]. The Two-fluid model [Amsden et al., 1978] is an attempt to circumvent this issue by incorporating discrete cross-sections from nuclear cascade models into its drag terms. This aside though, we should still hope to be able to, at least qualitatively, replicate some features of Relativistic Heavy Ion Collision (RHIC) studies.

As this code was devised as a test bed, and required to be as general as possible, the equations of motion have no terms for nuclear viscosity, Coulombic energy, surface tension or any of the suspected multi-field effects. As these are expected to be small compared to the large impact energies experienced, this should not affect the outcomes to any great degree. However, they do mean any observed clustering of particle matter and outflows will be physically meaningless. The momenta and energy distributions though, should be comparable to experimental measurement by deducing chemical compositions through assumptions of equilibria and reaction rates.

It has been stated clearly before, that this algorithm currently has no ability to account for the creation or destruction of particles. Neither can it keep track of RHIC specific variables, such as strangeness. As a result, an Equation of

State capable of phase transitions to quark plasmas is unnecessary. We will use instead, a simple equation of state, based only upon the nuclear fluid's internal state of energy.

Taking Myers & Swiatecki's application of a Thomas-Fermi treatment of an attractive Yukawa function, multiplied by a quadratic function of momentum, Amsden et al. [1975] deduce

$$P = \left\{ -\frac{2}{3} \left(\frac{\rho}{\rho_0} \right) - \frac{1}{3} \left(\frac{\rho}{\rho_0} \right)^2 + c \left(\frac{\rho}{\rho_0} \right)^{\frac{8}{3}} \right\} \rho_0 + \frac{2}{3} \epsilon \quad (2.69)$$

where ρ_0 is the unperturbed rest baryon number density. This gives us an explicit form for the pressure in terms of the known quantities of densities and thermal energies which we will use in later work.

Chapter 3

Discretizations and SPH

Methodology

The Smoothed Particle Hydrodynamical method, pioneered independently by Gingold and Monaghan [1977] and Lucy [1977] is well known. The basic premise involves dividing the computational space into free moving (Lagrangian) nodes or particles, which hold information about the surrounding field variables. In order to deduce a variable at a certain point, one interpolates between the surrounding nodes (particles) using direct summation and an appropriate kernel function to weight the separate contributions.

The SPH algorithm is an incredibly simple implementation, often involving less than 1000 lines of code. This is both its benefit and detraction, as minor errors in coding can still leave the code functioning fully, yet producing unphysical results (in contrast to a grid/finite difference scheme which is more likely to be brought down completely by a bug). Its simplicity allows for one relatively simple code to be used for a wide variety of calculations, by simply changing the initial and boundary condition modules.

3.1 Interpolation Theory

The interpolation used in SPH is an approximation to the identity

$$A(x) = \int A(x') \delta(x' - x) dx', \quad (3.1)$$

where $\delta(x' - x)$ is the Dirac delta function and A is any field variable. The δ function is replaced by a suitably behaved kernel function, $W(r, h)$. Here r is a radial vector with the particle's location as its origin, and h is a parameter described as the smoothing length, which controls the distance and volume a kernel is active for. It may change depending on the situation the particle is located in, but does not necessarily do so.

Substituting this functional kernel for the delta function gives the approximation to the field variable

$$\langle A(x) \rangle = \int A(x') W(x' - x, h) dx'. \quad (3.2)$$

A 'suitably behaved' kernel function is one which satisfies the conditions

$$\int W(x' - x, h) dx' = 1 \quad (3.3)$$

$$\text{and} \quad \lim_{h \rightarrow 0} W(x' - x, h) = \delta(x' - x) \quad (3.4)$$

Naturally, there are many choices for W which will fulfil these requirements, and the choice made can greatly affect the results. A Gaussian may seem a most natural and obvious choice, but having wings extending out to infinity requires a complete summation over $N - 1$ particles. In an effort to reduce the computational expense of these summations, which would be par-

ticularly prohibitive in three dimensions, most kernels used have 'compact support,' that is having zero value beyond some length scale (typically $2h$).

For reviews on kernel choices and their effects, please refer to Morris [1996], Monaghan [1992], Lattanzio and Monaghan [1983], Monaghan and Gingold [1983] and Fulbright et al. [1995] for a comparison between spherical and spheroidal kernels.

Assuming that we have a suitable kernel function, and noticing that ρdr is the element of mass, dm , the integral (3.2) can be approximated according to

$$\begin{aligned} \langle A(x) \rangle &= \int A(x') W(x' - x, h) dx' \\ &= \int \frac{A(x')}{\rho(x')} W(x' - x, h) dm \\ &\approx \sum_b v_b \frac{A(x'_b)}{\rho(x'_b)} W(x' - x, h) \end{aligned} \quad (3.5)$$

where v_b is the mass element associated with particle b . In non-relativistic applications it is the true mass element. In relativistic calculations, the mass is a frame dependent value. As a consequence, v_b represents the invariant number of baryons represented by the particle, recalling the assumption of simple baryonic composition and conservation.

From this derivation we get the derivatives of the field variable to be

$$\frac{d \langle A \rangle}{dt} = \langle \frac{dA}{dt} \rangle \quad (3.6)$$

$$\begin{aligned} \nabla \langle A \rangle &= \langle \nabla A \rangle \\ &\approx \sum_b v_b \frac{A(x'_b)}{\rho(x'_b)} \nabla W(x' - x, h) \end{aligned} \quad (3.7)$$

It should be noted that in using the above identities, the time and spatial

components of the smoothing length (which is variable) have been ignored. It would be expected that provided the length scale of variation in h is much greater than h , then these contributions would be negligible. The full effect that these terms may have on the general form of the error is unclear, although the applications presented in this work give reasonable, self-consistent results. Nonetheless, a fuller review of the possible effects is given in Section 4.5.

3.2 Discrete, Non-Dissipative, Lagrangian Equations of Motion

Having derived the equations of motion for a relativistic gas in Section 2.4, and the functional form of the particle equations, we can now discretise the continuum equations to produce the equations of motion for an SPH particle. In doing so, we attempt to utilise symmetry of forces where possible, reducing computational time, but more importantly, conserving linear and angular momentum by enforcing body-centred reactions.

Where appropriate, we present the equations using suitable computational units, either to reduce clutter in the equations, or to increase performance and stability in later numerical work.

The Lagrangian equations governing the system are

$$\frac{d}{dt}(D) = -D\partial_i v^i \quad (3.8)$$

$$\frac{d}{dt}(e) = \frac{1}{D}[(Pg^{00}v^i)_i - (Pg^{0i})_i] \quad (3.9)$$

$$\frac{d}{dt}(q_i) = \frac{-1}{D}[(\sqrt{-g}P)_i - \frac{\sqrt{-g}}{2}T^{\mu\nu}g_{\mu\nu,i}] \quad (3.10)$$

$$\frac{d}{dt}\mathbf{r} = \mathbf{v} \quad (3.11)$$

where $D = \sqrt{\eta}\gamma\rho_0$ and $\sqrt{\eta}$ is the determinant of the spatial metric after it has been decomposed into a 3 + 1 form [Siegler and Riffert, 1999].

It can be noted that by taking the appropriate limits ($g_{\mu\nu} \rightarrow \eta_{\mu\nu} = \text{diag}\{-1, 1, 1, 1\}$, and $\mathbf{v} \ll c$), these reduce to the easily recognised Euler equations of hydrodynamics.

$$\begin{aligned}\frac{d}{dt}(\rho) &= -\rho\partial_i v^i \\ \frac{d}{dt}(e) &= -\frac{1}{\rho}(Pv^i)_{;i} \\ \frac{d}{dt}(q_i) &= -\frac{1}{\rho}P_{;i}\end{aligned}$$

Taking equations (3.8) to (3.10) and reducing them to a symmetric SPH form, we find

$$\frac{d}{dt}(D_a) = \sum_b m_b (\mathbf{v}_a - \mathbf{v}_b) \cdot \nabla_a W_{ab} \quad (3.12)$$

$$\frac{d}{dt}(e_a) = \sum_b m_b \left\{ \frac{P_a}{D_a^2} (g_a^{00} v_a^i - g_a^{0i}) + \frac{P_b}{D_b^2} (g_b^{00} v_b^i - g_b^{0i}) + \Omega_{ab} \right\} \partial_i W_{ab|a} \quad (3.13)$$

$$\frac{d}{dt}(\mathbf{q}) = - \sum_b m_b \left\{ \frac{\sqrt{-g_a} P_a}{D_a^2} + \frac{\sqrt{-g_b} P_b}{D_b^2} + \Pi_{ab} \right\} \nabla_a W_{ab} + \frac{\sqrt{-g}}{2D} T^{\sigma\nu} g_{\sigma\nu,i|a} \quad (3.14)$$

Note that the $|_a$ indicates the preceding term is calculated with respect to the a^{th} particle's local Riemann Normal Coordinates.

Other than our inclusion of dissipation terms, Π_{ab} and Ω_{ab} , these are identical to equations (4.28) and (4.35) in Monaghan and Price [2001] which are derived using a Lagrangian method. The energy equation differs slightly, in that we have assumed a static metric, so we have no time derivative component of $g_{\mu\nu}$. These can be compared to the equations of Siegler and Riffert

[1999] to see that whilst the forms are similar, we have used a symmetrised pressure term and avoided the use of $\sqrt{-g}$ outside the summation.

It should be noted that keeping the SPH equations symmetric in a and b promotes, but does not guarantee, exact conservation of linear and angular momentum. The last term in equation (3.14) has a non-symmetric component, which could lead to the violation of these conservation laws. As we are looking at a system where the particle positions do NOT feed back into a metric calculation, the symmetries of the posed problem dictate how badly these conditions are violated. For instance, a direct radial infall on a Schwarzschild metric, will not conserve linear momentum, as the particles will accelerate towards the source, which will not reciprocate and accelerate towards them (being fixed in space). Similarly, in the same metric, an orbital problem will conserve angular momentum, but the linear momentum will oscillate as the particles complete each orbit. We need to be aware of this, and interpret results with this in mind.

3.2.1 Conservation Properties

It is important to examine just how this algorithm does conserve values. The usual conservation of mass becomes a more difficult proposition in a relativistic setting, as the measured mass is a frame dependant value. We do know however, that total baryonic rest-mass should be strictly maintained. This is easily accomplished by attributing to each particle a fixed number of baryons, v_b . Naturally, summing over the system will generate the total number of baryons within the system, which will remain unchanged.

$$\begin{aligned}
 N_{\text{baryon}} &= \int \rho dx \\
 &= \sum_b v_b \int W(x, h) dx \\
 &= \sum_b v_b
 \end{aligned} \tag{3.15}$$

$$\therefore \frac{d}{dt} N_{\text{baryon}} = 0 \tag{3.16}$$

Should a more complex equation of state be used, where particles can be created by a process of energy exchange, then a more complex conservation of mass-energy condition will need to be devised. Although this is would be a fruitful and useful extension, it is unnecessary here and will not be attempted in this thesis.

The energy of the system can similarly be calculated by summing over all particles.

$$E_{\text{tot}} = \sum_b v_b e_b \tag{3.17}$$

$$\begin{aligned}
 \frac{dE_{\text{tot}}}{dt} &= \sum_b v_b \frac{de_b}{dt} \\
 &= - \sum_b v_b \sum_a v_a \left\{ \frac{P_a}{D_a^2} (g_a^{00} v_a^i - g_a^{0i}) + \frac{P_b}{D_b^2} (g_b^{00} v_b^i - g_b^{0i}) \right\} \cdot \nabla_a W_{ab}
 \end{aligned} \tag{3.18}$$

$$\begin{aligned}
 &= -\frac{1}{2} \left(\sum_b v_b \sum_a v_a \left\{ \frac{P_a}{D_a^2} (g_a^{00} v_a^i - g_a^{0i}) + \frac{P_b}{D_b^2} (g_b^{00} v_b^i - g_b^{0i}) \right\} \cdot \nabla_a W_{ab} + \right. \\
 &\quad \left. + \sum_b v_b \sum_a v_a \left\{ \frac{P_a}{D_a^2} (g_a^{00} v_a^i - g_a^{0i}) + \frac{P_b}{D_b^2} (g_b^{00} v_b^i - g_b^{0i}) \right\} \cdot \nabla_a W_{ab} \right)
 \end{aligned} \tag{3.19}$$

By interchanging the a and b on the second summation in equation (3.19) and

noting the antisymmetry of $\nabla W_{ab} = -\nabla W_{ba}$, the second summation changes sign, reducing the right hand side to 0.

If we perform a similar operation upon the momentum equation in the i^{th} coordinate direction, we see

$$\frac{dq_i}{dt} = - \sum_b v_b \frac{\sqrt{-g_b}}{2D_b} T^{jv} g_{av,ilb} \quad (3.20)$$

In the case mentioned before, where we have spherical symmetry, i.e. the spatial coordinates are (r, θ, ϕ) , we can see this term will be zero for q_θ and q_ϕ due to the derivative of the metric being zero in the θ and ϕ directions, but not so for the radial component q_r .

3.3 Dissipation terms and their derivation

Most numerical methods will require numerical or artificial viscosity to help diffuse post-shock ringing, or to stop particle interpenetration. It should not be assumed that approaches which give good results in one-dimensional models will necessarily carry over to higher dimensional work. This is particularly true in the effort to stop particles from penetrating each other. It is easy to see how, with the extra degrees of freedom available to a given particle in three dimensions, the body centred repulsion forces between two particles will lead not to the desired effect of heating and stopping motion, but rather to a minor deflection of the original flight paths, and a continuing, penetrating flow. This undesired interpenetration leads to unphysical communication between regions of gas. We have also seen how difficult it is to involve physically complete theories of heat flux into the equations of motion, and how dangerous (in terms of stability) and inadequate it might be to arbitrarily

import an ad hoc diffusive term into the system.

We can, however, borrow ideas from Riemann solutions, taking each particle involved in the interaction to be the left and right states in a Riemann configuration. We have already deduced an expected dissipation form to be

$$\Lambda = Kv_{sig}(\mathbf{F}_r - \mathbf{F}_l),$$

where K is a parameter of order unity and \mathbf{F} is the state vector of conserved variables used previously. Continuing in this line of reasoning leads to the separation of two dissipative terms, the first (Π_{ab}) for the momentum equations, and a second (Ω_{ab}) for the energy.

$$\Pi_{ab} = -\frac{Kv_{sig}(\mathbf{q}_a - \mathbf{q}_b) \cdot \mathbf{j}}{\bar{N}_{ab}} \quad (3.21)$$

and

$$\Omega_{ab} = -\frac{Kv_{sig}(e_a - e_b)\mathbf{j}}{\bar{N}_{ab}} \quad (3.22)$$

In these expressions, we use $\mathbf{j} = \frac{\mathbf{r}_{ab}}{|\mathbf{r}_{ab}|}$ as the unit vector in the direction \mathbf{r}_{ab} [Chow and Monaghan, 1997], and \bar{N}_{ab} is some averaging of the two baryon number densities a and b . We have also upheld the symmetry properties (provided an adequate choice for \bar{N}_{ab} is made) which allow the characteristics of Equations 3.18 and 3.20 to be maintained. The term v_{sig} is an appropriately chosen signal speed, which we will deal with in the next section (Section 3.4).

Note that like the Riemann solvers, these dissipation terms involve jumps in the conserved variables of \mathbf{q} and e . This is unlike non-relativistic SPH which uses the thermodynamic variables \mathbf{v} and the thermal energy component ϵ in its dissipation terms.

To ensure that these dissipative terms remain positive definite, we trans-

pose the jump quantities to the line of sight values q^* and e^* by replacing γ in Equations 2.36 and 2.37 with

$$\gamma^* = \frac{1}{\sqrt{1 - (\mathbf{v} \cdot \mathbf{j})^2}} \quad (3.23)$$

giving

$$\mathbf{q}^* = \gamma^* \mathbf{v} \left(1 + \epsilon + \frac{P}{\rho_0}\right) \quad (3.24)$$

and

$$e^* = \gamma^* \left(1 + \epsilon + \frac{P}{\rho_0}\right) - \frac{P}{\rho} \quad (3.25)$$

The final form of the dissipative terms then are

$$\Pi_{ab} = -\frac{K v_{sig} (\mathbf{q}_a^* - \mathbf{q}_b^*) \cdot \mathbf{j}}{\tilde{N}_{ab}} \quad (3.26)$$

and

$$\Omega_{ab} = -\frac{K v_{sig} (e_a^* - e_b^*) \mathbf{j}}{\tilde{N}_{ab}} \quad (3.27)$$

Both the dissipation terms are set to zero if the particles are moving away from one another, and approach zero in the continuity limit of infinite particles.

3.4 Choice and Derivation of Signal Speeds

The choice of signal speed can be, and usually is, guided by the physics of the problem under consideration. In the non-relativistic case, one can see signal speeds are simply eigenvalues of the flux equations, produced by combining the sound (acoustic) speed and bulk (material) velocity. We would assume this also to be the case in the relativistic problem.

Unfortunately, unlike the Newtonian cases, we do not have the ability to simply add the two velocities together. We are limited by the fact that issues of simultaneity for any two particles are complicated by the finite speed of light and subsequent information travel.

With this in mind, we take two inertial frames, the lab frame (K), in which two particles are moving, and the frame (K'), in which particle a is at rest. At some time particle a emits a signal to the second particle b . In K' , we know this signal to travel at the rest sound speed, c_s . However to avoid computational difficulties associated with maintaining every particle's proper time, the numerical code works entirely in the lab frame. We therefore need to determine the signal speed in the lab frame (K), which we will denote as \hat{c}_s .




FIGURE 3.1: Frames of Reference for soundspeeds: In the laboratory frame (on the left) we see two particles, a and b , travelling with unique velocities, when a sends a signal to b . The right frame shows a at the origin, with the x -axis aligned with a 's observed velocity in the lab frame. Particle a still transmits to particle b , but at a different speed.

It would be possible to derive the following relationships using a completely arbitrary frame. But through the postulate of there being no special

observer, we can claim freedom with respect to the establishment of our frames without loss of generality.

We establish K' such that its x^1 axis is aligned with the vector \mathbf{v}_a , and the straight line connecting particle a and b , denoted by \mathbf{r}_{ab} , is in the $x^1 x^2$ plane. We also define θ to be the angle between the vectors \mathbf{r}_{ab} and \mathbf{v}_a . This then reduces the original 3-dimensional problem to one in only two dimensions. See Figure (3.1).

We describe the signal with a 4-vector,

$$k^\mu = \{\omega, \mathbf{k}\}, \quad (3.28)$$

where ω is the angular frequency and the wave vector, \mathbf{k} , of wave numbers in each of the coordinate directions. This is easily transformed between the two frames by simple Lorentz transforms, giving.

$$\begin{aligned} \omega' &= \gamma_v(\omega - k_1 v) \\ k'_1 &= \gamma_v(k_1 - \frac{\omega v}{c^2}) \\ k'_2 &= k_2 \\ k'_3 &= k_3 \end{aligned} \quad (3.29)$$

Since $\omega = c_s k$, and

$$\begin{aligned} k_1 &= \cos(\theta) \\ k_2 &= \sin(\theta) \end{aligned}$$

the above transforms reduce to

$$\omega' = \gamma_v \omega \left(1 - \frac{v \cos(\theta)}{\hat{c}_s}\right) \quad (3.30)$$

$$k'_1 = \frac{\gamma_v \omega}{\hat{c}_s} \left(\cos(\theta) - \frac{\hat{c}_s v}{c^2}\right) = \frac{\omega'}{c_s} \cos(\theta') \quad (3.31)$$

$$k'_2 = \frac{\omega}{\hat{c}_s} \sin(\theta) = \frac{\omega'}{\hat{c}_s} \sin(\theta') \quad (3.32)$$

$$k'_3 = k_3$$

Using the invariant $k^\mu k_\mu$, and 3.30, we get

$$\frac{\omega'^2}{\omega^2} = \frac{(1 - \frac{v^2}{c_s^2})}{(1 - \frac{v^2}{c^2})} \quad (3.33)$$

$$\Rightarrow 1 - \frac{c^2}{\hat{c}_s^2} = \gamma^2 \left(1 - \frac{c^2}{c_s^2}\right) \left(1 - \frac{v \cos(\theta)}{\hat{c}_s}\right)^2, \quad (3.34)$$

which can be solved for \hat{c}_s , giving (after scaling all velocities to c)

$$\hat{c}_s = \frac{1}{\gamma^2 \Upsilon - 1} \left[v \Upsilon \gamma^2 \cos(\theta) \pm \sqrt{v^2 \cos^2(\theta) \Upsilon^2 \gamma^4 - (\gamma^2 \Upsilon - 1)(1 + v^2 \Upsilon \gamma^2 \cos^2(\theta))} \right] \quad (3.35)$$

where $\Upsilon = (1 - \frac{1}{\gamma^2})$.

