

Pulsed Expansion of Plasma in a Magnetic Thruster

Thesis by
Fujio Inoya

In Partial Fulfillment of the Requirements
for the Degree of
Doctor of Philosophy

California Institute of Technology
Pasadena, California
1993

(Submitted and Defended August 5, 1992)

Copyright 1993
Fujio Inoya
All rights Reserved

Acknowledgements

The author would like to extend his deepest thanks to Professor F.E.C.Culick not only for allowing work on this subject to be undertaken, but for offering helpful suggestions and advice throughout the research effort. Furthermore, without his considerable psychological support and generous understanding, this work would never have been possible.

Thanks also go to Dr.A.R.Martin for having sent me useful publications, Dr.A.Nishiguchi for valuable discussions regarding his PIC algorithm, and Dr.C.D.Orth for helpful conversations regarding ICF. The author would also like to acknowledge the Caltech CADRE staff and the JPL-CRAY staff for installing ELLPACK on the respective computer systems, and in particular, the assistance of Mr.J.S.Miller of the JPL-CRAY was very much appreciated.

The major computational work for this research was conducted on the Caltech CADRE (VAX) system, the JPL-CRAY, and the GODDARD-CRAY (through JPL).

The author was supported primarily through the NSF Graduate Fellowship, Caltech fellowships, and graduate research assistantships supported in part by funds obtained from the Advanced Propulsion Group at JPL.

Abstract

The inertial confinement fusion (ICF) pulse rocket is an advanced space propulsion system, which, through intermittent nuclear fusion energy production isolated from the vehicle structure, is capable of both extremely high specific impulses and high thrust-to-weight ratios. Such rockets, if realized, should revolutionize space travel by making possible very robust interplanetary missions as well as interstellar flight. The thruster of the rocket, which converts an initially isotropically expanding ICF debris plasma into a directed pulse jet exhaust, is to be fabricated out of magnetic fields created by current coils attached to the vehicle. The proper operation of such a thruster therefore rests upon the successful redirection of an initially spherical plasma of high conductivity by a suitably configured vacuum magnetic field against which the plasma expands. But to date, there have been no detailed analyses to guarantee that the concept in the present form will function satisfactorily as envisioned to yield reasonable propulsive efficiencies.

Because of the highly dynamic behavior of the flow, which is bounded by an interface whose motion is unknown *a priori*, the first problem which must be investigated is that of the bulk flow under idealized conditions. In the work contained in this thesis, the plasma was assumed to be impermeable to the external fields, and the fields entered the debris dynamics only by way of applying a magnetic pressure force at the plasma-vacuum interface. The interface motion, bulk fluid profiles (when applicable), and resultant efficiencies were investigated for various parameter ratios and geometries. Such idealized bulk flow analyses are intended to serve as a basis for more detailed studies of how the flow will behave with a real plasma.

Numerical simulations of the bulk flow process were conducted under both the thin-shell and the classical hydrodynamic approximations. The thin-shell calculation has been pioneered by other authors, but the present work is more complete, and as for the

hydrodynamical calculations, application to the type of flows to be found in the magnetic thrusters of proposed ICF pulse rockets may be unique to this work, despite earlier claims. In the former approach, all of the plasma is assumed to be collected into an azimuthally symmetric perfectly conducting shell at the interface by virtue of the finite applied pressure at the interface. No fluid dynamics is considered under this approximation. These simulations showed that promising propulsive efficiencies could be obtained for a range of field-to-plasma energy ratios and thruster geometries, and the efficiencies reached a well-defined maximum for particular values of these parameters. However, because of the approximations used in this model, the efficiencies obtained do overestimate the real efficiencies. The thin-shell code is simpler to implement, and allows faster calculations and requires far less memory, than the more realistic hydrodynamic code, but the approximation made is not entirely accurate nor physical. In the second approach, the plasma is approximated by an unmagnetized perfectly conducting fluid obeying the laws of classical hydrodynamics. Here, we have a novel problem of a fluid expanding against a region of zero density, which nevertheless exerts a finite pressure on the fluid interface.

In both the two-dimensional thin-shell and hydrodynamic calculations, the vacuum magnetic pressure applying at the plasma-vacuum interface was calculated from the quasi-static Maxwell Equations. By assuming the plasma and field coil structures to be perfectly conducting, the magnetic field in the vacuum region, from which the magnetic pressure at the interface was computed, was calculated by prescribing the initial flux through the field coils to remain trapped between the expanding plasma surface and the surfaces of the field coil structures. Such a prescription, which can be explained through the presence of surface currents, is valid as long as we have ideal perfect conductors.

The hydrodynamic codes (both 1-D and 2-D) employed an advanced Classical Particle-In-Cell (PIC) scheme, and were successful at capturing the interface motion self-consistently (with pressure matching across the interface), and without iterations, via appropriate application of boundary conditions. The shock arising from the interface

deceleration was also captured correctly. The formation of a shell-like structure originating close to the interface was observed in simulations of flows with large expansion ratios that were carried out in two dimensions employing realistic thruster fields. But depending upon the pressure history at the interface, these "shells" did not necessarily stay at the interfacial region. When tested on such processes as free expansion into a vacuum or shock-tube problems, for which well-known theoretical solutions exist, the one-dimensional planar-geometry simulations gave results that matched well with the analytical calculations.

The qualitative features of the interface and its motion as found by the hydrodynamic simulations were similar to those obtained by the thin-shell simulations. Nevertheless, the physics of the internal flow was found to affect the performance of the thruster in ways not accountable by the thin-shell model. There were also implications that not all of the debris plasma may leave the thruster in one reflection. The substantial shock heating observed in the interfacial regions downstream of the inward-facing shock would help contribute towards maintaining high temperatures there for (possibly) achieving sufficient conductivities, provided the plasma stayed highly ionized. But because of the large expansion ratio experienced, the bulk temperature of an ICF debris plasma will fall below the ionization temperature from relatively early stages of expansion in the magnetic thrusters of currently proposed ICF pulse rockets, and the design parameters of these thrusters do not appear that promising.

Because of memory limitations imposed by computers, the maximum expansion ratio treatable by the two-dimensional hydrodynamic codes was limited, and initial plasma states rather far removed from those typical of situations in proposed thrusters had to be employed. This also lowered the efficiency values quite notably. The ignorance of real plasma properties such as finite conductivities further rendered the results of this work very optimistic. However, the primary goal of this work, which was to acquire intuition for the bulk flow and performance under idealized conditions, was accomplished. Furthermore, techniques for handling this type of problem were developed. Future work should

concentrate on treating more realistic parameters and on incorporating more precise plasma physics into the analysis, based on bulk flow results heretofore obtained.

Table of Contents

Acknowledgements	iii
Abstract	iv
1 Introduction	I-1
2 Topics in Nonchemical Space Propulsion	II-1
2.1 Some Facts on Interstellar Distance Scales	II-1
2.2 Fundamentals of Rocket Propulsion	II-2
2.3 Brief Survey of Nonchemical Space Propulsion Systems	II-8
2.3.1 Nuclear Thermal Rockets	II-8
2.3.2 Beamed Thermal Rockets	II-11
2.3.3 Electrothermal Rockets	II-13
2.3.4 Electrostatic Rockets	II-13
2.3.5 Electromagnetic Rockets	II-14
2.3.6 Photon Sails	II-15
2.3.7 Continuous Fusion Rockets	II-16
2.3.8 Antimatter Rockets	II-17
2.3.9 Interstellar Ramjets	II-19
References for Chapter 2	II-20
Appendix of Chapter 2	II-A1
Part 1 Tables for Chapter 2	II-A2
Part 2 Figures for Chapter 2	II-A3
Part 3 References for the Appendix of Chapter 2	II-A6

3	Introduction to ICF Pulse Propulsion	III-1
3.1	Nuclear Pulse Propulsion	III-1
3.2	ICF Pulse Propulsion	III-4
3.3	Proposed Designs and Earlier Works; High Field and Low Field Thrusters	III-10
	References for Chapter 3	III-16
	Appendix of Chapter 3	III-A1
	Part 1 Figures for Chapter 3	III-A2
	Part 2 Some Simple Estimates and Approximate Calculations	III-A6
	Part 3 References for the Appendix of Chapter 3	III-A25
4	Two-Dimensional Simulations Employing the Thin-Shell Approximation	IV-1
4.1	Problem Formulation	IV-5
4.1.1	Thin-Shell Dynamics	IV-5
4.1.2	Calculation of the Magnetic Pressure at the Interface	IV-9
4.1.2.1	Quasi-static Maxwell's Equations in a Vacuum	IV-9
4.1.2.2	Boundary Conditions for the Field Calculation	IV-11
4.2	Sample Run for a Single-Coil High-Field Thruster	IV-16
4.3	Thrust and Efficiency Calculations	IV-21
4.3.1	Calculation of Thrust in a Magnetic Thruster	IV-21
4.3.2	Calculation of Propulsive Efficiencies	IV-23
4.3.3	Efficiencies for a Rigid Paraboloidal Reflector	IV-25
4.3.4	Efficiencies in the Single-Coil Sample Run	IV-27
4.4	Comparison of Cases with Different Parameter Ratios for Single-Coil Thrusters	IV-29
4.4.1	Case 1	IV-30

4.4.2 Case 2	IV-32
4.4.3 Case 3	IV-32
4.4.4 Case 4	IV-33
4.4.5 Case 5	IV-34
4.4.6 Case 6	IV-35
4.4.7 Case 7	IV-36
4.5 Sample Run for a Multicoil Low-Field Thruster	IV-37
4.6 Summary	IV-41
References for Chapter 4	IV-45
Appendix of Chapter 4	IV-A1
Figures for Chapter 4	IV-A2
5 Introduction to PIC	V-1
5.1 Properties of Eulerian, Lagrangian, and PIC Methods	V-2
5.2 The PIC Scheme	V-6
5.2.1 Manipulation of Particle Quantities in PIC	V-6
5.2.2 PIC Algorithms	V-8
5.2.2.1 Classical PIC Methods	V-8
5.2.2.2 Full-Particle PIC Methods	V-10
5.2.2.3 Magnetic Field Calculations	V-11
5.2.3 SPH Methods	V-12
5.2.4 Time Scale Limitations	V-13
5.3 Nishiguchi-Yabe Algorithm	V-14
5.4 Artificial Viscosities and Shock Capturing	V-20
5.5 Ringing Instability	V-23
5.6 Summary	V-24
References for Chapter 5	V-29

Appendix of Chapter 5	V-A1
Part 1 Figures for Chapter 5	V-A2
Part 2 Reference for the Appendix of Chapter 5	V-A4
6 Slab-Geometry Hydrodynamic Simulations	VI-1
6.1 The Equations	VI-4
6.2 Some Useful Techniques	VI-7
6.3 Reduction of Numerical Problems	VI-11
6.4 Discussion of Selected Results	VI-14
6.4.1 Free Expansion into a Vacuum	VI-14
6.4.2 Expansion Against an Applied Pressure Vacuum Boundary	VI-17
6.4.3 Observation of the Ringing Instability	VI-20
6.4.4 Extended Period Runs	VI-22
6.4.4.1 Case 1	VI-25
6.4.4.2 Case 2	VI-26
6.4.4.3 Case 3	VI-26
6.4.4.4 Case 4	VI-27
6.4.5 Runs with Larger Expansion Ratios	VI-28
6.4.6 Shock-Tube Problems	VI-29
6.4.6.1 Case 1	VI-32
6.4.6.2 Case 2	VI-34
6.5 Summary	VI-34
References for Chapter 6	VI-38
Appendix of Chapter 6	VI-A1
Figures for Chapter 6	VI-A2
7 Two-Dimensional Hydrodynamic Simulations using Realistic Thruster Fields	VII-1

7.1 Formulation of the Algorithm	VII-3
7.2 Some Properties and Limitations of the Code	VII-7
7.3 Choice of Parameters	VII-16
7.4 Some Characteristic Parameter Ratios	VII-20
7.5 Discussion of Results	VII-23
7.5.1 Case 1	VII-25
7.5.2 Case 2	VII-36
7.5.3 Case 3	VII-37
7.5.4 Case 4	VII-41
7.5.5 Case 5	VII-43
7.5.6 Case 6	VII-44
7.5.7 Case 7	VII-45
7.5.8 Case 8	VII-46
7.5.9 Case 9	VII-47
7.5.10 Case 10	VII-48
7.5.11 Case 11	VII-50
7.6 Summary	VII-51
References for Chapter 7	VII-55
Appendix of Chapter 7	VII-A1
Figures for Chapter 7	VII-A2
8 Summary and Conclusions	VIII-1

Chapter 1

Introduction

Despite the difficulty of reaching the stars, there are no physical laws prohibiting humankind from traversing the great expanses of interstellar space, and one day we will most certainly become a starfaring civilization of the Galaxy. Today, interstellar flight no longer belongs to the realm of fantasy, but is a legitimate scientific and engineering challenge pushing against the limits of technology. As proponents of interstellar propulsion often state, the feasibility of interstellar flight is not a question of "can we?" but rather one of "when?" And all this rests upon how soon serious research of the physics underlying the operation of prospective interstellar propulsion systems is initiated.

The major problem with interstellar flight is that of propulsion. Of the space propulsion systems capable of very advanced missions such as ultrahigh performance interplanetary missions and possibly interstellar flight, but yet soundly based on currently understood principles of physics and within the grasp of near-future technology, the inertial confinement fusion (ICF) pulse rocket is one of the most promising. In an ICF pulse rocket, fusion debris plasma from an ICF microexplosion expands against, and is redirected by, a suitably configured vacuum magnetic field created by field coils attached to the vehicle (Chapter 3). Thrust is generated through electromagnetic interactions during this process. The physics involved in the expansion of a hot plasma against a vacuum magnetic field in the magnetic thrusters of ICF pulse rockets is generic to many other nonchemical and/or advanced propulsion schemes as well, because very high exhaust velocities are commonly obtained through the use of a hot plasma, and this generally necessitates some form of magnetic thruster. For example, certain types of antimatter

rockets may use magnetic thrusters not dissimilar to those of ICF pulse rockets. Furthermore, the subjects forming the foundation for these types of propulsion systems are topics of general interest in the arena of fundamental plasma physics. All this provides impetus for research on flows in the thrusters of ICF pulse rockets.

This thesis deals with numerical simulations of a hot plasma expanding against a vacuum magnetic field in configurations either of proposed magnetic thrusters for ICF pulse rockets or relevant to proposed designs.* It will be assumed that ICF will be successfully achieved and all analyses will start after the time when an ICF debris plasma has been created. Strong self-generated magnetic fields and plasma rotations (which can occur under the influence of magnetic fields varying in space and time, for example, when plasmas reflect off fields) that may be present in a real situation will not be considered. Possible ICF plasma instabilities will be ignored. A study of propulsive flows in these thrusters should be aimed at assessing the feasibility of thrust production by proposed designs, bringing to light problematic issues, and finding ways by which the propulsive efficiencies may be maximized. Currently, there is no assurance that the proposed designs will even function properly as envisioned, and much will depend upon design issues such as the choices of parameters and geometries. The fundamental operating principle of these thrusters, in which magnetic fields take the place of material walls, relies on the property that perfectly conducting fluids exclude external magnetic fields, and thus the entire concept of magnetic thrusters will not work if one does not have a highly ionized high-temperature plasma. Even before one begins treating plasma-field interactions, a major reconsideration

* Approximate analyses of the bulk flow that do not rely on numerical simulations were attempted. While crude estimates for bulk averaged quantities and plasma parameters could be made, modelings of realistic thruster geometries and flow situations using this approach often tend to get out of hand rather quickly if good accuracy is to be obtained. Part 2 of the appendix to Chapter 3 discusses a few examples of some simple estimations for numbers and approximate calculations.

of thruster parameters proposed to date may be necessary, reducing the plasma expansion ratio to keep the electrical conductivity at acceptably high levels.

The logical first step in initiating research in this field is to analyze the properties and motion of the bulk flow, a highly dynamic time-dependent process spanning a wide range of physical conditions, wherein the interface expands rapidly, deforming as it does so. Contrary to common belief, the study of bulk expansion of a fluid into a vacuum that exerts a pressure at the interface is not a well-studied phenomenon, differing from seemingly related problems in important aspects. Especially for flows in the thrusters of ICF pulse rockets, well-documented publications of fully hydrodynamic simulations apparently have not been published to date. The details of an earlier work, claimed to be the result of a hydrodynamical simulation of a flow in the magnetic thruster of an ICF pulse rocket is old and unavailable, to say the least. Plasma physics is a highly regime-dependent field, and without good intuition for the fluid properties (i.e., numbers) and flow geometries of the bulk flow, one cannot proceed. But at the same time, numerical analyses employing increasingly accurate approximations for describing a plasma become more and more restricted in the space and time scales that can be accommodated. It is therefore essential first to master the bulk flow hydrodynamics. Because of the complex nature of the bulk flow, details of real plasma behavior that are not part of the basic principles underlying the operation of a magnetic thruster for an ICF pulse rocket were omitted from the analysis. In particular, the plasma was approximated as an unmagnetized perfectly conducting medium, although a plasma does not behave exactly as a perfectly conducting fluid obeying the laws of classical fluid dynamics. Also employed were other simplifying assumptions such as azimuthal symmetry and the neglect of radiation and recombination.

In fact, before engaging in a fully hydrodynamic analysis, an even simpler approach was studied. Because of the deceleration of the interface and largeness of the velocity of interface expansion relative to the local thermal velocity, much of the plasma

will collect into a shell-like region, and the plasma may be very crudely approximated by a thin shell located at the interface with no internal fluid properties. This so-called thin-shell approximation, although not entirely realistic, allows one to circumvent the intensive memory requirements of fully hydrodynamic calculations for flows with very large expansion ratios. Making use of this property, a series of simulations employing the thin-shell approximation were run to obtain very rough estimates for the interface motion under a variety of realistic thruster setups, and the effects of varying parameter ratios and geometries on thruster performance were investigated (Chapter 4). Although the thin-shell approach has been pioneered by other authors, including a work dealing exactly with the problem of plasma expansion in the magnetic thruster of an ICF pulse rocket, the present work is believed to be one of the first thorough treatments of the entire flow, and with a comparison of various cases in parameter space.

Concerning the calculation of the magnetic field in the vacuum region, and thus the magnetic pressure at the interface, the plasma and field coil structures were assumed to be perfectly conducting, and a flux conserving prescription, valid under such a condition, was imposed. Namely, surface currents induced on the surfaces of perfect conductors prevent exterior magnetic fields from penetrating into them (a property of perfect conductors), and the initial flux through the field coils stays trapped between the surfaces of the plasma and the field coil structures throughout the plasma expansion process. Obviously, this phenomena can also be interpreted in terms of superposition of fields generated by the currents, including the induced surface currents; i.e., the superposed fields inside the initially unmagnetized perfect conductors cancel to zero, and the total flux between the plasma surface and the surfaces of the field coil structures stays fixed. In the actual computation, this condition was realized by letting the magnetic stream function value on the surfaces of the perfectly conducting plasma and field coil structures remain fixed during the entire flow process. Once the solution for the vacuum field is found, the magnetic pressure at the plasma interface can be calculated by taking the gradients of the field at the

interface. Because the motion of the interface is much slower than the speed of light, a quasi-static assumption was assumed valid for the field calculation.

Now the thin-shell approximation, by assuming all of the plasma to be collected into a thin shell without any internal fluid properties, and always located at the interface, is not only inaccurate, but also produces results with some unphysical flaws. In developing the increasingly complex, but more realistic hydrodynamical code for treating the plasma expansion, one of the central issues is to write the code so that it is able automatically to capture the interface between the perfectly conducting fluid and the vacuum magnetic field. The code should be able to advance the interface from time step to time step self-consistently so that a pressure match between the external magnetic pressure applying at the interface and the interfacial fluid pressure (with no magnetic field inside the plasma, the plasma pressure is just the fluid thermal pressure) will be met at all time steps, and without the need for trial-and-error iteration calculations, which can become difficult in multidimensions. This was handled through employing a Particle-In-Cell (PIC) scheme with appropriate prescriptions of boundary conditions at the interface. The PIC method, which combines the properties of both Eulerian and Lagrangian methods, is particularly suited for this purpose (Chapter 5). Shock-capturing techniques were also incorporated to handle shocks that arise from the deceleration of the interface.

The PIC algorithm was first tested in planar geometry, and its results were checked against analytical calculations (Chapter 6), with favorable outcomes. Valuable techniques for handling the flow correctly, as well as a better understanding of numerical artifacts that arise during the simulation, were obtained. As expected, a similarity in the behavior of the flow under various prescriptions was observed to exist between the one-dimensional planar geometry cases and their two-dimensional counterparts. Namely, the planar geometry calculations also proved themselves to be useful in predicting the multidimensional results

to a certain extent. Two-dimensional simulations** were, however, limited in the flow situations that could be handled because of memory limitations (Chapter 7). The flow still displays interesting features not uncoverable either by calculations employing the thin-shell approximation or by simple analytical estimates such as those in which the entire bulk is endowed only with uniform bulk-averaged quantities. For example, although the fluid is found to accumulate into a shell-like structure, this "shell" does not necessarily stay "attached" to the interface. Shock heating raises the interfacial temperature, contributing favorably (as long as the fluid is still fully ionized) towards keeping the electrical conductivity high there.

The work in this thesis is intended to serve only as a basis for further research, and is far from being an end in itself. In fact, although the bulk dynamics of the plasma motion, time-dependent bulk properties of the plasma, and propulsive efficiencies of the thruster may be found for idealized conditions by analyses of the type conducted in this work, some of the conclusions drawn, based on the simplifying assumptions used, may be overturned by a more accurate treatment of the plasma. For instance, a real plasma can allow external fields to diffuse into itself, undergo drifts across magnetic fields, and develop various instabilities, while at the same time, displaying a tendency to follow the field lines. Nevertheless, the idealized results obtained here may still serve the purpose of a baseline against which to compare more realistic solutions.

** In this work, "two-dimensional" flows will refer to physically three-dimensional flows with azimuthal symmetry. By assuming azimuthal symmetry, one does lose a freedom in direction, and some phenomena, such as certain instabilities that will be caught in a three-dimensional analysis, will not be able to be handled.

Chapter 2

Topics in Nonchemical Space Propulsion

Before embarking on a discussion of the ICF pulse propulsion system and the results of the various calculations that have been performed, it is appropriate to begin by presenting an introduction to the basic principles of space propulsion and a brief survey of the many nonchemical space propulsion systems that have been proposed to date.

2.1 Some Facts on Interstellar Distance Scales

To acquaint oneself with the distance scales involved in interstellar flight and therefore what will be required of an interstellar propulsion system, it may be best to look at a few examples. Table 2.1 gives the typical distance between major celestial bodies within our solar system. Large as these distances may seem, the solar system is only of minuscule dimensions on the scales of interstellar distances. As one heads outwards from the solar system, one may encounter Oort's Cloud at a distance of about 1.6 light years, but the nearest star (Proxima-Centauri) is currently 4.3 light years away from the sun. What this amounts to is that on a scale in which the earth's diameter is 1/13 th of an inch, the typical sun-to-Pluto distance will be one mile and Proxima-Centauri will be at a distance of 6850 miles from the sun. Figure 2.1 shows the distribution of stars in the immediate neighborhood of the sun. There are roughly ninety stars within a 20 light year radius of the sun and a few thousand within a 100 light-year radius. Double stars are common. Our Milky Way Galaxy is a spiral galaxy with an estimated diameter on the order of 10^5 light-

years (there is controversy as to the exact value of the diameter) and mass on the order of 2×10^{11} solar masses.

2.2 Fundamentals of Rocket Propulsion

The most fundamental equation governing rocket flight is the rocket equation, which gives the achievable velocity increment, ΔV , as a function of the exhaust velocity and mass ratio, where mass ratio is defined as the ratio of the initial mass of the vehicle to the final mass. The rocket equation in its simplest form is

$$\Delta V = u \ln(\mathcal{R}) , \quad (2.1)$$

where

$$\Delta V = V(t_{\text{final}}) - V(t_{\text{initial}}) , \quad (2.2)$$

u is the exhaust velocity defined relative to the rocket, and \mathcal{R} , the mass ratio, is

$$\mathcal{R} = M(t_{\text{initial}}) / M(t_{\text{final}}) . \quad (2.3)$$

More sophisticated forms of the rocket equation have been derived. For example, the equation can be written so that the useful payload mass (the final vehicle mass less the dead weight of the engine, tankage, etc.) appears instead of the final vehicle mass, with an effective exhaust velocity taking the place of the exhaust velocity.¹ For advanced propulsion systems, this effective exhaust velocity may be nonnegligibly lower than the exhaust velocity because of large structural dead weights. Although the classical rocket equation has been improved upon by several authors such as R.H. Goddard,² a more general form of the rocket equation, valid even when the rocket is traveling at relativistic speeds, can be obtained by applying special relativistic corrections to the addition of velocities. This relativistic form of the rocket equation has the form

$$\frac{\Delta V}{c} = \frac{\mathcal{R}^{2u/c} - 1}{\mathcal{R}^{2u/c} + 1} , \quad (2.4)$$

where ΔV is measured in the observer's frame, and u , by definition, is measured in the rocket vehicle's frame.³ The speed of light is denoted by c .

In rocket propulsion jargon, the exhaust velocity is often measured in terms of specific impulse (I_{sp}), which has the units of seconds. Specific impulse is related to exhaust velocity through

$$I_{sp} = \frac{\text{exhaust velocity}}{\text{gravitational acceleration @ sea level on earth}} \quad (2.5)$$

Thus the value of the exhaust velocity in meters per second divided by 9.8 gives the I_{sp} value.

Returning to the rocket equation (2.1), it can be seen that high final velocities of the rocket vehicle may be achieved by having either a high exhaust velocity or a high mass ratio. However, because the mass ratio appears in a logarithm, a rocket with a low exhaust velocity may require an astronomically high mass ratio to achieve the same final vehicle velocity a rocket with a higher exhaust velocity will be able to achieve with a reasonably low mass ratio. Furthermore, inspection of the rocket equation also shows that stopping at the destination (rendezvous mission) will square the required mass ratio, and returning will further square that mass ratio. The importance of high specific impulse for high ΔV missions cannot be overstated. Obviously, a low mass ratio also usually implies a high payload fraction. Table 2.2 illustrates the effect of specific impulse on mission performance. This shows that enormous mass ratios will be required for interstellar missions when values of specific impulse below those attainable by fusion propulsion are used. Multistaging allows higher values of ΔV to be achieved than single staging because stages are discarded as they burn out, but performance improvements brought about by multistaging fall short of the orders-of-magnitude improvements that will be made possible by selecting alternative energy sources (for example, nuclear versus chemical) and propulsion concepts (for example, solid core versus gas core).

The specific impulse is not the only important quantity determining rocket performance. Most notably, there are the issues of thrust, power, and propulsion system

mass. The thrust-to-weight (T/W) ratio and mass-to-power ratio (specific mass) are important parameters in determining the performance of a vehicle.⁴ While Newton's Second Law gives the relation between thrust, mass and acceleration of the vehicle, thrust can also be expressed as the product of the exhaust velocity and mass flow rate (\dot{m}) of the exhaust gases; namely,

$$\underline{F} = \dot{m} \underline{u} . \quad (2.6)$$

Ion rockets, for example, have a very high specific impulse relative to chemical rockets but have a much lower T/W ratio than the latter and are capable of only very low accelerations. However, for interplanetary missions to the outer solar system, electric rockets such as ion rockets do perform significantly favorably over chemical rockets despite their low T/W ratios because of their high specific impulses. Electric rockets can accelerate for long periods of time with low propellant consumption. Conventional versions of chemical rockets are theoretically limited to specific impulses of about 500 seconds, and with such low specific impulses, are practically incapable of accomplishing even moderately ambitious missions within the solar system. To take for example, the Apollo mission vehicles had a total liftoff mass of about 2900 tons, but of this, only about 100 tons were "useful" mass and the rest of the mass was mostly propellant mass. And all this was only for making it to the moon. But the high T/W ratio of chemical rockets, arising from high mass flow rates and low powerplant mass, does make them suited for launching vehicles off planetary surfaces. Many nonchemical propulsion systems do not possess this capability, and those that do (for example, nuclear thermal rockets) sometimes harbor safety and environmental problems.

In short, the T/W ratio is an indicator of the acceleration of the vehicle, determining the time required to achieve a given velocity, while specific impulse is a measure of the energy and propellant economy, i.e., the mass flow rate required to attain a given thrust. Often, the T/W ratio and specific mass quoted are those of the powerplant, and the T/W ratio for the entire vehicle will be much lower, and the specific mass, higher, if the

propulsion system mass is a small fraction of the total mass of the vehicle. Therefore, low specific-impulse systems with large propellant requirements are not always as advantageous as their powerplant T/W ratios may suggest.

Unfortunately, the property of high specific-impulse systems suffering from low T/W ratios and those with high T/W ratios suffering from low specific impulses is a prevailing trend among space propulsion systems. Simple algebra on Equation (2.8) shows that the T/W ratio is proportional to the inverse of the product of the specific mass and specific impulse, so to avoid this trend, a powerplant with a very different specific mass must be found. Nuclear pulse propulsion systems, through direct use of nuclear reactions, and in an external (usually) and intermittent way, are one of the fewer systems capable of providing both extremely high specific impulses and relatively high T/W ratios at the same time. They are low specific-mass drives.

Another quantity of some interest when discussing various propulsion systems using different energy sources is the mass-to-energy conversion efficiency of the energy producing reaction. The mass-to-energy conversion efficiency, ϵ , is here defined as the fraction of mass that is converted into energy in a reaction. For chemical reactions, ϵ is as low as 10^{-10} , while for nuclear fission, it is about 7×10^{-4} , and for nuclear fusion, it is typically in the neighborhood of 4×10^{-3} . Matter-antimatter annihilation reactions have $\epsilon=1$. The direct ejection of spent fuel is more efficient than keeping the spent fuel and ejecting separate propellant and accumulating "useless" mass onboard the spacecraft. Also, if power conversion systems are employed to convert the direct energy output from the fuel into, say, electrical energy, the efficiency and mass of the power conversion equipment degrade the performance of the vehicle. The most one can get out of any reaction is the direct channeling of all the energy produced by the mass deficit into axial kinetic energy of the reaction products. If it is assumed that the burnup fraction is unity (i.e., all of the fuel is utilized), all of the spent fuel is ejected, the exhaust stream consists only of spent fuel, and all of the energy from the mass deficit is converted into collimated kinetic energy of the

reaction products, then, by equating the energy corresponding to the mass deficit of the reaction to the kinetic energy, one obtains

$$u = c\sqrt{2\varepsilon - \varepsilon^2} \quad (2.7)$$

for the exhaust velocity. The finite divergence angle of the exhaust stream from the nozzle will cause the effective exhaust velocity to be lower. Now, if only part of the spent fuel is ejected, the available energy from the reaction will be channeled into a lower amount of mass (assuming that only spent fuel is ejected) so an exhaust velocity higher than that given by Equation (2.7) will be obtained. However, that part of the mass not ejected is wasted, and it can be seen that the final vehicle velocity achievable is not necessarily a monotonically increasing function of the exhaust velocity. For a given value of ε , indefinitely raising the value of specific impulse eventually degrades the performance of a rocket, and thereafter, higher specific impulses will be beneficial only if the energy producing reaction is switched to one with a higher ε . For example, using chemical reactions to heat a separate working fluid to obtain specific impulses well above 500 seconds or using a fusion reactor to drive a photon rocket are far from getting the most out of the fuel. In Figure 2.2, the rocket equation is plotted for exhaust velocities obtained by applying Equation (2.7) to various energy sources. The importance of using high specific impulse and thus high ε reactions for ambitious missions such as interstellar flight is clearly seen. The specific energies available from these reactions are listed in Table 2.3.

On the other hand, inert propellant can be added to spent fuel that is ejected. Examples of this are the addition of a cocoon of hydrogen gas flow around a pure fusion exhaust emerging from a magnetic fusion reactor in a continuous fusion rocket or the addition of expellant material around the target of each bomblet used in a nuclear pulse rocket. This will cause the energy available from the reaction to be distributed among a higher amount of mass than when only spent fuel is ejected, and the exhaust velocity will be lower than that given by Equation (2.7). But because the power supplied to the exhaust stream, i.e., the jet power, is given by

$$P = \frac{\dot{m}u^2}{2} = \frac{F_u}{2}, \quad (2.8)$$

it can be seen that when the mass flow rate is increased under constant power, the thrust is raised at the expense of lowered exhaust velocity. For the same mass ratio and power, this means quicker acceleration and shorter burn time than when the exhaust stream consists only of spent fuel particles and has a velocity given by Equation (2.7). Figure 2.3 shows a plot of the specific impulse as a function of the dilution ratio of the burned fuel in the exhaust stream for fusion rockets. The degradation of specific impulse with increasing inert propellant fraction in the exhaust is evident. One sees that with high fuel burnup fractions and minimal nonfuel expellant, a fusion rocket is a potent candidate for interstellar missions. Needless to say, for any particular mission, maximizing the final coast velocity is not necessarily the most desirable strategy.⁴ Compromises must be made between various factors such as trip duration, payload fraction, etc., and a detailed mission analysis must be carried out. However, that is beyond the scope of this thesis and will not be discussed except for mentioning that for interplanetary missions, substantial amounts of propellant (e.g., hydrogen) should be added to a direct fusion exhaust to raise the thrust at the expense of lowered specific impulse in order to shorten trip times, and for interstellar travel, the higher the specific impulse and the higher the final cruise velocity, the better. Only energy sources with high ϵ should be considered for interstellar propulsion.

Once a particular energy source is harnessed, ϵ becomes fixed, and the issues shift to those of power containment and conversion of isotropic power into directed jet power. The physics of thrust producing processes and engineering design considerations place limits on the performance of a propulsion system. For example, nuclear thermal rockets, although in part deliberately, are far from producing the high specific impulses that are theoretically possible with nuclear fission.

2.3 Brief Survey of Nonchemical Space Propulsion Systems

In this section, some representative nonchemical space propulsion systems that have been proposed to date will be reviewed very briefly. Introduced will be nuclear thermal rockets, beamed thermal rockets, electric rockets, photon sails, continuous fusion rockets, antimatter rockets, and interstellar ramjets. As nuclear pulse rockets will be described in the next chapter, they will not be presented in this section. Figure 2.4 presents rough estimates for the specific impulse, powerplant T/W ratio, and powerplant specific mass of these propulsion systems. For the photon sail, which by definition has infinite specific impulse, the specific mass and T/W ratio were used in positioning the device on the plot. The rectangle enclosing the performance region of antimatter rockets was constructed from very rough, upper and lower bounds drawn parallel to the vertical and horizontal axes of the plot.

2.3.1 Nuclear Thermal Rockets

The best-known type of rocket is probably the thermal rocket, which relies on some energy source to heat a working fluid which is then expanded out through a nozzle to produce thrust.⁵ The heating source can be chemical, nuclear, electric, solar radiation, laser radiation, etc. Because of material constraints, the use of material heat-exchanger type reactors and material nozzles places rather conservative limits on the temperatures that may be attained by the propellant gases. Therefore, the specific impulses of conventional nonchemical thermal rockets, although higher than those achievable by chemical rockets, generally fall quite short of the specific impulses achievable through electrostatic or electromagnetic acceleration of propellant or through direct exhaust of a fusion plasma via a magnetic nozzle.

In particular, the nuclear thermal rocket uses heat released from a high-power-density nuclear fission reactor to heat inert propellant, which is subsequently exhausted through a conventional convergent-divergent nozzle. Effectively, a fission reactor is being placed inside a thrust chamber. The use of nuclear reactions allows exhaust velocities much higher than those that can be obtained by chemical reactions to be attained. The combustion flame temperature places a limit on the temperature to which a gas may be heated by a chemical reaction. Also, the propellant is separate from the fuel generating the heat energy in a nuclear thermal rocket; this independence of the working fluid from the energy source allows the former to be chosen rather freely without being restricted by combustion characteristics. That is not the case for chemical rockets, although this is not to say that there are no concerns for propellant density, dissociation characteristics, chemical stability and compatibility, etc. For a rocket that uses a convergent-divergent nozzle to convert the thermal energy of hot gases in a thrust chamber into directed jet kinetic energy, the exhaust velocity varies as

$$u \propto \sqrt{\frac{T}{MW}}, \quad (2.9)$$

where T is the temperature of the gas in the chamber and MW is the average molecular weight. Therefore, most nonchemical thermal rockets which heat propellant through a heat-exchanger type process have been proposed with hydrogen as propellant, although hydrogen does have corrosive properties at high temperatures. The high heat capacity of hydrogen can also be advantageous. Equation (2.9) clearly shows that for thermal rockets, the temperature to which the working fluid is heated must be raised to obtain an increase in the exhaust velocity. Thus, to attain specific impulses higher than those corresponding to the structural temperature limits of solid reactor cores, the fuel will have to be in particulate form (alleviated structural integrity requirements and increased heat transfer areas), and for even higher specific impulses, temperatures at which the nuclear fuel will be in a molten or even gaseous phase must be considered. Depending upon the state of the nuclear reactor core, nuclear thermal rockets are classified as solid core, particle-bed core, colloid core,

liquid core, or gas core rockets. Hybrid designs have also been proposed. The specific impulse rises progressively from solid core rockets to gas core rockets, with the former capable only of offering values starting as low as 800 seconds and the latter capable of values possibly as high as 5000 to 7000 seconds (a more conservative estimate is about 3000 seconds, especially without radiators). The T/W ratio for a nuclear thermal rocket is high, and is generally in and around the range of 0.1 to 10. With the exception of solid core rockets, cavity type reactors are employed, and in such advanced designs the hottest parts of the core and working fluid are separated away from material walls. However, for gas core reactors, which operate at the upper end of the core temperature spectrum, radiation from the core becomes a limiting issue. Multiple core designs can improve the performance. Except for the solid core rocket and the nuclear light bulb (closed-cycle gas core) rocket, the propellant passes through nuclear fuel in a particle bed, molten, or gaseous phase, and fluid dynamical or centrifugal, mechanical core-rotation schemes must be used to retain preferentially the fuel particles in the core. For plasma core reactors, fuel retention may also be aided by electromagnetic means. In the nuclear light-bulb rocket, complete fuel containment is made possible by encapsulating the burning fuel and a buffer gas inside transparent walls through which radiation can escape.

When propellant gas temperatures are raised much higher than those typical of solid core rockets (e.g., about 3000°C or higher), frozen-flow losses, i.e., losses that are due to dissociation energy removed from a gas flow's not being returned (through recombinations) timely in the form of directed bulk kinetic energy, can become an important loss mechanism. Frozen-flow losses are very pressure dependent. Hydrogen does have large frozen losses. Frozen losses that are due to ionization become important above temperatures corresponding to specific impulses of about 2000 sec.

A propulsion concept somewhat related to the nuclear thermal rocket is the radioisotope heater rocket, which uses a radioisotope-bearing core to heat a propellant. The thrust obtained is typically low.

More exotic is the fission-fragment propulsion (or sometimes direct fission propulsion) system, which obtains thrust via the direct emission of fission fragments from a sheet of fissionable material. Radioisotope sails that obtain thrust from the emission of particles in a radioactive decay have also been proposed. Obviously, these two propulsion schemes do not fall into the category of thermal rockets but have very high specific impulses comparable to those of fusion rockets and low T/W ratios similar to those of electric rockets.

2.3.2 Beamed Thermal Rockets

In the beamed thermal rocket, energy to heat the inert propellant comes in the form of electromagnetic radiation beamed from a remote source. This beam can be, for example, a laser beam, microwave beam, or even solar radiation. These rockets have the merit of not having to carry their own power source and have no need for power conversion equipment. Typically, such a propulsion system consists of a collector (adaptive optics, nonadaptive rigidized optics, or inflatable optics, with the former two appropriate for laser radiation endowed with low angular divergence because of their coherence) that receives and focuses the incoming radiation, a cavity into which this beam is focused into and where energy transfer to the propellant occurs, and a conventional convergent-divergent nozzle through which the heated gas is expelled. Radiation energy is coupled to the propellant by various means, depending upon the type of beam, and these include seeding of the propellant, molecular resonance, inverse bremsstrahlung with a laser-supported plasma (laser-supported combustion wave heating), a microwave-supported self-sustaining plasma, and a solid heat-exchanger cavity, some of which are windowless designs. Heat transfer and

flow mixing are important design issues. Although in many designs the heating of the propellant occurs in a localized region away from the walls and aerodynamic windows have also been proposed, temperature limitations on material walls along with (re)radiation losses, which increase with temperature, become limiting factors in obtaining high specific impulses, as was the case with nuclear thermal rockets, and the specific impulses obtained by these rockets are similar to those achieved by nuclear thermal rockets. With laser thermal rockets, specific impulses as high as 4500 seconds may be theoretically possible, but lower values of 2500 seconds or less are more realistic. Because of the rather low energy densities in the beams (as compared to those present in a fission reactor) and the limitations imposed by the collector size, structural strength, and mass, the T/W ratio of beamed thermal rockets will be lower than those of nuclear thermal rockets, certain types of laser thermal rockets excepted, although higher than those of electric rockets. Generally, performance between chemical and electric rockets will be obtained. One application to which these rockets are suited for is orbital-transfer flight.

In laser thermal drives, repetitively-pulsed devices have been proposed along with the more conventional, continuous-wave devices. In a pulsed laser thermal rocket the laser beam typically enters the nozzle directly without going through an absorption chamber and explosively heats the propellant. This variant of the laser thermal rocket propels itself by a succession of laser-supported detonation waves (which become blast waves after the laser pulse terminates) formed in the nozzle. These repetitively-pulsed thrusters do not obey the conventional relations governing thermal rocket nozzle performance because ideal nozzle flow does not exist in these thrusters, and for example, the specific impulse will be proportional to the inverse fourth root of the propellant molecular weight.⁶ Use of magnetic nozzles in (pulsed) laser propulsion have been suggested, and preliminary experiments have yielded favorable results of enhanced axial collimation of the exhaust.⁷

2.3.3 Electrothermal Rockets

In the electrothermal rocket, electrical energy is used to heat an inert propellant, which is then exhausted through a conventional convergent-divergent nozzle to produce thrust. The electrical heating of the gas is accomplished through use of resistive heating elements in the resistojet and by passage of the gas through an arc discharge in the arcjet. The latter method can heat the propellant to higher temperatures and thus can achieve higher specific impulses by virtue of not being limited by the softening temperature of the heating elements. The same material limitations present in other conventional thermal rockets limit the specific impulse of electrothermal rockets to 1000 to 2000 seconds. Unlike other electric thrusters, the efficiency and performance of electrothermal thrusters are limited by convective (and some radiative) heat transfer loss from the hot propellant to the engine walls.

The electron-cyclotron-resonance thruster can be considered an electrothermal thruster as electromagnetic energy is used to excite cyclotron motion (basically thermal energy), which is then converted into axially directed flow via a $-\mu\nabla_{\parallel}B$ body force felt in the diverging magnetic fields of a magnetic nozzle.⁸ It is an example of an electrodeless device which has the advantage of circumventing electrode erosion problems.

Because of power generation and conversion/processing (an extra step necessary in electric propulsion systems) by heavy onboard equipment, electric propulsion systems generally suffer from high specific mass. Standard power sources are nuclear reactors and beamed energy, including solar radiation.

2.3.4 Electrostatic Rockets

While electrothermal rockets generally do not rely on electric forces to accelerate propellant, both electrostatic and electromagnetic rockets do, and therefore belong to

different classes of rockets from thermal rockets.⁹ Since the acceleration of propellant in electrostatic and electromagnetic rockets is not achieved by converting thermal energy of a gas into kinetic energy by flow out through a nozzle, the specific impulse is not limited by temperature. Very high specific impulses on the order of 10^4 seconds can be obtained. Also nonthermal acceleration means that the criterion for selecting the propellant to maximize the specific impulse will be different from that described in Section 2.3.1. However, the propellant flow must operate at very low densities so that electric effects will be the dominant accelerating mechanism. Characterized by high exhaust velocity, low mass flow rate, low thrust, and excessive mass of the power system equipment, the T/W ratio for these rockets tends to be very low in the 10^{-3} to 10^{-5} range.

In the electrostatic rocket, neutral propellant is ionized and then accelerated by an electrostatic field established between electrodes. The flow of ions out of the thruster must be neutralized by electrons from an electron emitter. The ions can be created via bombardment by electrons flowing from the cathode to anode inside the thruster (bombardment thruster), or by passing the propellant through a porous high-work-function material (contact thruster).

Other than the acceleration of ions (ion rockets), electrostatic acceleration of colloids is also possible (colloid rockets).

2.3.5 Electromagnetic Rockets

While electrostatic rockets use the $\underline{F} = q\underline{E}$ force to accelerate charged ions (or colloids) of one polarity, electromagnetic rockets use electromagnetic forces to accelerate a quasi-neutral plasma. Both continuous-operation and pulsed-operation devices can be built. A classic electromagnetic rocket is the MPD (magnetoplasmadynamic) rocket. In this rocket, the thruster comprises a coaxial outer and central electrode between which there is an annular region through which propellant flows. The propellant is ionized by passing

through an arc connecting the cathode and anode, and the arc current interacts with its self-induced magnetic field to produce a $\underline{J} \times \underline{B}$ body force, which then accelerates the plasma out of the thruster. MPD rockets generally produce higher thrusts than ion rockets, but are limited to somewhat lower specific impulses.

The use of the $\underline{J} \times \underline{B}$ force to obtain thrust is found in other electromagnetic drives such as the rail gun and mass driver, although these particular devices are often not envisioned in the form of an onboard thruster.

On the other hand, plasma in crossed electric and magnetic fields experiences an $\underline{E} \times \underline{B}$ drift, and this can be put to use as an acceleration mechanism for obtaining thrust. A propulsion device using this scheme is the MHD thruster.

2.3.6 Photon Sails

The photon sail does not obtain thrust through the ejection of propellant, but rather through the reflection of photons from a remote source. A photon carries momentum $h\lambda$ and thus can impart momentum to a body upon reflecting off its surface, but for reasonable accelerations to be felt, the ratio of the mass to the receptive area of the body must be very low. The propulsion system consists of a sail fabricated of materials such as thin plastic films coated with reflective materials, and appropriate rigging. The sails tend to be large and light structures. Typical solar sails that have been proposed for near earth or interplanetary missions have characteristic dimensions on the order of a few hundred meters to a few kilometers. Sail performance deteriorates as sail diameter increases, apparently because of the increase in cable mass. In some designs, centrifugal forces are used to keep the sails stretched (spin stabilization). Depending upon the radiation source, photon sails can be classified as solar sails, laser sails, etc. For low mass loss stars such as the sun, radiation pressure is very much stronger than pressure that is due to the stellar plasma wind, and a solar sail relies solely on radiation pressure to obtain thrust. However,

making use of the fact that the solar/stellar wind is a plasma, electromagnetic means such as an onboard current flow should be able to harness the solar/stellar wind for propulsion. In fact, such a concept was finally discussed recently under the name of "magsails."¹⁰ Photon sails are not necessarily an efficient means of propulsion when the vehicle speed is much less than the speed of light, but can compete favorably with electric rockets for interplanetary missions. Typically, outbound missions first approach the sun before accelerating outwards. Hyperthin and perforated high-temperature lightweight solar sails, if developed, could offer attractive performances, even for interstellar missions. Multistage laser sails with round-trip capabilities and aided by Fresnel lenses placed in the beam path have been proposed for interstellar missions.¹¹ Ultralight microwave wire-mesh sails may also be possible.¹²

Although not a radiation pressure sail, pellet stream propulsion in which a stream of material pellets originating from a remote source is intercepted by a pusher plate attached behind a vehicle, has been suggested.¹³ The pellets can be ionized before impinging upon the vehicle, and magnetic fields could then be used to aid in the reflection of the propulsive pellet stream.

2.3.7 Continuous Fusion Rockets

For obtaining very high specific impulses, the use of fusion is attractive because of the high temperatures of a fusion plasma. In the continuous fusion rocket, fusion plasma inside a magnetic fusion reactor is to be expelled to provide thrust ("continuous" because of the steady flow of plasma from the rocket nozzle). For example, a magnetic mirror reactor is an open-ended system relying on the property of conservation of an adiabatic invariant, the magnetic moment, to confine particles within the reactor. However, particles in the loss cone in velocity space will leak out the open ends of a mirror machine. By strengthening

the mirror ratio for one end of the cell over that for the other end, i.e., by making the field at one end stronger than that at the other end, preferential ejection of the reactor plasma will be made possible, leading to thrust production. With a toroidal reactor, a magnetic divertor could be used to extract plasma out of the reactor, and this plasma could then be allowed to flow out along the fields of a magnetic nozzle as in the case of a mirror machine rocket. The use of magnetic nozzles, made possible by virtue of the working fluid's being a conducting plasma, allows wall contact of a hot plasma to be avoided and alleviates material imposed temperature limitations. With fusion plasma temperatures on the order of 10^9 K, specific impulses above 10^5 seconds will be possible. However, magnetic fusion reactors will be very heavy machines because of the requirements for magnetic confinement, with excessively massive field coils, supports, shielding, and cryoplant. Therefore, the T/W ratio of continuous fusion rockets may possibly be as low as 10^{-4} to 10^{-5} . Suggestions have been made to develop and employ compact high-energy-density reactors. Throttling of the engine can be accomplished by adjusting the flow ratio of the fusion plasma to the inert propellant, as alluded to earlier.

Whereas the specific impulses of ion rockets render them unattractive as interstellar propulsion systems unless extremely slow interstellar arks are considered, multistage fission or fusion systems with high fuel burnup fractions and direct ejection of spent fuel may theoretically be able to reach nonnegligible fractions of the speed of light.¹⁴

2.3.8 Antimatter Rockets

A matter-antimatter annihilation reaction offers complete conversion of mass to energy and is the most ideal energy source for space propulsion, and especially for difficult high specific-impulse missions.

The annihilation of electrons and positrons results only in the immediate production of photons (γ -rays) and is not the favored reaction for propulsion applications. The

classical photon rocket that produces thrust by emitting a powerful photon beam is not practical with envisionable technologies of the near future. In such a rocket, an isotropic light source is placed at the focus of a concave mirror to produce a collimated beam of photons. However, the use of photons as "propellant" necessitates an enormous power input even for obtaining tiny thrust. The mirrors will have to have reflectivities far higher than those technically conceivable of today in order to prevent an instant evaporation at the power levels required, although exotic mirrors such as those fabricated out of electron gases have been proposed.

The term antimatter rockets usually means those that utilize intermediate particles produced by proton-antiproton reactions. Proton-antiproton annihilation reactions do not immediately result in the production of photons. Energetic pions and muons, some of them charged, are produced before decaying into photons, and thus, for example, proton-antiproton reactions can be used conveniently to heat an inert propellant that can then be expanded out through a conventional nozzle, or be used to create and maintain a plasma in a magnetic bottle without ever requiring a sustained thermonuclear reaction, and the plasma can then be extracted through a magnetic nozzle. Alternatively, the charged intermediate particles from a proton-antiproton explosion could be directly formed into a collimated propulsive stream by a magnetic nozzle. As such a variety of thrust producing schemes are possible, an extremely wide range of specific impulses and T/W ratios are obtainable. The maximum possible specific impulse will be higher than those of fusion rockets. One gram of antiprotons can produce an energy equivalent to 20kT of TNT, and only trace amounts of antimatter will be required for a typical mission unless a photon rocket is considered. Calculations have shown that virtually any mission with a final velocity less than $0.5c$ can be accomplished with a mass ratio of 5 or less.¹⁵ Although present antiproton production efficiencies must be improved by several orders of magnitude and storage problems must be solved along with thruster design issues, antimatter rockets, once built, will entirely revolutionize space travel from near-earth missions to interstellar flight.

2.3.9 Interstellar Ramjets

One of the most pronounced problems with interstellar flight is the enormous amount of propellant that must be carried onboard the vehicle to accomplish the mission in a reasonable amount of time, resulting in huge mass ratios. The interstellar ramjet circumvents this problem by collecting the fuel-propellant working fluid along the way.¹⁶ Interstellar gas, which is mostly hydrogen, is collected by a colossal magnetic scoop after it is ionized by some means, if it already isn't, and is funneled into a fusion reactor which produces an energetic exhaust out of it. However, because of the low densities of the interstellar medium, the parameters for the ramjet vehicle such as the required field strength and sheer size of the magnetic scoop become truly astronomical. Furthermore, the fundamental operating principles of the propulsion scheme have not been worked out properly. For example, the magnetic scoop will serve as a magnetic mirror reflecting charged particles away from the scoop instead of collecting them, and a self-sustaining p-p reaction may be difficult to achieve in a small system such as a starship reactor.

Several variants of the interstellar ramjet have been proposed such as the ram-augmented interstellar rocket, which carries its own fuel for the fusion reactor but obtains interstellar hydrogen for propellant; and the laser-powered interstellar ramjet, which obtains energy from a remote laser source to accelerate the incoming interstellar matter.

Although relativistic velocities will be possible with the interstellar ramjet, they are still closer to the realm of science fiction than engineering physics. However, using resources present in interstellar space to propel a vehicle is too intriguing to dismiss, and for example, it may be possible to utilize cosmic magnetic fields in some way to produce thrust.

References for Chapter 2

1. See, for example, R.S. COOPER, *J. of the Aerospace Sciences* **29** (1962), 1339.
2. See, for example, J.H. BLOOMER, "The Alpha Centauri Probe," p.225ff, in *Problems of Propulsion and Reentry : Proc. 17th Int'l Astronautical Congress, Madrid, 1966*, ed. M. LUNC (Gordon and Breach, New York, 1967).
3. See, for example, J. ACKERET, *Helvetica Physica Acta* **19** (1946), 103
or B.M. OLIVER, *JBIS* **43** (1990), 259.
4. See, for example, W.E. MOECKEL, *J. Spacecraft* **9** (1972), 863
or C. POWELL and O.J. HAHN, *JBIS* **26** (1973), 334.
5. There is a large body of literature in this field, including a book mainly on solid core designs, R.W. BUSSARD and R.D. DeLAUR, *Nuclear Rocket Propulsion* (McGraw-Hill, New York, 1958). For more advanced designs, an example of an introductory paper is R.S. COOPER, *Astronautics and Aeronautics* **4** (1966), No.1, 54.
6. G.A. SIMONS and A.N. PIRRI, *AIAA J.* **15** (1977), 835.
7. Y. ARAKAWA and K. YOSHIKAWA, *Space Power* **7** (1988), 17.
8. Private communication with F.E.C. CULICK.
9. There is a large body of literature in the active and established field of electric propulsion. The classic text here is R.G. JAHN, *Physics of Electric Propulsion* (McGraw-Hill, New York, 1968).
10. See, for example, D.G. ANDREWS and R.M. ZUBRIN, *JBIS* **43** (1990), 265.
11. R.L. FORWARD, *J. of Spacecraft and Rockets* **21** (1984), 187.
12. R.L. FORWARD, *J. Spacecraft* **22** (1985), 345.
13. C.E. SINGER, *JBIS* **33** (1980), 107.

14. D.F. SPENCER and L.D. JAFFE, "Feasibility of Interstellar Travel," p.409ff, in *Advanced Propulsion Concepts : Proc. 3rd Symposium on Advanced Propulsion Concepts, Cincinnati, Ohio, Oct. 1962*, vol. 1 (Gordon and Breach, New York, 1963).
15. See, for example, R.L. FORWARD, "Antiproton Annihilation Propulsion," AFRPL-TR-85-034 (1985).
16. R.W. BUSSARD, *Astronautica Acta* **6** (1960), 13.

Appendix of Chapter 2

Part 1: Tables for Chapter 2

(Page II-A2)

Part 2: Figures for Chapter 2

(Starts Page II-A3)

Part 3: References for the Appendix of Chapter 2

(Page II-A6)

Earth - Moon	1.3 light seconds
Earth - Sun	8.3 light minutes
Sun - Mars	13 light minutes
Sun - Jupiter	43 light minutes
Sun - Saturn	1.3 light hours
Sun - Uranus	2.7 light hours
Sun - Neptune	4.2 light hours
Sun - Pluto	5.5 light hours

Table 2.1 Typical Distances within the Solar System.

Specific Impulse (sec)	Mass Ratio
5×10^4	7.92×10^{26}
5×10^5	4.90×10^2
5×10^6	1.86
3.06×10^7 (photons)	1.11

Table 2.2 Mass Ratios for 40-Year Earth-Observer-Transit-Time Flyby Mission to 4 Light Years with a 1g Acceleration Period (lasting about 0.1 year) followed by a Coast Phase.¹

Reaction Type	Potential Specific Energy (Cal/g)
Chemical	3.6×10^3
Nuclear Fission	1.7×10^{10}
Nuclear Fusion	1.8×10^{11}
Matter-Antimatter	2.2×10^{13}

Table 2.3 Typical Specific Energies Available from Reactions.

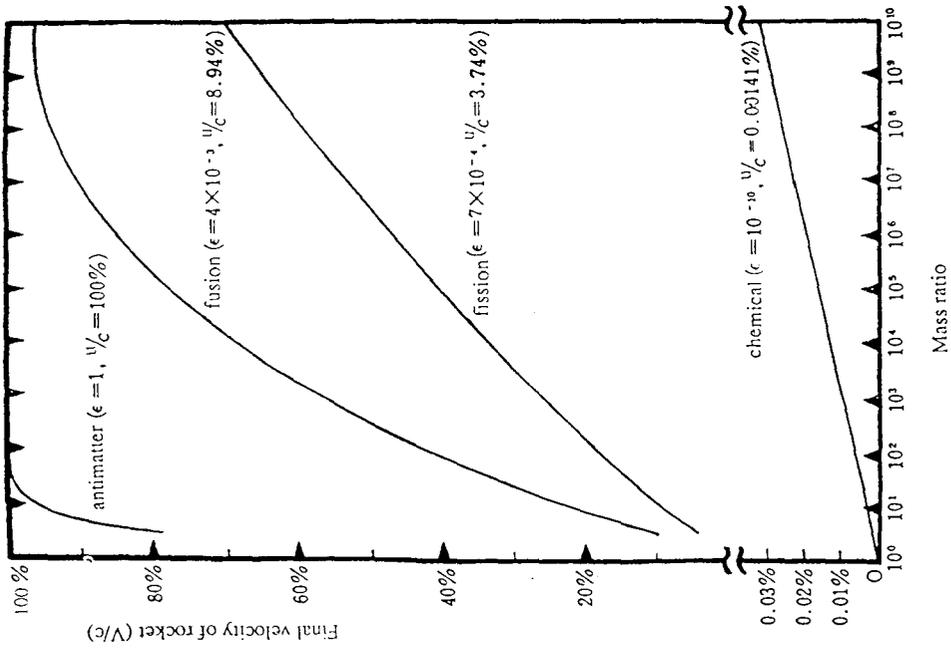


Figure 2.2 3

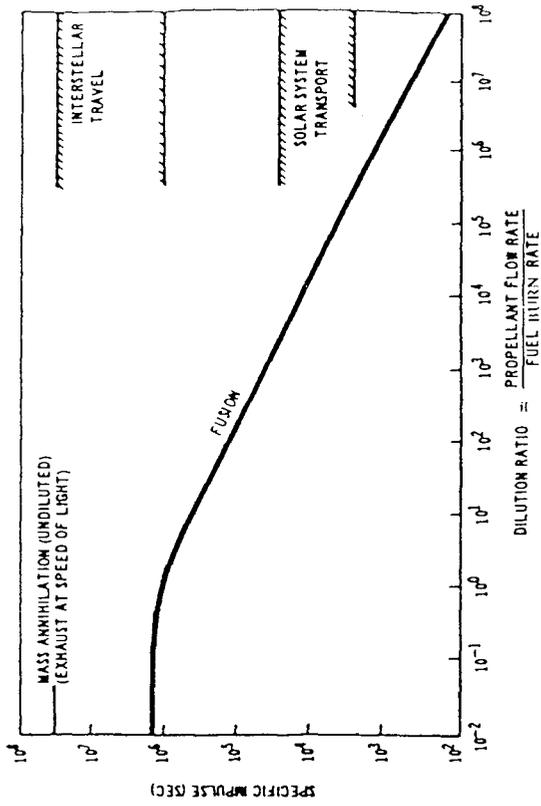


Figure 2.3 4

The asterisk on the vertical axis marks the I_{sp} corresponding to that of photons.

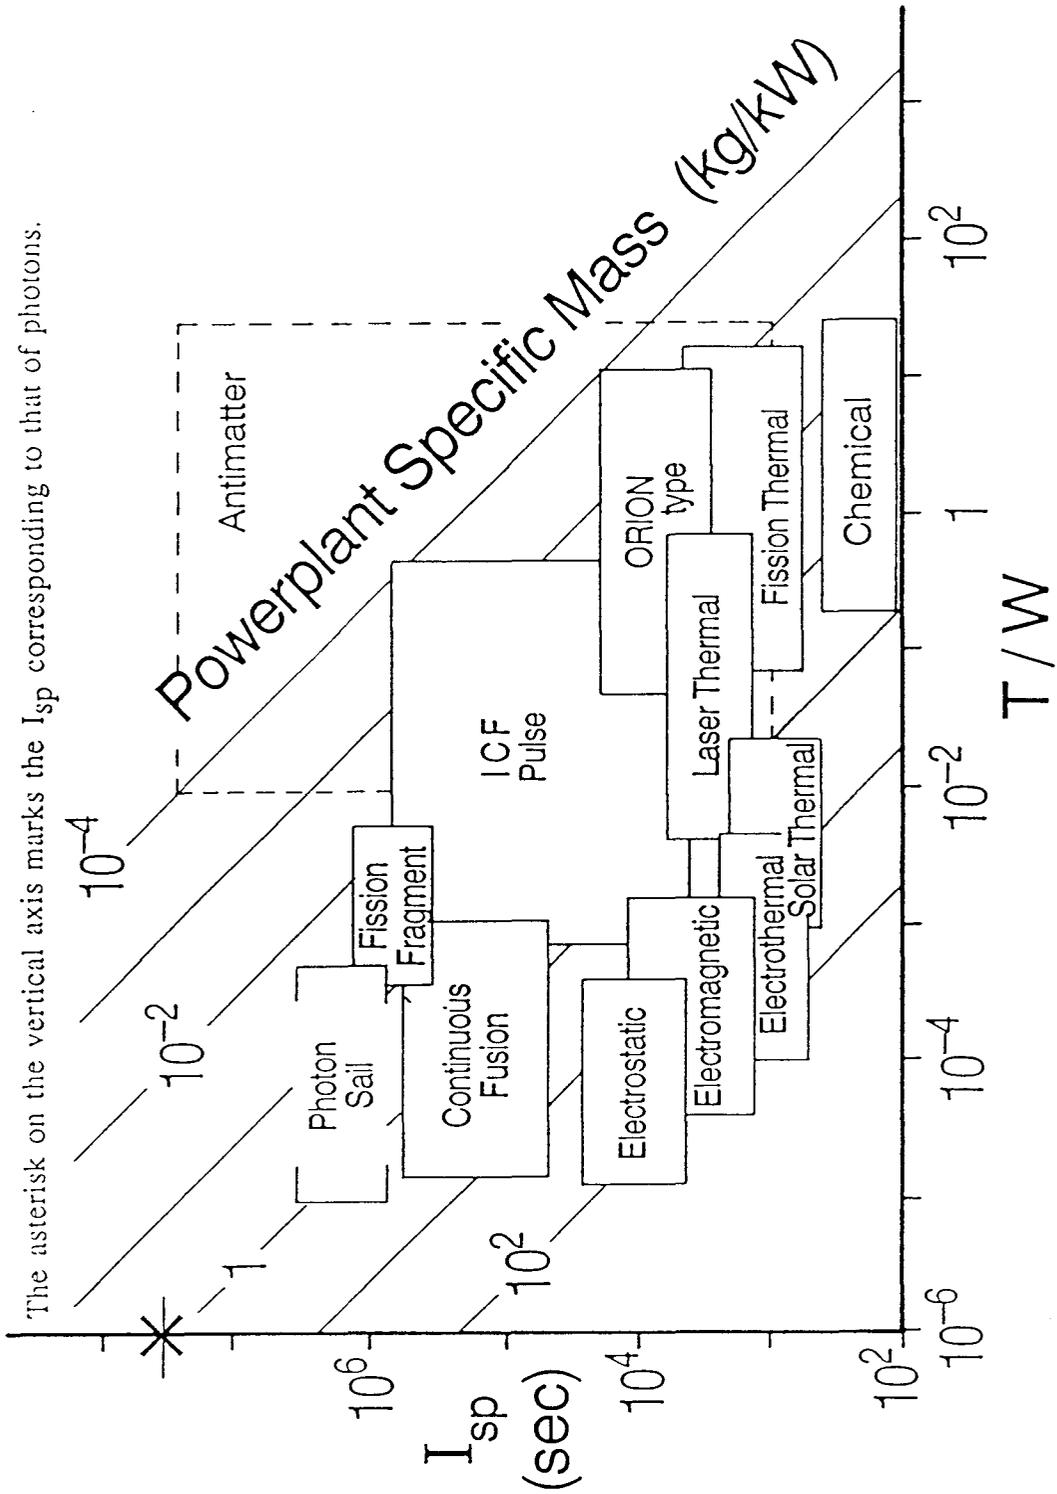


Figure 2.4

Part 3: References for the Appendix of Chapter 2

1. B.N. CASSENTI, *JBIS* 35 (1982), 116.
2. J.B. KALER, *Astronomy*, August 1991 issue, 26, (Kalmbach Pub., Waukesha).
3. Adapted from F. ISHIHARA, *Ginga Ryoko : Koseikan Hiko wa Kano ka*, (Kodansha, Tokyo, 1979).
4. Modified from M.W. HUNTER II, "Accessible Regions Beyond the Solar System," AAS-69-386, in *AAS Advances in Astronautical Sciences* (1969).

Chapter 3

Introduction to ICF Pulse Propulsion

3.1 Nuclear Pulse Propulsion

The nuclear pulse rocket, of which the ICF pulse rocket is a version, is a bomb propulsion system which obtains thrust from the intermittent interaction of the vehicle with debris from bomb(let)s exploded behind the vehicle. The concept was first introduced by Ulam of Los Alamos Scientific Laboratories in the mid 1940's.¹ See the literature for a history of the development of various nuclear pulse propulsion concepts.²

The direct use of nuclear reactions, with orders-of-magnitude higher mass-to-energy conversion efficiencies than chemical reactions, allows extremely high specific impulses to be achieved. Based on specific energies available from nuclear reactions, specific impulses as high as 10^6 seconds are theoretically possible with fusion.* At the same time, separation of the energy production process away from the vehicle and employment of intermittent pulses with short interaction times enable high T/W ratios to be obtained. In fact, T/W ratios on the order of unity or higher may be possible. This capability of offering both extremely high specific impulse and high T/W ratios simultaneously is a property that sets nuclear pulse rockets apart from many other space propulsion systems. Another advantage is that with a pulsed propulsion system, throttling

* Although limited fuel burnup fractions and fractional energy deposition into charged particles (only charged particles can be redirected by a magnetic field to provide thrust) have been taken into account, such high values of specific impulse may be difficult to achieve with currently envisioned designs of ICF pulse drives because a large portion of an ICF pellet mass must be in nonfuel materials. However, specific impulses on the order of 10^5 seconds should be well within the capabilities of ICF pulse rockets, and this should be adequately high for most purposes except interstellar missions.

(for thrust and power) can be readily achieved, and over an enormous range, by varying the pulse rate.

An early class of nuclear pulse rocket design employed a material pusher plate connected to the rest of the vehicle by struts and shock absorbers behind which nuclear explosions would be detonated. The expanding debris from the explosions would physically impact the ablation layer of the pusher plate and impart momentum to the vehicle. The bombs would typically be tamped with inert expellant not only to adjust the thrust level but also to absorb and smooth out the momentum of the explosion before impacting the pusher plate. This is the so-called external nuclear pulse drive and was the propulsion system for the Orion Project vehicle, which underwent serious research efforts in the late 1950's and early 1960's. Some typical parameters for a representative Orion vehicle are specific impulses in the range of 1800 to 6000 seconds (tailored for near-earth and interplanetary missions) and T/W ratios ranging from below 0.1 to as high as 15, with pulse rates ranging from below 0.1 Hz to about 1 Hz (limited by the time the impulse absorber system takes to return to its original position). Propulsion by a series of bombs detonated behind a pusher plate was demonstrated by a small experimental vehicle using chemical explosives.

In another class of nuclear pulse rockets, the nuclear explosions take place inside a large chamber. The chamber may be filled with propellant gas that is heated by the explosion, and the resulting shock waves and/or the bomblets may be surrounded by momentum absorbing expellant material. Both methods protect the chamber walls. In the so-called internal nuclear pulse drive, the chamber completely surrounds the explosion site except for an opening connected to a nozzle through which the energized propellant is expelled. Compared with the external pulse drive, better nuclear energy utilization and higher thrust are possible, but they are very much limited to lower specific impulses. Internal systems suffer from higher vehicle mass and more radiation heating of the vehicle than external systems. Unlike the external system, the internal system does not harness all

the potential advantages available to the nuclear pulse concept. Among the investigations of the internal pulse drive were those carried out under the Helios Project. The impulse absorbing chamber may also be cavity-shaped (e.g., hemispherical or bell-shaped) instead of entirely surrounding the explosion site; this would allow higher specific impulses to be attained than is possible with the entirely internally pulsed device and without as much focusing problems as the entirely externally pulsed device.

These concepts were initially proposed for nuclear fission, but the use of fusion is also possible.

Nuclear fission requires a critical mass and the bombs, and thus the pulse units also tend to become large. The vehicles must be very large systems to absorb the momentum efficiently (primary concern for moderate specific-impulse interplanetary vehicles) and energy (primary concern for high specific-impulse interstellar vehicles with high heat loads), and in fact, the efficiencies of nuclear pulse rockets go up with size. Besides, the idea of using material chambers or ablation-type pusher plates linked to shock absorbers to intercept debris from a nuclear explosion is a somewhat awkward concept and suffers from disadvantages such as radiation damage and limited specific impulse. However, with the introduction of the concept of ICF, it was realized that it would be possible to ignite small fusion bomblets or pellets. Furthermore, the relatively low energy of such ICF microbombs would make possible the replacement of material-made pusher plates by suitably configured magnetic fields of envisionable strengths as long as the debris plasma stays highly conducting.** The magnetic fields comprising the thruster would be created by current coils rigidly attached to the vehicle. This was the birth of the concept of the currently envisioned form of ICF pulse rockets. The use of magnetic fields instead of material pusher plates or chamber walls to interact with, deform, and redirect the hot plasma alleviates material limitations such as those regarding temperature and allows higher

** Reference 1 already suggested the possibility of using magnetic fields to aid in the protection of material pusher plates and thereby making possible higher specific impulses as well.

specific impulses to be achieved. The magnetic fields would also serve the purpose of protecting the material structures of the propulsion unit from impact by pellet debris, although there will still be some damage from neutrons and x-ray and γ -ray radiation. The fields, of course, provide the required "cushioning" of the impulse. One drawback with the design is that the beam drivers, field-coil systems with shielding, and radiators add to the mass of ICF pulse propulsion systems.

3.2 ICF Pulse Propulsion

Figure 3.1 shows a schematic of an ICF pulse propulsion system. The thruster consists of current coils (single or multiple), usually assumed to be superconductors, which create a magnetic field of suitable geometry. Except for the trusses that hold the field coils together and onto the rest of the vehicle, the required shielding, and optics, the thruster region can be made mostly void of material, and this is also recommended to minimize the solid angle of material subjected to hazardous radiation (and possible plasma particle impingement) from the explosion. Some designs with low initial field-to-plasma energy ratios have been proposed with hemispherical metal walls surrounding the thruster region to aid in the flux-conservation process to halt and redirect the plasma, but such material walls are not a fundamental necessity to the concept.

An ICF pellet is injected into the thruster along the central axis by either mechanical or electromagnetic means. The pellet path can be monitored and corrected as it passes down the long corridor. When the pellet reaches a specified position within the thruster, it is irradiated by energetic driver beams such as laser beams, electron beams, or ion beams, and undergoes an ICF explosion. The geometry of the field coils and the relative positioning of the explosion point should be such as to maximize the propulsive efficiency of the device, an important design criterion along with the proper selection of parameter

ratios such as the initial field-to-plasma energy ratio. In fact, such issues may not just affect the propulsive efficiencies but also determine the feasibility of the thruster.

Now, as the initially spherical (assumed) plasma starts expanding radially outwards, it sweeps the vacuum magnetic fields aside because a hot plasma is a very good conductor. A perfect conductor does not allow magnetic fields to diffuse through them. For a fully ionized plasma, the electrical resistivity varies as the inverse $3/2$ power of the temperature, and is independent of the density (Spitzer resistivity). The magnetic field in the thruster (and that which cancels to zero inside the perfectly conducting plasma) is the superposition of fields due to all the current distributions present in the system; i.e., both the plasma surface currents and the currents through the coils. The plasma, by pushing against the magnetic pressure force at the interface, does work, and thus loses energy and is decelerated. It should be noted that the energy of the plasma will be mostly in the form of kinetic energy by the time the plasma has expanded a small distance on the scales of the thruster dimension. Because of the shaping of the field lines, the interface deforms anisotropically, and the plasma traveling forward in the general direction of the field coils is halted and redirected in such a way that the bulk flow will be out of the rear of the thruster in an axially collimated pulse jet. The initial isotropic expansion of the plasma is converted into a net rearward axially directed flow via interaction with the magnetic fields, and in this process, momentum is transferred to the vehicle. With no upstream endplate, a small jet may leak out the open end in the forward axial direction, contributing negatively to thrust. The individual pulses are isolated phenomena with the time interval between them being very much longer than the time required for the debris from a single pulse to exit the thruster region. The maximum pulse rate is limited by the rate at which pellets can be injected into the thruster.

Clearly, the concept relies on the property that good conductors tend to prevent external magnetic fields from diffusing into them. The thrust is due to the integrated pressure of the compressed magnetic fields acting on the surface of the field-coil structure.

Or equivalently, it can be considered as being due to an action at a distance between the current in the thruster's field coils and the induced diamagnetic surface currents on the plasma. Thrust calculations will be explained in more detail in Section 4.3 of Chapter 4.

The power for the beam drivers is obtained from the expanding plasma by placing induction pickup coils in the thruster region. The temporal variation of the flux through the pickup coils arising from the compression and deformation of the fields induces an emf in those coils. Typically, the power needed to be extracted for the drivers is not a significant fraction of the plasma energy, because of the high energy gain of fusion pellets and the reasonable efficiency of beam drivers. The energy for the pulsed drivers may be stored in capacitor banks which could possibly be coupled to rotating machinery providing extended period storage.³ Energy storage via vacuum transmission lines charged by a diode or in inductive cavities have also been proposed.⁴ Engines may be started up inductively.

It can be seen that the thruster is functioning as both a magnetic thrust chamber and a magnetic nozzle, although the device has often been referred to as either a nozzle or a thrust chamber. Also, the thruster is serving as the power source for the propulsion system. However, the magnetic thruster of an ICF pulse rocket is fundamentally different from devices with the same name in electromagnetic propulsion because a volumetric $\underline{J} \times \underline{B}$ body force is not being used to accelerate plasma permeated by fields as in the latter. Instead, the kinetic energy of the plasma is borne out by an isotropic fusion explosion taking place within the thruster and the magnetic fields external to the plasma anisotropically decelerates, deforms, and redirects the plasma to produce thrust by applying a pressure force at the interface (the idealized situation of perfect conductivity is being assumed here). Also, the means by which the interacting currents and fields are produced are different between the two classes of devices. But then, of course, it is also important to realize that the magnetic pressure force acting on the interface of a plasma in the thruster of an ICF pulse rocket comes from the $\underline{J} \times \underline{B}$ force integrated over a thin volume straddling the interface (where the surface current is flowing).

Thrust vectoring may be achieved by exploding the pellets slightly off the central, longitudinal symmetry axis of the thruster.³

Figure 3.2 is a very simplified energy diagram for the flow process in a magnetic thruster during one pulse (assuming that the plasma blob leaves the thruster entirely in one bounce), and will be explained below. This figure is entirely out of scale, and radiation and recombination phenomena have been neglected. Also left out is consideration of energy lost to neutral particles and power required to feed the beam drivers. The plasma will be assumed to be unmagnetized and perfectly conducting.

Energy from an ICF explosion appears in the form of energy in charged particles, neutrons, and photons. The photon energy appears as a result of the interaction among the particles in a fusion fireball, and it has been suggested that a large fraction of the energy from an ICF pellet may come out in the form of electromagnetic radiation.^{***} The interplay between the many particles in a debris also results in an energy partition among species different from those quoted for isolated fusion reactions.^{3,4} Radiation energy losses, primarily bremsstrahlung, will be significant in the very early stages of plasma expansion while the plasma is still very hot but will not be a dominant loss factor after that stage. However, as a plasma gets magnetized in a real system, cyclotron radiation could play a role.

The total plasma energy is the sum of the internal and kinetic energies. During the rapid expansion of the plasma, most of the internal energy will be very quickly converted into bulk kinetic energy of the flow, and the temperature of the plasma will drop sharply. The total energy of the plasma gradually decreases as the plasma continues to do work in

^{***} Various estimates have been made for the fraction of energy from an ICF explosion that will be channeled into charged products. The LLNL studies such as those for the VISTA vehicle (Reference 5) have assumed about 1/4 of the energy from a DT explosion to go into charged products, while the Penn State calculations for the ICAN vehicle (Reference 9) have suggested that most of the energy from a DT pellet may come out in the form of photons.

pushing against the external magnetic pressure acting on the interface, and this is met by a concomitant increase in the energy of the vacuum magnetic field. The regions influenced by the deceleration of the interface will experience a temperature rise, and especially, there will be heating due to shocks. With this thermalization of the plasma, some of the kinetic energy is converted into internal energy rather than goes to the fields (and the interface won't be as fully expanded as it would be if no such conversions were allowed). This thermalization process does constitute an important mechanism in the deceleration/redirection of the fluid. The interface of roughly the forward half of the blob eventually gets halted by the fields and experiences a reversal in its direction of motion and subsequently gets accelerated out of the rear of the thruster by the recoiling fields. Other regions of the flow will experience only a "glancing angle" deflection, and the resulting global bulk flow will be a highly axially collimated rearward flow. In this process, the energy stored in the vacuum magnetic field decreases again as the excess energy is returned to the plasma and we see a rise in the kinetic energy of the plasma as it is pushed out of the thruster region. By the time the entire plasma blob has left the region of influence of the thruster fields, the magnetic field energy returns to its initial value, and the plasma energy returns to its initial value less the energy extracted by the induction pickup coils and the cumulative energy lost to radiation and neutral particles not already accounted for at the beginning. The internal energy, by the end of the thrusting process, generally will not go down as low as the minimal value it reached during the expansion phase.

Although circumventing many of the limitations associated with material thrust units, the use of magnetic fields to deform and redirect a plasma introduces new problems that appear quite formidable. Under the hydrodynamic approximation, a perfectly conducting fluid can be expected to exclude entirely the external magnetic fields; the vacuum magnetic field comprising the thruster will appear as an impermeable elastic wall. However, a plasma does not behave that ideally, even if it can be assumed to possess zero

resistivity. The plasma in an ICF thruster will have finite resistivities, and the deviation from idealized behavior can be large. For example, external magnetic fields can diffuse into the plasma and plasma polarizations (which can occur especially when distorted interface profiles arise from the development of instabilities) can lead to cross field drifts. Experimental investigations of plasma-field interactions (for which there is a large amount of literature) have often indicated increased confinement of plasma and increased plasma velocity in the direction normal to the confinement direction when magnetic fields were applied, but faster than accountable expansions and high plasma loss rates along and across the fields were also observed. Furthermore, the very large expansion ratios experienced by the plasma will cause the bulk temperature to fall below the ionization temperature in the conceived-of designs, and not only will recombination reactions and resistivities involving neutrals enter and complicate the process, the proper functioning of the device as a thruster itself is somewhat questionable with the currently chosen parameters. There is also the important and related issue of separation of a magnetized plasma flow from magnetic field lines, a necessary phenomenon for the production of thrust. Here, low conductivities may actually serve a useful purpose. Various plasma instabilities could also prevent the device from working. For example, the system under concern is susceptible to the Rayleigh-Taylor instability. This instability can manifest itself in systems where a light fluid supports a heavy fluid, the light fluid being the magnetic fields and the heavy fluid being the plasma expanding against the fields in this case. The time scales of field diffusions, drifts, and instability developments relative to the characteristic time scale of plasma expansion in a thruster, which is typically on the order of several to several tens of microseconds, will thus be important parameters in determining the feasibility and performance of a magnetic thruster.

3.3 Proposed Designs and Earlier Works; High Field and Low Field Thrusters

Since its inception, several designs of ICF pulse propulsion systems have been proposed. It has been customary to categorize the thrusters either as high-field or low-field.

In the high-field type, the initial magnetic field stores several times the initial energy of the plasma. Thus, the fields in the thruster region have enough energy to deform and decelerate the plasma without undergoing significant compression, and the plasma deforms significantly and turns away a large distance from the field coils. The field strength on the surface of the plasma typically changes by less than an order of magnitude during the expansion, and the deceleration process is relatively spread over the dynamic expansion time scale. A multiple-coil design with the coils arranged to form a nozzle-like configuration is not necessary (and in fact, not preferable, especially when using superconducting coils, in light of the hazardous radiation and neutrons present in the environment), but the field energy will not necessarily be concentrated only in the region where the plasma expands. Single-coil designs and twin-coil designs with the forward coil producing a stronger field than the rear coil have been proposed. The field lines will be such that they constrict to one side of the explosion point, enabling preferential ejection of the plasma for thrust production. This magnetic mirror geometry is capable of effective plasma flow restriction in the forward direction despite the lack of an end wall at the upstream end of the thruster.

In the low-field type, the energy stored in the initial magnetic field is much less than the initial energy of the plasma. The lower field requirements allow savings in mass of the field-coil system and structure.³ However it will generally be more difficult to achieve efficiencies as high as those obtainable with the high-field type because of differences in the dynamics of the interface. As mentioned earlier, the entire forward region of the thruster may be surrounded by a flux-conserving metallic wall as in the Daedalus design (discussed

below) to aid in the trapping of the magnetic flux between it and the advancing plasma interface. The plasma expands almost spherically, little impeded by the magnetic fields until the interface approaches the flux-conserving walls very closely, by which time the fields are highly compressed. The magnetic field rises very sharply, and eventually the interface motion is reversed a finite distance away from the walls. The deceleration process is concentrated almost entirely around the time the interface presides very closely to the walls, which also implies a less "cushioned" momentum imparting process than the high-field type. By nature, this design requires a multiple-coil configuration to obtain decent thrust efficiencies because only those parts of the interface advancing along the line of sight of the field coils are significantly affected by the fields. The exact geometrical layout of the coils will be an important consideration for the low-field design when seeking high propulsive efficiencies. By closely spacing the field coils, the flux-conserving metallic thruster wall will not be necessary. Unless very weak currents are used, the plasma generally will not leak out between the adjacent coils.

The studies out of Lawrence Livermore National Laboratories have dealt with the high-field design. One of the designs that received rather detailed attention was a DD fueled single-coil thruster with the explosion site located at one side of the plane of the current coil.³ Although the parameters chosen for this vehicle have become obsolete, being too optimistic (the parameters for the VISTA vehicle concept show a more up-to-date view⁵), they still serve as an estimate for the kinds of numbers we will be looking for in an ICF pulse rocket. In this rocket, the laser driver was assumed to obtain 33 MJ of input energy from the induction coils and to output a 2 MJ laser beam that would ignite a DD pellet in which about 15 mg of deuterium would be burned, yielding an estimated 1300 MJ of energy in charged products. Pulse rates on the order of 100 Hz were envisioned. The field coil was to have a radius of 6.5 meters and to carry a current of 22 MA (MA-turns by default), allowing the initial magnetic field to store roughly five times the energy of the

plasma. The proposed interplanetary spacecraft had a vehicle mass of roughly 500 tons (not including payload or propellant mass) with slightly more than half of this mass being taken up by the laser driver system, including the laser and radiators. The thruster system mass was about one third of the vehicle mass. Typical propellant consumption for an interplanetary mission was quoted to be on the order of several hundred tons. Pellet masses ranging from about 420 grams to 0.48 gram were to be obtained by varying the amount of expellant added, and specific impulses ranging from 5.2×10^3 seconds to 1.5×10^5 seconds were to be achieved. The maximum thrust of the engine was about 3×10^6 Newtons. A vehicle power-to-mass ratio of 110 W/gram and a maximum acceleration close to 0.1g was considered possible. In light of the more recent VISTA vehicle study, these values are optimistic (by a factor of five or more), and consequently, somewhat overestimated performance capabilities such as 9-day low-payload flights to Mars or 1-year flights to Pluto carrying 1500 tons of payload were quoted for a single-stage vehicle. With the DT fueled VISTA vehicle having a total initial mass of 6000 tons, a 3-month 100 ton payload round-trip mission to Mars was considered more realistic.

Whereas most other studies simply estimated the design parameters for thrusters by equating the initial plasma energy to the product of the average, initial thruster field energy density and the characteristic thruster volume, the work in Reference 3 presented the results of a simulation showing the profiles of the interface of an expanding plasma at evenly spaced time steps. Although claimed to be the results of a hydrodynamic calculation employing classified software, more detailed and convincing information regarding this calculation could not be obtained. This plot matched well with that obtained by one of the author's simulations carried out under the thin-shell approximation employing a similar parameter ratio and geometry. Favorable thrust efficiencies were quoted in general agreement with those obtained by the author's thin-shell simulations, but the calculation of the energy efficiency as the square of the thrust efficiency was not appropriate (see Section 4.3 of Chapter 4).

Figure 3.3 shows an earlier conception of an interplanetary ICF pulse rocket vehicle with a high-field twin-coil thruster.

The British Interplanetary Society selected the low-field type thruster for the propulsion system of the Daedalus starship.⁴ The Daedalus Project was a study conducted to design a conceptual, unmanned interstellar vehicle to accomplish a flyby mission to Barnard's Star at a distance of about 6 light years in 40 to 50 years. A terminal coast velocity of very roughly $0.1c$ was planned for a two-stage starship with an initial mass of 5.4×10^4 tons and a total engine burn time of approximately 4 years. A D^3He pellet was selected because of the high percentage of fusion energy available in charged particles. The pellet, an optimistic design, consisting mostly of fuel material (this is not the standard type of pellet usually envisioned as feasible in ICF literature) and with a small DT core to act as a trigger, was to be ignited by a relativistic electron beam. The first-stage ignition system was to store 2.7×10^9 J of energy. Each pellet for the first stage had a mass of about 2.8 grams and a radius of about 2 cm with an energy release of 3.2×10^{11} J even for a rather low pellet burnup fraction of about 15 % or so. For an interstellar mission, it is important to maximize the specific impulse, and a specific impulse of about 10^6 seconds was estimated for the vehicle. The thrust rating for the first-stage engine was 7.5×10^6 N at a pulse rate of 250 Hz. The radius of the hemispherical chamber for the first-stage engine was to be 50 meters, with the flux-conserving metallic wall having a thickness of about 1 mm and reaching a peak temperature of about 1600 K. The ICF explosion site was located at the center of curvature of the hemispherical thruster walls. A four-coil system wrapping around the chamber wall creates a cusp field, and the electron beams were to be injected down the (near radial) cusp to the ICF explosion site. The first-stage field-coil system stores roughly 4×10^9 J of energy with a characteristic initial magnetic field strength of 0.0333 T (as quoted). The converging fields downstream of the explosion site have a negligible effect in impeding the plasma motion because of its low strength. Of the dry

stage mass for the first stage measuring about 1700 tons, the propulsion system weighted roughly 950 tons of which about 125 tons were in the mass of the coils. The propellant mass to be carried by the first stage was about 4.6×10^4 tons. The second stage was a downsized version of the first stage and a total payload mass of 450 tons was planned. The vehicle specific mass was quoted at 10^{-8} kg/W, and the maximum acceleration of the vehicle was estimated at about 1.1 m/sec^2

An expanding perfectly conducting fluid will always be turned away a finite distance from a flux-conserving wall surrounding the thruster region except at the cusp separatrices. However, with a plasma, even with other considerations set aside, the finite Larmor radii of particles can result in plasma contact with the walls.

Elliott and Terry modeled the expansion of a plasma in the first-stage thruster of the Daedalus vehicle, using the thin-shell approximation.⁶ In this work, the expansion of the forward half of the plasma was simulated until shortly before turnaround (after which a flux-conservation calculation was employed to take the plasma out to the turnaround point) assuming all of the plasma to be collected into a thin perfectly conducting shell right at the interface. The plasma shell was found to expand spherically with minimal deceleration until it was within a few tenths of a meter from the walls, i.e., very close to the walls on the characteristic dimensions of the thruster, from which point sudden deceleration took place. The closest distance of approach to the walls under this approximate calculation ranged from about 0.5 cm to about 22 cm, with shell elements in regions of lowest initial field approaching the walls closest before feeling the strongest deceleration of all elements and thus also experiencing the highest pressure upon velocity reversal as well as undergoing velocity reversal at the earliest time. These distances of closest approach are very much larger than the Larmor radius of the alpha particle, implying no plasma wall contact, at least under this approximate and simplified analysis. Elliott states, quoting a work by Devaney,⁷ that both the flute instability which arises from the deceleration force at the interface and that which results from the bad curvature of the interface will not be

important during the early stages of the expansion process because of stabilization through disparate ion and electron Larmor radii causing charge separations to occur out of phase with the particle drift separation (finite Larmor radius stabilization). Approximate analyses of this type should be interpreted with caution though. For later stages of the expansion, Elliott relied on a simple analysis suggested by Poukey⁸ (although based on the thin-shell model and not a hydrodynamic model) to reason that the only gravitationally-driven (i.e., driven by the deceleration of the interface) modes that grow fast (i.e., fast enough to be dangerous) are stabilized by the finite Larmor radius effect. The danger of the curvature-driven (this is in a sense also “gravitationally driven” because the centrifugal force driving the instability may be considered an effective gravitational force) instability was left an open question. Conclusions drawn here by some of these earlier works regarding the danger of instabilities may be over optimistic (see Part 2 of the appendix to this chapter).

Figure 3.4 shows a cutaway view of the Daedalus starship. The table in this figure presents the geometries of the field-coil setups for both the first- and second-stage thrusters. Axial station refers to the distance of a field coil from the forward-most coil, where the rearmost coil, coil 4, is centered at the center of curvature of the hemispherical chamber.

Apart from thrust chamber designs, several novel approaches of initiating ICF explosions for space propulsion applications have been proposed. For example, Winterberg proposed staged thermonuclear microexplosions offering the possibility of igniting difficult aneutronic fusion reactions by refocussing energy from one explosion to the next.⁹ A recently initiated research effort to ignite ICF pellets using antiproton induced microfission is receiving serious consideration for the ICAN (ion-compressed antimatter-catalyzed nuclear pulse propulsion) concept.¹⁰ The antiproton source and accelerator will be a lightweight and low-energy system ideal for the application in mind.

References for Chapter 3

1. See, for example, C.J. EVERETT and S.M. ULAM, "On a Method of Propulsion of Projectiles by means of External Nuclear Explosives," Los Alamos Scientific Laboratory Report LAMS-1955 (1955). (Declassified, 1976).
2. A good review article with an extensive literature listing on the subject matter, although not up to date, is A.R. MARTIN and A. BOND, *JBIS* **32** (1979), 283.
3. R.A. HYDE, "A Laser Fusion Rocket for Interplanetary Propulsion," Paper IAF-83-396, 34th Congress of the Int'l Astronautical Federation, Budapest, Oct. 1983.
4. See, for example, Project Daedalus Study Group, "Project Daedalus : The Final Report on the BIS Starship Study," ed. A.R. MARTIN, *JBIS* **31** Supplement (1978), and A. BOND and A.R. MARTIN, *JBIS* **39** (1986), 385.
5. C.D. ORTH et al., "The VISTA (Vehicle for Interplanetary Space Transport Applications) Spacecraft - Advantages of ICF for Interplanetary Fusion Propulsion Applications," Lawrence Livermore National Laboratory Report UCRL-96676 (1987).
6. J.O. ELLIOTT and W.K. TERRY, *JBIS* **38** (1985), 120.
7. J.J. DEVANEY, "First Wall Magnetic Protection in an Inertially Confined Thermonuclear Reactor," Los Alamos Scientific Laboratory Report LA-8286-MS (1980).
8. J.W. POUKEY, *Phys. Fluids* **12** (1969), 1452.
9. F. WINTERBERG, *JBIS* **30** (1977), 333.
10. See, for example, G.A. SMITH, et al., "An Antiproton Driver for ICF Propulsion," to be published in the Proceedings of the 3rd Annual NASA Workshop on Advanced Propulsion (1992), and R.A. LEWIS et al., *Fusion Technology* **20** (1991), 1046.

Appendix of Chapter 3

Part 1: Figures for Chapter 3

(Starts Page III-A2)

Part 2: Some Simple Estimates and Approximate Calculations

(Starts Page III-A6)

Part 3: References for the Appendix of Chapter 3

(Page III-A25)

Schematic of ICF Pulse Propulsion

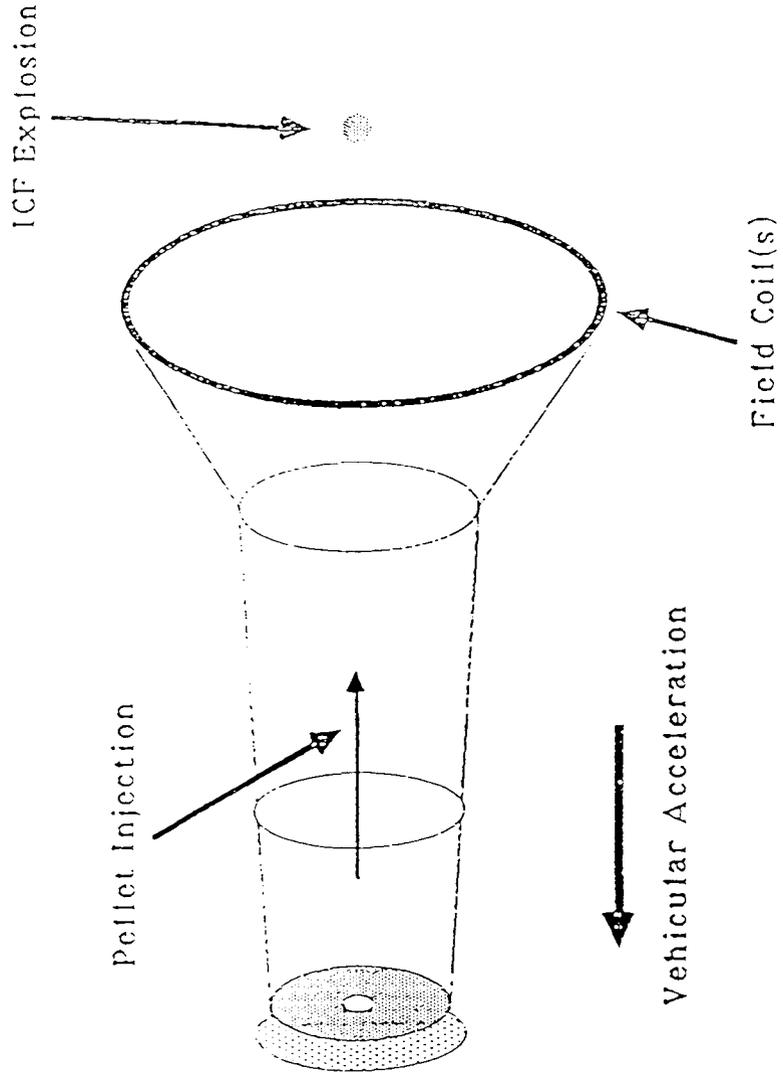


Figure 3.1

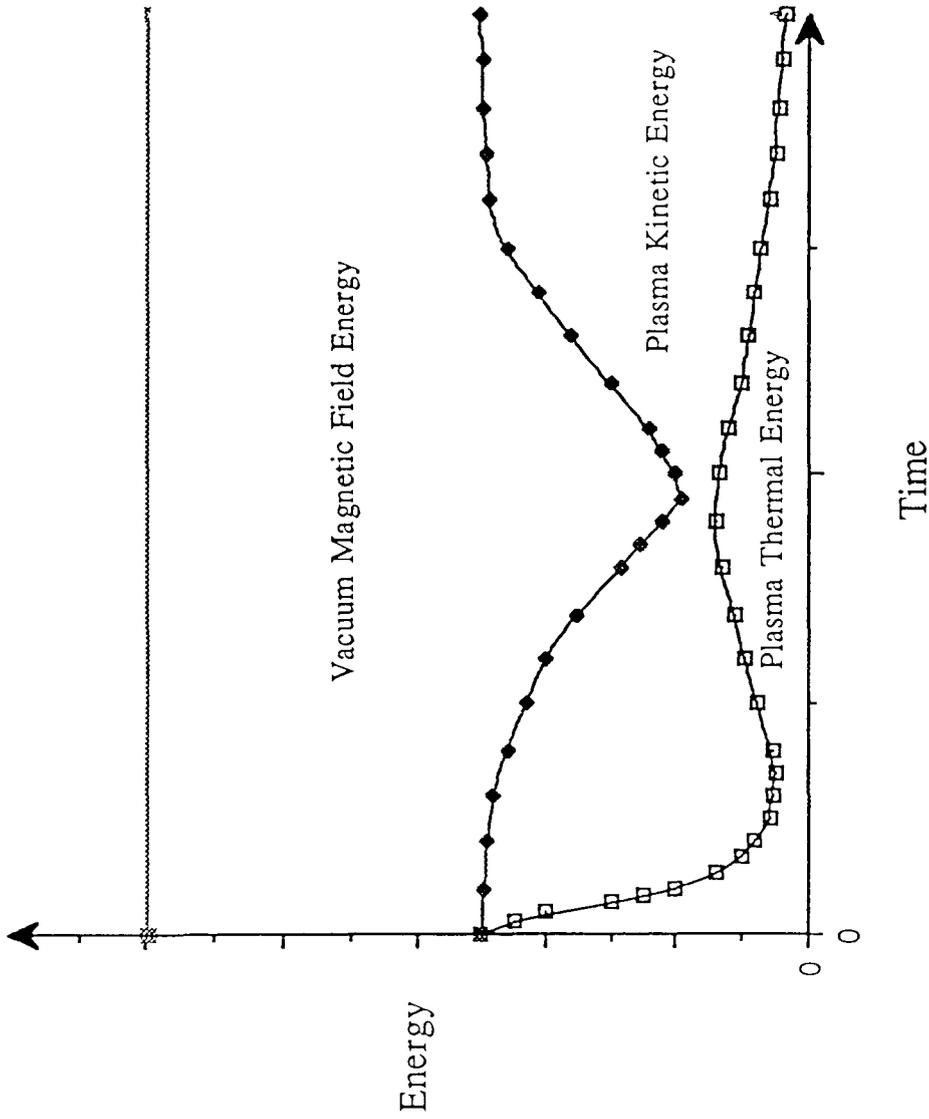


Figure 3.2

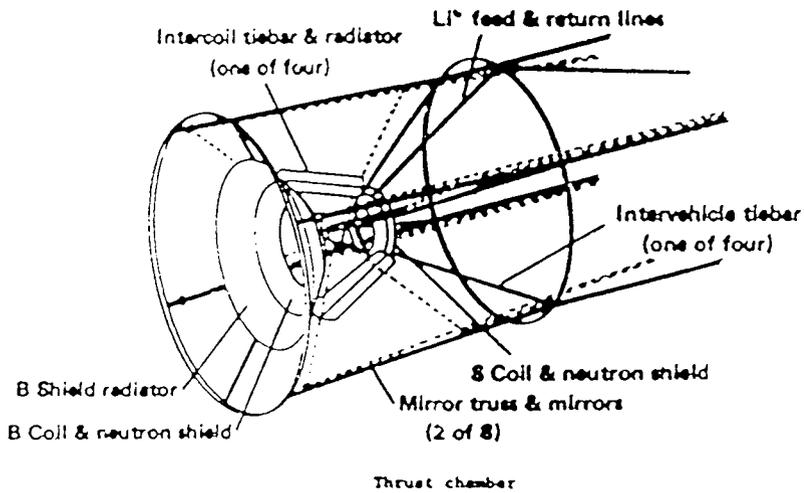
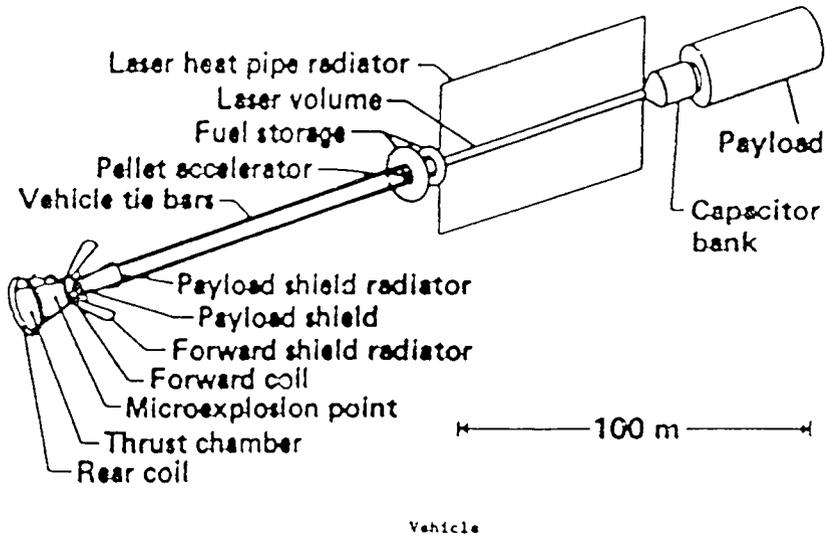
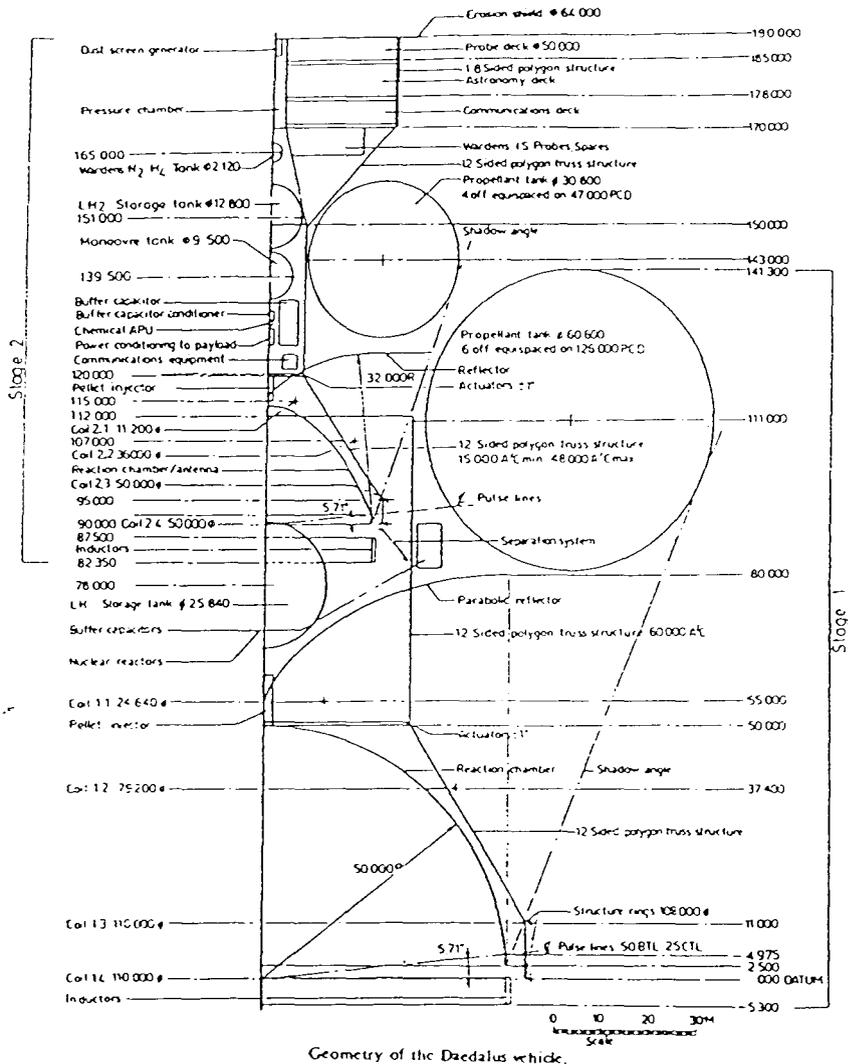


Figure 3.3 ¹

III-A5



Stage	Coil	Radius (m)	Axial Station (m)	Current (MA)
1	1	12.32	0.0	9.88
	2	39.60	17.6	1.05
	3	55.00	44.0	6.16
	4	55.00	55.0	-6.83
2	1	5.6	0.0	9.00
	2	18.0	8.0	0.95
	3	25.0	20.0	5.61
	4	25.0	25.0	-6.22

Thruster-Coil Parameters for the Daedalus Starship

Figure 3.4 2

Part 2: Some Simple Estimates and Approximate Calculations

Although not discussed in this thesis, a fair amount of work dealing with rough estimates of plasma parameters and approximate calculations of plasma motions and field variations in ICF pulse rocket thrusters and similar situations was attempted by the author. Because plasma behavior is very regime dependent, and because it is important to possess an intuition for the gross physical state and behavior of the flow, even when an accurate numerical simulation is carried out, approximate calculations not relying on heavy numerical analyses are important. In this appendix, a few of the estimates for numbers and approximate analyses will be introduced very briefly.

First, to open this section, some representative nuclear fusion reactions will be presented. The following three reactions are among the most commonly considered for controlled thermonuclear fusion :



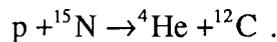
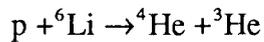
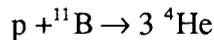
(these two DD reactions occur with about equal probability)



Appearing in parentheses are the energies deposited in the respective particles resulting from the reaction. In an actual fusion reaction, the interaction among the many particles present will result in a very nonnegligible fraction of the fusion energy being channeled into photons. In fact, copious production of x-rays and γ -rays may occur. Side reactions can also occur, and in general, the partition of energy among fusion products will differ from

the ratios quoted above for single reactions. This also means, for example, that some of the neutron energy may eventually be deposited into charged particles.

Of the reactions just quoted, the DT reaction is the easiest to ignite and has the highest burn rate at reasonable temperatures (e.g., 10 to 40 keV), but has the drawback of a high percentage of the fusion energy released appearing in the energy of neutrons. For propulsion applications involving the redirection of fusion debris by magnetic fields to obtain thrust, it is preferable to have most of the fusion energy emerge in the energy of charged particles which can be directed by the magnetic fields. Furthermore, as neutron radiation causes damage to the vehicle structure and payload, leading to heavy shielding requirements, aneutronic reactions producing no neutrons are attractive for use in space propulsion. This is especially so because of the extreme vulnerability of superconductors, which are the prime candidates for the current coils in the magnetic thrusters of fusion rockets, to such hazardous radiation. Shown below are a few examples of aneutronic fusion reactions, which unfortunately are not easy to ignite :



Note that the D ${}^3\text{He}$ reaction is not aneutronic because of DD self burn. Many other fusion reactions are possible, but not only are they generally difficult to ignite compared to reactions such as DT, not all are exothermic.

Regarding crude estimations of how plasma parameters vary with the expansion of the ICF debris, calculations were made under various assumptions on, for example, the initial values of the state variables and the forms of the flow profiles that develop. A simple example of such a calculation assumed that the plasma expands adiabatically, either as a uniform sphere or with all of the mass collected into a spherical shell of finite thickness at

the plasma-vacuum interface. In the case of an isentropic expansion of a uniform sphere, application of an elementary law of thermodynamics to an ideal gas (although we do not have an equilibrium process) quickly yields that the temperature of the gas, T , will vary with the radius of the gas sphere, R , as

$$T = T_0 \left(\frac{R_0}{R} \right)^2, \quad (\text{A3.1})$$

where the subscript zero denotes some reference state passed through by the system. If, on the other hand, we assume all of the gas to be collected into a uniform shell of thickness δ , where δ remains fixed in time or varies as some power of R , we obtain, for $\delta \ll R$,

$$T = T_0 \left(\frac{R_0^2 \delta_0}{R^2 \delta} \right)^{2/3}, \quad (\text{A3.2})$$

although this calculation contains some inaccuracies in the concept. In this work, a rather thick shell with δ on the order of several percent of R was often assumed.

The initial state of the plasma may be estimated from data on proposed ICF pellets, such as their mass, composition (fuel fraction, etc.), burnup fraction, radius, and state upon ignition. With such information, although not necessarily easy to obtain, one can estimate how plasma parameters such as densities (n_i , n_e ; where the subscript "i" denotes ions and the subscript "e" denotes electrons), pressures (p_i , p_e), temperatures (T_i , T_e), thermal velocities (v_{Ti} , v_{Te}), resistivity (η), plasma frequencies (ω_{pi} , ω_{pe}) and cyclotron frequencies (when the plasma is magnetized), collision frequencies (ν_{ii} , ν_{ei}), mean free paths ($\bar{\lambda}_{ei}$), Debye lengths (λ_{Di} , λ_{De}), Larmor radii (when the plasma is magnetized), and any characteristic ratios thereof vary with the expansion. For a fully ionized gas, the resistivity (Spitzer resistivity) may be expressed as

$$\eta \propto \frac{Ze^2 \ln \Lambda}{m_e} \left(\frac{m_e}{\kappa T} \right)^{3/2} \quad (\text{A3.3})$$

to within a numerical factor, where $\ln \Lambda$ is typically on the order of 10 to 20, and κ is the Boltzmann constant. The ionic charge is denoted by Z , and the charge and mass of an electron are denoted, respectively, by m_e and e . Furthermore, important parameters such as the time scales of growth of various instabilities and time scales and skin depths of field

diffusion into the plasma can be estimated. The importance of radiative losses may also be assessed.

Several estimates for these various parameters were made for typical conditions believed to exist in the magnetic thrusters proposed for ICF pulse rockets. The bulk temperature was found to drop drastically low, even well below the ionization temperature of the gas, as the plasma expanded out to the characteristic dimensions of the thruster. However, the flow generally (and at least in the important regions) tended to be comfortably in the quasineutral single-fluid MHD regime (low frequency and long wavelength: $u \ll c$, $\tau \ll L/c \gg 1/\omega_p$, $L \gg \lambda_D$, where τ , L , and u are, respectively, the characteristic time scale, length scale, and velocity present in the flow; additional conditions such as those involving the cyclotron frequency enter when the plasma becomes magnetized) for the most part during the important phases of the expansion flow. Excursions of electrons from ions was small. The collision frequency was high compared to the inverse of the characteristic expansion time of the bulk flow ($\tau \gg 1/v_{ii}$, $1/v_{ei}$), implying the applicability of the collisional fluid approximation. Ion and electron temperature equilibration was generally quick on the time scales of the expansion process. The neglect of physical viscosity also seemed quite acceptable during the expansion phase (generally acceptable under the presence of an isotropic distribution function; $\tau > \bar{\lambda}_{ei}/v_{Ti}$, $\bar{\lambda}_{ei}/v_{Te}$ for an unmagnetized plasma in which the condition on the Larmor radius is nonexistent), although possibly not to the extent of the situation in a magnetic confinement fusion reactor. Heat conduction effects during the thrust producing process did not appear significant. And, as is typical of a fusion type plasma, the average photon mean free path was very much longer than the characteristic dimensions of the ICF debris plasma.

Now, our system, in which a heavy fluid (in our case, the plasma) is supported against "gravity" by a light "fluid" (in our case, the magnetic field), is of the type susceptible to the Rayleigh-Taylor instability. In particular, we have the vulnerable

situation in which the plasma is supported by a magnetic field against inertial forces.³ Charge separations due to acceleration can induce the instability in a plasma. The typical curvature of our interface, which is convex outwards as one goes from the plasma side to the vacuum side (i.e., bad curvature) is such that perturbations will grow unstably. This may be understood by realizing that the centrifugal force the plasma flowing along the interface feels can serve as the driving force for the Rayleigh-Taylor (flute) instability. Similarly, the deceleration of the plasma by the vacuum magnetic field can also drive the Rayleigh-Taylor instability.

The growth rate of a flute instability may be given as⁴

$$\gamma = \sqrt{kg} , \quad (\text{A3.4})$$

where $k=2\pi/\lambda$ is the wave number of the flute mode, and g is the acceleration felt by the plasma that drives the instability. Taking the radius of the plasma blob to be R , and letting n be the mode of the instability, i.e., $n\lambda \approx 2\pi R$, the characteristic time scale of growth of a flute instability driven by bad curvature may be written approximately as

$$\tau_{\text{flute}} \approx \frac{2\pi R}{v_T \sqrt{n}} , \quad (\text{A3.5})$$

where v_T is the bulk velocity of the fluid along the interface, here represented by the thermal velocity (the Alfvén velocity is also appropriate for a magnetized plasma). If the characteristic time scale of growth of an instability is much longer than the characteristic time scale of the bulk expansion/redirection flow process, the instability will not be hazardous to the system. Now, simple estimates based on the approximate models outlined above revealed that the curvature-driven Rayleigh-Taylor instability could be hazardous to the successful operation of the magnetic thrusters concerned. The characteristic time scale of development of this flute instability during the early stages of expansion, at least under the assumptions made, tended to be much shorter than the characteristic expansion time scale of the plasma, even for the lowest $n=1$ mode. The flute instability could thus grow out of hand before the plasma expands significantly into the thruster. Not until after the plasma interface expanded to a few tens of centimeters did the instability growth rate start to

appear to become harmlessly small (provided the growth of the instability could be suppressed until then), in our case of a typical thruster. With the formation of a hot shell-like region at the interface, or if significant motions parallel to the interface arise as the interface deforms, matters worsen. Since it is known that the Rayleigh-Taylor instability is a potential hazard to ICF from during the formation phase of the ICF plasma, the curvature-driven instability requires major attention in a more accurate analysis.

Now, the (interfacial) plasma, by expanding against a vacuum magnetic field, generally experiences a deceleration (and redirection, depending upon the polar angle), and this causes the plasma to "see" an acceleration vector pointed from the vacuum side towards the plasma side. As alluded to above, this can also induce the Rayleigh-Taylor instability. The characteristic time scale of development of this gravitationally-driven instability, under the crude estimates carried out using Equation (A3.4), was shorter than the characteristic expansion time scale of the bulk flow during the starting phases of the expansion flow, even for the lowest $n=1$ mode. The growth rate tended to be quite lower than that for the curvature-driven instability at small radii (e.g., 1 cm), but became much higher than that for the curvature-driven one at larger radii (say, above several tens of centimeters in a typical thruster setup), although by that time, the instability may not be too dangerous if it could have been suppressed until then. The development time scale of the instability seemed to become comparable to the order of the characteristic expansion time scale of the bulk flow when the interface expanded out to a radius of a few tens of centimeters (again assuming the instability could be suppressed until then), and thereafter became longer than the latter, although not by orders of magnitude as in the curvature-driven case. The fact that the deceleration occurs somewhat suddenly in these thrusters around the time the interface experiences redirection, rather than evenly spread over the entire expansion process, may also possibly render this instability more dangerous. However, because the deceleration is very weak during the early phases of expansion when the interface radius is still small, this instability may also not be as hazardous as the crude estimates suggest. But in any case,

there is a nonnegligible possibility the instability driven by the interface deceleration, alone, could prevent the successful operation of a magnetic thruster, especially for high n modes, although apparently less dangerous than the curvature-driven mode.

In summary regarding Rayleigh-Taylor instabilities, both modes of the Rayleigh-Taylor instability appeared to be hazardous to the successful operation of magnetic thrusters, because their growth rates, according to crude estimates, could very well be high enough to allow perturbations to grow quickly out of hand, on a time scale shorter than that of the bulk expansion, while the plasma radius is still small. On the other hand, it also seemed possible that if these flute instabilities could be suppressed during the early phases of expansion, they, and especially the curvature-driven mode, may not seriously hinder proper expansion of plasma in the thruster. On the whole, it seemed more likely that the Rayleigh-Taylor instability will be one of the obstacles hindering satisfactory operation of the thruster. It should be stressed though, that these estimates, as well as many of the other estimates given in this appendix, have been based on very simple and crude models and rough guesses for parameters, and therefore in no way supply a definitive conclusion to the issue. The very motion of our interface, although typically exhibiting bad curvature, is very dynamic (with lots of free energy) and in a direction roughly normal to itself, and this may possibly change the associated stability picture as well, whether for better or for worse.

It is still possible that the situation regarding the development of flute instabilities may not be as discouraging as suggested above. For example, the finite Larmor radius stabilization, (especially) for large n values, may play a role in suppressing the instability. Shear in the magnetic field permeating a plasma can also stabilize the flute instability, so for current carrying plasmas, the instability may not be as dangerous as it first appears. However, because of the nature of our flow and the presence of currents, other instabilities inherent to plasmas (including fast electron flute modes, high frequency microinstabilities,

streaming instabilities, and universal instabilities) are also possible candidates for growth during the thrust producing process.

Another important quantity, when determining the feasibility and performance of a magnetic thruster is the time scale of diffusion of a magnetic field into a plasma. The entire magnetic thruster concept relies on the assumption that the characteristic time scale of field diffusion,

$$\tau_{\text{diff}} \approx \frac{4\pi L^2}{\eta c^2}, \quad (\text{A3.6})$$

(which may be derived from Ohm's Law, $\underline{E} + \frac{\underline{u} \times \underline{B}}{c} = \eta \underline{j}$, and $\nabla \times \underline{E} = -\frac{1}{c} \frac{\partial \underline{B}}{\partial t}$ and $\nabla \times \underline{B} = \frac{4\pi}{c} \underline{j}$; where B: magnetic field, E: electric field, j: current density, c: speed of

light) where here, L is the characteristic length of diffusion (on the order of the skin depth), is long compared to the characteristic expansion time scale of the plasma in the thruster region. As with other parameters, estimation of τ_{diff} for ICF debris plasma expanding in the thrusters of proposed ICF pulse rockets is difficult without accurate numerical results, and depending upon the assumptions used, such as on parameters, different conclusions can be drawn from the same type of analysis.

If conductivities are high, the field diffusion time scale will be longer than that of the expansion time scale of the plasma, and field diffusion into the plasma will not occur by much. Rough estimates based on the type of analysis outlined above indicated that it may be quite possible, even for the case of our plasma expanding in the magnetic thruster of an ICF pulse rocket, with the currently proposed pellet designs and thruster parameters, for no significant field diffusion into the plasma to occur during the important phases of the expansion process. This may be a little surprising in light of the large expansion ratio of the flow and the concomitant low temperatures, and thus rather low conductivities acquired by the fluid. However, it should be realized that the issue of field diffusion rests upon the relative time scale of field diffusion as compared to that of the bulk expansion, and the latter

is quite short in our case. But also, this simple analysis did not take into account the fact that our interface, instead of being primarily stationary, is moving rapidly against the vacuum fields, and the actual rate at which the plasma becomes permeated by the fields, with the conductivity not being too high, may be very much higher than was estimated here. The situation of the entire plasma blob becoming permeated by the "external" fields before the end of the thrusting process can not be totally ruled out. This is especially so, because even with the analysis that was carried out, which was very crude, the field diffusion time scale sometimes tended to linger around the characteristic expansion time scale of the bulk flow. Furthermore, accurate expressions for conductivities, appropriate for plasmas with the rather low temperatures expected of distended blobs in magnetic thrusters, were not employed. On the other hand, the formation of a narrow shell-like region near the plasma-vacuum interface (see Chapter 7, for instance) with a temperature very much higher than that of the bulk (a lot of the temperature rise is coming from shock heating, which is not incorporated into these simple analyses), hints at the possibility of the situation being more optimistic. As a temporary conclusion, it will be stated that the very crude estimates (which, for example, did not take into account the fact that there is a rapid movement of the interface in a direction roughly normal to itself, and towards the field region) carried out indicated that, despite the low temperatures achieved by the plasma, the characteristic time of magnetic field diffusion into the plasma may be long enough on the characteristic time scales of the bulk plasma expansion to be of harm; i.e., significant diffusion of the fields into the plasma may not occur during the important phases of the thrusting process.

But certainly, even if the fields do not diffuse significantly into the entire plasma blob during the primary portion of the thrusting process, some field diffusion into a fraction of at least the outer parts of the plasma blob will be inevitable. Even if we have a very highly conducting plasma with excellent coupling to the fields, all the way through the thrusting process, field diffusion into the interfacial regions will most probably not be

entirely negligible. This can be important because much of the plasma mass may very well be collected into the outer regions of the blob. Furthermore, the development of instabilities can drastically change the picture regarding field diffusion into a plasma (and also cross field drifts). Once field diffusion occurs into a plasma, the magnetic Reynolds number, which is a measure of field convection versus diffusion, will serve as a convenient quantity by which to evaluate the applicability of the $\eta \approx 0$ condition.

It will be remarked once again, that crude analyses for estimates of various plasma parameters during different stages of the expansion flow, such as those outlined above are not very accurate. They have been based on approximate models and rough guesses for parameters, including those for the initial state of the plasma. For example, rather unrealistic profiles as well as adiabaticity have been assumed to make the calculations tangible. Obviously, shocks, which as we shall see can play an important role in the actual flow process, cannot be handled by such analyses either. The difficulty of obtaining good data on ICF pellets and the state of a typical ICF plasma shortly after it is created, further lowered the reliability of the results obtained by the approximate calculations. These approximate calculations are intended only to serve the purpose of providing crude estimates, with the significance of relying only on simple analyses, and some of the temporary conclusions drawn from the results of these calculations may not be entirely correct.

To obtain a better handle on the physical situation involved, a variety of approximate calculations and estimates pertaining to flows and fields in, or of relevance to, ICF pulse rocket thrusters, other than those mentioned in the preceding paragraphs were also attempted. For example, the dynamics of the plasma interface motion and the temporal variations of density, pressure, temperature, and other plasma parameters of a blob

undergoing expansion were investigated using several different approaches. Estimates for the thickness of the shell-like region as a function of interface radius, albeit under an adiabatic assumption (which, as noted above, is not accurate because shocks are an integral part of, and play a major role in the formation of, the shell-like structures), were also made assuming crude models for the variation of applied interfacial magnetic pressure. One such result is quoted briefly in the beginning of Chapter 4. Such analyses can provide estimated expressions, in the form of a function of the interface radius, for the temporal variation of fluid quantities such as density, pressure, and temperature in the shell-like region (e.g., $\rho \propto R^{c1}$, $p \propto R^{c2}$, $T \propto R^{c3}$). With reasonable models, it was possible to obtain results in which the temperature of the "shell" rises with the expansion (recall the ideal gas equation and the fact that we have a distended expansion with pressure application at the interface). The estimates made here also supported the conclusions drawn above regarding the applicability of MHD, field diffusion time scales, etc.

Approximate expressions for realistic magnetic field variation (as well as for the initial field itself) in a magnetic thruster as a plasma expands were difficult to derive from theory, and models based on observations of numerical results often proved more useful. Shock tube problems and expansions in other geometries, such as in cylindrical geometries, with approximate prescriptions for interfacial magnetic pressure variation were also investigated analytically. Calculations employing approximate analytical techniques such as series expansions were attempted for such problems as two-dimensional expansion of a thin shell with magnetic pressure application at the "interface," and slab geometry hydrodynamical expansion of a plasma against a vacuum magnetic field, including the case of a resistive magnetized plasma. Approximate techniques to treat both the plasma and vacuum regions by one set of equations were also developed. The physics of field-plasma interactions relevant to ICF pulse propulsion were studied in addition to the crude attempts to investigate the interface motion, the hydrodynamics of the blob interior, and associated energy flows.

Unfortunately, when analyses that rather closely approximate the situation in an ICF pulse rocket thruster were attempted, calculations tended to become very difficult to pursue without resorting to highly numerical approaches, and in the general, the quality of the results that could be obtained for relevant problems were not high. It should be realized that we have a very time-dependent situation involving highly nonuniform fluid profiles and complicated geometries, which change significantly during the flow process. Problems readily treatable by analytical or "quasi-analytical" means were often unrealistic or did not have too much similarity with the situation in a magnetic thruster.

A large literature survey covering both theoretical and experimental works, was conducted in fields of fluid mechanics and plasma physics relevant to ICF pulse propulsion. A large amount of work exists on field-plasma interactions, including works treating cases of plasma expansion against a vacuum magnetic field (although these are not exactly in the context of ICF pulse rocket thrusters, and are usually experimental), and it is beyond the scope of this thesis to provide a list of such literature. Theoretical analyses of unsteady (pulsed) expansion of plasma against a vacuum magnetic field in situations where the primary direction of the flow is normal to the fields, apparently have not been carried out much, except for a few calculations which entirely neglect the details of the hydrodynamics of the fluid interior to the interface, and also employing only simplified field geometries and very crude prescriptions for pressure application at the interface.^{5,6} There have also been some works from several fields dealing with the classical hydrodynamics of an expansion of a fluid into a vacuum or a similar problem (although not the type of problem of a perfectly conducting fluid expanding against a vacuum magnetic pressure), but here, it will be considered sufficient to introduce a few books which may help offer insight into the subject matter. In this list that follows, it should be noted in connection with the first reference, that self similar analyses are generally not really

applicable to our problem in which characteristic scales exist, but the book is nevertheless of more than perfunctory value here.

- L.I. SEDOV, *Similarity and Dimensional Methods in Mechanics* (Academic Press, New York, 1959).
- K.P. STANIUKOVICH, *Unsteady Motion of Continuous Media* (Pergamon Press, New York, 1960).
- Ya.B. ZEL'DOVICH and Yu.P. RAIZER, *Physics of Shock Waves and High-Temperature Hydrodynamic Phenomena, vols. 1 & 2* (Academic Press, New York, 1966).

Finally, to close this discussion in the appendix, two simple examples of approximate calculations that do not rely on numerical simulations to model the motion of the interface will be presented. Here, the interface is of course the interface of a perfectly conducting plasma expanding against a vacuum magnetic field. In the two problems introduced here, all of the plasma is assumed to be collected into a thin shell at the interface, and the analyses do not concern themselves with the internal properties or dynamics of the fluid.

In the first analysis, the motion of an expanding plasma interface in a thruster that crudely resembles the Daedalus low-field thruster was calculated assuming a very simple flux conservation scheme. The plasma was assumed to expand spherically, centered about its explosion point, which is also the center of curvature of a flux conserving hemispherical chamber wall surrounding the region. The problem was, in fact, assumed to possess spherical symmetry. The effect of the magnetic pressure on the shell motion was assumed to be negligible until the interface approached the chamber walls closely. This is because, in a low-field thruster, the fields must be compressed highly before its strength rises high

enough to significantly affect the interface motion. Numerical simulations employing the thin-shell approximation revealed that the simplifications just listed are acceptable for the low-field design.⁷

The magnetic field was prescribed to vary as

$$B = \frac{K}{R^2 - r^2}, \quad (\text{A3.7})$$

based on a crude flux conservation model, and K is a proportionality constant (see below for notation). The momentum and energy relations for the shell (which is assumed to be comprised of elements independent of each other) may be written as

$$M \frac{dv}{dt} = -4\pi r^2 p_B, \quad (\text{A3.8})$$

where $p_B = B^2/8\pi$ is the magnetic pressure, and

$$E_1 = \frac{1}{2} M v^2 + \frac{4\pi K^2}{8\pi} \int_{r_1}^r \frac{r^2}{(R^2 - r^2)^2} dr, \quad (\text{A3.9})$$

by recalling that the energy loss of a shell (a shell can only have kinetic energy) arises from its doing work against the fields. Here, the variables were defined as follows;

R : radius of the hemispherical flux conserving chamber wall

r_1 : expansion radius at which the magnetic field starts exerting a nonnegligible influence on the plasma motion

r : radius of the spherical plasma shell

M : mass of the plasma shell

v : velocity of the plasma shell (radial only)

E_1 : energy of the plasma shell upon passing the radius r_1 (letting the initial energy of the plasma be E_0 , we have $E_1 \approx E_0$).

Now, assuming that the influence of the fields on the plasma motion is only notable when the interface is very close to the flux conserving chamber walls, one obtains, by combining the above equations,

$$R - r = \frac{1}{2Q} \left[1 + \sqrt{1 + \left\{ -2KtQ\sqrt{\frac{Q}{4M}} + \sqrt{\left\{ 2Q(R - r_1) - 1 \right\}^2 - 1} \right\}^2} \right], \quad (\text{A3.10})$$

where $Q \equiv \frac{8E_1}{K^2} + \frac{1}{R - r_1}$. The radius r_s , at which the shell reverses direction, will satisfy

$$E_1 + \frac{K^2}{8} \left(\frac{1}{R - r_1} + \frac{1}{4R} \ln|R - r_1| \right) = \frac{K^2}{8} \left(\frac{1}{R - r_s} + \frac{1}{4R} \ln|R - r_s| \right). \quad (\text{A3.11})$$

A result of this calculation is shown in Figure A3.1. In particular, the parameters were chosen, as listed below, to resemble the situation in a Daedalus first stage thruster :

$$R = 50 \text{ meters}$$

$$r_1 = 47.5 \text{ meters}$$

$$M = 2.8 \text{ grams}$$

$$E_0 = 3.2 \times 10^{18} \text{ ergs}.$$

The magnetic field strength was also chosen to be similar to that of the Daedalus engine. Both Figures A3.1a and b show the interface coordinate, as measured from the explosion center, as a function of time. The time was measured from the instant the interface passes through the radius r_1 . Figure A3.1b is just an enlargement of the curve in Figure A3.1a around the time the interface reverses velocity. With the choice of numbers quoted above, the shell reversed itself about 10 cm from the walls, and the maximum magnetic pressure rose to the order of 10^8 MPa. The results of this calculation were in decent agreement with the more sophisticated numerical simulations employing the thin shell model, and this also implies that the simplifying assumptions that went into this analysis were legitimate. For example, the deceleration of the shell occurred only when the shell approached the flux conserving walls very closely (but then very strongly). In general, such preliminary information about the gross expansion dynamics can be obtained to some extent from these simple estimates.

In the second example, the expansion of an initially spherical perfectly conducting shell into a uniform magnetic field was modeled, starting from when the shell has zero radius. Spherical expansion was not assumed (nor do we have such a situation), and the deformation of the interface was investigated (see also Reference 5). In the particular

calculation discussed here, the magnetic pressure applying at the interface, and decelerating the shell, was assumed to be that which would apply on the surface of a (perfectly conducting) sphere that expands at constant velocity, retaining its spherical shape, into a uniform magnetic field. An expression for such a magnetic pressure was given by Venezian⁸ as

$$p_B = \frac{9B_0^2}{32\pi} \frac{(1+\beta)^3}{(1-\beta)(1+2\beta)^2} \sin^2 \theta, \quad (\text{A3.12})$$

where $\beta \equiv \frac{1}{c} \frac{dr}{dt}$. Here, the radius to a shell element is being denoted by “r.” The magnetic pressure force at the interface was prescribed to apply only in the purely radial direction (in a polar coordinate system centered at the “explosion point”), as in the case of a sphere, regardless of the actual orientation of the interface. This of course is not realistic, but such approximations certainly make calculations more tractable. Again, all shell elements were assumed to be independent of each other. By applying either Newton’s Second Law ($F=ma$) or an energy equation ($p_B \Delta(\text{vol swept}) = -\Delta(\text{kinetic en})$) on the shell elements, and employing expression (A3.12) for the magnetic pressure, one obtains

$$\frac{dr}{dt} = \pm \sqrt{\frac{2K}{3} \left(\frac{3E_0}{MK} - r^3 \right)}, \quad (\text{A3.13})$$

to lowest order in β , where $K \equiv \frac{9B_0^2}{8M} \sin^2 \theta$. The total mass of the shell is being denoted

by M , the strength of the uniform magnetic field, by B_0 , and the initial (kinetic) energy of the shell, by E_0 ($E_0 = \frac{Mu_0^2}{2}$, where u_0 is the initial purely radial expansion velocity of the

shell). Figure A3.2 shows the result of solving the equation above for a particular choice of parameters somewhat resembling that for a high-field magnetic thruster. The location of the “interface” is plotted at successive instants of time. The “explosion point” is located at $x=y=0$, but the horizontal and vertical axes are not on the same scale, and the reader is reminded that the shell, during the early stages of expansion, is actually spherical in shape. The uniform field is being applied in the x direction here. For discussing the plot, the

parameters used in this calculation may be taken to be representing a case with $B_0 = 5 \times 10^4$ G, $M=1$ gram, and $u_0=10^9$ cm/sec.

As the expansion continues, the analysis showed correctly that the shape of the interface will deform drastically, in a manner similar to the generic way by which it really does so, as more accurate and realistic numerical simulations will also show. "Reflection" of the shell is clearly seen in the plot, especially for shell elements initially traveling in a direction normal (i.e., against) to the fields. In the near axial directions, where the flow is close to parallel to the natural direction of the fields, we see the interface continuing to flow outwards without receiving significant impedance from the fields. However, the geometry of the problem is such that we do not have a magnetic field that is asymmetrically shaped at the site of the plasma, and obviously there is no preferential ejection of the plasma to one side to produce thrust.

Overall, despite the simplifying assumptions used, these analyses do provide insight into the expansion flow process, especially regarding interface dynamics. However, when approximate methods were attempted on problems more closely approximating the real situation in a magnetic thruster, the difficulty of the calculation increased rapidly.

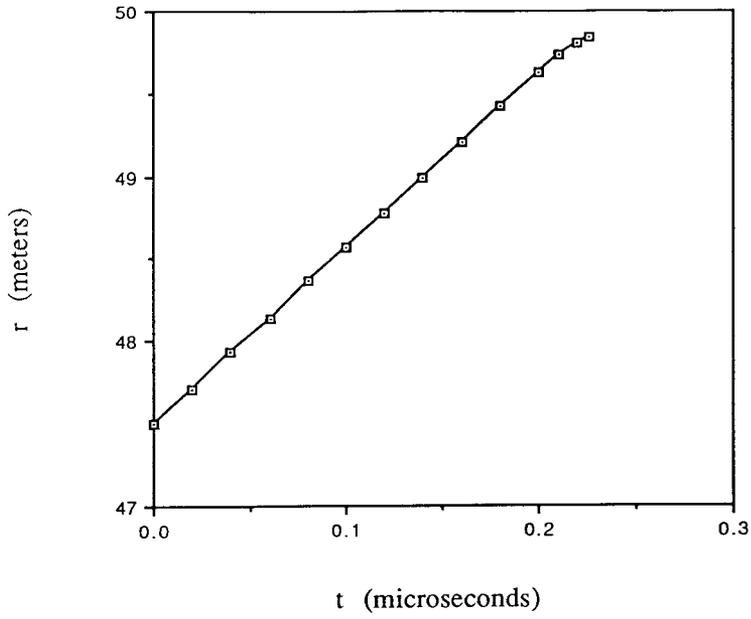


Figure A3.1a

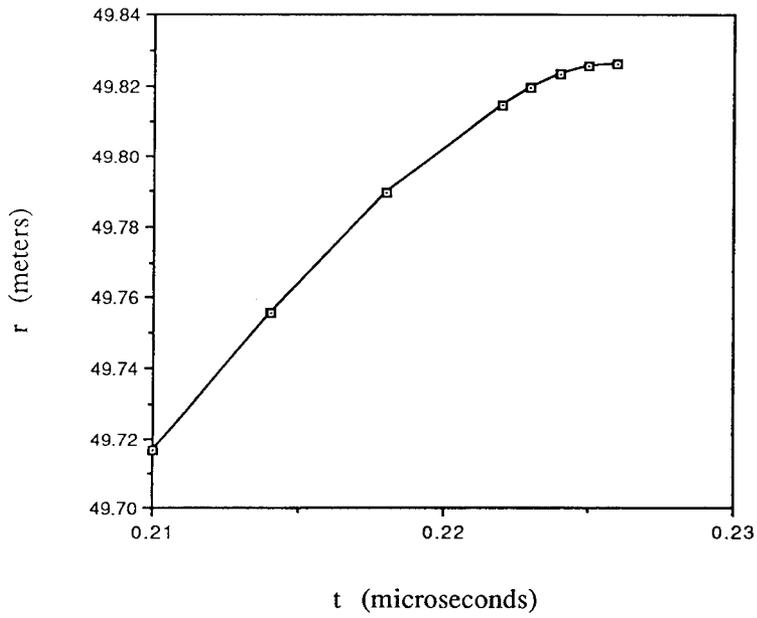


Figure A3.1b

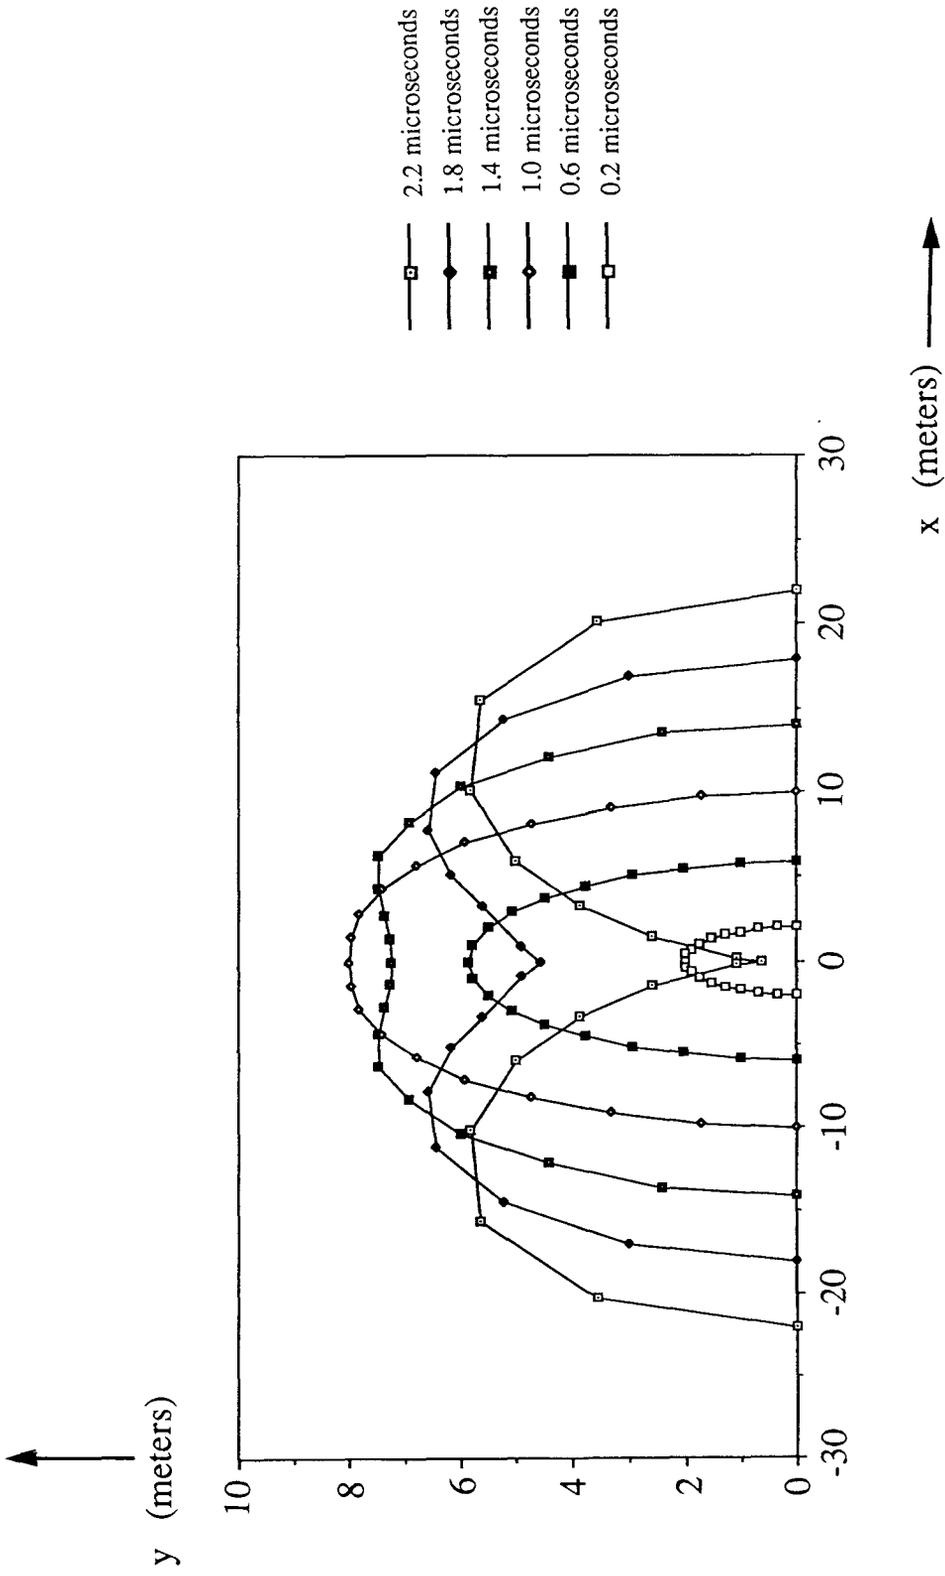


Figure A3.2

Part 3: References for the Appendix of Chapter 3

1. R. HYDE, L. WOOD, and J. NUCKOLLS, "Propulsion Applications of Laser-Induced Fusion Microexplosions," p.159ff in Proc. 1st Topical Meeting on the Technology of Controlled Nuclear Fusion, vol. 1, ed. G.R. HOPKINS, USAEC, CONF-740402-P1 (1974).
2. Project Daedalus Study Group, "Project Daedalus : The Final Report on the BIS Starship Study," ed. A.R. MARTIN, *JBIS* **31** Supplement (1978).
3. See, for example, G. SCHMIDT, *Physics of High Temperature Plasmas*, 2nd ed. (Academic Press, New York, 1979).
4. M. KRUSKAL and M. SCHWARZSCHILD, *Proc. Roy. Soc.* **A223** (1954), 348.
5. V.A. PILIPENKO, *Journal of Applied Mechanics and Technical Physics* **8** (1967), #2, 78.
6. J.W. POUKEY, *Phys. Fluids* **12** (1969), 1452.
7. See, for example, J.O. ELLIOTT and W.K. TERRY, *JBIS* **38** (1985), 120, or Chapter 4 of the present work.
8. G. VENEZIAN, Doctoral Thesis, California Institute of Technology (1965).

Chapter 4

Two-Dimensional Simulations Employing the Thin-Shell Approximation

The simplest way to model the expansion flow of a highly conducting plasma against a vacuum magnetic field in the thruster of an ICF pulse rocket is to assume all of the plasma to be collected into a thin perfectly conducting shell right at the fluid-vacuum interface from an early stage of the expansion. The thin shell is taken to be an infinitesimally thin entity with no internal material properties and endowed only with mass and velocity. The only force acting on it will be the magnetic pressure that is due to the vacuum magnetic fields.

The two-dimensional perfectly conducting thin-shell model is expected to model approximately the bulk expansion of a perfectly conducting fluid against a vacuum field because the finite deceleration of the interface and the enormous expansion ratio experienced by the plasma are expected to cause a shell-like structure, where much of the fluid mass will be collected, to form near the interface. This should be so (see Chapter 7) even for weak decelerations, and thus the thin-shell approximation should be reasonably valid from rather early stages of the expansion. The deceleration of the interface is caused by the finite magnetic pressure acting there. Actually, in a real plasma, the current induced in the volume of the interfacial region interacts with the diffused magnetic field through a $\underline{j} \times \underline{B}$ body force, and this serves as the force the field applies on the plasma. If, however, the plasma is approximated as an unmagnetized perfectly conducting fluid, then the induced current reduces strictly to a surface current, and the decelerating force becomes simply a magnetic pressure force applying externally and normally on the plasma interface. This can be seen by the usual method of shrinking the width of a “pill box” used in the

electromagnetic analysis to straddle the plasma-vacuum interface, to zero. The magnetic field that appears in the interfacial force's expression is of course the superposed field that is due to both the field coil's current and the currents on the plasma.

As the bulk fluid pressure decreases very rapidly with the expansion, the requirement for a pressure match across the interface between that of the fluid and the magnetic field causes a pressure rise to appear in the interfacial region. The mass accumulation behind the interface that arises from the deceleration of the interface (on a time scale much shorter than the time required for information of the interface deceleration to spread throughout the fluid blob) leads to a density rise there, and this density rise, together with a temperature rise, allows the required pressure rise to be met.

Very crudely, the thickness of a shell at the interface in an approximately spherical expansion can be estimated to be on the order of $\Delta R(t) = p_0^{1/\gamma} R_0^3 / (3 p_{\text{intf}}^{1/\gamma}(t) R^2(t))$, where $p_{\text{intf}}(t)$ and $R(t)$ are, respectively, the externally applied pressure at the interface and the characteristic radius of the interface at time t , and the subscript 0 indicates the state of the initial uniform spherical blob (see Chapter 7). This expression does not take into account any nonisentropic temperature rise such as those that are due to shocks, although shock formation is strongly associated with the development of a shell-like structure, as will be seen in Chapter 7. Substitution of rough estimates for numbers show that the shell-like structures will be very thin compared to the characteristic size of the blob. Spreading of the width of the shell generally will not be of much concern for the cases considered.

The thin-shell approximation allows simulations of plasma flow with realistic parameters and in realistic thruster geometries to be made approximately without requiring the extravagant computing resources typical of hydrodynamic codes, and should provide one with a rough feel for how the interface motion will proceed in a real situation. The investigation of a variety of thruster geometries and parameter ratios and their effects on propulsive efficiencies will be made possible with rather limited computational resources.

However, many limitations are brought about by representing the plasma by thin interfacial mass elements which have no provisions for taking into account the internal fluid properties of the expanding blob. The unphysical nature of the thin shells allows such unrealistic phenomena as shell elements falling onto the axis to occur (although by the time this happens, most of the thrusting process is over, at least for the first bounce). With no thermal energy included, the thin-shell-approximation model obviously cannot handle conversion between thermal and kinetic energies. Furthermore, in the real case, it will not be true that all mass will be concentrated in a shell at the interface, nor will the shell-like structure necessarily stay at the interface (see Chapter 7), and interior fluid dynamics not handled by the model can play an important role in determining the motion of the interface.

This approach of treating a plasma expanding against a magnetic field as a thin shell of mass elements has been taken by several authors as a simple and affordable way of treating the problem. For example, Poukey took this approach to solve approximately the problem of an initially spherical perfectly conducting plasma expanding into a uniform vacuum magnetic field.¹ Elliott employed the approach to simulate the expansion of a plasma in the Daedalus starship's low-field thruster.² His work is the only other thin-shell treatment of flows in the thrusters of ICF pulse rockets that is known to the author. However, Elliott treated only the upstream half of the plasma (i.e., that half initially expanding in the forward direction of the vehicle), and then only until the time that half of the interface reached its closest distance of approach to the thruster walls. One of the boundary conditions used by Elliott for solving the vacuum magnetic fields was not correct, and an error was detected in the version of his code published in his Master's thesis (an inconsistent definition of v_θ), but these problems are not expected to alter the results of his particular simulations very much. The use of Maxwell's Equations to compute numerically the interfacial magnetic pressure had to be stopped before the interface was turned around by the compressed fields because of computational difficulties. For

subsequent times, an approximate flux-conservation scheme making use of the trapping of magnetic flux between the conducting thruster walls and the plasma interface was employed to take the interface out until it was halted by the fields. This difficulty of calculating the field when it gets highly compressed within a narrow region was also experienced in the present work (not a problem in the high-field designs, because there, the interface does not approach the coils closely), and successful simulations of the flow past the time the interface elements turn around proved very painstaking for the low-field type thruster in which the interface comes very close to the thruster walls. In the case of the present work (using different softwares from Elliott for the calculation of the magnetic fields), the difficulty mainly resulted from the gradient of the magnetic stream function becoming too steep for a reliable calculation of the magnetic field on a coarse fixed grid.

In this work the thin-shell approximation has been reapplied to calculations of plasma expansion in the realistic geometries of thrusters proposed for ICF pulse rockets. The entire plasma sphere (not just the forward half) is treated until long after "reflection" of the shell elements by the recoiling magnetic fields. Although most of the impulse is transferred to the vehicle by the time the interface elements that will ever be turned back are turned back, it will still be important to carry out the analysis until the plasma is well out in the exhaust stream region, especially in a more realistic plasma simulation, because of issues such as field diffusions and detachment. Both the high-field and low-field thruster types were investigated. As found by Elliott, the shell in a low-field device expands with very little deceleration and with an almost perfectly spherical shape until it is almost at its maximum expansion location, very close to the field coils, upon which the interface elements are suddenly decelerated and reversed in velocity by the magnetic pressure which, by then, has acquired an extremely high value because of field compression. Compared to this, the deceleration and significant deformation of the interface take place more steadily throughout the expansion process in a high-field thruster, and turnaround of the interface occurs a large distance away from the field coils.

4.1 Problem Formulation

4.1.1 Thin-Shell Dynamics

The general formulation of the two-dimensional thin-shell-approximation scheme is based on the guidelines set forth by Poukey and Elliott.

The process is assumed to be azimuthally symmetric about the longitudinal, central thrust axis of thruster, and the thin plasma shell, taken to be perfectly spherical at the initial instant of the simulation, is broken up into azimuthal rings lying in planes perpendicular to the thruster's central axis. The Lagrangian mass elements are assumed to be independent entities with no internal forces, although in some cases, the details of the algorithm were such that neighboring elements could effectively affect each other's motion to a certain extent (see Figure 4.15, for example). A spherical coordinate system with the coordinate center located at the ICF explosion site is used. This is particularly convenient because the plasma starts expanding radially as a spherical shell. The polar angle θ will be measured away from the positive central axis of the thruster, facing forward. The initial θ coordinate of a ring element will be denoted by the variable θ_0 , and will be used to label the individual ring elements. Figure 4.1 shows these definitions and the coordinate system.

In this two-dimensional Lagrangian analysis, the independent variables are θ_0 and t , and the coordinates r and θ of the shell elements at a time t are expressed as functions of θ_0 and t , as

$$r = r(\theta_0, t) \quad (4.1)$$

$$\theta = \theta(\theta_0, t) , \quad (4.2)$$

with $\theta_0 = \theta(\theta_0, 0)$. The Lagrangian variables r and θ associated with the mass elements are not to be confused with the R and Θ variables used in the vacuum magnetic field

calculations on a fixed grid. The radial velocity and the angular velocity in the θ direction for the shell elements are given by

$$u = \frac{\partial r}{\partial t} \quad (4.3)$$

and

$$\omega = \frac{\partial \theta}{\partial t} . \quad (4.4)$$

The mass of an individual ring element, which subtends a polar angle of $\delta \theta_0$ at the initial instant, is given by

$$\delta M = \frac{M}{2} \sin \theta_0 \delta \theta_0 , \quad (4.5)$$

where M is the total mass of the plasma.

The only force acting on these Lagrangian mass elements is the magnetic pressure acting normally to the interface (magnetic pressures act normally to the field lines, and the field lines are tangential to the interface), and the motion of a shell element is governed by Newton's Second Law, which can be written as

$$\frac{\partial u}{\partial t} = r\omega^2 - \frac{4\pi p_B r^2}{M} \frac{\sin \theta}{\sin \theta_0} \frac{\partial \theta}{\partial \theta_0} \quad (4.6)$$

and

$$\frac{\partial \omega}{\partial t} = -\frac{2u\omega}{r} + \frac{4\pi p_B}{M} \frac{\sin \theta}{\sin \theta_0} \frac{\partial r}{\partial \theta_0} , \quad (4.7)$$

where p_B is the magnetic pressure acting on the outer surface of the shell element. For a successful simulation, the magnetic pressure values calculated as applying to the shell elements should be smoothed, and this was done by averaging over several (usually 5 when the total number of ring elements was a little over 90) neighboring elements. Nondimensionalization of these equations using

$$r' = r/l_*, t' = t/(l_*/v_*), u' = u/v_*, r'\omega' = r\omega/v_*, B' = B/B_*$$

(the primed quantities are the nondimensionalized variables) leads to

$$\frac{\partial u'}{\partial t'} = r'\omega'^2 - 4\pi K p_B' r'^2 \frac{\sin \theta}{\sin \theta_0} \frac{\partial \theta}{\partial \theta_0} \quad (4.8)$$

and

$$\frac{\partial \omega'}{\partial t'} = -\frac{2u'\omega'}{r'} + 4\pi K p_B \frac{\sin \theta}{\sin \theta_0} \frac{\partial r'}{\partial \theta_0}, \quad (4.9)$$

where

$$K = \frac{B_*^2 l_*^3}{M v_*^2}, \quad (4.10)$$

with l_* , v_* , and B_* being arbitrarily chosen constants, preferably characteristic quantities of significance to the process, and we have defined

$$u' = \frac{\partial r'}{\partial t'} \quad (4.11)$$

$$\omega' = \frac{\partial \theta}{\partial t'} \quad (4.12)$$

$$p_B' = \frac{|B'|^2}{8\pi}. \quad (4.13)$$

In the works discussed in this chapter, l_* was generally chosen to be on the order of the characteristic thruster dimension, v_* was chosen to be the initial expansion speed of the plasma, and B_* was defined as

$$B_* = \frac{I\pi}{c l_*}, \quad (4.14)$$

where I is the current through the coil.