Multiplying both numerator and denominator by $\frac{c_s}{\gamma^2}$ produces the following relation

$$\hat{c}_s = \frac{v \cos(\theta)(1 - c_s^2) \pm c_s \sqrt{(1 - v^2)[1 - v^2 \cos^2(\theta) - c_s^2 v^2 \sin^2(\theta)]}}{1 - v^2 c_s^2}. \quad (3.36)$$

If we decompose the velocity vector \mathbf{v} into its components parallel, v_{\parallel} , and perpendicular, v_{\perp} , to the signal direction \mathbf{r}_{ab} as suggested by the choice and

alignment of the frames, equation (3.36) becomes

$$\hat{c}_s = \frac{v_{\parallel}(1 - c_s^2) \pm c_s \sqrt{(1 - v^2)[1 - v_{\parallel}^2 - c_s^2 v_{\perp}^2]}}{1 - v^2 c_s^2} \quad (3.37)$$

It should be noted that written in this way, equation (3.37) for the signal speed is the same as the value λ_{\pm} deduced by Font et al. [1994] for their acoustic waves travelling within a relativistic medium.

One can immediately see that for the one dimensional case (where $v_{\perp} = 0$ and $v = v_{\parallel}$ necessarily), this reduces to the straight relativistic addition formula

$$\hat{c}_s = \frac{c_s \pm v_{\parallel}}{1 \pm c_s v_{\parallel}} \quad (3.38)$$

If we refer again to the General Relativistic High Resolution Shock Capturing work on an arbitrary metric of Font et al. [1994], the eigenvalues associated with the Jacobian matrix, B , introduced earlier (Section (2.5)) are

$$\lambda_0^i = \alpha v^i + \beta^i \quad (3.39)$$

$$\lambda_{\pm} = \frac{\alpha}{1 - v^2 c^2} [v^1 (1 - c_s^2) \pm c_s \sqrt{(1 - v^2) [\gamma^{11} (1 - v^2 c_s^2) - v^1 v^1 (1 - c_s^2)]}] - \beta^1 \quad (3.40)$$

with α and β^i representing respectively the lapse function and shift vector of the line element and superscripts correspond to directions.

By using this definition for our signal speeds (Equation 3.37) we ensure that the sound speeds remain causal, and so do not put at risk any assumptions made previously about the hyperbolicity of the system. We also ensure that the artificial viscosity terms (Π_{ab} and Ω_{ab}) do not lead to acausal effects.

Chapter 4

Numerics and Code Specifics

4.1 SPH: A Brief Overview

As has been previously mentioned, SPH is a Lagrangian particle method, where the field variables are attributed to small fluid elements or particles, which are then evolved through a set of hydrodynamical equations. By stipulating that these elements represent an equal and constant number of baryons, then there must be a direct relation between the number of particles per unit volume and the baryon number density. Being fundamentally an N-body program, the fluid elements are free to move through space according to the equations of motion. Inevitably this leads to some level of disorder in their locations at any given time, and we can conclude that recovering the density distribution from the particle distribution becomes a statistical process similar to recovering a probability distribution.

This is done by a smoothing kernel method, introduced previously in Section 3.1, which leads to the relations and functional forms of that section. The resolution of the method is controlled by the smoothing length, h . It is

natural then to apply some care when choosing both the kernel form, and the smoothing length value, to ensure that they give the desired resolution for the problem at hand. This choice is studied in Section 4.2.

Early applications of the SPH method selected a global value for h based on the overall density of the modelled structure. This however failed to utilise one of the great strengths of the SPH method. With the particles clustering with higher density, the dynamical range of resolution available can be increased by allowing locally varying smoothing lengths. There are a number of ways to implement this, and one needs to ensure that errors are not compounded, nor introduced by the implementation. This will be examined in Section 4.3.

With SPH's inherent Lagrangian nature, with particles being distributed as per the density distribution, it would be a shame not to utilise its full adaptive potential. Having built a method adaptive in space, and constrained by the usual numerical Courant conditions, adaptive time-stepping through individual particle timesteps is explored in Section 5.3.

4.2 Choice of Splines

As can be seen in equation (3.7), the gradients of the field variables are replaced by gradients of the kernel function. Therefore it is essential that the kernel function has derivatives at least to the order of the controlling hydrodynamical equations. For basic hydrodynamics, this translates to the first derivative needing to be smooth. For physical diffusion, at least the second derivative must exist. So long as these conditions, and the basic conditions previously asserted in Section 3.1 are abided by (being normalisation and approach to the delta function in the infinite limit) there exists a plethora of choices. An important consideration is the requisite of symmetry. This places additional

constraints, not only on the kernel choice, but the method of choosing the parameter h .

In order to sample a field value, a given particle must interrogate its neighbours, and there are two fundamental ways to do this. They are the concepts of computational 'gather' and 'scatter' of particle information. The 'gather' method is perhaps more intuitive, as it maintains the imagery and concept of each particle being a fluid element. The field variable, for instance baryon density, is then calculated at a given locale by assuming the observer is active and samples all surrounding particles, weighting their contribution by the kernel W .

$$\rho_a = \sum_b v_b W(r_{ab}, h_a) \quad (4.1)$$

Note that, for ρ_a the smoothing length used for all neighbours is h_a , that of particle a .

In the 'scatter' model, the particle is smeared out in space, according to the kernel's distribution. The local value of a variable is then obtained by simply summing the contribution from each surrounding particle.

$$\rho_a = \sum_b v_b W(r_{ab}, h_b) \quad (4.2)$$

Note that, unlike the gather model (Equation 4.1), here we use h_b for the density at particle a .

Provided the smoothing parameter, h , is constant for all particles, then these two models are identical. However, if one wishes to vary the resolution, these two methods will produce differing neighbouring particle lists. This is highlighted in Figure 4.1 when the particle in the upper left has a larger smoothing length than O , and is included in the 'scatter' model, but not in the




FIGURE 4.1: Schematic of gather and scatter

'gather' interpretation. As a consequence, exact conservation noted in Section 3.2.1 will not be maintained. It is therefore essential to somehow balance these two methods, using some form of hybrid.

This can be achieved by either symmetrising the smoothing length, or balancing the kernels [Hernquist and Katz, 1989], as follows:

$$W_{ab} = W(r_{ab}, \hat{h}), \quad (4.3)$$

where \hat{h} is some averaged smoothing length given by any one of the many possible ways of averaging two values. Most common are the simple geometric mean

$$\hat{h} = \frac{h_a + h_b}{2} \quad (4.4)$$

[Evrard, 1988] or even

$$\hat{h} = \sqrt{h_a h_b}. \quad (4.5)$$

The other option available is to average the kernel contributions from each particle,

$$W_{ab} = \frac{1}{2} \{W(r_{ab}, h_a) + W(r_{ab}, h_b)\} \quad (4.6)$$

[Hernquist and Katz, 1989]

It is argued that the latter method offers more refined control over the number of neighbours contained in the summation, as Evrard's symmetrisation can only offer control over the *minimum* number. Some interactions can be allowed, even though $2h_n$ may be less than r_{ab} , provided $2h_b$ is large enough to cause the average of the smoothing lengths to be greater than r_{ab} . In this thesis, the number of neighbours used in each particle summation will certainly be monitored, but not rigidly enforced or controlled, (Hernquist & Katz rigidly control the number of neighbours to be 50, then on subsequent runs allow a tolerance of ± 1 or 2) so we use Evrard's symmetrisation method. This is justified by preliminary testing showing no discernable difference in dissipation or response between the two methods.

Having decided how to balance the contributions and ensure symmetry, the next choice is that of a kernel form.

Schüssler and Schmitt [1981] indicate that kernels which have a null gradient at the origin ($\nabla W(r=0) = 0$) can lead to artificial clumping, and consequently should be replaced with a functional with a non-zero derivative across the origin. This artificial clumping effect, although not seen in any of our tests, is actually desirable in our application. As particles 'on top' of one another do not repulse or attract each other if $\nabla W(r=0) = 0$, they behave as a single, conglomerate particle. This can be utilised to avoid artificial heating near boundaries, caused by ghost particle collisions and heating. This will be covered in more detail in Section 5.1.

Traditionally, a spherically symmetric kernel was used, although for many applications, particularly those involving anisotropic compression or expansion, it has been argued that spheroidal kernels may be better [Fulbright et al., 1995]. Whilst this certainly has some meritorious effects, the added complica-

tions are deemed unnecessary, and arguably inappropriate, for this work. Not the least of these concerns is the reduction in the scheme's ability to conserve angular momentum through the loss of the body-centred reaction forces. The application of spheroidal kernels also gives rise to the question of exactly how best to deform the original spherical unit. This is generally done through either a simple analysis of the divergence of the fluid at study, or application of the tensor, $\frac{\partial v^i}{\partial x^j}$. In relativistic applications though, this would be further complicated by relativistic length contraction in the direction of motion, and even tidal curvature effects from a nearby gravitational source. Combined, or individually, these contributions would lead to a given deformation of the kernel appearing non-sensical and arbitrary as one moved between frames in the analysis.

Early work and discussions about Relativistic SPH have also argued that the kernel should be invariant under frame transformations [Lun, 2001]. It is this author's opinion that this also is an unnecessary complication. If we look at a similar finite difference scheme, with a laboratory centred symmetric grid, and place ourselves in a different frame, such as one co-moving with the fluid under consideration, the grid will appear arbitrary and of poor choice. But so long as one is aware of which frame calculations are performed in, it is a simple matter to transform variables into any other frame. For this reason we avoid the expense of firstly finding a suitable, invariant kernel function, and secondly, implementing it.

In support of this contention is the use of Riemann Normal Coordinates, which are effectively a flat space centred on the particle in question. Whilst curvature terms are integral for mapping one RNC to another particle, Section 2.3 shows that the Riemann tensor, $\tilde{R}_{\alpha\beta\mu\nu}$ is of the order of the coordinate patch length scale squared, δ^2 . As we place a flat space kernel upon each particle,

the errors associated with this as opposed to a curved space kernel can be expected to be of the order of its size, h^2 . Provided that the interpolation error of the kernel is of the same order, we will not be introducing any more errors.

The choice for the kernel is, in theory, infinite. In fact, work by Morris [1996] establishes how one can derive sets of kernels, which can be summed together, giving a kernel with whatever properties one desires. Although this work is focussed mainly on the stability of SPH (particularly as the particles are exposed to negative stresses), his findings on the desirability of higher order kernels are no less applicable.

Morris finds that lattice positioned particles approximating an isothermal gas admit instabilities for large ratios of h to particle spacing. These are induced by variations in the kernel's gradient function as the particles relax, and the growth of these instabilities is greatly inhibited by use of higher order polynomial spline kernels. This trend is observed right up to the Gaussian limit, where the propagation becomes negligible. This suggests that the method responds best to kernel implementations where the Fourier decomposition in $|r|$ falls off quickly.

With this foresight, one might conclude that the Gaussian kernel (or super-Gaussian [Monaghan and Gingold, 1983]) is the most natural choice. However, this leads to all $N-1$ particles contributing to each summation loop, when it is clear that only the close neighbouring particles would affect the motion of a given particle if we neglect the effects of gravity. Combining this effect with the unsupportable computations that would be required in a 3-dimensional calculation, and we conclude that at least a kernel with compact support is required.

Compact support can be introduced by simply truncating the Gaussian, this however means that the Fourier decomposition will no longer fall away

as a Gaussian itself, and the benefits of the Gaussian curve are lost. The work by Lattanzio and Monaghan [1983] introduced polynomial spline curves in an effort to combat this effect.

These spline kernels can be broken into two broad classes: those optimised to ensure stability through the smoothness of their curve and its derivatives, and those optimised for accuracy over interpolating across a grid. The first group use higher order polynomial splines (quadratic, cubic, quartic, etc.) to increase their smoothness and continuity of higher derivatives. This increases the computational effort in deducing them and using them, or at least memory if an adequate look-up table is devised, but the interpolation error is not decreased by the increased smoothness. The second group use a different set of splines to increase their interpolation accuracy over particles arranged on a regular grid. Such kernels are really only of much use if it is ensured that there are enough particles involved in the summations to ensure that the error induced by the Monte Carlo integration is maintained at a smaller scale than that induced by the kernel.

Another attribute of some of the higher order interpolation kernels is their sacrifice of positive definiteness (see Figure 4.2). Steinmetz and Müller [1993] point out that under certain situations this can lead to negative densities and energies, particularly if the sampling of particles is too coarse. This situation can easily arise in a rarefaction fan or similar, rapid expansion process. In a relativistic application, this can lead to there being no solution for the equation of state. Whilst it is certainly an argument that this situation can be avoided simply by employing a switch in the software to detect and correct a violation of this form, a more robust application can be implemented by use of a positive definite kernel.

Taking into consideration the desire for a general code capable of sim-

ulating relativistic flows yet avoiding as much as possible any reliance on exact solutions, the decision to use the 'common' B-spline kernel of Lattanzio and Monaghan [1983] has been made. This allows a fast and proven kernel with the requisite continuous derivative, and avoids unnecessary complications associated with non-spherically symmetric kernel choices or negative densities.

$$W(s, h) = \frac{1}{\pi h^3} \begin{cases} \frac{3}{2}(\frac{2}{3} - s^2 + \frac{1}{2}s^3); & 0 \leq s < 1 \\ \frac{1}{4}(2 - s)^3; & 1 \leq s < 2, \\ 0; & s \geq 2 \end{cases} \quad (4.7)$$

where s is the separation distance given in smoothing lengths, h . The $O(h^3)$ spline mentioned in the same paper was also assessed, but lead to no noticeable improvement and was deemed unsuitable, as it was no longer positive definite.

It should be noted here that there is no requirement that only one kernel is used. Optimised kernels can be used for density interpolation, with other choices applied to calculations involving the pressure force, for instance. This thesis, however, is looking to produce an algorithm in its most general form, capable of being optimised for a given application at a later date, and so uses the same kernel for all particle calculation loops.

The kernel is used not only in the integration of field variables, but also in the code's neighbour counting regime, in an effort to avoid the possibility of artificial shocking caused by a particle moving from a dense to rarefied region (or vice versa). A 'top-hat' implemented counting kernel may modify the smoothing length too quickly, allowing for the possibility of false shocking. This modification is detailed in Section 4.3.




FIGURE 4.2: Kernels of higher orders from Lattanzio and Monaghan [1983], normalised

4.3 Smoothing Length Modification

As has been previously noted, it is desirable to use the full adaptability of the SPH method. Early works in SPH have used a fixed value for h , usually based upon the initial particle spacing, and many of the error estimates rely heavily upon this assumption. Gingold and Monaghan [1978] introduced a basic varying h in their cloud collapse work, where the smoothing length was treated as a global parameter and was adjusted according to the simulation's central density variation. However, neither of these basic methods utilise the full adaptive resolution provided by endowing each particle with an indi-

vidual smoothing length, and the ability to change its interaction radius in response to the local environment.

A succinct analysis offered in Steinmetz and Müller [1993] highlights just how large the dynamical range can be for a given simulation by looking at the two number density extrema. The minimum density, ρ_{min} , is found when all particles are greater than $2h$ apart, and is found to be simply the particle's self contribution of

$$\rho_{min} = \nu W(0, h) \quad (4.8)$$

Conversely, when all the particles are compressed into one point, we get

$$\rho_{max} = \sum_i^N \nu_i W(0, h) \quad (4.9)$$

From these we see immediately that the theoretical range is N , although as is correctly pointed out by the authors, this is particularly optimistic. In order to avoid sampling errors, a minimum number of particles is required (generally considered to be of the order 10) and at the other extreme, the resolution offered by all particles collapsed to a point can only reveal density distributions proportional to the kernel itself. Therefore the dynamical range reduces to the still impressive $\sim \frac{N}{50} \dots \frac{N}{100}$

In the work of Gingold and Monaghan [1982], two proposals are offered. The first uses the relation

$$\sum_i^N e^{-\frac{(r_i - r_j)^2}{h_j^2}} = \kappa = \text{constant} \quad (4.10)$$

which can be solved for h_j iteratively, and is basically the requirement that the number of neighbours of the j^{th} particle remains constant. (Although this

expression, only holds for Gaussian kernels.)

The simpler option is to treat h as a variable, capable of being evolved along with any other through

$$\frac{dh}{dt} = \sum_i \frac{\partial h}{\partial \mathbf{r}_i} \cdot \mathbf{v}_i. \quad (4.11)$$

Although both these approaches are meritorious, the first is difficult to apply in its present form, and both can lead to instabilities through oscillatory corrections to the value of the smoothing parameter.

If we recognise the reciprocal relation between the number density and h , enforced by the particle distribution reflecting the mass distribution, then we are led to the relation

$$h(\mathbf{r}) \propto \frac{1}{\rho^d} \quad (4.12)$$

where d is the number of spatial dimensions. Unfortunately, this is a cyclic relation, $h(\mathbf{r})$ is required to calculate $\rho(\mathbf{r})$, which in turn is required to deduce h .

Wood [1981] and subsequently Miyama et al. [1984] use the value of the number density, ρ , at previous time step to specify h at the current time-step. This is however prone to instabilities mentioned above, which can be overcome by using a predicted or smoothed value of the density on the right hand side of relation (4.12) [Steinmetz and Müller, 1993] This method though, appears to be both application and kernel dependant, not to mention labour intensive.

By taking the time derivative of relation (4.12) and combining with the

continuity equation (in 3-dimensions), we can derive h 's evolution to be

$$\frac{dh}{dt} = \frac{1}{3}h\nabla \cdot \mathbf{v}, \quad (4.13)$$

which is the method implemented by Evrard [1988]. It is clear that this implementation relates the change in h to its current value. Should a particle from a higher density region (small h) find itself in a lower density region, its ability to adapt relative to a local particle undergoing the same divergence is greatly hindered. We would prefer that the technique used to modify the smoothing length is based more on the particle's surroundings.

Another technique is that employed by Borge et al. [2000] in their Regularised SPH formalism, which quantises the values of h , allowing only multiples of powers of two times some base value, h_0 . These are chosen by fitting a piecewise-smooth curve through the density profile. Whilst showing very promising results, particularly in Magneto-Hydrodynamical simulations [Borge et al., 2001], it is deemed too computationally intensive and not general enough for the applications in this thesis.

More appropriate than these methods of h evolution is the concept of nearest neighbour limits, where the length of the neighbour list is maintained at a constant, or nearly constant value. (An exact maintenance is computationally expensive for arguably little gain, and many algorithms therefore work from the closely related density profile.)

So, rather than matching the smoothing length to the density profiles (or some related smooth profile), we can relate it to the number of neighbours, which through the theoretical arguments presented in Section 4.1 should offer a direct relationship.

This method evolves h by averaging the current smoothing length and that

implied by the neighbour list length, N^n (the number of particles inside the interaction radius of $2h$ at timestep n) indicated by

$$h^{n+1} = \frac{h^n}{2} \left(1 + \left(\frac{N_{desired}}{N^n} \right)^{\frac{1}{3}} \right) \quad (4.14)$$

[Hernquist and Katz, 1989]

An appropriate minimum limit to h is usually enforced, depending on the problem posed. Unfortunately this algorithm, as with those mentioned previously, can suffer from oscillatory behaviour, particularly where there is rotation in the problem. It can overcorrect for the apparent density changes (tricked by particles moving closer together, but not on colliding orbits), leading to the number of neighbours being excessive, then too low on successive timesteps. This artificial oscillation in the neighbour count (and consequent assumptions about density) results in false clumping and 'ringing.' It can be reduced by a smaller timestep, but this is not a desirable alternative due to the already time-intensive nature of modern simulations. It is better to modify the root cause, which is due to the effective 'top-hat' potential used in standard neighbour counting regimes. This allows a minor extension (or reduction) in the interaction radius defined by h to dramatically affect the number of particles involved in the summation.

This case can be easily envisaged by visualising a particle with an interaction radius of $2h$ which is enforced rigidly by counting all particles within this radius as whole particles. Suppose that this particle had one neighbour less than the optimum number, and so wishes to increase its smoothing length. If there were a cluster of five particles just outside this counting potential of $2h$, then the increase in h will lead to an over supply of 4 neighbours. On the next time-step, the particle will reduce its smoothing length to below what it was

when we started this example, which will naturally be too few neighbours.

If instead, we reduced the weighting of particles near the $2h$ cut-off, then a small increase in h , which causes the cluster of particles to be included, will not result in such an oscillation being established, and will serve to stabilise the algorithm. This is achieved by removing the influence of the effectively discontinuous 'top-hat' potential, and replacing it with a smoothly decreasing one.