Equations (4.8) through (4.10) show that one of the two parameters that characterize a thin-shell expansion process is the initial ratio of the magnetic field energy to the plasma energy (in the thin-shell model, the plasma energy is simply the kinetic energy of the shells). This ratio will be denoted by ER in subsequent discussions. The other parameter can be termed the geometry of the thruster, and in a thruster with only a single current coil, this "geometry" can be characterized by the ratio of the radius of the current coil to the distance between the center of the current coil and the explosion site. This ratio will be defined as the aspect ratio of a single field-coil thruster and will be denoted by AR .

In calculating the initial magnetic field energy stored in a thruster, an approximate method will suffice. Since the magnetic field energy that is due to a system of rigid currents is given in integral form as

$$W_B = \frac{1}{2} \int LI dI , \quad (4.15)$$

where L is the inductance of the system, the energy contained in the vacuum magnetic field of a thruster can be found by estimating the inductance of the thruster's field-coil setup. For a long solenoid in which the magnetic field energy is concentrated within the device where the field lines are running parallel, Equation (4.15) becomes

$$W_B = \frac{1}{2} LI^2 \quad \text{with} \quad L = \frac{4\pi N^2 A}{c^2 \ell} , \quad (4.16)$$

where N , A , and ℓ are, respectively, the number of turns, cross-sectional area, and length of the solenoid. Use of Equation (4.16) to calculate the energy contained in the fields due to, say, a single current loop is not entirely appropriate because of the difference in the geometry of the fields between a long solenoid and a single current loop. Nevertheless, the equation can still be used as an approximation with appropriate choices for the variables. In this work, the magnetic field energy contained in a single current-coil thruster was estimated by using Equation (4.16) with

$$\begin{aligned} A &= \pi R_c^2 \\ \ell &= R_c \\ N &= 1 \end{aligned} ,$$

where R_c is a characteristic radius of the thruster chosen to be the coil radius, so that

$$W_B = \frac{2\pi^2}{c^2} R_c I^2 . \quad (4.17)$$

Note that without any plasma present, the characteristic field strength varies as the current. This convention is, of course, not definitive, and depending upon the geometry of the thruster, is not necessarily always the best. It was found that the initial field energy stored in the thruster calculated as $\frac{B_*^2}{8\pi} \frac{4\pi}{3} R_c^3$ or $\frac{(2B_*)^2}{8\pi} \frac{4\pi}{3} R_c^3$ (some characteristic thruster dimension can replace R_c) substantially underestimates the real field energy. Note that B_* is equal to half the field value at the loop center and is roughly equal to the field value at the explosion point in a thruster with a value of AR equaling one.

Now, in defining the ratio of the initial field energy to the initial plasma energy, the field energy may be computed using the characteristic field at the coil's center, but alternately, it may also be computed using the field at the explosion point. The latter takes into account the "effective" field energy at the location of the plasma and is thus the more physical of the two choices since what really matters is the field that the plasma pushes against. On the other hand, the former definition decouples the dependence on aspect ratio from the ratio of the field energy to the plasma energy, although it must be kept in mind that even if the ratio of the field energy to the plasma energy defined this way is kept fixed, varying the aspect ratio will vary the effective (i.e., physically meaningful, as opposed to defined) ratio of the field energy to the plasma energy. Actually, as long as the geometry of the thruster is specified along with a ratio of the field energy to the plasma energy defined under *some* convention, all necessary information for characterizing an expansion flow turns out to be provided in the case of a thin-shell problem. In this work, the initial ratio of the field energy to the plasma energy, ER, was defined by specifying the field energy relative to the coil position, and in particular, as given in Equation (4.17). The initial plasma energy in the thin-shell model is simply $Mv_0^2/2$, where M is the total mass of the shell and v_0 is its initial (radial) expansion speed.

4.1.2 Calculation of the Magnetic Pressure at the Interface

4.1.2.1 Quasi-static Maxwell's Equations in a Vacuum

To obtain the magnetic pressure at the interface, Maxwell's Equations were solved in the vacuum region bounded by the perfectly conducting plasma surface, the current-coil surfaces, and infinity. Maxwell's Equations for a vacuum region are

$$\nabla \times \underline{\mathbf{E}} = -\frac{1}{c} \frac{\partial \underline{\mathbf{B}}}{\partial t} \quad (4.18)$$

$$\nabla \times \underline{\mathbf{B}} = \frac{1}{c} \frac{\partial \underline{\mathbf{E}}}{\partial t} \quad (4.19)$$

$$\nabla \cdot \underline{\mathbf{E}} = 0 \quad (4.20)$$

$$\nabla \cdot \underline{\mathbf{B}} = 0 , \quad (4.21)$$

but if these full equations are employed, it will be necessary to follow the calculations on the characteristic time scales of propagation of information at the speed of light. This would be impractical as well as unnecessary because the time scale of expansion of the plasma is very much longer than such time scales.

For purposes of computing the magnetic field for incorporation into the equations of motion for the plasma shells, the vacuum fields can be considered to adjust themselves instantaneously to the new interface location at each instant of time. This is called the quasi-static approximation, and the Maxwell Equations reduce to

$$\nabla \times \underline{\mathbf{B}} = 0 \quad (4.22)$$

$$\nabla \cdot \underline{\mathbf{B}} = 0 . \quad (4.21)$$

These equations can be solved conveniently by making use of the magnetic stream function.¹ The contours of constant magnetic stream function value coincide with the magnetic field lines and thus makes the approach physically intuitive. The magnetic stream function ψ is defined through

$$B_r = -\frac{1}{r^2 \sin\theta} \frac{\partial \psi}{\partial \theta} \quad (4.23)$$

$$B_\theta = \frac{1}{r \sin\theta} \frac{\partial \psi}{\partial r} \quad (4.24)$$

so that Equation (4.21) is automatically satisfied. The requirement that Equation (4.22) be satisfied leads to

$$\frac{\partial^2 \psi}{\partial r^2} - \frac{1}{r^2} \cot\theta \frac{\partial \psi}{\partial \theta} + \frac{1}{r^2} \frac{\partial^2 \psi}{\partial \theta^2} = 0 . \quad (4.25)$$

In order to solve this equation numerically on a finite grid but yet rigorously accommodate the boundary at infinity, the coordinate transformation

$$\xi = \beta \arctan(\alpha r) \quad (4.26)$$

was chosen to map the infinite physical domain onto a finite computational domain. The constants α and β were selected so as to provide good resolution where most needed.

Using relation (4.26), Equation (4.25) becomes

$$\alpha^2 \frac{\partial^2 \Psi}{\partial \theta^2} - \alpha^2 \cot \theta \frac{\partial \Psi}{\partial \theta} + \beta^2 \alpha^2 \sin^2 \frac{\xi}{\beta} \cos^2 \frac{\xi}{\beta} \frac{\partial^2 \Psi}{\partial \xi^2} - 2\beta \alpha^2 \sin^3 \frac{\xi}{\beta} \cos \frac{\xi}{\beta} \frac{\partial \Psi}{\partial \xi} = 0, \quad (4.27)$$

and the expressions for the components of the nondimensionalized magnetic field become

$$B_r' = - \frac{\frac{\partial \Psi'}{\partial \theta}}{\tan^2 \frac{\xi}{\beta} \sin \theta} \quad (4.28)$$

$$B_\theta' = \frac{\beta \cos^2 \frac{\xi}{\beta} \frac{\partial \Psi'}{\partial \xi}}{\tan \frac{\xi}{\beta} \sin \theta}. \quad (4.29)$$

4.1.2.2 Boundary Conditions for the Field Calculation

Next, we shall consider the boundary conditions. The stream function set equal to an arbitrary constant defines a family of curves. In this work, we shall choose that constant of integration so the stream function will be zero along the central axis of the thruster. By making this choice, the stream function at infinity will also have zero value because the magnetic field line running along the central axis "wraps around" infinity. Now, when a perfectly conducting pellet is injected into a magnetic thruster along the central axis, displacing the fields as it does so, and a perfectly conducting plasma blob is formed out of it, the value of the magnetic stream function on the surface of the plasma may also be considered to be zero. Finally, we must compute the magnetic stream function value on the surface of the field-coil structure. If the field coil structure itself is perfectly conducting, surface currents induced on the coil structure will prevent the penetration of external fields into its interior. By assuming both the plasma and field coil structures to be perfectly conducting (an idealization), and choosing the prescription that the flux through

the field coils stays fixed with time as the plasma expands from negligible radius, the values of the magnetic stream function on the surfaces of the coil structures (as well as the plasma) will remain constant throughout the plasma expansion process. Therefore, it is necessary only to calculate the magnetic field at the initial instant and thereby obtain the stream function value that will be tied to the field-coil surface throughout the simulation. The exact expression for the magnetic field that is due to a single current loop can be found by using the vector potential defined through

$$\underline{B} = \nabla \times \underline{A} . \quad (4.30)$$

Employing the Coulomb gauge,

$$\nabla \cdot \underline{A} = 0 , \quad (4.31)$$

one obtains, by substituting Equation (4.30) into one of Maxwell's Equations,

$$\underline{A}(\underline{x}) = \frac{1}{c} \int \frac{\underline{J}(\underline{x}')}{|\underline{x} - \underline{x}'|} d^3 \underline{x}' . \quad (4.32)$$

The closed form expression for the magnetic field that is due to a single current loop of radius R_c is then found to be, in a spherical coordinate system (R, Θ, Φ) which has its origin at the loop center,³

$$A_\Phi(R, \Theta) = \frac{4IR_c}{c\sqrt{R_c^2 + R^2 + 2R_cR\sin\Theta}} \left\{ \frac{(2 - k^2)K(k) - 2E(k)}{k^2} \right\} , \quad (4.33)$$

where $k^2 = \frac{4R_cR\sin\Theta}{R_c^2 + R^2 + 2R_cR\sin\Theta}$ and $E(k)$, and $K(k)$ are complete elliptic integrals.

Figure 4.2 shows the spherical coordinate system used in this field calculation. The components of the magnetic field in this coordinate system can be computed as

$$B_R = \frac{1}{R\sin\Theta} \frac{\partial}{\partial \Theta} (A_\Phi \sin\Theta) \quad (4.34)$$

$$B_\Theta = -\frac{1}{R} \frac{\partial}{\partial R} (A_\Phi R) \quad (4.35)$$

$$B_\Phi = 0. \quad (4.36)$$

The solution to the magnetic field in (R, Θ, Φ) coordinates can be transformed to expressions in a spherical coordinate system centered at the explosion site by simple

trigonometry. In the case of multiple current loops, superposition of the vector fields created by each current loop can be used.

The finite difference form of the elliptic partial differential Equation (4.27) was solved on a rectangular ξ - θ coordinate grid at each time step with the boundary conditions just described. For this calculation, existing elliptic equation solving software packages, EPDE1 (CERNLIB) and ELLPACK, were used.

With EPDE1, the solution to ψ for an equation of the form $a \frac{\partial^2 \psi}{\partial X^2} + b \frac{\partial^2 \psi}{\partial Y^2} + c \frac{\partial \psi}{\partial X} + d \frac{\partial \psi}{\partial Y} + e\psi + f = 0$ is outputted at grid points, but the software does not have provisions for calculating any of the derivatives of ψ . The method of solution is finite differencing on a rectangular grid, and the boundary of the computational domain must lie entirely in the first quadrant. The boundary is specified by a set of points at which boundary values are to be supplied and between which the boundary is assumed to be straight and the value of ψ is taken to vary linearly unless it is a Neumann boundary. If a boundary point lies on or too close to a grid line, that point is moved slightly off the grid line, although this deforms the boundary a little. One shortcoming of EPDE1, apart from some problems with the software, is the rather low limit it places on the grid resolution (there is also a limit on the allowed number of boundary points).

The grid must also be uniform in the space in which the finite difference equations are being solved. To comply with these limitations, a 100 by 30 grid in ξ - θ space was usually selected for representing the upper physical half-plane. Under these conditions, the calculation of the gradients of the stream function at the interface became difficult and required careful curve fittings. This caused some inaccuracies to enter the simulation. The delicacy of having to deal with thin-shell elements complicated matters. However, the use of a uniform grid in rectangular ξ - θ coordinate space, which implies inaccurate field calculations at very large r , is not of as much concern as it may appear to be, because by the time the shell elements leave the immediate vicinity of the thruster, the effect of the fields on the motion of the shell elements becomes minimal. The reliance on a coarse underlying

grid, on the other hand, also limited how small, and to some extent, what shape, the cross section of the field-coil structure could be made. In these numerical simulations, the surface artificially chosen to represent the surface of the perfectly conducting field-coil structure, namely, one on which a magnetic stream function value that stays fixed throughout the expansion process is prescribed, was chosen to be as small as possible under the restrictions imposed by the underlying grid system. For the case of a single coil, the surface of the coil structure's cross-sectional area was formed by connecting with straight lines in rectangular ξ - θ coordinate space, four points roughly centered about the intended position of the coil. Since such a surface will not coincide exactly with any stream function contour that is due to an infinitesimally thin current loop, the magnetic stream function value assigned to be constant over the surface was taken to be the average of the stream function values that would be present along the surface when only an infinitesimally thin loop was present. Because of the largeness and awkward shaping of the cross-sectional area representing the actual field-coil structure's cross section, the resultant field will not be exactly the same as the one that is due to an infinitesimally thin current loop, but the similarity could be, and was, made quite large. Furthermore, because the input data had to be chosen rather carefully at each time step for the code to work, and in particular, the grid and boundary points could not be too close together because of the finite differencing procedure employed, a difficulty arose from having to run multiple time steps successively without being allowed to start each time step anew after inspecting the would-be input data. To add to this, the software's predetermined boundary shifting prescription sometimes introduced problems of its own.

Because of the limitations with the EPDE1 software, calculation of the quasi-static vacuum magnetic field was later switched to a more robust and less restrictive elliptic equation solver called ELLPACK.⁴ ELLPACK is a high level software package for solving elliptic boundary value problems and is characterized by many optional ways by which an equation may be solved. It is implemented as a FORTRAN preprocessor that

converts a user-supplied program written in ELLPACK language to FORTRAN, and is supported by a large software library to be called by the generated program. The ELLPACK language itself is an extension of FORTRAN, and it allows one to state and solve elliptic partial differential equations in two-dimensions on general domains or in three-dimensions on rectangular domains. Difficult problems can be solved by using problem solving modules within the ELLPACK framework. The fact that ELLPACK is allowed to be tailored by the user proved helpful. ELLPACK has the capability of calculating the solution function and its various derivatives at any point off the grid mesh points (via interpolations), allows nonuniform and variable grids to be used for more reliable calculations, allows more freedom in the choice of domain boundaries with respect to the grid, and can take boundaries defined in a parameterized way. Also, the use of a larger number of grid cells (and boundary points) than possible with EPDE1 is allowed, especially through authorized modifications of the source code. However, the calculation still failed, for example, when a grid coordinate and a boundary coordinate got too close during the course of a simulation, and to minimize these problems, the ELLPACK source code was modified, and furthermore, prescriptions for shifting grid lines were included in the calling program. The relative positioning of boundaries and grid lines still required care. Boundary parameterizations and grid spacings could not be varied erratically either, and too small a distance between boundary points relative to the grid spacing also caused problems. In addition, the use of a nonuniform grid was restricted to when the number of grids at all time steps was held fixed. Successful runs to late time steps generally required 64-bit precision.

Now, the algorithm for a thin-shell simulation proceeds as follows :

- 1) solve the magnetic stream function in the vacuum region using the boundary for the present time step, and calculate the magnetic pressures at the location of the shell elements,

2) solve the equation of motion for the Lagrangian shell elements using the magnetic pressure forces just obtained, and advance the position of the shell elements (this new position becomes the boundary for the next time step).

The specific thruster geometries and parameters used in the simulations were based on proposed designs by Hyde and the British Interplanetary Society (see Chapter 3). For the base-line of the high-field thruster designs, a single-coil device with an approximately 650 cm radius coil through which roughly 22 MA of current flows was chosen. The value of AR was taken to be $\tan(1 \text{ radian}) \approx 1.56$, and a total plasma mass of 0.9 gram and an initial expansion speed of 10^8 cm/sec were selected. These numbers result in a value of ER equaling about 15.3. The baseline parameter values were also used as the characteristic values employed in nondimensionalizing the variables in the various runs presented in this chapter. For simulations of the low-field thruster, a four-coil cusp-field system with the coil locations and current strengths as given in the Daedalus Report was used as a basis, although only approximately.

4.2 Sample Run for a Single-Coil High-Field Thruster

All simulations for the high-field situation were made with the single-coil design. Runs were made with both ELLPACK and EPDE1 for the field calculations, although all simulations for single-coil thrusters whose results are presented in this chapter are those that used ELLPACK. With ELLPACK, a standard second-order-accurate central finite differencing scheme was selected. For purposes of comparison, the geometry and size of the cross section representing that of the field-coil structure (on whose surface the magnetic stream function is tied to throughout the simulation) in the ELLPACK runs were chosen to be the same as those employed in the EPDE1 runs. The grid on which the field calculations

were carried out with ELLPACK were also kept similar to those employed in the EPDE1 calculations except for the use of increased resolution near "coil" surfaces and other boundaries of the system to enhance the reliability of the computations. When runs employing EPDE1 and ELLPACK were compared, the qualitative aspects of the plots remained the same, with less difference, such as in the divergence angle of the exhaust contour, between the two than the difference brought about by whether particles are pushed in cylindrical coordinates or spherical coordinates. This implied the decent reliability of the simulations. Finer grid resolutions were tested, but without marked improvement in the results. Ninety-one shell ring elements, each subtending equal polar angles at the initial instant, were used in all the runs presented, and they were numbered from 1 to 91, going counterclockwise from $\theta=0$ to π . In all simulations, the uniform shell at the initial instant was taken to be perfectly spherical, with a radius of approximately 3.93 cm.

The result of running a code employing the thin-shell approximation for a single-coil high-field thruster with the baseline parameters is presented in Figure 4.3. Figure 4.3-i displays the field lines at the initial instant. The ICF explosion site is marked by a small bump on the central axis. The shell expands, initially isotropically, against the vacuum magnetic fields, and the magnetic field contours are shown every 2×10^{-6} seconds starting 10^{-6} seconds after the initial instant in Figures 4.3-ii through -ix. The interface at later time steps appear bumpy with indentations, but these are only artifacts of the way the contour plotting is carried out and the way the stream function value is interpolated in regions outside the computational domain. The actual shell contours on which $\psi=0$ holds are perfectly smooth, as seen in the interface contours (i.e., the location of the shell elements) of Figure 4.4. The derivatives of the magnetic stream function were also being calculated smoothly along the interface curve. In Figure 4.4a, the interface contours are being plotted every 2×10^{-6} seconds (so the separation between the individual contours at early phases of the expansion is roughly 200 cm), starting with the initial instant, while in Figure 4.4b,

they are being plotted every 4×10^{-6} seconds (in the rest of Section 4.2, "one time step" will refer to 10^{-7} seconds).

One of the deficiencies of the thin-shell model alluded to earlier is evident in these plots. Because of the absence of hydrodynamical pressure, mass elements in regions where the recoiling fields push the material back towards the central axis can fall straight onto the central axis. With a real fluid, finite internal pressures will prevent such a process from occurring. Elliott, working only with the Daedalus-type low-field thruster, had to terminate his simulations well before reaching the stage at which this phenomenon may be observed. Since this is a phenomenon brought about through the use of hypothetical thin-shell elements, the prescription of what to do with the mass elements as they collide with the central axis must be chosen artificially. Because of azimuthal symmetry about the central axis, it is unphysical for particles to fall through the central axis. Although mass elements that fall onto the central axis could be made to reflect off the central axis, they were simply prescribed to stay on the central axis, retaining the axial velocity they had upon colliding with the axis (this allows conservation of momentum, although not kinetic energy, which is the only form of plasma energy present in a thin-shell model). After shell elements fall onto the central axis, they cease to contribute to the thrusting process. Under this prescription, one sees no part of the blob expanding in the upstream direction through the constriction of the fields in this high-field run, but there does exist a stream of (a fair number of) mass elements flowing freely with no deceleration along the central axis. The exact prescription of how the particles fall onto the central axis and move thereafter had a somewhat visible effect, along with round-off errors, on the shape of the blob at late stages of the flow.

Now, particles were pushed in physical cylindrical coordinates (i.e., Cartesian in cross section) instead of in spherical polar coordinates to prevent them from following false trajectories. When particles are pushed in spherical coordinates, the difference between motion along a constant radius arc and a straight line tangent to it introduces spurious

deviations in particle paths away from the real trajectories. Inaccuracies will be especially large near the origin where a small displacement in physical space can correspond to a large displacement in the θ coordinate. This will be discussed again in a later chapter.

In Figures 4.5 a through g, we see the shell element locations in physical space every 10 time steps for a) the 15th through 25th shell elements, b) the 25th through 35th shell elements, c) the 35th through 45th shell elements, d) the 45th through 55th shell elements, e) the 55th through 65th shell elements, and f) the 65th through 75th shell elements. In this run, most elements initially traveling in the forward direction are reflected backwards and fall onto the axis. Mass elements initially traveling more or less normally away from the central axis are significantly redirected to flow in a rearward axial direction, and this is resulting in a rather strong axial collimation of the overall flow. Even for this high field a device, there is minimal deceleration of the far downstream side of the plasma interface (i.e., shell elements traveling rearward in the near-axial direction from the start) where the fields are weak, but this does not entirely invalidate the thin-shell approach because the large expansion ratio and nonzero interfacial pressure, even if very low, will still cause a thin shell to develop, although it will not necessarily stay at the interface (see Chapter 7).

It can be seen that significant deceleration of the shell elements and their redirection occur mostly while the shell elements are localized within a short distance of the location at which they are halted (and during a short time relative to the characteristic expansion time scale of the bulk, although compared to the low-field thruster case, these properties are far less pronounced). The magnetic fields are acting almost as a rigid wall (against which the particles reflect) placed at the location where the shell elements are redirected, and this property generally held true for all cases run. Note that (especially with no rebound of the forward portion of the "plasma") the entire "plasma" is undergoing inertial motion after it has left the general region (of influence) of the magnetic thruster, indicating that the thrusting process has mostly completed by that stage.

Figures 4.6 and 4.7 show, respectively, the temporal variation of the radial coordinate and the polar angle θ of a) the 10th shell element, b) the 20th shell element, c) the 30th shell element, d) the 40th shell element, e) the 50th shell element, f) the 60th shell element, g) the 70th shell element, and h) the 80th shell element.

Figures 4.8 and 4.9 show, respectively, the x and y component velocities, where the x axis is the central axis pointing in the direction of $\theta=0$ (the direction of vehicle acceleration), and the y axis is along $\theta=\pi/2$, of a) the 10th shell element, b) the 20th shell element, c) the 30th shell element, d) the 40th shell element, e) the 50th shell element, f) the 60th shell element, g) the 70th shell element, and h) the 80th shell element, versus time. Particle velocity reversal, fall onto the central axis, and axial collimation are clearly observed in these plots. Shell elements at relatively small θ are experiencing the earliest deceleration. When complete inertial flow is established, the x and y component velocities become constant.

Figure 4.10 shows the nondimensionalized magnetic pressure along the interface as a function of θ in radians, every forty time steps, starting with the initial instant. Unlike in Chapter 7, the nondimensionalized magnetic pressure plotted is p_B' as given in Equation (4.13). Note that the magnetic pressure decreases sharply on the downstream side of the interface. As noted earlier, the calculation of field pressure at later time steps become somewhat inaccurate because of the coarseness of the grid for the field calculation at large distances away from the explosion site.

4.3 Thrust and Efficiency Calculations

4.3.1 Calculation of Thrust in a Magnetic Thruster

With jet propulsion systems, the primary objective is the production of thrust, and a fundamental issue is the efficiency with which thrust is produced. With a pulsed magnetic thruster in which part of the jet may leak in the forward direction, thrust and efficiency calculations differ somewhat from those for conventional nozzles.

Thrust production in a thruster comprised of magnetic fields can be thought of in terms either of electromagnetic forces acting over a distance or magnetic pressures acting on field-coil surfaces. Currents create magnetic fields. The field that is due to an element of current of length $d\mathbf{l}$ through which a current I is flowing, is

$$d\mathbf{B} = \frac{I d\mathbf{l} \times \mathbf{x}}{c|\mathbf{x}|^3} \quad (4.37)$$

at a location separated from the element of current by the vector \mathbf{x} . This is Biot-Savart's law (which, after taking the curl of its integral form, becomes the differential form of Ampere's Law). Now, when a current distribution is placed in a field, say, due to another current distribution, each current distribution will feel a net force. This net force felt by a current passing through a magnetic field is

$$d\mathbf{F} = \frac{I d\mathbf{l} \times \mathbf{B}}{c} \quad \text{or} \quad \mathbf{F} = \frac{1}{c} \int \mathbf{J}(\mathbf{x}) \times \mathbf{B}(\mathbf{x}) d^3x . \quad (4.38)$$

The $\mathbf{J} \times \mathbf{B}$ force can be expressed as the sum of a magnetic pressure and tension term as

$$\mathbf{J} \times \frac{\mathbf{B}}{c} = \frac{c}{4\pi} (\nabla \times \mathbf{B}) \times \frac{\mathbf{B}}{c} = -\frac{\nabla B^2}{8\pi} + \frac{(\mathbf{B} \cdot \nabla) \mathbf{B}}{4\pi} , \quad (4.39)$$

where the first term is the magnetic pressure acting normally to the field lines, and the second term is the magnetic tension acting along the field lines. Both are body forces, so if one does not have magnetic fields and currents coexisting inside a volume, but instead, has a surface current flowing on a perfect conductor with fields existing to its exterior, then only the force exerted on the boundary will arise, the pressure being $B^2/8\pi$.

The sign of the force between currents is such that parallel currents attract and antiparallel currents repel. For example, a single current loop has a tendency to burst radially apart, and this can also be interpreted in terms of magnetic pressure. However, self-fields alone obviously result in no net force acting on the center of mass of a current loop. In the magnetic thruster of an ICF pulse rocket, an axially symmetric conducting plasma expands in a location offset from the plane of the coil (i.e., a plane of symmetry of the thruster's magnetic field), and the current through the field coils of the thruster and the induced diamagnetic plasma currents interact so that an equal and opposite axial force is felt by both the plasma and the vehicle. It should be clear by now that one way to calculate the thrust on a vehicle to which the field coils are attached is to integrate the magnetic pressure over the surface of the coils; i.e.,

$$\underline{F}(t) = \oint_{\text{coil surface}} \frac{B^2}{8\pi} \hat{n} dA . \quad (4.40)$$

The average thrust, which is the average rate at which net momentum is imparted to a vehicle, can be calculated, for a pulsed device, by integrating the instantaneous thrust $\underline{F}(t)$ over the pulse period τ (inverse of the pulse frequency), and dividing that by τ ; i.e.,

$$F_{\text{avg}} = \frac{\int_0^\tau \underline{F}(t) dt}{\int_0^\tau dt} . \quad (4.41)$$

The total impulse delivered by one pulse can be multiplied by the pulse frequency to give the average thrust.

However, in the numerical simulations, the magnetic field could not be calculated satisfactorily on the surface of a current "coil" by differencing data on neighboring grid points because the values of the magnetic stream function are not given accurately on the low-resolution grid and computations of its derivatives in the immediate vicinity of the oddly shaped "coil" surface with sharp corners are very unreliable. For example, the magnetic field lines slightly exterior to the "coil" surface tended to arch around the broad and nearly flat surfaces of the "coil" and the magnetic pressure forces concentrated at the sharp corners. Such pressure distributions do not lend themselves to accurate treatment in

numerical computations. Several different "coil" cross section shapes were investigated, but marked improvements were not possible.

Therefore, a more indirect means of calculating thrust than using Equation (4.40) was chosen (although not all losses, such as the usually negligible momentum loss that is due to anisotropic photon radiation, will be handled correctly). By assuming that all momentum lost by a plasma blob is transferred to the vehicle, the net impulse delivered to the vehicle by one blob from the beginning of the plasma expansion to a particular instant in time can be calculated by taking the difference between the net axial momentum of the plasma blob at that instant and that at the initial instant; i.e.,

$$\left(\begin{array}{l} \text{Net impulse} \\ \text{delivered to vehicle} \\ \text{from time 0 to t} \end{array} \right) = \oint_{\text{blob}} v_x dm \Big|_0 - \oint_{\text{blob}} v_x dm \Big|_t . \quad (4.42)$$

With an initially isotropic expansion, the first term on the right-hand side of this equation vanishes. The instantaneous thrust can be obtained by taking the time derivative of the cumulative impulse delivered to the vehicle.

4.3.2 Calculation of Propulsive Efficiencies

As we have seen in Chapter 2, thrust is generally presented as

$$F = \dot{m} u_{\text{exh}} , \quad (4.43)$$

where \dot{m} is the mass flow rate and u_{exh} is the axial exhaust velocity corrected for any angular divergence of the exhaust stream. For a pulse propulsion system, this concept still holds on a time-averaged sense. If the initial expansion velocity after complete conversion of internal energy to kinetic energy is v_0 and if rethermalization of kinetic energy (the unfavorable aspect of this process is that the rethermalized energy may not all be returned to the kinetic energy of an efficiently directed flow in timely fashion) can be considered negligible (both concerns are not applicable in a thin-shell model), the maximum impulse deliverable by one pulse of mass M will be Mv_0 . Thus, if the pulse frequency is f , the

maximum possible average thrust obtainable will be Mv_0f , and for a device with a thrust efficiency of ϵ_{th} , defined through

$$\epsilon_{th} = \left(\begin{array}{l} \text{Cumulative impulse} \\ \text{delivered to vehicle} \\ \text{by one pulse.} \end{array} \right) / (Mv_0) , \quad (4.44)$$

i.e., the average thrust will be $Mv_0f\epsilon_{th}$.

The efficiency of a thruster can also be discussed in an energywise way by measuring the fraction of the initial plasma's total energy that ends up in the kinetic energy of the rearward axially directed flow, which is the only useful component of the jet. It was thus decided to define a jet kinetic energy efficiency through the formula

$$\epsilon_{ke} = \frac{\lim_{t \rightarrow \infty} \int_0^\pi v_x^2(\theta_0, t) \left\{ -\frac{v_x(\theta_0, t)}{|v_x(\theta_0, t)|} \right\} dm(\theta_0)}{\int_0^\pi v_0^2 dm(\theta_0)} . \quad (4.45)$$

This was considered more convenient and appropriate for our purposes than some of the other definitions of propulsive efficiencies found in the literature. The numerator in this expression is proportional to the kinetic energy that is due to the axial component of the (shell) particles' velocities with the forward jet's contribution counted negatively, and the denominator is proportional to the initial energy of the (shell) particles. By letting the velocities in an equation of the form (of course, when dealing with a fluid, we will have dependence on r_0 as well) of Equation (4.45) be the sum of the directed bulk and random thermal velocities, thermal losses can theoretically be taken care of. This form, when generalized to a fluid, can take into account radiation losses as well as conversions between internal (except for the energy stored in nontranslational modes at the initial instant) and kinetic energies.

As an aside, the efficiencies for an initially uniform and isotropic expansion can be written more generally (omitting the limits $t \rightarrow \infty$), as

$$\epsilon_{th} = \frac{-\int v_x dm}{v_0 \int_M dm} \quad (4.44')$$

and

$$\epsilon_{ke} = \frac{-\int v_x^2 \left(\frac{v_x}{|v_x|} \right) dm}{v_0^2 \int_M dm} . \quad (4.45')$$

This shows that the square of the thrust efficiency does not in general equal the jet kinetic energy efficiency.

4.3.3 Efficiencies for a Rigid Paraboloidal Reflector

Now, the most efficient thruster geometry for a bomb propulsion system is a rigid paraboloid (against which the expanding debris experiences an elastic collision) with the explosion site located at its focus. Particles emanating from the focus of a paraboloid and intercepted by the paraboloidal walls are redirected so that they exit parallel to the central (i.e., optical) axis. Assuming an isotropic point (shell) source at the focus, a simple calculation shows that the thrust efficiency for a paraboloidal reflector will be

$$\epsilon_{th} = \frac{\lim_{t \rightarrow \infty} \int_0^\pi \{ |v_0(\theta_0)| \cos \theta_0 - v_x(\theta_0, t) \} dm(\theta_0)}{\int_0^\pi |v_0(\theta_0)| (1 + \cos \theta_0) dm(\theta_0)} . \quad (4.46)$$

Although this type of form can take into account conversions between kinetic and thermal energies in translational modes, it cannot take into account radiation losses because photon momentum has not been incorporated into the derivation, but this should have negligible impact. Now, using Equation (4.5) for dm (i.e., $dm = \frac{M}{2} \sin \theta_0 d\theta_0$) and assuming that the paraboloidal wall subtends a polar angle θ_s when viewed from the focus, Equation (4.46)

becomes

$$\epsilon_{th} = \frac{0 - \frac{1}{M} \left\{ \int_0^{\theta_s} (-v_0) \frac{M}{2} \sin\theta_0 d\theta_0 + \int_{\theta_s}^{\pi} v_0 \cos\theta_0 \frac{M}{2} \sin\theta_0 d\theta_0 \right\}}{v_0 + 0} \quad (4.47)$$

$$= \frac{1}{4} (1 - \cos\theta_s)(3 + \cos\theta_s) .$$

The thrust efficiency correctly reaches unity for a paraboloidal nozzle of infinite extent. The jet kinetic energy efficiency as defined by Equation (4.45) becomes, for a rigid paraboloidal reflector with an isotropic explosion site located at its focus, from which a uniform shell of total mass M emanates, initially spherically, with a radial velocity of v_0 ,

$$\epsilon_{ke} = \left\{ \begin{array}{l} \frac{\int_0^{\theta_s} \frac{1}{2} v_0^2 \frac{M}{2} \sin\theta_0 d\theta_0 - \int_{\theta_s}^{\frac{\pi}{2}} \frac{1}{2} (v_0 \cos\theta_0)^2 \frac{M}{2} \sin\theta_0 d\theta_0 + \int_{\frac{\pi}{2}}^{\pi} \frac{1}{2} (v_0 \cos\theta_0)^2 \frac{M}{2} \sin\theta_0 d\theta_0}{\frac{1}{2} M v_0^2} \\ = \frac{1}{2} \left(\frac{4}{3} - \cos\theta_s - \frac{\cos^3\theta_s}{3} \right) \quad \text{for } 0 \leq \theta_s \leq \frac{\pi}{2} \\ \\ \frac{\int_0^{\theta_s} \frac{1}{2} v_0^2 \frac{M}{2} \sin\theta_0 d\theta_0 + \int_{\theta_s}^{\pi} \frac{1}{2} (v_0 \cos\theta_0)^2 \frac{M}{2} \sin\theta_0 d\theta_0}{\frac{1}{2} M v_0^2} \\ = \frac{1}{2} \left(\frac{4}{3} - \cos\theta_s + \frac{\cos^3\theta_s}{3} \right) \quad \text{for } \frac{\pi}{2} \leq \theta_s \leq \pi . \end{array} \right. \quad (4.48)$$

This efficiency also tends to unity as θ_s approaches π radians. These theoretical thrust and jet kinetic energy efficiencies for a rigid paraboloidal nozzle are plotted in Figures 4.11a and b, respectively. When assessing the efficiencies of a thruster, a possible figure of merit is the angle of subtention (as viewed from the focus/explosion site) of an idealized paraboloidal reflector wall required to give itself the same efficiency as that of the thruster under concern. This "effective paraboloidal angle of subtention" obtained by matching the thrust efficiencies of a magnetic thruster and a rigid paraboloidal nozzle with the explosion point located at its focus will in general not be equal to that obtained by matching the jet

kinetic energy efficiencies of the two. In fact, depending upon the parameters and geometries of the thruster, the departure between these two effective paraboloidal angles of subtention can be quite large, implying the reflection process in such a thruster to be quite different from that in an ideal, rigid paraboloidal thruster. The deviation of the behavior of a magnetic thruster from that of a rigid paraboloidal nozzle with the explosion point located at its focus may be crudely quantified by the percent deviation between these two angles. This percent deviation will be defined as the excess of the former angle over the latter angle multiplied by 100 and divided by the arithmetic average of the two angles. This quantity will be denoted by DPA. As will be seen subsequently, the propulsive efficiencies decreased with increasing DPA, vice versa,* and the maximum propulsive (both thrust and jet kinetic energy) efficiencies were obtained when DPA reached a minimum.

4.3.4 Efficiencies in the Single-Coil Sample Run

Returning to the simulation results of our baseline thruster, the cumulative impulse delivered to the vehicle (directly proportional to the thrust efficiency) by a single explosion is shown in Figure 4.12a as a function of time, while in Figure 4.12b, the instantaneous thrust is plotted as a function of time, as derived from the data of Figure 4.12a. The instantaneous thrust is strongly peaked around the time the plasma shell is most heavily decelerated and reverses its direction of motion. In about twice the time it takes to reach this stage, the thrust is again down to low values, and thereafter, the thrust gradually trails down to negligible values. For the present case, the influence of the fields on the motion of the shell elements becomes quite negligible beyond a distance from the field coil roughly equal to the diameter of the coil. The fact that particles whose velocities are reversed tend to fall onto the axis in the thin-shell-approximation model does contribute to the rather early

* To reach this conclusion, cases to the same side of the case for maximum efficiency in Figures 5.13 and 5.14 were compared with each other.

termination of the impulse transfer between the plasma and the vehicle. The magnetic fields extending out in the exhaust stream do help collimate the exhaust, but on the whole, those fields do not impede the rearward flow very much unless we have an extremely high field device. After the bulk of the impulse is transferred to the vehicle, Figure 4.12a still shows the net impulse delivered to the vehicle, slowly increasing with time instead of completely leveling off. This is due to round-off errors and is not real. The true value of the net impulse delivered to the vehicle by one pulse should be taken to be the value at about the point at which the curve of Figure 4.12a turns over completely from a steep rise to a very shallow rise. The large oscillations present in the thrust curve (all false) are due to having calculated the thrust by taking the time derivative of the cumulative impulse curve without incorporating any averaging procedures. Most of the noise arises from particles falling onto the central axis. The thrust efficiency, taking into account the fact that the final shallow rise in the impulse curve is false, is about 0.655. The average thrust at a pulse rate of 100 Hz would be about 6.2×10^4 Newtons.

Figure 4.12c is a plot of the development of the jet kinetic energy efficiency with time. Again, the shallow rise of this curve after a turnover from a steep rise, is numerical. Taking this into account, the jet kinetic energy efficiency is about 0.515.

Finally, these efficiencies were compared with the performance of a rigid-walled paraboloidal nozzle with the explosion point located at its focus. It turns out that a rigid paraboloidal nozzle will give a thrust efficiency of 0.655 when the paraboloidal walls subtend an angle of 79.9 degrees as viewed from the focus, and a jet kinetic energy of 0.515 when the walls subtend an angle of 72.9 degrees. The value of DPA for the baseline case is 9.16 %.

4.4 Comparison of Cases with Different Parameter Ratios for Single-Coil Thrusters

In this section, the effects of using various ratios of initial field energy to initial plasma energy (ER) and various aspect ratios (AR) will be compared for single field-coil thrusters. The adjustable parameters in a run that uses the thin-shell approximation are the plasma mass, initial plasma velocity, current strengths in the coils, coil radii, geometrical setup of the coil system, and location of the explosion site relative to the field-coil system. But in a nondimensionalized analysis, only the two parameter ratios, ER and AR, characterize the process for a single-coil thruster.

In Figure 4.13, we see the effect varying the value of ER has on the performance for a single-coil thruster with a value of AR of $\tan(1 \text{ radian}) \approx 1.56$. The horizontal axis is the $\log_{10}(\text{ER})$ for a thruster normalized by the value of ER for the base-line thruster ($\text{ER}_0 \approx 15.3$). In Figure 4.13a, the vertical axis is the thrust efficiency, while in Figure 4.13b, it is the jet kinetic energy efficiency. Both efficiencies are peaking at values of ER somewhat higher than that of the base-line value. In Figure 4.14, we see the effect varying the value of AR has on the performance of a single-coil thruster with a value of ER of approximately 15. The horizontal axis is the arctangent of the value of AR for the thruster, in degrees. In Figure 4.14a, the vertical axis is the thrust efficiency, while in Figure 4.14b, it is the jet kinetic energy efficiency. Both efficiencies are peaking at values of AR in the vicinity of the base-line value. The values of AR at which the efficiencies are maximized are expected to change with the value of ER.

We will now look at some of the results of several individual cases. In Cases 1 through 4, the value of ER was varied with the value of AR held fixed at the base-line value. In Cases 5 through 7, the value of ER was held fixed at the base-line value, while the value of AR was varied.

4.4.1 Case 1

In this run, the value of ER is 0.01 times that of the base-line value of approximately 15.3. If the field-coil radius and current strength are kept the same as in the base-line case, this can be achieved, for example, by keeping the mass of the plasma (M) the same as that of the base-line case and increasing the initial expansion velocity (v_0) an order of magnitude, or by keeping the initial velocity the same as that of the base-line case and increasing the mass two orders of magnitude, etc. For purposes of discussion it will be assumed that the thruster field coil is 650 cm in radius, coil current is roughly 22 MA (MA-turns by default), plasma mass is 0.9 gram, and initial expansion velocity is 10^9 cm/sec. For the purpose of comparing results, all other single-coil cases will also be assumed to have these values of coil radius, current, and plasma mass.

Figure 4.15 shows the interface contours in physical space (upper half-plane) every 10^{-6} seconds, starting with the initial instant when the interface appears as a dot on the axis.

Figure 4.16 shows the location, at ten times this frequency, of a) the 15th to 25th shell elements, b) the 25th to 35th shell elements, c) the 35th to 45th shell elements, d) the 45th to 55th shell elements, and e) the 55th to 65th shell elements. No shell elements have fallen on the axis. Case 1 is a rather low field device and the fields, except for along (and close to) the line of sight of the field coil, do not have enough strength to decelerate the flow drastically. Along (and close to) the line of sight of the field coil, the field compression arising from flux conservation does result in velocity reversal of the shell elements, but there again, because of the narrowness of the recoiling region, the shell elements that are once redirected to travel (virtually radially) inwards, instead of falling onto the central axis, start moving outwards again (without having acquired much of a θ direction velocity). Thus an oscillation of the elements within the chamber beyond the first bounce is occurring. The shell elements that reverse direction come very close to the field

"coil" surface before turning around. Shell elements traveling in a direction that would have resulted in oblique incidence upon rigid thruster walls, had they been present, are redirected (as if incident on such rigid walls) with a nonnegligible θ -component velocity.

Figure 4.17 shows the nondimensionalized interfacial magnetic pressure as a function of θ in radians every 10^{-6} seconds, starting with the initial instant. Unlike the case of higher field thrusters, the magnetic pressure averaged over the interface peaked during only a short period of time when the shell elements were suddenly decelerated and redirected.

Figures 4.18a, b, and c show, respectively, the cumulative impulse delivered to the vehicle, the instantaneous thrust, and the jet kinetic energy efficiency, as functions of time. It can be seen that most of the thrusting is concentrated in a narrow interval of time around the epoch during which interface elements are reversed in their direction of motion. The negative thrust observed is entirely due to noise in the impulse curve, and is not real. Since the thrusting process is not finished by the end of the simulation, the efficiency values obtained at the end of the plots are not final, but the thrust efficiency of about 0.195 (measured at the end of the simulation) can be obtained from a paraboloidal nozzle with the explosion site located at the focus if the paraboloid subtends an angle of 37.4 degrees at the focus, and similarly, the jet kinetic energy efficiency of about 0.104 (at the end of the simulation) can be obtained by a paraboloidal nozzle subtending about 27.1 degrees at the focus. The tentative value of DPA using these results is 31.9 %. The propulsive efficiencies of these rather low field single-coil thrusters are not very high. In general, for thrusters having low values of ER, as in this case, a multiple-coil design becomes necessary for achieving decent propulsive efficiencies because, as we have seen, many of the shell elements not along the line of sight of a coil are not affected heavily by the fields.

4.4.2 Case 2

In this case, $ER \approx 3.83$, one quarter of the base-line value. The following discussion will be made, assuming that the case represents one in which the initial expansion velocity of the shell is 2×10^8 cm/sec.

Figure 4.19 shows the location of the shell elements in physical space every 10^{-6} seconds, starting with the initial instant. Shell elements are seen falling onto the central axis without re-expanding outwards again.

Compared to the lower field situation of Case 1, the predominant thrusting period and the (related) period during which the interfacial magnetic pressure stays high have now spread over a greater fraction of the characteristic time span of the bulk expansion. The thrust efficiency and jet kinetic energy efficiency at the end of the simulation were 0.477 and 0.324 (numerical inaccuracies are expected to be resulting in a 1 to 2 % underestimation here), and these efficiencies can be obtained by a rigid paraboloidal nozzle subtending 63.5 degrees and 52.4 degrees, respectively, at the focus. The value of DPA for the thruster is 19.2 %, although it should be kept in mind that the DPA values quoted in this chapter are only DPA values valid under the thin-shell approximation, with all its inaccuracies.

4.4.3 Case 3

In this case, $ER \approx 61.2$, four times that of the base-line value. The following discussion will be made, assuming that this case represents one in which the initial expansion velocity of the shell is 5×10^7 cm/sec.

The plots of Figure 4.20 show the location of the shell elements in physical space every a) 2×10^{-6} seconds, and b) 4×10^{-6} seconds, starting with the initial instant. With values of ER this high, the magnetic field begins to exert a notable influence even on the

downstream side of the interface, leading to significant impedance of the flow even in those regions. The exhaust flow (or puff) is tapered, and the interface profiles are showing a better apparent collimation than in the base-line case, but it must be kept in mind that many more of the shell elements are falling onto the central axis in this case, some of which continue to travel in the forward direction. Also, compared to cases for lower ER, the anisotropy of the interface deformation is becoming less pronounced, and the expansion is starting to look similar to an expansion into a uniform field. The field in this case is so strong that the interface is pushed downwards towards the axis by the recoiling fields, even at large distances away from the coil (e.g., on the downstream side of the blob), resulting in the flattened shape of the interface. Although not shown, the interfacial magnetic pressure distribution with θ is also getting flatter than before, especially upon turnaround of the shell elements, and the deceleration process is becoming even more broadly distributed over the expansion time scale.

The thrust efficiency and jet kinetic energy efficiency for the thruster are about 0.637 and 0.463, respectively. The effective paraboloidal nozzle subtention angles are 78.2 degrees for the thrust efficiency and 67.2 degrees for the jet kinetic energy efficiency. The value of DPA for the thruster is 15.1 %. If the value of ER is raised very much further, the efficiencies, which have been rising with ER, start to decrease again (see Figure 4.13).

4.4.4 Case 4

In this run, $ER \approx 153$, one hundred times that of the base-line value. The following discussion will be made, assuming that the initial expansion velocity of the shell is 10^7 cm/sec.

Figure 4.21 shows the interface profiles in physical space with the contours plotted every 2×10^{-6} seconds, starting with the initial instant.

Figure 4.22 shows the location at the same time steps for a) the 15th to 25th shell elements, b) the 25th to 35th shell elements, c) the 35th to 45th shell elements, d) the 45th to 55th shell elements, e) the 55th to 65th shell elements, f) the 65th to 75th shell elements, and g) the 75th to 85th shell elements. The interface does not expand very much on the characteristic dimensions of the thruster before it is redirected. The redirection process is such that the shell elements are pushed down towards the central axis with an axial velocity that is small in magnitude relative to the magnitude of the velocity in the direction normal to the axis. All elements, including those expanding outwards in the downstream direction from the start, eventually end up on the central axis and stop contributing to thrust production. The resultant efficiency is quite low.

A thrust efficiency of about 0.268 was obtained, and such a value can be obtained by a rigid paraboloidal nozzle with the explosion site located at the focus if the paraboloid subtends an angle of 44.7 degrees at the focus. Similarly, the jet kinetic energy efficiency reached about 0.079, and this value can be obtained by a paraboloidal nozzle subtending about 23.4 degrees at the focus. The value of DPA for the thruster is 62.6 %.

4.4.5 Case 5

In Case 5, the value of AR for the thruster is $\tan(75^\circ) \approx 3.73$. For discussion purposes, the pellet mass, initial expansion velocity, radius and current strength of the current coil, will all be assumed to be the same as those quoted for the base-line design, and only the axial distance between the coil center and the explosion site will be assumed changed.

Figure 4.23 shows the interface profiles in physical space, every 2×10^{-6} seconds in Figure 4.23a, and every 4×10^{-6} seconds in Figure 4.23b, starting with the initial instant. This is a case in which the explosion site is placed close to the plane of the coil, and the constriction of the interface towards the axis in the downstream region is evident.

Figure 4.24 shows the shell element locations every 5×10^{-7} seconds for a) the 5th to 15th shell elements, b) the 15th to 25th shell elements, c) the 25th to 35th shell elements, d) the 35th to 45th shell elements, e) the 45th to 55th shell elements, f) the 55th to 65th shell elements, g) the 65th to 75th shell elements, and h) the 75th to 85th shell elements.

The thrust efficiency at the end of the simulation reached about 0.384, and such an efficiency can be obtained by a paraboloidal nozzle with the explosion site located at the focus if the paraboloid subtends an angle of 55.3 degrees at the focus. Similarly, the jet kinetic energy efficiency at the end of the simulation registered about 0.284, and this is a value that can be achieved by a paraboloidal nozzle subtending about 48.2 degrees at the focus. These efficiencies are markedly lower than those achieved by the base-line thruster. The value of DPA using these numbers is 13.7 %. The asymmetry in the fields of such a high AR device is not sufficient to produce high efficiencies. The interface profiles in Figure 4.23 show very good axial collimation of the flow, but again, not evident from the figure is the fact that a lot of shell elements have fallen onto the axis (traveling in both the forward and rearward directions).

4.4.6 Case 6

In this case, we have $AR = \tan(45^\circ) = 1$.

Figure 4.25 shows the interface profiles in the upper half of the physical plane every 2×10^{-6} seconds, starting with the initial instant.

Figure 4.26 shows the shell element locations every 5×10^{-7} seconds for a) the 5th to 15th shell elements, b) the 15th to 25th shell elements, c) the 25th to 35th shell elements, d) the 35th to 45th shell elements, e) the 45th to 55th shell elements, f) the 55th to 65th shell elements, g) the 65th to 75th shell elements, and h) the 75th to 85th shell elements. The shell elements expanding in the downstream direction are very little affected

by the fields, and the elements initially expanding generally in the upstream direction are redirected with a reversed axial velocity. This contrasts with the situation when the value of AR was higher.

The thrust efficiency at the end of the simulation of about 0.616 can be obtained by a paraboloidal nozzle with the explosion site located at the focus, if the paraboloid subtends an angle of 76.1 degrees at the focus, and similarly, the jet kinetic energy efficiency at the end of the simulation of about 0.451 can be obtained by a paraboloidal nozzle subtending about 65.8 degrees at the focus. The resultant value of DPA is 14.5 %.

4.4.7 Case 7

In this case, $AR = \tan(33^\circ) \approx 0.649$.

Figure 4.27 shows the interface profiles in physical space every 3×10^{-6} seconds, starting with the initial instant. The bunching together of the contours at the upstream end of the blob close to where they meet the axis is due to numerical round-off errors, and can be removed by adjustments of the code.

Figure 4.28 shows the shell element locations every 5×10^{-7} seconds for a) the 5th to 15th shell elements, b) the 15th to 25th shell elements, c) the 25th to 35th shell elements, d) the 35th to 45th shell elements, e) the 45th to 55th shell elements, f) the 55th to 65th shell elements, g) the 65th to 75th shell elements, and h) the 75th to 85th shell elements. With the explosion site placed this far downstream of the field coil, the upstream portion of the shell is strongly decelerated and pushed down onto the axis after velocity reversal, but the downstream side expands almost freely. At intermediate angles, the shell elements are redirected with a high axial component velocity.

A thrust efficiency of about 0.567 was obtained, and this value can be achieved by a paraboloidal nozzle with the explosion site located at the focus, if the paraboloid subtends an angle of 71.5 degrees at the focus. Similarly, the jet kinetic energy efficiency reached

about 0.372, and this is a value for a paraboloid subtending about 57.5 degrees at the focus. These numbers yield a value of DPA equaling 21.7 % for the thruster.

4.5 Sample Run for a Multicoil Low-Field Thruster

The thin-shell approximation was also applied to the Daedalus (first-stage) -type low-field thruster (see Figure 3.4) using EPDE1 for the vacuum magnetic field calculations. The same number of plasma shell ring elements as used in the previous simulations for single-coil thrusters was employed. As discussed earlier, a partial simulation of flow in the Daedalus (first-stage) thruster under the thin-shell approximation was also carried out by Elliott and Terry.²

Figure 4.29 shows the initial cusp-type field geometry of the Daedalus first-stage thruster in the upper physical half-plane as obtained through an exact analytical calculation (using the techniques described earlier). In Figure 4.29a, only the field lines passing equidistantly between coils 3 and 4 (see the table in Figure 3.4) are plotted, while in Figure 4.29b, the magnetic field vectors are plotted at select locations. As a minor point, the cusp in the magnetic field through which the ICF driver electron beam is intended to be injected, failed to pass through the center of the 50 meter hemispherical thruster where the ICF explosion site is to be located, but this could have been due to inaccuracies in the calculation.

In Figure 4.30, we see contours formed by perfectly conducting shell elements at various stages of "plasma" expansion in a thruster whose field-coil configuration roughly resembles that of the Daedalus thruster, but without a metallic flux-conserving wall surrounding the chamber: "roughly," because, as mentioned previously, the cross-sectional geometries enclosed by the surfaces on which the stream function is specified to stay fixed throughout a simulation do not exactly mimic those of real field-coil structures, especially

in that they are unrealistically large and awkwardly shaped. The "coil structures" used in the simulations were approximately centered about the position of the coils in the original design, except that in the run of Figure 4.30-i, they were roughly centered about the "real" coils' θ coordinate and a radial coordinate of 50 meters. With approximate prescriptions of the magnetic stream function on such unrealistic surfaces and the use of a coarse underlying grid, precise simulations of flow in a low-field thruster, in which the plasma interface approaches the field coils very closely, become impossible.

At the initial instant, the shell was assumed to be perfectly spherical with a radius of 3.93 cm and a purely radial expansion velocity of 1.5×10^9 cm/sec. The total pellet mass was taken to be 2.844 grams, in accord with the Daedalus parameters. In Figure 4.30-i, the interface is shown every 5×10^{-8} seconds, starting 5×10^{-8} seconds from the initial instant. Plotting was stopped when the interface along the axis reached approximately 38 meters. In this run, the inner surfaces of all four coil structures were approximately placed at a 45.3 meter radius. During the phases of expansion depicted in Figure 4.30-i, the motion of the interface is not influenced very much by the fields, but by the last time steps shown in that plot, the interface is somewhat flattened in the direction perpendicular to the thrust axis. For a successful continuation of the expansion simulation as the shell approaches the coils/walls more and more closely, the time step employed in the computations had to be made progressively smaller. Otherwise, for example, the interface could even cross the "coil" surface in one time step.

Although not consistent with the setup for the simulation shown in Figure 4.30-i, the remainder of the simulations whose results are presented in Figure 4.30 were begun with a perfectly spherical shell of 42 meter radius still assumed to be possessing the initial purely radial expansion velocity of 1.5×10^9 cm/sec. The interface along the central axis was prescribed to move in an undecelerated fashion. In these runs, the inner surfaces of the "coil structures" were placed at a radial distance of 50 meters for coils 1 and 2, and 54.95 meters for coils 3 and 4. In Figure 4.30-ii, the interface is shown every 5×10^{-8}

seconds for a period of 4.5×10^{-7} seconds after the first 5.5×10^{-7} seconds have passed since the 42 meter mark. Velocity reversal of shell elements was already occurring at this stage, but only as the interface approached the "coil" surfaces extremely closely. In Figure 4.30-iii, we see the interface every 10^{-7} seconds, starting 6×10^{-7} seconds after the last time step plotted in Figure 4.30-ii. Radial velocity data indicated significant turnaround to be occurring. In Figure 4.30-iv, the interface is plotted every 10^{-7} seconds, starting 7×10^{-7} seconds after the last time step plotted in Figure 4.30-iii. In Figure 4.30-v, the interface is being plotted for a period of 3×10^{-8} seconds, starting 2.1×10^{-7} seconds after the last time step plotted in Figure 4.30-iv. Velocity reversal was observed up through the downstream region slightly behind the rearmost coil by virtue of the magnetic field's extending into that region also being compressed somewhat, and probably even more so because of some "pull" from the neighboring elements in the forward region turning around. This indicates that even some of the interface elements behind the rearmost coil are not expanding in an entirely unrestrained manner. Note that the entire interface has acquired a rather flattened shape by this stage.

An oscillation (temporal and spatial) in the velocity of the interface shell elements, believed numerical, was observed during the intermediate phases of the onset of velocity reversal, and neighboring shell elements acquired alternating negative and positive velocity values. Eventually, all interface elements that were initially traveling in the forward direction of the hemisphere experience a uniform velocity reversal except for the near-axial elements, which, in the absence of walls and due to the weakness of the fields, leak out as a forward jet. Although it is true that oversized conductor surfaces were used to represent the field coils, reflection of the interface was observed without fingers of "plasma" leaking through between the adjacent coils, implying that sufficient current strengths were chosen in the original design.

Extending the simulation much beyond the last stage depicted in Figure 4.30 proved difficult because as mentioned previously, very large gradients in the stream function that

cannot be treated very accurately arose as the magnetic fields got highly compressed into a narrow region between the interface and the coils.

Simulations in which a flux-conserving wall was assumed to exist around the thruster region were also attempted, but again, because of limitations on how the flux-conserving surface could be specified, the effective wall resulted in being very thick, and the field geometry deviated, even more so than those of the wall-less cases, from that of the Daedalus design. With a hemispherical flux-conserving wall surrounding the thruster region, there is obviously no forward jet. The interface acquired a more spherical appearance. Velocity reversal was first observed most prominently in regions of highest initial field, but deceleration was strongest at later stages in regions of weakest initial field, where interface reflection occurred last and most suddenly. This result is consistent with Elliott's observations and is also characteristic of this type of expansion-deceleration process, fully hydrodynamic flows included (see Section 6.4 of Chapter 6 and Section 7.5 of Chapter 7).

The closest approach distance of the interface to the walls ranged from a few centimeters to many tens of centimeters. Again, velocity reversal first manifested itself numerically at alternating shell element locations, until finally, a smooth reversal over the entire forward half of the expansion was achieved. Shell elements in the rear half of the hemisphere also experienced some reduction in the radial expansion velocity, but nowhere is the motion close to coming to rest. In fact, in these low-field devices, the impedance to the flow by the magnetic pressure itself of the trailing fields behind the thruster is not very significant, because with low field devices, high field compression is necessary for redirecting the plasma, and this does not occur in those regions. Generally, shell elements did not acquire much of a polar component velocity during the redirection process. This is understandable because the thruster wall is hemispherical, and the interface which emanates from the wall's center of curvature is redirected only when it approaches the walls very closely. This implies a lower efficiency than if the exhaust had been allowed to acquire a

higher axial component via the reflection process, and also hints at the possibility of some of the plasma's not leaving the thruster in one bounce.

4.6 Summary

The expansion of a highly conducting plasma blob against a vacuum magnetic field may be very crudely modeled, with total neglect of interior fluid dynamics, by the thin-shell model in which all of the plasma is assumed to be collected into a thin shell at the interface. Although a previous work ² has applied this approach to simulate plasma expansion in the thruster of an ICF pulse rocket, the present work is a more complete undertaking in that unlike the earlier work, the entire plasma debris is treated (not just a section) and the simulation is taken until the debris is well downstream of the thruster (not just until turnaround). Also, different thruster parameters were tested, and their performances (e.g., propulsive efficiencies) were assessed.

The thin-shell approximation is not entirely accurate because it is not true that in a real expansion of a perfectly conducting fluid against a vacuum magnetic field, all of the mass will be collected into a thin shell right at the interface. Because of the deceleration of the interface, a shell-like region of high mass concentration will form, but as we shall see in Chapter 7, this "shell" will not necessarily stay at the interface. Furthermore, the very large fraction of the volume within the interface exclusive of the "shell" will have a low (relative to the "shell" region), yet finite, density. No internal fluid dynamics or internal properties such as fluid pressure can be taken into account in the thin-shell model. On the whole, the approximations used render the efficiency estimates obtained from this model overidealized and optimistic.

One of the consequences of taking the thin-shell approach is that unphysical phenomena can occur. Most notably, some of the reflected shell elements can fall onto the

central thrust axis, and once this happens, realistic calculations for those "fallen" elements cannot be made. This also results in the characteristic shape taken on by the interface in the thin-shell simulations at late stages of the expansion. With particles falling onto the axis, the apparent collimation of the exhaust is not always a good indicator of the propulsive efficiencies, either.

Despite its deficiencies, the thin-shell simulation allows simulations of acceptable accuracy (within the framework of treating the plasma as a perfectly conducting single-species fluid) to be made without placing excessive demands on computing resources, making possible the running of a variety of cases for a parametric study of the behavior of the interface motion and the resultant thruster efficiencies.

The two parameters determining thruster operation and design under the thin-shell approximation were found, both in theory and through simulations, to be the ratio of the initial field energy to the initial plasma energy (ER) and the geometric setup of the thruster. The latter may be expressed in terms of the aspect ratio (AR) for a single-coil thruster.

For the single-coil thruster, promising propulsive efficiencies were obtained for a range of thruster parameters, ER and AR, in the high-field regime. Of the cases investigated in this work, the highest thrust efficiency attained was 69% and the highest jet kinetic energy efficiency attained was 56%, both for the case of $ER \approx 30$ and $AR \approx 1.56$. The propulsive efficiencies monotonically tapered off to lower values for higher values of ER or AR as well as for lower values of these two parameters. But as seen in Figures 4.13 and 4.14, the efficiencies stay fairly respectable for a rather broad range of ER and AR. It is believed that the value of AR at which the maximum efficiencies are obtained will vary with ER, and vice versa. Interface motion suspiciously (especially in light of the types of codes and computing resources required to carry out a hydrodynamical simulation for a flow in the thruster of an ICF pulse rocket) similar to those of a work⁵ that was claimed to be due to a hydrodynamical simulation was obtained.

Now, if one has a rigid paraboloidal nozzle with the explosion point located at its focus, all particles intercepted by the wall will exit purely axially in the rearward direction, and those that aren't intercepted obviously will travel unimpeded. The effective angle of subtention of such a nozzle, as seen from its focus, which would give the rigid nozzle the same thrust efficiency as a magnetic thruster under concern, and that which would give the same jet kinetic energy efficiency, were computed and compared. It was found that the propulsive efficiencies of the thruster were maximized when the discrepancy between the two effective paraboloidal angles of subtention (DPA) was minimized. This result is reasonable in light of the fact that a paraboloid with the explosion point located at its focus is the most efficient reflector.

For the low-field type, a thruster based on the Daedalus design was investigated, but taking the expansion much beyond velocity reversal of the upstream side of the interface was difficult. The downstream side of the interface exhibited an expansion that was not entirely spherical (somewhat flattened), apparently mostly because of some influence from neighboring upstream elements. As the results from the forward half of the interface show, the fields in the downstream region are not strong enough (they never get highly compressed) to influence the interface motion as much as was observed. But one can see that a treatment of only the forward hemisphere, as was carried out in an earlier work,² will not give a complete picture, even for the Daedalus-type thruster. Also, depending upon the setup, an outward rebound of the reflected interface elements of the forward hemisphere could occur for low-field designs (with a fluid blob and not a shell, this can occur even for a high-field design; see Chapter 7).

The production of thrust peaked around the time the interface elements were very strongly decelerated and redirected. With high-field thrusters, the deceleration process occurred gradually throughout the expansion process and without the interface approaching the field coils too closely. But with low-field thrusters, the deceleration of the interface could be achieved only when the interface elements approached the field coils very closely

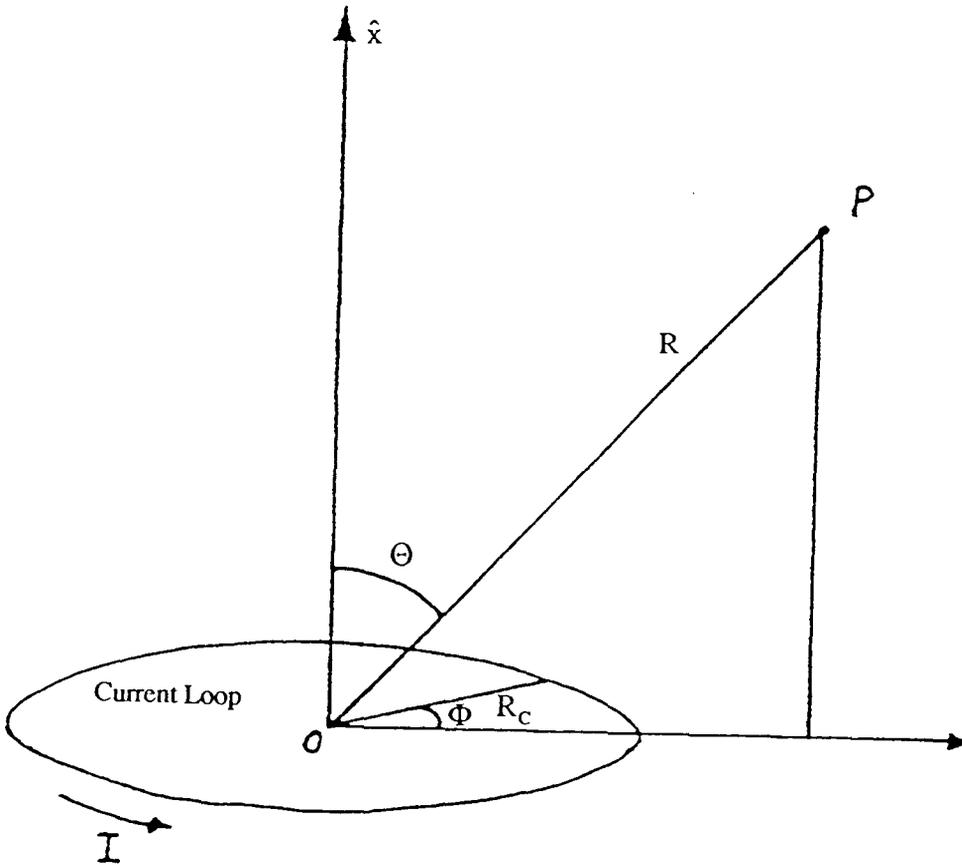
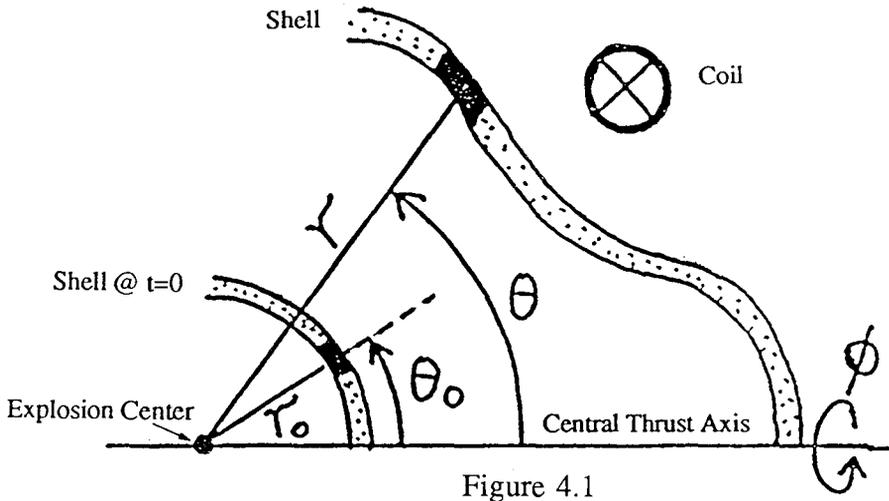
because of the need to compress the fields highly to obtain a high enough interfacial magnetic pressure (in agreement with the results of Reference 2). Interface deceleration in the low-field thrusters thus occurred suddenly and very strongly when the interface approached the field coils very closely. It can therefore be seen that low-field thrusters will require a multicoil design. In general, the deceleration/redirection of the interface elements was found to be concentrated while the interface was localized in a narrow region in space, even for quite high field designs, and the process was akin to reflection from a rigid wall.

References for Chapter 4

1. J.W. POUKEY, *Phys. Fluids* **12** (1969), 1452.
2. J.O. ELLIOTT and W.K. TERRY, *JBIS* **38** (1985), 120.
3. J.D. JACKSON, *Classical Electrodynamics, 2nd ed.* (John Wiley and Sons, NewYork, 1975).
4. J.R. RICE and R.F. BOISVERT, *Solving Elliptic Problems using ELLPACK* (Springer-Verlag, NewYork, 1985).
5. R.A. HYDE, "A Laser Fusion Rocket for Interplanetary Propulsion," Paper IAF-83-396, 34th Congress of the Int'l Astronautical Federation, Budapest, Oct. 1983.

Appendix of Chapter 4

Figures for Chapter 4



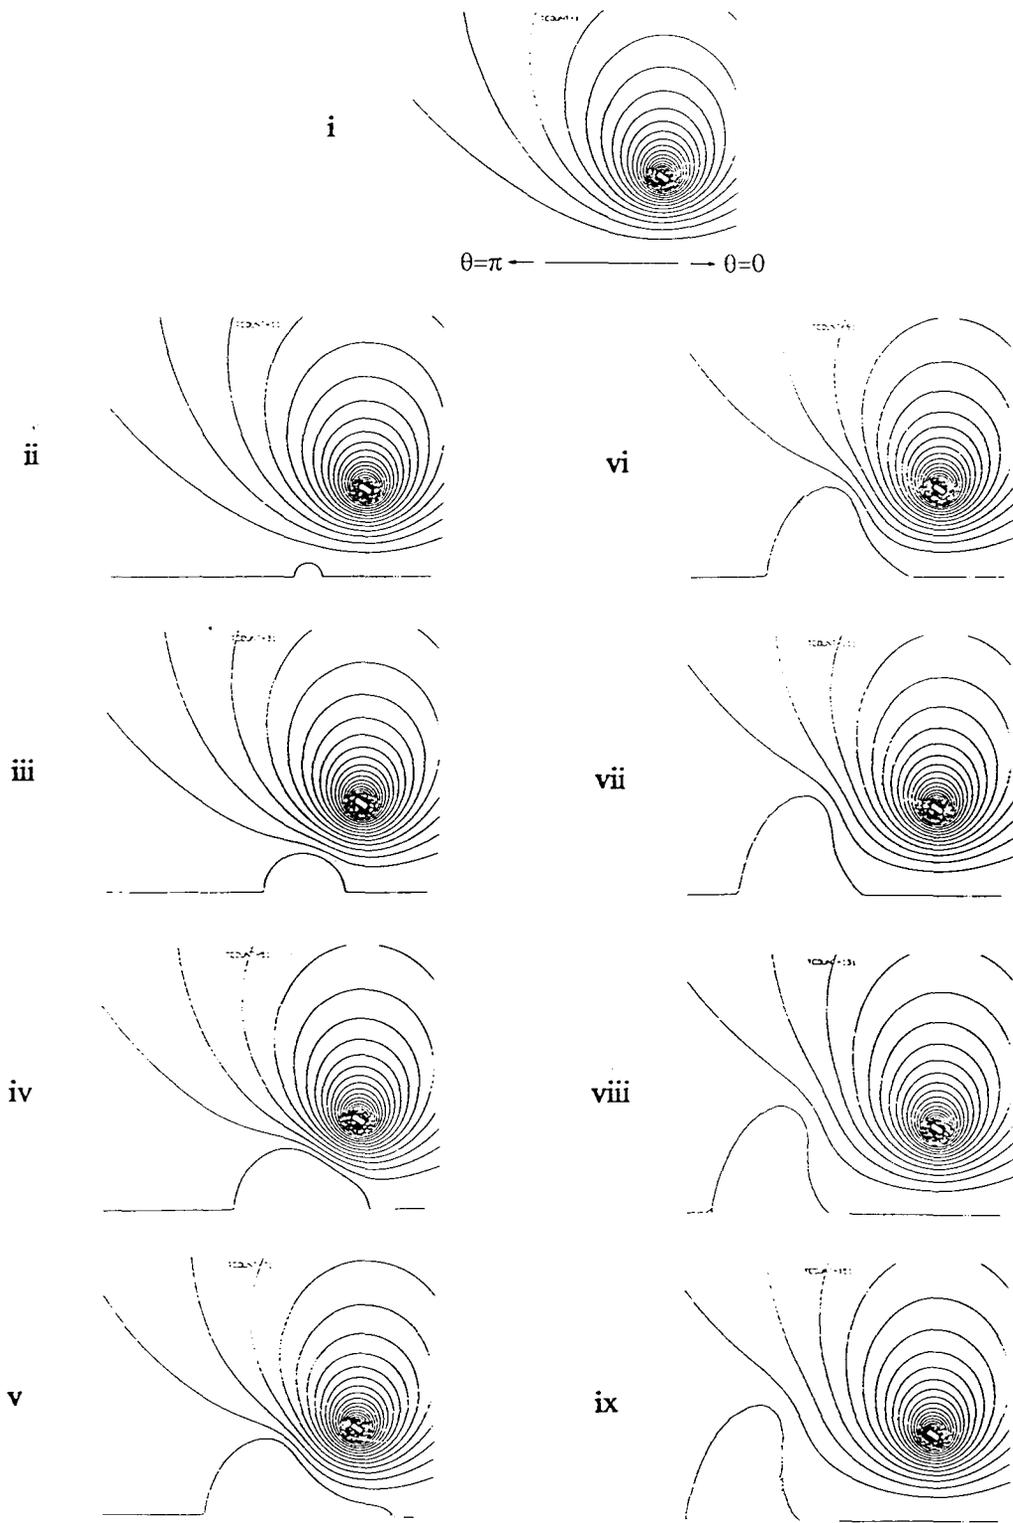
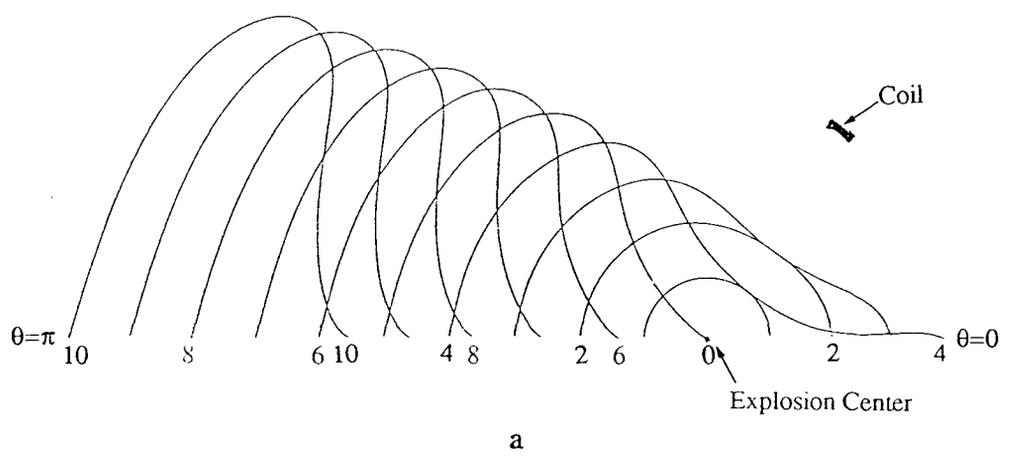
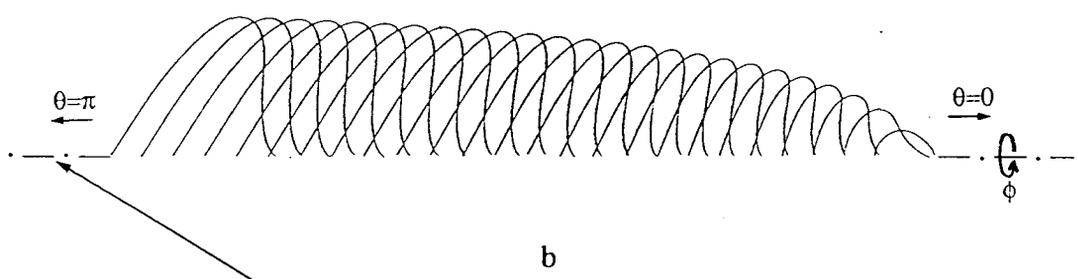


Figure 4.3



(As a convention, in Chps. 4 and 7, interface locations plotted have been numbered successively with time.)



(Convention for all plots of the interface in Chps. 4 and 7.)

Figure 4.4 Location of shell elements for the baseline case (ER=15.3. AR=1.56).

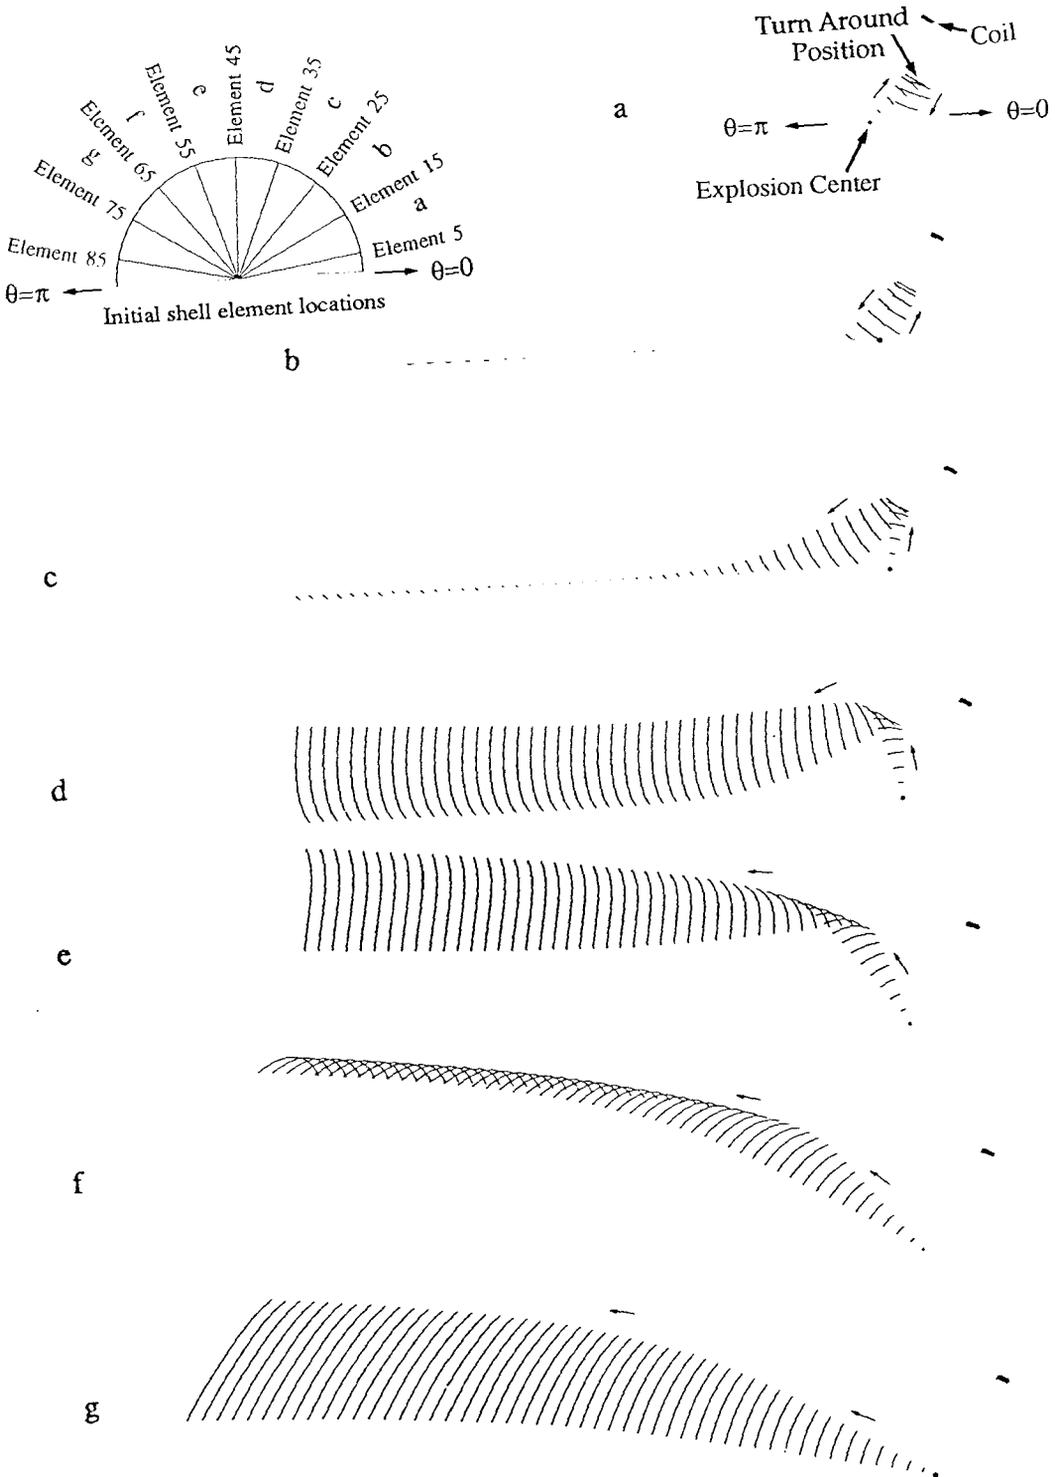


Figure 4.5

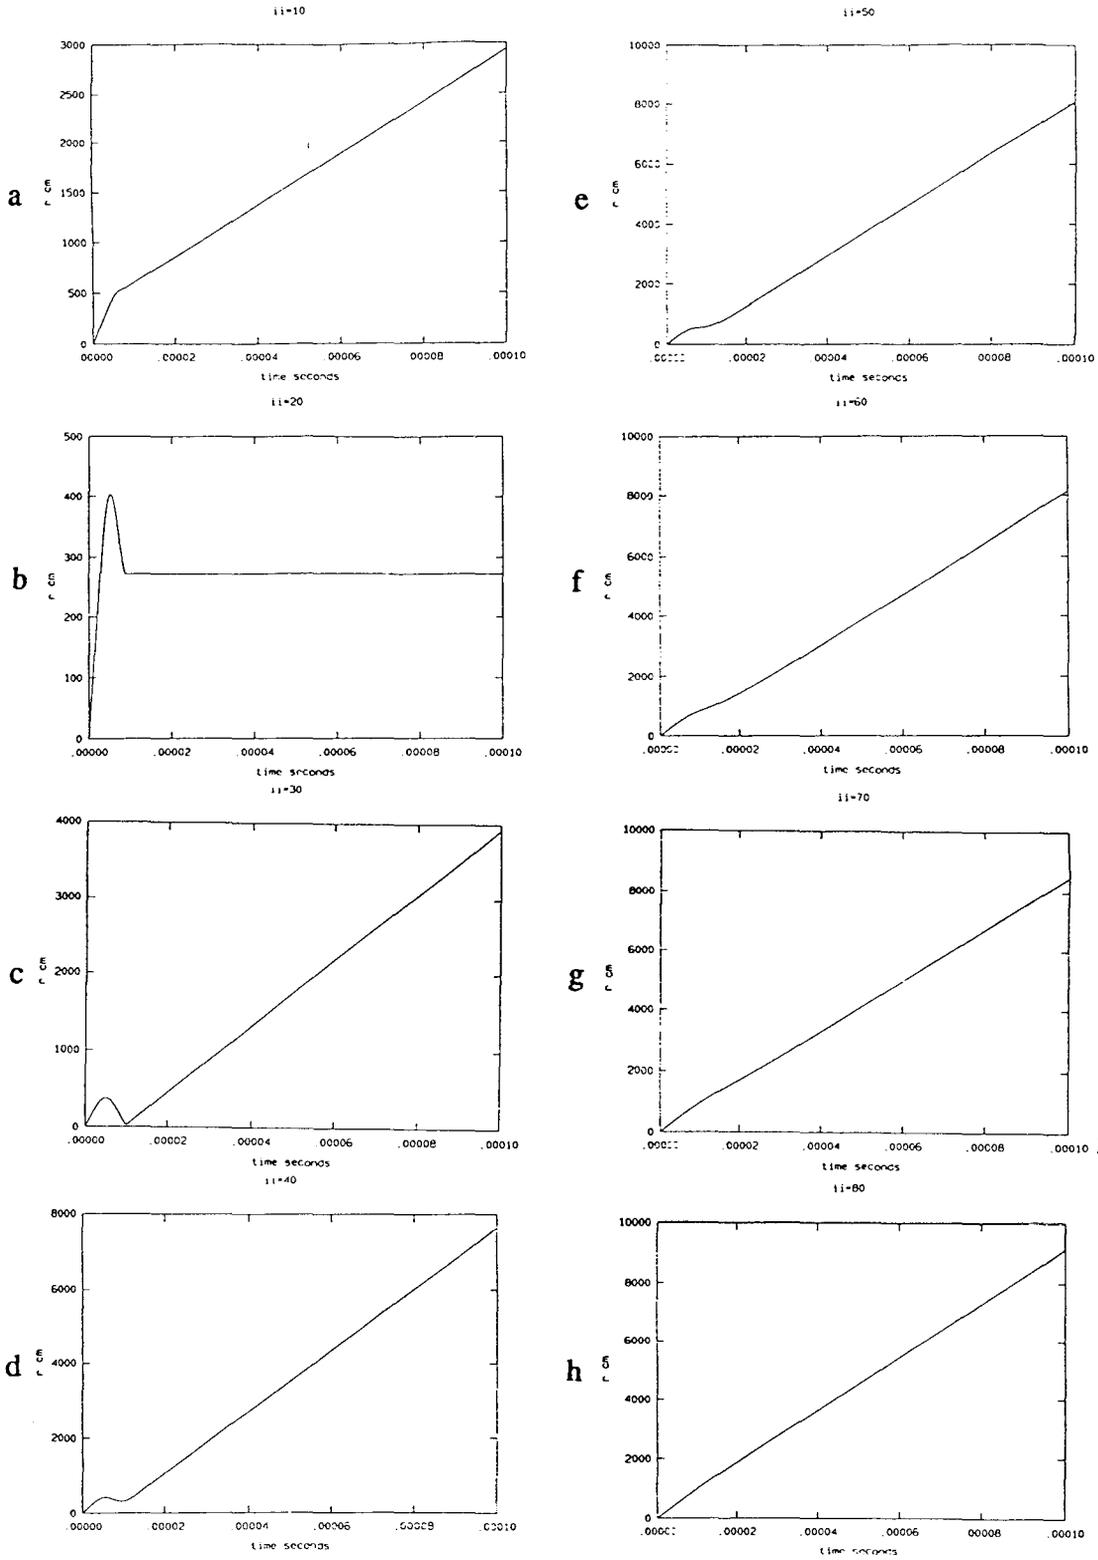


Figure 4.6 Radial coordinate of shell elements in cm (vertical axis) v.s. time in seconds (horizontal axis).

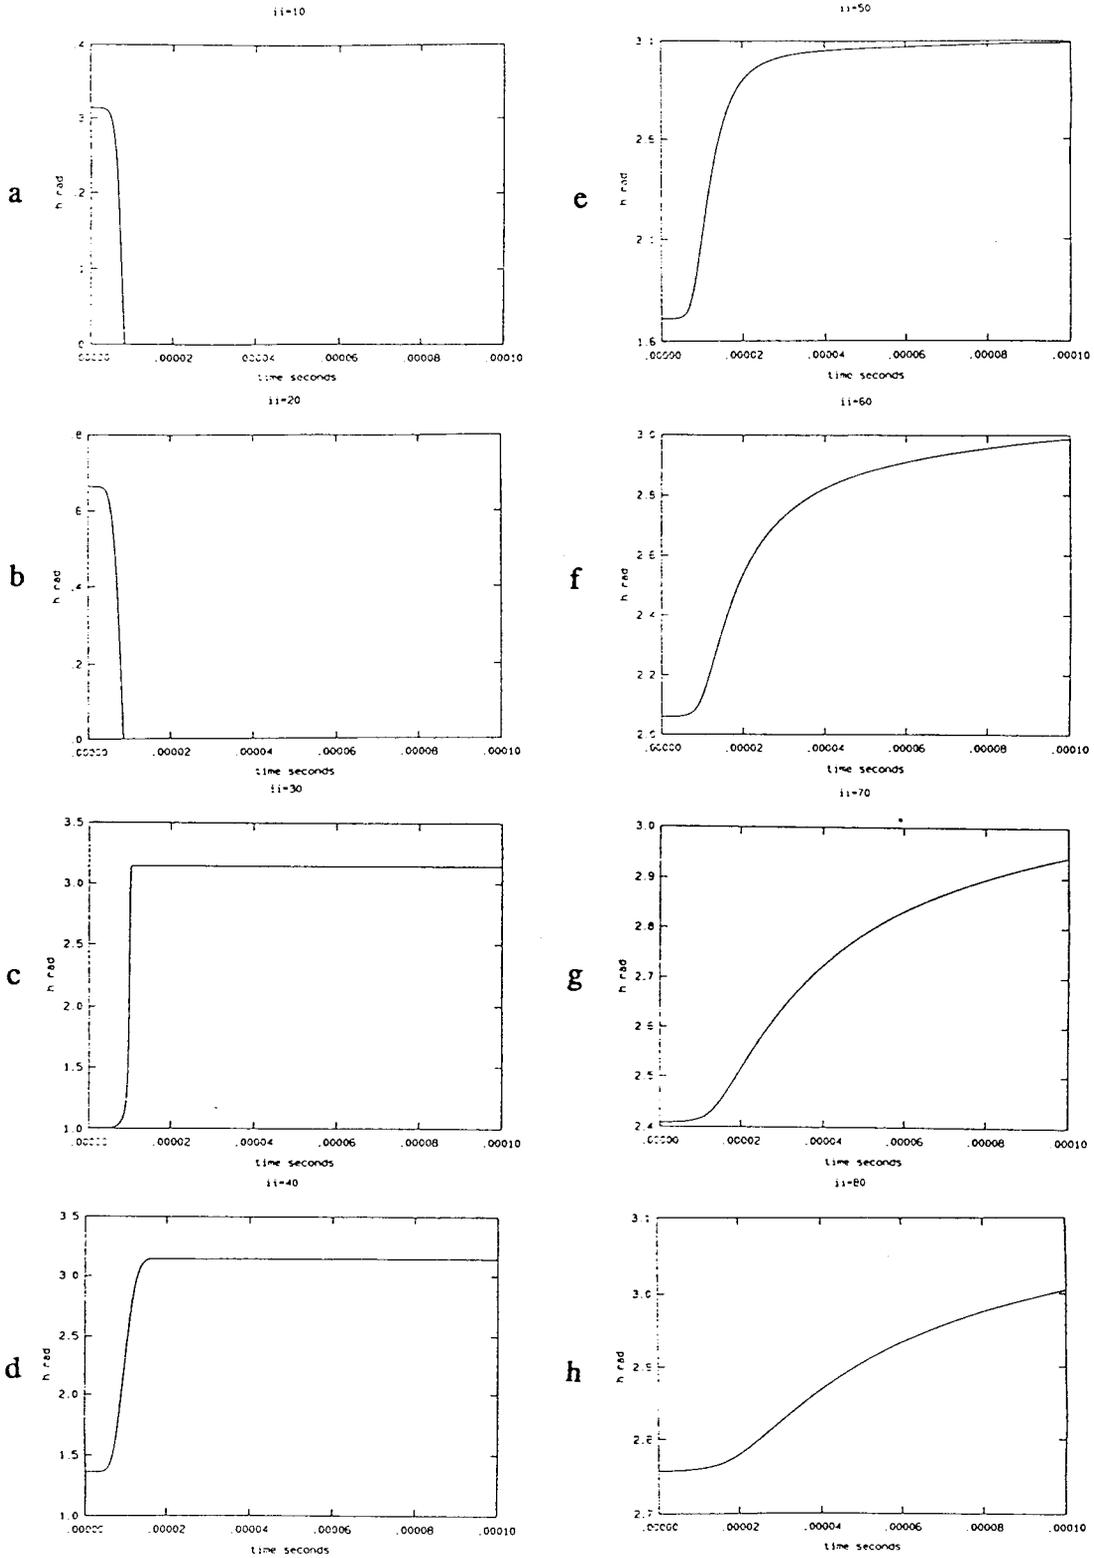


Figure 4.7 Polar angle (θ) of shell elements in radians (vertical axis) v.s. time in seconds (horizontal axis).

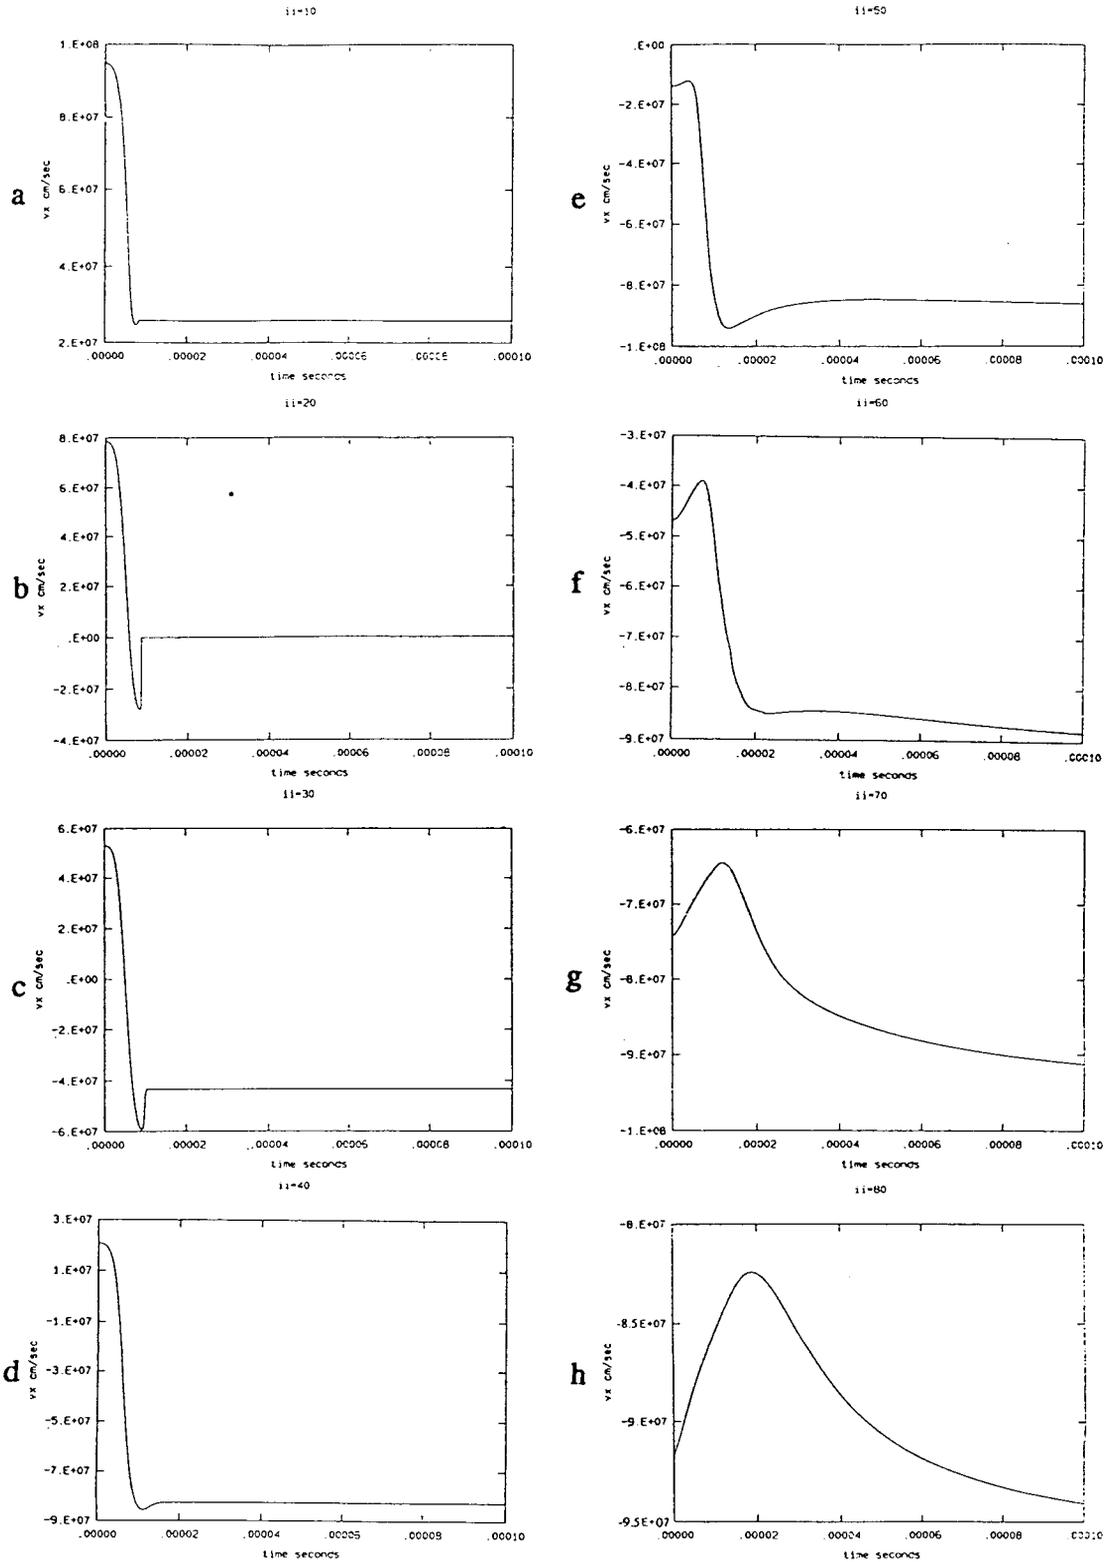


Figure 4.8 X-component (axial) velocity of the shell elements in cm/sec (vertical axis) v.s. time in seconds (horizontal axis).

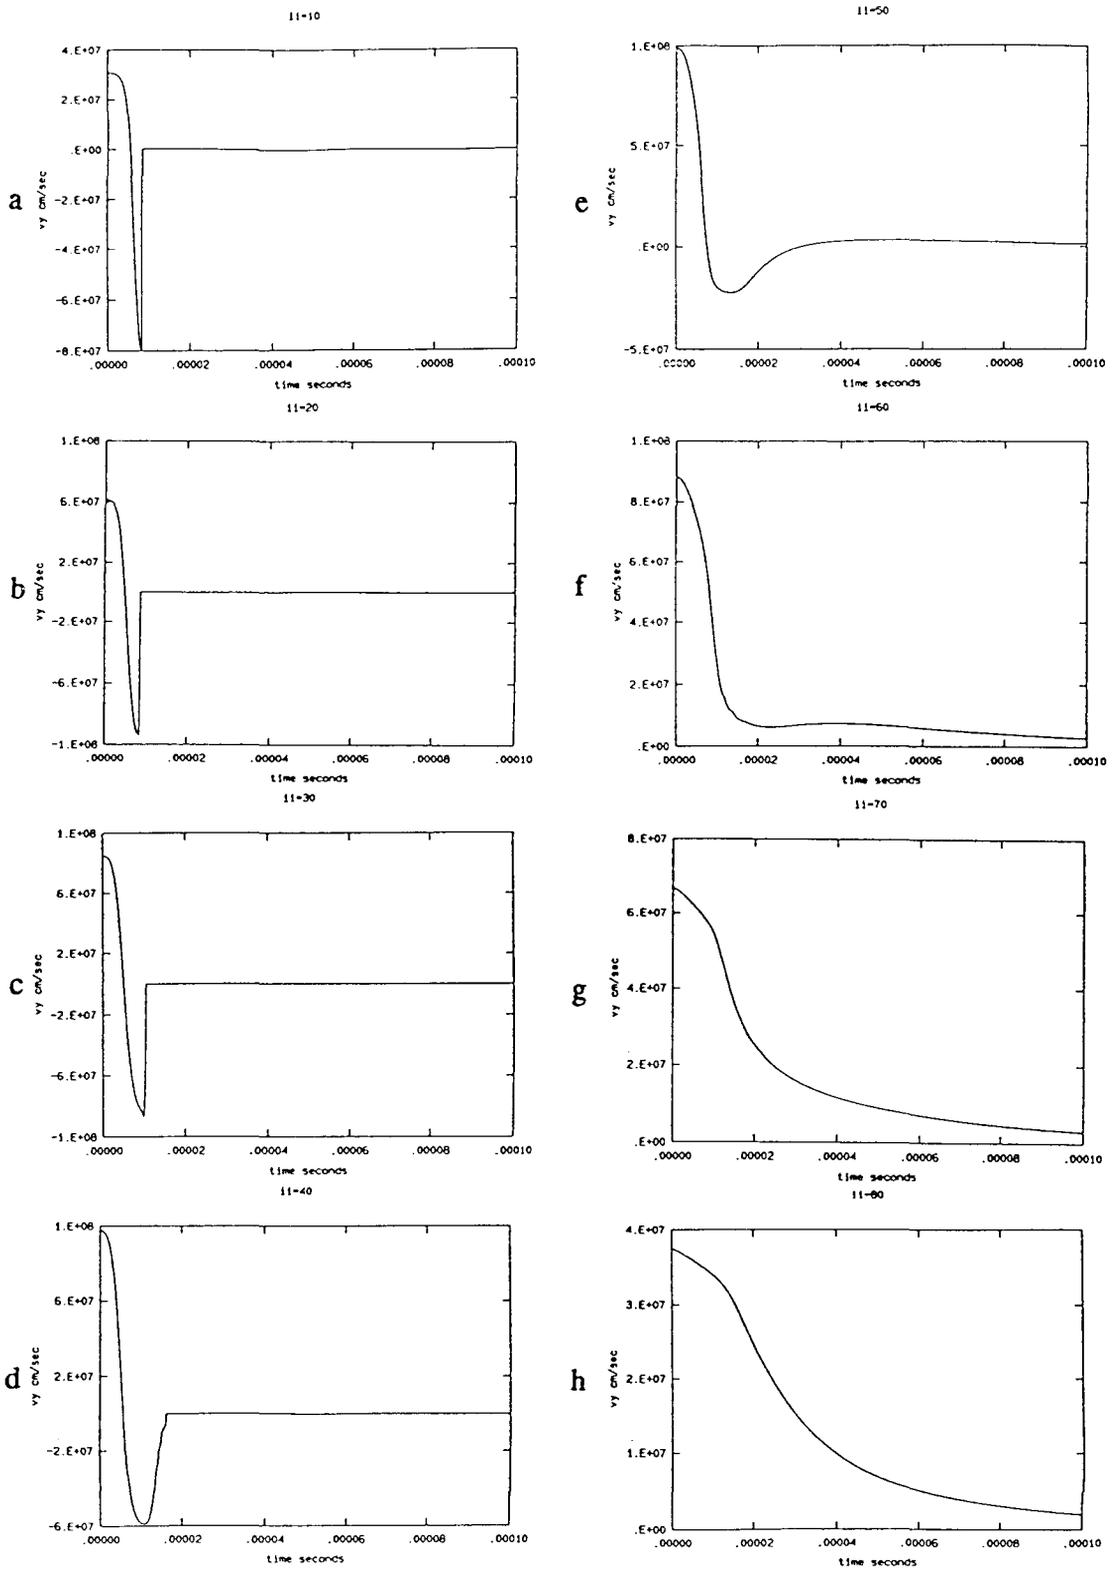


Figure 4.9 Y-component velocity of the shell elements in cm/sec (vertical axis) v.s. time in seconds (horizontal axis).

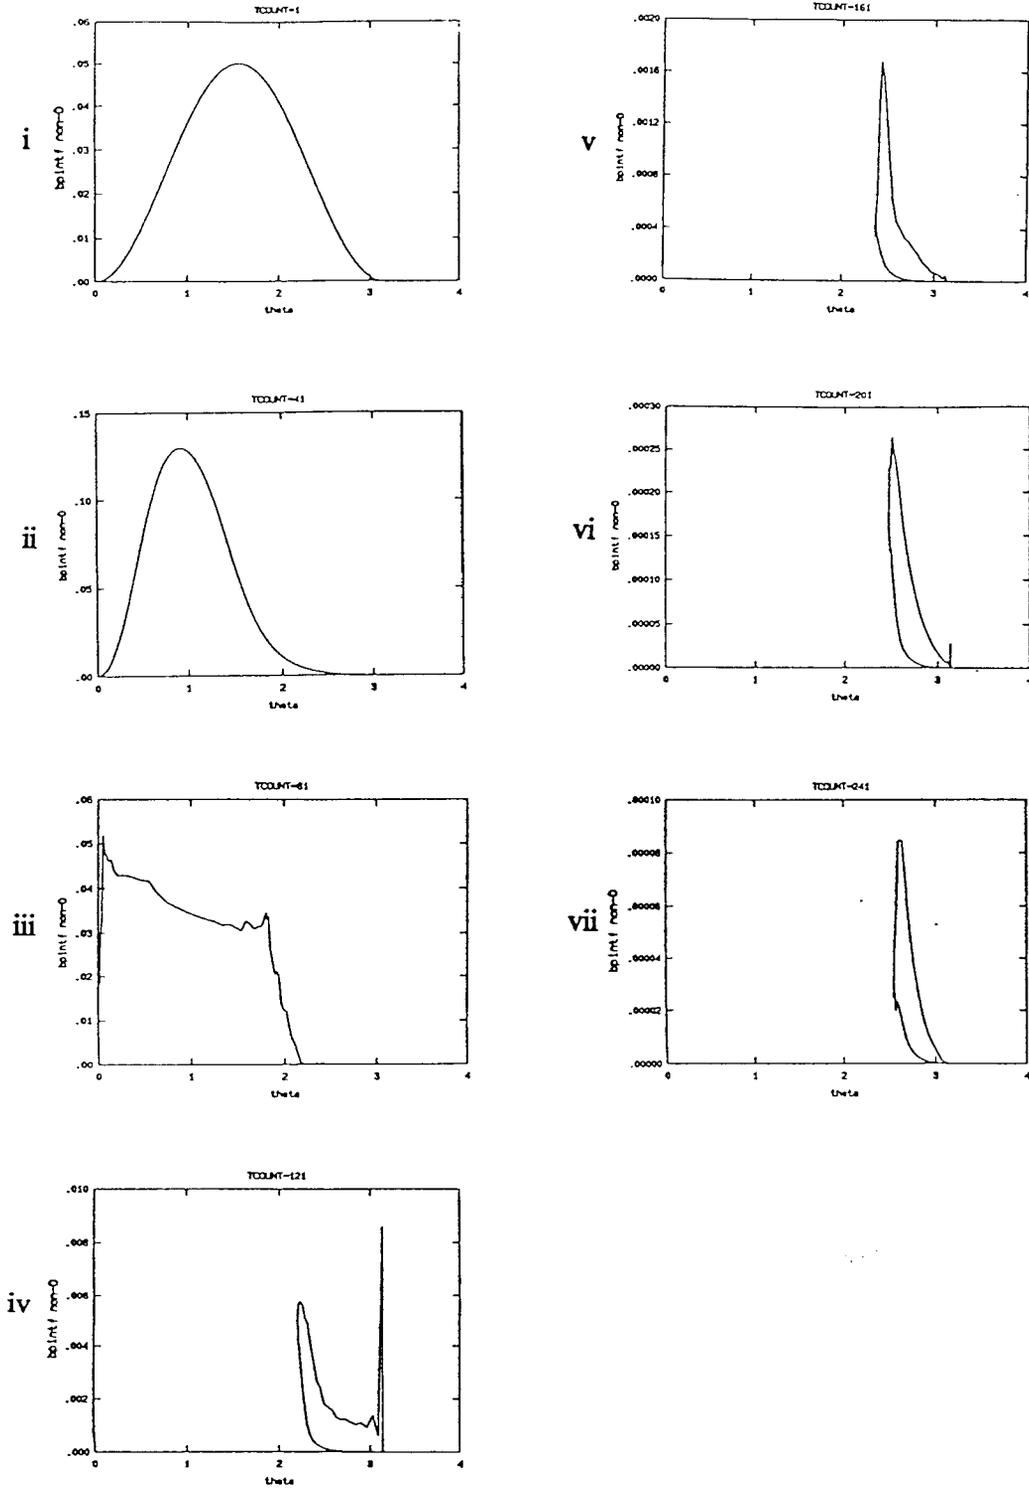


Figure 4.10 Nondimensionalized interfacial magnetic pressure (p_B') (vertical axis) v.s. polar angle (θ) in radians (horizontal axis).

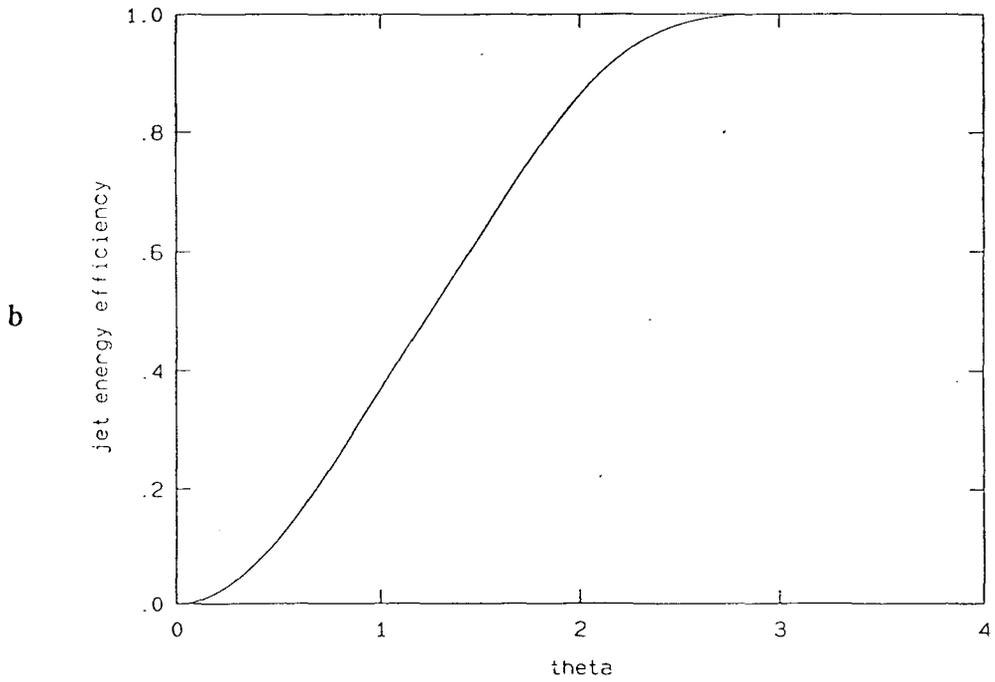
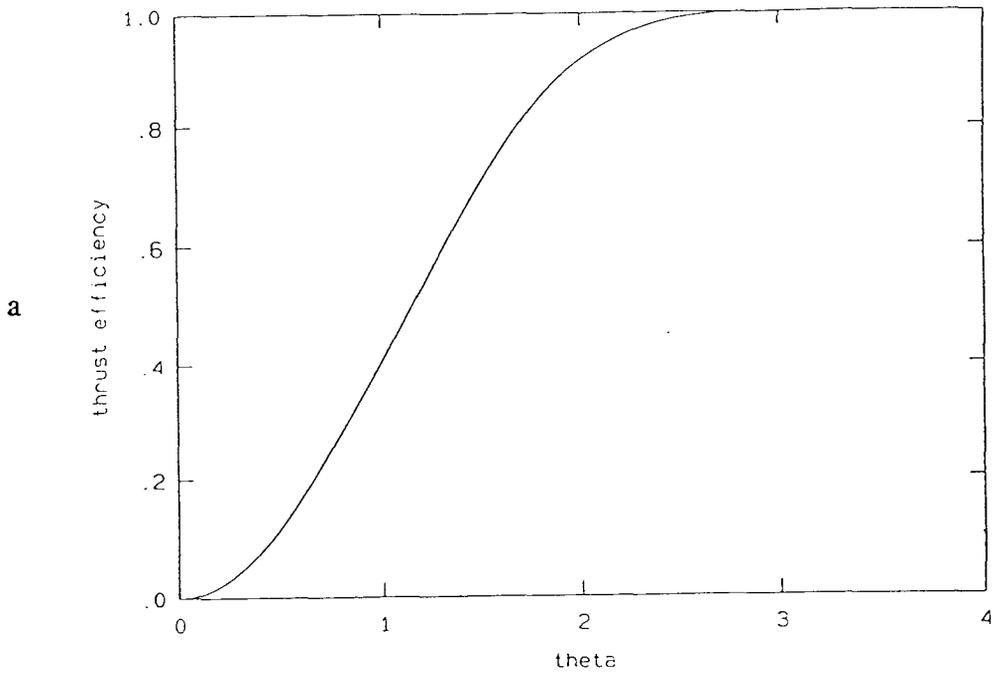


Figure 4.11

IV-A12

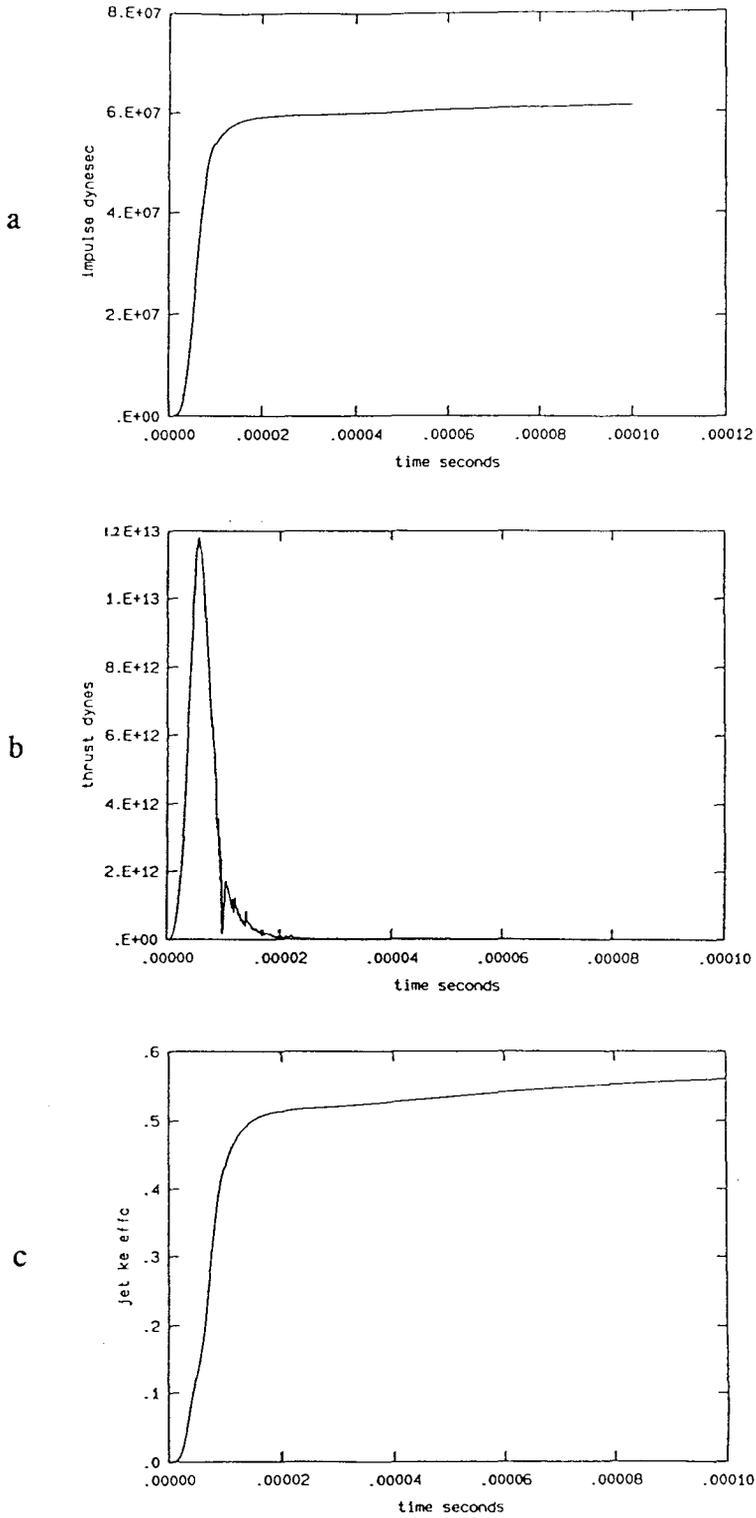
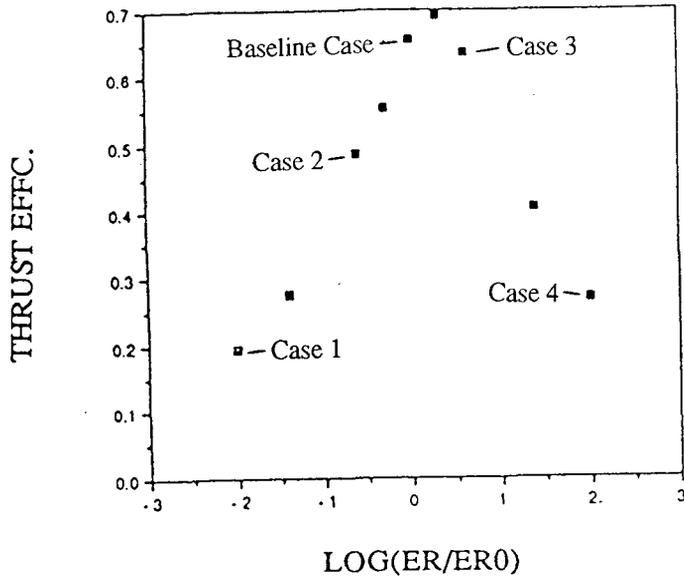


Figure 4.12

a



b

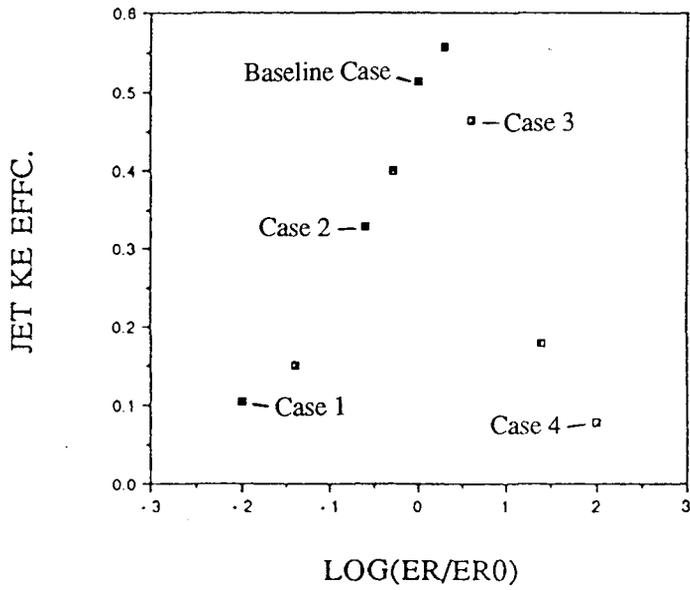
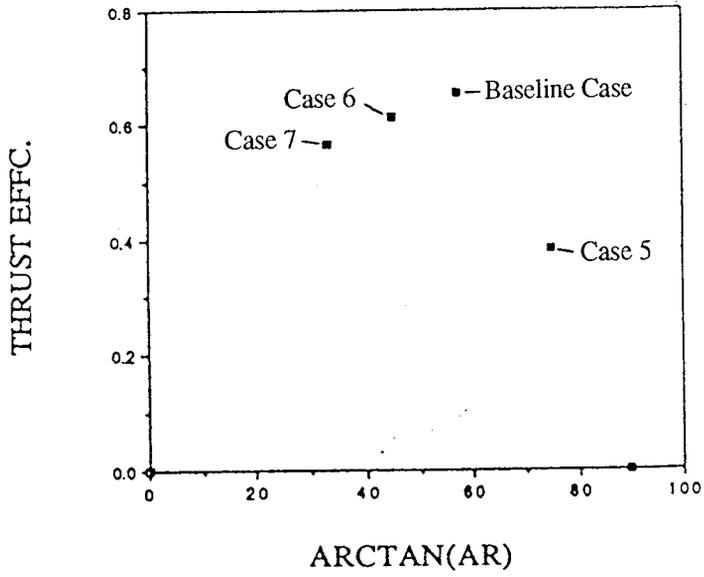


Figure 4.13

a



b

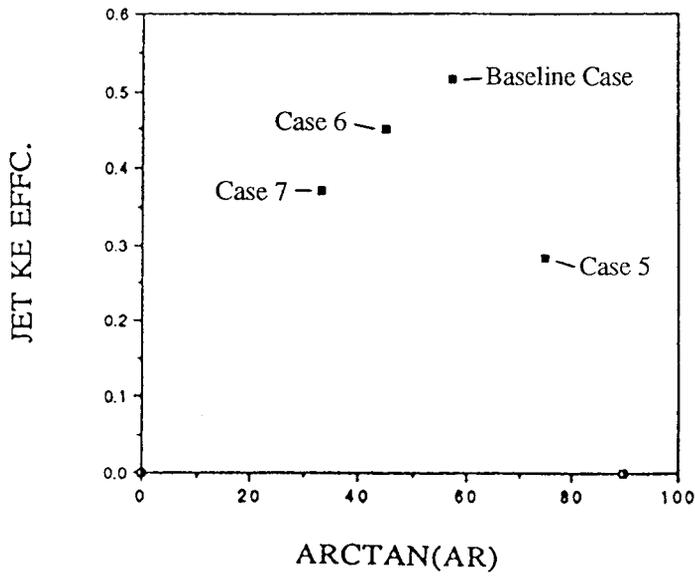


Figure 4.14

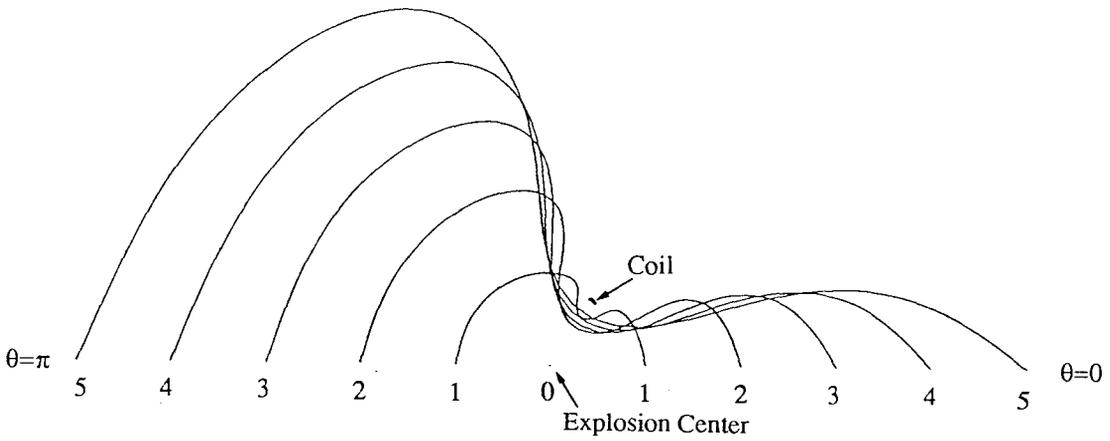


Figure 4.15 Location of shell elements for Case 1 ($ER \approx 0.153$, $AR \approx 1.56$).

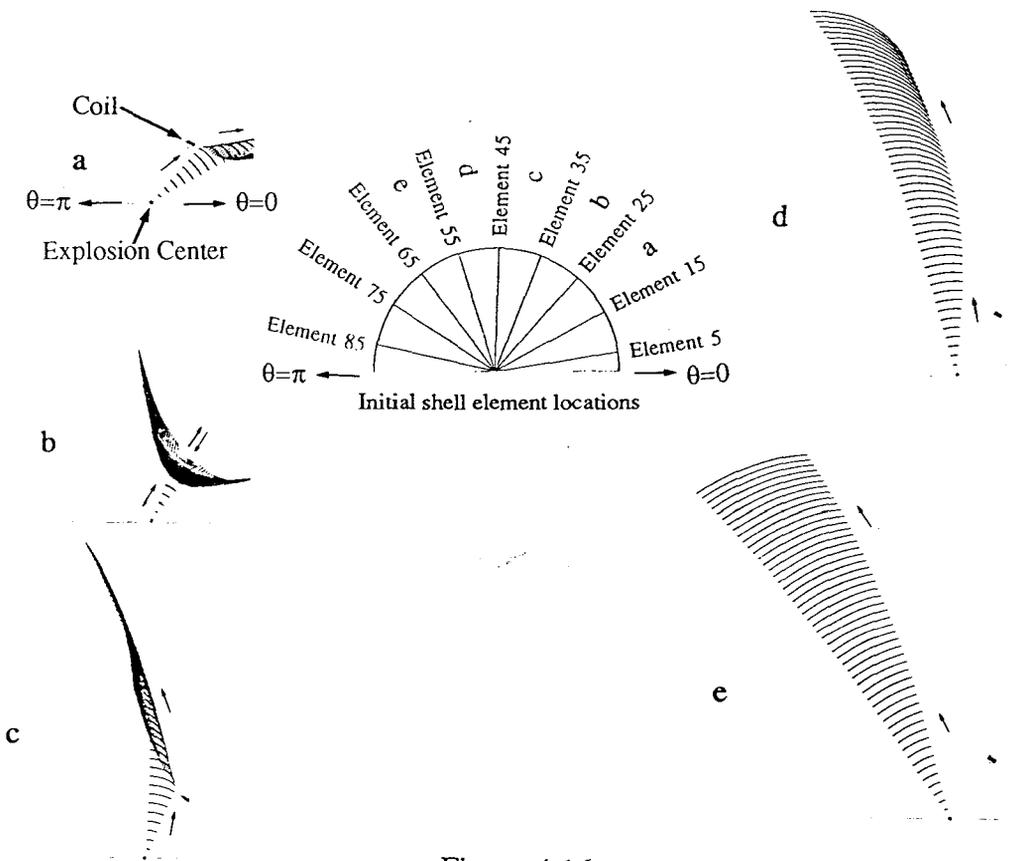
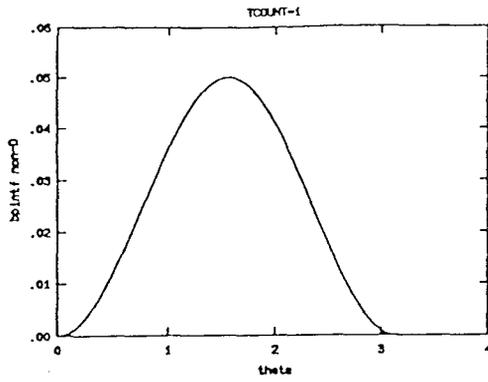
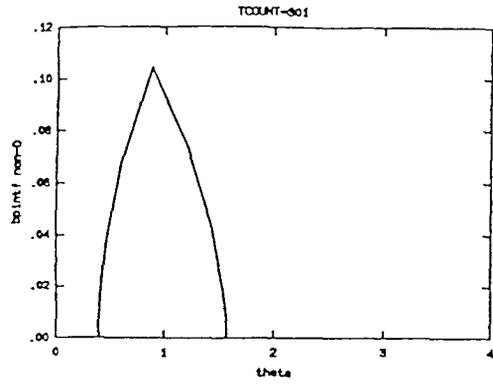


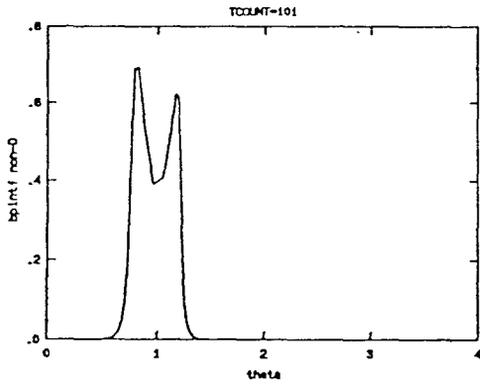
Figure 4.16



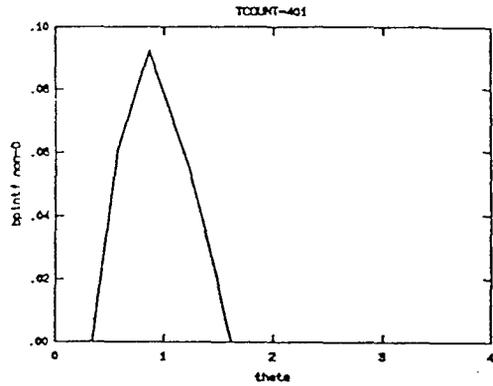
i



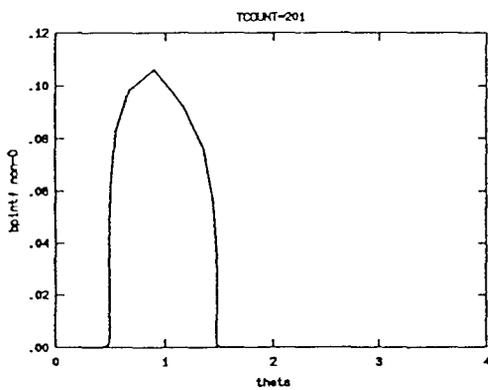
iv



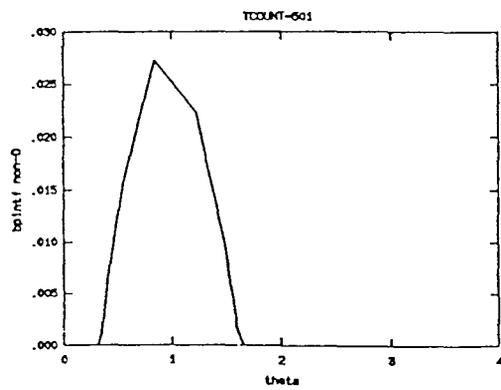
ii



v



iii



vi

Figure 4.17 Nondimensionalized interfacial magnetic pressure (p_B') (vertical axis) v.s. polar angle (θ) in radians (horizontal axis).

IV-A17

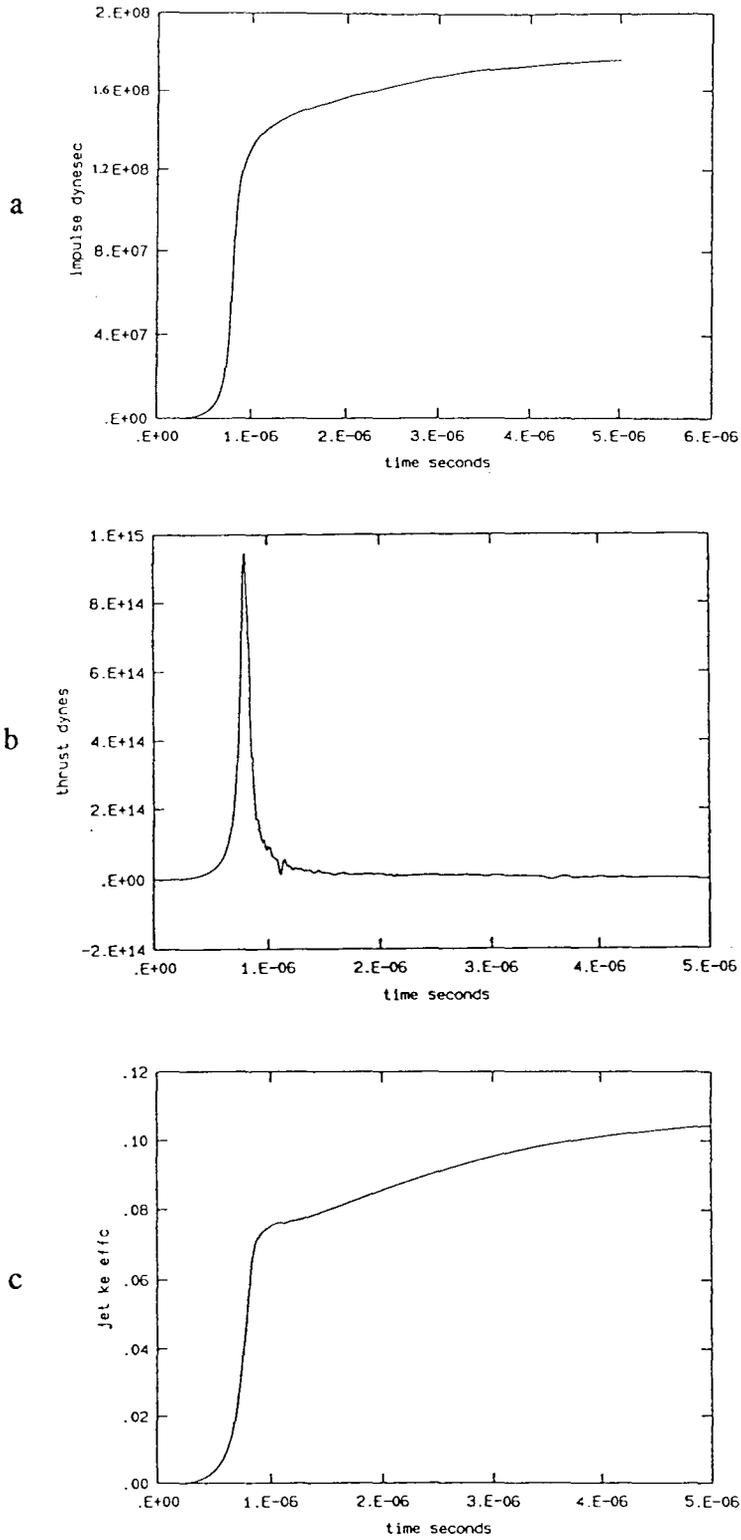


Figure 4.18

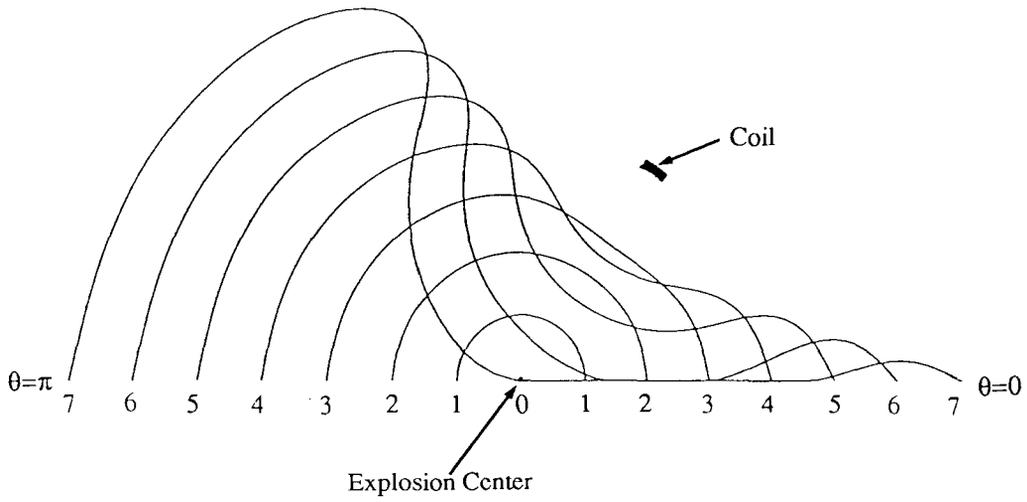
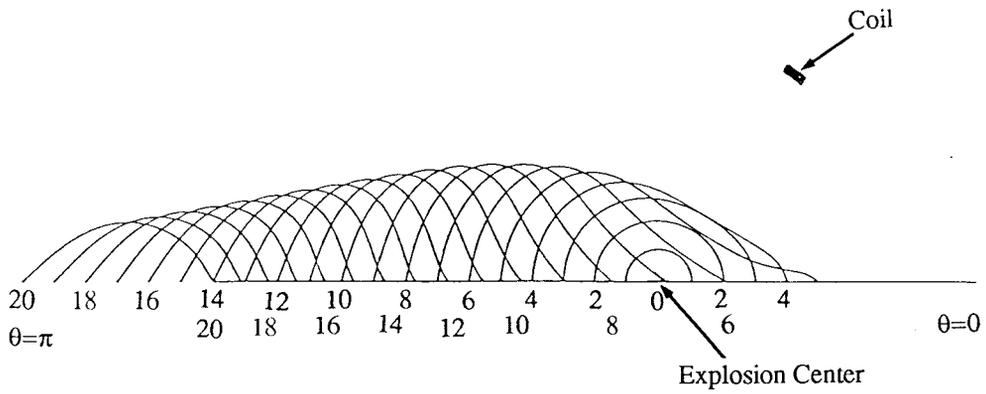


Figure 4.19

Location of shell elements for the Case 2 ($ER \approx 3.83$, $AR \approx 1.56$).



a



b

Figure 4.20 Location of shell elements for the Case 3 (ER=61.2, AR=1.56).

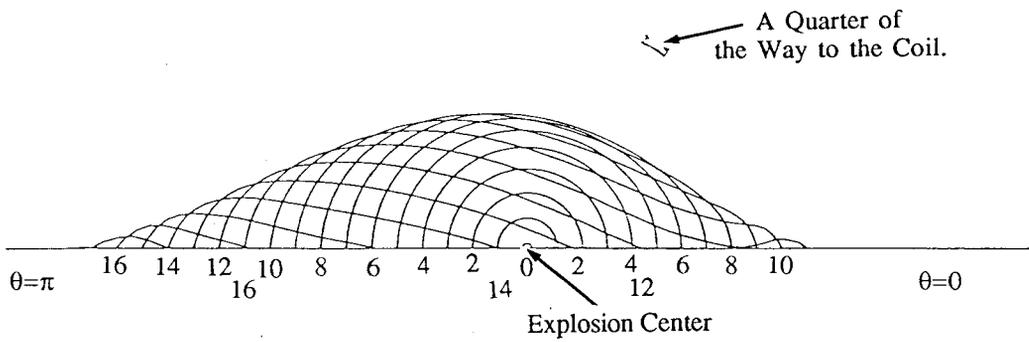


Figure 4.21 Location of shell elements for the Case 4 (ER ≈ 153, AR ≈ 1.56).

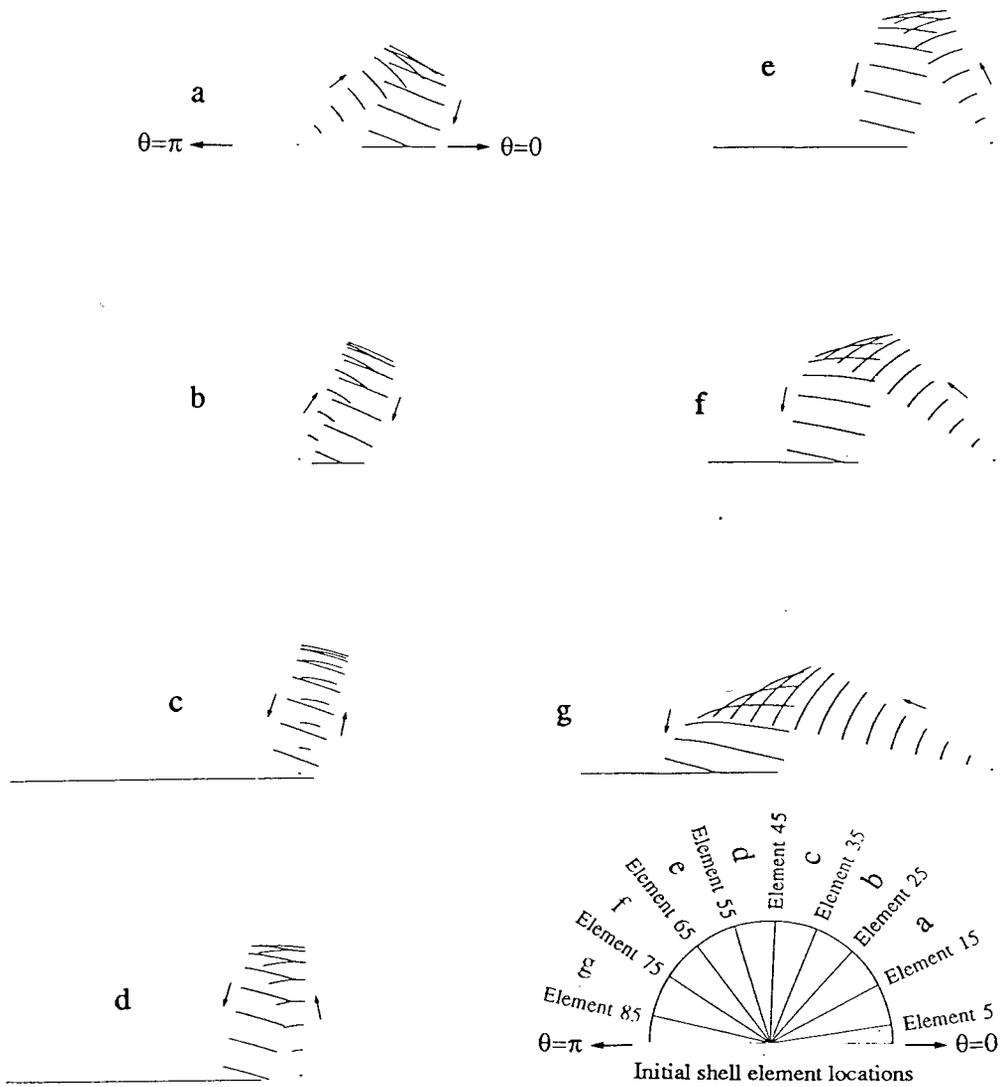


Figure 4.22

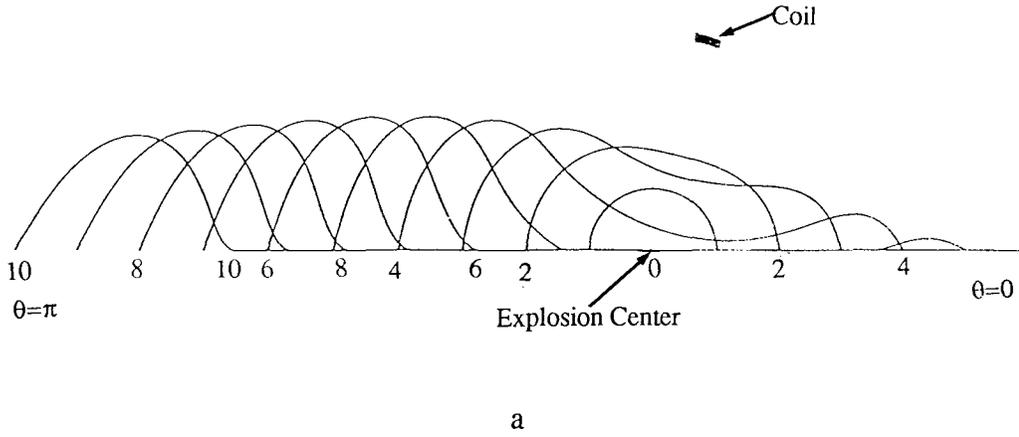


Figure 4.23 Location of shell elements for the Case 5 ($ER \approx 15.3$, $AR \approx 3.73$).

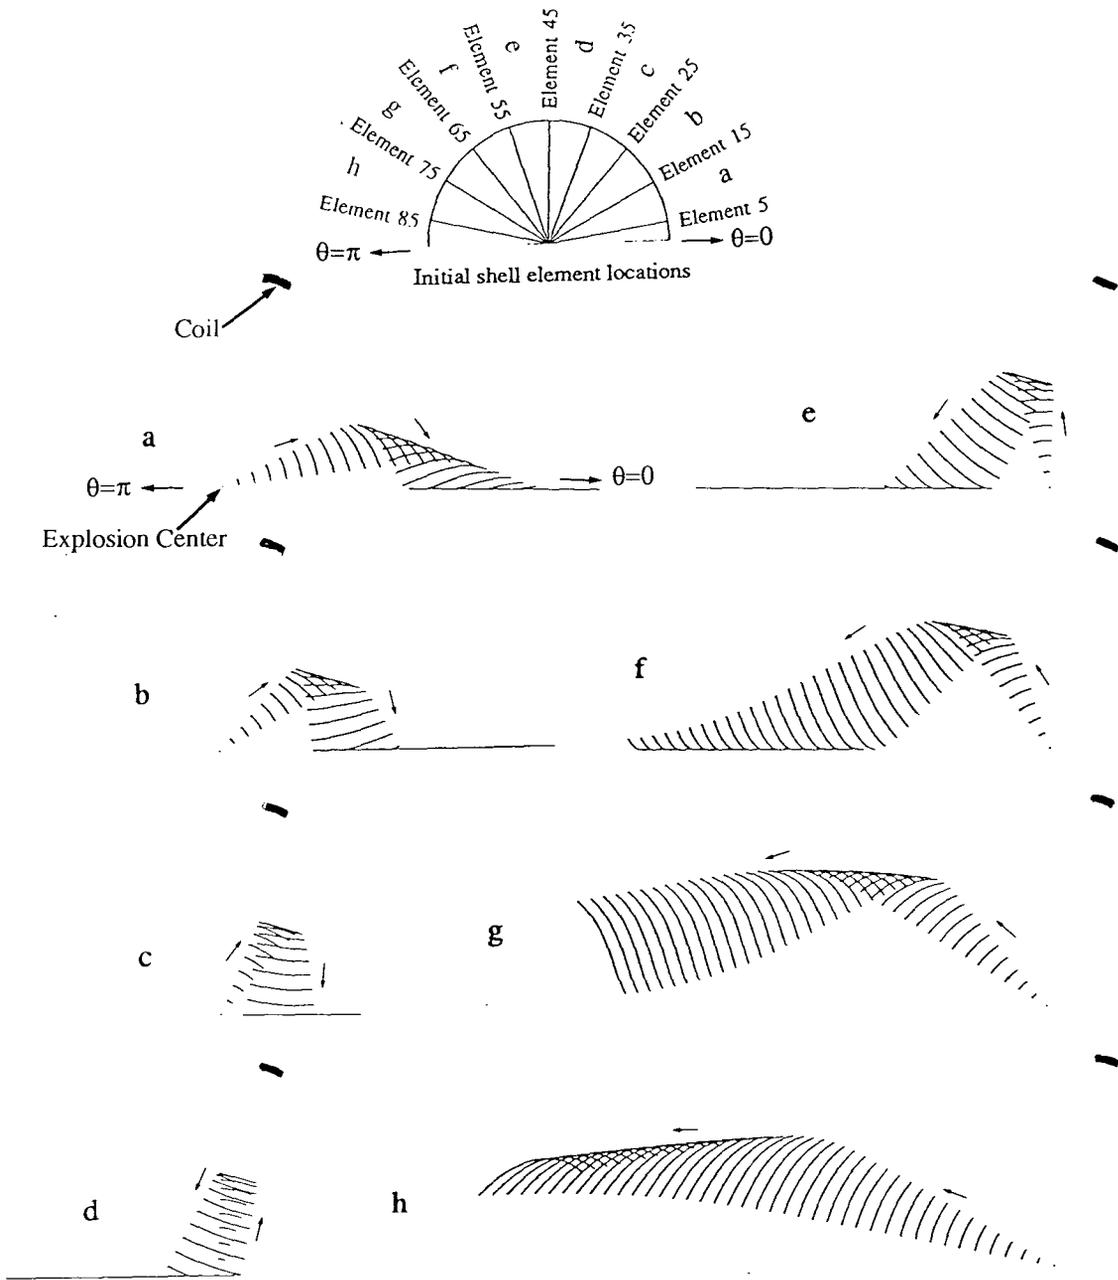


Figure 4.24

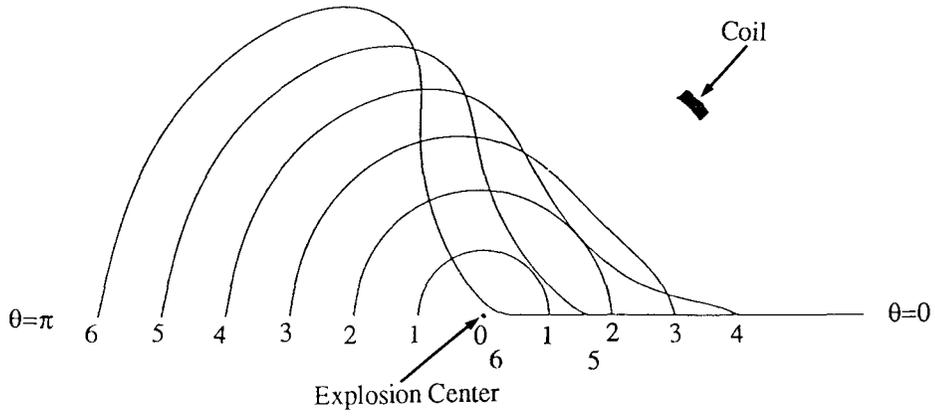


Figure 4.25

Location of shell elements for the Case 6 ($ER \approx 15.3$, $AR = 1$).

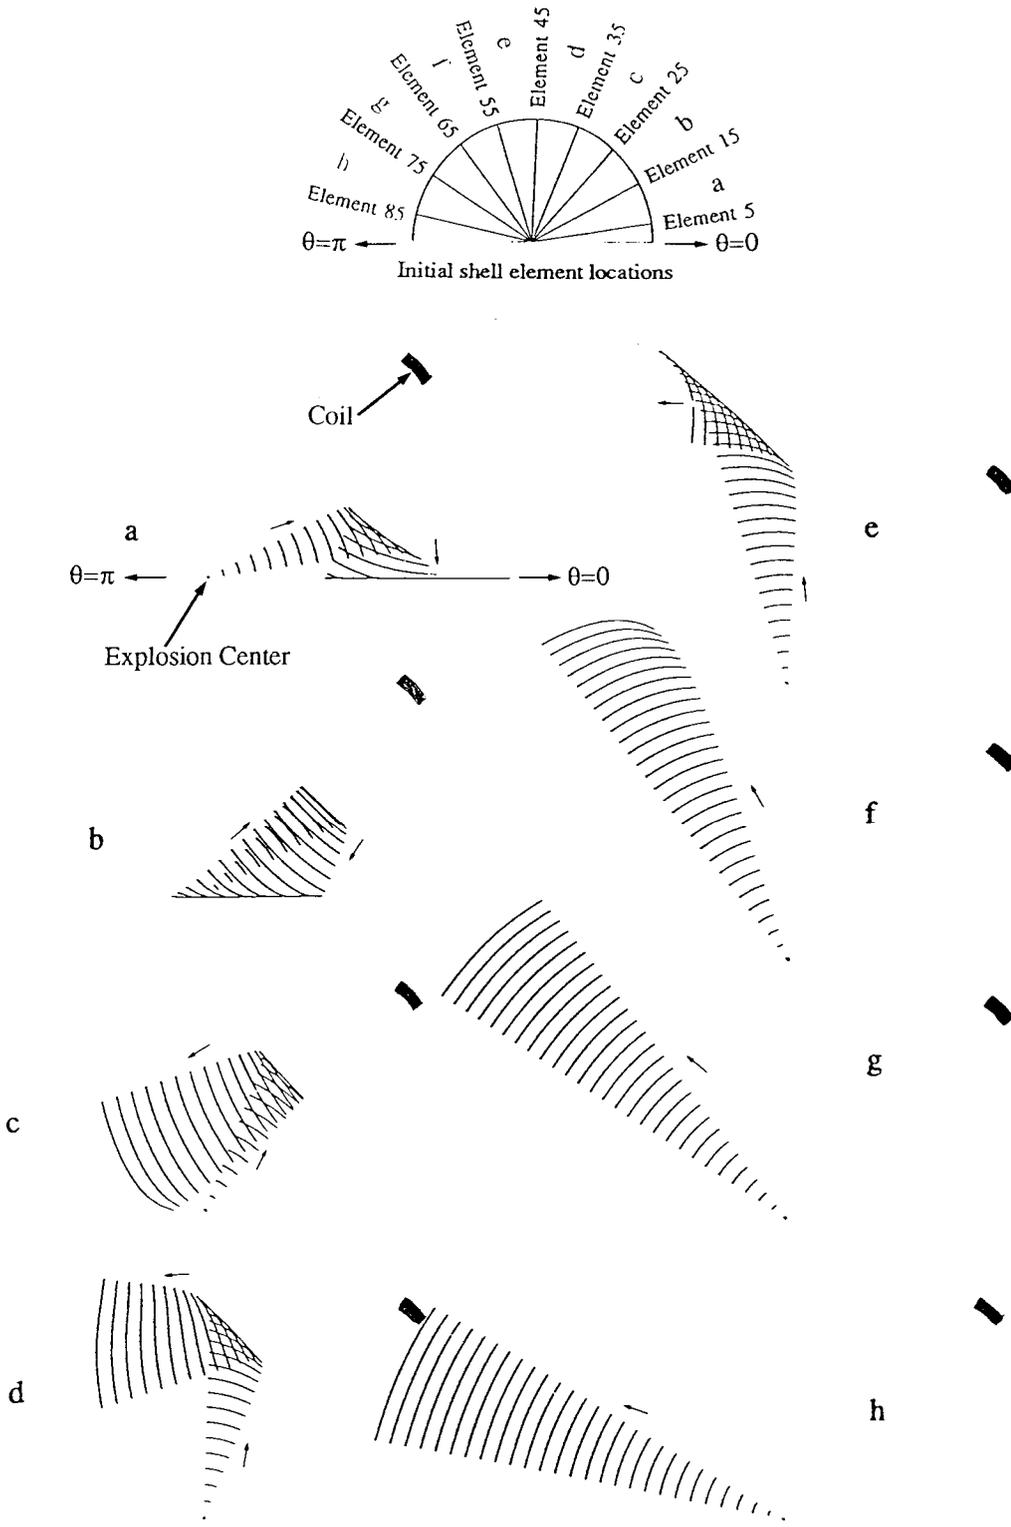


Figure 4.26

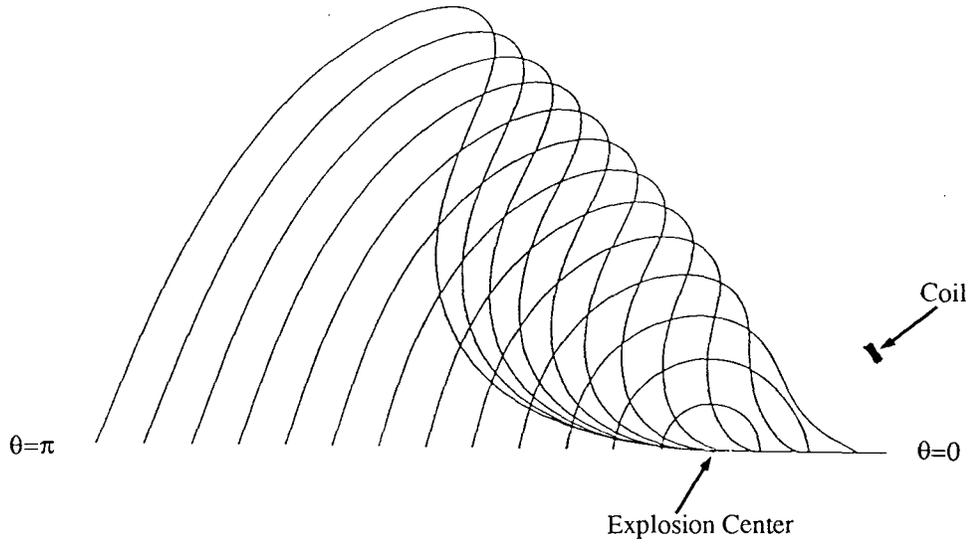


Figure 4.27 Location of shell elements for the Case 7 ($ER \approx 15.3$, $AR \approx 0.649$).