To illustrate this we show the usual weighting used in a standard SPH algorithm, W_{std} , with that suggested by Thacker et al. [2000] and implemented in this code, W_{count} .

$$W_{std}\left(\frac{r}{h}\right) = \begin{cases} 1; & 0 \leq \frac{r}{h} < 2 \\ 0; & \frac{r}{h} \geq 2 \end{cases} \quad (4.15)$$

$$W_{count}\left(\frac{r}{h}\right) = \begin{cases} 1; & 0 \leq \frac{r}{h} < 1.5 \\ \pi W_{ab}(4\frac{r}{h} - 1.5); & \frac{r}{h} \geq 1.5 \end{cases} \quad (4.16)$$

where W_{ab} is the usual kernel function employed for the particle summation loops.

The number of neighbours recognised as being used in a given summation cycle changes from the previous methods of

$$N_a = \sum_b 1$$

to

$$N_a = \sum_b W_{count}\left(\frac{r_{ab}}{h_{ab}}\right) \quad (4.17)$$

We then attempt to maintain this value as a constant number of neighbours, $N_{desired}$. Setting $N_{desired}$ higher, stabilises post-shock oscillations, but smears out

shock-fronts. A lower value resolves sharper discontinuities, but is less able to damp out oscillations. The actual counting procedure involves little work, as dynamical linked lists are used in the particle loops, and the lengths of these are required for many other subroutines. In implementing this sort of modification, one needs to ensure that 'false shocking' caused by a particle finding itself in a drastically different environment and having too few particles to complete the summation adequately, is avoided. This is controlled by applying Courant conditions on each time step.

Complimenting the idea of limiting timesteps by Courant conditions, we also attempt to control the advent of 'false shocking' by controlling the amount the smoothing length, 'h,' can be adjusted in any one timestep. This is done by the algorithm again described by Thacker et al. [2000]

Setting $s = (\frac{N_{\text{actual}}}{N_i})^{\frac{1}{3}}$ (for 3 dimensions), we use

$$h_i^n = h_i^{n-1}(1 - a + as) \quad (4.18)$$

where a is a weighting coefficient described by

$$a = \begin{cases} 0.2(1 + s^2); & s < 1 \\ 0.2(1 + \frac{1}{s^2}); & s \geq 1 \end{cases}$$

4.4 Associated Errors and Stability of the Algorithm

If we assume that the smoothing length is a fixed parameter, then the integral

$$\langle f(\mathbf{r}) \rangle = \int f(\mathbf{r}') W(\mathbf{r}' - \mathbf{r}, h) d\mathbf{r}' \quad (4.19)$$

can be expanded around $\mathbf{r} = \mathbf{r}'$ as a power series in h , giving

$$\langle f(\mathbf{r}) \rangle = f(\mathbf{r}) + \nabla f \cdot \int \mathbf{s} W(\mathbf{s}) d\mathbf{s} + \frac{1}{6} h^2 \nabla^2 f \int \mathbf{s}^2 h^3 W(\mathbf{s}) d\mathbf{s} + O(h^3) \quad (4.20)$$

[Gingold and Monaghan, 1982]

As the kernel is chosen to be an even function, and the spatial integrals occur over all space, the second term is found to be zero. The last integral is independent of h , so the smoothing approximation is shown to be of the order h^2 , or $O(\frac{1}{N^{1/3}})$.

If we assume complete disorder in the particle positions, the integral approximation becomes a Monte Carlo estimate, and consequently the error grows to $O(\frac{1}{\sqrt{N}})$. As the particles do remain ordered to some extent, but certainly deviate from a grid [Gingold and Monaghan, 1978], the true error is likely to be somewhere between these two bounds. Monaghan [1985] quotes Wozniakowski's attempt to qualify this by assuming the particles were quasi-ordered. This results in an error estimate of $O(\frac{(\log N)^d}{N})$.

Just how disordered the positions of the particles in a given simulation can become is an interesting issue ultimately at the centre of any method's error estimates. Due to the clear parallels between the kernel interpolation and statistical sampling, many applications have employed randomised initial grids to lay the particles upon. As a given simulation evolves, it naturally will maintain a random position matrix (if initialised with random positions). More modern models tend to place the particles upon some ordered grid, such as a body-centred cubic structure like this one. Although free to move in a disordered fashion, a truly random distribution is unlikely to result from such an initialisation.

4.5 Effects of Omitting ∇h Terms

When one implements the variable smoothing length algorithm, whether or not the smoothing algorithm remains an $O(h^2)$ accurate one is unclear. No matter upon which symmetrisation one chooses, the evolution equations will all contain $\frac{dh}{dt}$ or ∇h terms. One would expect that these terms would be smaller than the dominating hydrodynamical terms by $O(N^{-\frac{p}{2}})$, where p is generally considered to be ~ 2 , and so should be inconsequential for simulations with a high enough resolution [Gingold and Monaghan, 1982].

There may be, however, situations where this may not be the case, and violations in energy conservation may be of the order of 10% [Hernquist, 1993]. It should be noted that these large conservation of energy violations noted by Hernquist occurred in extreme circumstances, where the entire simulated fluid (governed by a polytropic equation of state) participated in a compressive collision and subsequent expansion, and that the thermal energy component was used as an integrable variable. Similar effects are noted by Nelson and Papaloizou [1994], who also integrate the thermal energy variable. For lesser grazing collisions, the violations would be expected to be less. Secondly, in the relativistic case, it is the total energy equation which is integrated forward, which has shown to be more stable than the integration of the thermal component used in both Hernquist's and Nelson & Papaloizou's models.

Having noted this though, any high resolution SPH simulation should be examined closely. The effects on these ∇h terms of integrating the total energy equation rather than the thermal energy needs to be studied. For the purposes of this work, we will not attempt to include these smoothing length derivatives, relying on the inverse relation between N and the error estimates and caution, to guide us.

Chapter 5

Application Specifics

Having prepared the equations of motion, the SPH formalism, and established how the basic algorithm will function, there are now some issues specific to the application which need to be addressed. These include determining how to implement boundary and initial conditions (Sections 5.1 and 5.2), finding adequate, stable timestepping algorithms (Section 5.3), efficient memory management techniques (Section 5.5) and which method will be used to solve for transient variables, such as the pressure, on each time-slice (Section 5.4).

The algorithm in its current form uses a leapfrog predictor-corrector method, where the conserved, integrable variables of ρ , q , and e are maintained on the integer time step, and the hydrodynamical variables (P , v , ϵ) are solved on the half-steps.

5.1 Boundary Conditions

One of the clear benefits of a particle method over a grid method is the ease in which free boundaries can be handled. Where possible, we use no bound-

ary conditions, letting the particles have access to the whole computational domain. Whilst this is fine for the collision experiments, and even the gravitational infall problems, the hydrodynamic test problems require a fixed box of known dimensions.

In this situation, we employ a rectangular box which can be mapped easily and quickly into rankspace (see Section 5.5). Two of the dimensions are periodic, and the third uses ghost-particles to maintain its structure. Periodicity is maintained by simply mapping the leftmost particles into pseudo-positions against the right boundary (and vice versa), and modifying the location vectors accordingly. This remapping is actually carried out in rankspace and, other than the restriction of the size of the computational domain in terms of h_{max} , requires little effort.



FIGURE 5.1: Maintaining a Periodic Boundary by remapping particles

Although Figure 5.1 shows a simple 1-dimensional application, care does need to be taken with higher dimensions, and exactly how one handles the edges and more particularly the vertices of a rectangular box. Care needs to be taken in choosing the dimensions of the grid such that slight mismatches between particles placed on the grid and the boundaries of the computational box are not exacerbated by the triple remap of particles near edges (across, up, and diagonally up and across if the particle is near the upper edge), and then a reflection in the ghost boundary (if the particle is near the end of the

computational cell). This many transformations can create a small clump of particles at the vertices of the computational box, which can in turn cause spurious waves as the extra gradients dissipate into the computational space.

Avoiding this situation does unfortunately place a restriction on the dimensions of the box. It is important that no particle have a smoothing length, or interaction circle, which exceeds the dimensions of the computational cell. Should this occur, multiple mirror mapping would need to be implemented to avoid this particle 'seeing' beyond the edge of the box. Some possible situations arising from this form of particle vector remapping are shown in Figure 5.2. The upper image shows a computational box of l length, with a particle, denoted by X . As X is near a boundary, its sphere of influence, governed by $2h$ can 'see' over to the other side. The particles located within $2h - (l - X)$ of the left boundary are added to X 's neighbour list with an additional tag indicating that the location vector \mathbf{r}_{Xb} does not point into the box, but rather is located outside the computational zone.

The middle figure shows a shaded region of particles, which must be counted twice in the particle summation loop. The first time they are contribute to X is through a normal interaction within the box. The second time, they will have had their location vector altered to place them in the psuedo cell to the right of the computational box.

Finally, we can see the effects of X 's smoothing length equalling $\frac{l}{2}$. This situation would be unstable as the whole computational box must be double counted, and X itself would contribute to the summation, possibly twice. It should be noted that as the number of spatial dimensions increases, this double counting effect increases four and eight-fold.

This periodic boundary implementation is applied to the y and z (x^2 and x^3) axes. The other axis is given ghost particles, which allow the computa-




FIGURE 5.2: 1-D example of remapping particles

tional box to describe asymmetrical problems in this direction. Initial testing showed that a body-centred cubic (BCC) lattice was quite stable if contained by these boundaries, and so the particles are initially laid out in this fashion. It should be noted that this places another restriction upon the dimensions of the box, in that the BCC lattice should fit precisely across it for maximum stability. If this condition is not met, the particles need time to relax into a hydrodynamical equilibrium. Whilst this in itself is not a problem, the simulated wave understudy will interact upon a disordered set of particles. This serves to increase numerical scattering, which may have some effects, or mask

some effects, which we might be interested in

The basic principle of ghost particles is to place an effective mirror at the boundary, where any particle on a path destined to leave the computational box will 'see' an identical particle inward bound on the other side of the boundary. Blindly reflecting these particles can lead to artificial heating, so a little care is taken in this reflection process, using knowledge of the initial system configuration.

Two approaching particles will respond to each other as described by their kernel gradient function and their smoothing parameter. As the particles come into range the kernels will exert a repulsive force, which in most cases will result in the particles heating, deflecting and passing each other. True head on collision will be rare, as small errors in the kernel function will impart a small sideways force on the particles, which is magnified as the encounter progresses. (If only one-dimension is available, then they will heat and repel.) If one of the approaching particles happens to be a ghost of the other, then there is no avoiding a collision at the boundary. Any deflection experienced translates the particle along the wall, where it artificially impacts and heats adjacent particles. This can lead to acoustic waves from the boundary rippling into the computational box.

We combat this by using a ghost reflection derived from the initial lattice layout. Instead of a single ghost particle approaching head on, we use (in the 3-dimensional BCC example) four daughter ghost particles, displaced one half of a grid spacing (as defined by the smoothing length) in each direction. Each of these daughter particles is given one quarter the baryon number of the original boundary approaching particle, with the other variables distributed accordingly (See Table 5.1). The result of this is rather than numerical errors pushing a directly approaching particle along the boundary, they serve to keep

it approaching head-on. A two-dimensional example, using the 2-D analogue of body-centred-cubic is shown in the right hand image of Figure 5.3. Instead of four daughter ghost particles, each particle births two daughters of half baryon number.

Attribute	Particle	Ghost Particle
x	$\{x_{\text{boundary}} - x, y, z\}$	$\{x_{\text{boundary}} + x, y \pm \frac{1}{2}\Delta y, z \pm \frac{1}{2}\Delta z\}$
v	$\{v_x, v_y, v_z\}$	$\{-v_x, v_y, v_z\}$
q	$\{q_x, q_y, q_z\}$	$\{-q_x, q_y, q_z\}$
baryon number	v	$\frac{1}{4}v$
number density	ρ	ρ
Pressure	P	P
Energy	ϵ	ϵ

TABLE 5.1: Division of properties from original particle to each of its 4 daughters

In early stages of a simulation, the true particles inside the boundary should be in their original BCC layout. If we were to use the traditional particle reflection, one can see that the pattern of particles is broken at the boundary (Figure 5.3, left diagram). But by using offset daughter particles, each grid position is now occupied by four, quarter weight particles in three dimensions, or two daughters in 2-D, effectively continuing the grid over the boundary. Should one implement a kernel with non-zero gradient at the origin as Schüssler and Schmitt [1981] suggest, then these daughter particles will strongly repulse each other, destabilising the grid. If, on the other hand, a zero gradient functional is used, these particles do not affect each other in terms of evolution, and so the configuration is stable. Note that a zero gradient kernel only affects the evolution equations, the summation approximation will see all four (two) daughters in the one location (indicated by the larger circle in Figure 5.3) and combine them accordingly.




FIGURE 5.3: Ghost Particle Configurations: Ghosts are depicted by dashed lines. Primary reaction vectors only shown. Note in the right hand example, how the BCC grid is continued.

5.2 Initial Conditions

The primary consideration with any particle code of this sort is to establish the particle positions according to the density distribution. This means that physical dimensions and particle spacings need to be carefully chosen. In order to avoid the need to relax the system before any simulation, it is important that these spacings fit neatly into the computational box so as the boundary conditions applied continue the grid pattern.

For collision problems, the particles are initially laid out in their BCC grid, before a radial distance calculation cleaves off the outer particles, leaving the desired sphere (or length-contracted ellipsoid if travelling at a relativistic speed). The particle hydrodynamical attributes are then allocated.

For gas tube problems, the computational cell is chosen such that the high density lattice spacing fits exactly across the box. The number of particles in this lattice is chosen so that the low density half fits as closely as possible across the same box, and there are enough particles in it to avoid the problems raised previously with regards to there being too few particles within the

interaction radius (particles 'seeing' themselves in the mirror) or h needing to be greater than l in order to maintain enough neighbours in their neighbour listings (refer to Figure 5.2). This process will be explained in more detail in the pertinent chapters (Section 6.3).

5.3 Timestepping Conditions

The previous chapter highlighted how beneficial it is to fully use the natural, adaptive nature of the SPH particles. Having spatial resolution without temporal adaption would be simply untenable in terms of the prohibitively small timesteps that would be required for a collision model, when for much of the run (approach and re-expansion) a larger timestep than that required for the actual contact phase could be used. This raises the issue of how to choose the timestep whilst maintaining stability.

Again we can refer to the much older techniques of grid and finite difference codes, and the existence of Courant, Friedrichs and Levy (CFL) conditions deduced from sound propagation. We use the usual two velocity and acceleration conditions, and a third 'spatial' CFL condition which becomes particularly necessary to avoid penetration when there are two stable, colliding regions, otherwise uninhibited by timestep controls. This last condition is simply blocking the particle from moving too far (in terms of its interaction sphere) in a given timestep. Each particle calculates these limits for the region around itself, and then the extreme is chosen to ensure stability across the whole computational space.

These Courant conditions are constrained by a set of parameters to further reduce the timestep from the theoretical safe limit. These parameters are initialised at the start of a simulation and remain constant.

The timestep used, dt , is given by

$$dt = \min\{C_v dt_v, C_a dt_a, C_h dt_h\} \quad (5.1)$$

with

$$dt_v = \min_i \frac{h}{v_{sig_i}} \quad (5.2)$$

$$dt_a = \min_i \sqrt{\frac{h_i}{|q_i|}} \quad (5.3)$$

$$dt_h = \min_i \frac{h_i}{\max_j |v_i - v_j|} \quad (5.4)$$

and typical values for C_v , C_a and C_h as 0.2, 0.35 and 0.35 respectively.

Whilst the range for tuning these parameter is immense, other than restricting the spatial condition in times before collisional impact, their effects were not studied in any detail. It should also be noted that Thacker et al. [2000] use the same Courant conditions, but use 0.4, 0.25 and 0.2 as their limiting parameters C_v , C_a , and C_h . As any value under unity should be stable, there is little difference in the actual effects of these changes. Their individual values are dependant upon the kernel used, and the time-stepping rule used. The dominant limiting factor is usually dt_v , and it has been further limited in our scheme in an attempt to reduce overshooting due to the frequent supersonic motions and large gradients involved. The only time dt_h needs to be invoked is when two otherwise stable balls of nuclear fluid are travelling towards each other. In this case the other time conditions dt_v and dt_a see no change in their environment (as the approaching fluid is initially outside interaction range) and will specify arbitrarily large timesteps as suitable, allowing the nuclei to pass right by each other in one timestep unless the dt_h constraint is invoked.

5.4 Solution for Pressure using a Non-Linear Equation

The final equation required is the equation of state (EoS), giving pressure as a function of the density and internal thermal energy, $P = P(\rho, \epsilon)$. Three choices of EoS are used in this work.

For most of the hydrodynamical calculations, the perfect fluid assumption is made. With this, we use an appropriate, constant Γ value and the relation

$$P = (\Gamma - 1)\rho_0\epsilon. \quad (5.5)$$

This form is not appropriate for colliding nuclei studies. For reasons specified in Section 2.6.1, we use the equation of state given in Amsden et al. [1975] and used in their later numerical works [Amsden et al., 1977] [Harlow et al., 1976].

$$P = \left\{ -\frac{2}{3}\left(\frac{\rho}{\rho_0}\right) - \frac{1}{3}\left(\frac{\rho}{\rho_0}\right)^2 + c\left(\frac{\rho}{\rho_0}\right)^{\frac{8}{3}} \right\} \rho_0 + \frac{2}{3}\epsilon \quad (5.6)$$

where ρ_0 is the unperturbed rest baryon number density.

Note that this equation does not allow for particle production, which is a complexity beyond this application at the moment. It is certainly apparent that this equation of state is not the most recent, however it is explicit, and in physical variables already in use in the main algorithm. For a complete application purely for studying Heavy Ion Collisions, not only would a more suitable field-based equation be used, but one would also use different attributes for the particles in the simulation. Whilst an interesting study, this was never intended as the primary purpose of this generalised algorithm, and so will not be entertained. However, this algorithm is suitably general, and

the application of such a state equation is a logical extension.

As in any Relativistic Hydrodynamical code, the variables can be split into the conserved variables evolved forward across timeslices (D, e & q), and those deduced at each step (v, P, ϵ, γ) as functions of the conserved variables. Unlike equation (5.6), equation (5.5) is implicit, and requires a specialised subroutine to obtain the hydrodynamical variables from the conserved, integrable variables at each time step. This is done by creating $f(P)$ from the equation of state and other constraint equations, and solving for its root with a Newton-Raphson method or other similar method.

Accordingly, we define $f(P)$ to be

$$f(P) = [\Gamma - 1] \rho_0(P) \epsilon(P) - P \quad (5.7)$$

where

$$\rho_0(P) = \frac{D}{\gamma(P) \sqrt{\eta}} \quad (5.8)$$

$$\epsilon(P) = \frac{\alpha \sqrt{-g}}{\gamma(P)} e - \frac{P \sqrt{\eta}}{\gamma(P) D} (g^{00} + \gamma^2(l^j)) - 1 \quad (5.9)$$

$$\gamma(P) = \frac{1}{\sqrt{1 - \left(\frac{q}{\sqrt{-g}} \right)^2 \frac{1}{e - \frac{P}{\gamma} g^{00}}}} \quad (5.10)$$

From these relations we generate $f(P)_a$ as a highly non-linear function of the location of particle a on the background metric, and the integrated hydrodynamic variables only, as follows:

$$\begin{aligned}
 f(P) = & \sqrt{-g}e - \frac{q^2}{\sqrt{-g}(e - \frac{P}{D}g^{00})} - \\
 & - \frac{P\sqrt{\eta}}{\alpha D} \left\{ g^{00} + \frac{\Gamma}{\Gamma-1} \right\} - \\
 & - \frac{1}{\alpha} \sqrt{1 - \frac{q^2}{\sqrt{-g^2}(e - \frac{P}{D}g^{00})}}
 \end{aligned} \tag{5.11}$$

We use the definition of q (equation (2.37)) and ϵ (equation (5.5)) to determine the flat space relation

$$q = v\left(\epsilon + \frac{P}{\rho}\right) \tag{5.12}$$

The minimum pressure will occur when all the energy of the particle is in the momentum, or when $|v| = 1$. This tells us the minimum pressure value possible for a given energy and momentum must be

$$P_{min} = \rho(|q| - \epsilon) \tag{5.13}$$

The monotonicity of $f(P)$ within the range P_{min} to ∞ ensures uniqueness of the solution.

Knowing the minimum value for the pressure, and the monotonicity of $f(P)$ allows the safe application of a Newton-Raphson method to solve for the pressure variable.

5.5 Neighbour Searching and Rank Space

As with many particle methods, a great amount of computational time is spent on finding out the neighbouring particles on a given computational node. Traditionally this has been done through the ubiquitous linked list

routine. This method has some shortcomings however, not the least of which is the requirement of a grid of some form to divide the particles into cells. In simulations where there may be large voids, or large variations in h , this can lead to difficulties in defining the grids and complications in terms of memory management.

As was alluded to in Section 4.2 a 'scatter' algorithm can give different nearest neighbour lists to a 'gather' model, and how important it is to avoid this and use symmetrical neighbour lists. The generation of these lists needs to be carefully considered, as this is the most computationally intensive component of the algorithm, and non-symmetrical listings have been shown to lead to violations in the conservation laws.

This problem has previously, and typically, been tackled by laying a grid over the computational region and scanning each cell for particles. If one defines the grid spacing to be $2h_{max}$, and systematically sweeps the grid, then a particle's neighbours must be in the particle's home cell, or at most the next cell or cells along in the sweep (previous cells being checked and summed beforehand). However, if the computational space has large variations in h , or is an awkward shape, as depicted in Figure 5.4, then this method is computationally expensive. Either cells contain far too many particles, increasing queries which result in no contribution to the summation loops ($r_{ab} > \hat{h}$), or memory space is wasted in allocating empty cells.

These situations can be avoided by maintaining a regular computational box, and giving all particles the same h value. This is very limiting in terms of physical spaces, but we can attain just this configuration if we transfer our computational space into its 'ranked' equivalent, and replace the aforementioned linked-list search, with ranked-lists.




FIGURE 5.4: An awkward computational domain, showing the memory cells used in linked listing, many of which will be empty

5.5.1 Ranklisting

As mentioned previously, linked lists are the most common form of nearest neighbour searching to be found in particle codes. In its simplest and briefest form, creating a linked list involves dividing the computational space into cells of a given size, usually determined by the twice the largest smoothing length. A sweep is then carried out, attributing each particle to one of these cells. Then, to find the interacting neighbours of a given particle, one need only look in that particle's home cell, and possibly the neighbouring cells. (All other particles assuredly more than $2h$ away and consequently not included in any interaction.)

The linked listing described here is most efficient when all the particles have the same h , and each cell of computational space has only one or two particles within it. The method's effectiveness decreases when there are large areas with no particles, or regions with h much smaller than the largest value used to assign cell sizes. This can lead to empty cells, or cells with many

more particles within them than is computationally efficient and is a problem exacerbated by increasing spatial dimensions. We can illustrate this with Figure 5.5 where we can see 10 particles unevenly spaced across the domain.




FIGURE 5.5: A 1D simulation, showing particle positions (denoted by letters) and the resultant density curve

Typically one would lay a grid over this space, defined by the largest smoothing length, i.e. that of particle a . This grid is shown in Figure 5.6. Note how, due to large differences in smoothing length, cell #2 is empty, whereas cell #5 has 6 particles in it.




FIGURE 5.6: The Linked List memory map for Figure 5.5

If the computational space can be modified or manipulated in such a way that the optimal configuration is achieved, then the problems associated with

wasted particle queries will be alleviated. This can be accomplished by shifting the computation into 'rank-space,' where each particles is ordered by its x, y and z values, and its rank amongst all other particles in these three directions, gives the rank coordinate of the particle in rank-space. One can immediately see that however complex and transmuted the original workspace becomes, it will always reduce down to a rectangular cell, of n_x, n_y , and n_z dimensions, with no voids, and a constant particle density of one particle per rank cell.




FIGURE 5.7: The same computational space as in Figure 5.5, transferred to rankspace. Particle names are shown in lower case, the x coordinate is shown in upper case (a is at $x = A$) The mapping system is shown by placing upper and lower bound maps onto the top and bottom of the cell respectively

In this space, we have replaced the spatial identification of a particle (x, y, z) with an integer based call (i^{th} in the x direction, j^{th} in the y , k^{th} in the z). The ordered sweep used in a linked list sweep is no longer available to us (as we do not know in which direction neighbours will be situated), so we also rank the smoothing lengths of each particle. We can then avoid double summations by only examining particles with a greater (or lesser) h .

This rank space can then be broken down in the usual fashion, creating a linked list of the interacting and neighbouring particles. The only difference required, is that a map of the particles' true spatial locations must now be

kept. This corresponds to mapping the outside faces of the data cube. This is then used, coupled with the home particle's position and smoothing length value, to reduce the complete rank box down to a number of cells surrounding the home particle's location which need to be searched. (This refined sub-box need not be symmetric, and will in fact reflect the asymmetries of the density (h) distribution.) The particles residing in each of these cells are assembled into a list of neighbours, which is used in the particle summation loops.

If we were using the linked lists of Figure 5.6 and wished to find the neighbours of particle c , then we would check the inhabitants of c 's cell, and the next along, giving a neighbour list of $\{b, d, e, f, g, h, i\}$. Many of these particles are in a higher density region, and so will not be required to interact with particle c . This leads to spurious enquiries and inefficiency.

Figure 5.7 shows the ten particles of the earlier example now placed into rankspace. In the upper image, the memory blocks for an occupation of one particle per cell is shown. Only one map is then required, indicated by the single coordinate map along the top. To deduce the neighbour list as before, particle c examines its location (C) and its smoothing length (h_c) to find an interaction radius of $C \pm h_c$. This is compared to the coordinate map, to see that only b, d and e fall within this radius.

The lower image shows the same situation, but here memory efficiency has been increased at the expense of searching efficiency, by using a cell occupation parameter ($\frac{N_{\text{particles}}}{N_{\text{cells}}}$) of 2. In this situation two coordinate maps, specifying the lowest and highest values in each cell. Particle c then must query the particles in the cells where $X_{\text{low}} > C - h_c$ and $X_{\text{up}} < C + h_c$. In such a simple example as this, the same list of particles (b, d and e) will be generated as before. However, it is easy to see how increasing the occupancy of the rankcells results in more false positives and spurious queries, but smaller memory demands.

Chapter 6

Preliminary (Shocktube) Testing

In this chapter we show the code's ability to deal with hydrodynamical shock waves. The importance of these waves, both as physical systems to be understood and as numerical test cases, is introduced in Section 6.1. Their analytical solutions are dealt with in the following Section 6.2. Having established all these foundations to the problem, Section 6.3 explains the routine used to establish the initial particle positions and hydrodynamical variables. We then conclude this chapter by looking at the algorithm's performance in modelling these phenomena.

6.1 The Shock Tube as a Hydrodynamical System

Sod [1978] introduced the shock-tube test as a powerful analytical tool in assessing hydrodynamical codes and their capabilities. The shock tube as used in Sod's paper, and later revised to relativistic attributes, is fundamentally a special class of Riemann problem.

A Riemann problem is governed by the conservation equations

$$\Phi_t + F(\Phi)_x = 0, \quad (6.1)$$

and two initial states separated by a discontinuity in one of more of the state variables.

$$\Phi(x) = \begin{cases} \Phi_l; & x \leq 0 \\ \Phi_r; & x > 0, \end{cases}$$

This initial state breaks down into some configuration of three, non-linear waves:

Rarefaction waves \mathcal{R} : where all the state variables are continuous and undergo a smooth transition between the before and after states.

Shock Waves \mathcal{S} : These are indicated by a discontinuous jump in the density and pressure. All state variables are thereby discontinuous across a shock front.

Contact Discontinuities \mathcal{C} : This wave travels at the fluid velocity and exhibits a jump in the density variable only. Pressure and Velocity (necessarily) are smooth over a contact discontinuity.

The differences between rarefactions and shockwaves are dictated by the pressure of the before and after states. Shocks raise the pressure, and abruptly accelerate the particles, whilst rarefactions result in a drop in pressure, and smooth particle accelerations.

After the breakdown of the initial state, two waves (either shock or rarefaction) move into the unperturbed left (L) and right (R) states. Between these two waves, two new states will appear (L^* and R^*), which will be separated by the contact discontinuity. If we use arrows to indicate a leftward or rightward

travelling waves, then the possible configurations of this initial state's decay are given by the three possibilities

- a) $L \quad \mathcal{R}_- \quad L^* \quad C \quad R^* \quad S_- \quad R$
- b) $L \quad S_- \quad L^* \quad C \quad R^* \quad S_- \quad R$
- c) $L \quad \mathcal{R}_- \quad L^* \quad C \quad R^* \quad \mathcal{R}_- \quad R$

in direct analogy to the Newtonian dynamics cases

The first scenario, (a), is the shocktube decomposition. Here a shock front impacts the lower energy state on the right, sweeping up material in a dense shell demarked by the shock and the contact discontinuity, while a rarefaction fan sweeps into the high energy state on the left.

The second decomposition, (b), is two colliding cold streams, and the third, (c), is an explosive expansion.

We will use two standard shock test cases in the analysis of this algorithm's performance. The first is that introduced by Sod (1978). This is a non-relativistic shock/rarefaction pair separated by a contact discontinuity, used to show the correct resolution of basic hydrodynamical shocks and the correct Newtonian limits.

The second shock is a relativistic shock where the fluid is accelerated to around $0.75c$. Whilst the speeds of the gas reached are only mildly relativistic, the particles swept up in the dense shell behind the shockfront experience a density increase of up to an order of magnitude. In a thermodynamic sense then, this is an extremely relativistic shock.

6.2 Exact Solution

In order to solve the relativistic shock-tube cases, one needs to calculate the locations of the three waveforms and the intervening states, L^* and R^* . Thankfully, the solutions are self similar in $\frac{x}{t}$, allowing initial solutions to be propagated forward in time by simple scaling. This is best highlighted by showing the shocktube in the $x - t$ plane.




FIGURE 6.1: Characteristics of the Shocktube on the $x - t$ plane

In order to avoid complexity caused by over-generalisations, we will work only with the shock where L^* is the higher pressure initial state, expanding into the R^* lower energy state. The problem reduces to 5 states separated by the waves introduced earlier. Reading from left to right:

1	L^*	Known state in a Riemann problem
2	Rarefaction Fan	
3	Plateau	Separated from 4 by CD
4	Plateau (Shocked Shell)	Separated from 5 by S
5	R^*	Known state in a Riemann problem

For a more complete and thorough examination of this problem, see the papers of Marti and Müller [1994], and Thompson [1986]. These works are summarised briefly in the next few sections.

6.2.1 The Rarefaction

Denoted in Figure 6.1 as Region 2 and bounded by χ_1 and χ_2 is the rarefaction fan. The self similar solution dictates that the rarefaction is isentropic through the stipulation that entropy is conserved along field lines. We can show this by transforming the fluid equations into the variable $\chi = \frac{x}{t}$ in the laboratory frame (Eulerian), giving

$$s_{,\mu} u^\mu = 0 \quad (6.2)$$

$$\frac{ds}{dt} u^0 + v \cdot \nabla s = 0 \quad (6.3)$$

which, as we can reduce to one-spatial direction

$$\Rightarrow \frac{ds_0}{dt} u^0 + v \frac{ds}{dx} = 0 \quad (6.4)$$

$$(v - \chi) \frac{ds}{d\chi} = 0 \quad (6.5)$$

In general, v will not equal χ , so entropy must be conserved. If we follow

the same transformation with the equations of motion, we derive

$$(v - \chi) \frac{d\rho_0}{d\chi} + \rho_0 \gamma^2 (1 - v\chi) \frac{dv}{d\chi} = 0 \quad (6.6)$$

$$(1 - v\chi) c_s^2 \frac{d\rho_0}{d\chi} + (v - \chi) \rho_0 \gamma^2 \frac{dv}{d\chi} = 0 \quad (6.7)$$

where

$$c_s^2 = \frac{1}{h} \left(\frac{\partial P}{\partial \rho_0} \right)_s \quad (6.8)$$

In order to avoid trivial solutions, the determinant of the above coefficient matrix must be null, which gives us the relation

$$c_s = \pm \frac{v - \chi}{1 - v\chi} \quad (6.9)$$

Note that this is simply the one-dimensional equivalent of Equation 3.37.

Due to the nature of the Riemann problem the left hand state is stationary. So solving Equation 6.9 for $v = 0$ gives $\chi_1 = -c_s$, where the minus sign comes from our initial establishment of the problem. This tells us that the head of the rarefaction fan travels into State 1 at the local sound speed.

We can solve for the differentials of the density and the velocity by taking Equation 6.9 and substituting it into Equation 6.6. This then gives us

$$\gamma^2 dv \pm \frac{c_s}{\rho_0} d\rho_0 = 0, \quad (6.10)$$

which in turn gives

$$\ln \left(\frac{1+v}{1-v} \right)^{\frac{1}{2}} \pm \int \frac{c_s}{\rho_0} d\rho_0 = \text{constant} \quad (6.11)$$

As we are using a polytropic equation of state, we can write the sound speed as a function of ρ_0 alone, giving

$$c_s^2 = \frac{k\Gamma(\Gamma-1)\rho_0^{\Gamma-1}}{(\Gamma-1) + K\Gamma\rho_0^{\Gamma-1}} \quad (6.12)$$

and allowing the integral in Equation 6.11 to be solved through the rarefaction fan.

We now have the relation

$$\left(\frac{1+v}{1-v}\right)\left(\frac{\sqrt{\Gamma-1}+c_s}{\sqrt{\Gamma-1}-c_s}\right)^{\pm\frac{2}{\Gamma-1}} = \text{constant} \quad (6.13)$$

which gives us the fan, but cannot be used yet, as we do not know χ_2 , the tail of the fan.

6.2.2 The Intermediate States

In order to know where the fan stops, we need to know State 3, L^* . We move to a frame which is comoving with the shock, and introduce the subscripts of s to denote shock values, A for ahead of the wave, and B for behind. With this notation, the relativistic Rankine-Hugoniot jump conditions are

$$[\rho_0 u^\mu] n_\mu = 0 \quad (6.14)$$

$$[T^{\mu\nu}] n_\mu = 0 \quad (6.15)$$

The square parenthesis denote the jump in the observed quantity

$$[Q] = Q_A - Q_B \quad (6.16)$$

and $n_\mu u$ is the normal to the hypersurface across which the discontinuity occurs. In the co-moving frame, this will be given by

$$n^\mu = \gamma_s \{v_s, 1, 0, 0\}. \quad (6.17)$$

The mass flux across the front will be

$$f = \gamma_s D_A (v_A - v_s) = \gamma_s D_B (v_B - v_s) \quad (6.18)$$

Note: we are using the notation of previous chapters.

Understanding that in this frame, the mass, momentum and energy fluxes must be continuous (or else there would be a build up at the shock front) we can deduce the jump conditions to be

$$[v] = -\frac{f}{\gamma_s} \left[\frac{1}{D} \right] \quad (6.19)$$

$$[P] = \frac{f}{\gamma_s} [q] \quad (6.20)$$

$$[vP] = \frac{f}{\gamma_s} [e] \quad (6.21)$$

We then take Equation 6.20, and using the definition of q and solving for v_B deduce

$$v_B = \frac{h_A \gamma_A v_A + \frac{\gamma_s (P_B - P_A)}{f}}{h_A \gamma_A + (P_B - P_A) \left(\frac{\gamma_s v_A}{f} + \frac{1}{D_A} \right)} \quad (6.22)$$

We now can use this relation and the condition across the contact discontinuity to deduce the plateau states 3 and 4 through

$$v_{L^*}(P) = v_{R^*}(P) \quad (6.23)$$

Having solved this, we also know χ_3 , the speed of the Contact Discontinuity, as it must travel at the fluid velocity, v .

6.2.3 The Shock Front

All that is left is to deduce the density of the shell in region 4 and the velocity of the shock itself, v_s . This requires the Taub adiabat Taub [1978], in analogy to the application of the Hugoniot adiabat in Newtonian shock theory. The Taub theory allows the jump condition for the relativistic enthalpy to be deduced as

$$[h^2] = \left\{ \frac{h_B}{\rho_B} + \frac{h_A}{\rho_A} \right\} [P]. \quad (6.24)$$

As the Equation of State is polytropic, we can remove the post shock density term, ρ_B through

$$\rho_B = \frac{\Gamma P_B}{(\Gamma - 1)(h_B - 1)}, \quad (6.25)$$

which leaves us with

$$h_B^2 \left(1 + \frac{(\Gamma - 1)(P_A - P_B)}{\Gamma P_B} \right) - h_B \frac{(\Gamma - 1)(P_A - P_B)}{\Gamma P_B} + h_A \frac{P_A - P_B}{\rho_A} - h_A^2 = 0 \quad (6.26)$$

This is a simple quadratic in the post-shock enthalpy as a function of the (known) post-shock pressure. It can be shown to always have one and only one physical root, which allows the mass flux f to be found through

$$f^2 = - \frac{[P]}{\left[\frac{1}{\rho} \right]} \quad (6.27)$$

This in turn gives v_s the shock velocity (χ_4) through Equation 6.18.

We have now deduced the complete analytical solution to the relativistic Riemann Shock.

6.3 Initial Conditions

It has been previously stated that we attempt to keep the unperturbed particles in as regular a formation as possible. This is an attempt to reduce scatter in the signal caused by a wave impacting onto random distributions. In order for this to occur, we need the initial particle distribution to be hydrodynamically relaxed. This can be achieved two ways. The first is to allow the particles time (usually a few sound crossing times) to reassemble into relaxed state. This has the disadvantage of producing a disordered interaction face to the incoming wave, which increases scatter, and may hide important performance issues. In doing this, one would also have to relax each half of the computational zone independently. Whilst not a concern in itself, establishing the initial configuration after such a time would be difficult as discontinuities need to be resolved over at least a few particles. Exactly how this would be done to a disordered set without artificially introducing heat or other spurious artefacts is unclear.

The second option involves placing the particles into a pseudo-relaxed position. Using the Split-daughter boundary conditions and periodic boundaries introduced in Section 5.1, an initial equilibrium particle distribution with no hydrodynamical gradients will maintain its grid structure for a number of sound crossing times.

In order to implement these conditions, it is important that the initial particle distribution fits within the box evenly and smoothly, with no strain on the particle displacements.

We begin by looking at a 10:1 density shock.

The 10:1 variation in density corresponds (in 3-Dimensions) to a variation in particle spacing of $10^{1/3}$, or approximately 2.15. We must consider the

memory constraints of computational applications, and so attempt to run the simulation with as few particles as possible in the directions parallel to the shock front. With this in mind, we establish the diaphragm in the y, z plane at $x = 0$. The lower density region will have fewer particle across its bounds by a factor of 2.15, but we wish the initial particle spacings to fit within the box. We begin in the high density region, laying the particles down on a body centred grid of 15 particles square. The smallest integer ratio giving a particle density ratio of ~ 2.15 is 15:7 with a 0.33% misfit of the two grids. This does create a small amount of particle settling, but it is a small effect, noticeable as a minor increase in scatter near the boundaries, and one which we will not worry ourselves with too much. For a more specific application, more care could be taken in establishing better performing spacings.

The initial particle spacing is an input parameter, and is used to specify the resolution in the shock propagation direction, x . The y and z boundaries (periodic) are established by the number of particles across the initial grid spacing and this spacing parameter. The initialisation procedure involves laying down the grid from $x = -0.5$ until close the the origin, where the particle spacing smoothly changes across to the lower density values in the right hand region. The upper x -boundary is chosen to fit the grid in smoothly at around $x = 0.5$.

The smoothing procedure is a simple exponential, smearing any discontinuity out across a few particle spacings in the x direction. To accomplish this, we simply define the particle position to be in one of three zones, the left steady state (A_L), right steady state (A_R), or a region of transferral in between. These regions are defined by particle rank position, rather than x -coordinate. i.e. the smeared region of transferral extends for 10 particles either side of the origin in the x direction. This is because particles can only interact over a

certain range, controlled by their smoothing parameter, which is established by numbers of neighbours. It is therefore more logical to set up this region as numbers of particles.

The field variables $\langle A \rangle$ are then smoothed over the range by the rule

$$A(x) = \frac{A_L + A_R e^{\frac{x}{\Delta x}}}{1 + e^{\frac{x}{\Delta x}}} \quad (6.28)$$

where Δx is the particle separation chosen to maintain the relation between particles and desired density distribution. This is the same method used by [Chow and Monaghan, 1997] in their 1-D application where it worked quite well, and no need was seen to develop it further.

The number of baryons represented by each particle, ν_b , is the last variable attributed to each particle. The number of particles to be used in the simulation is decided by the desired resolution and the available memory. (see Section 6.4.2) After establishing the value for the field variable ρ , the baryon number density, the representative baryon number is chosen to satisfy this condition. This is a simple calculation, through taking a single bcc cell, with a volume $\Delta x \Delta y \Delta z$ and noting that it contains two computational particle (one in the centre of the box, and 8-octants in each corner).