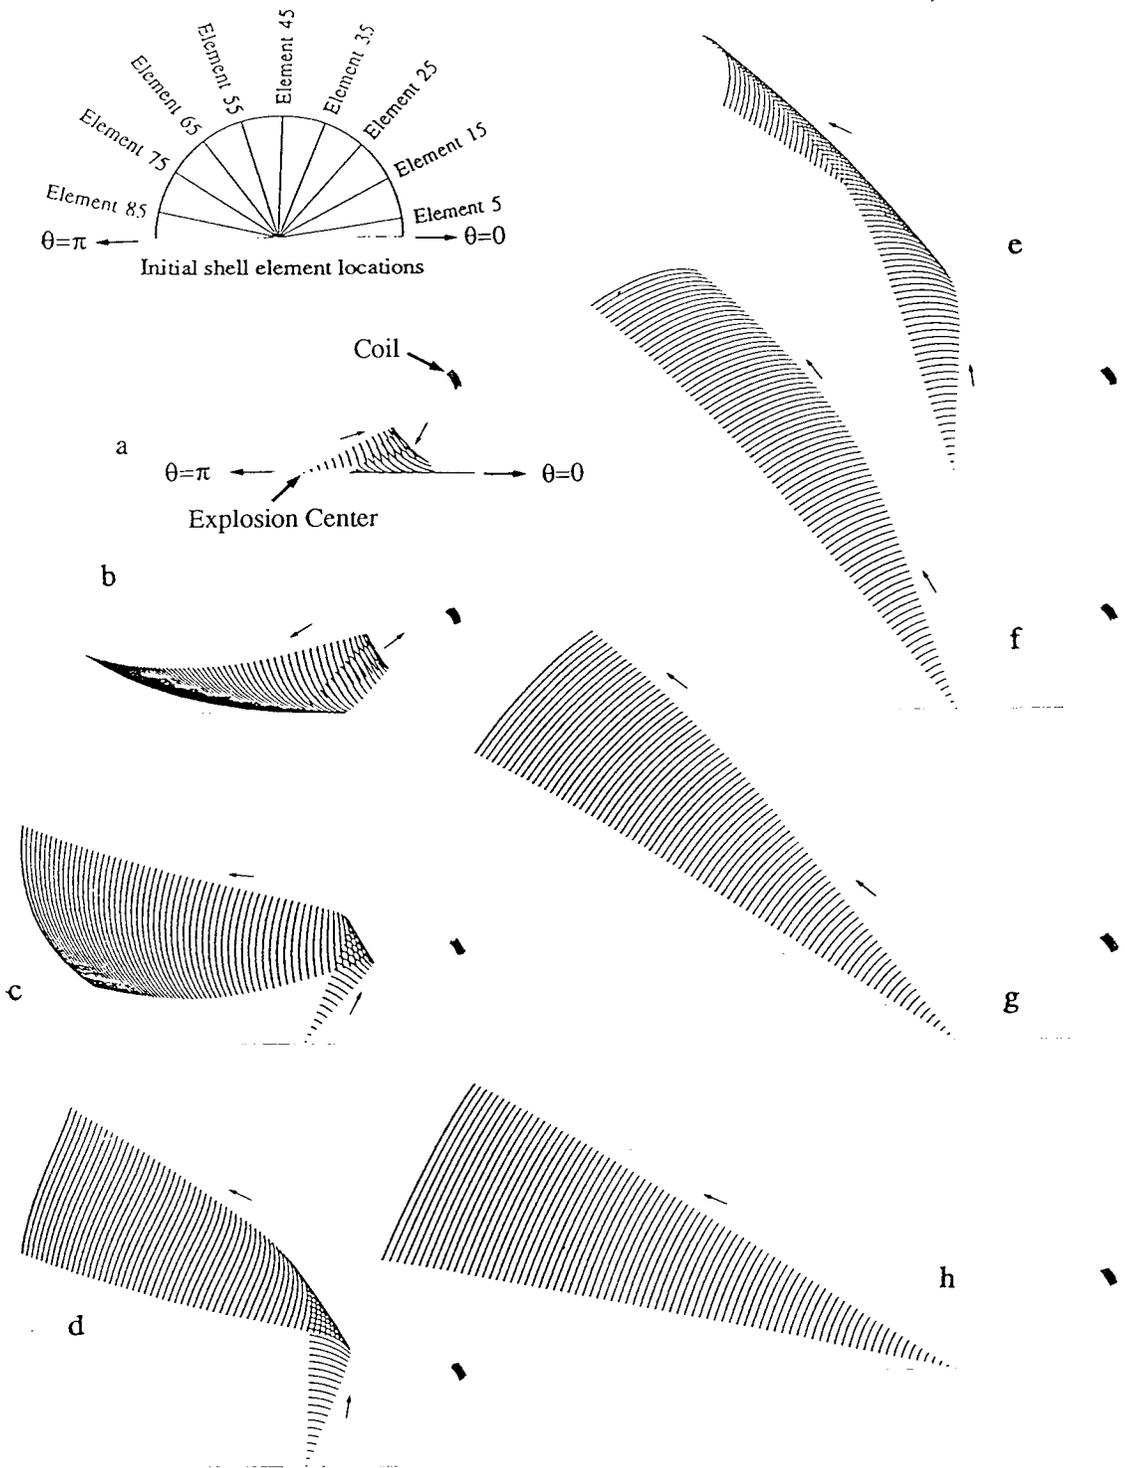


Figure 4.28

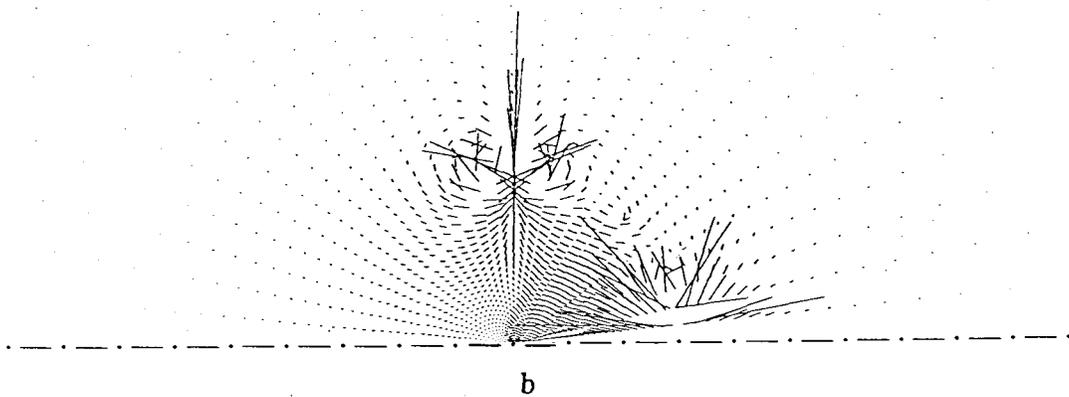
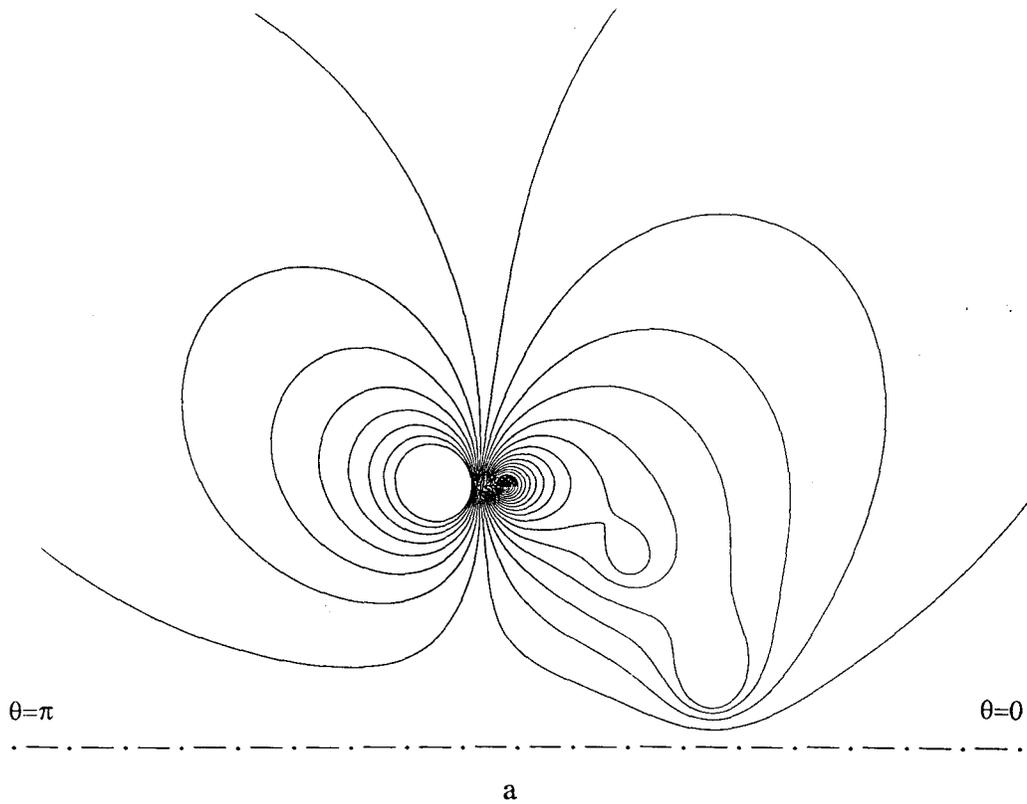


Figure 4.29— Magnetic field contours (a) and magnetic field vectors (b) for the Daedalus first stage thruster.

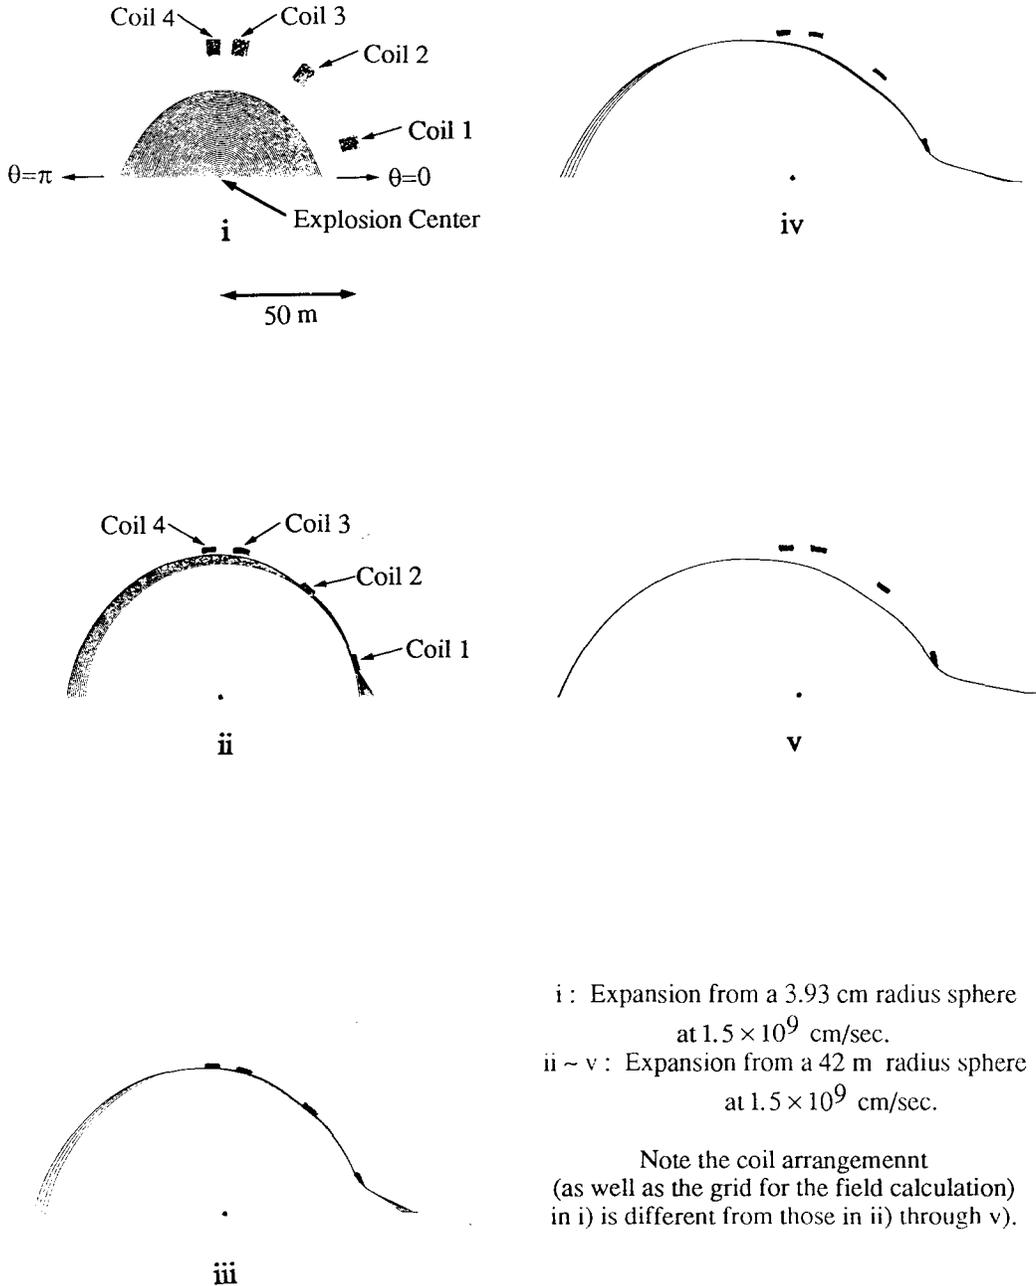


Figure 4.30 Location of shell elements (all same scale) at various stages of expansion in a low field thruster based on the Daedalus design.

Chapter 5

Introduction to PIC

Now, it is evident that the thin-shell approximation is not an entirely realistic nor accurate way to model plasma expansion against a vacuum magnetic field, even if it can be assumed that the plasma will behave as a perfectly conducting fluid obeying the classical laws of hydrodynamics. For instance, it is not true that in a real such expansion, all of the fluid will be collected into a thin shell located precisely at the interface. The thin-shell model does not take into account any internal fluid dynamics or properties of the material assumed to be collected into a shell, and as a consequence, unphysical processes such as reflected shell elements falling onto the central axis can occur.

Therefore, it is essential that a fully hydrodynamical analysis be conducted in order to simulate the expansion flow of plasma more realistically. In this thesis, the Particle-In-Cell (PIC) scheme was selected for the hydrodynamical simulation calculations. One of the primary reasons for this choice was the ability of PIC to track interfaces automatically with unknown motions without the difficulties associated with Lagrangian codes in complex flow situations.

In the following sections, a brief introduction to the PIC scheme will be presented with particular emphasis on the Nishiguchi-Yabe version ^{1,2}, which was chosen for the present work.

5.1 Properties of Eulerian, Lagrangian, and PIC Methods

The PIC method was developed by Harlow and his co-workers at Los Alamos Scientific Laboratories in 1955³ as an alternative to purely Eulerian or Lagrangian schemes, and in particular, to treat phenomena such as multidimensional flows with large distortions or slippages, colliding surfaces, etc., for which existing techniques were not adequate. In recent years, PIC codes have been met with a revival of interest, especially in the domain of plasma physics.^{4,5,6} Although early versions of PIC algorithms suffered from rather high numerical diffusion and noise, more recent developments employing finite-sized particles, higher-order area-weighting methods, or highly Lagrangian approaches have succeeded in reducing these problems.

The PIC technique involves the use of both computational particles that move in a Lagrangian manner and an Eulerian grid. The elements of fluid are represented by computational particles, preferably possessing finite spatial extent in the computational domain, which can be thought of as collections of real fluid particles to which are assigned locally averaged fluid quantities or some distributions thereof. These particles carry quantities associated with the fluid elements across the Eulerian background grid mesh. In the original PIC scheme, the background grid was fixed, but later developments have included methods in which the background grid moves at the local fluid velocity to curtail excessive demands on computer memory by increasing the grid resolution only where necessary. The particle quantities are advanced in time either in an Eulerian/Lagrangian manner (Classical PIC) or in a purely Lagrangian manner (Full-Particle PIC),⁵ and information is communicated back and forth between the computational particles and the underlying grid. Some of the advantages of the PIC scheme over purely Eulerian or Lagrangian schemes will be discussed below.

In conventional Eulerian methods, fluid quantities are assigned to grid points of a coordinate mesh at rest and over which the fluid streams. The Eulerian equations are used

to describe the evolution of grid quantities at these points. Computational particles are not used to represent fluid elements, but since numerical diffusion is eliminated for fluid quantities that are permanently assigned to computational particles which move at the local fluid velocity, Eulerian codes suffer from high numerical diffusion. This numerical diffusion has the characteristic of being high where there are large gradients or where the fluid velocity relative to the grid is high. Recent developments in Eulerian algorithms have included methods such as the Flux-Corrected-Transport (FCT) scheme ^{7,8,9} in which an antidiffusive step is incorporated in an attempt to subtract off numerical diffusion explicitly, but such methods can get rather cumbersome, especially in multidimensional flows. Eulerian codes suffer from (the related problem of) advective term instabilities.^{6,7,10} The advective term in the Eulerian equations, defined here as the term that appears under the divergence when the equations are cast in conservative form, pose problems both in terms of stability of the finite difference equations and numerical diffusion. For example, it is a well-known fact from stability analyses that a finite difference representation of the differential equation

$$\frac{\partial U}{\partial t} = -v_x \frac{\partial U}{\partial x} \quad (5.1)$$

as

$$\frac{U_j^{n+1} - U_j^n}{\Delta t} = -v_x \frac{U_{j+1}^n - U_{j-1}^n}{2\Delta x}, \quad (5.2)$$

where n is the time step index and j is the spatial grid-cell label, is unconditionally numerically unstable.⁶ Many methods, the Two-Step Lax-Wendroff method ¹¹ to name but one, have been developed to alleviate the problems of advective term instabilities and excessive numerical diffusion. Furthermore, Eulerian codes, because they do not use computational particles to represent the fluid, are not apt at handling material interfaces. This means that tracking unknown and moving interfaces may require many iterations per time step, with substantial difficulties in multidimensions. Other shortcomings of the Eulerian method include the possibility of negative densities arising during a simulation and

the difficulty of resolving small features in a flow moving with the fluid through a larger region due to the use of finite-sized grid cells to which fluid quantities are assigned.

However, Eulerian codes are capable of treating complex flow situations and have good conservation properties. Codes written in such a way so as to conserve rigorously physically conserved quantities perform better than those that are not, and thus the equations to be solved should be cast in conservative form. For the case of slab geometry, the conservative form can be presented generally as

$$\frac{\partial U}{\partial t} + \frac{\partial(AU)}{\partial x} = 0 \quad (5.3)$$

under the absence of source or sink terms. The volumetric variable U here represents the conserved quantity. Through the use of an Eulerian grid system, rigorous conservation can be achieved. To see this, we refer to Equation (5.3). The finite difference form of this equation can be written as

$$\frac{U_j^{n+1} - U_j^n}{\Delta t} + \frac{(AU)_{j+\frac{1}{2}}^n - (AU)_{j-\frac{1}{2}}^n}{\Delta x} = 0 . \quad (5.4)$$

If there are a total of N grid cells and we expand Equation (5.4) for $j=1$ through N , sum, and prescribe suitable boundary conditions, the desired conservation property is obtained through cancellation in pairs. For example, if Equation (5.3) were the continuity equation, empty cells would be assigned zero densities and one would obtain

$$\sum_{j=1}^N \rho_j^{n+1} = \sum_{j=1}^N \rho_j^n , \quad (5.5)$$

which represents conservation of mass between time steps.

Boundary conditions come in a variety of forms depending upon the type of boundary: vacuum boundary, applied pressure boundary, reflective wall, boundary across which fluid flows, etc. The pressure on a boundary between a fluid filled cell and a force-free vacuum cell should be set to zero, and the pressure on a boundary at which an external pressure is applied should be set equal to the applied pressure. The velocity assigned to an exterior cell bordering a vacuum or applied pressure boundary should generally be chosen equal to that assigned to the adjacent fluid cell in which the calculation is being performed at

the time. At reflective walls, the fluid velocity normal to the wall must be set to zero, and the pressure gradient across the wall should be made to vanish. The latter can be accomplished by letting the fictitious cell just outside the wall have exactly the same pressure as the mirror cell just inside the wall. Boundaries representing symmetry centers act effectively as reflective walls. For boundaries across which fluid flows (such as boundaries artificially imposed to restrict a computational domain), the pressure and velocity assigned to fictitious cells just outside the boundary should generally be chosen equal to those assigned to the adjacent fluid cells bordering the boundary (although if the flow out of such a boundary is supersonic, the nature of the boundary condition is obviously not crucial).

In the Lagrangian method, the fluid is subdivided into a large number of finite zones that characterize the fluid elements and follow the local fluid motion. Each of these computational particles obey the Lagrangian equations. However, no underlying computational grid is used, and it is more difficult to conserve rigorously physical quantities that are conserved in reality. It is also difficult to handle complex flow situations in multidimensions involving large distortions and shear or colliding interfaces because of such problems as mesh entanglements. Also, Lagrangian codes tend to be more memory intensive than Eulerian codes. However, because physical interfaces automatically appear as interfaces to collections of computational particles in Lagrangian codes, capturing of interfaces becomes possible without having to perform iterations to look for them. The term "capture" is being used here in the same context as in "shock capturing," which will be discussed later in this chapter.

PIC codes circumvent or alleviate many of the aforementioned problems inherent in either Eulerian or Lagrangian codes by combining the traits of Eulerian and Lagrangian methods. Firstly, the employment of computational particles that undergo Lagrangian motion results in the disappearance of advective terms from the finite difference equations. PIC codes are therefore spared the advective term instabilities and excessive numerical

diffusions that plague pure Eulerian codes. The use of Lagrangian computational particles also allows unknown moving interfaces to be automatically tracked and contact surfaces to be clearly delineated. On the other hand, the use of an Eulerian grid facilitates the calculation of gradients and forces and also makes possible the rigorous conservation of physically conserved quantities. And unlike the case for pure Lagrangian codes, complex multidimensional flows with high distortion and shear can be treated without difficulty.

But PIC does have problems of its own. The biggest shortcomings are those of numerical noise and numerical diffusion. However, as mentioned earlier, both problems can be mitigated through the use of finite-sized computational particles (FSPs) and area weighting. The FSPs are generally chosen to have the same size as those of the grid cells over which they lie* and allow efficient coarse graining of particle quantities and suppress short-wavelength density variations to provide fluid distributions that more closely approximate the smooth original distribution. The coarse graining prevents spurious physics from creeping in when going from continuous space to discrete space. Usage of FSPs and interpolations permits the number of particles required per cell to be quite low. The fact that the finite size of computational particles allows fractional portions of a particle to overlap a grid cell implies that larger density variations may be handled than when only point particles are used.

5.2 The PIC Scheme

5.2.1 Manipulation of Particle Quantities in PIC

The way in which a PIC scheme works is best illustrated by giving examples.

* In the case of a nonuniform grid, the FSPs should be given the same size as the smallest cell over which they lie.

Suppose we have a swarm of FSPs whose center coordinates are given by the position vectors \underline{x}_i . The shape factor, which is a function describing the normalized spatial mass distribution within each FSP, will be denoted by $S(\underline{x} - \underline{x}_i)$, with

$$\int S(\underline{x} - \underline{x}_i) d\underline{x} = 1 . \quad (5.6)$$

Then, the number density at a particular position \underline{x} will be given by

$$n(\underline{x}) = \sum_{i=1}^N S(\underline{x} - \underline{x}_i) , \quad (5.7)$$

and the average velocity by

$$\langle \underline{v} \rangle = \frac{\sum_{i=1}^N \underline{v}_i S(\underline{x} - \underline{x}_i)}{\sum_{i=1}^N S(\underline{x} - \underline{x}_i)} . \quad (5.8)$$

The charge density at the point \underline{x} may be written as

$$\sigma_f(\underline{x}) = \int \rho_p(\underline{x}') S(\underline{x} - \underline{x}') d\underline{x}' \quad (5.9)$$

for a continuous distribution. In an approximation using a finite number of point particles, this becomes

$$\sigma_p(\underline{x}) = \sum_i q_i \delta(\underline{x} - \underline{x}_i) , \quad (5.10)$$

and when particles with finite extent are used instead, one gets

$$\sigma_f(\underline{x}) = \sum_i q_i S(\underline{x} - \underline{x}_i) . \quad (5.11)$$

Square functions and Gaussians are commonly chosen shape factors. By taking \underline{x} to be a grid coordinate, a relation between grid and particle quantities is obtained. As an aside, comparison of Equations (5.10) and (5.11) points to the possibility of calculating derivatives, and thus force terms, via Fourier Transforms (except when dealing with the magnetic field), using quantities accumulated at grid mesh points, instead of relying on the slower finite difference methods.^{4,6}

Other ways of relating grid and particle quantities also exist. For example, in the more primitive Nearest-Grid-Point (NGP) method,^{5,6} the centers of the original grid cells are connected so as to form a parallel secondary mesh of cells, and all particles that fall into

one of these secondary cells are associated with the original grid point located at the center of that cell. Effectively, all particles contained in a particular cell of the secondary grid are considered to be positioned at the center of that cell. There is no spatial averaging with this method, and one can write

$$n(\underline{x}) = \sum_{\mathbf{g}} S(\underline{x} - \underline{x}_{\mathbf{g}}) \rho_{\text{NGP}}(\mathbf{g}) , \quad (5.12)$$

where

$$\rho_{\text{NGP}}(\mathbf{g}) = \sum_{i \in \mathbf{g}} 1 . \quad (5.13)$$

The summation in Equation (5.13) is over all particles closest to the grid point \mathbf{g} of the original grid. This method, although plagued by rather high noise and diffusion due to discontinuous transport across cell boundaries (there is no area weighting), does not suffer from zeroth-order numerical diffusion,⁵ i.e., the changing of the grid solution from one time step to the next even when the particles do not move.

5.2.2 PIC Algorithms

5.2.2.1 Classical PIC Methods

PIC algorithms can be categorized into two major classes : Classical PIC and Full-Particle PIC.⁵ In Classical PIC, fluid quantities (specific cellwise-averaged quantities) are distributed on an Eulerian grid mesh to characterize the fluid variables. The particles which represent the fluid elements describe only temporarily the fluid to model the convection process; namely, permanent memory carried by a particle is limited to its mass and position and a particle carries area weighted momentum and energy information only during motion. Momentum and energy experience numerical diffusion because of the transfer of information back and forth between the grid cells and particles. Mass, being permanently assigned to the computational particles (i.e., it is a Lagrangian variable), experiences no

numerical diffusion. The two phases of calculation involved at each time step of a Classical PIC code can be summarized as follows :

first phase - finite differencing of the Eulerian equations less the advective term to obtain temporary cellwise fluid quantities on the grid,

second phase - Lagrangian motion of the computational particles to take care effectively of the material transport effects neglected in the first phase (i.e., fluid motion is realized).

The hydrodynamic Eulerian conservation equations for mass, momentum, and total energy (sum of the kinetic and internal energies), less the viscosity and heat conduction terms as well as any electromagnetic terms, are given, respectively, in conservative form, as

$$\frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \underline{v}) = 0 \quad (5.14)$$

$$\frac{\partial (\rho \underline{v})}{\partial t} + \nabla \cdot (\rho \underline{v} \underline{v}) = -\nabla p \quad (5.15)$$

$$\frac{\partial (\rho \epsilon)}{\partial t} + \nabla \cdot (\rho \epsilon \underline{v}) = -\nabla \cdot (p \underline{v}) . \quad (5.16)$$

Because the conservation of mass is automatic through the use of computational particles to which masses are assigned permanently, only the momentum and energy equations need be considered. The first step of a Classical PIC scheme can therefore be represented by the following Eulerian finite difference equations written in two-dimensional Cartesian coordinates :

$$\tilde{u}_{i,j} = u_{i,j}^n - \frac{\Delta t}{\rho_{i,j}^n \Delta x} \left(p_{i+\frac{1}{2},j}^n - p_{i-\frac{1}{2},j}^n \right) \quad (5.17)$$

$$\tilde{v}_{i,j} = v_{i,j}^n - \frac{\Delta t}{\rho_{i,j}^n \Delta y} \left(p_{i,j+\frac{1}{2}}^n - p_{i,j-\frac{1}{2}}^n \right) \quad (5.18)$$

$$\tilde{\epsilon}_{i,j} = \epsilon_{i,j}^n - \frac{\Delta t}{\rho_{i,j}^n} \left\{ \frac{\left[(p^n \bar{u})_{i+\frac{1}{2},j} - (p^n \bar{u})_{i-\frac{1}{2},j} \right]}{\Delta x} + \frac{\left[(p^n \bar{v})_{i,j+\frac{1}{2}} - (p^n \bar{v})_{i,j-\frac{1}{2}} \right]}{\Delta y} \right\} , \quad (5.19)$$

where here, u and v are, respectively, the x and y components of the Eulerian fluid velocity \underline{v} , and ϵ is the specific total energy

$$\varepsilon = \frac{1}{\gamma - 1} \frac{p}{\rho} + \frac{|\underline{v}|^2}{2} . \quad (5.20)$$

The barred velocities are defined through

$$\bar{\underline{v}} = \frac{\underline{v}^n + \tilde{\underline{v}}}{2} \quad (5.21)$$

for improved stability of the calculation and rigorous conservation of energy.³ The tilde ($\tilde{\sim}$) quantities for the grid cells represent physical variables at the end of the first step (the Eulerian finite differencing step) of the PIC calculation. They are only temporary solutions because the advective terms have been neglected in these finite difference equations.

Next, the fluid quantities are assigned to the FSPs via area weighting of the tilde grid-cell values. The FSPs are given a fraction of the total cell quantity in proportion to their area overlap over the cell. The particles are then moved using area weighted tilde velocities (equivalent to linear interpolation among neighboring grid-point values). After the particles are moved, new values are assigned to the grid-cell fluid variables by repartitioning the physical quantities from the FSPs back onto the grid via area weighting. In short, if an FSP moves, one adds or subtracts from the relevant cells proportionate shares of the physical quantity, consistent with the fractional amount of the particle moved into the new cell. The fluid quantities associated with Eulerian grid points are specific quantities, but the summation over all fractional FSP areas contained in a cell in the postmotion picture yields cellwise totals. One converts the cellwise totals to cellwise-averaged specific quantities via dividing through by the total FSP area or mass overlapping the cell.

5.2.2.2 Full-Particle PIC Methods

The other general class of PIC algorithms is the Full-Particle PIC.¹² The Full-Particle PIC scheme is a modification of Classical PIC to represent fully a Lagrangian fluid. It was developed in an effort to reduce the problems of numerical diffusion and noise,

although advanced Classical PIC schemes have also been developed to reduce these problems without resorting to a fully Lagrangian approach. In Full-Particle PIC, the computational particles carry all fluid quantities (momentum and energy as well as mass) on a permanent basis, so not only the diffusion of mass, but also the major source of diffusion of momentum and energy are eliminated. The particle data describe the fluid completely from cycle to cycle, not just during convection. The fluid quantities assigned to mesh points are temporarily accumulated averages from particle quantities, and the grid is only for computational convenience. Because of the highly Lagrangian approach, Full-Particle PIC is more memory intensive than Classical PIC. It is also more difficult to conserve both total momentum and energy rigorously with the Full-Particle scheme. Furthermore, because of the collisionless nature of FSPs, multistreaming and multitemperature problems can occur, although this problem can be remedied. In Classical PIC, this problem was avoided by attaching momentum and energy to the particles only during transport.

5.2.2.3 Magnetic Field Calculations

Now, even with a Full-Particle code, the magnetic fields will be defined only on a background Eulerian grid and will be advanced in time relying entirely on finite differencing of the equations. The magnetic fields present in a plasma can be taken into account by incorporating the magnetic stress tensor into the conservation equations. That is, the fluid pressure appearing in the equations should be replaced by the sum of the scalar fluid (i.e., thermal) pressure and the magnetic stress tensor, whose i,k th component may be written out as ¹³

$$P_{ik} = \frac{1}{4\pi} \left(B_i B_k - \frac{1}{2} B^2 \delta_{ik} \right). \quad (5.22)$$

In perfectly conducting MHD, Faraday's Law and Ohm's Law combine to give

$$\frac{\partial \underline{\mathbf{B}}}{\partial t} = \nabla \times (\underline{\mathbf{v}} \times \underline{\mathbf{B}}). \quad (5.23)$$

The particle velocity and magnetic field can be advanced in time using the momentum equation together with Equation (5.23).¹⁴ Total flux conservation can be achieved by properly summing Equation (5.23) over the grid.

5.2.3 SPH Methods

At this point, a relative of PIC called Smoothed Particle Hydrodynamics (SPH) should be introduced.¹⁵ This method is similar to PIC except that it is gridless. It involves the motion of a set of points to which masses are assigned and whose velocities and thermal energies are known at any time. The computational particles possess a spatial extent defined through the smoothing kernel, and this allows a better approximation to the actual continuous density distribution to be obtained than when using only a discrete set of point particles. The forces on the fluid elements needed for moving the particles are calculated from information carried by the particles. Therefore particles must be present wherever field calculations are to be carried out. For example, the field

$$\underline{A}(\underline{x}) = \int \underline{A}(\underline{x}') \delta(\underline{x} - \underline{x}') d\underline{x}' \quad (5.24)$$

will be approximated by

$$\langle \underline{A}(\underline{x}) \rangle = \sum_{k=1}^N m_k \frac{\underline{A}_k}{\rho_k} w(\underline{x} - \underline{x}_k, h) , \quad (5.25)$$

where N is the total number of particles, $w(\underline{x} - \underline{x}_k, h)$ is the interpolating kernel, which satisfies

$$\int w(\underline{v}, h) d\underline{v} = 1 \quad (5.26a)$$

and

$$\lim_{h \rightarrow 0} w(\underline{v}, h) = \delta(\underline{v}) , \quad (5.26b)$$

where h is the support (i.e., characteristic width), and $\underline{A}_k \equiv \underline{A}(\underline{x}_k)$ and $\rho_k \equiv \rho(\underline{x}_k)$. Under this approximate representation, the density becomes

$$\langle \rho(\underline{x}) \rangle = \sum_{k=1}^N m_k w(\underline{x} - \underline{x}_k, h) , \quad (5.27)$$

and the equation of motion,

$$\frac{d\mathbf{v}_i}{dt} = -\frac{1}{\rho} \nabla p = -\left[\nabla \left(\frac{p}{\rho} \right) + \frac{p}{\rho^2} \nabla \rho \right], \quad (5.28)$$

becomes

$$\frac{d\langle \mathbf{v}_i \rangle}{dt} \approx -\sum_k m_k \left[\frac{p_k}{\rho_k^2} + \frac{p_i}{\rho_i^2} \right] \nabla_i w(\mathbf{x}_i - \mathbf{x}_k, \mathbf{h}). \quad (5.29)$$

The energy equation can be represented in a similar fashion. Artificial viscosities can be introduced to prevent collisionless penetration of different streams of particles. SPH has the merit of being easy to implement in three dimensions and was, in fact, first developed to treat situations in which fluid masses move arbitrarily in the absence of boundaries.^{16,17} However, some difficulties are met when modeling low density regions next to, and interacting with, high density regions.

5.2.4 Time Scale Limitations

Finally, it should be mentioned that the time step used in a numerical simulation must be such that it satisfies the Courant-Friedrichs-Lewy (CFL) condition¹⁸ in order for the solution to the finite difference equations to converge stably onto the true solution of the original partial differential equations. The CFL condition states that the ratio of the grid spacing to the time step must be higher than the highest characteristic velocity present in the system as allowed by the equations being solved. For example, if no plasma equations are involved, this characteristic velocity could be the acoustic velocity or the flow velocity. Directionality considerations enter in multidimensions. Here, it will be noted that PIC is stable against local violations of the CFL condition and that the appearance of vacuum cells do not hinder calculations.

Another limitation on how large the time step may be made in a simulation using finite difference equations comes from the criterion that the internal energy, or equivalently, the pressure, not go negative.³ This often places a more stringent upper limit on the time

step than the CFL condition, but the form of the criterion is similar to that of the CFL condition.

As an aside, limitations on the maximum time step employable in a numerical calculation has severe consequences, especially when treating plasmas. In a plasma, a wide range of time scales are involved, ranging from the periods of plasma oscillations and cyclotron gyrations to the characteristic time scales of bulk motion. If one were to employ the particle picture, thereby rigorously following the particle dynamics, the time step would have to be chosen shorter than the characteristic time scale of the fastest phenomenon allowed by the particle equations. If this were done, it would become impractical to handle the far slower MHD phenomena. The slowest of MHD flow phenomena must then be handled separately by a bulk flow code. Attempts to fill in the gap between plasma particle codes and lowest frequency MHD codes have been made with some difficulty. Hybrid codes such as those in which the ion momentum equation is combined with an inertialess electron momentum equation are examples of such attempts.^{19,20}

5.3 Nishiguchi-Yabe Algorithm

The particular PIC scheme used in this work was the one developed by Nishiguchi and Yabe,^{1,2} of which the second-order method is dubbed SOAP (Second-Order-Accurate Fluid Particle Scheme). This scheme is a Classical PIC method employing FSPs and higher-order interpolations for changes in momentum and energy due to transport. It is an example of a Classical PIC algorithm that tackles the problem of noise and numerical diffusion in a highly Eulerian manner by improving the handling of the advective term. Other than being characterized by significantly reduced numerical diffusion without requiring as much memory as a Full-Particle approach, because it is a Classical scheme, it also exhibits ease of conservation of momentum and energy and ease of application to

complex flow situations. In problems involving large expansion ratios that necessitate large variations in density to be treated, the algorithm can be modified to incorporate nonuniform and moving grids to improve the local resolution only where needed. Computational particles can be created or annihilated as necessary. The scheme has been successfully applied to plasma simulations applicable to laser fusion.

The primary feature of the SOAP algorithm is its ability to give a second-order-accurate treatment of the advective term in space (first-order-accurate in time) by assigning a distribution of quantities within each FSP. To see how this is done, the concept of numerical diffusion will be briefly explained. Numerical diffusion is an unphysical diffusive term that enters into a numerical analysis as the result of a finite difference algorithm. Depending upon what is being transported, numerical diffusion can be categorized into numerical mass diffusion, numerical viscosity (momentum diffusion), and numerical heat conduction (energy diffusion). In Figure 5.1, we see a single FSP presiding over four neighboring grid cells in two-dimensional space. The diagram to the right depicts the situation after the FSP has undergone Lagrangian motion from the state at the left. The velocity by which the FSP is moved is obtained by area weighting the tilde velocities as ³

$$\underline{v}_p = A_1 \tilde{v}_{i,j} + A_2 \tilde{v}_{i+1,j} + A_3 \tilde{v}_{i,j+1} + A_4 \tilde{v}_{i+1,j+1} , \quad (5.30)$$

where the As' (see also Equation (5.31)) and Bs' (see Equations (5.32) and (5.33)) are the partial cross-sectional areas** of the FSP. The total area of the FSP is being normalized to 1 for convenience (i.e., $A_1 + A_2 + A_3 + A_4 = 1$ and $B_1 + B_2 + \dots + B_9 = 1$). A physically conserved quantity, i.e., momentum or energy, that is assigned to the i,j th cell is being denoted by $U_{i,j}$. All $U_{i,j}^{\text{old}}$ s' are specific quantities (i.e., total quantity contained in a cell divided by the total normalized mass of FSPs overlapping that cell, where the latter is

** Actually, they should be considered as mass because area weighting is carried out in a "space," not necessarily physical, in which the mass distribution across the cross section of each FSP is uniform. In other words, the "area" in the "cross-sectional area weighting" scheme is not necessarily a physical area, but an area in the "space" just described.

equivalent to the total fractional FSP cross-sectional area contained in the cell in a space in which the cross-sectional area weighting is carried out), and are the tilde variables. For the sake of simplifying the discussion, we will take $U_{i,j}$ to represent momentum. Now, a simple method of prescribing the transferring of momentum from the grid to the FSP and back is the following :

$$U^* = A_1 U_{i,j}^{\text{old}} + A_2 U_{i+1,j}^{\text{old}} + A_3 U_{i,j+1}^{\text{old}} + A_4 U_{i+1,j+1}^{\text{old}} \quad (5.31)$$

$$\begin{aligned} U_{i,j}^{\text{new}} &= B_1 U^*, & U_{i+1,j}^{\text{new}} &= (B_2 + B_3) U^* \\ U_{i,j+1}^{\text{new}} &= (B_4 + B_7) U^*, & U_{i+1,j+1}^{\text{new}} &= (B_5 + B_6 + B_8 + B_9) U^* . \end{aligned} \quad (5.32)$$

Here, the total momentum carried by the FSP at the end of the Eulerian phase is U^* . The momentum is prescribed to be uniformly distributed over the FSP in the space depicted in the figure, so after the motion of the FSP, the FSP under consideration will supply the i,j th cell with a momentum of $B_1 U^*$, the $i+1,j$ th cell with $(B_2 + B_3) U^*$, etc. The total momentum contained in the i,j th cell after FSP motion will be obtained by taking the sum of $U_{i,j}^{\text{new}}$ contributions from each FSP. Dividing this total by the total sum of fractional FSP areas (i.e., normalized mass) overlapping the i,j th cell in the postmotion picture will give the specific $U_{i,j}$ for the new time step.

However, with this algorithm, it is evident that even without any particle motion, the grid solution changes from one time step to the next. As alluded to earlier, this is called zeroth-order diffusion. Zeroth-order diffusion can be removed by memorizing the locations of cell boundaries in an FSP before motion. For the example depicted in Figure 5.1, this procedure becomes

$$\begin{aligned} U_{i,j}^{\text{new}} &= B_1 U_{i,j}^{\text{old}} \\ U_{i+1,j}^{\text{new}} &= B_2 U_{i,j}^{\text{old}} + B_3 U_{i+1,j}^{\text{old}} \\ U_{i,j+1}^{\text{new}} &= B_4 U_{i,j}^{\text{old}} + B_7 U_{i,j+1}^{\text{old}} \\ U_{i+1,j+1}^{\text{new}} &= B_5 U_{i,j}^{\text{old}} + B_6 U_{i+1,j}^{\text{old}} + B_8 U_{i,j+1}^{\text{old}} + B_9 U_{i+1,j+1}^{\text{old}} . \end{aligned} \quad (5.33)$$

This area-weighting scheme is called a first-order scheme because the error introduced by the finite difference algorithm is first-order in the grid spacing. This can be seen as follows. Taking the x component of the momentum equation as an example,

$$\frac{\partial(\rho u)}{\partial t} = -\frac{\partial p}{\partial x} - \frac{\partial(\rho u^2)}{\partial x} \quad (5.34)$$

is effectively represented approximately by a PIC algorithm employing the method of Equation (5.33) as ¹

$$\frac{(\rho u)_i^{n+1} - (\rho u)_i^n}{\Delta t} = -\frac{p_{i+\frac{1}{2}}^n - p_{i-\frac{1}{2}}^n}{\Delta x} - \frac{(\rho \bar{u} \langle \bar{u} \rangle)_{i+\frac{1}{2}} - (\rho \bar{u} \langle \bar{u} \rangle)_{i-\frac{1}{2}}}{\Delta x}, \quad (5.35)$$

where

$$\langle \bar{u} \rangle_{i-\frac{1}{2}} \equiv \begin{cases} u_{i-1} & \text{for } \bar{u}_{i-\frac{1}{2}} > 0 \\ u_i & \text{for } \bar{u}_{i-\frac{1}{2}} < 0. \end{cases} \quad (5.36)$$

A Taylor expansion of Equation (5.35) yields

$$\frac{\partial(\rho u)}{\partial t} = -\frac{\partial p}{\partial x} - \frac{\partial(\rho u^2)}{\partial x} + \frac{\partial\left(\frac{1}{2}\rho|u|\frac{\partial u}{\partial x}\right)}{\partial x} \Delta x + O(\Delta x^2, \Delta t). \quad (5.37)$$

When this is compared with the original differential equation, one finds that the lowest-order error in Δx that has crept in is the third term on the right-hand side of Equation (5.37). This term acts as a viscosity and has the property of being strong where the fluid velocity relative to the grid is high and where the velocity gradient is high. This is the numerical viscosity. The representation of the flow on a finite resolution grid using finite difference equations and a particular distribution of quantities within the moving FSPs has given rise to round-off errors of a particular order, which are dissipative and act effectively as a diffusion term. The numerical heat conduction term can be found in much the same way.

The first-order numerical diffusion (viscosity and heat conduction) may be removed by resorting to a second-order scheme in which the specific momentum (i.e., velocity) and specific energy are assigned higher-order distributions within the FSPs. A distribution given by the following equation will yield a second-order result :

$$U(\xi, \eta) = \left\{ (U_{i+1,j+1} - U_{i+1,j})(2\eta - \delta\eta) + U_{i+1,j} \right\} (2\xi - \delta\xi) - \left\{ (U_{i,j+1} - U_{i,j})(2\eta - \delta\eta) + U_{i,j} \right\} (2\xi - \delta\xi - 1), \quad (5.38)$$

where ξ and η are coordinates measured in the horizontal and vertical directions from the lower left-hand corner of our FSP in Figure 5.1. The contributions from the individual FSPs to the U of each cell may be obtained by integrating the $U(\xi, \eta)$ s in ξ - η space over the respective overlapping areas.

In one dimension, Equation (5.38) reduces to

$$U(\xi) = 2\xi(U_i - U_{i-1}) + (1 + \delta\xi)U_{i-1} - \delta\xi U_i, \quad (5.39)$$

where $\delta\xi$ is the length of the (unit length) FSP that is contained in the $i-1$ th cell. The right-hand diagram of Figure 5.2 shows the distribution of the physical quantity U within an FSP as dictated by Equation (5.39). This distribution is maintained within an FSP during its motion. One sees from the left-hand diagram of Figure 5.2 that because cells are assigned only average values, a stepwise distribution for U exists. This results in errors on the order of $\frac{\partial U}{\partial x} \Delta x$, but this error is significantly reduced by taking the method illustrated in the right-hand diagram. With this second-order scheme, the equation corresponding to Equation (5.35) becomes

$$\frac{(\rho u)_i^{n+1} - (\rho u)_i^n}{\Delta t} = -\frac{p_{i+\frac{1}{2}}^n - p_{i-\frac{1}{2}}^n}{\Delta x} - \frac{(\rho \bar{u}^2)_{i+\frac{1}{2}} - (\rho \bar{u}^2)_{i-\frac{1}{2}}}{\Delta x}, \quad (5.40)$$

and taking the Taylor expansion of this equation, one obtains

$$\frac{\partial(\rho u)}{\partial t} = -\frac{\partial p}{\partial x} - \frac{\partial(\rho u^2)}{\partial x} + O(\Delta x^2, \Delta t). \quad (5.41)$$

Note that the first-order numerical viscosity term has disappeared.

In the approach just discussed, an attempt was made to eliminate the origin of numerical viscosity by assuming a particular distribution of quantities within the FSPs. It is also possible to try to subtract numerical viscosity in finite difference form, but it is not

easy to cancel accurately the original numerical viscosity this way, especially in multidimensions, and this can result in high noise.

Multidimensional simulations in cylindrical or spherical coordinates generally proceed in the same manner as in Cartesian coordinates, but some points are worth mentioning. For simplicity, azimuthal symmetry will be assumed. In the case of cylindrical coordinates, the cross section of the grid along a plane containing the central axis looks identical to a Cartesian grid plane, but the actual grid cells are toroidal in shape and a $1/r$ factor enters into the conservation equations. When working in spherical coordinates, it is still possible to solve the differential equations in cylindrical coordinates by projecting the grid cells of the spherical mesh onto the cylindrical mesh, but the finite difference equations will be different from the standard cylindrical equations because of the difference in geometry of the cells between the two grid systems.^{1,2} That approach was not taken in this thesis. The other approach to solving problems on a spherical polar grid is obviously to use the equations written out in spherical coordinates.¹ A coordinate plane in which the r and θ grid lines form a rectangular mesh will be employed, and the mass distribution within each FSP will be taken to be uniform in this plane instead of in physical space, and cross-sectional area weighting will also be carried out on the rectangular r - θ grid.^{***} This prevents fictitious mass fluxes from arising during radial particle motion and also makes the transformation of physical quantities between cells and particles easier.² Although interpolations for particle velocities are undertaken on the rectangular r - θ grid, the advancement of particle positions with time should be carried out in physical cylindrical coordinates instead of in physical spherical coordinates to prevent fictitious forces such as numerical centrifugal forces and Coriolis forces from entering. The problem with advancing particle positions in spherical coordinates can be recognized most clearly if one

^{***} In slab-geometry simulations, cross-sectional "area (actually "length" in 1-D)" weighting will be carried out in physical "space."

considers simulating the motion of a particle undergoing inertial motion with a velocity purely in the θ direction at the moment under concern. If the particle position is advanced to the next time step in spherical coordinates, it will travel along a constant r trajectory and end up at a different position from the point where a straight trajectory in physical space in the direction of the original velocity would have taken it. This false curving of a particle trajectory can be considered as a manifestation of fictitious forces and will be especially problematic near the origin of a spherical coordinate system.

5.4 Artificial Viscosities and Shock Capturing

Up to this point, we have been concerned mainly with the removal of spurious diffusion terms not present in the original differential equations. However, such diffusion terms can play positive roles in numerical simulations. For example, numerical viscosities are useful in stabilizing PIC calculations. In fact, the Eulerian part of Classical PIC is unconditionally unstable, and diffusive effects are helping suppress this problem. Short-wavelength numerical noise and instabilities are attenuated by the presence of diffusive terms such as viscosities because they convert the kinetic energy of noise oscillations to heat. Furthermore, diffusion, and especially viscosity, play an important role in the treatment of shocks.

When a shock is treated theoretically as a mathematical discontinuity, two different sets of solutions, one for each side of the discontinuity, must be matched self-consistently across the shock so that the Rankine-Hugoniot relations are satisfied. To do this numerically, especially in multidimensions, would be difficult. The shock motion is unknown to start with, and many iterations per time step would be required to find the correct shock motion and solutions to the fluid variables on both sides of the shock. Such a method of looking for the correct solution of a shock is called shock fitting.

A real shock has a finite thickness (on the order of a few mean free paths) by virtue of diffusion. In fact, it is the viscosity (and to some extent heat conduction) ¹⁸ that prevents a compression wave from steepening indefinitely until the unphysical situation of a multivalued state arises. This suggests the possibility of using numerically or artificially introduced viscosities to spread a shock discontinuity out into a smooth transition region of finite width (i.e., a few grid cells wide), thereby making possible the treatment of states on both sides of the shock region, as well as through the shock, by one continuous solution. This will eliminate the necessity of carrying out cumbersome iterations to find a shock, because shocks will automatically appear as near-discontinuities at roughly the correct location with approximately the correct speed, strength, and jump conditions. This method of automatically handling shocks through the use of diffusive effects is called shock capturing.

In numerical work relying on shock-capturing methods, an artificial viscosity term is often added to the finite difference equations, especially if the numerical viscosity inherent to the algorithm is not high enough in the regions of interest. Artificial viscosities should have the property that they act only where there is compression (i.e., should not admit sudden decompressions). As viscosities, they must damp out short-wavelength oscillations. Also, the resultant shock thickness should be more or less independent of the shock strength, ¹⁸ and the shock should, preferably, extend over only a few grid zones.

A popular artificial viscosity is the von Neumann-Richtmyer viscosity, which has the form ^{18,21}

$$q_B = \begin{cases} b\rho\left(\frac{\partial u}{\partial x}\right)^2 \Delta x^2 & \text{for } \frac{\partial u}{\partial x} < 0 \\ 0 & \text{otherwise} \end{cases}, \quad (5.42)$$

where "b" is a constant, typically chosen to be on the order of unity. The thickness of the transition region can be adjusted by varying the coefficient "b." The quadratic dependence on $\partial u/\partial x$ enables the thickness of the shock transition region to be independent of the shock strength. With ordinary viscosities, the viscous stress is proportional to the rate of

shear, and the shock thickness is proportional to the shock strength. However, if only this second-order viscosity is used, large numerical oscillations arise behind the shock, high coefficients become necessary to suppress long-wavelength oscillations, and the shock becomes rather wide. Trailing postshock oscillations are more effectively suppressed by the first-order viscosity discussed below. The use of high viscosity coefficients allows the suppression of large-amplitude oscillations, but the overall results tend to get overly distorted from excessive diffusion, and the profiles can also develop other spurious features. Compromises must be made regarding the strength of viscosities to be employed in a code. Another example of an artificial viscosity is the Landshoff type, which has the form ²²

$$q_A = \begin{cases} -apc_s \left(\frac{\partial u}{\partial x} \right) \Delta x & \text{for } \frac{\partial u}{\partial x} < 0 \\ 0 & \text{otherwise} \end{cases}, \quad (5.43)$$

where c_s is the sound speed and "a" is a constant which, again, is typically chosen to be on the order of unity. This first-order viscosity more readily enables the suppression of long-wavelength oscillations and overshoots in narrow shock regions, and allows for thinner transition regions, but is weak at strong shock fronts. The accuracy by which the Rankine-Hugoniot relations are obeyed is not much affected by the type of viscosity employed. Following other examples, a linear combination of the von Neumann-Richtmyer and Landshoff viscosities was chosen wherever an artificial viscosity was used in this thesis. Artificial viscosities (q's) can be incorporated into the equations by replacing the pressure term p everywhere by $p+q$.

Whereas viscosity allows all quantities to vary smoothly through a shock transition region, heat conduction by itself cannot make all quantities vary continuously through that region. Temperature will vary smoothly through a shock, but part of the pressure and density transitions will occur through a discontinuity for a strong shock.¹⁸ Therefore, the presence of numerical heat conduction is generally not favored. It only distorts the flow without the merits of viscosity.

5.5 Ringing Instability

The ringing instability is a numerical instability peculiar to PIC calculations and is encountered when certain criteria are met.²³ This instability, characterized by large density fluctuations with particles bunching into clumps, arises from the employment of both particles and grids. It is primarily a finite-grid instability that is due to aliasing. Aliases are different particle solutions which are indistinguishable when projected onto the grid on which the physical quantities are defined (note that there are far more particles than there are grid cells). Because the number of computational particles is finite, the number of alias particle modes supported will also be finite. Multiple particle modes that are aliases of the same grid solution introduce additional resonances into the dispersion relation, and through nonlinear wave interactions, this can result in a growing mode instability. Given a wave number k , determined by the parameters of the code, there will be a range of fluid velocities for which the instability can grow. Stability analyses have shown that the growing instability will occur in a range of flow velocities below a critical limit much lower than the local sound speed. Since the instability growth rate is only weakly dependent on the number of aliases, the growth rate will also have very little dependence on the number of computational particles employed. Implicit differencing should suppress the ringing instability, but the use of higher-order interpolations with extended support²³ and the addition of viscosities can also weaken or suppress, or at least substantially slow, the growth rate of the instability.

5.6 Summary

The central issue of this work was to treat the novel problem of a hydrodynamic expansion of a fluid into a vacuum that applies a finite pressure at the fluid-vacuum interface. Unlike other problems involving a fluid(plasma)-vacuum interface, the primary direction of motion of the interface in the problems handled in this work is normal to the interface, instead of being along it. The PIC scheme was chosen for the hydrodynamic calculations, first and foremost because it was deemed capable of automatically capturing these interfaces with unknown and dramatic motions. The use of Lagrangian computational particles would make this possible when coupled with suitable application of boundary conditions and prescriptions at the interface (discussed in the following chapter). But unlike the Lagrangian scheme, the use of an underlying Eulerian grid in PIC prevents mesh entanglements from occurring in complex multi-dimensional flows. The use of a grid also enables the rigorous conservation of physically conserved quantities, which is a necessity for obtaining good results.

The Eulerian method will generally require many iterations per time step to find and keep track of a moving interface, and this can become very difficult in multi-dimensions. Eulerian codes also suffer from negative densities, advective term instabilities, and numerical diffusions, but these problems are either circumvented or reduced with PIC. And unlike the case of SPH, in PIC, particles need not be present where field calculations must be performed.

Of the PIC algorithms, an advanced Classical scheme was selected. Such a scheme, in contrast to the competing Full-Particle scheme, allows a reduction in numerical diffusion and noise without the intensive memory requirements and other problems that plague the Full-Particle approach, which is highly Lagrangian.

Following is a list of representative equations that were solved in the hydrodynamic part of the numerical simulations carried out in this work, as selected from the main text of this chapter.

With a PIC scheme, the conservation of mass is automatic through the use of Lagrangian computational particles, and the continuity equation does not appear. The calculation of the momentum and energy conservation equations is split into two steps. The first step is the finite differencing of the equations less the advective term, on an Eulerian grid. For the case of two-dimensional Cartesian coordinates, this may be written as

$$\tilde{u}_{i,j} = u_{i,j}^n - \frac{\Delta t}{\rho_{i,j}^n \Delta x} \left(p_{i+\frac{1}{2},j}^n - p_{i-\frac{1}{2},j}^n \right) \quad (5.17)$$

$$\tilde{v}_{i,j} = v_{i,j}^n - \frac{\Delta t}{\rho_{i,j}^n \Delta y} \left(p_{i,j+\frac{1}{2}}^n - p_{i,j-\frac{1}{2}}^n \right) \quad (5.18)$$

$$\tilde{\epsilon}_{i,j} = \epsilon_{i,j}^n - \frac{\Delta t}{\rho_{i,j}^n} \left\{ \frac{\left[(p^n \bar{u})_{i+\frac{1}{2},j} - (p^n \bar{u})_{i-\frac{1}{2},j} \right]}{\Delta x} + \frac{\left[(p^n \bar{v})_{i,j+\frac{1}{2}} - (p^n \bar{v})_{i,j-\frac{1}{2}} \right]}{\Delta y} \right\}, \quad (5.19)$$

where here, u and v are, respectively, the x and y component Eulerian fluid velocities (i.e., v_x and v_y), and ϵ is the specific total energy

$$\epsilon = \frac{1}{\gamma - 1} \frac{p}{\rho} + \frac{|\mathbf{v}|^2}{2}. \quad (5.20)$$

The barred velocities are defined through

$$\bar{\mathbf{v}} = \frac{\mathbf{v}^n + \tilde{\mathbf{v}}}{2}. \quad (5.21)$$

The tilde (\sim) quantities for the grid cells represent physical variables at the end of the first step. In the two-dimensional hydrodynamic simulations conducted in this work, a spherical polar coordinate version of these equations was employed. In some cases, the adiabatic condition (isentropic, if uniform) was imposed, and took the place of the energy equation displayed above.

Boundary conditions are generally prescribed as explained in the main text on pp. V-4 to V-5.

After the first step, the fluid quantities are assigned to the FSPs, from the grid cells, via area weighting of the tilde grid-cell values, and the finite-sized computational particles (FSPs) are then moved physically in a Lagrangian fashion, using area-weighted tilde velocities. This second step takes care of material transport effects that were neglected in the first step. For the particle of the example shown in Figure 5.1, this area weighted velocity will be given as

$$\underline{v}_p = A_1 \tilde{v}_{i,j} + A_2 \tilde{v}_{i+1,j} + A_3 \tilde{v}_{i,j+1} + A_4 \tilde{v}_{i+1,j+1} , \quad (5.30)$$

where the A 's and B (to appear later) s ' are the partial cross-sectional areas of the FSP overlapping the respective cells. The total area of the FSP is being normalized to 1 for convenience (i.e., $A_1 + A_2 + A_3 + A_4 = 1$ and $B_1 + B_2 + \dots + B_9 = 1$).

Once the physical quantities are transported along with the computational particles, the fluid quantities are repartitioned, from the particles, back to the cells, and this completes the set of calculations for advancing fluid quantities from one time step to the next. If a physically conserved quantity, i.e., momentum or energy, that is assigned to the i,j th cell is represented by $U_{i,j}$, a zeroth-order scheme may be described by the following equations for the situation depicted in Figure 5.1 :

$$U^* = A_1 U_{i,j}^{\text{old}} + A_2 U_{i+1,j}^{\text{old}} + A_3 U_{i,j+1}^{\text{old}} + A_4 U_{i+1,j+1}^{\text{old}} \quad (5.31)$$

$$\begin{aligned} U_{i,j}^{\text{new}} &= B_1 U^* , & U_{i+1,j}^{\text{new}} &= (B_2 + B_3) U^* \\ U_{i,j+1}^{\text{new}} &= (B_4 + B_7) U^* , & U_{i+1,j+1}^{\text{new}} &= (B_5 + B_6 + B_8 + B_9) U^* . \end{aligned} \quad (5.32)$$

All $U_{i,j}^{\text{old}}$ s ' are specific quantities (i.e., total quantity contained in a cell divided by the sum of fractional cross-sectional areas of FSPs overlapping that cell), and are the tilde variables. Mass is prescribed to be uniformly distributed over an FSP in the space in which the cross-sectional area weighting is carried out, although as

discussed earlier in this chapter, this space need not be physical space. Therefore the cross-sectional area weighting scheme is equivalent to a mass weighting scheme. The total quantity contained in the i,j th cell after FSP motion is obtained by taking the sum of $U_{i,j}^{new}$ contributions from each FSP. Dividing this total by the total sum of fractional FSP areas (i.e., normalized mass) overlapping the i,j th cell in the post motion picture will give the specific $U_{i,j}$ for the new time step.

An area-weighting scheme that is first-order accurate may be represented by the following equations, for the example depicted in Figure 5.1 :

$$\begin{aligned}
 U_{i,j}^{new} &= B_1 U_{i,j}^{old} \\
 U_{i+1,j}^{new} &= B_2 U_{i,j}^{old} + B_3 U_{i+1,j}^{old} \\
 U_{i,j+1}^{new} &= B_4 U_{i,j}^{old} + B_7 U_{i,j+1}^{old} \\
 U_{i+1,j+1}^{new} &= B_5 U_{i,j}^{old} + B_6 U_{i+1,j}^{old} + B_8 U_{i,j+1}^{old} + B_9 U_{i+1,j+1}^{old} .
 \end{aligned} \quad (5.33)$$

On the other hand, a distribution given by the following equation yields a second-order calculation :

$$\begin{aligned}
 U(\xi, \eta) &= \left\{ (U_{i+1,j+1} - U_{i+1,j})(2\eta - \delta\eta) + U_{i+1,j} \right\} (2\xi - \delta\xi) \\
 &\quad - \left\{ (U_{i,j+1} - U_{i,j})(2\eta - \delta\eta) + U_{i,j} \right\} (2\xi - \delta\xi - 1) ,
 \end{aligned} \quad (5.38)$$

where ξ and η are coordinates measured in the horizontal and vertical directions from the lower left-hand corner of our FSP in Figure 5.1. The contributions from the individual FSPs to the U of each cell is obtained by integrating the $U(\xi, \eta)$ s in ξ - η space over the respective overlapping areas.

Apart from the various order numerical viscosities that are introduced by employing the area-weighting schemes discussed above, artificial viscosities were also used, especially for automatically capturing shocks. In particular, a linear combination of the von Neumann-Richtmyer viscosity and the Landshoff viscosity was employed. In slab geometry, these artificial viscosities are defined, respectively, through

$$q_B = \begin{cases} b\rho\left(\frac{\partial u}{\partial x}\right)^2 \Delta x^2 & \text{for } \frac{\partial u}{\partial x} < 0 \\ 0 & \text{otherwise} \end{cases} \quad (5.42)$$

and

$$q_A = \begin{cases} -a\rho c_s\left(\frac{\partial u}{\partial x}\right) \Delta x & \text{for } \frac{\partial u}{\partial x} < 0 \\ 0 & \text{otherwise} \end{cases}, \quad (5.43)$$

where c_s is the sound speed and "b" and "a" are constants typically chosen to be on the order of unity. These viscosities, q 's, may be incorporated into the analysis by replacing the pressure term, p , in the equations, by $p+q$.

References for Chapter 5

1. A. NISHIGUCHI, Special Research Report, Osaka University (8 March, 1982).
2. A. NISHIGUCHI and T. YABE, *J. Comp. Phys.* **52** (1983), 390.
3. See, for example, F.H. HARLOW, "The Particle-in-Cell Computing Method for Fluid Dynamics," in : *Methods in Computational Physics vol.3*, eds. B. ALDER, S. FERNBACH, and M. ROTENBERG (Academic Press, New York, 1964).
4. See, for example, J.N. LEBOEUF, T. TAJIMA, and J.M. DAWSON, *J. Comp. Phys.* **31** (1979), 379.
5. See, for example, J.U. BRACKBILL, D.B. KOTHE, and H.M. RUPPEL, *Computer Physics Communications* **48** (1988), 25.
6. See, for example, T. TAJIMA, *Computational Plasma Physics : with Applications to Fusion and Astrophysics* (Addison-Wesley, Redwood City, 1989).
7. See, for example, J.P. BORIS and D.L. BOOK, *J. Comp. Phys.* **11** (1973), 38.
8. D.L. BOOK, J.P. BORIS, and K. HAIN, *J. Comp. Phys.* **18** (1975), 248.
9. J.P. BORIS and D.L. BOOK, *J. Comp. Phys.* **20** (1976), 397.
10. See, for example, G.A. SOD, *J. Comp. Phys.* **27** (1978), 1.
11. See, for example, P.J. ROACHE, *Computational Fluid Dynamics, revised printing*, (Hermosa Publishers, Albuquerque, 1976).
12. See, for example, J.U. BRACKBILL and H.M. RUPPEL, *J. Comp. Phys.* **65** (1986), 314.
13. See, for example, K. MIYAMOTO, *Plasma Physics for Nuclear Fusion* (MIT Press, Cambridge, 1980), or generally any other book on plasma physics or electromagnetism.
14. See, for example, F. BRUNEL et al., *J. Comp. Phys.* **43** (1981), 268.
15. See, for example, J.J. MONAGHAN, *Computer Physics Communications.* **48** (1988), 89.

16. L.B. LUCY, *Astron. J.* **82** (1977), 1013.
17. R.A. GINGOLD and J.J. MONAGHAN, *Mon. Not. Roy. Astron. Soc.* **181** (1977), 375.
18. See, for example, R.D. RICHTMYER and K.W. MORTON, *Difference Methods for Initial Value Problems, 2nd ed.* (Wiley-Interscience, New York, 1967).
19. See, for example, T. TAJIMA, J.N. LEBOEUF, and J.M. DAWSON, *J. Comp. Phys.* **38** (1980), 237.
20. D.W. HEWETT, *J. Comp. Phys.* **38** (1980), 378.
21. J. von NEUMANN and R.D. RICHTMYER, *J. Applied Phys.* **21** (1950), 232.
22. R. LANDSHOFF, Los Alamos Scientific Laboratory Report LA-1930 (1955).
23. J.U. BRACKBILL, *J. Comp. Phys.* **75** (1988), 469.

Appendix of Chapter 5

Part 1: Figures for Chapter 5

(Starts Page V-A2)

Part 2: Reference for the Appendix of Chapter 5

(Page V-A4)

Illustration of an area weighting procedure showing a single FSP centered about (\bullet) , before (a), and after (b) motion.

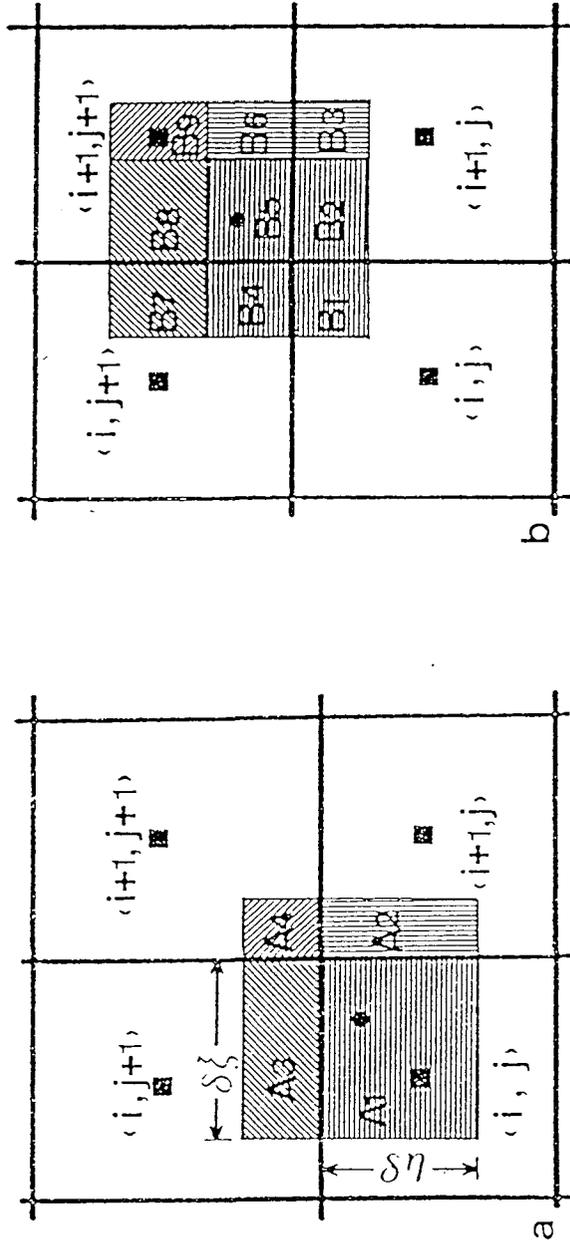
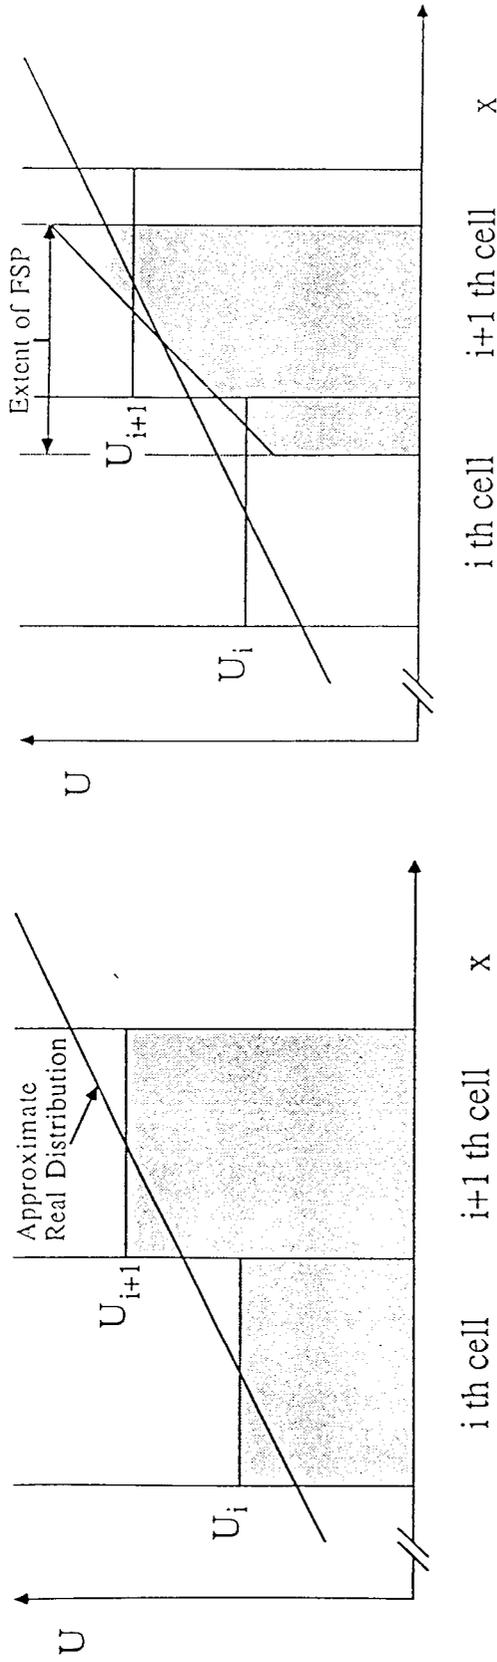


Figure 5.1 ¹



Step-wise distribution due to finite gridding.

a

Higher-order interpolation within each FSP.
(Second-order interpolation for one FSP shown.)

b

Figure 5.2

Part 2: Reference for the Appendix of Chapter 5

1. A. NISHIGUCHI, Doctoral Thesis, Osaka University (1985).

Chapter 6

Slab-Geometry Hydrodynamic Simulations

This chapter treats hydrodynamic PIC simulations in one-dimensional slab (i.e., planar) geometry. For reasons that will be stated in Section 6.2, calculation of real magnetic pressures was not carried out for the work described in this chapter. The slab-geometry codes were written and run as a step towards developing a two-dimensional hydrodynamic PIC code that would simulate the expansion of a conducting fluid in the realistic geometry of a thruster for an ICF pulse rocket. Namely, the techniques studied and gained in slab geometry will be applied to the more complex, multidimensional ICF pulse rocket simulations. With a slab-geometry code, it is possible to develop and test the basic techniques necessary for handling the expansion of a perfectly conducting fluid against a vacuum magnetic field without being plagued by complications that would be encountered in a multidimensional code. One can also check the performance of a simulation rather easily against theory, because exact analytical solutions exist for several slab-geometry problems. A feel for the fluid profiles to expect, as well as an estimate for the general accuracy of the algorithm, even if in one dimension, should help when working in multidimensions where exact analytical solutions are not readily available. The way the flow behaves under various prescriptions was studied, as an analogy was expected to exist between the slab geometry situations and the real thruster situations in multidimensions, where the thruster flows must be examined and compared against each other as important parameters are varied. The analogy was indeed found to exist. Related to this, an analysis of the results obtained in one dimension will aid in interpreting multidimensional results through supplying knowledge regarding, for example, limitations of the algorithm and typical numerical features that are prone to arise during the simulation. However, with a

slab-geometry situation, it is true that it is difficult to select values of parameters comparable to those characterizing a real two-dimensional thruster situation. There can also be some differences in the qualitative features of the flow between the slab geometry and multidimensional cases, because of the difference in the dimensionality of the respective expansion flow processes.

One of the major objectives of this thesis was to develop a code that would automatically capture the interface motion of bulk flow in the magnetic thruster of an ICF pulse rocket without the use of any trial-and-error iterations. This technique, along with the related issues of the handling of vacuum interfaces, was studied. The slab-geometry codes were also used to test the capturing of shocks and other discontinuities, such as contact surfaces, by investigating the effects of artificial viscosities and various orders of numerical viscosities. This was important, as shocks can form in an expansion flow process involving a decelerating interface.

The problems that were run in slab geometry were free expansion into a (force-free) vacuum, expansion into a vacuum with a finite pressure applied at the fluid interface, and shock-tube problems involving the breakup of arbitrary discontinuities. The results of the first and last simulations can be checked against analytical results, and the second problem resembles that of expansion of an unmagnetized perfectly conducting fluid against a vacuum magnetic field. In all cases, good agreement with theory was obtained. Concerning the automatic tracking of interface motion, the code was able to produce, at all time steps, an interfacial fluid pressure that rather closely matched the artificially prescribed external pressure applied at the interface, indicating correct capturing of the interface. The position of the interface and the profiles of the physical variables (by default, "profile" will refer to a distribution of a physical quantity with respect to the spatial coordinate) up to the interface were accurately found. In this capability, the code turned out to be quite robust, being able to capture the interface for a wide range of externally applied interfacial pressure prescriptions, including cases in which the applied interfacial pressure was deliberately

formulated to diverge actively away from the interfacial fluid pressure value. Even if the initial pressure mismatch across the interface between the internal fluid pressure and the externally applied pressure was large, good interface tracking could be observed. Shocks were successfully captured, and contact surfaces clearly delineated.

With PIC, the ratio of the initial fluid density to the lowest fluid density that can be treated is limited by the number of computational particles per cell used to represent the initial density. Therefore, the low densities that arise during simulations of flows with large expansion ratios become difficult to treat accurately, and eventually become impossible to treat. The use of nonuniform, moving grids is a way to circumvent this problem. Also, the discretization using finite-sized grid cells and a finite number of particles results in PIC's not being able to model the flow entirely correctly even for higher densities. In this work, ways of mitigating these problems, except for the use of nonuniform, moving grids, were considered.

One major reason extended period simulations were attempted was because the simulation of flows with large expansion ratios in the thrusters of ICF pulse rockets was the final goal of this work. This necessitated the quality of the calculations to be high, and sometimes required the suppression of numerical instabilities. Efforts were made to suppress these instabilities while maintaining the accuracy of the results (minimization of numerical distortion). Techniques such as the combining of various order numerical viscosities and artificial viscosities were tested. The tradeoff between diffusive distortion and numerical oscillations was also investigated, regarding the treatment of shocks and contact surfaces.

Previous works have employed PIC to treat various hydrodynamical problems including free expansion (especially in one dimension) into a vacuum, but many of the results obtained in the present work were of very high quality because of the use of an advanced scheme and intricate attention paid to the details, especially at interfaces. As for the calculation of an unsteady hydrodynamic (as opposed to simpler approximations)

expansion *against* a vacuum region that exerts a finite pressure upon the fluid-vacuum interface, it is believed that not many precedents exist. This is regardless of the numerical algorithm used, PIC or not. This problem should not be considered the same problem as one dealing with a plasma that does not flow primarily against (i.e., in a direction normal to) the vacuum magnetic field which it bounds, or one not involving any fluid dynamics for the interior regions of the blob.

6.1 The Equations

The equations governing the flow of an unmagnetized perfectly conducting fluid are those of classical hydrodynamics.* The continuity, momentum, and energy equations cast in conservative form were presented in the previous chapter (Equations (5.14) through (5.16)). The gas constant γ will be taken to be 5/3 in this thesis, as is appropriate for a fully ionized gas. Physical viscosity, heat conduction, and radiation were all neglected, as the major purpose of the work was to develop techniques for handling bulk flow with automatic interface capturing for situations in which such physics are not estimated to play major roles except for during limited epochs, e.g., at the very beginning of the expansion process. In slab geometry, Equations (5.14) through (5.16) reduce to

$$\frac{\partial \rho}{\partial t} + \frac{\partial(\rho u)}{\partial x} = 0 \quad (6.1)$$

* If the simplest form of the MHD equations is taken and it is further assumed that the plasma is perfectly conducting and unmagnetized, then the equations will reduce to those of classical hydrodynamics. However, the term MHD is a somewhat vague term for an approximation of a real plasma, and can encompass a range of approximations. In fact, if more than a minimal number of terms are retained in the MHD equations, even the equations for an initially unmagnetized perfectly conducting plasma will differ from those of classical hydrodynamics. This betrays the property that a plasma, even under such idealized conditions, will not behave as an unmagnetized perfectly conducting fluid obeying just the laws of classical hydrodynamics. Therefore the classical hydrodynamic analyses will not be referred to as "MHD analyses for a perfectly conducting unmagnetized plasma" in this work, unless the phrase "use of the simplest form of MHD equations" is substituted for "use of MHD equations."

$$\frac{\partial(\rho u)}{\partial t} + \frac{\partial(\rho u^2)}{\partial x} = -\frac{\partial p}{\partial x} \quad (6.2)$$

$$\frac{\partial(\rho \varepsilon)}{\partial t} + \frac{\partial(\rho \varepsilon u)}{\partial x} = -\frac{\partial(\rho u)}{\partial x} . \quad (6.3)$$

Using dimensionless quantities (the primed quantities) defined through

$$\rho' = \rho/\rho_*, \quad p' = p/p_*, \quad u' = u/u_*, \quad \varepsilon' = \varepsilon/\varepsilon_*, \quad x' = x/\mathcal{X}_*, \quad \text{and } t' = t/t_*,$$

where ρ_* , p_* , u_* , ε_* , \mathcal{X}_* , and t_* are arbitrarily chosen characteristic values obeying the relations

$$u_* = \sqrt{\gamma \frac{p_*}{\rho_*}} \quad (6.4)$$

$$\varepsilon_* = \frac{1}{\gamma - 1} \frac{p_*}{\rho_*} \quad (6.5)$$

$$t_* = \frac{\mathcal{X}_*}{u_*} \quad (6.6)$$

among themselves **, and taking only the part that is involved in the Eulerian step of Classical PIC, one obtains

$$\rho' \frac{\partial u'}{\partial t'} = -\frac{1}{\gamma} \frac{\partial p'}{\partial x'} \quad (6.7)$$

$$\rho' \frac{\partial \varepsilon'}{\partial t'} = -(\gamma - 1) \frac{\partial(\rho' u')}{\partial x'} . \quad (6.8)$$

The continuity equation (conservation of mass) is automatically satisfied through the use of Lagrangian computational particles. After dropping the primes, the finite difference representation for these equations becomes

$$\tilde{u}(i) = u(i) - \frac{\Delta t}{\gamma \rho(i)} \frac{p^+(i) - p^-(i)}{\Delta x} \quad (6.9)$$

$$\tilde{\varepsilon}(i) = \varepsilon(i) - \frac{(\gamma - 1)\Delta t}{\rho(i)} \frac{p^+(i)\bar{u}^+(i) - p^-(i)\bar{u}^-(i)}{\Delta x} , \quad (6.10)$$

where

$$p^+(i) = \frac{p(i) + p(i+1)}{2} \quad (6.11)$$

** In this chapter, ρ_* and p_* were chosen to be the initial bulk values and \mathcal{X}_* was simply chosen to be the grid cell width. Thus, for example, a velocity value of one in the plots corresponds to the initial sound speed in the bulk, and each unit of distance along the horizontal coordinate axis corresponds to one cell length.

$$\tilde{p}^-(i) = \frac{p(i) + p(i-1)}{2}, \quad (6.12)$$

and similarly for $\tilde{u}^+(i)$ and $\tilde{u}^-(i)$. The tilde quantities are values at the end of the Eulerian step, and the barred velocities are as defined in Equation (5.21).

Some simulations were run using the adiabatic assumption, wherein

$$\frac{p}{p_0} = \left(\frac{\rho}{\rho_0} \right)^\gamma \quad (6.13)$$

holds for each Lagrangian fluid element, although this will not allow shocks to be handled appropriately except in the acoustic wave limit. The quantities p_0 and ρ_0 represent the pressure and density values at some reference state passed through by the particular Lagrangian fluid element under concern. For a uniform system, the adiabatic condition becomes the isentropic relation. With this simplification, the sole conservation equation that need be dealt with in finite difference form becomes

$$\frac{\partial u'}{\partial t'} = -\frac{1}{\gamma-1} \frac{\partial \rho'^{\gamma-1}}{\partial x'}. \quad (6.14)$$

All planar-geometry simulations assumed a slab of gas with a rigid wall, or equivalently, a mirror symmetry plane^{***}, located at $x=0$ and with infinite extent and complete uniformity in the y and z directions. The extent of the slab was considered finite in the x direction, and making use of symmetry properties, only the $x \geq 0$ region was treated. However, "ghost" FSPs with their centers located in the $x < 0$ region but partially overlapping the $x \geq 0$ region had to be considered for a correct calculation. It should be added that (refer back to Chapter 5) in the slab-geometry simulations, cross-sectional area (actually "length" in 1-D) weighting was carried out in physical geometry.

^{***} Symmetric for density and pressure and antisymmetric for velocity.

6.2 Some Useful Techniques

In this section, some techniques necessary for obtaining successful simulations will be discussed. Simulations were usually started from initial states in which the density and pressure values were uniform almost everywhere, but fell off in the last one or two cells adjacent to the vacuum interface or arbitrary discontinuity. The initial bulk fluid pressure value was simply prescribed to be proportional to the γ power of the initial density value. Although an initial density profile that is perfectly uniform right up to a vacuum interface can be obtained, such an exact prescription, for example, for simulating a free expansion of an initially uniform gas held by a diaphragm, was not necessary. Such a distribution can be set up by using FSPs with nonuniform size, but as FSPs smaller than a grid-cell size do not spread over neighboring cells, inaccurate data and high noise result, and successful simulations become difficult to achieve. With a uniform grid in slab geometry, all FSPs are given the same size and mass. The initial velocity was usually selected to be zero everywhere or rising linearly with distance from zero value at $x=0$. If the fluid is initially assumed to be at rest, an initial pressure gradient prescribed across the interfacial region allows fluid motion to be initiated. A state with uniform density and pressure and a linear velocity distribution was considered to be a reasonably acceptable state from which to start the expansion simulation of an ICF pulse rocket plasma flow or any approximation thereof.

Better results with smoother profiles and less noise can obviously be obtained by increasing the number of particles per cell. The simulations whose results are presented in this chapter typically employed 50 particles per cell at the initial instant. Now, during the simulation of free expansion into a vacuum, computational particles can break off the bulk, resulting in empty cells forming between the interface of the bulk and the broken-off particles, and cause problems. To get around this problem, particles that broke off the interface were not treated by the equations and were also prescribed not to travel too far from the bulk interface.

Although FSPs allow fractional particle areas to overlap a grid cell and thus make possible the treatment of lower densities than would be possible with the same number of point particles, the accuracy of the calculation for cells in which the FSP number density drops low is still poor. The quantizations resulting from the use of computational particles prohibit decent treatment of densities lower than the value corresponding to roughly one particle per cell. Similarly, relative inaccuracies mount if the applied pressure value at the interface falls to the order of, or below, the pressure fluctuations inherent in the discretization algorithm. Furthermore, physical quantities for a cell neighboring a vacuum interface when the flow is primarily normal to the interface, as calculated from FSP information in that cell, are in general very inaccurate even if the boundary is an applied pressure interface, for the same reasons of finiteness of grid cells and particles and the resultant stepwise motion of the interface. Two cases for which the values of physical quantities calculated for such a cell are particularly bad are when a previously empty cell acquires a finite density to become the outermost filled cell and when an empty cell adjacent to the outermost filled cell is about to acquire a finite density by overflow of particles from the presently outermost filled cell. Figure 6.1 illustrates how physical quantities at the interface can fluctuate unrealistically if the interface cell is defined as the outermost filled cell at each instant. At time t , the interface density is given by ρ_a , but if the interface moves out by one cell during the transition to the next time step, the interface density will be given by ρ_c at time $t+\Delta t$, which we see is much less than ρ_a . There is too much mass accumulation in the outermost filled cell at time t and a false density hump may exist at the interface. The real interface is more extended. At the next time step, there is very little FSP area overlapping the "newly formed" interfacial cell, and the values of physical variables assigned to that cell from particle data will be particularly bad. This unreliability of physical variable calculations in a cell adjacent to a vacuum cell becomes intolerable if one tries to match the applied pressure at the interface to the pressure of the outermost fluid cell. Therefore, to mitigate these problems, although certainly not completely, the use of finite

difference equations to calculate physical quantities for a bulk fluid cell bounding a vacuum cell was usually dropped, and the fluid cell adjacent to that vacuum bounding cell was interpreted as the interfacial cell. Namely, the second from the outermost bulk fluid filled cell was defined as the interfacial cell, and the external pressure was also applied on the boundary between the outermost bulk fluid filled cell and this "second from outermost" cell. When needed, such as during the calculation of finite difference equations for an adjacent cell, the outermost fluid cell bounding the vacuum was assigned velocity and specific energy values using information from adjacent fluid cells.

It was generally found that in simulations modeling the expansion of a gas into a vacuum that applies a finite pressure at the interface, the two or so fluid cells adjacent to the actual outermost fluid filled cell display fluid pressures in the neighborhood of the externally applied interfacial pressure. However, the numerical fluctuation of fluid quantities in those cells made it impossible to reach, for all time steps and all cases, a consensus as to which cell, or average of which cells, as counted from the outermost filled cell, should be interpreted as the real interface cell in terms of pressure value matching. Instead, it turned out that the fluid pressure crudely approximated by the eye as the interfacial pressure by observing the plots best matched the externally applied interfacial pressure at all time steps. Numerical oscillations are effectively smoothed out by the process of approximating with the eye. When the pressure rise region near the interface was not narrow, the peak of that pressure rise region was usually the real interfacial fluid pressure, and when a very narrow (for example, a grid cell's width) interfacial pressure spike existed above a broader pressure rise region, the base of the spike was usually the real interfacial fluid pressure. All this implies the difficulty that would have been encountered if at each time step, the correct interface location had to be found by iterations until a pressure match across the interface was obtained. Noise could also throw off any iteration attempts.

In the slab-geometry simulations approximating the expansion of an unmagnetized perfectly conducting fluid into a vacuum magnetic field, the parameter ratios and geometries in no way resembled those expected to be found in the thrusters of ICF pulse rockets. Therefore, there was no point in assuming a real external current setup and employing a field routine to calculate the magnetic pressure acting upon the applied pressure interface. It would only have added to the computational requirements of the code. Thus the externally applied interfacial pressure for such slab-geometry codes were simply chosen to vary inversely as some power of the distance between the fluid-vacuum interface and some fixed imaginary wall in the vacuum region. Depending upon the value of this power, this interfacial pressure variation can be considered to model crudely that due to an external magnetic field under the constraints of flux conservation between the interface and the fixed wall.

For the cases presented in this chapter, the following prescription was used for the external pressure applying at the interface unless otherwise noted :

$$p_{\text{Bintf}}(t) = p_{\text{Bref}} \frac{x_L - x_{\text{intf}0}}{x_L - x_{\text{intf}}(t)} , \quad (6.15)$$

where $p_{\text{Bintf}}(t)$ and p_{Bref} are, respectively, the externally applied interfacial pressures at time t and at the initial instant, $x_{\text{intf}}(t)$ and $x_{\text{intf}0}$ are the coordinates of the interface at time t and at the initial instant, and x_L is the coordinate of the imaginary outer wall. Numerous different choices for p_{Bref} were experimented with, but it was most commonly chosen to be in the neighborhood of the initial pressure of the fluid cells in the interfacial region. For an adiabatic code, the following relation, although not necessarily realistic, was used in combination with Equation (6.13) instead of Equation (6.15) :

$$\rho_{\text{Bintf}}(t) = \rho_{\text{Bref}} \frac{x_L - x_{\text{intf}0}}{x_L - x_{\text{intf}}(t)} , \quad (6.16)$$

where the variables are defined similarly as in Equation (6.15). Especially for free-expansion simulations, carefully choosing ρ_{Bref} through the use of interpolations and corrected densities that attempt to negate the inaccuracy of density calculations at the

interface was found to improve the results, such as reducing interfacial spikes and producing interfaces extending closer to the theoretical locations. Several other modifications in the interfacial regions, such as the use of corrected interface coordinates and interfacial pressure gradients as well as refined ways of representing the momentum equation, were required to obtain the best results. In general, the exact prescription of interfacial quantities at the start of a simulation had an influence on the subsequent stability of the results at the interface.

For capturing shocks, a linear combination of the von Neumann and Landshoff artificial viscosities was added to the equations to aid the numerical viscosity in the process and also to allow fine tunings via variation of the artificial-viscosity coefficients. Although a range of the artificial-viscosity coefficients "a" and "b" (see previous chapter) were tested, the author often found the use of combinations such as $a=0.5$ and $b=4.0$ or $a=1.0$ and $b=2.0$ preferable for the cases run. Increasing these coefficient values further did not help too much. Using values such as $a=2.0$ and $b=2.0$ or $a=2.0$ and $b=4.0$ or $a=4.0$ and $b=2.0$ yielded results that did not vary much from one another but were significantly distorted, i.e., excessively diffusive, compared to the $a=1.0$ and $b=2.0$ case.

6.3 Reduction of Numerical Problems

In simulations treating expansion into a vacuum with low (relative to the initial bulk value) or zero applied pressure at the boundary, a density spike followed by noise oscillations was often observed at the interface. In the case of a classical fluid, such features are entirely numerical (although in the case of a plasma, such features may be real¹). The numerical spike can be due to a density discontinuity that exists across the interface even for the case of free expansion because of the finiteness of FSP quantizations and grid-cell resolutions. Discontinuities of variables such as density cause numerical spikes and

oscillations to appear because algorithms are effectively attempting finite-series fits to profiles. Depending upon the type of code, the profiles for some of the other variables could be quite smooth in the interfacial regions. Unless the adiabatic relation was employed, it was possible to observe a smooth match between the externally applied pressure and the fluid pressure across the interface by virtue of the continuity of the sum of the thermal and external pressures. If an adiabatic relation is used, the pressure is calculated from the density, which always has a discontinuity due to the use of computational particles, so there will be a discontinuity in the pressure as well.

Also, inaccuracies in the interfacial velocity calculations contribute to problems. Cellwise quantities are distributed onto individual FSPs by area weighting but this means that for FSPs overlapping the outermost fluid filled cell, the physical quantities assigned to that cell must be used in the area-weighting calculations (see, for example, Equations (5.31) through (5.33) and (5.38) of Chapter 5). We have already discussed the point that this outermost filled cell does not possess very accurate values for the physical variables, and that for the velocity, the value of the adjacent fluid cell is used. For a free expansion or a low applied boundary pressure problem with minimal interface deceleration, this results in a lower-than-real velocity getting assigned to the outermost cell. Not only will the velocity value cut off below the theoretical value at the location of the interface given by the simulation; the observed interface for such simulations will not expand as far as predicted by hydrodynamic theory because there is a density below which PIC cannot treat. Interpolations between finite-sized particles and grid cells do not allow the true spectrum of particle velocities to be obtained in the interfacial region, and mass, that in reality would be spread over several cells at a very low density, piles up short of the real interface location.

Numerical viscosities can also play a role in introducing not only unreal profiles but also retarded interface motions. For free expansions into vacuum, the best results were obtained when numerical viscosity was reduced to a minimum, and even by going as far as subtracting estimated numerical viscosities in finite difference form. However, in the

interest of the suppression of excessive noise and numerical oscillations at, for example, shocks, interfacial regions, or other discontinuities, artificial viscosities were sometimes added even if the standard condition of $\partial u/\partial x < 0$ (see Equations (5.42) and (5.43) of Chapter 5) was not met. In codes employing higher-order algorithms with reduced numerical viscosities, the forceful addition of artificial viscosities at suspect locations was vital to the suppression of the ringing instability; i.e., it was often beneficial to add them also wherever the flow velocity dropped to values that would cause the ringing instability to occur. Such procedures allow codes to run longer before they fail as a result of numerical instabilities growing out of hand, although obviously only at the expense of increased distortion. Unphysical breaks in profiles also appear where the extra viscosities are turned on or off. Very much larger variations in the diffusivity were obtainable by varying the order of the numerical viscosity present in the algorithm than by varying the artificial-viscosity coefficients, especially because raising the artificial-viscosity coefficients too high caused the code to fail. It must be remembered that artificial viscosities are unphysical entities originally developed just for applying where needed to handle shocks. Zeroth-order numerical viscosity (the one-dimensional equivalents of Equations (5.31) and (5.32) of Chapter 5) usually entirely suppressed, or at least substantially reduced the growth rate of, the ringing instability and other numerical oscillations for periods long enough as regards the purposes of this chapter, but this could not be done with artificial viscosities.

Simulations employing the algorithm with first-order numerical viscosity (the one-dimensional equivalent of Equation (5.33) of Chapter 5) and those employing the second-order type (Equation (5.39) of Chapter 5) gave similar results in the bulk, although the latter gave better results in the interfacial pressure rise region. Whether the zeroth-order algorithm was used in the interfacial region or not altered the overall results of the simulation. A code possessing the capability of switching between various orders of numerical viscosities depending upon the region of the flow was considered desirable and was often used in this work, but a break in the profiles appeared at locations where the

order of the numerical viscosity was switched. Such a code used the second-order scheme to obtain profiles with the least diffusive distortion in all regions except where numerical oscillations become large, and there, zeroth-order diffusion was used to suppress noise and instabilities and to thereby prevent premature failure of the code. Under this prescription, the zeroth-order algorithm was often employed in the several cells behind the interface as well as where the fluid velocity fell to values that permit the growth of the ringing instability.

As numerical oscillations at a discontinuity can be reduced by smearing out the discontinuity by viscosities, a logical next step in reducing noise and numerical spikes at an interface or discontinuity was to specify the initial transition region representing the discontinuity or interface to be spread out over several grid zones with a gradual variation of density and pressure values through it. As expected, the interfacial spike in density and pressure could be eliminated, and noise oscillations at discontinuities could be substantially reduced, by spreading the initial transition region over a few cells.

6.4 Discussion of Selected Results

6.4.1 Free Expansion into a Vacuum

The first problem to be discussed is the adiabatic hydrodynamic expansion of a fluid into a vacuum with zero applied pressure at the fluid-vacuum interface. An exact analytical solution to this problem exists in slab geometry. The theory of characteristics provides the solution for a rarefaction-fan region spreading in an initially uniform region in plane isentropic flow.² To quote the results, one obtains the constancy of Riemann invariants along characteristics; i.e.,

$$d\left(u \pm \frac{2c_s}{\gamma - 1}\right) = 0 \text{ along } \frac{dx}{dt} = u \pm c_s . \quad (6.17)$$

It follows from Equation (6.17) that in a rarefaction-fan region in which the rarefaction wave travels in the negative x direction into a uniform region at rest, while satisfying

$$\frac{x - x_{\text{inf}0}}{t} = u - c_s , \quad (6.18)$$

one has

$$c_{s0} = \frac{\gamma - 1}{2} u + c_s , \quad (6.19)$$

where the subscript 0 denotes values for the uniform initial state. Combining the last two equations, one obtains

$$u = \frac{2}{\gamma + 1} \left(c_{s0} + \frac{x - x_{\text{inf}0}}{t} \right) . \quad (6.20)$$

Use of Equation (6.19) and the isentropic relation yields

$$\rho = \rho_0 \left(1 - \frac{\gamma - 1}{2} \frac{u}{c_{s0}} \right)^{\frac{2}{\gamma - 1}} \quad (6.21)$$

$$p = p_0 \left(1 - \frac{\gamma - 1}{2} \frac{u}{c_{s0}} \right)^{\frac{2\gamma}{\gamma - 1}} . \quad (6.22)$$

For a free expansion into vacuum, the rarefaction fan extends out spatially until the density reaches zero. By setting c_s to zero at a vacuum interface in Equation (6.19), one sees that the free-expansion interface, under the hydrodynamic approximation, moves out constantly at a speed of $2c_{s0}/(\gamma - 1)$.

Figures 6.2a and b show, respectively, the theoretical profiles of density and velocity for plane isentropic free expansion of a fluid into a vacuum at the moment the theoretical hydrodynamic interface is situated at $x'=500$. At the initial instant, the fluid uniformly fills the region up to $x'=200$ and is at rest. Once the diaphragm at $x'=200$ is removed, the inward propagating rarefaction wave sets the fluid into motion and a linear velocity profile is established in the disturbed region.

Figures 6.3a and b show the density and velocity profiles obtained from an adiabatic zeroth-order hydrodynamical PIC simulation for the same problem at the same time step as that of Figure 6.2. Figures 6.4a and b show, similarly, the results from a second-order PIC simulation. The difference in the diffusivities between the zeroth- and second-order algorithms is evident. The zigzagging profile in front of the bulk of the fluid is due to computational particles' having broken off the interface. When the second-order result is overlaid on the analytical result, one finds that the plots are indiscernible except at the interfacial region and a narrow region about the rarefaction wave head, where there is some rounding off in the profile of the PIC simulation as a result of numerical diffusion.

In earlier works, such processes as free expansions were often treated just to demonstrate the capability of the PIC code, but here, especially in preparation for running more difficult simulations involving large expansion ratios and requiring pressure matchings across a vacuum boundary, extended period runs with high quality results were strived for.

In a free (or nearly free) expansion into a vacuum, the density ratio between the initial bulk density and the density at the interfacial regions is very large, but as stated earlier, PIC requires a minimal number of computational particles to be present per cell. In an attempt to better treat the low density regions that arise at the interface, the creation of extra computational particles was tested. This was so that the large number of computational particles required per cell for an accurate calculation will always be present even in cells of low density. When the sum of partial FSP lengths (or cross-sectional areas, if in two dimensions) contained in a particular cell dropped below a value equivalent to a few particles per cell, every FSP in that cell was split into several FSPs, each with the same size as before the breakup but lower in mass in such a way that the total mass was conserved. These newly created particles were distributed, slightly offset from each other, about the location of the original particle. If these additional computational particles were created when the density fell below a certain limit, too many particles tended to be

generated. Unfortunately, the creation of new computational particles resulted in even more noise and spurious profiles to appear in the solutions than when without, because cells with particle numbers low enough for particle generation to occur tended to possess very inaccurate data of the physical quantities. Still, the capability of the codes with particle generation to prevent vacuum cells from forming where they shouldn't, is illustrated in Figure 6.5. In this figure, a step function assigned a value of 0.95 in cells for which the FSP density is nonzero and a value of zero in empty cells, is overlaid on the density distribution for free expansion into vacuum obtained by a code with generation of additional particles. It can be seen that no breakoff of particles from the bulk is occurring and that very low densities not treatable without particle generation are being handled.

6.4.2 Expansion Against an Applied Pressure Vacuum Boundary

In this section we examine the problem of a fluid expanding against a region of zero density (i.e., a vacuum) that nevertheless exerts a finite pressure on the fluid-vacuum interface. Here we have a finite density jump across the fluid-vacuum interface, but the total pressure is continuous. The hydrodynamic interface no longer propagates at the free-expansion velocity but moves in such a way that the fluid pressure at the interface is matched with the externally applied interfacial pressure at each instant. This is because finite pressure discontinuities cannot exist across infinitesimally thin, and therefore massless, interfaces across which mass does not flow. The interface will be decelerated and the information of the deceleration will propagate inwards, relative to the fluid flow. The region passed over by this information will develop a distribution of physical variables which deviates from that for free expansion. This region extends inwards from the interface to the rarefaction fan region. The rarefaction fan profile will be the same as for the case of free expansion into a vacuum, except that it will not extend as far out as in the free expansion case, being truncated by the "outer region" just described. That such a profile

will develop after a sudden release of the gas (codes were often supplied with a finite pressure jump across the interface at the initial instant) can be seen easily by employing the theory of characteristics in the $x-t$ plane. In the "outer region" that deviates from free-expansion flow, the pressure and density profiles rise and the velocity decreases as the interface is approached, for the standard cases presented in this chapter. The simulations were able to produce generally correct profiles for all variables with the right wave propagation velocities.

At this point, a remark will be made regarding a difference between the interface dynamics of the thin-shell model and that of the hydrodynamic model. In the thin-shell model, the thin shell at the interface contains mass (in fact, the entire mass of the plasma), and the dynamics of the shell therefore follows Newton's Second Law ($\underline{F} = m\underline{a}$). The only force applying on the shell as a whole is the magnetic pressure force applying against it from the vacuum side, and this finite pressure force thus always causes a deceleration. An exception is the case where the motion of neighboring elements can, through mutual influence, for example, force a shell element to even accelerate against the pressure force, as seen in Figure 4.15. However, in the case of the hydrodynamic model in which a fluid fills a blob, the interface is only a massless boundary separating the fluid from the vacuum fields. Because of this, Newton's Second Law obviously cannot be applied on the interface to determine its dynamics. In fact, as mentioned elsewhere, there will always be a complete match of pressure forces across this massless interface. Now, the external force applying on the plasma is still the repulsive force felt between the surface currents flowing on the plasma interface and the currents through the thruster's field coils. But a finite force acting against the interface from the vacuum field side no longer immediately implies a deceleration of the interface. It is certainly possible for the fluid interface to move outwards at constant velocity, or even accelerate outwards, under the influence of a finite force acting against the fluid at the interface. For a simple example of this, one need only imagine a cylinder filled with a gas and with a moving piston serving as one of its end walls. In fact,

our problem is exactly such a system with the material piston replaced by a magnetic piston which moves in a way such that the value of the fluid pressure applying at the piston surface will vary as a prescribed function of the piston location. It is the interplay between the fluid dynamics (including wave phenomena) within the blob and the externally applied interfacial pressure that now determines the interface dynamics.

Figure 6.6a shows a density profile obtained by an adiabatic code treating the expansion of a fluid against a vacuum that exerts a pressure on the fluid interface according to a prescription of the form given by Equation (6.16) combined with the relation $p \propto \rho^\gamma$. At the initial instant, the fluid uniformly fills the region up to $x'=200$ and is at rest. The value of x_L' in this run was chosen to be 300. The code employed a second-order-accurate algorithm except in the twelve cells behind the interface, where a zeroth-order algorithm was used. In Figure 6.6b, we see the density distribution for the same problem at the same time step as shown in Figure 6.6a but that which is due to a code that used a second-order-accurate scheme everywhere. An investigation of the pressure plots reveals that the interfacial fluid pressure value closely matches the interfacial applied pressure value throughout the simulation. This implies that the interface is being captured correctly. However, the effects of spurious viscosities of different orders and their various combinations are also apparent, and this should be taken into account when interpreting the results of the simulations.

The plots in Figure 6.7 show the pressure profiles at equally spaced time steps computed with a code employing the full energy equation and with an applied interfacial pressure of the form given by Equation (6.15) with x_L' again chosen to be 300. The second-order-accurate algorithm was used everywhere and the artificial-viscosity coefficients were chosen to be $a=1.5$ and $b=2.0$. Due to a difference in the values selected for the initial applied interfacial pressure, the heights of the interfacial pressure peaks in Figures 6.6a and 6.7 are very different. The horizontal lines shown in each plot of Figure 6.7 indicate the values of the externally applied pressures at each of the plotted time steps.

Because of the small local sound speed in the vicinity of the interface compared to the outward flow speed there (which is on the order of magnitude of the sound speed in the initial slab of fluid), the region influenced by the deceleration of the interface has not spread much from the interface even by the end of the simulation.

6.4.3 Observation of the Ringing Instability

In some simulations, the fluid was given a nonzero initial velocity distribution that is zero at $x'=0$ and increases linearly towards the interface. With a uniform initial density distribution, such a flow has a simple solution for the region that has not been affected by waves emanating from the interface. In the isentropic case, if the initial density is given by

$$\rho'_{t'=0} = 1 \text{ everywhere in the fluid} \quad (6.23)$$

and the initial velocity profile is specified as

$$u'_{t'=0} = \frac{x'}{A}, \quad (6.24)$$

where A is a constant, then the solution at subsequent time steps for the region that does not know of the existence of an interface is

$$\rho'(x', t') = \frac{A}{t' + A} \quad (6.25)$$

$$u'(x', t') = \frac{x'}{t' + A}. \quad (6.26)$$

Namely, the density remains spatially uniform although decreases in value with time, and the velocity profile remains linear, with a slope that decreases with time. The individual Lagrangian fluid elements are maintaining their initial velocities.

When these simulations were run, an instability was found to develop in a region far from the interface and rather close to the origin where the flow velocity was low, i.e., far subsonic. It was later realized that this instability could also be observed under similar velocity conditions even if a different density profile, such as one that was due to starting the fluid from rest, prevailed. If simulations of cases with applied interfacial pressure were

run long enough so that the interfacial fluid velocity was decelerated to low enough magnitudes, the instability appeared in the interfacial regions as well. The appearance of the instability, although not its location, was very sensitive to the amount of viscosity present in the code.

In Figure 6.8 we see the result of a simulation of the problem just described with $A=200$ and a second-order-algorithm used everywhere. The density profile is plotted at evenly spaced time steps. By changing the initial velocity distribution, the position and growth rate of the instability could be changed. Reducing the time step did not eliminate the instability nor alter its position of occurrence or growth rate. Furthermore, the location and growth rate of the instability were rather independent of the density, and thus also the number of particles per cell, but the location always coincided with a region where the flow velocity had a particular value, namely, of about $1/6$ th to $1/5$ th of the initial sound speed in these simulations. For example, with $A=200$, the instability started manifesting itself around $t'=30$ to 40 at $x'=40$ to 45 , while with $A=100$, these numbers changed to $t'=45$ and $x'=25$ (the instability amplitude was also smaller and the growth rate lower than when $A=200$), and with $A=67$, they became $t'=55$ and $x'=20$, but the instability did not persist. The instability was found to spread gradually into regions of lower fluid velocity but the oscillations generally cut off sharply at the higher velocity end of about $1/5$ th of the initial sound speed. Careful observations revealed that the noise from which the instability grows out of was arranged into periodic standing-wave pattern beats throughout the fluid region and that these ripples do not grow except where the above-mentioned flow velocities prevailed. The small-amplitude noise oscillations (not visible on the scale of the plots presented), which eventually organize themselves into beats were such that the adjacent peaks and troughs were located in neighboring cells, indicating a coupling phenomenon. The periodic structure was washed away when artificial viscosities were introduced.

Instability analyses of the small perturbation type did not show these instabilities, and artificially inserted modulations of fluid profiles tended to get damped out, indicating a

stable flow process. However, all properties of the ringing instability were met by the instability observed in these simulations.

6.4.4 Extended Period Runs

As noted earlier, the parameter ratios used in the slab-geometry simulations differ very much from those associated with flows in the thrusters of ICF pulse rockets, especially in that the slab-geometry simulations dealt with expansion ratios very much smaller than those characteristic of real thruster flows. Also, with real multidimensional flows, there is an additional freedom regarding flow directions, and the details of the flows treated by the slab-geometry simulations will be different from those of flows in the thrusters of ICF pulse rockets. Nevertheless, the generic physics of the bulk flow phenomena will be similar between the two, and attempts were made to run several simulations of the problem involving a vacuum interface at which a finite pressure is applied, until long after the interface reaches its maximum expansion distance for the first time.

In these simulations, the deceleration of the interface causes a shock to develop as the result of a steepening of a compression wave emanating from a region close behind the interface. The captured near-discontinuity representing the shock was found to satisfy approximately the correct jump conditions and to move at roughly the correct speed relative to the fluid flow, and its thickness hardly changed with time. This inward-facing shock marks the inner boundary of the region influenced by the deceleration of the interface. As expected, a temperature rise was observed in this region due to conversion of some of the kinetic energy of the flow into thermal energy. Because of shock heating, this temperature rise is higher than would be possible under an isentropic calculation (although such a remark has some problems because a situation in which a shock develops can not be treated by an isentropic analysis, especially through the shock).

A large portion of the kinetic energy decrease experienced by the fluid is due to its doing work against the externally applied interfacial pressure. Eventually the interface comes to a halt, and the interface was generally observed to stay nearly stationary for a very long time (in terms of the characteristic time scales of the initial expansion process) around this maximum expansion position, after which it would start moving backwards, only to move outwards again at a much later time. A closer observation revealed that the interface was actually oscillating with a small amplitude during the time it appeared to be almost stationary. Even at its closest approach to the coordinate x_L' , the interface was still several grid-cell widths away from that coordinate, and inaccuracies due to coarse grid resolutions were not notably affecting calculations of the interface dynamics. Plots clearly showed that the internal fluid dynamics such as the reflection of waves from the center and the interface, and rethermalizations play a major role in determining the motion of the interface, contrasting dramatically from the case of thin-shell simulations, which have no provisions for taking into account such hydrodynamical phenomena.

Discussed in this section will be simulations for two different initially applied interfacial pressure values and two different prescriptions for the rate at which the applied interfacial pressure rises with interface coordinate. When comparing the performance of different designs of thrusters for ICF pulse rockets, one must study the influence of various parameter ratios on the plasma flow. Although planar-geometry flows differ from the multidimensional flows that will be present in the thrusters, it was considered worthwhile to study, in slab geometry, the generic influence different prescriptions of interfacial pressure variation have on the fluid flow. The various pressure prescriptions were chosen to model different thruster designs crudely, such as low-field and high-field. Codes were also tested with different orders of numerical viscosities and various coefficients of artificial viscosities and combinations thereof.

In simulations employing the higher applied interfacial pressure, the interface was found to decelerate notably from early time steps and come to a halt a large distance away from x_L' , while in simulations with the lower applied interfacial pressure, the interface expanded out to larger distances with less deceleration and approached x_L' very closely, whereupon it experienced a very sudden and large deceleration. These observations are similar to the conclusions drawn with the two-dimensional thin-shell simulations. The recoil and oscillatory motion of the interface was more pronounced (e.g., larger amplitude of oscillation) for the case with higher interfacial pressure. The recoil of the interface became more evident with the case of lower applied pressure (this is not an exactly appropriate nomenclature because as the interface approaches its maximum expansion location, the interface pressure for the “lower applied pressure case” rises to extremely high values) when x_L' was increased so that a larger expansion ratio was obtained. The profiles of the physical quantities also exhibited marked differences between the high and low applied pressure cases. For example, in the higher applied pressure case, rather gently sloping profiles of fluid variables were obtained, but in the case of lower applied pressure, a shock propagated inwards from the interface as the interface suddenly decelerated to a halt. The compression wave in the former case did not develop quickly into a shock. And generally, simple considerations show that the shock will form closer to the interface for the case of lower applied pressure.

In all of the following simulations, the fluid in the initial state is uniform (with $\rho'=p'=1$) and at rest, with the fluid-vacuum interface situated at $x'=200$. In Cases 1 through 3, p_{Bref}' was set to 0.03125, while in Case 4, it was set to approximately 0.4159. The first three cases represent situations of low applied interfacial pressure, and the last case represents a situation of high applied interfacial pressure. All simulations were run with codes employing the full energy equation.

Parenthetically, it may be remarked that the situation of having a perfectly conducting plasma region separated permanently from a vacuum magnetic field region, and with the interface undergoing an oscillation forever (under the idealized assumption of no damping mechanisms present) is unrealistic. Under the hydrodynamic approximation taken here, a stationary state is possible if an entirely uniform fluid at rest is pressure matched with the applied external pressure at the interface. However, thermodynamically, an equilibrium state is one which is homogeneous, isotropic, and stationary. The state of a perfectly conducting uniform plasma at rest, bounded by a vacuum magnetic field is clearly not an equilibrium state, and the plasma will slowly diffuse across the fields until a truly uniform and stationary equilibrium state is reached. In reality, fluids which do not allow external magnetic fields to diffuse into them do not exist. That plasmas with the resistivity term set to zero can cross fields may even be seen (with a little work) using the simplified equations of MHD theory, if more than the minimal number of terms are retained.

6.4.4.1 Case 1

In this run, x_L' was set to 300 and the second-order-accurate algorithm was employed everywhere except in the 25 cells closest to the interface where the zeroth-order algorithm was used. Artificial viscosity was added in the standard way, i.e., where $\partial u/\partial x < 0$, with $a=0.5$ and $b=4.0$.

Figures 6.9a, b, c, and d show, respectively, the nondimensionalized pressure, density, velocity, and temperature distributions at evenly spaced time steps. The nondimensional temperature was defined as the quotient of the nondimensional pressure and density, i.e., p'/ρ' . The velocity profile was plotted with the outermost filled cell and all empty cells to its right assigned the same velocity as that of the second from outermost filled cell. An inward-facing shock is clearly observed with a density, pressure, and temperature increase, as well as a velocity decrease, on its downstream side. We see a

drastic deceleration of the interface and the flow behind it. The spurious break in the profile roughly 25 grid-cell zones inwards from the interface is due to the switch in the diffusivity of the algorithm at that location. After 2000 time steps (where each time step was set equal to 0.1 unit of nondimensionalized time for all cases 1 through 4), energy conservation was still holding to within 1 %. The development of the ringing instability caused the code to fail rather early.

6.4.4.2 Case 2

In this run, all parameters were kept the same as in Case 1 except that for zeroth-order numerical viscosity was used everywhere for the sole purpose of strongly suppressing numerical instabilities and thereby allowing the simulation to run very extended periods at the expense of a very diffusive profile.

Figures 6.10a, b, and c show, respectively, the nondimensionalized pressure, density, and velocity distributions at evenly spaced time steps, except for a change in time interval between plots at $t'=1000$ and $t'=2000$. The plots are showing the reflection of shock waves and rarefaction waves from both the center ($x'=0$) and interface, as well as the interaction among them. Although not entirely evident upon first glance unless one looks at the velocity plots, a small-amplitude oscillation of the interface coordinate is occurring. If the code is run this long, even zeroth-order numerical diffusion is not sufficient to suppress numerical instabilities at late time steps.

6.4.4.3 Case 3

In this run, all parameters were kept the same as in Case 1 except that x_L' was changed to 500 and zeroth-order numerical viscosity was employed everywhere.

Figures 6.11a, b, and c show, respectively, the nondimensionalized pressure, density, and temperature distributions at evenly spaced time steps. Note that compared to the case in which x_L' was 300 and the expansion ratio smaller, the amplitude of the interface oscillation is much larger.

Figures 6.12a, b, c, and d show, respectively, the normalized (with respect to area) total internal energy of the fluid slab, the sum of the normalized total internal and kinetic energies of the fluid slab, the normalized cumulative work done by the fluid against the externally applied interfacial pressure, and the normalized energy of the entire system (fluid plus the external force agent), all as functions of time. The plot of the fluid's internal energy shows the correct features of decrease with time as the internal energy is converted into kinetic energy of the expansion, and subsequent increase that is due to rethermalization arising from interface deceleration, with a significant contribution coming from shock heating. The plot of the cumulative work done by the fluid reflects the history of the interface motion. A dip in this curve indicates an inward motion of the interface, and vice versa. The cumulative work can be considered as "excess" energy stored in the external force field applying the pressure at the interface, and was set to zero when the interface was at its initial position. Its sum with the combined kinetic and internal energies of the fluid is a conserved quantity in the absence of losses. The final figure (Figure 6.12d) therefore shows the accuracy with which energy conservation is being satisfied by the code.

6.4.4.4 Case 4

This case is the same as Case 1 except for the use of a different p_{Bref}' value and the employment of zeroth-order numerical viscosity not only in the 25 cells nearest the interface but also where $u' < 0.3$ (i.e., where $u < 0.3c_{s*} = 0.3c_{s0}$).

A stronger deceleration of the interface than observed in the weaker applied pressure cases above is clearly reflected in the velocity profiles during the early phases of

expansion. Namely, the velocity is significantly lowered at the interface instead of continuing to rise towards the interface. The interface is also exhibiting a large-amplitude oscillation, quite unlike in the case of the lower applied pressure runs.

Figures 6.13a, b, c, and d show, respectively, the nondimensionalized pressure, density, velocity, and temperature distributions at evenly spaced time steps, except for there is a change in the time interval between plots at $t'=100$. Figures 6.13a', b', c', and d' show the results of a run (for one time step only) differing only in that zeroth-order numerical viscosity, in addition to being applied where $u' < 0.3$, was employed in the 5 (instead of 25) cells nearest the interface. The difference in the results between these two runs illustrates the unreliability of the simulations at late time steps due to distortions arising from nonuniformly applied excessive numerical viscosities.

Figures 6.14a, b, c, and d show, respectively, the internal energy of the fluid slab, the sum of the internal and kinetic energies of the fluid slab, the cumulative work done by the fluid, and the energy of the entire system, all in normalized units, as functions of time.

6.4.5 Runs with Larger Expansion Ratios

In the flow of plasma in a thruster of an ICF pulse rocket, the expansion ratio experienced by the plasma is enormous because it expands from a sphere with a radius on the order of a few centimeters to a distended blob with a characteristic dimension of a few to a few tens of meters. As alluded to earlier, the maximum expansion ratio treatable depends upon the number of particles employed and the mass distribution profiles that develop. If one were not to use a nonuniform, moving grid, this will require an enormous number of computational particles to be employed as well as a very low number of cells to be filled by particles at the initial instant. But as expected, the quality of the results

obtained by the simulations were not very good when only a few cells (as low as 3 were tested) were occupied by particles at the initial instant.

Also, the initial ratio of the bulk fluid pressure to the externally applied interfacial pressure is quite high for the rocket thruster situation, but most of the initial internal energy of the plasma is converted into kinetic energy very early on in the expansion process, and simple estimates show that not long after expansion starts, the interfacial pressure will be much higher than the average bulk fluid pressure behind the interfacial region. The slab-geometry codes were tested with enormous interfacial pressure mismatches at the initial instant, but were found to be capable of quickly converging the interfacial fluid pressure upon the externally applied interfacial pressure. Keeping track of applied interfacial pressures as low as one part in 10^4 of the initial bulk fluid pressure posed no problems.

Figures 6.15a through c show the results of a simulation in which a slab of fluid, initially at rest with uniform density and pressure ($\rho'=1$, $p'=1$) and occupying the region $0 \leq x' \leq 30$, expands against an externally applied interfacial pressure that maintains a value of $p'=0.285$ regardless of the interface position. Each bulk fluid cell was filled with 500 particles initially and the initial interface was smeared over 4 cells with a gradual variation of physical variables through the transition region. Profiles are plotted every 10 dimensionless time units for nondimensionalized pressure (Figure 6.15a), density (Figure 6.15b), and velocity (Figure 6.15c). Modestly high expansion ratios on the order of a thousand were tested without any problems in similar runs, and smooth profiles with good pressure matchings across the interface were obtained.

6.4.6 Shock-Tube Problems

The shock-tube problem involving the breakup of an arbitrary discontinuity serves as an ideal problem by which to test a hydrodynamic code because it requires the treatment of shocks, contact surfaces, rarefaction waves, and the reflection of waves off end walls

and mutual wave interactions. Not only is testing the capabilities of a code in handling these fluid phenomena important; the correct treatment of shocks is a must in flows with a decelerating interface such as flows in the thrusters of ICF pulse rockets. The codes for treating shock-tube problems will differ from those employed in the examples above because fluid will fill both sides of the discontinuity and because there will be an absence of vacuum boundaries.

In an arbitrary discontinuity problem, the two states on both sides of the initial discontinuity do not in general satisfy the Rankine-Hugoniot relations :

$$\frac{\rho_1}{\rho_2} = \frac{(\gamma - 1)\frac{p_2}{p_1} + (\gamma + 1)}{(\gamma + 1)\frac{p_2}{p_1} + (\gamma - 1)} \quad (6.27)$$

$$u_1^2 = \frac{p_1}{2\rho_1} \left\{ (\gamma - 1) + (\gamma + 1)\frac{p_2}{p_1} \right\} \quad (6.28)$$

$$u_2^2 = \frac{p_1}{2\rho_1} \frac{\left\{ (\gamma + 1) + (\gamma - 1)\frac{p_2}{p_1} \right\}^2}{\left\{ (\gamma - 1) + (\gamma + 1)\frac{p_2}{p_1} \right\}}, \quad (6.29)$$

where the subscripts 1 and 2 refer to the two states, one to each side of the shock, and the velocities are measured in the shock's frame. Therefore, the two states will be connected by a combination of a rarefaction-fan region, a contact surface, and a shock. Depending upon the arbitrary discontinuity, two shocks or two rarefaction waves may propagate away. The contact surface, under approximations of no molecular diffusion, continues to separate material that was on one side of the initial discontinuity from that which was on the other side. It manifests itself only as a discontinuity in density and temperature and not in pressure or velocity. The contact surface will exist in the region between the shock and the rarefaction fan (or between the two shocks or rarefaction fans), which is a region in a stationary state with uniform pressure and velocity.

One of the problems investigated was that with a jump in the pressure and density across the initial discontinuity but not in the temperature (Case 1). The fluid was taken to

be at rest and uniform on either side of the discontinuity. With this setup, the density drop across the contact surface is towards the shock. Plots of the temperature profile correctly displayed a region of temperature rise between the shock and contact surface. Plots of entropy showed very high entropy production at the shock but virtually none elsewhere, as it is in reality.

Among the other arbitrary discontinuity problems tested was that with an initial density profile that is uniform everywhere but with a pressure jump across the initial discontinuity (Case 2). Here again, the initial fluid was taken to be at rest everywhere. In this case, the density across the contact surface rises towards the shock by just the right amount to match the value on the downstream side of the shock. The temperature decreases across the contact surface towards the shock.

For both problems, the PIC simulations gave results that were in very good agreement with theory. The jump conditions and wave velocities were all obtained with good accuracy. For example, the pressure ratio across the shock calculated via the Rankine-Hugoniot relations using the density ratio obtained from the plots, agreed to within a few percent of the corresponding pressure ratio read off the plots. The velocity of the flow into, and out of, the shock, as calculated from the Rankine-Hugoniot relations using the pressure ratio read off the plots, came to within the order of a percent of the velocities measured off the plots. Similarly, the distributions of variables across the rarefaction-fan region, as obtained from the simulations, showed equally accurate agreements with theory, and the variations of quantities across the contact surface were also correctly obtained.

The tradeoff between an increase in the widths of near-discontinuities, brought about through the use of higher amounts of viscosity or wider widths of the initial transition region, and a reduction in the noisy oscillations at those discontinuities was investigated. Prescribing the initial near-discontinuity to be spread over several (generally 2 to 10) grid cells substantially reduced unfavorable numerical features such as postshock

oscillations or spurious dips and spikes at the contact surface and at the edge of the rarefaction-fan region. When interparticle distances were varied gradually through the initial transition layer, even if only one cell wide, improved and less diffusive results were obtained over cases in which the initial particle separations were varied in a stepwise fashion from cell to cell. Also, decreasing the jump in quantities across the initial discontinuity helped reduce noise at subsequent time steps. Although specifying the initial discontinuity as a gradual transition region several cells wide had large effects, especially in controlling the near-discontinuity's thickness and the associated noise features, artificial viscosity had good control over noise at the shock. Varying the artificial-viscosity coefficient values from, for example, $a=0.5$ and $b=4.0$ to $a=1.5$ and $b=0.5$, could literally remove postshock oscillations, but other spurious features appeared and long-wavelength distortions worsened. As alluded to earlier, one choice of coefficients that worked quite favorably was roughly $a=1.0$ and $b=2.0$ to 2.5 , and many of the arbitrary discontinuity simulations were run with $a=0.85$ and $b=2.5$. When the combinations of artificial-viscosity coefficients, "a" and "b," which gave successful simulations for the problem, were examined in parameter space, they very approximately fell in the region given by $a \geq 0.16b$ on average.

Again, energy was found to be extremely well conserved in the simulations. In a typical run, the combination of the internal and kinetic energies summed over the fluid slab was conserved to within roughly 0.007 to 0.008 % in 1000 time steps (equivalent to 100 nondimensionalized time units).

6.4.6.1 Case 1

In this simulation, the fluid fills the region $0 \leq x' \leq 400$ and is bounded by rigid walls at both ends. Initially there is a diaphragm at $x'=200$, and the nondimensionalized densities to its left and right are 1.0 and 0.2, respectively, and the nondimensionalized

pressures, 1.0 and $0.2^{5/3}$, respectively. Fifty particles per cell were used to represent a dimensionless density value of one. Increasing the particle number per cell above this value hardly improved the results. The initial discontinuity was spread over two cells, and a second-order-accurate algorithm was used everywhere. Artificial-viscosity coefficient values of $a=0.85$ and $b=2.5$ were used, and artificial viscosity was added not only where $\partial u/\partial x < 0$, but also where large gradients in pressure and density were present **** for the sole purpose of suppressing numerical oscillations, even at contact surfaces. The codes were generally set up so that the widths of the captured shocks and contact surfaces, which hardly change during the simulation, will be a few grid cells wide.

In Figure 6.16a, the nondimensionalized pressure profile is plotted at $t'=0$ and 100. Figures 6.16b through d show profiles of nondimensionalized density, velocity, and temperature in similar fashion (note that in some of these and the following plots, the lowest ordinate value plotted is not necessarily zero). In these plots, we see a shock advancing into a uniform undisturbed region to the right, and in the uniform flow region to the left of the shock, we observe the presence of a contact discontinuity. The true location of a contact surface, when it is spread over several cells, is at the high density end of the transition region. The uniform flow region ends at the tail of the rarefaction-fan region, whose head is propagating into the still undisturbed uniform region to the left.

Vast improvements in the results are being obtained over cases in which the artificial viscosity is added only where the standard criterion is met (i.e., where $\partial u/\partial x < 0$) and in which the initial discontinuity is not spread over multiple cells.

**** Care must be taken when adding artificial viscosities in nonstandard ways such as this.

6.4.6.2 Case 2

This run treats the situation in which, initially, the density is uniform across the entire shock tube spanning $0 \leq x' \leq 400$ but the pressure is discontinuous across $x'=200$ with a pressure ratio of 5 to 1. The initial discontinuity was not spread over multiple grid-cell zones at the initial instant, and artificial viscosity was added only under the standard condition. Artificial-viscosity coefficients of $a=0.5$ and $b=4.0$ were used. Other parameters for the code were chosen to be the same as those for Case 1.

Figures 6.17a, b, c, and d show, respectively, the profiles of nondimensionalized pressure, density, velocity, and temperature, at $t'=5$ and 85. Again, we see the correct profiles develop for each variable, although they are qualitatively different from those of the previous case. From the right, we see an undisturbed region, a shock, a uniform flow region, a contact surface, another uniform flow region, a rarefaction-fan region, and at the very left, another undisturbed region.

6.5 Summary

For the hydrodynamic calculation of flows in the thrusters of ICF pulse rockets, the PIC method was chosen. This was because for this type of flow, the capture of an unknown fluid-vacuum interface with a dynamic motion is required, and PIC was considered to be capable of this task. But before working on the two-dimensional flows in realistic thruster geometries, it was first felt necessary to develop the basic techniques of interface capturing and to test the capturing of shocks via numerical and artificial viscosities, both in the simpler slab geometry. The capability of the algorithm to handle other relevant flow phenomena was also tested. Also, in slab geometry, the important issue of checking the accuracy of the algorithm against exact analytical results was

undertaken, and the limits of the capability of the algorithm were investigated. Spurious numerical features of the results were identified. This was all carried out in preparation for the development of a more complex two-dimensional code for treating realistic thruster flows.

The Nishiguchi-Yabe algorithm was retested in this work for a variety of flow phenomena of relevance to flows in the thrusters of ICF pulse rockets. Generally, the quality of the results could be made very good. For planar free expansions into a vacuum, the second-order-accurate scheme gave results that were hardly discernible from analytical results. The first-order-accurate algorithm gave results that were, in general, quite comparable in quality to those that were due to the second-order-accurate algorithm, and the second-order algorithm was not considered an absolute necessity, at least for the type of problems involved here. The zeroth-order-accurate algorithm was visibly diffusive, but nevertheless proved useful in suppressing numerical instabilities which otherwise caused premature code failure when attempts were made to run simulations of flow processes vulnerable to the development of such instabilities for long periods of time. For example, the ringing instability turned out to be dangerous when low-diffusivity algorithms were used. As always, there was a tradeoff between decreased numerical oscillations and increased diffusiveness of the profiles. In order to obtain highly accurate results with minimal numerical diffusion but yet be able to run the code extended time steps without being plagued by the development of large-amplitude numerical instabilities, it was found convenient to use different order algorithms for different regions of the flow. However, a break in the profiles of the physical variables appeared where a switch in the orders of the algorithm was made, and had to be tolerated

Shock and contact-surface capturing proceeded in good agreement with theory. The use of artificial viscosity was found to be helpful here, and ranges of artificial-viscosity coefficients that worked well were investigated.

Next, PIC is not capable of handling densities that are very low compared to the initial density, should they arise during a simulation, unless nonuniform, moving grids and/or the generation of additional computational particles (under the constraint of conserved total mass) are employed. This is because PIC requires a minimum number of computational particles to be present per cell for the calculation to proceed successfully. However, when particle generation was attempted at low density interfaces, the results were not very good, with very high noise, and this was attributed to the low quality of physical variable data in the low density interfacial cells.

The expansion of a perfectly conducting fluid against a vacuum magnetic field is one in which the fluid interface faces a zero density vacuum that nevertheless exerts a finite pressure upon the interface. It is believed that this type of problem has not been treated in the past using a fully hydrodynamic approach for explosive flow situations. In this problem, a pressure match between the fluid pressure and the magnetic pressure must be met at each instant of time across the interface. This requires interface capturing unless interpolations are carried out at each time step to look for the interface, which can become impractically difficult in multidimensions. In this chapter, the external pressure applied at the interface was supplied through a simple prescription, because calculating it as a magnetic pressure, using Maxwell's Equations, was not necessary when the setup itself was far removed from that of any real thruster. Concerning the automatic capture of the interface, the algorithm was found to be very capable of the task. Capturing was accomplished by suitable prescription of boundary conditions at the interface. Basically, the standard prescription of applying the external pressure at the cell boundary was used. However, because of the use of finite resolution grid cells, and to some extent, the use of a far lower number of finite-sized computational particles than there are real particles, the fluid quantities computed (from particle data) in a cell adjacent to a vacuum cell were very inaccurate when the interface motion was normal to the interface (which is always the case

in a one-dimensional flow unless the interface is at rest). Therefore, such "outermost fluid filled cells" were left out of the calculations and the next fluid cell was considered as the effective interface cell. Also, because of the noise inherent in the profiles of the fluid variables at the interface, it was found that the interface fluid pressure was best determined by observing the plots with the eye. Following these rules, a good match between the interfacial fluid pressure and the applied pressure was obtained at all time steps even if the pressure mismatch at the initial instant was large, and even if the applied pressure prescription was purposefully designed to make pressure matchings difficult. This indicated a robust and accurate interface-capturing capability of the algorithm.

Simulation results showed that with interfacial pressure prescriptions resembling those of a low-field thruster, the interface expanded with little deceleration until it reached a position close to its turnaround point, upon which it experienced sudden deceleration. With a "high-field" type pressure prescription, the interface deceleration proceeded gradually (and noticeably) from earlier in the expansion phase, and the interface did not expand as much (as in the "low field" case) either. These qualitative observations are similar to those made with the thin-shell simulations of Chapter 4, and also with those from the two-dimensional hydrodynamic simulations of Chapter 7, indicating the commonality in the behavior of the interface under the different approximations and situations. However, the interface was found to remain nearly stationary at its maximum expansion location for very long periods of time (on the time scales of the expansion) for these one-dimensional planar geometry flows, when the expansion ratio was small and a large amount of mass was present in the interior regions (and not just all near the interface). With a high amount of mass contained in the interior regions (in these runs for small expansion ratio flows) and with an expansion velocity not significantly higher than the sound speed, the fluid dynamics of the interior had a substantial influence on the motion of the interface.

References for Chapter 6

1. For the case of a plasma, such features can be partly real. See Ch. SACK and H. SCHAMEL, "*Plasma Expansion into Vacuum - A Hydrodynamic Approach*," *Physics Reports* **156**, #6 (Dec., 1987).
2. See, for example, Ya.B. ZEL'DOVICH and Yu.P. RAIZER, *Physics of Shock Waves and High-Temperature Hydrodynamic Phenomena, vol. 1* (Academic Press, New York, 1966).

Appendix of Chapter 6

Figures for Chapter 6

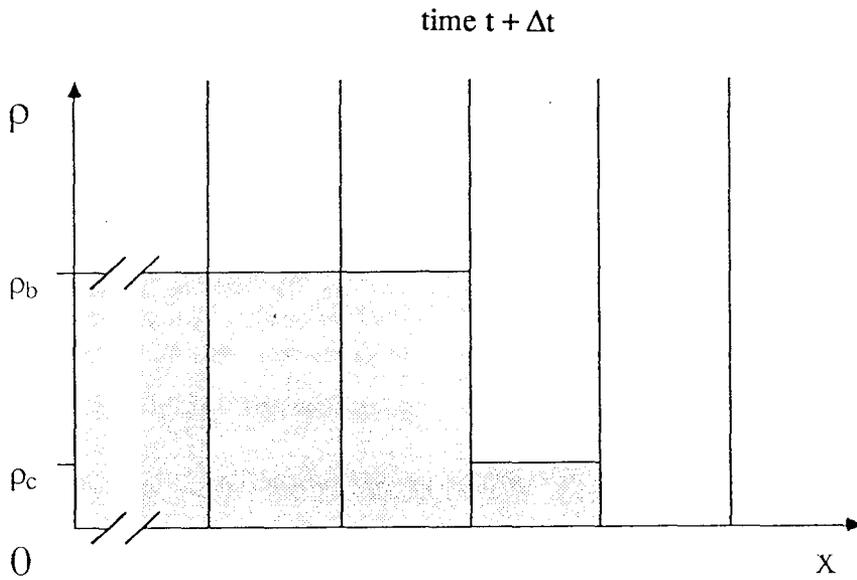
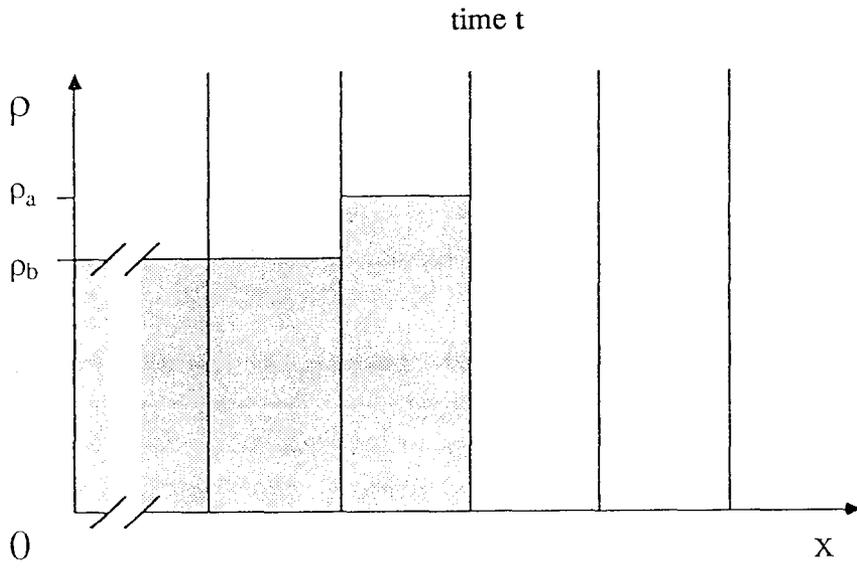


Figure 6.1

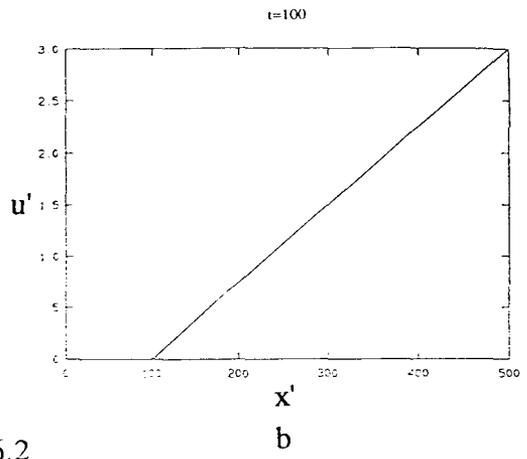
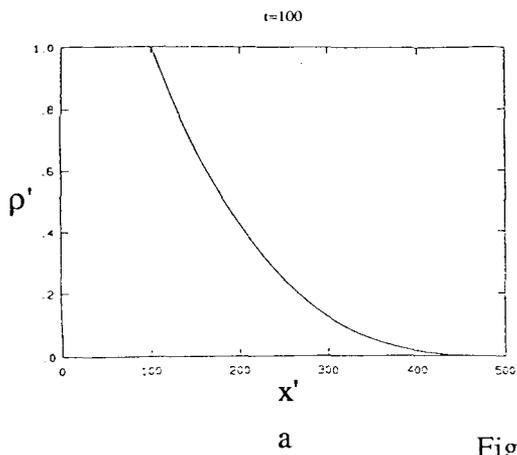


Figure 6.2

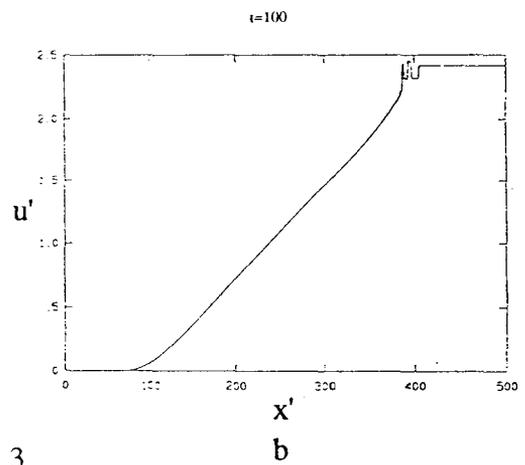
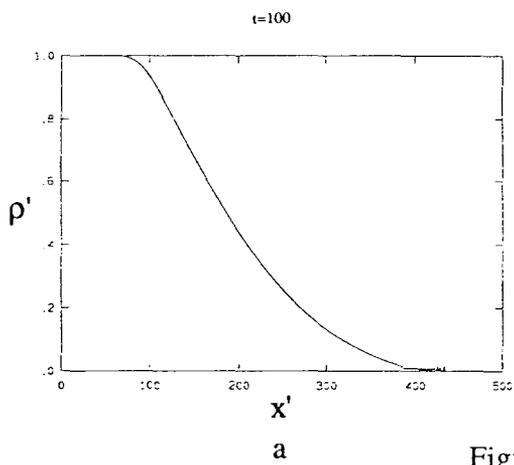


Figure 6.3

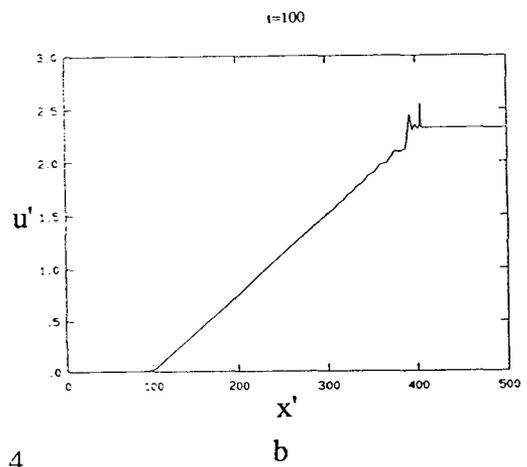
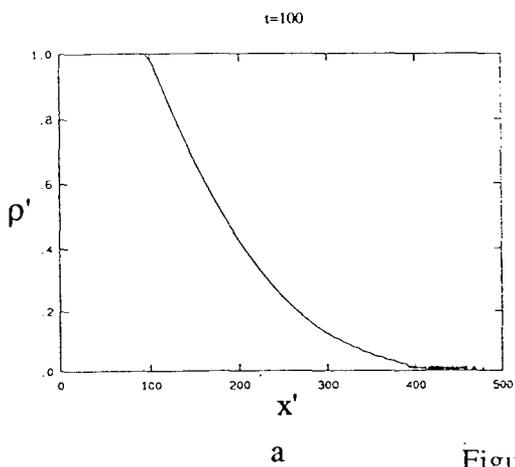


Figure 6.4

t=85

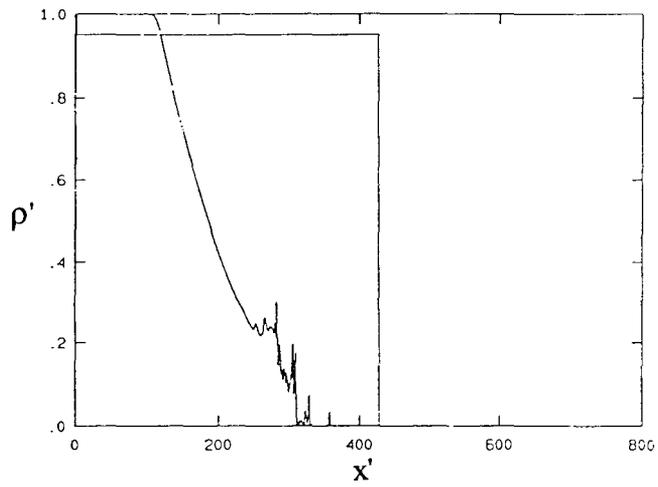


Figure 6.5

t=100

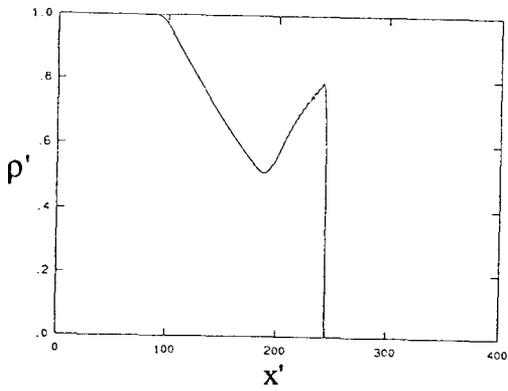


Figure 6.6a

t=100

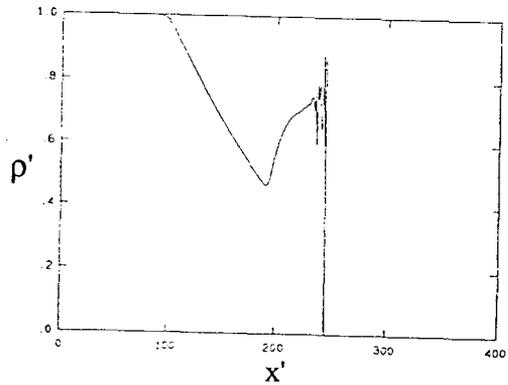


Figure 6.6b

VI-A5

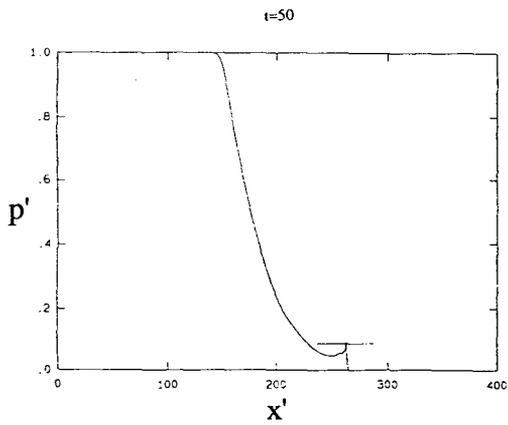
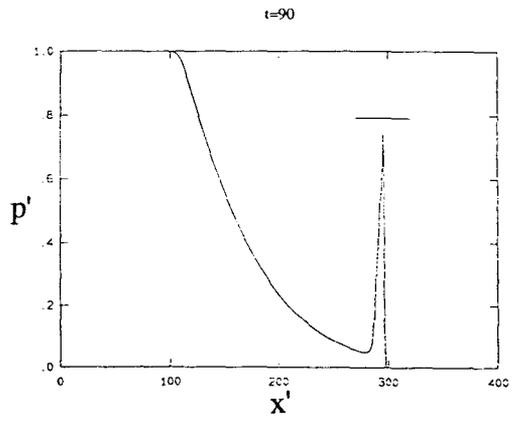
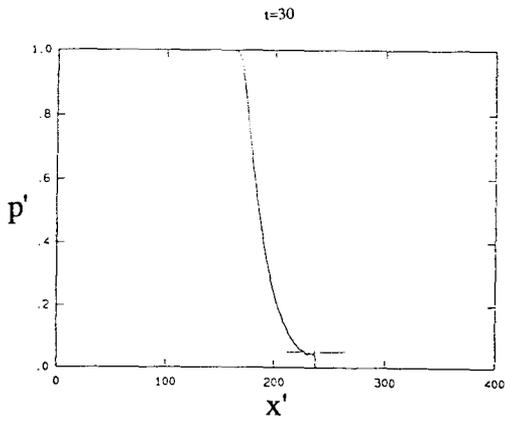
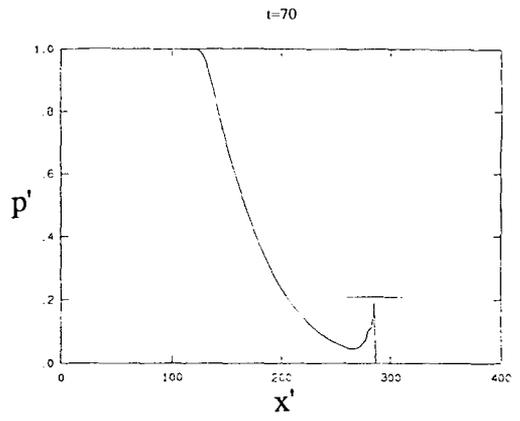
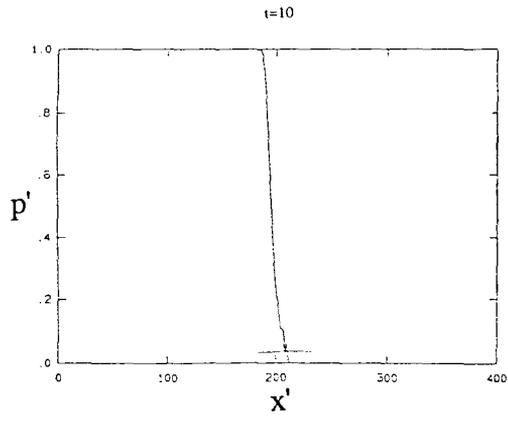