6.4 SPH Performance and Results

It is important that any relativistic code, although optimised for relativistic flows, should be able to handle subtler conditions. With this in mind, we will reproduce firstly a non-relativistic shock flow. As one-dimensional shock tests have already been performed using SPH, we will not reproduce it here, opting instead to show the 3-dimensional algorithm's performance.

One-dimensional test cases of relativistic shock have also been performed, however they will be reproduced here as a comparative example. With only one spatial dimension, colliding particles always move along radial orbits, so the diffusion terms do not need to be as hard as they do in higher dimensions, to produce an equivalent amount of stopping force and avoid unwanted inter-particle penetration. By showing this 1-D case, and following with 3D equivalent simulations, one can see the interpenetration and diffusive performance differences more easily.

6.4.1 Diffusion Parameter Choices

Choosing the parameter K for the artificial diffusion terms may appear to be a black art, the choice of which is capable of hiding, or causing, a number of effects. For each initial conditions various values for the artificial viscosity parameter were trial, and used as a form of similarity check. Starting from values clearly too small, deduced from the code's inability to resolve the density shell at all and crashing due to the density jump becoming unresolved, through until the values calculated for each region reached some form of constant value. Particularly in the 1 dimensional studies, it was found that increasing K beyond this point would produce instabilities in the supposed quiescent states of the simulation. This was found to be a good indicator of how to set the artificial diffusion parameters. Closest results to analytical solutions were gained by using a parameter just smaller than that which would destabilise a static region, and around which small variations in it would not affect the density value for the five resolved regions.

Increasing the value of K from some initial estimate, until the value of the density shell in subsequent simulations stabilised, produced a value of $K =$

1.8. The maximum value which allowed correct modelling of the quiescent end regions was 2.1, with severe diffusion across the contact discontinuity. Consequently, the high-resolution 1D simulations use a diffusion parameter value of 1.85.

In 3-dimensional cases this distinction is no longer as clear, as it became possible to reduce the density of the post-shock shell through injudicious application of artificial viscosity schemes, something not able to be done in 1-D tests. Best results were still obtained by running a slightly smaller value than that which created spurious density spikes in the quiescent gas. These spurious effects appear as single particle's density jumping and is most likely caused by inappropriate time step controls. As they occurred only after increasing the amount of artificial viscosity above the sensible limits discovered by stable runs, these were not investigated further. Sensible results were gained for values of $K < 6$ which is very close to the 1-dimensional value raised to the power 3 ($1.85^3 \approx 6.3$) and the code was not found to be particularly sensitive to variations in K .

Some time was spent in analysing possible extensions to these terms, particularly in closely watching effects of fluid and sound speeds. Whilst these applications were fruitful in generating better numerical approximations to the analytical solution, there was no clear distinctions to be made in how to best choose the parameters. So they were abandoned for a lesser tuned version, but whose results could be trusted in the situation where there was no analytical solution to match the output data to.

6.4.2 Sod's Shock

Sod's shock tube problem is an accepted initial test for any hydrodynamical code. It is well understood, and replicated here. The initial states are described in Table 6.1 with an equation of state of perfect fluid, and a value of $\frac{5}{3}$ is used for Γ .

TABLE 6.1: Sod Shock Tube Test Conditions

	Left state	Right State
Pressure	1	0.1795
Density	4	1

The first dimension to be ascertained is the lateral grid $y - z$ dimensions. For this problem we will be modelling a 4 : 1 density drop, which corresponds to $(4^{1/3} \approx 1.58 \approx \frac{15}{10})$ to the nearest integer ratio. The physical y and z dimensions are left free, set by the condition that the x spacing covers a physical range of unity.

The initial placement of particles uses the body-centered cubic array, which appears as four particles placed at each vertex of a cubic cell, and one directly in the centre. This same configuration can also be described as a series of uniform square sheets in the $y - z$ plane, stacked along side each other in the x -direction. Each sheet has a regular, square lattice of particles on it, and every other sheet is off-set one half a spacing in the y and z directions. The particles were laid out on this body-centred grid, of 15×15 particles in the $y - z$ plane and the spacing of the $y - z$ sheets as an input parameter. This reduces to 10×10 particles in each sheet in the lower energy state and a similar increase in x spacing after the discontinuity at $x = 0$ is crossed. This corresponds to a density drop of 4:1 and yields an x resolution of ~ 1000 particles in the x -direction, modelling the region $-0.5 \leq x \leq 0.5$.

The actual numbers of particles filling the computational box in each direction (x, y, z) is

$$627, 15, 15 : 353, 10, 10$$

$$141180 : 35295 = 176475 \text{ particles}$$

This highlights the memory difficulties of these sorts of tests. Ultimately, we would prefer most particles to be situated around the shockfront to resolve the rapid acceleration experienced there, but are resigned to have most memory used in allocating particles to the higher density, yet relatively quiescent, high density state. This problem is only exacerbated in later tests of 10:1 density drops.

The three waves mentioned in Section 6.1 can clearly be seen in Figures 6.2 through to 6.4. The rightward travelling shock and the leftward rarefaction fan are clear characteristics. Note how the contact discontinuity marks the end of the shocked density shell, but is unidentifiable in the pressure plot. The clear difference in the rarefaction waves, as against the shock wave are also clear in these images.

If we plot the velocity curve, we can immediately see the rapid acceleration of the shock front, but also notice much more scatter in the plateau regions. This region appears to be poorly resolved, although whilst there is a large amount of scatter, the average velocity of this region is correct. This scatter is a direct result of post-shock ringing or oscillation, and can be controlled by increasing numerical diffusivity, but has been included nonetheless to highlight the difference in scatter between different variables.

The Plots shown in Figures 6.2, 6.3, 6.4 and 6.5 are created by sampling down the middle of the computational box. This is done simply to reduce the




FIGURE 6.2: 3-dimensional simulation of Sod's Shock tube showing Baryon Number Density against x location




FIGURE 6.3: 3-dimensional simulation of Sod's Shock tube showing Thermal Energy against x location




FIGURE 6.4: 3-dimensional simulation of Sod's Shock tube showing Pressure against x location




FIGURE 6.5: Sod's Shock tube Velocity distribution, showing considerable scatter in the plateau regions. The noise on the righthand boundary is an errant Boundary condition not affecting the central simulation

actual number of particles in the image, and no selection process is carried out. All plots show the same individual particles. Some of the plots show regions where there appear to be no particles. This is a result of this method of plotting, which ignores particles outside of some radial distance from the x-axis. These gaps can be removed by simply increasing this sampling radius, usually by a small amount, to capture the 'missing' particles in the plot. Increasing the radius also increases the total amount of data within the plot, increasing the size of the image file unnecessarily. Figure 6.5 also shows an acoustic wave coming in from the boundary. This is caused by a mismatch between the particle spacings and the dimensions of the computational cell in the lower density regions. This non-continuance of the initial particle placement grid is exacerbated at the corners, where two periodic boundary conditions meet. For this application, the particles spacings should differ by $4^{1/3}$ or ≈ 1.587 . The grid ratio of 15 : 10 yields a 6% discrepancy or strain on the particle placements, which is enough to cause this wave on the boundary. For later simulations, this particle spacing to computational cell mismatch is much reduced, and these waves do not appear.

6.5 Relativistic Shock Wave

The first relativistic example we will employ is the shock tube studied in detail in one-dimension by Hawley et al. [1984], Schneider et al. [1993], Marti and Müller [1996] and in three-dimensions by Siegler and Riffert [1999]. This is similar in its initial conditions to the Sod tube shown previously, but more extreme. As before, the higher density and pressure state is to the left, and lower energy states to the right. When the two states are brought together, a shock wave propagates into the low density gas to the right, and

a rarefaction fan propagates leftwards into the higher density state. Between these two waves, there is a contact discontinuity where the pressure and velocity are continuous, but the density drops. In order to aid comparison between one and three dimension simulations, the 1-dimension outputs in Figures 6.6 through to 6.9 have been included.

TABLE 6.2: Relativistic Shock Tube Test

	Left state	Right State
Pressure	13.3	0.000001
Density	10	1

The artificial viscosity was the largest value that would allow stable quiescent regions. This is dependent upon the number of particles, and for this particular configuration came to be ~ 1.8

6.5.1 3-D Shocks

We have outlined previously the significant increase in demand placed upon the dissipation terms in the extension into higher spatial dimensions. There is a lot more opportunity with the added degrees of freedom for the particles to interpenetrate, and for artificial or false communication to occur. Therefore we would not expect that what occurs in 1D necessarily carries over into 3D.

To highlight this fact, we show two shock simulations, one calculated in a one-dimensional code, the other 3-dimensional application. The greatest differences between the one and three dimensional performances are evidenced in the baryon number density plots which are shown in Figures 6.10 and 6.11.

The images shown in Figures 6.10 and 6.11 show the two codes after 2000 timesteps. The code on the left is a one dimensional application, with a total of 3900 particles. This corresponds to ~ 350 in the lower density region. The




FIGURE 6.6: 1D simulation of the Relativistic Shock showing Baryon Number Density, Distribution




FIGURE 6.7: 1D simulation of the Relativistic Shock showing Pressure Distribution




FIGURE 6.8: 1D simulation of the Relativistic Shock showing Thermal Energy Distribution




FIGURE 6.9: 1D simulation of the Relativistic Shock showing Velocity Distribution




FIGURE 6.10: 1-D baryon number density versus x location, with identical particle resolution in the x direction as Figure 6.11




FIGURE 6.11: 3-D simulation showing baryon number density versus x location, with identical particle resolution in the x direction as Figure 6.10

right hand graph is the 3-dimensional simulation, with $\sim 190,000$ particles, which also corresponds to ~ 350 sheets in the x direction.

$$\begin{aligned} x_l, y_l, z_l : x_r, y_r, z_r \\ 771, 15, 15 : 354, 7, 7 \\ 173472 : 17347 \text{ particles} \end{aligned}$$

The issue raised in the study of the non-relativistic shock, with regards to memory allocation is made again here. Most particles in these tests are required in the non-active, or quiescent, regions of higher density, making this a very memory inefficient test for this application. One would like to use more of the particles in resolving the shock front, and surrounding regions, but this would require variable particle baryon number representations, v_b . This is certainly something which would benefit from concerted study in the future. In fact, this modification has been implemented in the MHD work of Borge et al. [2001] with promising results.

Immediately apparent is the difference in diffusion of both the shock front and the contact discontinuity. The apparent increase in sharpness of the three-dimensional shock is due to the initial smearing application in the particle's initial set up routine. As has been mentioned, the smearing is calculated by a function of the right and left steady states, and an exponential involving the particle separation. In the one dimensional case, the particle spacing changes by the full factor of 10 as one goes from the high density to the low density state. In three dimensions, this is only $10^{1/3}$, or ~ 2.15 in the x direction. Therefore, in order to resolve the shock over ~ 12 particles, as it is in both of these cases, the one-dimensional application resolves the shock over a larger spatial distance. It should be noted though, that merely counting the numbers

of particles across a shock front is a poor guide to the relative performance of the codes, as both the particle spacings and smoothing lengths are completely dynamic variables.

The diffusion evident on the contact discontinuity in 3-D is acceptable, although clearly more than that seen in the one-dimensional case. This is due to the large amount of diffusion required to inhibit post-shock ringing and over estimation of the shell's density. Interestingly, or perhaps co-incidentally, in order to gain similar resolution and smoothness, the 3-dimension diffusion parameter is very close to that of the one-dimension to the power of three.

It is suggested from similar tests in one dimension, that the diffusion over the contact discontinuity is reduced by increasing the number of particles, unfortunately an impossibility at the moment due to memory restrictions on single processor machines.

6.5.2 Effects of the Number of Neighbours

We can examine the extra diffusion caused by increasing the numbers of desired neighbours. As shown in Section 6.3 the algorithm calls for the input of a 'desired number of neighbours,' which is used to control the evolution of h , the smoothing length. Increasing this number forces more particles to be included in the particle summations, which smears out shocks, and equally the postshock oscillations. Reducing this number will increase the resolution of a shock wave, but be less likely to diffuse oscillations.

Except for variations in the number of desired neighbours in the particle summation loops, the three simulations of relativistic shocks shown in Figures 6.12, 6.13 and 6.14 use identical set-up algorithms and parameter configurations.

The most prominent feature immediately noticeable is the difference in the scheme's ability to capture the shockfront. In the upper image, where the desired number of included neighbours is only 35, the value of the peak of the density spike is overestimated. The simulation still produces a reasonably smooth profile however, with minimal oscillations. It is this ability which means one needs to run convergence tests to ensure that the simulation is producing correct results.

The middle images shows the scheme correctly capturing the peak value, using 57 nearest neighbours. This has introduced a minor oscillation in the plateau region immediately behind the shock. The shock front itself is also more diffused.

The final image highlights how the peak value stabilises to that of around the analytic solution. The extra diffusion introduced by summation loops of 80 neighbours smooths out the post shock oscillations completely, but also smears the shockfront noticeably.

The second difference is the resolution of the rarefaction fan. With less than optimum numbers of neighbours, one can see that the head of the fan is poorly resolved, with too much apparent diffusion. It is also clear that there is some density wave within the rarefaction fan. Looking carefully at this particular run, it is noted that this wave travels as a downwind soundwave, i.e. travelling at the speed $\frac{v+c_s}{1-vc_s}$ where v is the fluid velocity (accelerating to the right) and c_s is the local sound speed. The soundwave originates at the origin at $t=0$, and is possibly the initial grid mismatch leading to a perturbation at the initialisation of the shockfront. With added neighbours (and the inherent diffusion) this peak is dissipated into the surrounding fluid. Note that the integral $\int_{x_1}^{x_2} \rho dV$ is maintained and correct. It is this conservation which enforces the incorrect rarefaction head in the low dissipation models.

As the numbers of neighbours increases, and this soundwave is diffused more and more, the rarefaction fan becomes increasingly better defined. However, there is a small glitch noted in the 80 neighbour simulation at the rarefaction tail. It is believed that this is caused by the h modification algorithm acting too aggressively in its control of the smoothing length as the density begins to increase again and the fluid stops accelerating.




FIGURE 6.12: Varying the desired number of neighbours: 35 nearest neighbours included




FIGURE 6.13: Varying the desired number of neighbours: 57 nearest neighbours




FIGURE 6.14: Varying the desired number of neighbours: 80 nearest neighbours

Chapter 7

Relativistic Heavy Ion Collisions (RHIC)

One of the important experimental applications of relativistic fluid dynamics is that of Heavy Ion Collisions. This is the practice of accelerating atomic nuclei up to very high speeds, before colliding them into atomic targets of some description. The resulting spray of subatomic debris is then analysed to deduce information about the nuclear interaction. This chapter opens with an historical overview of these accelerators and their capabilities. Section 7.2 then looks at the basic events that constitute a nuclear collision, and this is followed with an analytical look at ways these processes might be modelled (Section 7.3). It has been discussed previously how collisional studies use the equation of state as an input, and this concept is looked at in more detail in Section 7.4. Section 7.5 introduces what hydrodynamical phenomena one would hope to see in a well performing model, which leads us to examining some of the advantages to using a particle method when studying collisions (Section 7.6).

Having fully established what precisely one would expect to see in a hydro-

dynamical simulation of a nuclear collision, the following sections examine algorithm specific topics, such as establishing the initial conditions (Section 7.7) and how the final outputs and performance assessment of Section 7.9 are carried out (Section 7.8).

7.1 An Historic Overview

All matter consists of atoms, the basic building block of the universe. If we look only at the nucleus, ignoring the easily disassociated electrons, then we see a collection of neutrons and protons, generally referred to as nucleons. Current theory states that each of these nucleons is a bound grouping of three of the most fundamental subatomic particles, quarks. The positively charged protons are made up by two up flavoured quarks and one down, whereas the neutrons consist of two down and a single up. These sub-nuclear bags of quarks are held together by the intermediary mesons known collectively as gluons.

Being fundamentally quantum particles, the available energy states for any trinity of quarks will depend entirely upon the energy density of the state it finds itself in. For lighter nuclei, this means that the nucleons maintain some 'personal identity,' occupying distinct energy states. As the mass of the nucleus increases, the nucleons become packed more closely, increasing the number of available states and simultaneously decreasing the spacing between them. Eventually, the available states become so numerous and so closely spaced that the nucleons lose their discrete qualities, and behave approximately as a fluid. This coherent state is often referred to as 'nuclear matter.'

Up until nuclei with a nucleon number $A \sim 210$ (somewhere around lead

and bismuth) these configurations are stable. Above this stability limit, the Coulombic repulsion of the like charged protons cause the nuclei to undergo fission. However, sufficiently large numbers of nuclei ($A \sim 10^{57}$) can be stabilised by gravity, allowing the creation of Neutron Stars and other compact objects. [Stock, 1989]

Understanding the behaviour of this fluid is important to accurately model compact objects, supernova explosions (particularly the core-bounce scenario) and even the nature and consequent evolution of the Big Bang. The principle piece of information required for these models is just how stiff nuclear matter is. That is, how quickly does the pressure rise as a function of an increasing density? A Type II supernova may experience number densities upwards of 0.6 fm^{-3} , four times the rest density of a nucleus ($0.15 \text{ baryons per cubic femtometer (fm}^{-3})$ corresponds to $\sim 2.7 \times 10^{14} \text{ g cm}^{-3}$) [Stock, 1989]. The interior of neutron stars may be higher still. Clearly some form of experiment was needed to get to these densities and deduce how nuclear matter behaved.

It was not until 1974 when the Lawrence Berkeley Laboratory's Bevalac became the first machine capable of replicating the densities required to explore the nuclear equation of state. Fundamentally a synergy of the Bevatron synchrotron and a linear accelerator, the Bevalac accelerated a beam of ions (originally only capable of up to $A=40$) up to energies of 2 GeV per nucleon. This beam of accelerated particles then was allowed to impact into a stationary target [Gutbrod and Stöcker, 1991]. Basic improvements over time allowed the Bevalac to use beams of nuclei with nucleon numbers up to that of uranium, although the analysis of the debris from a collision event was still basically a photographic system until the development of the Plastic Ball Spectrometer.

This was developed by Poskanzer, Ritter and Gutbrod, and was capable of

measuring the mass, charge, energy and angle of emission at a high enough sampling rate to collect nearly all the particles from a succession of events. This instrument showed that nuclei were indeed capable of stopping each other and would not necessarily tunnel through each other. This stopping effect was indicated by the complete disappearance of projectile and target nuclei, which was replaced by a pile of nuclear debris at rest in the centre of mass frame [Gutbrod and Stöcker, 1991].

Evidence gathered from these experiments suggests that not only is the equation of state for nuclear matter quite stiff, but may have a number of metastable higher energy states, or even multiple phase transitions. One of these proposed phases is the exotic quark-gluon plasma, whose existence indirectly allows for the existence of 'quark stars,' a new class of compact objects.

7.2 Phases of a Nuclear Collision

The physical description of a nuclear collision can be divided into four distinct phases or time periods. The first of these is the initial contact and heating phase, where the projectile nucleons first transfer some of their kinetic energy into the target nucleus. As these motions are well above any conceivable local sound speed, shocks propagate outwards from the initial point of contact. The finite size of target nuclei means that these shockfronts will be quite complex curves, and not describable in terms of simple curves of revolution.

When a nucleon first impacts another, considerable amounts of its kinetic energy are lost. Goldhaber and Busza (quoted in Wong [1994]) refer to this as the 'nuclear stopping power.' Nucleon-nucleus and nucleus-nucleus collisions reveal that a substantial amount of the projectile nucleon's initial ki-

netic energy is lost through this nucleon-nucleon collision interaction [Wong, 1994]. This energy is converted into the production of large numbers of particles (mostly pions). Simply put, the initial kinetic energy, in the form of longitudinal momentum, is converted into energy of the hadronic matter in the centre of mass of the system. This creates a fireball (the second phase), the magnitude and degree of which depends on the equation of state, and its ability to provide nuclear stopping.

The fireball zone translates much of the initial energies into transverse and longitudinal motions which drive the third phase, a rapid expansion. This phase is not seen in nucleon-nucleon collisions, implying that these systems behave far from that of a macroscopic system capable of being understood with a thermodynamic model [Hung and Shuryak, 1998]. However, more of this collective flow is seen as the energies and mass of the colliding nuclei increase. This implies that macroscopic behaviour increases with the energies, and with macroscopic behaviour comes the ability to understand the collisions in terms of thermodynamics. Much of the information from the previous fireball phase is lost through scattering events during this phase of thermal expansion.

The final phase is the chemical freeze-out. Once the density drops to less than around 30% of the rest density, the fluid disassociates into its constituent, non-interacting hadrons. The exact temperatures and density this occurs at is dependant upon the individual species created in the collision and their individual response to their density and temperature environments. This phase then becomes a bit of a bug-bear for the theoretical modellers and most use some form of simple 'freezeout' temperature, T_f , below which thermal and chemical equilibrium are assumed. [Teaney et al., 2001]. Naturally, at this stage, any hydrodynamic scheme will become inappropriate.