Figure 6.7

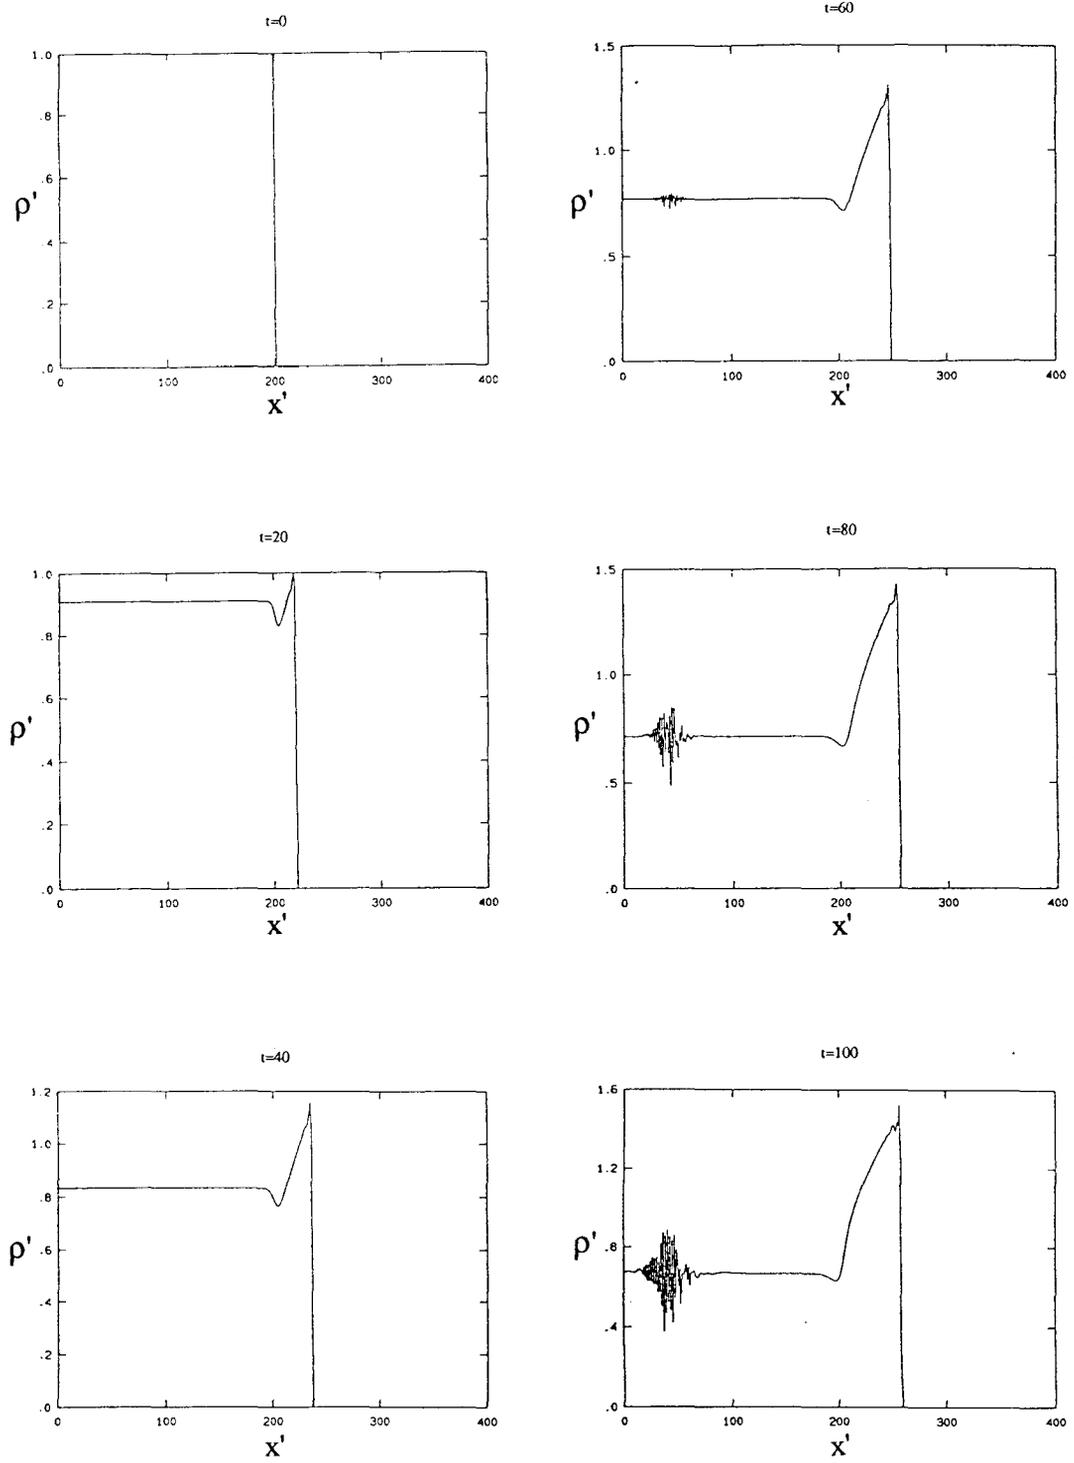


Figure 6.8

VI-A7

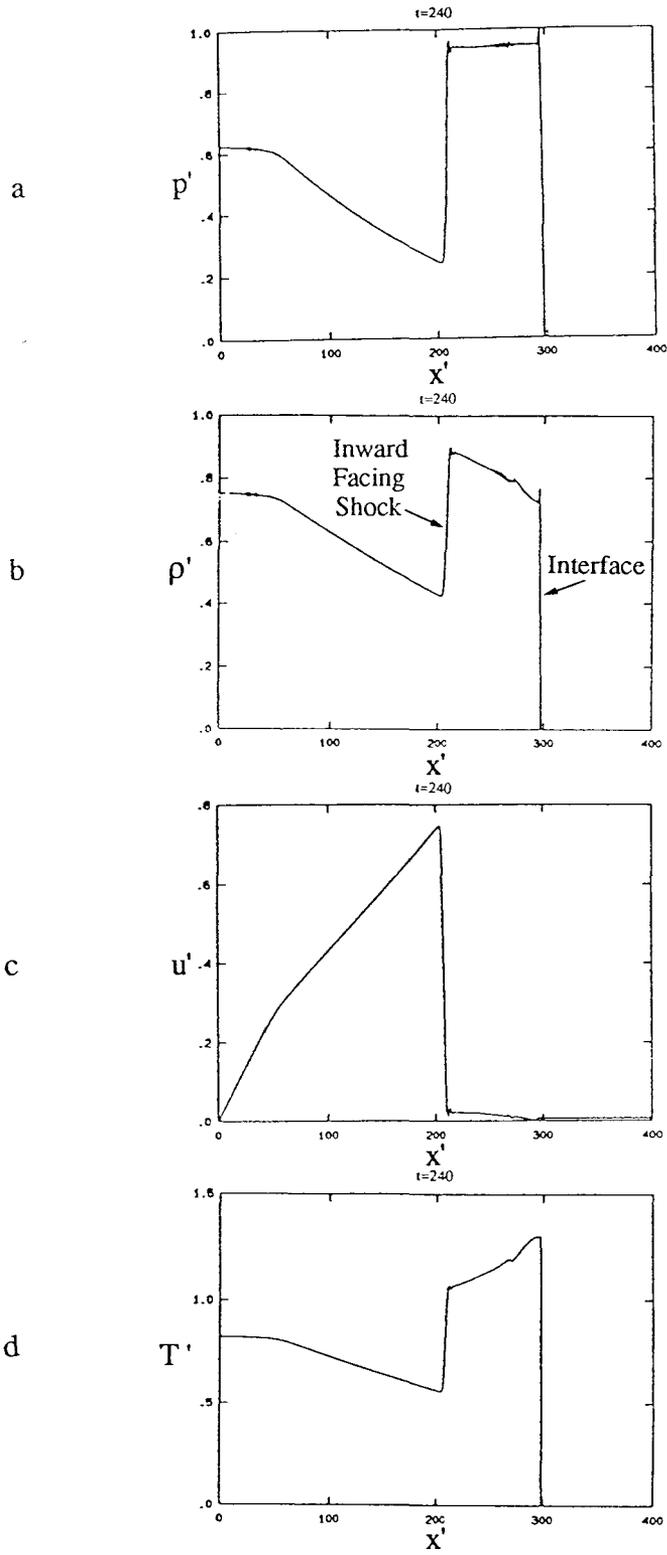


Figure 6.9

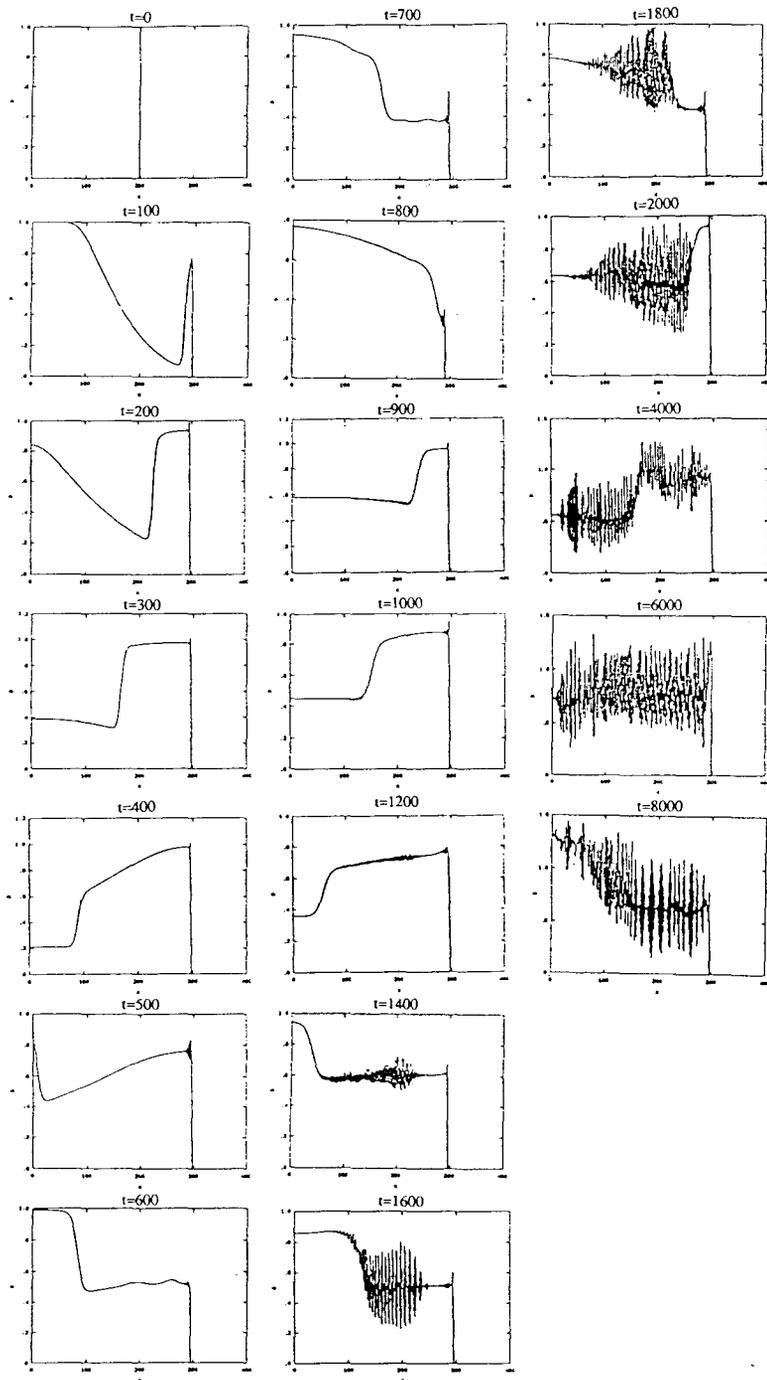


Figure 6.10a Nondimensionalized pressure (p') v.s. nondimensionalized x-coordinate (x').
 (vertical axis) (horizontal axis)

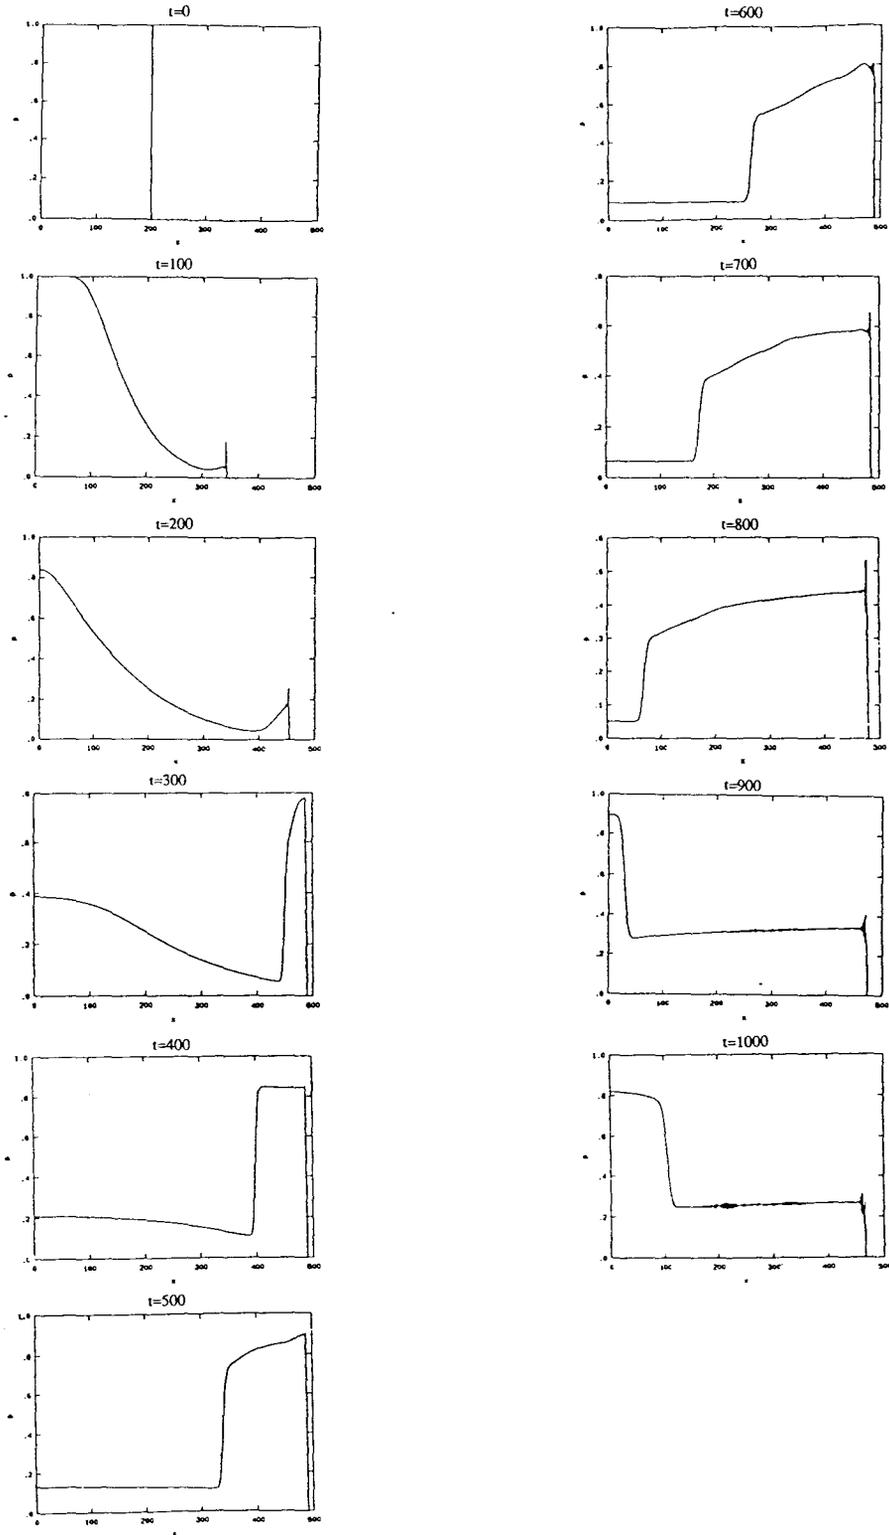


Figure 6.11a Nondimensionalized pressure (p') v.s. nondimensionalized x -coordinate (x').
 (vertical axis) (horizontal axis)

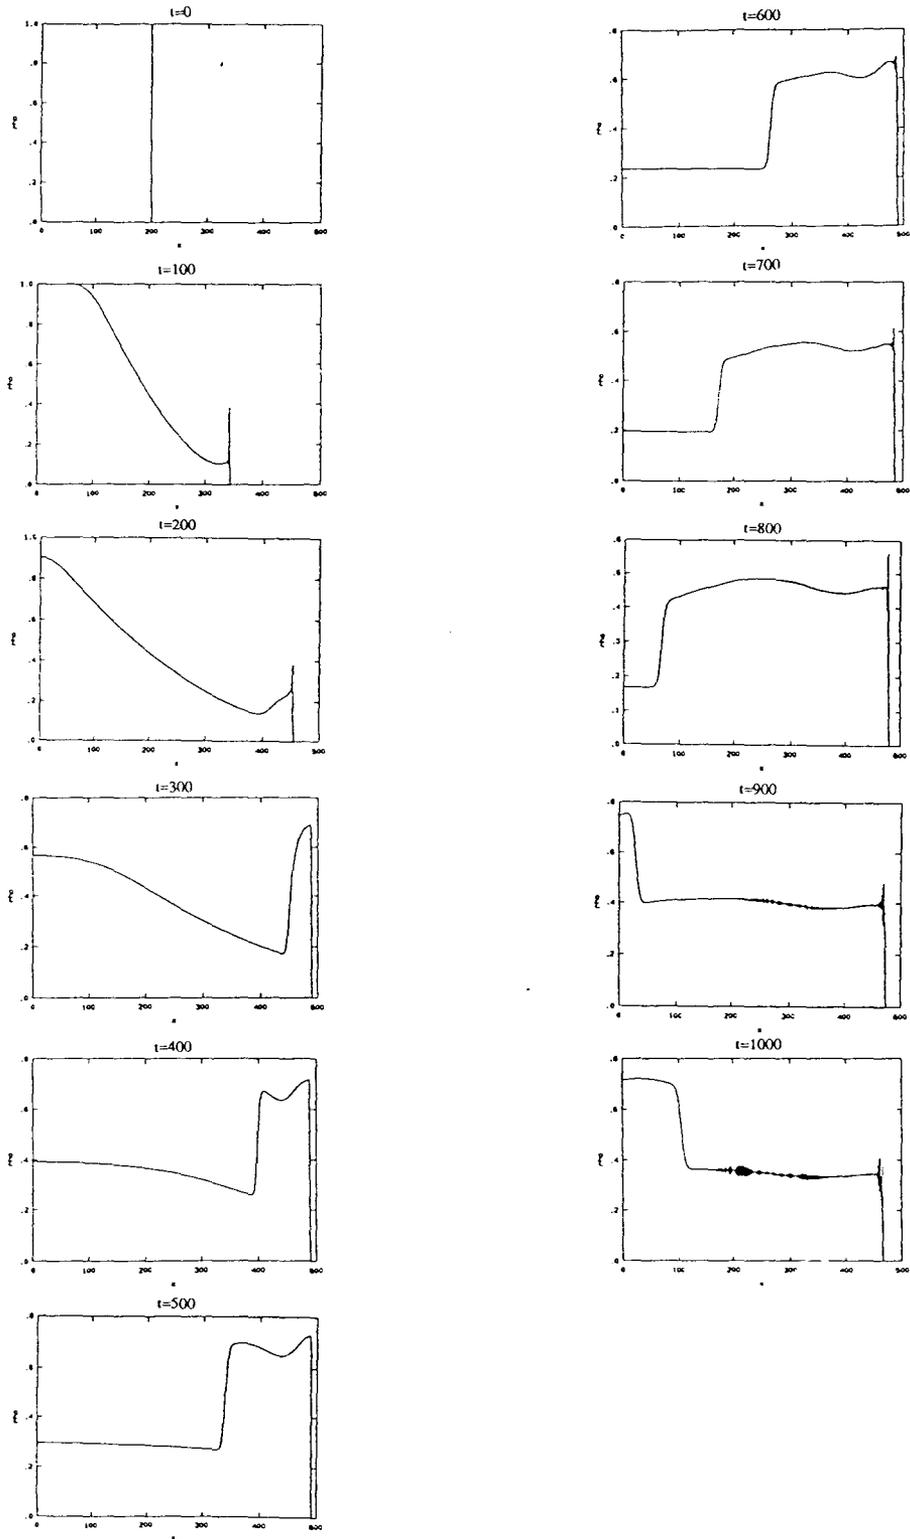


Figure 6.11b Nondimensionalized density (ρ') .v.s. nondimensional x-coordinate (x').
 (vertical axis) (horizontal axis)

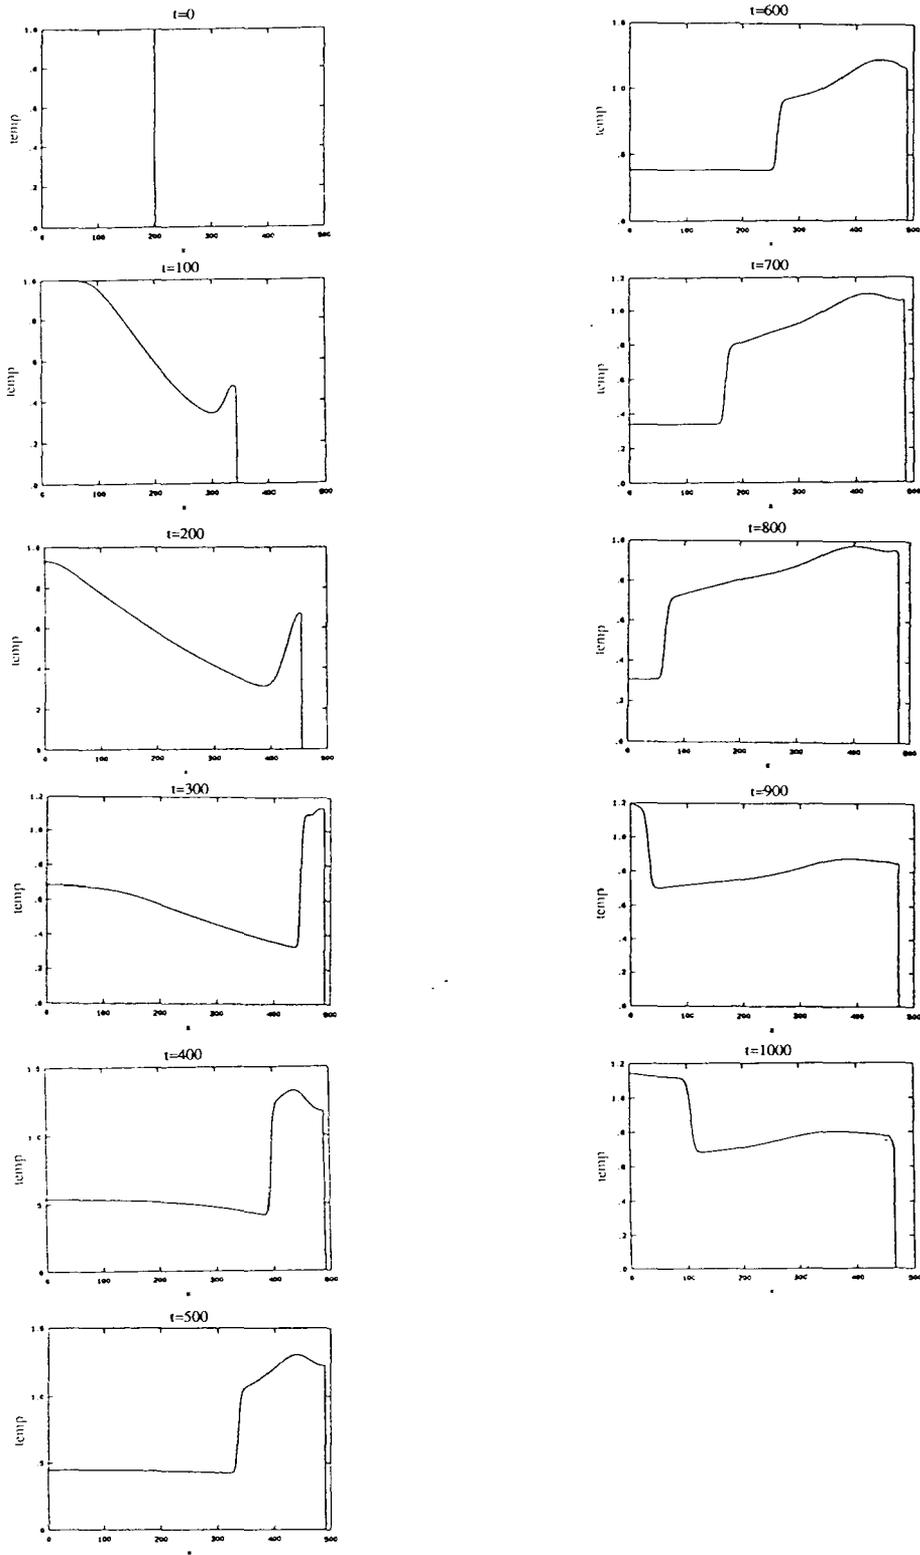
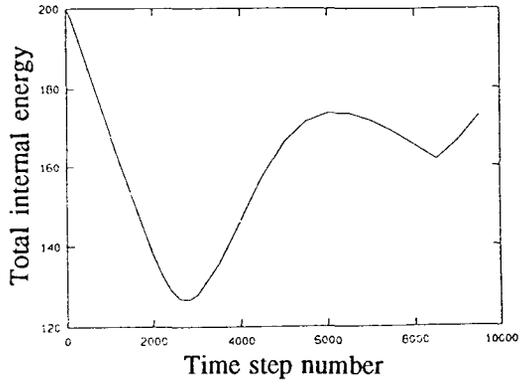
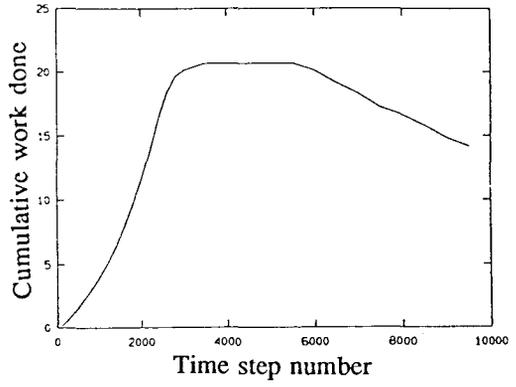


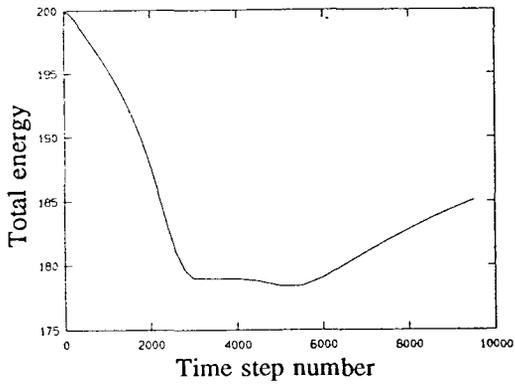
Figure 6.11c Nondimensionalized temperature (T') .v.s. nondimensionalized x -coordinate (x').
 (vertical axis) (horizontal axis)



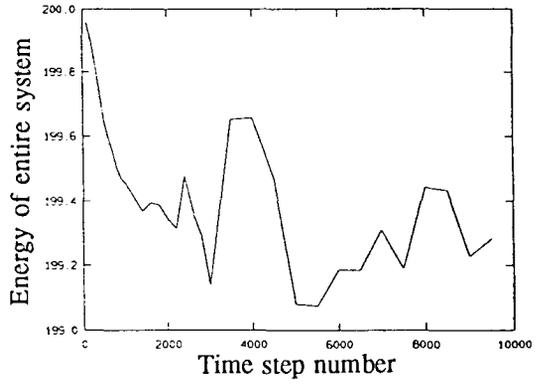
a



c



b



d

Figure 6.12

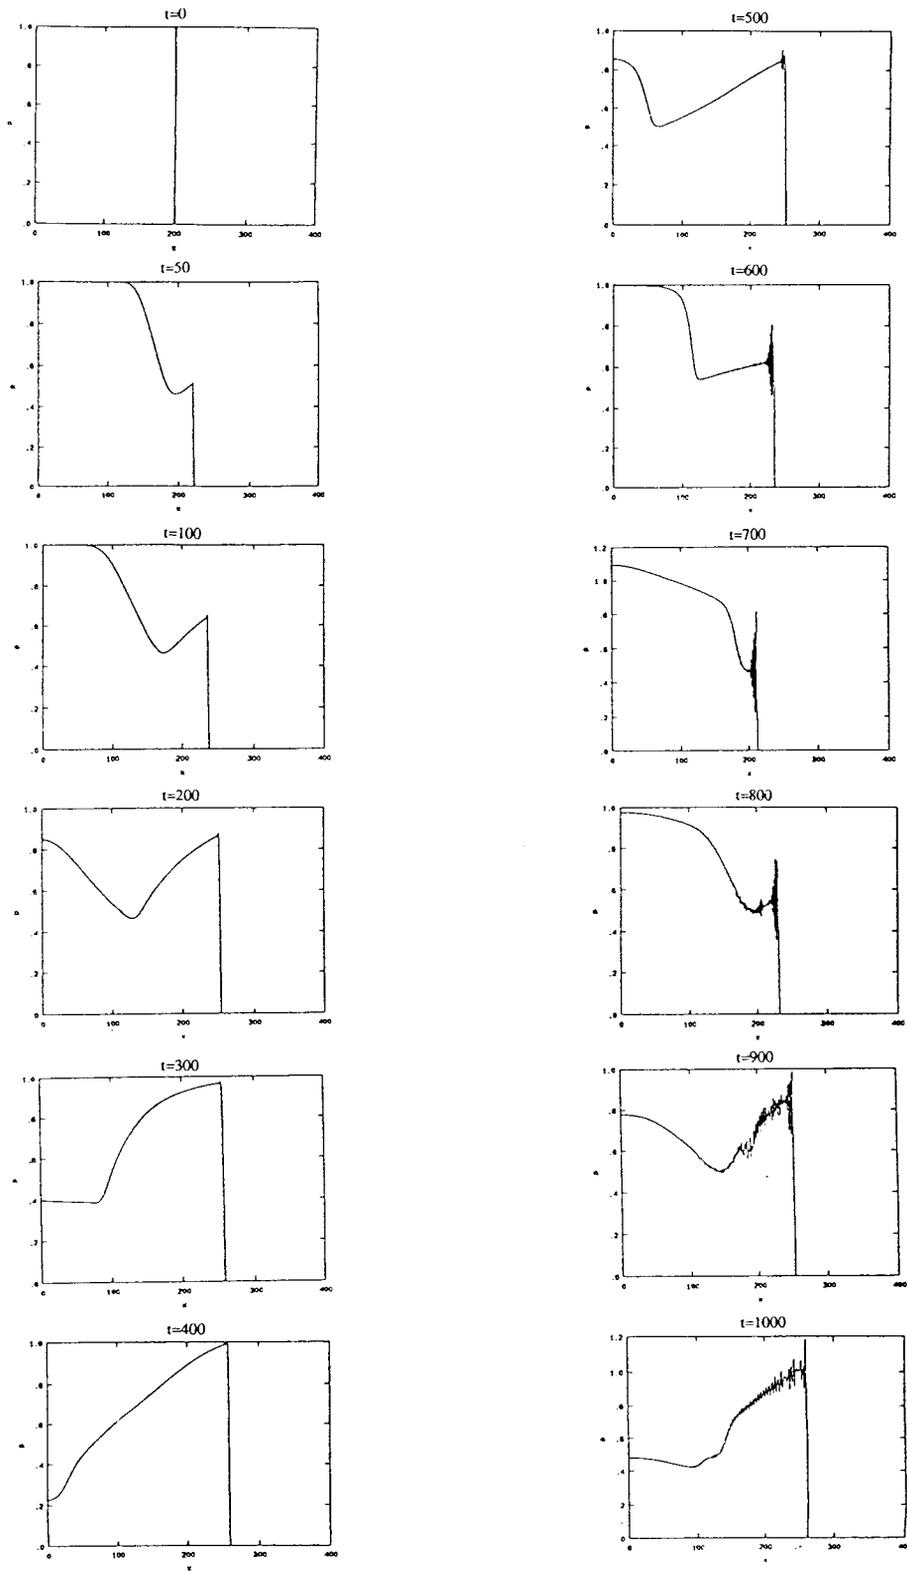


Figure 6.13a Nondimensionalized pressure (p') v.s. nondimensionalized x-coordinate (x').
 (vertical axis) (horizontal axis)

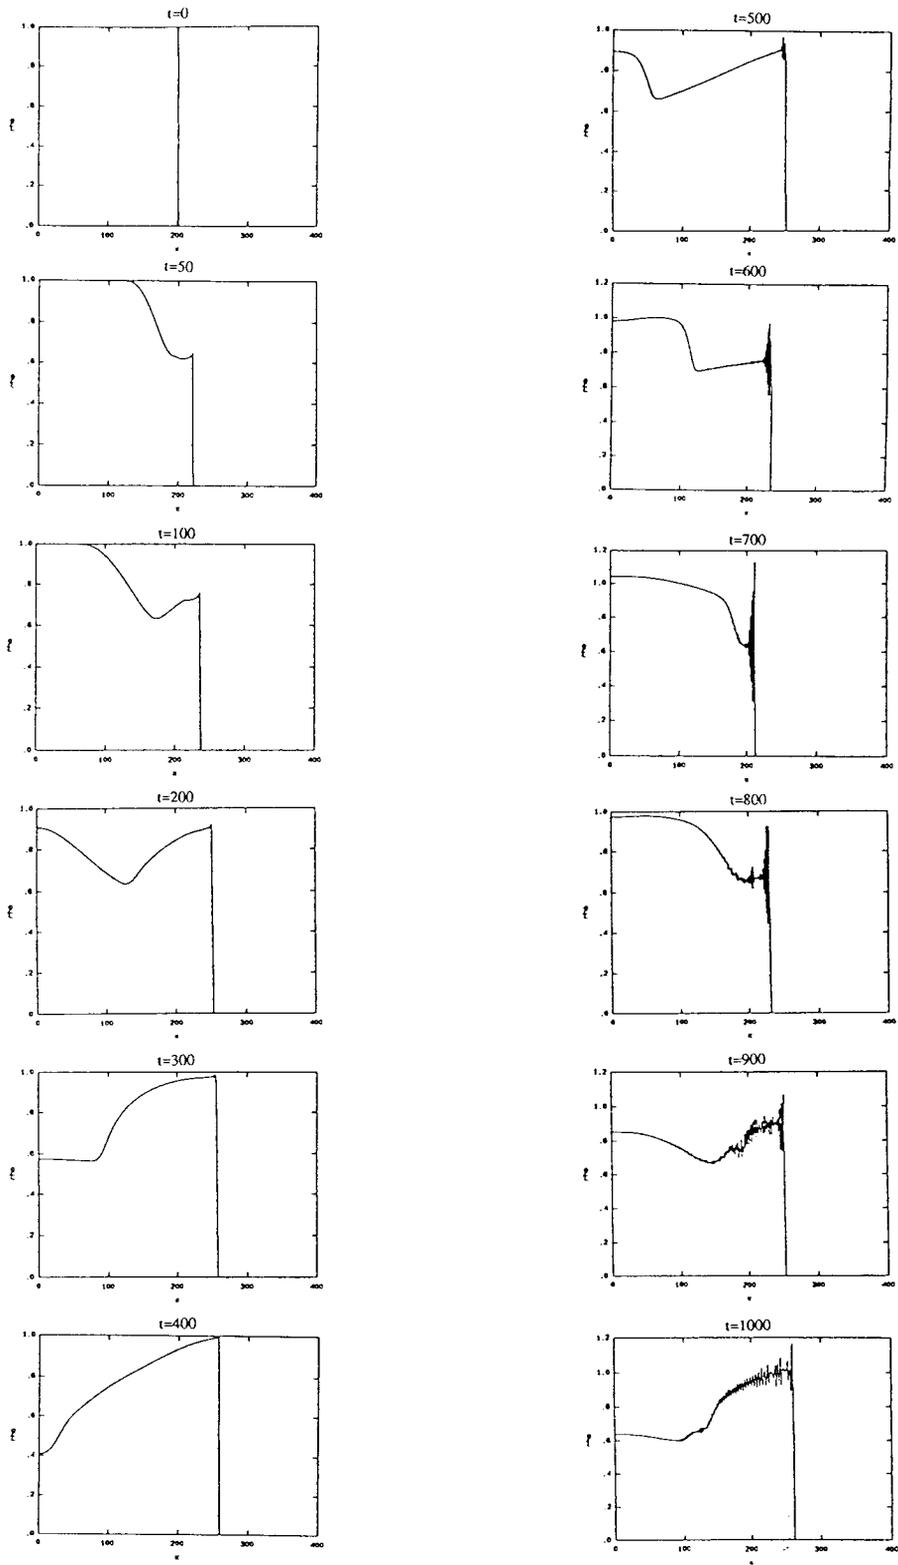


Figure 6.13b Nondimensionalized density (ρ') .v.s. nondimensional x-coordinate (x').
 (vertical axis) (horizontal axis)

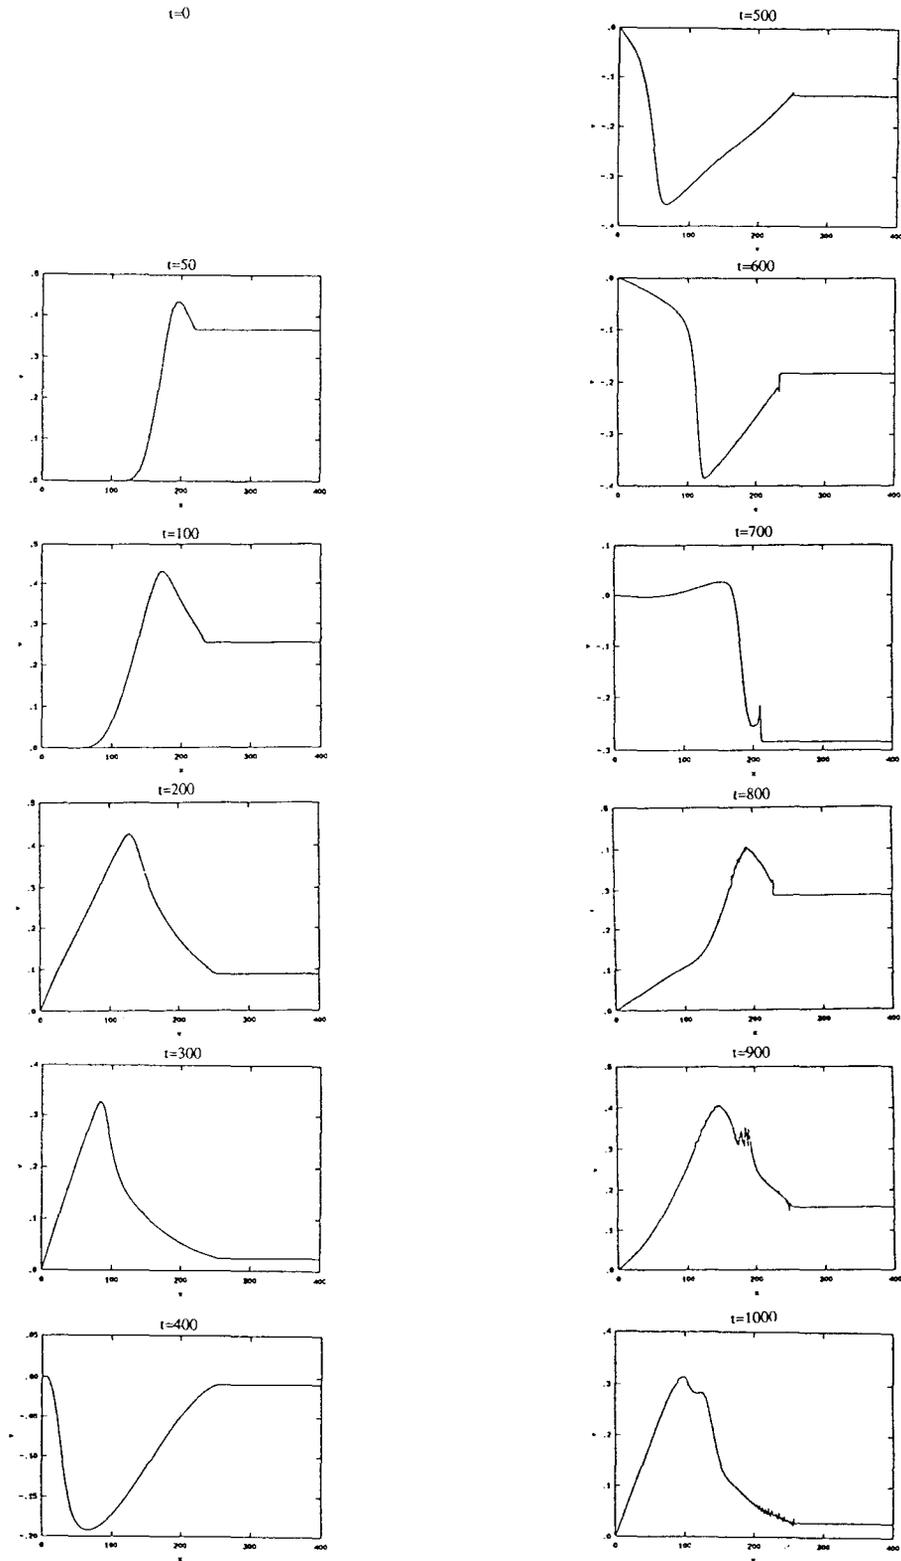


Figure 6.13c Nondimensionalized velocity (u') .v.s. nondimensionalized x-coordinate (x').
 (vertical axis) (horizontal axis)

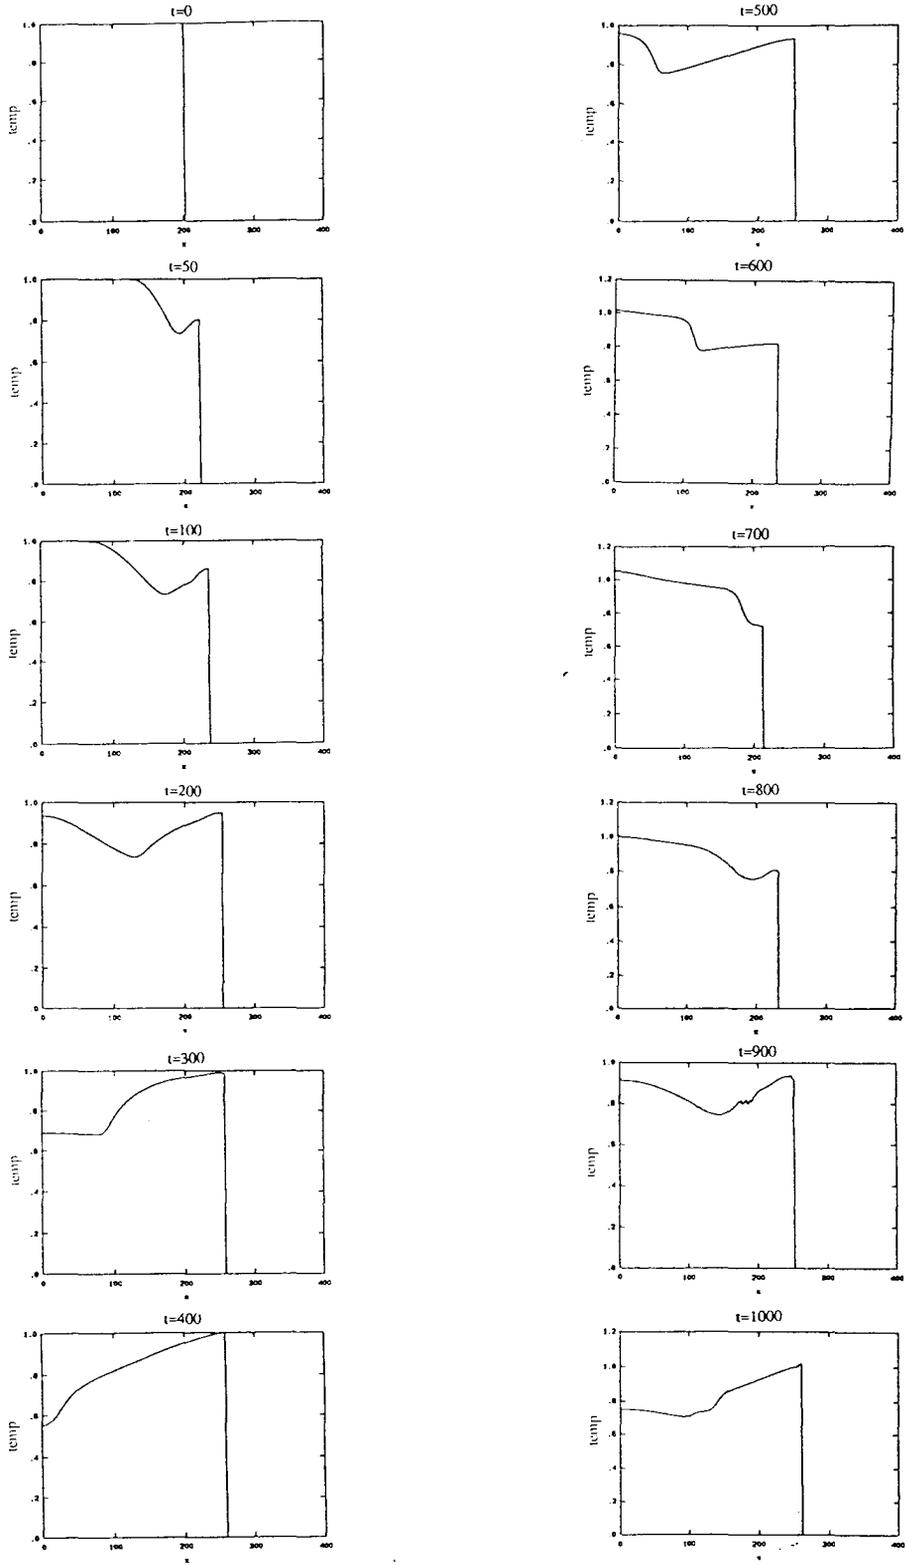


Figure 6.13d Nondimensionalized temperature (T') .v.s. nondimensionalized x-coordinate (x').
 (vertical axis) (horizontal axis)

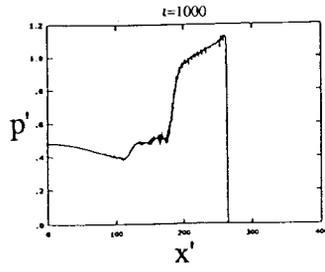


Figure 6.13a'

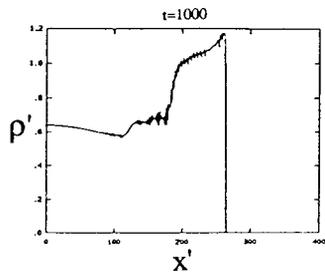


Figure 6.13b'

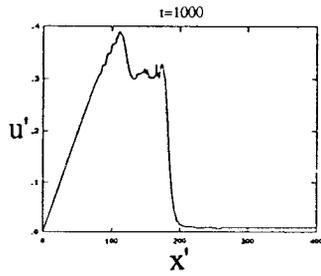


Figure 6.13c'

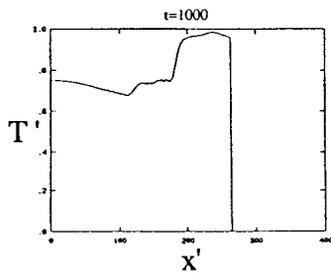
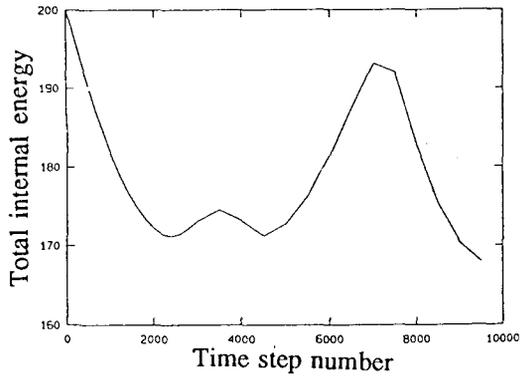
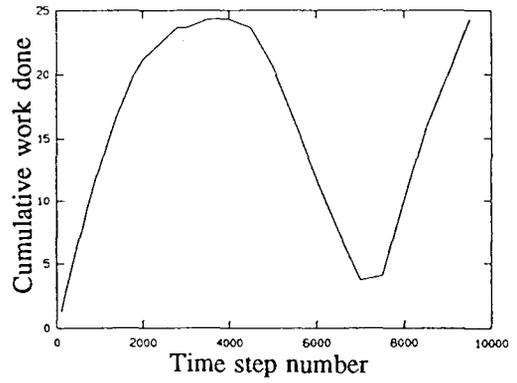


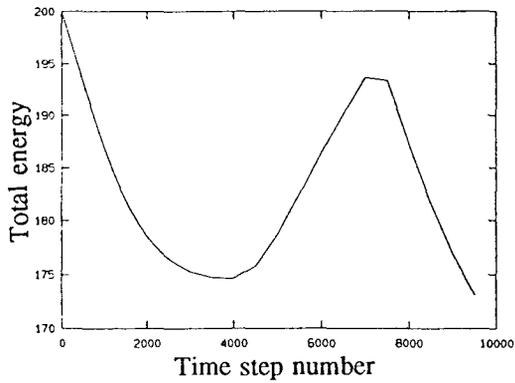
Figure 6.13d'



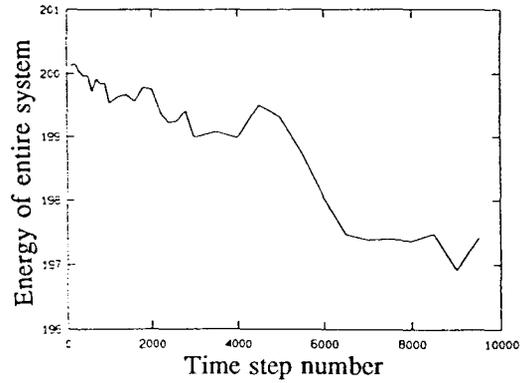
a



c



b



d

Figure 6.14

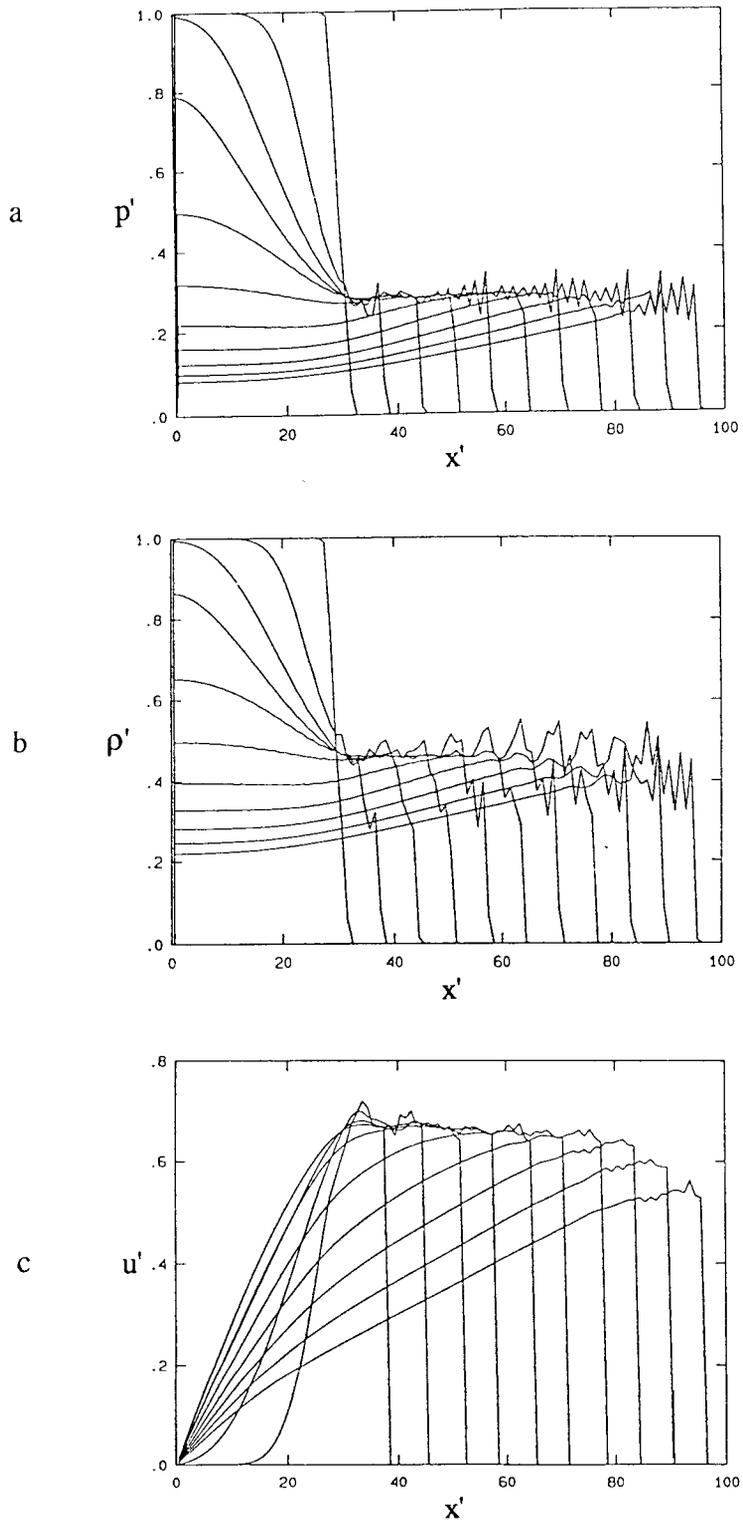


Figure 6.15

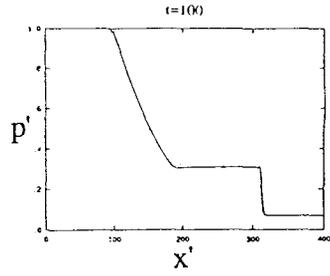
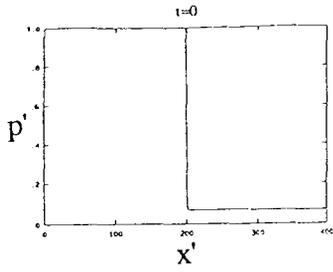


Figure 6.16a

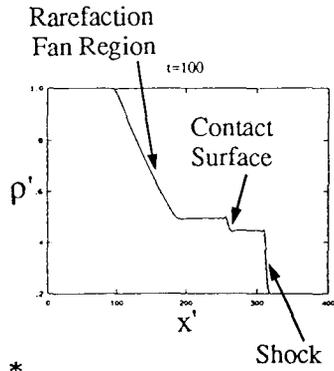
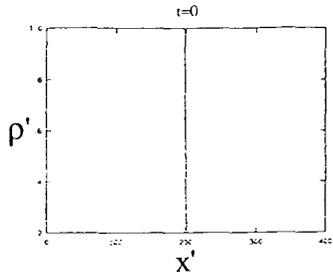


Figure 6.16b *

$t=0$

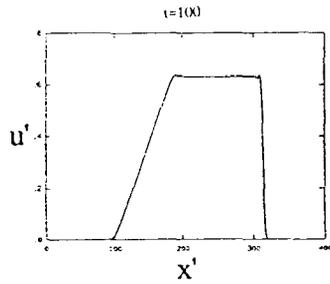


Figure 6.16c

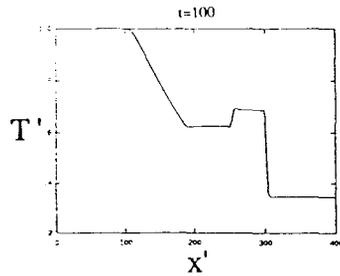
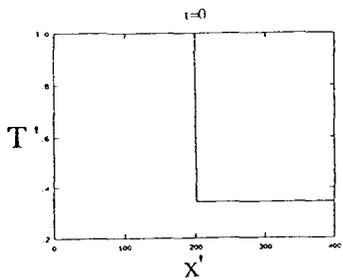


Figure 6.16d *

* Note the horizontal axis is not necessarily at the zero of the quantity for the vertical axis.

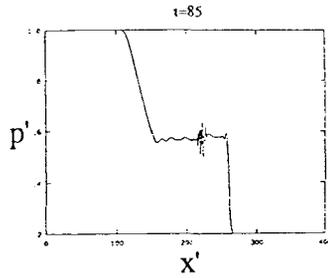
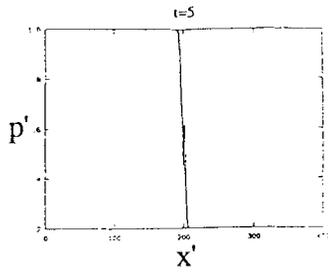


Figure 6.17a *

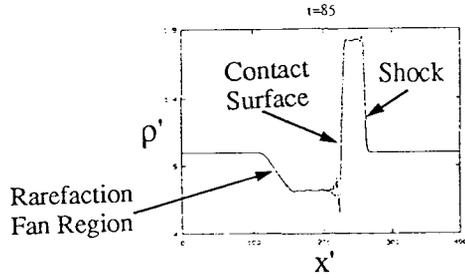
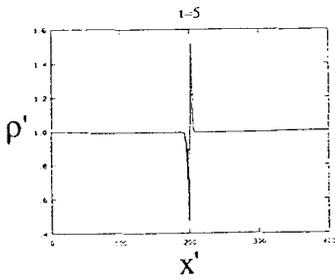


Figure 6.17b *

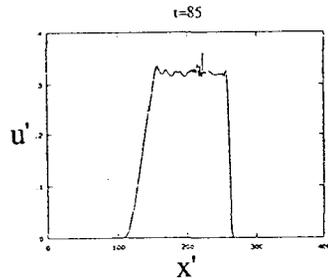
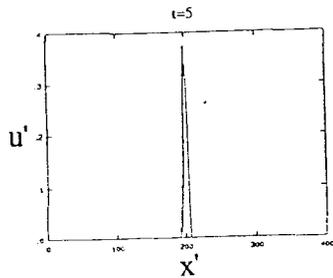


Figure 6.17c

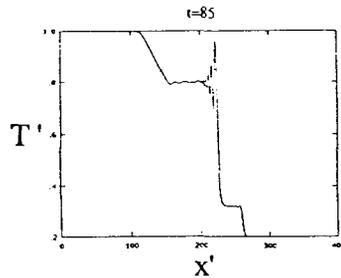
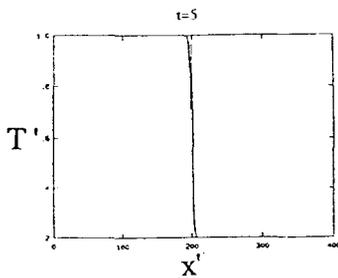


Figure 6.17d *

* Note the horizontal axis is not necessarily at the zero of the quantity for the vertical axis.

Chapter 7

Two-Dimensional Hydrodynamic Simulations using Realistic Thruster Fields

This chapter will deal with two-dimensional fully hydrodynamic simulations of plasma bulk flow in the realistic geometry of a magnetic thruster for a proposed ICF pulse rocket. In all calculations presented in this chapter, it is assumed that an initially spherical and unmagnetized perfectly conducting fluid obeying the classical laws of hydrodynamics expands, under conditions of azimuthal symmetry, against a vacuum magnetic field created by the thruster's field coils. The particular thruster setup chosen was a single current-coil design with an aspect ratio of $\tan 1$ (radian) and a coil radius of 650 cm. The explosion point is on the longitudinal central axis of the thruster about which azimuthal symmetry is assumed, as is the center of the thruster's field coil whose plane is perpendicular to the axis. Consistent with the simplifying assumptions used in earlier chapters, radiation, recombinations, physical viscosities and heat conduction are neglected.

The PIC technique employed in the planar-geometry hydrodynamic simulations has been combined with the techniques for calculating the magnetic pressure at the interface between the plasma and the vacuum field, which were employed in the simulations carried out in two-dimensions using the thin-shell approximation. In fact, the calculation of the vacuum magnetic field is carried out the same way as in Chapter 4, on the same rectangular ξ - θ grid used in the base-line thin-shell-approximation simulations there, and the same cross section of the coil structure (on which surface the magnetic stream function is held fixed) is also used. The effect of the induced surface currents is fully taken into account by assuming flux conservation between perfect conductors; i.e., that the magnetic flux present between the central axis and the surfaces of the perfectly conducting field coil structures

from before the plasma is introduced stays trapped between the plasma surface and the surfaces of the field coil structures as the perfectly conducting plasma starts expanding. The grid employed for the PIC calculation is independent of this ξ - θ grid. For the PIC calculation, the same version used in the planar-geometry simulations was selected. Numerical diffusion is controlled by the first-order-accurate algorithm as experience with the one-dimensional calculations suggested that it was sufficiently accurate while at the same time effective at suppressing numerical instabilities. The addition of artificial viscosities over and above the numerical viscosity implicitly present in the algorithm was not performed. For improved results, options for switching between different orders of numerical diffusivities and the addition of artificial viscosities could be considered.

As fixed grids were employed in an effort to minimize extraneous numerical difficulties, memory limitation of the computers forced the use of parameters that were rather removed from those that would describe the situation in thrusters proposed for ICF pulse rockets. However, the goal of developing a two-dimensional code that automatically and self-consistently tracks, i.e., captures, the motion of a highly dynamic interface between a perfectly conducting fluid and a vacuum magnetic field has been accomplished. Shocks were also successfully captured.

Now, one earlier work ¹ familiar to the author presented a plot of the interface profiles, at regular time intervals, for an ICF debris plasma expanding in the magnetic thruster of a proposed ICF pulse rocket. The work is the only one known to the author that is referred to as having conducted a (perfectly conducting) hydrodynamic analysis of a flow in the magnetic thruster of an ICF pulse rocket. The plot has been claimed to be due to the use of a classified radiation-hydrodynamic code for handling nuclear bomb explosions, except for the addition of a magnetic field routine for supplying the magnetic pressure at the interface. But other than the fact that the parameters for which the classified code was written for were not exactly appropriate for the ICF pulse rocket situation, little information

has been made available (stated as lost) regarding the calculation. However, this plot of the interface profiles could be replicated extremely closely by simulations using the thin-shell approximation conducted by the author, when similar parameters were used. The plotting of the interface in Reference 1 is terminated before the time the interfacial shell elements started falling onto the central axis in the author's thin-shell simulations. Furthermore, as a code that handles the expansion of a conducting fluid against a vacuum magnetic field cannot be created by taking a code for treating (the different problem of) nuclear bomb explosions in an atmosphere (the usual nuclear explosion code) and simply supplying an "*interfacial*" magnetic pressure through a magnetic field routine, the classified code must have been a somewhat specialized code for handling nuclear explosions into a vacuum, unless other necessary modifications were also made. Lastly, tremendous memory requirements as well as long CPU hours, even on today's largest supercomputers, are necessary to run these codes for handling the very large expansion ratios displayed in the plots. Sophisticated moving grid codes with novel adaptive gridding schemes may also be required. Since the work of Reference 1 precedes the present work by about a decade, it is surprising to learn that the use of such large and powerful advanced computers (and codes) was already possible back then.

7.1 Formulation of the Algorithm

Because we have a fluid that is initially spherical and expanding purely radially, a spherical polar coordinate system centered at the sphere's center lends itself most naturally to the calculations. However, the presence of a coordinate singularity at the origin of a spherical coordinate system does introduce numerical difficulties, especially in association with the treatment of computational particles touching the central axis. The flow in the region near the central axis, and especially at the coordinate origin, becomes very difficult to model accurately. A cylindrical coordinate system could be employed as an alternative.

The convention for the spherical coordinate system used in this chapter is the same as that employed for the shell motion in the chapter on two-dimensional simulations carried out under the thin-shell approximation, with $\theta=0$ pointing along the central thrust axis towards the front of the vehicle and $\theta=\pi$ pointing aft in the direction of the main exhaust.

A uniform grid in rectangular r - θ space normalized by Δr and $\Delta\theta$ in the two perpendicular directions was employed for the area-weighting calculations, and because the size of an FSP is set equal to that of the grid cells over which it lies, the FSPs were also generally all of the same cross-sectional area in this plane. As alluded to in Chapter 5, a cross-sectional area-weighting scheme must obviously take into account a factor proportional to the mass of an individual FSP, and thus the weighting scheme is effectively a mass-weighting calculation, but it is often implicitly assumed that the mass distribution within each FSP is uniform in the space in which the cross-sectional area-weighting scheme is carried out. The mass of an individual FSP was set proportional to its exact initial physical volume with the same proportionality constant for all FSPs. For simplicity, the initial density profile was chosen to be uniform, and this was achieved by a uniform distribution of the FSPs in rectangular r - θ space. The initial pressure was also taken to be uniform within the fluid. The precise details of the distributions of physical quantities in the early phases of expansion of an ICF debris plasma are not expected to affect significantly the properties of expansion flow out on the characteristic length scales of the thruster.

That part of the momentum and energy equations which are solved via finite differencing on the Eulerian grid are

$$\rho \frac{\partial \underline{v}}{\partial t} = -\nabla p \quad (7.1)$$

and

$$\rho \frac{\partial \epsilon}{\partial t} = -\nabla \cdot (p\underline{v}) , \quad (7.2)$$

respectively. In spherical coordinates, these become

$$\rho \frac{\partial v_r}{\partial t} = -\frac{\partial p}{\partial r} \quad (7.3)$$

$$\rho \frac{\partial v_\theta}{\partial t} = -\frac{1}{r} \frac{\partial p}{\partial \theta} \quad (7.4)$$

and

$$\rho \frac{\partial \varepsilon}{\partial t} = -\frac{1}{r^2} \frac{\partial (p v_r r^2)}{\partial r} - \frac{1}{r \sin \theta} \frac{\partial (p v_\theta \sin \theta)}{\partial \theta}, \quad (7.5)$$

where

$$v_r = \dot{r} \quad (7.6)$$

$$v_\theta = r \dot{\theta} \quad (7.7)$$

and

$$\varepsilon = \frac{v^2}{2} + \frac{1}{\gamma - 1} \frac{p}{\rho}. \quad (7.8)$$

The equations can be nondimensionalized by defining dimensionless variables (primed quantities) as follows :

$$\rho' = \rho/\rho_*, \quad p' = p/p_*, \quad v' = v/v_*, \quad \varepsilon' = \varepsilon/\varepsilon_*, \quad r' = r/l_*, \quad \text{and} \quad t' = t/t_*,$$

where the quantities ρ_* , p_* , v_* , ε_* , l_* , and t_* are arbitrarily chosen density, pressure, velocity, specific energy, length, and time values. The values of v_* , ε_* , and t_* were chosen to be defined through

$$v_* = \sqrt{\gamma \frac{p_*}{\rho_*}} \quad (7.9)$$

$$\varepsilon_* = \frac{1}{\gamma - 1} \frac{p_*}{\rho_*} \quad (7.10)$$

$$t_* = \frac{l_*}{v_*}. \quad (7.11)$$

The value of v_* is equal to the sound speed in a collisional gas of neutrals with a pressure of p_* and a density of ρ_* . It equals the ion acoustic speed in a plasma if p_* is the total pressure of the plasma (i.e., $p = p_i + p_e$) and if $\gamma_e = \gamma_i (\equiv \gamma)$, although the latter is not necessarily true. With these substitutions, Equations (7.3), (7.4), and (7.5) become

$$\frac{\partial v_{r'}}{\partial t'} = -\frac{1}{\gamma \rho'} \frac{\partial p'}{\partial r'} \quad (7.12)$$

$$\frac{\partial v_{\theta}'}{\partial t'} = -\frac{1}{\gamma' \rho'} \frac{\partial p'}{\partial \theta} \quad (7.13)$$

$$\frac{\partial \epsilon'}{\partial t'} = -\frac{\gamma-1}{\rho' r'^2} \frac{\partial}{\partial r'} (p' v_r' r'^2) - \frac{\gamma-1}{\rho' r' \sin \theta} \frac{\partial}{\partial \theta} (p' v_{\theta}' \sin \theta) . \quad (7.14)$$

Concerning the motion of the interface, the major difference between thin-shell calculations and hydrodynamical PIC calculations is that in the latter, the interface advances only in a stepwise manner because of the use of an underlying Eulerian grid. A grid cell either has particles in it (i.e., "filled") or is empty, and the continuous motion of a real interface is replaced by a motion in which the interface, in one time step, either moves by one grid-cell width or does not move at all. The interface cells, in analogy with the planar-geometry case, were defined to be the cells closest to the interface that do not bound empty cells along any of its boundaries or corners, and the finite difference equations describing the hydrodynamics were not solved in these interface cells. Best results were generally obtained when the interface "curve," as inputted to the field solver to represent the fluid-vacuum field boundary, was formed by connecting the midpoints of the "outer" boundaries of the interface cells. The code was designed to search for the interface points, given a checkerboard of filled and empty cells, and then to connect them together in correct order.

Because the plasma is being approximated by an unmagnetized perfectly conducting fluid, there will be no magnetic tension or pressure body forces, and the only force the magnetic field exerts on the plasma will be the $B^2/8\pi$ force acting on the interface (refer to Sections 4.1.2 and 4.3.1 of Chapter 4), calculated quasi-statically under the prescription of flux conservation between perfect conductors. This externally applied magnetic pressure was specified to apply normally to the interface "curve" defined above. Magnetic pressure forces applying on the sides of an interface cell were separated into components and summed up separately to be used in the respective components of ∇p in the hydrodynamic equations for those cells. To help prevent anomalous calculations of the interfacial magnetic pressure that can arise from the coarseness of the underlying grid, from applying,

an averaging of the interfacial magnetic pressure was carried out over several neighboring interface cells before being assigned to the individual cells.

7.2 Some Properties and Limitations of the Code

The major difficulties in these simulations arise at the interface and at the coordinate origin. The success of the calculations depends largely upon whether or not the code can satisfy a pressure match across the interface at all time steps by correctly keeping track of the interface motion. With hydrodynamic computations for the planar geometry described in the previous chapter, it was found that setting the interfacial thermal pressure of the perfectly conducting fluid exactly equal to the externally applied pressure at the interface at the start of the simulation was not necessary to obtain a decent pressure match across the interface at subsequent time steps. What is more important is that the parameters be chosen in such a way that the code be able to match the interfacial pressures at subsequent time steps, under such limitations as on the grid resolution, and in the two-dimensional codes, efforts were not made to match the exact interfacial fluid pressure with the polar angle (i.e., θ) dependent interfacial magnetic pressure at the start of the simulation.

The choices of the pressure mismatch across the interfacial region and the fluid velocity distribution at the start of the simulation do affect how smoothly and quickly a good pressure match can be achieved across the interface. If the initial fluid pressure is not significantly higher than the maximum initial magnetic pressure, difficulties can arise at the interface because in a three-dimensional expansion with a large expansion ratio, the bulk pressure of the plasma falls very rapidly with the expansion while the vacuum magnetic pressure acting externally on the interface does not. This produces an enormous pressure difference between the interfacial pressure and the bulk pressure of the fluid as the expansion proceeds, leading to the formation of a narrow region of fluid pressure rise near

the interface, which can be difficult to simulate (see below). It should be noted, in passing though, that the interfacial magnetic pressure, even where it rises with the expansion, does not rise by orders of magnitude during a typical expansion unless one is concerned with the stages of flow in a low-field thruster when the interface is very close to its velocity reversal location. On the other hand, if the initial fluid pressure is very much higher than the maximum initial interfacial magnetic pressure, but not excessively high, such that the bulk fluid pressure during the expansion ends up on the order of the typical interfacial pressure, good results at the interface can be obtained with relative ease. Also, other factors play a role in obtaining a quick and successful matching of the interfacial pressure during the simulation. For example, starting the fluid from rest helps. Now, when difficulties were met in matching the pressures across the interface in a time much shorter than the characteristic expansion time scale of the plasma, a deep negative pressure spike, due entirely to numerical difficulties, sometimes appeared at the interface, at the locations of highest field. This negative interfacial pressure spike would appear when the deceleration at the interface started becoming noticeable and disappear when, after running many time steps, the fluid pressure in the neighborhood of (i.e., right behind) the interfacial region rose high enough to approach the interfacial magnetic pressure value so that a pressure match across the interface could be more readily achieved. The width of the negative interfacial pressure spike was about the same as that of the numerical interfacial density spike and stayed approximately constant with time. This negative spike was not occurring because the simulation of the bulk fluid was not satisfying the criteria to prevent negative internal energies from occurring, and in fact, could not be eliminated by reducing the time step employed in the calculations.

Next, in simulations employing grid-cell dimensions that are rather large on the length scales of the gradients, the coarseness of the grid's spatial resolution may hamper the observation of a good pressure match across the interface because pressure values assigned to cells are only cellwise averages. For example, if the width of the interfacial

pressure rise region is on the order of, or narrower than, the width of the grid cells, then the discrepancy between the true peak pressure and the cellwise average we see in the PIC results can be quite large, leading to an underestimation of the interfacial fluid pressure. However, this also means that an apparent mismatch between the fluid and magnetic pressures across the interface may simply be due to a sharp and narrow rise in the real fluid pressure not being represented correctly because of the cellwise-averaging procedure, and not because the code is incapable of giving correct results even within the limitations of the grid resolution. Care is also required in reading off the interfacial fluid pressure value from the plots because the sharp drop-off of the pressure profile at the interface may actually be several grid cells wide at some polar angles. Furthermore, numerical spikes and oscillations, which are characteristic features present at discontinuities and interfaces can make accurate read-offs from plots difficult.

Also, the fact that an interface represented via an underlying grid system may remain stationary for many time steps while the real interface moves, implies that a good match between the interfacial fluid and vacuum magnetic pressures may not be observed even when the interface motion is being captured correctly to within the resolving capabilities of the code (note that all the while the interface is stationary, the magnetic field stays unchanged). Having the "cushion" layer at the interface that is not counted as fluid cells helps mitigate this effect. It was also found that it generally takes many time steps, during which time the interface moves over multiple grid zones in the direction of the flow, before the interfacial fluid pressure slowly converges onto the changing interfacial magnetic pressure, even on a time-averaged sense. Again, that was largely due to the difficulty of representing the real distribution of physical variables at the interface accurately, using PIC with a coarse grid, and not an indication that the general bulk flow was not being treated correctly.

In a PIC simulation, the maximum density that can be obtained at a particular radius is given by the value that would be achieved were all computational particles to be packed

into one grid-cell zone at that radius. If the entire plasma mass, M_{tot} , is assumed to be collected into a spherical shell of radius R with a thickness of one radial grid zone, Δr , then this density will be given as

$$\rho_{\text{max}} = \frac{M_{\text{tot}}}{4\pi R^2 \Delta r}, \quad (7.15)$$

and if an isentropic relation is assumed for the flow, one obtains the following relation for the maximum treatable pressure at that radius :

$$P_{\text{max}} \leq P_0 \left(\frac{R_0^3}{3R^2 \Delta r} \right)^\gamma. \quad (7.16)$$

In deriving this relation, it has been assumed that the plasma starts out from a spherical blob of radius R_0 , uniform pressure p_0 , and uniform density. For the most stringent criterion, one can choose R to be on the order of the characteristic radius at which the plasma interface is turned back by the magnetic fields. This shows that for PIC calculations, the use of large grid cells can become a limiting factor in obtaining high densities and thus pressures out at large radii for an adiabatic flow. However, if relation (6.13) is not imposed, such a concern is generally not serious for the cases run in this work (especially because of the appearance of shocks). If relation (7.16) were to be satisfied, a prohibitive number of grid cells would have been required.

High noise can also plague plots where the particle density is low and the accuracy of PIC, low. In fact, the pressure, density, and temperature rises observed in regions of the fluid affected by the finite interfacial pressure application are sometimes numerical. In particular, the very narrow spikes observed at the interface at polar angles where the field is weak, such as in the vicinity of $\theta=0$ and π , or even at higher field regions at early time steps, are almost entirely numerical, and is characteristic of numerical simulations involving vacuum interfaces that exert no, or relatively low, pressures on the fluid boundary. The lower the interfacial magnetic pressure gets relative to the fluid pressure of the bulk in the neighborhood of the interface, the taller the spikes get. The broad rises in the profiles of the physical variables are real, and moreover, sometimes the peaks of the profiles are not

attaining the real heights as mentioned above, but this affects only the details of a narrow region around the sharp peak and the distribution of physical variables in the bulk are believed to be given quite realistically. On the whole, the capturing of the interface appears to be proceeding successfully in terms of cell averaged profiles and interface locations.

Because no physical heat transfer from the exterior was assumed, results were compared with codes employing the adiabatic relation of Equation (6.13). Use of this relation will obviously prevent the problem of negative pressures from arising because pressure will then be calculated from the density, which is based on an actual number count of computational particles. Smoother, less noisy profiles for the distribution of physical variables could be obtained with codes employing relation (6.13), but such codes could not get the fluid pressure at the interface to rise to levels that codes which used the "full" energy equation (7.2) could, and were thus less capable of yielding pressure plots that showed a pressure match across the interface. With codes that do not rely on the adiabatic relation, quite a nonnegligible temperature rise above what is possible under the adiabatic assumption was observed in regions of the fluid where the deceleration of the interface was being felt. This indicates heating due to numerical viscosities (and artificial viscosities, if present). In a real fluid, these viscosities will be replaced by physical viscosities, with viscous effects being high where there are large gradients in velocity, and this heating of the fluid can be real. As expected, the temperature rise was generally found to be highest at locations where the interfacial magnetic field is strongest. Shocks form during the course of the expansion because of the deceleration of the interface (in a plasma, the classical hydrodynamic shock will be replaced by a collisionless shock) and it is this (real) shock heating of the fluid that is playing a large role in raising the temperature at the interfacial region, but as noted earlier, shocks cannot be handled properly with an adiabatic code. Simple estimates show that any raising of the plasma temperature in the outer regions may be vital for the device to function properly as a thruster.

As an aside, by changing the inequality of relation (7.16) into an equality and by replacing Δr with the physical shell thickness, ΔR , which is to be assigned an approximate functional form in terms of R (based on data), one may obtain, with the help of the ideal gas equation of state, an approximate expression for the temperature of an interfacial shell as a function of its radius for the isentropic case. Such a calculation shows that it is even possible for the temperature of the interfacial shell to rise with radius, as was mentioned earlier.

It should be noted that the temperature rise occurs where the density and pressure are also rising. This contrasts with the situation present in an explosion into a medium such as a nuclear explosion in an atmosphere or a supernova explosion in the interstellar medium, wherein a low density region in the interior has very high temperatures. In such problems, a shock travels outwards into the external medium, followed by a contact surface, which separates the contents of the explosion debris from the fluid of the medium that has been run over by the shock. In the explosion of a plasma against a vacuum magnetic field, the interface is in a sense a contact surface, holding all of the debris mass within, and the interface deceleration causes a shock wave to start traveling inwards, relative to the outward flow, from the vicinity of the interfacial region. This explains the qualitative difference observed between the two classes of phenomena.

As stated in the previous chapter, a fixed-grid PIC code will require an impractical number of computational particles to treat even modest expansion ratios ("modest" in the context of plasma expansion in the thruster of an ICF pulse rocket). The number of grid cells required to treat fully an expansion flow in a typical ICF pulse rocket will also not be small. Nonuniform, moving grids will allow a reduction in the number of grid cells and particles,² but as they will introduce spurious numerical problems such as false compressions or expansions, especially as the particle number density drops low, they were not employed in the present work. Instead, parameter ratios were altered from those

that approximate conditions in the thrusters of proposed ICF pulse rockets so that "thruster-filling" expansions will be treatable with a fixed-grid code. Namely, the ratio of the characteristic dimension of the expansion to the radius of the uniform spherical body of gas at the beginning of the simulation was taken to be very much smaller than appropriate for a realistic ICF pulse rocket thruster simulation. If the field-coil parameters are assumed to correspond to those of proposed thrusters, this implies starting the simulation with a plasma sphere whose radius is much too large for the blob to be considered uniform in density and pressure. As will be discussed later, such simulations give different results from those that allow a uniform plasma of almost negligible radius (relative to the characteristic dimensions of the thruster) to expand into the thruster. The exact development of the interface shapes with time, the precise distribution of fluid quantities within the blob, and thus also the propulsive efficiencies, will all differ, but the characteristic qualitative features of the flow will be the same, as will the techniques for handling the hydrodynamics.