7.3 Modelling a Nuclear Collision

The choice for theoretical modelling of such a complicated event such as these nuclear collisions can be broadly broken into two methods. If the ratio of mass to kinetic energy of the nucleons is small, one can dispense with the relativistic methods completely, and use a quantum dynamical, non-relativistic approach. This will involve the solution to a many-body Schrödinger Equation, which would require many other assumptions and simplifications.

If, however, the momenta are significantly higher, then the de Broglie wavelengths are small. This allows quantum effects to be ignored, and a relativistic theory can be applied. Relativistic theories capable of being applied to extremely relativistic collisions all fit within 4 groups:

Cascade Models

These models approximate the collision as a series of two-body collisions assumed to be in free space. The resultant scattering is naturally restricted by Pauli exclusion principles, but other than this, the collisions remain distinct. This approximation is really only valid if the mean free path of the colliding nucleons is much larger than the range of the nuclear force, and due to the assumption of free-space collisions, can not reproduce any collective flows. No assumptions need be made about local thermal or chemical equilibrium, but the cross-sections for each possible interaction need to be measured experimentally and used as a basic input. These methods can however accurately depict non-equilibrium effects, such as those after the thermal expansion phases, where there is no time for any form of equilibrium to occur. These methods will also predict the freeze-out surface well, where the density and coupling are greatly reduced.

Hydrodynamics

At the other end of the spectrum to the cascade models, are hydrodynamical models. The requisite assumption of there existing either a large interaction (smaller interactions are resolved with the cascade models) or small bombarding energies is the most contentious for the application of this model. Amsden et al. [1977] contend that the mean free path required to stop a nucleon at 250 AMeV (250 MeV per nucleon) is 3 femtometer (fm) compared to the mean free path for nucleon collisions at normal nuclear density of ~ 1.7 fm. This means that the ability of a hydrodynamical code should be adequate at lower bombarding energies, but as the bombarding energy and the stopping distance increase, its performance should deteriorate. As is also pointed out in Amsden et al. [1977], these simple assumptions ignore the effects of coherent collective fields, such as pion condensates, which would further reduce the stopping distance. Despite this argument, later hydrodynamical codes used to study the ion collisions, such as Aguiar et al. [2000], have shown good results.

The other assumptions inherent to applying a general fluid algorithm to this application are that there is enough time for the fluid to reach local thermodynamic equilibrium and there are many degrees of freedom available. The second of these is safe, in that initially there are 200+ nucleons, which rapidly become a soup of many more subatomic particles.

In terms of the possible extent that thermal equilibrium could be reached, we can look at the time of the collision, which is of the order of the size of a nucleus divided by the speed of light $\sim \frac{1 \times 10^{-15}}{3 \times 10^8} \approx 10^{-23}$ s. An exchange of a pion between two nucleons requires $\sim 10^{-24}$ s [Amsden et al., 1977] or a tenth of the collision time.

The hydrodynamical codes reproduce the macroscopic behaviour of the nuclear fluid, such as the fireball phase and resultant collective motions. They cannot deal with times after the freeze-out phase, as the fluid is completely decoupled and so thermodynamical treatments become inappropriate.

As opposed to the cascade models, the hydrodynamical models require an equation of state as a basic input.

Two-Fluid Hydrodynamics

The Two-fluid approximation attempts to combine the two previous methods, by maintaining the bulk flow properties of the hydrodynamical models with some of the cascade effects. This is done by solving the equations of fluid motion, but with added terms coupling the target and projectile through energy and momentum transfer as they penetrate each other. This is done through a drag function, which takes into account the cross-sections of the cascade models. For a detailed synopsis and presentation, refer to Amsden et al. [1978].

Hybrid Models

The last group are the most recent and use thermo-dynamic modelling to reproduce the bulk motion of the impact. This supplies the basic fluid properties of momentum, density and thermal energies, which are then used as an input for a cascade model used to model the reaction from the expansion phase onwards [Teaney et al., 2001].

7.4 The Equation of State as an Hydrodynamical Input

The fireball stage and its resultant spray of material is entirely dependent upon the compressibility or nuclear stopping power of the nuclear fluid, which is a function of the equation of state. Unlike the cascade models, which use cross-sections as their inputs, a hydrodynamical code requires this relation between pressure and mass-energy density as its input.

Following the paper of Nix et al. [1982] we identify that in the rest frame of the nuclear fluid, the total energy per nucleon will be given by the ground state energy per nucleon, ϵ_0 , plus a thermal term

$$\epsilon(\rho_0, \mathcal{T}) = \epsilon_0(\rho_0) + \mathcal{T} \quad (7.1)$$

where \mathcal{T} is the thermal contribution and is a function of the number density.

The pressure is then given by

$$P = \rho_0^2 \left\{ \frac{\partial \epsilon(\rho_0, \mathcal{T})}{\partial \rho_0} \right\}_S \quad (7.2)$$

$$= \rho_0^2 \left\{ \frac{\partial \epsilon_0(\rho_0)}{\partial \rho_0} + \left(\frac{\partial \mathcal{T}}{\partial \rho_0} \right)_S \right\} \quad (7.3)$$

with S denoting derivatives performed at constant entropy per baryon. [Amsden et al., 1977]

The ground state energy is given by

$$\epsilon_0 = a \left\{ \frac{\rho_0}{\rho_0} \right\}^{\frac{2}{3}} - b \left\{ \frac{\rho_0}{\rho_0} \right\} + c \left\{ \frac{\rho_0}{\rho_0} \right\}^{\frac{5}{3}} \quad (7.4)$$

with $a = 19.88 \text{ MeV}$, $b = 69.02 \text{ MeV}$, and $c = 33.46 \text{ MeV}$. Taking the rest baryon

number to be 0.1346 fm^{-3} , we can then deduce

$$P = \left\{ -\frac{2}{3} \left(\frac{\rho}{\rho_0} \right) - \frac{1}{3} \left(\frac{\rho}{\rho_0} \right)^2 + c \left(\frac{\rho}{\rho_0} \right)^{\frac{2}{3}} \right\} \rho_0 + \frac{2}{3} \epsilon \quad (7.5)$$

[Harlow et al., 1976]

As has been previously noted, the disregard of all terms to do with Coulombic attraction and surface forces experienced by the nuclear fluid, mean that this is by no means a complete simulation. Whilst these energies are small relative to the kinetic energy of the collision, it does preclude an accurate representation of the coalescence of nuclear debris into particles after the fireball. Making assumptions about equilibria and nuclear reaction rates and using information from the temperature and baryon number densities, one could interpret the momentum distribution to deduce theoretical particle spectra from these simulations. These could be used as a comparison against the experimental data in an effort to refine the equation of state. This, however, is beyond the scope of this thesis.

In using this form of the equation of state, we have also neglected the production of particles through the nuclear stopping interactions, and the resultant radiative energy losses. Whilst this could be accommodated in a later version of this algorithm, it is again beyond the scope of this work.

7.5 The Hydrodynamic Phenomena of RHIC

As has been pointed out in the discussions on equations of state and their applications here, the performance of this algorithm in its current guise will be purely phenomenological. A dedicated RHIC SPH application would use more appropriate variables and coordinates, and certainly a more complex

EoS. Even the seemingly safe choice of conservation of baryons used to establish the evolution equations in this work, could fail in the fireball phase of a heavy-ion collision. Here, thermal pion production could result in a null-baryon number, pion dominated phase region, and a completely inappropriate fluid description. [Aguilar et al., 2000]

By our choice of equation of state and its lack of particle production phases, we are avoiding this eventuality, but at the same time condemning ourselves to be unable to resolve any effects dependant upon this transition. The ability to convert the longitudinal momentum into hadrons affects the so called compressibility function of the nuclear matter, its ability to perform nuclear stopping reactions. By deliberately excluding these effects, any conclusions to be drawn from the fireball must be seen purely as qualitative, with little predicative power. In other words, whilst compression will produce a fireball region, our ignoring the effects of phase changes, particularly the generation of a quark-gluon plasma, causes its magnitude and long term effects to be ill-defined.

For a similar reason, determining the freeze-out surface from this algorithm could be done, with experience in reducing our primitive variables into species spectra. However, one expects that this surface would be at best a very complicated structure and highly dependant on the nature of the fireball and the amount of energy it deposits into the fluid. By knowingly simplifying this early fireball stage, any follow on effects such as the freeze-out surface, would also be greatly affected.

These effects are due to the choice of the Equation of State which we have made, and the desire to maintain as general an algorithm as is possible. There is no obvious reason why more complex equations of state may not be applied, and the resultant flows studied in more detail by later variations of

this algorithm. (See Section 7.6)

These choices do not detract from the ability of the code to model the large collective flows seen in heavy projectile impacts. Although the exact qualitative amount of energy deposited is unresolvable, conservation of energy is maintained. So there is a clear energy increase in the central regions, which will manifest itself as a thermal expansion phase. The overall response of the nuclear fluid will also exhibit compressional heating and expansion effects, with resolvable shock fronts, all effects of relativistic hydrodynamics, as opposed to specific nuclear dynamics.

The collective flows mentioned above are the macroscopic motions of nuclear matter during a collision, and as such are handled best by a hydrodynamical description. The debris from a nuclear collision can be described by 'in-plane' motion, or 'out of plane'. Here the plane referred to is the reaction plane, deduced by sphericity analysis or similar experimental reduction, and is basically the plane defined by the projectile's longitudinal momentum before and after impact. Note that for head-on collisions, this momentum vector will maintain a constant direction, and so there is no unique reaction plane. This is what allows 2-dimensional codes to model heavy-ion collisions. However, once an impact parameter is introduced, the loss of symmetry provides a before and after momentum vector, defining the reaction plane.

In the ejecta after the collision, there are three main flows detected by experiments. Two of these are in the reaction plane, and a third is perpendicular to it. Within the reaction timeframes, the effect of nuclear stopping and regional heating should also be apparent. These two closely linked effects also necessitate a distinction between nuclear particles to be drawn: 'participant' and 'spectator' particles. Put simply, nuclear matter which has experienced the fireball and been exposed to these pressures can be said to be participants.

Much of the fluid will have been slowed, stopped and then its momentum reversed by the hardness of the colliding nuclei. This fluid will still be involved in the subsequent re-expansion, but will not have undergone any reactions or processing within the fireball region. This is the spectator fluid.

Side-Splash

The first of the collective, post-collision motions contained within the reaction plane is known as 'side-splash'. As the name suggests, side-splash is an asymmetric outflow of participant fluid, spraying out from the fireball region in an azimuthal distribution.

Bounce-off

Having collided and transferred energy to target, some projectile spectators will be given a side-ways push from the collision zone, and continue along upon new, deflected paths. This flow will also be contained within the reaction plane.

Squeeze-out

The final collective flow is perpendicular to the reaction plane, and occurs early on in the reaction. It involves participant fluid, shock-heated in the fireball region, being squeezed out laterally. It is this fluid which, experimentally, can reveal the most information about the nuclear equation of state at its higher densities. This fluid is ejected from the fireball region and almost immediately freezes out, without passing through an interacting region of hadrons. It is this fluid which can give a clear and distinct signal from the fireball region to the detectors.

Regional heating and nuclear stopping

The latter of these effects is perhaps a little obvious. Somehow the material must have clear regions of unmixed fluid (spectators which will be picked up on the experimental detectors). The nuclei must become compressed enough to stop the interpenetration of the target and projectile. (Although at high energies and light ions, this effect becomes less distinct as the quantum effects of the individual participants becomes more apparent and the nucleons become more transparent to each other)

This stopping results in regions of heating, which are currently used as the initial condition for many nuclear codes. Other than the central fireball, consisting entirely of participant fluid, two spectator regions of high density, shock heated by the fireball and the mass of the nuclei colliding, should also appear. This appears as three ellipsoids of high density described in spatial order along the beam axis as : the core of the projectile (spectator), the fireball region (participant), and the core of the target (spectator).

7.6 Advantages of Using a Particle Method for RHIC Modelling

Neglecting the issues raised before about the complex nature of the nuclear flow, there are a number of reasons one would use a particle method for these problems over a Particle-In-Cell (PIC) or finite difference scheme.

The reliance of many grid schemes on exact or approximate Riemann solvers means that their application to Heavy Ion collisions is compromised. Riemann Solutions for these kind of equations of state are poorly understood, and this becomes exacerbated in the three dimensions required to adequately

model collisions.

The issue of efficiency of memory use for 3 dimensional calculations has been raised, and the application to these sorts of collisions highlights the point. In order to gain useful insights and to be able to compare results with experimental detectors, the models must be able to run through the thermal expansion phase until freeze-out. This means that for much of the collision, a large empty space must be modelled if one is using a finite difference, or grid-based scheme. This memory demand must be weighed against the need for fine resolution in the fireball phase, to adequately model all interactions in the region of compressed matter. An SPH application automatically distributes its particles into this dense region, increasing resolution, and if used in conjunction with a rank-space neighbour algorithm (Section 5.5), will not waste memory in empty spaces. In fact, no extra memory is required as the particles expand outwards to the physical size of the detectors, as the rankspace datacube remains a constant size.

Another issue is correct modelling of the freeze-out surface, across which there is a transition from fluid to individual particles. This occurs when the baryon density drops so low as to preclude any more interactions between the nucleons. With a grid method, this will require some form of triangulation and interpolation between nodes to determine where to turn off nuclear interactions. For a full explanation of the amount of computational effort required to perform this task, see the appendix of Hung and Shuryak [1998]. With an SPH algorithm, the same can be achieved much more efficiently by attaching a numerical tag to each particle, which dictates whether or not the packet of fluid depicted is chemically active or not. The rule for this chemical inertness could be as simple as the rather naive: Once the baryon number density drops to $\sim 0.2 \text{ fm}^{-3}$, the particle is understood to have passed through the freeze-out

phase, or as complicated as one needs, using the thermodynamic variables to predict a species spectrum for the represented fluid, and controlling the evolution with regards to freeze-out through this. Visualising the freeze-out surface is also simplified, as one would have knowledge of where each particle was and its momentum when it passed through this phase, giving an instant picture of the surface and its evolution.

The final benefit comes from being able to model a full collision interaction. The grid distortion or remapping required to model the pre-contact phase through to the fireball mean that many current codes use a separate generator to produce a pre-fireball energy distribution [Aguiar et al., 2001], and then evolve this through the expansion phases. There is no need to do this with a particle code, as the grid entanglements and associated difficulties caused by colliding or reconnecting flows simply do not exist.

7.7 Initial Conditions

In accordance with different particle accelerators, the velocities of the projectile particles need to be deduced from the energy per nucleon classification of the accelerator. This is frequently given in the form, for example, 200A MeV, which corresponds to an energy of 200 MeV per nucleon. Having deduced the required projectile velocity, the decision is made as to the computational frame which will be used. Earlier simulations of these collisions had the target nucleus stationary in the lab-frame. However after the development of the Plastic Ball Spectrometer, it became more common to study the centre of mass frame.

The first models run are moderate collisions between ^{20}Ne and ^{238}U at energies of 250 and 400 MeV per nucleon. This corresponds to the energies of

the colliders available in the early 1980's.

Energy	Lab Velocity	γ
250AMeV	0.6149c	1.26
400AMeV	0.7219c	1.44
2.1AGeV	0.9516c	3.25
5.0AGeV	0.9876c	6.36
11.0AGeV	0.9969c	12.71

TABLE 7.1: Impactor speeds for stationary target for given energy per nucleon

Energy	Lab Velocity	Impact Velocity (approach speed)
2.1AGeV	0.719c	0.9516c
11.0AGeV	0.924c	0.9969c

TABLE 7.2: Impact velocities for centre of mass frames for given energy per nucleon

The basic nuclei are established in their own rest frame by first laying down a cube of particles with specified spacing and side length of $2r_0$. The value of r_0 is found by a simple calculation assuming the the number of nucleons are known (which atomic species is being modelled) and the rest baryon number density, taken to be 0.1346 fm^{-3} . This initial cube is then cleaved, whereby any particles greater than r_0 from the centre of the cube are discarded. We are then left with a sphere of radius r_0 consisting of N particles laid out in a body-centred grid structure.

Knowing how many particles are in this representation, and the number of physical baryons (A) allows us to allocate the baryon number, v_b , to each particle and its number density variable and smoothing length value. All the thermodynamic variables can then be allocated knowing the equation of state and basic nuclear matter rest values.

We now have a ball of particles representing a given nuclei at rest and equilibrium. The laboratory velocity, chosen as an input parameter to represent the given energy of a collision, gives the special relativistic contraction factor, γ . (See Table 7.1) This is required because the nuclei are moving with respect to the laboratory frame, and so will undergo length contraction in the direction of motion. We chose the x direction for the beam axis, along which the projectile will travel, and the y axis for any offset if a non-central collision is simulated. As a consequence of this, only the x-axis will undergo contraction. To perform this operation, the coordinates for each ball (spherical in its rest frame) are contracted by a factor γ in the x direction, resulting in the original sphere appearing in the laboratory frame as an ellipsoid, circular in the y-z plane, and foreshortened in x.

Some simulations of higher mass particles are calculated in the centre of mass frame, where initially both the target and projectile are moving. In this case, both nuclei need to be length contracted and have their computational variable adjusted to the laboratory frame. These cases are shown in Table 7.2 for common energies per baryon.

To this new configuration, we attribute the pre-defined velocity and momenta variables, and adjust the baryon number density by γ such that it corresponds to the computational frame rather than the nuclei's rest frame.

The final stage in the setting up of the collision simulation is to position the nuclei such that they are close to each other on approaching trajectories. This is a simple translation operation placing the centre of mass of the nuclei at $\pm(\frac{1}{\gamma}r_0 + 3h_0)$, where r_0 is the rest frame nuclei radius, and h_0 is the attributed smoothing length. A slide showing this phase is shown in Figure 7.1. Initialising the projectile so close to the target nucleus is done for the simple reason that we do not wish to use computational time modelling the flight of the nuclei




FIGURE 7.1: The initial configuration showing the length contracted projectile. ^{20}Ne at 250 MeV colliding with ^{238}U target

through free space.

As the contraction is only experienced in the x-direction, it is arguable that a spheroidal kernel would offer the best resolution in this case. It has already been explained why these kernel types have not been employed, but we are left with the choice of smoothing length value to be used in the ellipsoid representation. We have chosen to leave its value as that used to define the original sphere of particles. Consequently, the particles will be exposed to more neighbours than they were in the sphere, and possibly information from regions of the nucleus which they should not be able to communicate with. Fortunately, this problem is circumvented by the form of the sound speed we are employing, which takes all the particles' velocities into consideration and disallows acausal information travel. The particles will still have too many particles in their summation loops, and as a consequence will adjust their smoothing lengths in an attempt to rectify this. This is particularly evident for

the particles on the surface of the sphere, who are limited in the numbers of neighbours they can physically have access to. This means that the projectile nucleus is not in a relaxed state, and so to avoid unnecessary adjusting of smoothing lengths and positions pre-collision, we place the nuclei close to one another at the start of the simulation.

7.8 Data Visualisation

As well as the restrictions placed upon higher dimensional modelling by the requirement of computer memory and speed, visualisation of the data sets can also become a problem.

In an effort to produce images of a manageable size yet with enough detail, a central slice of data from the reaction plane and of a few particles in width will be shown, with all z locations of the particles projected down onto a now 2D slice in the $x : y$ plane. As the x axis corresponds to the incoming trajectory, and any non-central impact parameter results in a shift in the y direction, this simple projection allows the major features of the modelled interaction to be seen. The desired variable is plotted using colour, delineated by a linear binning of the range in question. This is done by taking the maximum and minimum values of the desired variable, dividing this evenly into the resolution range required (i.e. 12 ranges) and then placing each particle into one of these bins, depending on the value of its variable. These bins are then allocated a colour, through the spectrum with black and white at the extremes, and the particles are plotted on the $x : y$ projection slice in their appropriate colours.

The original plotting routine created each particle by a point, then a number of concentric rings around it. Many of these rings are evident in the printouts

and have nothing to do with the simulation. They are merely a printing artefact.

7.9 Performance and Results

7.9.1 Neon : Uranium Collisions

The collision between a ^{20}Ne nucleus and a ^{238}U nucleus was used as a test bed for Nix and Strottman [1981]'s Particle-In-Cell method for relativistic hydrodynamics, and also for Amsden et al. [1978]'s two-fluid PIC model. It is presented here as a 3-dimensional SPH calculation for various impact parameters (something unable to be done adequately with less than three dimensions).

Figure 7.2 shows thermal energy contours from the central reaction plane region for a collisional impact parameter (b) of 0.1 and projectile energies 2.1 AGeV, plotted as described in the preceding section. The slide show the particles after 150, 300, 450 and 600 timesteps, which correspond to $1.5 \text{ fm}c^{-1}$, $3 \text{ fm}c^{-1}$, $4.5 \text{ fm}c^{-1}$ and $6 \text{ fm}c^{-1}$.

The important features to note are firstly the curved bow-shock structures penetrating into the target nucleus. The projectile is flattened out into a disk-like structure, both through Lorentz contraction ($\gamma = 3.25$ from Table 7.1), and its impact compression. The second slide shows the central fireball region due to the shockheating, as well as the still defined projectile nucleus. This frame also clearly depicts the side-splash region forming, producing a spray of ejecta connected to the central dense regions. The third frame shows the shocks travelling through the target. The projectile is nearly stopped, and is stopped in the final image, where the shockfronts have run right through the

target. Note the asymmetry of the side-splash, as the right-hand side of the target has cooled already.

Figure 7.3 shows an identical collision model, except the projectile beam is offset to give an impact parameter of 0.5. Here the initial shocks show the same structure, but the regions of side-splash are much more evident. The bounce-off can also be seen as a small projection in the lower left of the main mass.

Figure 7.4, shows an impact parameter of 0.9. The bounce-off is clearly evident, and shown to be a spectator mass by the fact that in this frame, one can see the collision induced shockfront still traversing the projectile remnant. Note also the complex structure and mixing evident in the exposed region of side-splash.

Both Figures 7.3 and 7.4 show different base colours for the earlier timeslice. This is most probably caused by a single particle (unseen in these spatial slices) having an extremely low thermal energy, most likely ejected in the sidesplash and cooled. As there is no continuity in terms of depicted colours between the two timeframes, it does not detract nor enhance the information depicted in these plots. Although the main target mass is depicted with a different colour, we know that these energies (being purely spectator and completely untouched by and shock waves of other form of energy transfer) must be the same in the two images.

The extreme Lorentz contraction of the projectile, seen in both Figure 7.3 and 7.4 should also be noted.




FIGURE 7.2: NeU collision sequence ($b=0.1$), showing thermal energy distribution within the central region. Spatial scales are 10^{-15} m




FIGURE 7.3: NeU collision sequence ($b=0.5$), showing thermal energy distribution within the central region




FIGURE 7.4: NeU collision sequence ($b=0.9$), showing thermal energy distribution within the central region

7.9.2 Au:Au collisions

Experiments into Heavy Ion collisions advanced rapidly as developments in colliders allowed heavier ions to be used, and at higher energies. Many of the collective flow effects, most notably squeeze out, are only seen at the higher energies and larger physical sizes of heavy projectiles such as gold and uranium. Smaller, lighter ions retain too much identity in their constituent nucleons, for these phenomena to really be seen. In order to keep the collisions well within the bounds of thermodynamics, the situations of pure macroscopic behaviour such as collective flow need to be visited. A collision involving two gold atoms at GeV energies will produce just this sort of environment. We have also employed a change in frame to that of the centre of mass. This means that the resulting dataset is more consistent with the current practice of accelerating two beams of particles in opposite directions on the same trajectory before allowing them to collide within the detector, such as is used in Plastic Ball Spectrometer experiments.

Figure 7.5 shows the thermal slicing used in the previous section, with the baryon number density detail pictured below. The left images are after 300 timesteps, or approximately $1 \text{ fm}c^{-1}$, and the right images are at 1500 timesteps, or $4.9 \text{ fm}c^{-1}$, well after the fireball phase and into the thermal expansion. The scales are equal on both axes to avoid any distortion.

Note how the left hand images show the central fireball region effectively thermalised, but still with the distinct, ellipsoidal regions of differentiated density. Note also how close to the surface the reacting region is, effectively exposed. In the upper thermal slice, the propagating shock fronts extending outwards from the fireball are also evident.

As the simulation then runs through the fireball phase, and into the thermal




FIGURE 7.5: AuAu collision sequence

expansion, we can see evidence of nuclear stopping, as the expansion of both particles begins. The three ellipsoids, the central participants, and the two spectator regions are very evident in this frame. There is also a distinctive asymmetry, appearing as a warp in the fireball disk and the baryon number density peaks. This does however exhibit a certain mirror symmetry about the origin, so is most likely caused by a minor difference in the two ion's constituent SPH particle's initial locations. This is a head on ($b=0.0$) collision, so there should be no asymmetry extending from this cause.

Chapter 8

Performance on a Curved Metric

Up until this stage, all calculations have been performed in regions where the effects of gravity can safely be neglected. This has been accomplished by using the Minkowski metric,

$$g_{\mu\nu} = \text{diag}\{-1, 1, 1, 1\}. \quad (8.1)$$

Many interesting hydrodynamical systems, particularly those studied under the astrophysical umbrella, will naturally require gravity, and so it is logical to extend the algorithm to function on a curved space metric.

We must first examine how a background metric can be justified, and the inherent assumption we will use in such an application. This is done in Section 8.1. We then introduce the space-times we will use in Section 8.2. Sections 8.3 and 8.4 go on to generate exact geodesic equations and how they will be used to assess the performance of the algorithm. It then becomes important that the reader understands some of the physics involved in a radial infall problem, and these are outlined in Section 8.5. The following sections look at the algorithm's initial conditions (Section 8.6), and then the performance of

the algorithm itself in Section 8.7.

8.1 Motivation

In the general relativistic understanding of the physics governing the motion of masses, there are no forces. Rather, particles merely respond to the curvature of the spacetime they find themselves in. Finding these curvatures for a given mass distribution is no mean feat and requires the solving of ten coupled, partial differential equations, the Einstein equations.

As a consequence of this computational difficulty, a relativistic version of the Newtonian N-body codes which calculate the gravitational potential at each step and evolve their particles accordingly is not yet possible. This is not to say that no meaningful simulation can be done, but it does place some constraints.

Just as some gravitational simulations use a background gravitational potential which is unaffected by the low mass test particles orbiting within it, we will use a static background spacetime. This is a valid approximation provided that the mass-energy density of the fluid being modelled is significantly smaller than that of the energy curving the spacetime.

This situation is precisely what is expected to be found in the region around compact objects such as neutron stars or black holes.

Whilst eventually it is desired that this algorithm be applied to the modelling of neutron stars themselves and the dynamics of black hole accretion disks, this work is only looking at situations where the particles' self-gravity can be ignored. Consequently we will look at the simple test case of radial infall of a cold gas cloud into a stationary, static black hole. By specifying the spacetime to be static and stationary, we have an analytical solution against

which to check the code's performance in integrating the equations of motion correctly. We also have specified that the gas is cold dust. This effectively removes the hydrodynamics from the gas' behaviour, leaving only the particles' response to the curvature. Without specifying this constraint, the process of gravitational infall results in heating and expansion of the gas, which resists the straight forward geodesic infall predicted by the analytical solution. As well as this, without particle self-gravity, there is no restoring force to counter the gas' desire to expand as it heats. Without this restoring force, the gas expands freely under the influence of the temperature gradients, resulting in an unphysical representation.

8.2 The Schwarzschild and Kerr Spacetimes

The Schwarzschild space-time represents a spherically symmetric vacuum space-time surrounding a non-charged, non-rotating black hole singularity. Although it is a vacuum solution, provided we keep the modelled fluid's mass-energy to much less than that of the singularity, there should be no significant error in the simulation. The solution is unique, which means that although a variety of coordinate systems can be used to describe the space-time, all are locally isometric to the Schwarzschild space-time.

The classical Schwarzschild coordinates are given by

$$ds^2 = -\left(1 - \frac{2M}{r}\right)dt^2 + \left(1 - \frac{2M}{r}\right)^{-1}dr^2 + r^2(d\theta^2 + \sin^2\theta d\phi^2) \quad (8.2)$$

Apart from the trivial singularities of all polar coordinates at $\theta = 0$ and $\theta = \pi$, this form suffers from a coordinate singularity at $r = 2M$ as well as the physical singularity at the origin $r = 0$.

One way of avoiding this coordinate singularity is to transform the metric to its isotropic form. This is done through the radial coordinate relation

$$r = \bar{r}\left(1 + \frac{M}{2\bar{r}}\right)^2, \quad (8.3)$$

where r is the previous Schwarzschild radial component.

This gives the new space-time description

$$ds^2 = -\left(\frac{1 - \frac{M}{2\bar{r}}}{1 + \frac{M}{2\bar{r}}}\right)^2 dt^2 + \left(1 + \frac{M}{2\bar{r}}\right)^4 \{d\bar{r}^2 + \bar{r}^2 d\theta^2 + \bar{r}^2 \sin^2 \theta d\phi^2\} \quad (8.4)$$

This form of the metric still suffers from singularities at the poles of the coordinate system, but removes the coordinate singularity from the event horizon at $r = 2M$ and suffers only from the unavoidable physical singularity at $r = 0$.

For reasons of simplicity, much of the numerical code has been written in cartesian coordinates, and so we will transpose this metric again. This is achieved through simple tensor transformation rules where,

$$g'_{\mu\nu} = \frac{\partial x^\alpha}{\partial x'^\mu} \frac{\partial x^\beta}{\partial x'^\nu} g_{\alpha\beta} \quad (8.5)$$

We are then left with the final form of the cartesian, isotropic metric:

$$ds^2 = -\left(\frac{1 - \frac{M}{2\bar{r}}}{1 + \frac{M}{2\bar{r}}}\right)^2 dt^2 + \left(1 + \frac{M}{2\bar{r}}\right)^4 \{dx^2 + dy^2 + dz^2\}, \quad (8.6)$$

where

$$\bar{r}^2 = x^2 + y^2 + z^2 \quad (8.7)$$

and

$$x = \bar{r} \cos \phi \sin \theta \quad (8.8)$$

$$y = \bar{r} \sin \phi \sin \theta \quad (8.9)$$

$$z = \bar{r} \cos \theta \quad (8.10)$$

$$(8.11)$$

In this form we have described the simplest, non-Euclidean space-time, that of the non-rotating blackhole. Practical observation suggests that most astronomical bodies rotate. As a consequence the Schwarzschild description, whilst used for most tests, is probably not likely to be an accurate description of the spacetime in the vicinity of a blackhole. The Kerr metric describes the extended space-time around a chargeless blackhole with rotation, and is given below in Lindquist-Boyer coordinates $x^\alpha = \{t, r, \theta, \phi\}$

$$ds^2 = (r^2 + a^2 \cos^2 \theta) \left(\frac{dr^2}{r^2 - 2Mr + a^2} + d\theta^2 \right) + (r^2 + a^2) \sin^2 \theta d\phi^2 - dt^2 + \frac{2Mr}{r^2 + a^2 \cos^2 \theta} (a \sin^2 \theta d\phi - dt)^2, \quad (8.12)$$

where a is the angular momentum per unit mass. Note how setting the angular momentum to zero returns the previous Schwarzschild metric as displayed in equation (8.2).

We then take this new space-time and convert to Kerr-Schild coordinates $\{\bar{t}, x, y, z\}$ through the coordinate transform:

$$x + iy = (r + ia) \sin(\theta) \text{cis}\left(i \int (d\theta + \frac{adr}{r^2 - 2Mr + a^2})\right) \quad (8.13)$$

$$z = r \cos \theta \quad (8.14)$$

$$\bar{t} = \int (dt + \frac{(r^2 + a^2)dr}{r^2 - 2Mr + a^2}) - r \quad (8.15)$$

In these new coordinates, we have the metric

$$ds^2 = dx^2 + dy^2 + dz^2 - d\bar{t}^2 + \frac{2Mr^3}{r^4 + a^2z^2} \left\{ \frac{r(xdx + ydy) - a(xdy - ydx)}{r^2 + a^2} + \frac{zdz}{r} + d\bar{t} \right\}^2, \quad (8.16)$$

where

$$r^4 - (x^2 + y^2 + z^2 - a^2)r^2 - a^2z^2 = 0 \quad (8.17)$$

which reduces in the Schwarzschild limit of no angular momentum ($a = 0$) to

$$ds^2 = dx^2 + dy^2 + dz^2 - d\bar{t}^2 + \frac{2M}{r} \left\{ \frac{(x dx + y dy)}{r} + \frac{z dz}{r} + d\bar{t} \right\}^2, \quad (8.18)$$

Note that this metric has cross-terms of $dx^i dt$ which correspond to a shift vector, β^i . The equations of motion as described in previous chapters were derived with the assumption that the shift vector was zero, precluding this from of metric. It should be a small matter to include these terms in the equations of motion and allow for application of Kerr-type metrics, however it is beyond the scope of this thesis for the time being.

8.3 Exact Solution

As a comparative tool, we would like analytical solutions of the Schwarzschild orbits. As mentioned in the previous section, the metric given in Schwarzschild radial components (equation (8.2)) has some coordinate irregularities that make it unsuitable for the numerical application. It does however, lend itself to neat analytical geodesic solutions, which will allow us to monitor the performance of the code. As the only difference between the two metrics of equation (8.2) and equation (8.4) is the transformation of the radial coordinate given in equation (8.3), we can readily apply these geodesic solutions to our numerical output.

Given an appropriate Lagrangian, equations of motion can be derived. As the initial conditions used specify a cold, dust equation of state ($P = 0$) with no initial velocity, we can assume that the particles should simply fall along the radial geodesics of the spacetime. Taking

$$ds^2 = g_{\mu\nu} dx^\mu dx^\nu \quad (8.19)$$

as equation (8.2) with the assumption of constant θ and ϕ , then we have the Lagrangian

$$\mathcal{L} = \frac{1}{2} g_{\mu\nu} \frac{dx^\mu}{d\tau} \frac{dx^\nu}{d\tau} \quad (8.20)$$

$$= \frac{1}{2} \left\{ -\left(1 - \frac{2M}{r}\right) \left(\frac{dt}{d\tau}\right)^2 + \left(1 - \frac{2M}{r}\right)^{-1} \left(\frac{dr}{d\tau}\right)^2 \right\} \quad (8.21)$$

$$= \frac{1}{2} \left\{ -\left(1 - \frac{2M}{r}\right) \dot{t}^2 + \left(1 - \frac{2M}{r}\right)^{-1} \dot{r}^2 \right\}, \quad (8.22)$$

with τ an affine parameter which can be taken to be the proper time of the

particles, and f' taken as the derivative of $f(\tau)$ with respect to this parameter.

From this we deduce the two canonical momenta to be

$$\hat{p}_t = \frac{\partial \mathcal{L}}{\partial \dot{t}} = \left(1 - \frac{2M}{r}\right) \dot{t} \quad (8.23)$$

$$\hat{p}_r = \frac{\partial \mathcal{L}}{\partial \dot{r}} = \left(1 - \frac{2M}{r}\right)^{-1} \dot{r} \quad (8.24)$$

and the Hamiltonian of the system is given by

$$\mathcal{H} = -\hat{p}_t \dot{t} + \hat{p}_r \dot{r} - \mathcal{L} = \mathcal{L} = \text{constant}, \quad (8.25)$$

ensuring the energy of the system comes entirely from the kinematics. Consequently we can rescale τ such that for the time-like geodesics we are interested in,

$$\mathcal{H} = -\frac{1}{2}. \quad (8.26)$$

The derivative of the time momenta, \hat{p}_t gives us

$$\frac{d\hat{p}_t}{d\tau} = \frac{\partial \mathcal{L}}{\partial t} = 0 \quad (8.27)$$

implying that

$$\hat{p}_t = \text{constant} = \left(1 - \frac{2M}{r}\right) \dot{t} = E \quad (8.28)$$

The equations of motion then governing this system reduce to

$$\frac{dt}{d\tau} = \frac{E}{\left(1 - \frac{2M}{r}\right)} \quad (8.29)$$

and through substituting equation (8.29) into $\mathcal{L} = -\frac{1}{2}$

$$\left(\frac{dr}{d\tau}\right)^2 = \frac{2M}{r} - (1 - E^2) \quad (8.30)$$

These are the equations (96) from Chapter 2 of Chandrasekhar [1983]'s work, which we now follow.

In order to find a solution to these equations, we transform to a new parameter space through

$$r = r_i \cos^2(\eta), \quad (8.31)$$

where r_i is the initial radial distance, and connected to E through equation (8.30) and the initial condition $\dot{r} = 0$ when $r = r_i$ to give

$$r_i = \frac{2M}{(1 - E^2)}. \quad (8.32)$$

The equations of motion in terms of this parameter η reduce to

$$\frac{dr}{d\tau} = -\sqrt{\frac{2M}{r_i}} \tan\left(\frac{\eta}{2}\right) \quad (8.33)$$

$$\frac{dt}{d\tau} = \frac{E \cos^2(\frac{\eta}{2})}{\cos^2(\frac{\eta}{2}) - \cos^2(\frac{\eta_H}{2})} \quad (8.34)$$

$$\frac{dr}{d\eta} = -r_i \sin\left(\frac{\eta}{2}\right) \cos\left(\frac{\eta}{2}\right) \quad (8.35)$$

Here we have introduced the negative root in equation (8.33) (as we are considering infalling particles) and the parameter η_H which corresponds to the η value when the particle arrives at the event horizon, $r = 2M$, given by

$$\eta_H = 2 \sin^{-1}(E). \quad (8.36)$$

Equations (8.35) and (8.33) combine to give

$$\frac{d\tau}{d\eta} = \left(\frac{r_i^3}{8M}\right)^{\frac{1}{2}}(1 + \cos \eta) \quad (8.37)$$

which is integrated assuming $\tau = 0$ when $r = r_i$ to give

$$\tau = \left(\frac{r_i^3}{8M}\right)^{\frac{1}{2}}(\eta + \sin \eta) \quad (8.38)$$

We can now show

$$\frac{dt}{d\eta} = E\left(\frac{r_i^3}{8M}\right)^{\frac{1}{2}} \frac{\cos^4(\frac{\eta}{2})}{\cos^2(\frac{\eta}{2}) - \cos^2(\frac{\eta_H}{2})} \quad (8.39)$$

and

$$t = E\left(\frac{r_i^3}{8M}\right)^{\frac{1}{2}} \left\{ \frac{1}{2}(\eta + \sin(\eta)) + \eta(1 - E^2) \right\} + 2M \ln \left\{ \frac{\tan^2(\frac{\eta_H}{2}) + \tan^2(\frac{\eta}{2})}{\tan^2(\frac{\eta_H}{2}) - \tan^2(\frac{\eta}{2})} \right\} \quad (8.40)$$

Together with equation (8.32), equation (8.40) allows us to specify the coordinate time and location of any particle as functions of its proper time and initial position.

8.4 Generating the Exact Solution

In order to display the performance of the SPH algorithm in this application, a plotter is required which is capable of deducing the location of a particle at a given time. The code itself outputs all the relevant particle attributes at each timestep, as well as information regarding the coordinate time of the simulation. To generate the exact comparison, each particle is treated individually. Its initial \bar{r} is determined and transformed into Schwarzschild

coordinates, from which we deduce the value of E through

$$E^2 = 1 - \frac{2M}{r}. \quad (8.41)$$

We can then use a Newton-Raphson root-finding method to determine η from the root of

$$f(\eta) = E\left(\frac{r_i^3}{8M}\right)^{\frac{1}{2}}\left\{\frac{1}{2}(\eta + \sin(\eta)) + \eta(1 - E^2)\right\} + 2M \ln\left|\frac{\tan^2(\frac{\eta_H}{2}) + \tan^2(\frac{\eta}{2})}{\tan^2(\frac{\eta_H}{2}) - \tan^2(\frac{\eta}{2})}\right| - t. \quad (8.42)$$

Having found the current η , it is simple then to deduce r from equation (8.31) and to convert this back to the isotropic radial coordinate, \bar{r} .

For much of the analysis, only one particle is required to be traced. The graphs, such as Figures 8.3, 8.4 and 8.5, which show \bar{r} as a function of coordinate time t or η show the evolution of just one particle (particle number n) from a given simulation.

8.5 The Physics of the Schwarzschild Infall Problem

The spherically symmetric space-time described here is portrayed from the point of view of an observer at a great distance from the source of the curvature, i.e. the black hole. The curvature of the space-time increases indefinitely until the origin, which is singular. At some point outwards from this there exists an event horizon, across which no information can escape the black hole. There is nothing physically different about this location, although the Schwarzschild

coordinate system is singular at this point ($r = 2M$). The isotropic coordinate frame is symmetric about the horizon ($\bar{r} = \frac{M}{2}$) and how the geodesics behave here is important.

The equations of the radial geodesics along which particles will move were derived in the last section, and the major features of the paths are described in Table 8.1.

	η	τ	t
Initial ($r = r_i$)	0	0	0
Event Horizon ($r = 2M$)	$\eta_H = 2 \sin^{-1} E$	$(\frac{r_i^3}{8M})^{\frac{1}{2}} (\eta_H + \sin(\eta_H))$	∞
Singularity ($r = 0$)	π	$(\frac{r_i^3}{8M})^{\frac{1}{2}} \pi$	**

TABLE 8.1: Critical Points in the Infall Problem

It is important to note that although the proper time for a particle to cross the event horizon and to continue on down to the singularity is finite, for the distant observer infinite time will pass before the particle reaches the event horizon. As the curvature increases the lapse function collapses to zero at the throat, effectively halting evolution in this coordinate frame.

8.6 Initial Conditions

In order to reproduce geodesic infall, all thermodynamical properties of the SPH particles must be removed. This is done by simply specifying that the isotropic pressure is zero, and all terms associated with the artificial viscosity routines must also go to zero. If this is not done, the gravitational energy of the clump of gas causes heating, and a tendency to expand, as it falls into the gravitational well. Without any particle self-gravity, this heating produces a force, lifting the particles from their initial geodesics, and making any comparison with analytical solutions impossible.

The initial conditions for this test use similar routines to those in the RHIC tests (See Section 7.7). A ball of SPH particles, created by removing the corners from a large, body-centred-cubic structured grid is placed at some radius ($r = 200M$) and allowed to evolve.




FIGURE 8.1: Initial configuration, $n=1100$ particles showing the code output on the upper panel, and the analytical solution on the lower

Figure 8.1 shows the initial set-up of a cold ball of approximately 1100

particles situated at $\bar{r} = 200M$. The frame below shows the exact solution, which (naturally) is no different at $t = 0$.

8.6.1 A Note on Time-Stepping

As the normal time-stepping algorithm is inherently connected to the thermodynamic properties of the gas, there is no time-step control applied here. The maximum time-step allowed is an enforced, and arbitrary, $\Delta t = 1.5M$. This is fine for stability, as there are no thermodynamics, but it causes large errors in regions of large curvature, such as is found close to the black hole. These errors can be reduced by forcing a smaller time-step, but would only resurface again later in the simulation due to the re-scalability of the back-ground metric. As any later simulation will utilise the full hydrodynamical nature of the SPH algorithm and its variable time-stepping controls, developing variable time-steps here would be fruitless as this is merely a test of correct response to background curvature.

8.7 Performance

As mentioned before, the event horizon cannot be passed using this coordinate system. This is not a physical problem though, as once crossed, no information can be gotten back from this horizon, effectively removing any modelling across it anyway. All meaningful models must be run in the space well to the outside of this horizon. Whereas the Schwarzschild metric used to generate the exact solution is undefined at this point, with a division by zero, the isotropic metric used by the code is well defined, but symmetric across the event horizon. The result of this is that should any particle erroneously

contact the event horizon (through an inappropriate timestep) the code will be unable to continue and crash. Up until this point though, one would expect no problems from the metric, although the dilations anywhere near this horizon do become extreme. This is highlighted in Figure 8.2, which depicts both the computational output (above) and the exact solution (below) for $t \approx 3000M$. The axes are justified, as in Figure 8.1, to show the large tidal distortions occurring.

The following graphs (Figures 8.3 to 8.5) present the data for only one particle, plotted as a series of circles, and the exact solution, plotted as a solid line. Figure 8.3 shows the particle's infall as a function of the coordinate time t . This detail is also shown in Figure 8.4 where the η coordinate is used as the independent variable. The value of η_H is shown on the right as a broken line.

Each circle represents the location of the n^{th} particle at a given coordinate time. Looking at the spacings between these points, which are fixed at constant Δt , it is apparent how the timestep will lead to failure of the algorithm eventually. Early on in the model run, the particle is in a region of weaker curvature and the timestep is adequate, as indicated by the line of the exact solution bisecting the circles. As the curvature, and hence acceleration, of the particle increases, we can see the circles begin to leave the curve. This can be rectified by using smaller timesteps, but due to the nature of the space-time, any fixed timestep will eventually be too great for the curvature, and will fail.

Absolute errors associated with the particle's location are calculated and are shown in Figure 8.5. The nature of the space-time means that errors in location will compound, as seen by the exponential growth. The rate of change of the curvature with respect to radial distance, commonly known as tidal forces, is so great that small errors in the particle's location result in different acceleration forces, which leave the particle in further error after




FIGURE 8.2: Schwarzschild Infall, $t=3000M$, showing large tidal distortions of the initially spherical cluster of particles

each timestep. The error remains quite acceptable until such time as the constant timestep is unable to handle the massive distortions that the space-time experiences. Eventually, the particle will overstep and wrongly cross




FIGURE 8.3: Single Particle Infall as a function of coordinate time, t .

the event horizon, crashing the code. Whilst this can easily be combatted by using a smaller timestep, it cannot be resolved by any application of a constant timestep, as eventually any time-step will become too large for the curvature.

The failure of a constant timestep routine is indicated most clearly by ex-





FIGURE 8.4: Single Particle Infall as a function of the parameter η

amining the performance of the code in the η coordinates, as seen in Figure 8.4. Here the asymptotic approach to the singularity is seen, and the inadequacy of the timestep as the distance between datapoints (circles) grows as the curvature increases. Clearly this is the region where we need a smaller

timestep to correctly resolve the rapid accelerations. This should be an easy thing to implement, as information about time evolution is contained within the lapse function, which can be incorporated into the hydrodynamic timestep controls. The choice of $D = \sqrt{\eta\gamma\rho_0}$ also contains information that will control dynamical timesteps, as increasing curvature appears as a density gradient in the computational variables, which should induce smaller timesteps.

Figure 8.5 shows a ramping of the error as the particle approaches the throat. This large error is not indicated on the other plots, and it is unclear whether it is a case of small denominators inducing round-off errors, or the fact that the analytical solution is coping with a coordinate failure in this region.

As pointed out previously, this region of space is one with large distortions and where one would not usually be attempting simulations. The errors induced here can easily be avoided by not running simulations at too small a radius, where the chosen timestep will be inadequate, or by using a hydrodynamical timestep control which takes this curvature into account.

**FIGURE 8.5:** Relative error of particle location

Chapter 9

Conclusions and Further Applications

This work has shown that the SPH method can be applied to problems where the relativistic effects of high velocities and/or background space-time curvature are non-trivial. One of the clear benefits of the SPH method is its simplicity and ease of application. The three examples shown in this thesis: Shock-Tubes, Relativistic Heavy Ion Collisions and Gravitational Infall problems all use the same code. The only modules required to be changed are the module containing the metric functions (Minkowski flat-space or isotropic Schwarzschild), the chosen equation of state and the initial conditions. Thus, the algorithm has the scope to be applied to a wide range of problems with relatively little effort.

The algorithm does not rely on exact or approximate Riemann Solvers to resolve the shocks, and considering also the general nature of the code, its ability to capture shocks and resolve density spikes is acceptable. The post-shock ringing can be tuned out by changing the artificial viscosity parameter

or numbers of desired neighbours, or selecting specific viscosity measures for a given problem. Many of the oscillations are actually diminished by increasing the resolution, something that the inefficiencies of modelling 1-D shocktubes with a 3D code makes prohibitive and is more of an issue of poor test case availability than a performance issue.

With the excellent performance of many high-resolution shock-capturing codes, attempting to make SPH perform in this role would be ignoring the real performance strengths of the method. The SPH method can resolve fronts and waves within the fluid without pre-warning of what they will be. It is best used in this application, where a simple code will resolve all the hydrodynamical waves in the fluid, allowing further study with the benefit of the knowledge of what wave types to expect, when and where and the consequent application of a suitable method with the required resolution and stability.

When there are free-surfaces, the particle method has clear advantages over other methods which must somehow enforce a free boundary to the grid.

The Relativistic Heavy Ion Collision calculations show that relativistic SPH calculations perform satisfactorily, at least on the qualitative front. In particular, the avoidance of Riemann Solvers means the use of equations of state more complex than that of the ideal relativistic gas is a trivial matter. The complex fluid motions, particularly those of initial contact and grazing collisions, are handled easily by the particle simulation, which models recombining fluids without the expense of grid remapping, and automatically places more particles into this region to resolve finer scale motions.

Whilst a detailed comparison with experimental results would require the use of non-general, more elaborate equations of state which are beyond the

aims of this thesis, certain advantages of the SPH method for this application do become clear.

These include:

- The efficiency of using particles, placing the computational work where the fluid is. This includes the efficient modelling of expansion processes through the rank-list neighbour routines, ensuring that the computational box is free to expand or contract with no extra memory requirement.
- The non-reliance on Riemann Solvers with their non-triviality of adding extra, more complicated physics such as non-ideal or empirical equations of state
- The ability to satisfy and incorporate other requirements peculiar to HIC, such as calculations of the freeze-out surface which involves a transition from coherent fluid behaviour to something closer to that of non-interacting, individual particles.

It is on this last point that the method could best be improved. In doing so, much of the general nature of the original application will be lost, and as such the extensions are best done to suit the specific study at hand.

The algorithm has also shown it can perform well, responding to simple background metrics (those with a zero-shift vector), allowing it to perform hydrodynamical calculations in the region of compact objects without any post-Newtonian additions or laborious pseudo-tensor calculations. The algorithm is written in such a way as to incorporate features used in modern numerical relativity studies, such as shift vectors. This allows it to be used in testing 'new' metrics describing different regions of space.

It is in this area of metric applications that the author believes most benefit can come. The most obvious application is the modelling of thick, dusty accretion disks around black-holes. In order to do this, a suitable internal boundary condition must be devised to capture particles as they cross the event horizon and remove them from the simulation. There would be no reason to attribute their mass to the source, as the application and suitability of the background space-time means that the SPH particles are much less massive than the source anyway.

Changing the equations of motion so that the particles can respond to a metric with non-zero shift vector would allow the most physically apt metric for accretion and orbital problems (the Kerr) to be applied.

Possibly the clearest advantage in this algorithm is its use of efficient memory allocation achieved by performing most calculations in rank-space. In doing this, the effective computational cell remains the same size, and of constant occupation density. This should easily allow load-balancing on a parallel application. It also has the added benefit of simplifying the creation of load-balanced oct-trees for applications with particle self gravity, in that each node in the tree should remain constant. The structure of the tree then can be built once on rank-space, and particles then allocated according to their rank identities, not their physical location.

In conclusion, the relativistic SPH algorithm has shown itself to be a robust and capable application, with clear benefits in both ease of writing and application of complex physics. This thesis provided a simple, general algorithm that has the potential to be applied to a wide range of modelling scenarios.

Bibliography

Aguiar, Hama, Kodama, and Osama. Event-by-event fluctuations in hydrodynamical descriptions of heavy-ion collisions. arXiv:hep-ph/0106266 v1, Jun 2001.

Aguiar, Kodama, Osada, and Hama. Smoothed particle hydrodynamics for relativistic heavy ion collisions. arXiv:hep-ph/0006239, June 2000.

Amsden, Bertsch, Harlow, and Nix. Relativistic hydrodynamic theory of heavy ion collisions. *Physical Review Letters*, 35(14):905-908, 1975

Amsden, Goldhaber, Harlow, and Nix. Relativistic two-fluid model of nucleus-nucleus collisions. *Physical Review C*, 17(6), 1978.

Amsden, Harlow, and Nix. Relativistic nuclear fluid dynamics, 1977.

A. Anile. *Relativistic Fluids and Magneto-Fluids*. Cambridge Uni-Press, 1989.

Borve, Omang, and Trulsen. Smoothed particle hydrodynamics with time varying, piecewise profiles. *American Astronomical Society Meeting*, 197 (107.22), 2000.

Borve, Omang, and Trulsen. Regularized smooth particle hydrodynamics: A new approach to simulating magnetohydrodynamic shocks. *Astrophysical Journal*, 561:82-93, 2001.

- Bravina, Brandstetter, Greiner, Zabrodin, Belkacem, Bleicher, Bass, Ernst, Hofman, Soff, Stöcker, and Greiner. Local thermal and chemical equilibration and the equation of state in relativistic heavy ion collisions. *Journal of Physics G:Nuclear Particle Physics*, 25:351-361, 1999.
- L. Brewin. Riemann normal coordinates. Department of Mathematics Preprint, Monash University, Clayton, Vic., Australia 3600, 1997.
- Chandrasekhar. *An Introduction to the Study of Stellar Structure*. University of Chicago, 1939.
- Chandrasekhar. *The Mathematical Theory of Black Holes*. The International Series of Monographs on Physics. Oxford University Press, 1983.
- Chow and Monaghan. Ultrarelativistic sph. *Journal of Computational Physics*, 134:296-305, 1997.
- Eckart. The thermodynamics of irreversible processes: I the simple fluid. *Physical Review*, 58, 1940.
- Evrard. Beyond n-body - 3d cosmological gas dynamics. *Monthly Notices of the Royal Astronomical Society*, 235:911-934, 1988.
- Font, Ibanez, Marquina, and Marti. Multidimensional relativistic hydrodynamics: Characteristic fields and modern high-resolution shock-capturing schemes. *Astronomy and Astrophysics*, 282(1):304-314, 1994.
- Fulbright, Benz, and Davies. A method of smoothed particle hydrodynamics using spheroidal kernels. *Astrophysical Journal*, 440:254-262, 1995.
- Gingold and Monaghan. Smoothed particle hydrodynamics: Theory and

application to non-spherical stars. *Monthly Notices of the Royal Astronomical Society*, 181:375-389, 1977.

Gingold and Monaghan. Binary fission in damped rotating polytropes. *Monthly Notices of the Royal Astronomical Society*, 184:481-499, 1978.

Gingold and Monaghan. Kernel estimates as a basis for general particle methods in hydrodynamics. *Journal of Computational Physics*, 46:429-453, 1982.

Glassgold, Heckrotte, and Watson. *Annals of Physics*, 6, 1959.

Gutbrod and Stöcker. The nuclear equation of state. *Scientific American*, 265: 32-39, 1991.

Harlow, Amsden, and Nix. Relativistic fluid dynamics calculations with the particle-in-a-cell technique. *Journal of Computational Physics*, 20:119-129, 1976.

S. Hawking and G. Ellis. *The Large Scale Structure of Space-Time*. Cambridge Monographs on Mathematical Physics. Cambridge at the University Press, 1973. pp.162.

Hawley, Smarr, and Wilson. A numerical study of non-spherical black hole accretion. ii-finite differencing and code calibration. *Astrophysical Journal Supplement*, 55:221, 1984.

Hernquist. Some cautionary remarks about smoothed particle dynamics. *Astrophysical Journal*, 404:717-722, 1993.

Hernquist and Katz. Treesph: a unification of sph with the hierarchical tree method. *Astrophysical Journal Supplement*, 70:419-446, 1989.

- Hiscock and Lindblom. Generic instabilities in first-order dissipative relativistic fluid theories. *Physical Review D*, 31(4), 1985.
- Hung and Shuryak. Equation of state, radial flow, and freeze-out in high energy heavy ion collisions. *Physical Review C*, 57(4):1891–1906, 1998.
- Israel. Nonstationary irreversible thermodynamics: A causal relativistic theory. *Annals Of Physics*, 100:310–331, 1976.
- Israel and Stewart. Transient relativistic thermodynamics and kinetic theory. *Annals of Physics*, 118:341–372, 1979.
- P Laguna, W Miller, and W Zurek. Smoothed particle hydrodynamics near a black hole. *The Astrophysical Journal*, 404:678–685, 1993.
- Landau and Lifshitz. *Fluid Mechanics*, volume 6 of *Course of Theoretical Physics*. Pergamon Press, 1959. Translated by Sykes and Reid.
- Lattanzio and Monaghan. A refined particle method for astrophysical problems. *Astronomy and Astrophysics*, 149:135–143, 1983.
- Lucy. A numerical approach to the testing of the fission hypothesis. *Astronomical Journal*, 82:1013–1024, 1977.
- Lun. Private Discussion, 2001.
- Marti and Müller. The analytical solution of the riemann problem in relativistic hydrodynamics. *Journal of Fluid Mechanics*, 258:317–333, 1994.
- Marti and Müller. Extension of the piecewise parabolic method to one-dimensional relativistic hydrodynamics. *Journal of Computational Physics*, 123:1–14, 1996.

- Marti and Müller. Living reviews in relativity, vol 2.
<http://www.livingreviews.org/Articles/Volume2>, 1999.
- C. Misner, K. Thorne, and J. Wheeler. *Gravitation*, chapter 11.6, pp 285. W.H. Freeman and Co, New York, 1973a. 22nd printing.
- C. Misner, K. Thorne, and J. Wheeler. *Gravitation*. W.H. Freeman and Co, New York, 1973b.
- Miyama, Hayashi, and Narita. Criteria for collapse and fragmentation of rotating, isothermal clouds. *Astrophysical Journal*, 279:621–632, 1984.
- Monaghan. Particle methods for hydrodynamics. *Computational Physics Reports*, 3:71–124, 1985.
- Monaghan. Smoothed particle hydrodynamics. *Annual Review of Astronomy and Astrophysics*, 30:543–574, 1992.
- Monaghan and Gingold. Shock simulation by the particle method sph. *Journal of Computational Physics*, 52:374–389, 1983.
- Monaghan and Price. Variational principles for relativistic sph. Epsilon Laboratory, Department of Mathematics and Statistics Preprints, Monash University, Clayton, Victoria, AUSTRALIA 3600, 2001.
- Morris. A study of the stability properties of smooth particle hydrodynamics. *Publications of the Astronomical Society of Australia*, 13:97–102, 1996.
- Nelson and Papaloizou. Variable smoothing lengths and energy conservation in smoothed particle hydrodynamics. *Monthly Notices of the Royal Astronomical Society*, 270:1–20, 1994.

- Nix and Strottman. Fluid dynamical description of relativistic nuclear collisions. for presentation at the Second International Colloquium on Drops and Bubbles, Monterey, California, Nov 19-21 1981.
- Nix, Strottman, Yariv, and Fraenkel. Coherent collective flow versus independent two-nucleon collisions in central high energy heavy ion collisions. *Physical Review C*, 25(5):2491-2497, 1982.
- R. Oechslin, S. Rosswog, and F. Thielemann. Conformally flat smoothed particle hydrodynamics: Application to neutron star mergers. *arXiv:gr-qc/0111005*, 2001.
- Olson and Hiscock. Stability, causality and hyperbolicity in carter's 'regular' theory of relativistic heat conducting fluids. *Physical Review D*, 41:3687, 1990.
- Schneider, Katscher, Rischke, Waldhauser, Maruhn, and Munz. New algorithms for ultra-relativistic numerical hydrodynamics. *Journal of Computational Physics*, 105:92-107, 1993.
- Schüssler and Schmitt. Comments on smoothed particle hydrodynamics. *Astronomy and Astrophysics*, 97(2):373-379, 1981.
- S. Siegler and H. Riffert. Smoothed particle hydrodynamics simulations of ultra-relativistic shocks with artificial viscosity. *arXiv:astro-ph/9904070*, 1999.
- Sod. A survey of several finite difference methods for systems of non-linear hyperbolic conservation laws. *Journal of Computational Physics*, 27:1-31, 1978.
- Steinmetz and Müller. On the capabilities and limits of smoothed particle hydrodynamics. *Astronomy and Astrophysics*, 268:391-410, 1993.

- Stock. The physics of dense nuclear matter from supernovae to quark gluon plasma. *Nature*, 337:319–324, 1989.
- Synge. *The Relativistic Gas*. North Holland Publishing, 1957.
- Taub. Relativistic fluid mechanics. *Annual Review of Fluid Mechanics*, 10:301, 1978. As quoted in [Marti and Müller, 1999].
- Teaney, Lauret, and Shuryak. A hydrodynamic description of heavy ion collisions at the sps and rhic. arXiv:nucl-th/0110037 v2, Dec 2001.
- Thacker, Tittley, Pearce, Couchman, and Thomas. Smoothed particle hydrodynamics in cosmology: a comparative study of implementations. *Monthly Notices of the Royal Astronomical Society*, 319:619–648, 2000.
- Thompson. The special relativistic shock tube. *Journal of Fluid Mechanics*, 171:365–375, 1986.
- Witten, editor. *Gravitation: An Introduction to Current Research*. Louis Wiley and Son, 1962.
- Cheuk-Yin Wong. *Introduction to High-Energy Heavy-Ion Collisions*. World Scientific, 1994.
- Wood. Collapse and fragmentation of isothermal clouds. *Monthly Notices of the Royal Astronomical Society*, 194:201–218, 1981.

'An expert is a person who has made all the mistakes that can be made in a very narrow field,' Niels Bohr

'If I had only known, I would have been a locksmith', Einstein

'This World will never know Peace until the last politician is strangled with the entrails of the last priest,' Voltaire