As we have seen, the use of finite-sized grid cells to which cellwise-averaged quantities are assigned and the employment of a modest number of finite-sized computational particles lead to PIC simulations that do not represent a real fluid flow exactly. The use of finite time steps further aggravates the problem. But while these are problems typical of numerical simulations and generally do not prevent obtaining successful results, the existence of a coordinate singularity does amplify the incapability of code to model the flow accurately. The largest difficulty here was associated with the treatment of FSPs that touch the central axis.

In the simulations, FSPs, all with the same cross-sectional area, were first uniformly distributed up to the interface on the rectangular r - θ plane as long as they at least partially overlapped the region that corresponds, in physical space, to the upper half-plane above the central thrust axis. This domain given by $r \geq 0$ and $0 \leq \theta \leq \pi$ will be referred to

here as the "region above the central axis." Because of rotational symmetry about the central axis, the FSPs form a section of a torus extending some predetermined angle $\Delta\phi$ around the physical central axis, and FSP cross-sectional areas that are intersected by $r=0$ and/or $\theta=0$ or π cannot be considered as extending outside the region above the central axis in the rectangular r - θ plane. Therefore, FSPs intercepted by $r=0$ and/or $\theta=0$ or π were truncated by those respective boundaries. The center of a "truncated" FSP cross section was redefined to be at the center of its area above the central axis in the rectangular r - θ plane, and that portion was rotated $\Delta\phi$ about the physical central axis to create the whole FSP.

The assignment of velocities to these truncated FSPs was difficult. For example, it is easy for all particles in a cell bounded by $r=0$ and/or $\theta=0$ or π to leave that cell creating the unphysical situation of an empty cell's forming where it shouldn't. Or, depending upon the prescription, all truncated particles can readily end up remaining attached to the truncating boundaries throughout the simulation. To minimize these problems, the velocity of a truncated FSP was usually determined by carrying out a linear interpolation involving the distance of the truncated particle's center from the truncating boundaries, the cellwise average velocities, and symmetry/antisymmetry properties. The specific energy of a truncated FSP was also prescribed in a similar manner. As a truncated FSP's center moved away from its truncating boundaries, the FSP cross-sectional area was allowed to grow, with the particle maintaining contact with the truncating boundaries until it reached the same dimensions, in rectangular r - θ space, as that of a nontruncated FSP. After that, the FSP was allowed to detach from the truncating boundaries. This way, all FSPs detached from the boundaries possessed the same cross-sectional area in rectangular r - θ space. The mass of all FSPs were held fixed with time.

However, even with these and other refinements to the algorithm, cells bounded by the domain boundaries described above (especially the origin) could still be a source of difficulties such as with regard to the taking of gradients or the assignment of quantities to

particles, especially at low particle densities. Because of this difficulty in handling the flow near the coordinate singularity, anomalous dips or spikes in the pressure, density, and temperature could easily appear near the origin.

In order to reduce numerical difficulties involving false variations of physical variables from cell to cell near the origin (and to some extent near the entire central axis) as particle number densities approach very low values, and also to suppress excessive numerical noise, cellwise quantities were averaged over several (usually 3 or 5) neighboring cells. The averaging scheme employed for the velocity and specific energy involved linear interpolations incorporating symmetry/antisymmetry properties where necessary. Averaging over neighboring cells in the polar direction was usually found to be beneficial and considerably smoothed-out noise in that direction. However, the typical plasma expansion/deformation flow in a magnetic thruster is predominantly radial, and the grid resolution was very coarse in that direction. Therefore, averaging cellwise quantities over neighboring cells in the radial direction up to close to the interface tended to smother out the details of the spatial profiles of the physical variables too much. The smoother, less noisy results obtained did not compensate for the heavy distortion of the profiles with concomitant false wave propagation velocities. But when averaging in the radial direction was carried out only in the few radial grid zones closest to the origin, a broad unphysical dip or break in the radial profiles of the physical variables appeared at the general location where the averaging was stopped. On the whole, it appeared preferable not to carry out averaging in the radial direction. Of the examples presented in Section 7.5, all except for Cases 1, 2, 6, and 11 employed averaging of cellwise quantities over neighboring cells only in the polar direction. Cases 1, 2, and 6 employed averaging of cellwise quantities in the polar direction as well as in the first 10 grid zones in the radial direction. Case 11 did not employ averaging of cellwise quantities in either the radial or polar direction.

The prescription of a uniform mass distribution within each FSP in rectangular r - θ space results in a density that is not entirely uniform in physical space when the FSPs are

placed uniformly on the rectangular r - θ plane. It can be seen from the density and pressure plots that at the initial instant, a spike in these variables exists adjacent to $r=0$ and that the bulk values in cells bordering $\theta=0$ or π are not the same as those in cells not bounding $\theta=0$ or π , although the pressure distribution may be smoothed out artificially by an averaging scheme. This problem can be entirely removed by employing a mass distribution within each FSP that is uniform in physical space, thus normally implying a physical-volume weighting scheme, but this method was already abandoned for reasons given above.

7.3 Choice of Parameters

The modeling of an expanding plasma in the magnetic thruster of an ICF pulse rocket is initiated, assuming that an ICF pellet has been successfully exploded and that a small spherical high-temperature plasma with a radius approximately on the order of the pellet size exists. Instabilities and strong self-generated magnetic fields that might be present in an actual ICF debris plasma as well as details of the plasma debris structure involving the exact and real distributions of physical quantities, have all been neglected.

After a survey of the literature and some simple calculations, it was estimated that an acceptable initial state for an ICF plasma expanding into a thruster would be a sphere of 1 cm radius with a mass of 0.5 gram and a temperature of 25 keV having uniform properties. Such a state can be obtained, for example, if the pellet is comprised entirely of $A=2.5$, $Z=1$ material (A and Z are, respectively, the average atomic weight and charge of the ions) with about 44 mg of DT fuel of which roughly 1/3 burns with 3.5 MeV per reaction available to the charged particles (which is not necessarily realistic). It was assumed that the electron and ion temperatures are equal. At this "initial state," the entire plasma energy was taken to be in the form of internal energy.

Numerical simulations were carried out with an initial radius of 50 cm to allow the plasma to expand on the scales of the characteristic dimensions of the thruster, undergoing significant decelerations and deformations, all within limitations of memory and reasonable CPU time. Taking the ICF plasma as having a uniform density and pressure when its outer radius is 50 cm is not realistic, but for the purpose of simply developing a code that successfully handles the type of flow concerned, this was not considered to be a problem. In reality, by the time the plasma has expanded out to a 50 cm radius, the large expansion experienced and the finite interfacial magnetic pressure will have long caused a highly nonuniform distribution to have developed, although in a typical thruster, the interface is still expected to possess a spherical shape. A uniform distribution will develop only if the characteristic sound crossing time across the blob is much less than the characteristic expansion time. A nonuniform profile can persist if the local sound velocity is much less than the flow velocity.

The state of the plasma at the 50 cm radius, which will be referred to here as the "starting state," was approximately derived from the state at which the radius was 1 cm and the temperature, 25 keV, by assuming that the plasma expands isentropically, remaining spherical and uniform throughout. Under this simplified (though not necessarily accurate) assumption, the temperature varies with the radius of the sphere as

$$\left(\frac{T}{T_0}\right) = \left(\frac{R_0}{R}\right)^2, \quad (7.17)$$

where R_0 and T_0 are the radius and temperature values at some reference state passed through by the system. This yields a temperature of 10 eV when the radius of the sphere is 50 cm. Equilibration of ion and electron temperatures was assumed. Neglecting radiation, most of the energy of the plasma is converted into kinetic energy by the time the radius is 50 cm. The work done against the magnetic field during this early phase of expansion is not significant for proposed thruster designs.

Next, the distribution of the purely radial velocity at the starting state was taken to be linearly rising with radius from zero at the origin. The bulk-averaged velocity defined through

$$v_{\text{avg}} = \sqrt{\frac{2 \text{KE}_{\text{tot}}}{M_{\text{tot}}}}, \quad (7.18)$$

where KE_{tot} is the total kinetic energy of the blob (at the instant under concern), was given a value of 2.4×10^8 cm/sec at the starting state, roughly consistent with the isentropic expansion prescription stated above. This velocity is of the same order of magnitude as the sound speed in the fluid at the time it was started off from rest. The density profile resulting from a free expansion is not uniform, but if it is assumed that a spherical blob has a uniform density and a linear distribution of a purely radial velocity, then the velocity of the interface, v_{inf} , is related to the bulk-averaged velocity through

$$v_{\text{inf}} = \sqrt{\frac{5}{3}} v_{\text{avg}}. \quad (7.19)$$

As for the numerical values, the density of the 50 cm radius sphere will be 9.55×10^{-7} g/cm³, the total pressure, 7.17×10^6 dynes/cm², and the sound speed, 3.54×10^6 cm/sec (note its lowness compared to the bulk expansion velocity at the 50 cm radius, and this will prohibit uniform profiles from developing). The base-line case was defined to be one with a starting state involving these values for the fluid variables. These parameter values were also chosen as the characteristic density, pressure, and velocity values, ρ_* , p_* , and v_* , and along with $L_* = 500$ cm and $t_* = 1.4 \times 10^{-4}$ sec, were used to nondimensionalize the variables appearing in the plots presented in this chapter, except for the labeling of the radial distance on the plots is not in units of L_* , but is rather in units of the radial grid-cell width, Δr .

The temperature was nondimensionalized in the following manner. The ideal equation of state for a plasma is

$$p = n_i k T_i + n_e k T_e. \quad (7.20)$$

Now, if it is assumed that $T_i = T_e \equiv T$, one obtains

$$T = \frac{A}{1+Z} \frac{m_p}{k} \frac{p'}{\rho_*}, \quad (7.21)$$

where m_p is the proton mass. This becomes, for our case, with T measured in units of eV,

$$T [\text{eV}] = 7.8 \frac{A}{1+Z} T', \quad (7.22)$$

where $T' \equiv T/T_*$ was defined as the nondimensionalized temperature, also given through $T' \equiv p'/\rho'$.

A temperature of 10 eV at a radius of only 50 cm is not very promising as it is already under the ionization potential of hydrogen, but it should also be realized that a uniform expansion was assumed in taking the sphere out from a 1 cm radius to a 50 cm radius, and that in reality, very nonuniform profiles develop, and thermalizations (especially due to shocks) arising from the deceleration of the interface could substantially raise the temperature of the plasma in the important interfacial regions. However, it is still estimated that even with shock heating, the cooling of the bulk will be rapid and severe enough because of the (near) free expansion experienced, that the plasma will not only have rather high resistivities, but will also no longer be fully ionized once it has expanded to the dimensions involved in the simulations described in this chapter.

As with the planar-geometry simulations, if only a small number of radial grid zones were filled with particles at the starting state, good results were not obtained because of the coarseness of the grid, and because cells bounding the interface and the origin are prone to numerical difficulties. To ensure acceptable memory requirements with uniform gridding in the radial direction, the sphere at the starting state of approximately 50 cm radius was represented by 14 radial grid zones for simulations discussed in this chapter. In the polar direction, the grid-cell boundaries were all spaced $\pi/30$ radians apart.

In most of the simulations presented in this chapter, the separation between FSP centers in the radial direction at the starting state was $\Delta R/200$, although $\Delta R/500$ was sometimes used (similar results obtained), and the separation in the polar direction was $\Delta\theta/4$. The greater the number of particles per cell, the lower the density the code will be

able to handle, and with less noise, so this implies, "theoretically," that in the calculation of an expanding plasma, use of a larger number of particles will enable one to treat the flow longer. However, with a larger number of particles, the CPU time required per time step increases, and fewer time steps can be run given the same total CPU time.

7.4 Some Characteristic Parameter Ratios

For calculations employing the thin-shell approximation which started with a sphere of negligible radius, the two parameters that determined the behavior of the expansion were the initial ratio of the magnetic field energy to the plasma energy, ER , and the geometry of the thruster (field coil / explosion site) setup, which reduces to AR , the aspect ratio, for the single-coil case. However, when modeling the plasma expansion using the full hydrodynamic equations without assuming material concentration into a thin shell possessing no internal fluid properties, other parameters can also affect the flow behavior. For example, the initial distribution of physical variables within the fluid sphere and the initial ratio of the kinetic energy to the internal energy can alter the dynamics of the flow even if the two parameters ER and AR are not varied. Nevertheless, all real ICF debris plasmas may be considered to be in a state with similar values of these additional parameters after only a small expansion has taken place. But as the simulations discussed in this chapter start out with a blob whose size is quite large on the scales of the thruster's dimensions, the parameters just mentioned can have a large effect on the subsequent hydrodynamics and interface motion.

The parameter ER was defined in an earlier chapter as the ratio of the field energy to the plasma energy when the plasma sphere is negligibly small compared to the size of the thruster. Here, the plasma energy should be defined as the total plasma energy, i.e., the sum of the kinetic and internal energies. The field energy was calculated approximately as

the energy contained in the volume of the thruster region, although it should be kept in mind that the plasma in a high-field thruster does not expand to fill the entire volume of the thruster. However, in these hydrodynamic simulations in which the expansion calculations are started from a uniform blob of 50 cm radius, it is more sensible to define the ratio ER at the state at which the plasma radius is 50 cm (i.e., the starting state), especially because in some of the cases run, the plasma was taken to be at rest at the 50 cm radius state. This was done. It did turn out, though, that for all cases discussed in this chapter, the difference between the magnetic field energy contained in the thruster when there is no plasma present and when the plasma has a 50 cm radius, was negligible, and also, that the difference in plasma energies between the two states was very small except in cases of very high fields where the difference was about 10 %.

Other than ER, the ratio of the characteristic interfacial magnetic pressure to the bulk fluid pressure at the starting state will also be used as a parameter characterizing the different cases run, and will be denoted by PR. Because the interfacial magnetic field strength varies substantially with polar angle, the characteristic interfacial magnetic field pressure at the starting state was arbitrarily specified to be the maximum value along the interface. The value of PR indicates the mismatch between the bulk and interface pressures at the start of the simulation, but because it was found that the (peak) interfacial magnetic pressure changes by less than an order of magnitude during a typical high-field expansion process while the bulk fluid pressure goes down roughly as the cube of the expansion radius, PR also contains information regarding the ratio of the applied interfacial pressure to the bulk pressure at subsequent stages of the expansion.

When a spherical perfect conductor with a radius of 50 cm exists in a thruster with a 650 cm radius current coil, the volume of the displaced magnetic fields is not negligible, and a significant surface current will be present on the sphere's interface. In this state, the interfacial magnetic field varied closely as the square of the current coil's current strength and thus also the square of the magnetic stream function value on the current coil's surface.

This contrasted with the case in which no plasma was present, because then, the magnetic field strength varied directly as the strength of the current in the field coil (see Chapter 4, for example). This means that PR will be proportional to the fourth power of the current in the current coil while the field energy in ER will vary as the square of the field coil's current. Because the plasma energy is a sum of the kinetic and internal energies and also because of the difference in the way by which the field varies with current when the plasma radius is negligible and when it is not, the ordering of PRs and ERs among different cases can even differ.

Next, the plasma really does work against the magnetic pressure acting on the interface, so the ratio, field energy/plasma energy, with the field energy calculated using a typical interfacial magnetic field acting on a 50 cm radius spherical interface would also be quite appropriate. The typical magnetic pressure acting on the 50 cm radius interface was arbitrarily chosen to be two thirds of the peak value at that radius. This pressure was multiplied by the volume of the thruster to obtain a newly defined field energy for the thruster. The field-to-plasma energy ratio at the starting state defined, using such a field energy, will be denoted by ER_i .

The maximum expansion distance of a plasma interface can be crudely estimated by dividing the plasma energy at the beginning of the expansion by some magnetic pressure value. Here, the interfacial magnetic pressure at angles where the interface turn around occurs should be used for the magnetic pressure instead of the thruster volume averaged value, because it is the interfacial magnetic pressure, and not the volume-averaged thruster field energy, which determines the expansion characteristics of the plasma.

7.5 Discussion of Results *

In this section, a few of the results from some of the cases run will be presented and discussed to illustrate the highlights of the flow phenomena. For all cases, the starting state was given a uniform density and pressure. The velocity at the starting state was either set to zero everywhere or was specified to be purely radial and linearly increasing towards the interface from zero value at the origin.

While working with simulations employing the baseline parameters (see Case 3 below), it became apparent that if parameters are chosen in such a way that the code encounters difficulties in handling the flow during the course of its run, spurious results such as negative interfacial pressure spikes, excessive numerical oscillations, and the formation of empty cells (or nearly empty cells interpreted by the computer as empty cells) could arise. Another problem was the inability to treat through all important phases of the expansion and redirection flow process for large expansion ratios, due primarily to memory limitations, but the limitation on CPU time was also a problem here. In order to avoid some of these problems and to better understand the generic behavior of the fluid under various choices of parameters as well as to explore the capabilities of the code, the code was run over with many different parameters. This was done at the expense of parameters deviating, even more than those of the base-line values, from those appropriate for simulations of flows in the thrusters of proposed ICF pulse rockets. The first simulation to be described is precisely a run that was made to observe good hydrodynamic calculations

* Although dimensionless variables may be scaled to represent any absolute magnitude, the variables in this chapter were quoted with particular choices for their absolute magnitudes to provide a better feel for the physical situation that might be represented by the set of parameters employed in a simulation. Furthermore, there are phenomena such as radiation, recombination, resistivities, etc., which depend explicitly upon the absolute magnitudes of the variables representing the state.

not obscured by exaggerated numerical artifacts, under the hardware-limited handling capacities of the code.

Crude estimates for the growth rate of Rayleigh-Taylor instabilities were made for the two-dimensional hydrodynamic runs of this chapter, in a fashion similar to those presented in the second part of the appendix to Chapter 3. The flute instability due to the bad curvature of the interface did not appear to be too dangerous, with estimated characteristic growth times typically a few tens of times longer than the characteristic expansion time scale of the plasma, at least for the lowest and slowest growing $n=1$ mode. However, estimates indicated that the time scale of growth for the flute instability due to the deceleration of the interface could be shorter than the characteristic time scale of expansion until the late phases of the expansion, by which time the interface would be expanded out to the characteristic dimensions of the thruster. This implies this mode of the instability to be a potential hazard to the runs discussed in this chapter (see Chapter 3). But because of the uneven distribution of deceleration strengths over the expansion period, and the inaccuracy of the data read off from noisy plots, these simple estimates are not entirely reliable either. We do not observe the development of plasma instabilities in the simulations, because for one thing, no detailed plasma physics which gives rise to such instabilities is incorporated into the code. However, the Rayleigh-Taylor instability analogous to those (flute instabilities) which arise in plasmas can occur in classical fluids as well. Numerical viscosities could have possibly been preventing (i.e., suppressing) these instabilities from manifesting themselves during the simulations, even had conditions been as such as to have permitted their growth during the expansion process.

7.5.1 Case 1

The parameters chosen for Case 1 can be considered to represent, for example, a case in which the uniform plasma of 50 cm radius at the starting state is at rest with a mass of 0.6 gram and a pressure of 7.67×10^9 dynes/cm². The sound speed at this starting state is approximately 1×10^8 cm/sec and this gives an interface expansion velocity on the same order as that of the base-line case. The current in the field coil was chosen to correspond to roughly 28.5 MA under this normalization. The pressure difference across the interface at the start of the simulation allowed the interface to initiate an outward expansion. The values of ER, ER_i, and PR for this case are roughly 17, 3.85, and 0.004, respectively.

Figure 7.1 shows the plasma-vacuum interface contours at evenly spaced time steps in the upper physical half-plane bounded along the bottom by the central thrust axis. The interface in three-dimensional physical space is obtained by rotating these contours around the central axis. As a convention, the $\theta=0$ direction was taken to be to the right in all plots of the interface contours. The innermost contour is that of the 50 cm sphere from which the simulation is started, and the contours are being plotted every 1.4×10^{-7} seconds. The separation between the innermost two contours is less than those between the others because in this simulation, the fluid starts out from rest, and a fair number of time steps are required before a realistic velocity profile sets up numerically. The interface was still slowly moving out at the end of the simulation, and at $\theta=0$, the interface at the last time steps (not shown) appeared to be oscillating with a small amplitude around a fixed location, although this is most likely a numerical phenomenon. Compared to the base-line case, the effect of the fields is stronger in this case, and the interface contours are more flattened towards the central axis in agreement with the general trend observed in the simulations employing the thin-shell approximation. Also, the interface in the direction of $\theta \approx 0$ is clearly less expanded than the almost freely expanding interface at $\theta \approx \pi$ (although even at

$\theta \approx \pi$, the internal fluid profile does differ substantially from that of free expansion because of the effect, especially at rather early epochs, of the finite interfacial pressure).

Figure 7.2 shows the nondimensional interfacial magnetic pressure (normalized the same way as the fluid pressure) as a function of polar angle in radians every 2.8×10^{-7} seconds, starting with the beginning of the simulation. As the expansion proceeds, the magnetic pressure at the interface becomes increasingly narrowly peaked in the direction of the field coil (this is more evident in the plots for Case 3). The field lines at polar angles generally in the direction of the field coil (roughly the direction in which the field is highest) get compressed as the interface expands in that direction, and the magnetic pressure rises with the expansion at such polar angles. However, along polar angles not in the general direction of the field coil (direction in which the field is relatively weak), the interfacial magnetic pressure decreases with the expansion of the interface (it should be realized that there, the interface is generally increasing its distance away from the field coils). The range of polar angles for which the latter behavior is observed depends not only upon the geometry of the field-coil setup and the relative positioning of the ICF explosion site, but also on the ratio of the field energy to the plasma energy. The lower the effect of the fields on the interface motion, the wider the range of angles at which the latter behavior is observed.

Figures 7.3 through 7.7 show the nondimensionalized pressure, density, radial velocity, velocity in the polar direction, and temperature as functions of radial distance along cells centered at a) $\theta=0.5\Delta\theta$, b) $\theta=9.5\Delta\theta$, c) $\theta=19.5\Delta\theta$, and d) $\theta=29.5\Delta\theta$. All plots are being made every 2.8×10^{-7} seconds, initiating with the starting state. It can be seen that the information of the deceleration of the interface does not propagate significantly into the interior of the blob during the characteristic expansion time scale, and that a large volume of the fluid remains unaffected by the interface deceleration. In this particular simulation, the cellwise quantities were averaged over neighboring cells not only in the polar direction but also in the radial direction in the first 10 radial grid zones, and this

switch in averaging procedures at $r=10\Delta r$ has resulted in, as mentioned earlier, the rather broad dips in the pressure and density profiles and the break in the slope of the radial velocity profile in the vicinity of that coordinate. The false and problematic profiles around $r=0$, such as spikes and dips in pressure, density, and temperature, and breaks in the slope of velocity, are due to the difficulty of treating FSPs touching the origin previously mentioned in Section 7.2. However, the mass tied up in the density spike at the origin is minimal. Narrow numerical spikes are often observed at the interface, but these are not to be confused with the usually somewhat broader real rises in pressure, density, temperature, and/or velocity that manifest themselves there, as the interface continues to expand out. Both the direction and absolute magnitude of the polar direction velocity, v_θ , displayed in the plots are very unreliable while the magnitude of v_θ is much smaller than that of the radial velocity, especially if the feature is in the form of a narrow spike, and are mostly numerical in origin. This was evidenced by comparing the results with those that were due to a free-expansion simulation. However, because this inaccuracy in the calculation of v_θ occurs mostly while $|v_\theta| \ll |v_r|$, it is not very crucial, although even at larger magnitudes of v_θ , the profiles close to the central axis are often not real, if they are narrow.

The deceleration of the interface is clearly shown in the radial velocity plots. Once the expansion starts, the radial velocity value at the interface quickly levels off and starts decreasing. As the deceleration of the interface begins and the bulk pressure and density immediately behind the interface fall low, a pressure and density rise region emerges from the interfacial region. As a characteristic analysis will show, the deceleration of the interface causes a compression wave that travels inwards relative to the outward flowing fluid to form, and with time, this steepens into an inward facing shock wave. This will usually occur some distance away from the interface. The steepening compression wave is first observed as a small density and pressure hump close to the interface. The shock will mark the inner edge of a region in which the density rises conspicuously above the rest, and hereafter referred to as the “shell.” The shell will obviously not necessarily have

uniform properties within. Now, depending upon the polar angle, the shock, and thus the shell, either travels away from the interfacial region, or doesn't. This is because the way by which the externally applied interfacial pressure varies with the expansion differs depending upon the polar angle. Once the shock forms, it is the shock, positioned a finite distance away from the interface, across which the outflowing fluid experiences its primary deceleration, and in later stages, redirection as well. The region of finite width just outside the shock with lowered (and later, reversed, at some angles) radial velocity is the shell. The strength of deceleration of the fluid across the shock, relative to the dynamics (e.g., strength of acceleration or deceleration) of the interface is important in determining the type of profile that is obtained here.

For a range of polar angles pointing roughly in the direction of the field coil (e.g., $\theta \approx \pi/3$), the interface decelerates (strongly) as it expands, and this generally results in the shock, which abruptly decelerates the flow behind the interface, not being able to increase its distance from the interface very much with time. The shock stays close to the interface, maintaining a relatively fixed distance from the interface, and the conspicuous density and pressure rise (a large fraction of it occurring across the shock) stays confined to a narrow region situated against the interface. Namely, the shell (and the pressure peak region) extends up to the interface, and the shell appears to "stay" at the interfacial region. Mass accumulates into the shell as the interface deceleration continues. The radial velocity drops sharply with radius across the shock from its peak value, a short distance behind the interface, and the radial velocity of the interface decreases steadily as the expansion proceeds. However, once the interface slows drastically or starts moving inwards, the pressure and density maxima can obviously increase their distance (i.e., travel away) from the interface, even at these polar angles.

At polar angles which point away from the general direction of the field coil, the density and pressure hump just mentioned starts "separating away" from the interface while retaining a relatively constant height. The pressure gradually tapers downwards from the

peak at the shock to the now lower value at the interface. The general qualitative features of the density profiles are similar to those of the pressure profiles, but the density rise region is narrower than that of the pressure, and falls off more sharply a shorter distance away from the shock on its downstream side. A low density region forms between the interface and the density hump, which is bounded on its inner side by the shock. It should be apparent by now that the double-peaked structure observed in the density and pressure plots of our three-dimensional expansion problem of a finite blob is really due to the outer spike at the interface's being numerical and the inner one's being caused by a real shock. The temperature stays more or less uniform, or even rises noticeably (part of the way) towards the interface in a region far broader than the shock transition region (unlike the pressure and density profiles, and revealing to some extent, the deceleration and heating history), on the downstream side of the shock, and this may be helpful in preventing field diffusions into the plasma. Now, in the early phases of interface deceleration, the region of decreased radial velocity spread inwards from the interface, relative to the fluid flow, but by the time the shell is observed increasing its distance from the interface, the radial velocity in this region exterior to the shock acquires a profile that rises with radius towards the interface. A real double-peaked radial velocity profile comes into existence. This region stretching from the shock towards the interface, and possessing a negative pressure and density gradient and a positive radial velocity gradient (i.e., a rarefaction region), commonly tends to arise from an outward acceleration of the interface. The shock decelerating the outward fluid flow, and which forms as a result of the interface deceleration at an earlier epoch, now, with the acceleration of the interface, noticeably increases its distance away from the interface with time. The shell is thus observed traveling away from the interfacial region. Note that the shell is here not comprised of the same material. Mass is flowing through it. The noise and large numerical oscillations plaguing the profiles of the variables (especially that of v_r) in the low density region downstream of the shock are due to the low number of particles per cell there.

Strong decelerations of the interface tend to be observed when the applied interfacial pressure increases with the expansion of the interface (although the two phenomena are not necessarily equivalent, and for example, it is possible for the interface to decelerate even as the applied interfacial pressure decreases with the expansion), and interface accelerations often occur when the applied interfacial pressure decreases at a rather high rate with the expansion of the interface. In fact, at polar angles for which the shell moves away from the interfacial region, the pressure profile between the shock and the interface, in a sense, very crudely reveals the pressure history at the interface. So a trend was often observed in which the shell tended to "stay" at the interface for angles at which the interfacial pressure increased with increasing coordinate of the interface (although, again, these two phenomena are not necessarily equivalent), and to "separate away" from the interface at angles where the interfacial pressure decreased strongly with increasing interface coordinate.

A few remarks should be made concerning the observation of the shock. The radial locations of the near-discontinuous rises in pressure, density, and temperature, and the sharp decrease in radial velocity were all found to coincide, and the near-discontinuity had a width of a few (2 to 3) grid cells, and this thickness stayed fairly constant throughout the simulation. The flow into the near-discontinuity was supersonic, while the flow out of it was subsonic. These are all characteristic features of a shock, but how well the Rankine-Hugoniot relations are satisfied should be investigated. Unfortunately, this turns out to be a difficult task to carry out because of the high noise component in the plots at the shock and the narrowness and sharp tapering of the density and pressure rise regions on the downstream side of the shock.** The difficulty is further aggravated by the physical variables' not always rising to their true values, and it was almost impossible to deduce from the plots, the "true" heights of the peaks that were due to the shock. However, it is

** For example, with a planar normal shock in a gas with $\gamma=5/3$, an infinitely strong shock will have a density ratio of 4, so any density ratio observed across such a shock with a value higher than this, we know is partly numerical.

possible to substitute the values/ratio of a physical variable (for which a relatively accurate measurement can be made off the plot) across the near-discontinuity, as read off from the plots, into the Rankine-Hugoniot relation to determine the values/ratio of another physical variable across the shock, and to see if the latter result is compatible with the plots for that variable. By taking this approach, it was found that the relation of physical variables across the near-discontinuity in the plots could be explained by the Rankine-Hugoniot relations. For Case 1, the Mach number of the shock was generally in the 2 to 3+ range at most polar angles. This relative constancy of shock strength with angle was observed to some extent in many of the runs, but there was also some trend of the shock showing a tendency of being strongest at the highest field region where the deceleration is strongest during the simulations. Regarding the type of shock that was occurring, the radial coordinate of the shock was roughly constant with polar angle (a spherical shock), at least while the interface deformation was not extremely severe, and thus the shock appeared to be approximately normal to the direction of flow. This property often prevailed, unless the deformation of the interface became severe. The fluid velocity in the polar direction generally did not vary through the shock for the cases investigated.

In the present run, the fluid pressure is significantly higher than the interfacial magnetic pressure at the starting state, but because of the large expansion experienced, the bulk fluid pressure falls well below the typical interfacial magnetic pressure by the end of the simulation. However, the use of a large (50 cm radius) initial sphere, with uniform density, and an expansion ratio that is not enormous is leading to a density distribution in which the bulk density is not negligible compared to the peak density of the region heavily affected by the application of a finite interfacial pressure. The interfacial fluid pressure first undershoots the applied magnetic pressure, but is rising towards the latter's value at the end of the simulation. It is evident that the interfacial fluid pressure is attempting to keep track of the applied pressure there.

Next, Figures 7.8 through 7.12 are contour plots of fluid pressure, density, radial velocity, velocity in the polar direction, and temperature in the rectangular r - θ plane. The plots are again presented every 2.8×10^{-7} seconds. The lower horizontal boundary of these plots marks $\theta=0$, while the upper horizontal boundary corresponds to $\theta=\pi$, and the vertical boundary to the left represents $r=0$. The r and θ axes are scaled linearly. At each time step for which these plots were made, a maximum and minimum value of the physical variable for which contours would be plotted were selected, and between them were chosen nine equally spaced contour values. The maximum and minimum contour values were chosen (except in the v_θ plots) so as to leave peaks that are clearly numerical artifacts, out of the plotting range. However, they do not necessarily represent the exact theoretical maximum and minimum values of the quantity at each of the time steps. These maximum and minimum values are listed as "max" and "min" beside the individual plots.

Figures 7.13a and b show, respectively, the cumulative impulse delivered to the vehicle (proportional to the temporal thrust efficiency) and the (temporal) jet kinetic energy efficiency as functions of time (note that these efficiency curves are taken out to later times than either the interface contour plots or the radial profile plots of the physical variables, but the calculation beyond 2×10^{-6} seconds is not reliable due to the formation of nearly empty cells within the fluid boundaries). The definition of these quantities are as described in the chapter on two-dimensional calculations carried out under the thin-shell approximation, but the jet kinetic energy efficiency here is referenced to the total plasma energy at the starting state, because a continuous conversion between internal and kinetic energies exists once the expansion is initiated. The thrust efficiency can be obtained by dividing the cumulative impulse delivered to the vehicle by the maximum impulse the plasma can transfer to the vehicle. Denoting the total plasma energy at the start by E_{tot} , the maximum impulse transferable to the vehicle, which is when the entire plasma energy is converted into kinetic energy and the exit flow has a uniform and purely axial velocity field with no forward jet, is $M_{\text{tot}} v_0$, where $v_0 = \sqrt{2 E_{\text{tot}} / M_{\text{tot}}}$. The thrust efficiency calculated this way using the

cumulative impulse value at the end of the simulation was roughly 2.5 %. The thrusting process has not completed by the end of the simulation, so the efficiency values obtained are not final.

The instantaneous thrust can be calculated by taking the time derivative of the impulse being delivered to the vehicle. This amounts to finding the slope of the cumulative impulse curve. In this run, the instantaneous thrust rose to about 3.9×10^{12} dynes during the early stages of the expansion and then to about 2.7×10^{12} dynes at a later stage (although based on an impulse curve whose profile is not entirely real, due to numerical artifacts). Recalling how thrust was calculated for these magnetic thrusters, one realizes that under fixed geometry and field-coil current strengths, the thrust and impulse delivered to the vehicle depend only upon the interface radii and shape (and its evolution) as long as the plasma currents are limited to surface currents, as is the case with a perfectly conducting unmagnetized fluid. However, the efficiency values also depend upon the internal properties of the entire fluid blob.

Depending upon internal flow conditions, the propulsive efficiency curves occasionally exhibited a profile wherein the curves rose sharply towards a local maximum and then decreased, after which they leveled off in some cases, only to start rising again later. When the thrusting process is over, the propulsive efficiency curves will level off at their respective final values. The first maximum of the propulsive efficiency curves, when it manifested itself, appeared at about the time the compression wave traveling inwards relative to the fluid started becoming clearly visible on the pressure and density plots, which was also about the time the radial velocity started developing a double-peaked structure at relevant angles. A decrease with time of the cumulative impulse being delivered to the vehicle (a negative slope of the impulse curve) means negative thrust, but an observation of the geometrical profile of the plasma interface and the location of the field coil relative to the interface reveals that the negative thrust obtained is not real and is due to numerical inaccuracies. Therefore this says that the form of the curves displaying the

temporal development of the cumulative impulse is not always being produced accurately. The same applies to the curves of jet kinetic energy efficiency, although a negative slope of these curves does not necessarily imply negative thrust. It should be kept in mind that the thrust producing mechanism can be explained entirely in terms of the interaction between the current through the field coil and the plasma surface currents, which is repulsive.

However, there generally will be an overall reduction of efficiencies over a model in which no conversion of kinetic energy to thermal energy is allowed. Across a shock or other temperature-raising profile, part of the kinetic energy of the inflowing fluid is converted into internal energy, i.e., energy in random thermal motions including nontranslational excitation modes, so if the conversion back to kinetic energy (radiation losses included) by the end of the thrusting process is not complete, and if it is not carried out in a way such that the resultant directed flow is as if no kinetic to thermal conversion had taken place, there will be a reduction in the propulsive efficiencies. This can be considered as a form of frozen flow loss, but in using this terminology, it should be kept in mind that for redirection and collimation of flow in a magnetic thruster, flow in the state of an ionized plasma is necessary, contrasting with the case of flow in a nozzle with material walls. Note that the energy given to the field will be returned to the plasma when the field returns to its original configuration.

Finally, Figure 7.14 shows the total plasma energy as a function of time. In all the total energy plots of this chapter, the vertical axis has been labeled in terms of ϵ' divided by approximately 240π . Each unit of time on the horizontal axis corresponds to 1.4×10^{-9} seconds for the present case. The small rise in the curve at the beginning is caused by an underestimation of the radial velocity due to the code requiring a finite number of time steps to achieve a theoretical velocity profile from a state at rest. The time at which the break in the total energy curve's slope occurs from shallow to steep, and thus marking the onset of strong deceleration by the fields, is coincident with the time the compression wave traveling against the outward fluid flow starts to manifest itself prominently over a wide range of

polar angles. With no radiation assumed, the only energy loss from the plasma is due to its doing work against the vacuum magnetic field, and the sum of the plasma and (quasi-static) vacuum field energies remain constant with time. Therefore, as the interface moves out and the fields get compressed, the field energy increases exactly by the same amount the total plasma energy decreases. As the fields recoil, the plasma energy given up to the fields due to the compression of the latter gets returned to the plasma, and the plasma energy curve starts to rise (see Case 9). But because the starting states of our simulations were taken to be the state when the plasma radius is 50 cm, the curve can eventually rise to values above that of the starting state if the recoiling interface is pushed down to low heights above the central axis. Contrasting with the thin-shell model, the fluid model allows some of the kinetic energy of the decelerating plasma to be converted into heat, and this leads to some interesting phenomena's being observed. The spikes seen near the end of Figures 7.13b and 7.14 are numerical.

In summary, a rather narrow shell-like region was found to form in the expansion process, which starts out like a free expansion but has a decelerating force acting at the interface. As long as there is a finite external pressure applying at the interface, no matter how low, a shell-like structure will generally form in the cases concerned, unless the expansion ratio is very low, because of the rapid fall-off ($\propto 1/R^3$) of the bulk density with expansion. However, this shell-like region, defined to be a region of substantial density rise above the rest and where a large fraction of the total plasma mass is concentrated, does not necessarily reside against the interface. The shell-like structure is the density and pressure rise region immediately downstream of the shock. The thicknesses of the density and pressure rise regions were found to stay thin for reasonably realistic choices of parameters, and the density rise region's thickness stayed relatively constant with time for many of the cases run, although this can depend upon the fluid dynamics and thus on the parameters selected for the run.

7.5.2 Case 2

In order to discern numerical features present in the plots from real features, to better understand the physics of the flow, and to obtain a better feel for the accuracy of the results, a simulation treating free expansion was run, using the parameters of Case 1 except for the absence of an external pressure applying at the interface. The averaging scheme for the cellwise quantities was also taken to be the same as that used in Case 1.

Figures 7.15 through 7.18 show, respectively, the nondimensionalized pressure, density, radial velocity, velocity in the polar direction, and temperature as functions of radius, with plots being presented at the same time steps as in the figures of Case 1. Comparison of the plots of Cases 1 and 2 helps highlight the real features of the flow in Case 1 that are due to the application of pressure at the interface. Note that the spikes at the interface are still observed for the free-expansion case, substantiating their numerical origin. The shock and shell present in the cases with finite applied interfacial pressure do not appear here, proving that they are real features. The profiles interior to the location to which information of the finiteness of the applied interfacial pressure, and in particular, the deceleration of the interface, has propagated in the runs with applied interfacial pressure are identical with those from the free-expansion runs. An outwardly rising, linear profile of the radial velocity sets up in the region overrun by the inward-traveling rarefaction wave, and with reflections of this wave, eventually the density in the interior of the bulk at late time steps acquires a very crudely uniform distribution as expected. However, the free-expansion calculations are not entirely reliable from early time steps because nearly empty cells form by virtue of zero applied pressure at the boundary, and for one thing, noticeable numerical heating is observed at the interface where the particle density is low and the inaccuracy high.

In this free-expansion simulation, the interface velocity stabilized at roughly 1.8 times the sound speed of the uniform blob at the starting state. However, as mentioned

earlier, PIC codes cannot treat densities below a particular limit, and the extended region of very low density leading to the force-free vacuum interface cannot be handled properly. Firstly, the interface in the plots are not observed at the theoretical locations predicted under the hydrodynamical approximation. The PIC simulation's interface is not expanded out as much as it should be, and this discrepancy increases with time. Furthermore, even at the location of the interface given by the PIC calculation, the peak of the simulated velocity profile is cut off short of the theoretical value at that location. Thus the apparent velocity of the free-expansion interface, as obtained from a PIC calculation, underestimates the theoretical value quite notably. As we have always seen, the physics at the interface can be handled more accurately if the density there is higher.

Based on this free-expansion simulation, the errors in the cumulative impulse and jet kinetic energy efficiency calculations for the cases with finite interfacial pressure were estimated to be on the order of a percent or two, and the inaccuracy in the energy conservation of the code, less than about 15 %.

7.5.3 Case 3

In this run, the starting state plasma is ascribed with the base-line parameters described earlier. The current in the single field coil was taken to be very approximately 22 MA. The ER, ER_i , and PR values for this case are roughly 4.2, 0.56, and 1.5, respectively. The ratio of the kinetic energy to the internal energy at the starting state is about 2.6×10^3 .

Figure 7.19 shows the evolution of the interface contours in physical space with plots being made every 1.4×10^{-7} seconds, starting with that of the initial 50 cm radius sphere. The last contour is plotted 1.96×10^{-6} seconds after the start of the simulation.

In many of the cases run, including this run, the largeness of the expansion ratio caused empty cells to arise in the interior regions of the fluid, where, in reality, there

should be a finite density. This eventually caused the code to fail and is the reason that some of the simulations whose results are presented in this chapter could not be run to their intended final time steps.

It can be seen that the portion of the interface located roughly along the line of sight of the coil, where the interfacial field is strongest, is decelerated the most, and that in these relatively high field thrusters, the deceleration and significant deformation of the interface occurs steadily throughout the expansion and thus at large distances away from the coil. The expansion in regions close to the central axis is nearly unimpeded for the present case. Under the hydrodynamic approximation, the flow exactly along the axis will be undecelerated because the central axis is coincident with a field line, but in our simulations, even the fluid cells closest to the axis are off-centered from the axis, so at stronger field strengths, retardation of the interface at $\theta \approx 0$ will be clearly observed. Actually, in a real plasma where particles possess finite Larmor radii, particles whose orbits are centered on the axis can be reflected as they move into regions of higher (e.g., constricted) field because of the conservation of an adiabatic invariant, the magnetic moment.

In Figure 7.20, the nondimensionalized magnetic pressure along the plasma interface is plotted as a function of polar angle. The plots are presented every 2.8×10^{-7} seconds, starting with the beginning of the simulation. As it turns out, both the absolute magnitude and distribution of magnetic pressure on the surface of a perfectly conducting sphere of 50 cm radius are rather similar to those on the surface of such a sphere with a radius of only a few cm. The maximum interfacial magnetic pressure goes up by only a factor of about four during this simulation, and as alluded to earlier, this is typical of expansion in high-field thrusters in which the initial field stores enough energy to decelerate and redirect the plasma without too much field compression.

The velocity plots revealed that the flow stays predominantly radial, even as it is decelerated significantly and redirected. The magnitude of the velocity in the polar direction, at least in this run, never approached that of the radial velocity. Namely, the

fluid elements are being decelerated and reversed in flow direction primarily along the radial direction in which they flowed outwards. Not only is this type of flow rather inefficient in terms of thrust production, it also implies the possibility of fluid initially flowing into the forward hemisphere not leaving the thruster in one "bounce," and rebounding several times while it thermalizes and eventually diffuses out across the thruster on a thermal time scale. Overall, the redirection process, complicated by having a blob filled with fluid, is not resulting in a highly axially collimated flow, but longer simulations taking the flow process into later stages of the expansion will be needed to obtain a more complete picture of the flow. Nevertheless, it is noticed that the general geometry and motion of the interface is similar to those obtained from simulations employing the thin-shell approximation. The imparting of a θ -component momentum to the fluid by the redirection process is far higher at $\theta \approx 2\pi/3$ than at $\theta \approx \pi/3$ because of the oblique angle of the flow vector against the natural contours of the fields at the former location. In fact, depending upon the parameters chosen, redirection with a rather high velocity component in the polar direction can occur at $\theta \approx 2\pi/3$.

Figures 7.21a and b show, respectively, the cumulative impulse delivered to the vehicle and the jet kinetic energy efficiency, as functions of time. Both curves are still rising at the end of the simulation, indicating that the thrusting process has not yet completed by that time. This is consistent with the observation that the interface did not come to a halt and reverse velocity at any polar angle during the run. The instantaneous thrust at late stages of the simulation was estimated at about 1.2×10^{13} dynes. The thrust efficiency registered a value of 9.1 % at the end of the simulation. Considering the fact that an unrealistic starting state was chosen for the expansion, the efficiency values obtained are not bad. With not all material collected into a shell and not all of the plasma energy at any instant being in the form of kinetic energy, the hydrodynamic model will yield lower efficiencies than the thin-shell-approximation model for the same pellet mass and energy, interface profile and radii, field-coil current strengths and radii, and geometrical setup of the

thruster. As noted earlier, the conversion of internal energy to kinetic energy will be mostly complete from very early on in the expansion in a real ICF thruster, and even in contrived situations where a 50 cm plasma starts out from rest, incomplete internal-to-kinetic energy conversions will not severely lower the thrust efficiency compared to the thin-shell case. It should be noted, though, that the thermalization of kinetic energy due to interface deceleration at later stages of the expansion, and especially shock heating, does have some effect in lowering the efficiency over models in which all of the plasma energy is always in the form of kinetic energy. This is mainly because of incomplete conversion back from thermal energy to kinetic energy during the time the plasma is still in the thruster's field of influence. Even if it is supposed that the conversion back is "complete," the total kinetic energy available in axially directed form will generally not be as high as had there been no conversion. However, it is important to realize that this thermalization constitutes an integral part of the deceleration/redirection process of the fluid, and a realistic hydrodynamic analysis cannot be discussed without it. Now, as an example, the internal energy, depending upon the angular slice, went up by two orders of magnitude during the deceleration process in this run, but with the internal energy usually being a small fraction of the kinetic energy, the efficiency loss will generally not be large unless the blob starts out from rest at a large radius. On the other hand, although a significant amount of mass is accumulated into a rather narrow shell-like region, the fact that a nonnegligible amount of mass remains in a region unaffected by the application of finite interfacial pressure (i.e., the low density but large volume region interior to the shock) does have a notable effect in lowering the efficiency from values obtainable by a thin-shell simulation, especially at the same (i.e., comparable) stage of expansion (although this means that the thrusting process in these hydrodynamic runs may continue on for quite some more time). In the case under concern, about 1/3 of the mass along the direction of $\theta=2\pi/3$ was not contained in the shell-like region during the run. A lot of the mass was ending up not being collected into the shell-like region because the simulation was started from a uniform sphere with a radius

that was large on the scales of the thruster. If the expansion were begun more realistically with a uniform blob of much smaller radius (e.g., 1 cm radius instead of 50 cm radius), the efficiencies obtained would have been higher than those quoted in this chapter. Namely, the choice of an unrealistic starting state for the simulations was the single largest factor contributing to the low efficiencies obtained by the fluid expansion calculations of this chapter. Still, even if a realistic initial condition were chosen, the hydrodynamic simulations are not expected to give efficiencies as high as those obtainable from simulations employing the thin-shell approximation. This is because of the involvement of the hydrodynamics of a fluid with a distribution of mass and velocities across the entire blob, rather than having just the dynamics of a shell containing all of the mass and always located precisely at the interface, exactly where the external pressure force is directly applied. Furthermore, if a forward jet emerges in a hydrodynamical simulation, it will not all be along the central axis. On the whole, as long as we have a "well-behaved" plasma, the efficiency values quoted in the earlier chapter on simulations carried out under the thin-shell approximation are expected to be a better estimate of the efficiencies that would be obtained by a real ICF pulse rocket thruster than are the values quoted in this chapter.

Finally, Figure 7.22 shows the temporal variation of the total energy of the plasma. Each unit of time on the horizontal axis corresponds to 1.4×10^{-9} seconds.

7.5.4 Case 4

In Sections 7.5.4 through 11, a quick run-through of various cases will be made. The cases have been ordered approximately by the level of decelerating influence the fields have on the expanding plasma, from weak to strong. It can be seen that the parameter ER_i has the strongest influence in determining this ordering of cases and that PR has a secondary influence. PR contains information regarding the internal energy of the fluid at the starting state, but not the kinetic energy, which can also be present from the start and be

much higher than the internal energy. ER itself does not characterize very much the type of flow that will be obtained. All this is because the interfacial pressure is what the plasma does work against, and in the case of our two-dimensional hydrodynamic simulations, ER does not relay this information well, while ER_i and PR do. The interface often does not expand to fill a thruster, and the total field energy contained in a thruster (note that in a high-field design, the fields do not get highly compressed, so the fractional variation of energy stored in the fields is not drastic during an expansion process) is not always a good indicator of the actual energy expended in displacing the fields, which is the field energy that really matters. And if the interface does expand to fill a thruster, the fields get so highly compressed upon velocity reversal of the interface that the energy stored in the fields by that stage becomes very much higher than that at the initial state (which means that in the very low field thrusters such as that of the Daedalus, the field-to-plasma energy ratios will not be very important parameters, while the geometry of the coil arrangement will). Actually in the (high-field) thin-shell simulations too, it is really ER_i that matters, but because in those simulations the plasma at the initial state (when the parameters ER and ER_i are defined) was chosen negligibly small compared to the dimensions of the thruster, ER and ER_i ended up being proportional to each other with the same proportionality constant for all cases (the latter was smaller than the former by a factor of about 15), and one could talk about ER_i and ER in the same way.

Case 4 can be considered to be representing a situation in which the starting state's 50 cm radius sphere of uniform density and pressure is at rest with a pressure value of 1.6×10^{10} dynes / cm^2 . Under this normalization, the mass of the plasma is 68 grams and the current in the field coil is 22 MA. The sound speed in the bulk at the starting state is 1.4×10^7 cm/sec. The parameters ER, ER_i , and PR are approximately 5, 0.67, and 0.0007, respectively. The state specified in this run is very far removed from that to be found in a thruster of an ICF pulse rocket, especially because the plasma born with zero

bulk velocity and a radius of 50 cm is endowed with a very low temperature. However, pressure matching across the interface can be achieved with relative ease although PR was chosen a little on the low side.

Figure 7.23 shows the interface contours in physical space every 7×10^{-7} seconds, starting with that of the 50 cm sphere. Note that despite the rather high ER value, the interface motion is far less influenced by the fields than cases with lower ER value but higher ER_i (or PR) value.

The thrust and jet kinetic energy efficiencies were still low at about 1.5 % (peak value during the simulation), but the major thrusting phase has not yet been reached by the end of the simulation. The inward-traveling density and pressure rise regions are beginning to appear only towards the end of the simulation, and significant interface deceleration is only starting to set in.

7.5.5 Case 5

This run can be taken to represent a case in which the uniform plasma sphere of 50 cm radius at the starting state has a pressure of 3.44×10^9 dynes / cm² and a linear profile of the purely radial velocity stretching from zero value at the origin to 1.26×10^9 cm/sec at the interface. The mass of the plasma is 15 grams and the current in the field coil, 103 MA. The sound speed in the 50 cm sphere is 1.4×10^7 cm/sec. The ER, ER_i , and PR values for this case are 0.19, 0.55, and 1.5, respectively, and the ratio of the kinetic energy to the internal energy at the starting state is approximately 2.6×10^3 . Note that this case has almost the same ER_i and PR values as the base-line case (Case 3), although the ER value is very different, and, in fact, all results for this run, including the propulsive efficiencies, turned out to be very similar to those of the base-line run.

Figure 7.24 shows the interface contours in physical space at 7×10^{-8} second intervals, starting with that of the 50 cm sphere. The radial velocity at the interface was

found to stay at a rather constant value during the simulation except in the direction of the field coil. The (false) narrow protrusions seen emanating from the interface in the polar direction in Figure 7.24 are due to the formation of a very low density region behind the interface at angles pointing away from the coil as a rather broad density rise region travels away from the interface at those angles.

Figures 7.25a and b show, respectively, the temporal variations of the cumulative impulse delivered to the vehicle and the jet kinetic energy efficiency. Using the cumulative impulse delivered to the vehicle by the end of the simulation, although only a temporary value, a thrust efficiency of about 10 % was obtained. But because the interface was still continuing to expand at the end of the run, even at angles where the field is strongest, and the propulsive efficiency curves, still rising sharply, the final propulsive efficiencies are expected to be well above this ballpark.

7.5.6 Case 6

This run can be taken to be one in which the mass of the plasma is 0.6 gram, and the bulk pressure of the uniform 50 cm sphere is 7.67×10^9 dynes/cm² at the starting state. The plasma in this state is taken to be at rest, with a sound speed in the bulk measuring 10^8 cm/sec. The current in the field coil is 22 MA. The ER, ER_i, and PR values for this case are roughly 10.5, 1.36, and 0.0014, respectively.

Figure 7.26 shows the interface contours in physical space every 3.5×10^{-7} seconds, starting with that of the 50 cm sphere. Note that the interface is less decelerated than in the stronger field run of Case 1, which used the same plasma conditions for the starting state.

Recalling that the shell was defined to be a region of conspicuous density rise, the shell, where it stays "confined against" the interface, maintained a relatively constant thickness with time. The shock did not move notably away from the interface. The shock

in the expansion flow is strong, and a rather large negative gradient (with radius) of the radial velocity can exist from the shock to the interface. This is consistent with the strong deceleration due to the high fields at the relevant angles, and with theory. Where the shell moves away from the interface, the "shell" really does not correspond to the interfacial shell of the thin-shell model, but it was found to grow in thickness as $\delta R \propto R^\alpha$ with $\alpha > 1$ during its formation phase, where δR is the shell thickness, and to maintain a more or less constant thickness with time, once clearly established. The spreading of width was not large. This implies a shock that is not too strong, and is consistent with the fact that the shock formed earlier during the expansion phase when there was a rather weak deceleration of the interface, before "separating away" from the interfacial regions. As mentioned earlier, the characteristic width of the pressure rise region is broader than that of the density shell except at polar angles where the shell does not move away from the interface.

For this run, the Mach number of the shock, as estimated from the plots, fell in the 1.5 to 5 range, and often around 3, and did not vary drastically during the simulation.

The tentative thrust efficiency at around $t = 1.2 \times 10^{-6}$ seconds was approximately 1.4 %, and that at the end of the simulation was about 2.5 %. The slope of the impulse curve indicated an instantaneous thrust of about 2.2×10^{12} dynes at late stages of the expansion. The maximum instantaneous thrust peaked at about 5×10^{12} dynes during the simulation.

7.5.7 Case 7

This case can represent a situation in which the uniform 50 cm radius sphere at the starting state has a pressure of 3.44×10^9 dynes / cm² and a linearly rising (with radius) distribution of the purely radial velocity that reaches a peak value of roughly 1.6×10^8 cm/sec at the interface. Under this normalization, the mass of the plasma is 15 grams and the current in the field coil is approximately 70 MA. The sound speed in the bulk of the 50

cm sphere is 1.4×10^7 cm/sec. The ER, ER_i, and PR values for this case are 5, 6.6, and 0.3, respectively. The ratio of the kinetic energy to the internal energy at the starting state is about 45.

Figure 7.27 shows the interface contours in physical space every 2.8×10^{-7} seconds, starting with that of the 50 cm radius sphere. When parameters are chosen so that the interfacial magnetic pressure has a strong influence in decelerating the interface from small radii, the interface generally develops a rather flattened profile (see Cases 9 and 10 for exaggerated examples), as mentioned earlier. The flattened shape of the interface results from the strong interfacial magnetic pressure causing the expanding plasma to be quickly decelerated and redirected over a wide range of polar angles. On the other hand, if the effect of the field is weaker, strong interfacial deceleration occurs only in the direction of the coils and only as the interface approaches the coils rather closely.

Figures 7.28a and b show the temporal variations of the cumulative impulse delivered to the vehicle and the jet kinetic energy efficiency. Tentative propulsive efficiencies of 5 to 7 % were obtained at the end of the simulation.

The curve of the total plasma energy for this run was beginning to level off towards the end of the simulation, indicating that velocity reversal was about to occur. It should be pointed out that the inflection point in the plasma energy curve occurs roughly around the time the break in the slopes of the propulsive efficiency curves are observed. The instantaneous thrust lowered by about 30 % from the local peak value (of about 1×10^{14} dynes) after this point and more or less maintained that value during the later phases of the simulation.

7.5.8 Case 8

In this simulation, all parameters were kept the same as those of Case 5, except for the current in the field coil, which was changed to roughly 180 MA. The ER, ER_i, and

PR values for this case are 0.58, 5.2, and 14, respectively, and the ratio of the kinetic energy to the internal energy at the start of the simulation is about 45.

Figure 7.29 shows the interface contours in physical space every 7×10^{-8} seconds, starting with that of the 50 cm radius sphere. Most of the forward half of the interface has already been turned back by the recoiling fields at the end of the simulation. Note that the interface is being held back significantly even at $\theta=0$.

Figures 7.30a and b show the temporal variations of the cumulative impulse delivered to the vehicle and the jet kinetic energy efficiency. An instantaneous thrust value of about 6×10^{15} dynes was being developed during the simulation, and a thrust efficiency of about 11 % was reached towards the end of the simulation.

7.5.9 Case 9

The parameters chosen for this run can be considered to be representing a case in which the uniform 50 cm radius sphere at the starting state has a pressure of 1.6×10^{10} dynes / cm² and the same velocity distribution as that of Case 8. The mass of the plasma is 68 grams, and the current in the field coil is roughly 385 MA. The sound speed in the 50 cm sphere is 1.4×10^7 cm/sec. The ER, ER_i, and PR values are roughly 0.6, 24, and 70, respectively, and the ratio of the kinetic energy to the internal energy at the starting state is about 2.6×10^3 .

Figure 7.31 shows the interface contours in physical space every 7×10^{-8} seconds, starting with that of the 50 cm radius sphere. Reflection of the interface is observed. The relative importance of the values of ER, ER_i, and PR in determining the characteristics of the thruster, i.e., the dynamics of the interface motion, is clearly illustrated here. For example, in the present case, ER is low but the flattened interface profile being obtained is characteristic of a very high field device, as the ER_i (and PR) values are suggesting. Not only the (initially) forward flow but also the (initially) rearward flow (such as at $\theta \approx 2\pi/3$)

are being heavily affected and shaped by the fields, and in fact, unlike the cases in which the effect of the field is much weaker, the behavior of the flow (i.e., the radial profiles of variables) at $\theta \approx 2\pi/3$ is more similar to that at $\theta \approx \pi/3$ than those at $\theta \approx 0$ or π .

In Figure 7.32 we see the nondimensionalized magnetic pressure along the interface plotted as a function of polar angle every 7×10^{-8} seconds.

Figures 7.33 through 7.37 show the nondimensionalized pressure, density, radial velocity, velocity in the polar direction, and temperature as functions of radial distance along cells centered at a) $\theta=0.5\Delta\theta$, b) $\theta=9.5\Delta\theta$, c) $\theta=19.5\Delta\theta$, and d) $\theta=29.5\Delta\theta$. The time interval between the plots is the same as in Figure 7.32. Notice that these profiles are quite generic regardless of the specific parameter ratios chosen for the runs. Also observed in this simulation is the previously stated fact that the false negative interfacial pressure spike at $\theta \approx \pi/3$, even if present at earlier time steps, disappears as the simulation proceeds.

Figures 7.38a and b show the temporal variations of the cumulative impulse delivered to the vehicle and the jet kinetic energy efficiency. The propulsive efficiencies are still rising at the end of the simulation, but the tentative thrust efficiency calculated using the cumulative impulse value at the end of the simulation came out to about 8.3 %. The instantaneous thrust during the simulation reached values ranging from about 3.4×10^{16} dynes to 6.3×10^{16} dynes.

Figure 7.39 is a plot of the total plasma energy as a function of time, with each unit of time on the horizontal axis corresponding to 3.5×10^{-10} seconds. The curve is seen rising as the interface motion is reversed by the recoiling fields.

7.5.10 Case 10

The parameters chosen for this run can be considered to be representing a case in which the uniform 50 cm radius sphere of 15 grams at the starting state has a pressure of

3.44×10^9 dynes / cm², a linear distribution (with radius) of the purely radial velocity that rises to a value of 1.26×10^9 cm/sec at the interface, and a sound speed of 1.4×10^7 cm/sec. The strength of the current in the field coil is approximately 480 MA. The ER, ER_i, and PR values for this case are 4.2, 266, and 720, respectively. The ratio of the kinetic energy to the internal energy at the starting state is roughly 2.6×10^3 .

Figure 7.40a shows the interface contours in physical space every 1.75×10^{-8} seconds, starting with that of the 50 cm radius sphere. Figure 7.40b-i shows the first five contours of Figure 7.40a, while Figures 7.40b-ii through b-v show, respectively, the sixth through ninth contours (the tenth contour has not been plotted separately) of Figure 7.40a. Even the interface (initially) traveling outwards adjacent to $\theta=0$ is being reversed in flow direction. The peak of the thrusting process is over by the end of the simulation. It is expected that an outward rebound of the plasma portion initially traveling in the general direction of the field coil, but moving inwards by the end of the simulation, will occur, and that the plasma will not all flow out of the thruster in one bounce. At late time steps, the very low densities prevailing near the interface at $\theta \approx \pi$ and deep in the interior regions are limiting the reliability of the calculations.

In Figure 7.41, we see the nondimensionalized interfacial magnetic pressure as a function of polar angle every 3.5×10^{-8} seconds. With these thrusters in which the effect of the field on the plasma interface motion is very high, the maximum interfacial magnetic pressure occurs not along the line of sight of the field coil, but almost around $\theta \approx \pi/2$, and the interfacial magnetic pressure profile does not become narrowly peaked about its maximum. As the interface gets pushed down flatly against the central axis, the interfacial magnetic pressure acquires a somewhat uniform value over a broad range of angles. Despite the rather awkward shape of the interface with its highly apparent zigzags brought about through the use of a low-resolution underlying grid (actually the interface shape seen in Figure 7.40 is not exactly what is inputted to the field solving software), the field calculations were found to be decently reliable.

After interface reflection, the region of reversed radial velocity was seen spreading inwards from the interface, while the pressure and density rise regions, which were confined to the interface during the expansion phase, gradually started separating away from the interface. In this run, $|v_\theta|$ did acquire high values at late stages of the simulation, especially at $\theta \approx 2\pi/3$ (as high as 20% of $|v_r|$ at this angle).

Figures 7.42a and b show the temporal variations of the cumulative impulse delivered to the vehicle and the jet kinetic energy efficiency. The instantaneous thrust registered roughly 10^{16} dynes during a relatively early phase of the expansion, and about 1.5×10^{16} dynes during a later phase. The tentative thrust efficiency using the peak impulse delivered to the vehicle during the simulation period was about 6 %. The efficiencies of these extremely high field thrusters are generally not that high. The fluid gets pushed down too much against the central axis of the thruster (i.e., redirected with a downwards flow towards the axis instead of highly axially in the rearward direction) over a broad range of angles, and without too much preference of the $\theta \approx 0$ direction over the $\theta \approx \pi$ direction. Also, the smallness of the expansion ratio in the very high-field cases has an effect in lowering their efficiencies, because having started out from a uniform blob with a radius as large as 50 cm, a fluid distribution substantially different from a thin-shell type ends up developing. Even in this run, some of the fluid in the interior regions of the blob has not yet been influenced by the deceleration of the interface at the end of the simulation.

7.5.11 Case 11

This case can be considered to be one in which we have a 0.6 gram plasma whose pressure, when the plasma is a uniform 50 cm radius sphere at rest, is 7.67×10^9 dynes/cm². The sound speed in this 50 cm sphere is 10^8 cm/sec, and the current in the field coil is roughly 86 MA. The ER, ER_i, and PR values for this case are approximately 155, 320, and 0.33, respectively.

The expansion ratio for this case has been made very small, and unlike the case of a real ICF debris plasma expansion, a rather large fraction of the plasma energy remains in the form of internal energy. By the end of the simulation, the density and pressure rise region occupies a significant volume of the entire blob, although it hasn't developed yet in the $\theta \approx \pi$ direction. As alluded to earlier, the internal fluid dynamics with the propagation of sound waves and shock waves can have a marked influence on the interface motion here.

Figure 7.43a shows the interface contours in physical space every 8.4×10^{-8} seconds, starting with that of the 50 cm sphere. In Figure 7.43b, we see each interface contour plotted separately. The thrusting process is not over by the end of the simulation. At angles of highest field, a region of very low velocity forms behind the interface and the interface remains almost immobile for a long time (on the scales of the dynamic time of the expansion) during the (interface) velocity reversal process (see also Chapter 6). Even for these extremely high field cases, the overall deceleration of the flow close to $\theta \approx \pi$ is small.

A rather good match (estimated error typically less than 15 %) between the interfacial fluid pressure and the interfacial magnetic pressure was observed at the individual values of polar angle in this run. The Mach number for the shock, as estimated from the plots, generally fell in the 1.5 to 5 range, and often around 3. The shock has not appeared yet at $\theta \approx \pi$ by the end of the simulation. The efficiency values obtained were on the order of those of Case 10.

7.6 Summary

The two-dimensional simulations discussed in this chapter treated the flow of an unmagnetized perfectly conducting fluid in a single-coil magnetic thruster with AR=1, under the hydrodynamic approximation. Limitations imposed by finite difference

calculations did not always allow the best results to be observed, but interface capturing did appear to be proceeding successfully (or at least in the very process of doing so), as well as was the capturing of shocks. It was found that a narrow shell-like region does form at all polar angles, unless a very small expansion ratio was chosen, because of the deceleration of the interface resulting from the finite applied pressure at the interface, and also the rapid decrease of the bulk density and pressure with the expansion. The deceleration of the interface causes an inward-facing shock to develop behind the interface (in the predominantly radial flows that prevailed, the shock appeared to be mainly normal to the direction of flow), and this shock marks the inner edge of the "shell." This shell-like region, defined to be a region where the density is substantially higher than the rest, however, did not necessarily "stay" at the interface during the expansion flow process. Generally, where the interface continued to decelerate (strongly), which was usually where the applied interfacial pressure rose with the expansion, the "shell" stayed against the interface, but at other polar angles, commonly where the interface started to accelerate after the shell formed, the "shell" moved inwards, away from the interface. The pressure and temperature also rose in the "shell," like the density, primarily due to flow passage through the shock, but more downstream of the shock, the pressure and density decreased towards the interface at angles where the shock traveled noticeably away from the interface. The temperature, however, after rising at the "shell," maintained a high value, all the way up to the interface. Of course, as alluded to above, the primary contribution to the temperature rise of the fluid between the shell and the interface was coming from shock heating. The high temperature at (and downstream of) the shell is a welcoming feature of the flow for propulsion (magnetic thruster) applications, because *as long as* the plasma is highly ionized, resistivity decreases with increasing temperature.

Due to limitations in the computer memory available, the expansion simulations were started from an unrealistic state of a uniform blob possessing a very large (in terms of the thruster dimensions) radius. This resulted in a large amount of mass being present in

the interior regions not influenced by the finite applied pressure at the interface, and this was the major reason low efficiencies were obtained during the simulations of this chapter (simulations also often terminated early, before the thrusting process was over, because of hardware imposed limitations), especially in relation to the efficiencies obtained with the thin-shell model. Other less important factors contributing to the lowness of the efficiencies presented in this chapter, relative to those obtained through simulations employing the thin-shell approximation, included the conversion of kinetic energy into internal energy (which is not allowed in the thin-shell model). In the thruster design studied, the flow pattern was also found not to be the most efficient, with fluid elements traveling initially into the forward hemisphere having a tendency of being reflected primarily back along the direction they came from (without receiving too much θ -component momentum upon reflection). A more axially collimated outflow is desirable. Overall, as long as the simplifying assumptions used in this work may be considered valid, i.e., that the plasma behaves as a perfectly conducting fluid obeying the classical laws of hydrodynamics, the efficiencies obtained by the thin-shell simulations of Chapter 4 are believed to be giving a more realistic estimate of the real efficiencies under idealized conditions, than those obtained in this chapter employing unrealistic parameters.

The qualitative features of the motion and deformation of the interface were similar to those obtained through the thin-shell approximation. Flow restriction and confinement in the forward direction appeared to be good (with the thin-shell calculations, this is more difficult to observe because of particles falling onto the central thrust axis). However, it was the field-to-plasma energy ratio, ER_i , with the field energy calculated using the pressure acting on the interface, that determined the behavior of the interface. This was because it is the field pressure at the interface that the plasma must actually work against, and the field-to-plasma energy ratio, ER , defined with the magnetic field energy contained in the entire thruster, is not a good indicator of the flow dynamics. Actually, the same holds for the thin-shell case, but there, due to having started the plasma expansion from a

very small radius, the two definitions of field-to-plasma energy ratio reduced to quantities proportional to each other. In fact, it should be recalled that Equations (4.8) through (4.10) in Chapter 4 were derived assuming the initial plasma radius to be negligible on the scales of the thruster dimensions, and therefore no characteristic initial plasma size existed in the analysis. This led to ER's being a valid and important quantity determining the behavior of the thruster.

In order to treat flows that more closely resemble those to be found in the thrusters of ICF pulse rockets, the initial state at which uniform properties (uniform for simplicity, under lack of good ICF debris data) are to be assumed must be made very small relative to the characteristic size of the thruster. Of course, a possible alternative is to start the simulations from large radii and with the correct (nonuniform) internal distributions, but this will be more difficult. Future work should concentrate on modifying the algorithm to incorporate nonuniform, moving grids to enable the treatment of larger, more realistic expansion ratios and to include the effects of finite resistivities and to upgrade the code to MHD. The latter is especially important because the plasma will become quite resistive as it cools drastically with the expansion, and will thus allow field diffusions to occur. In subsequent work, neglected physics of radiation and recombination should also be considered. Relaxation of the assumption of azimuthal symmetry, and working in full three dimensions will also be an undertaking of importance. This is especially so because some important instabilities and some other plasma phenomena, including plasma rotations, appear only by virtue of having a fully three-dimensional flow geometry.

References for Chapter 7

1. R.A. HYDE, "A Laser Fusion Rocket for Interplanetary Propulsion," Paper IAF-83-396, 34th Congress of the Int'l Astronautical Federation, Budapest, Oct. 1983.
2. A. NISHIGUCHI and T. YABE, *J. Comp. Phys.* **52** (1983), 390.

Appendix of Chapter 7

Figures for Chapter 7

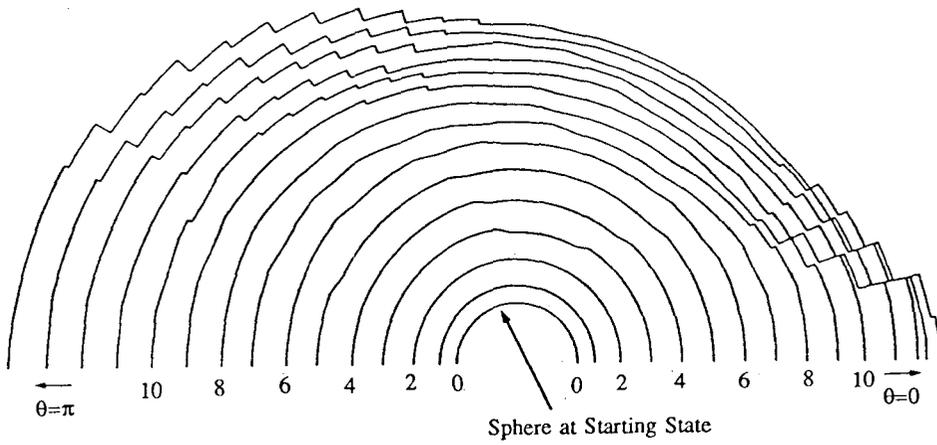
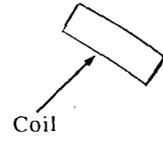


Figure 7.1 Locations of the interface at evenly spaced time intervals for Case 1.
($ER \approx 17$, $ER_i \approx 3.85$, $PR \approx 0.004$)

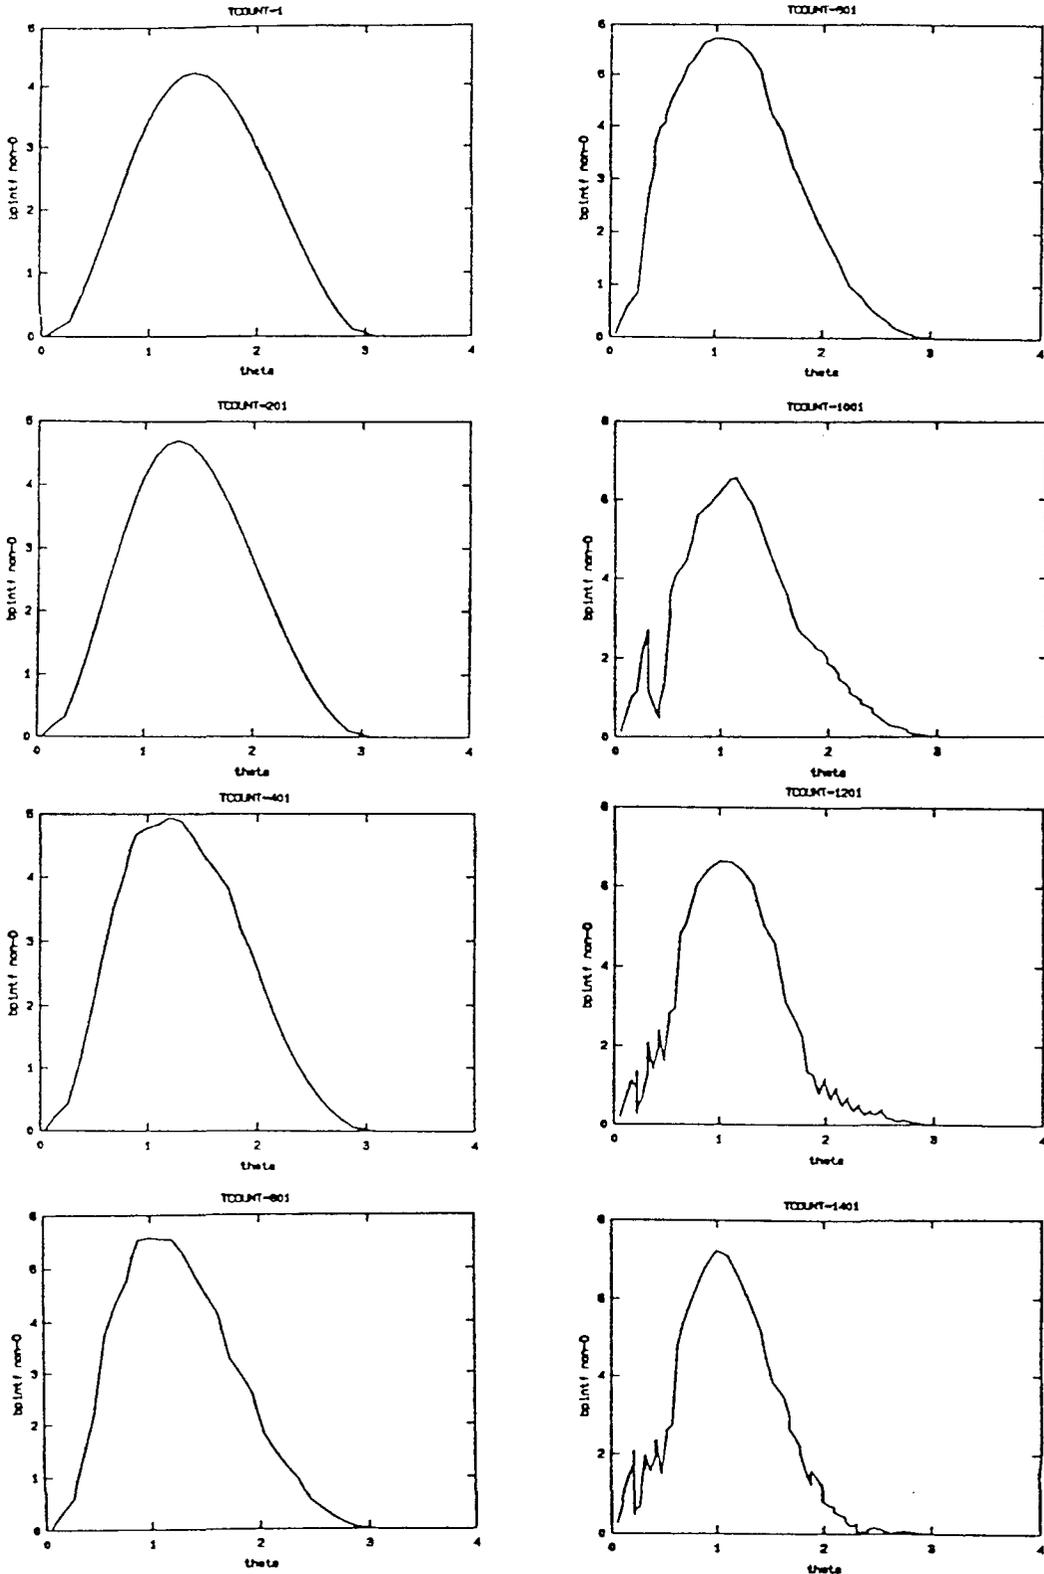


Figure 7.2 Nondimensionalized interfacial magnetic pressure (p_B') (vertical axis) v.s. polar angle (θ) in radians (horizontal axis). (As a convention, figures are arranged top to bottom, and then left to right.)

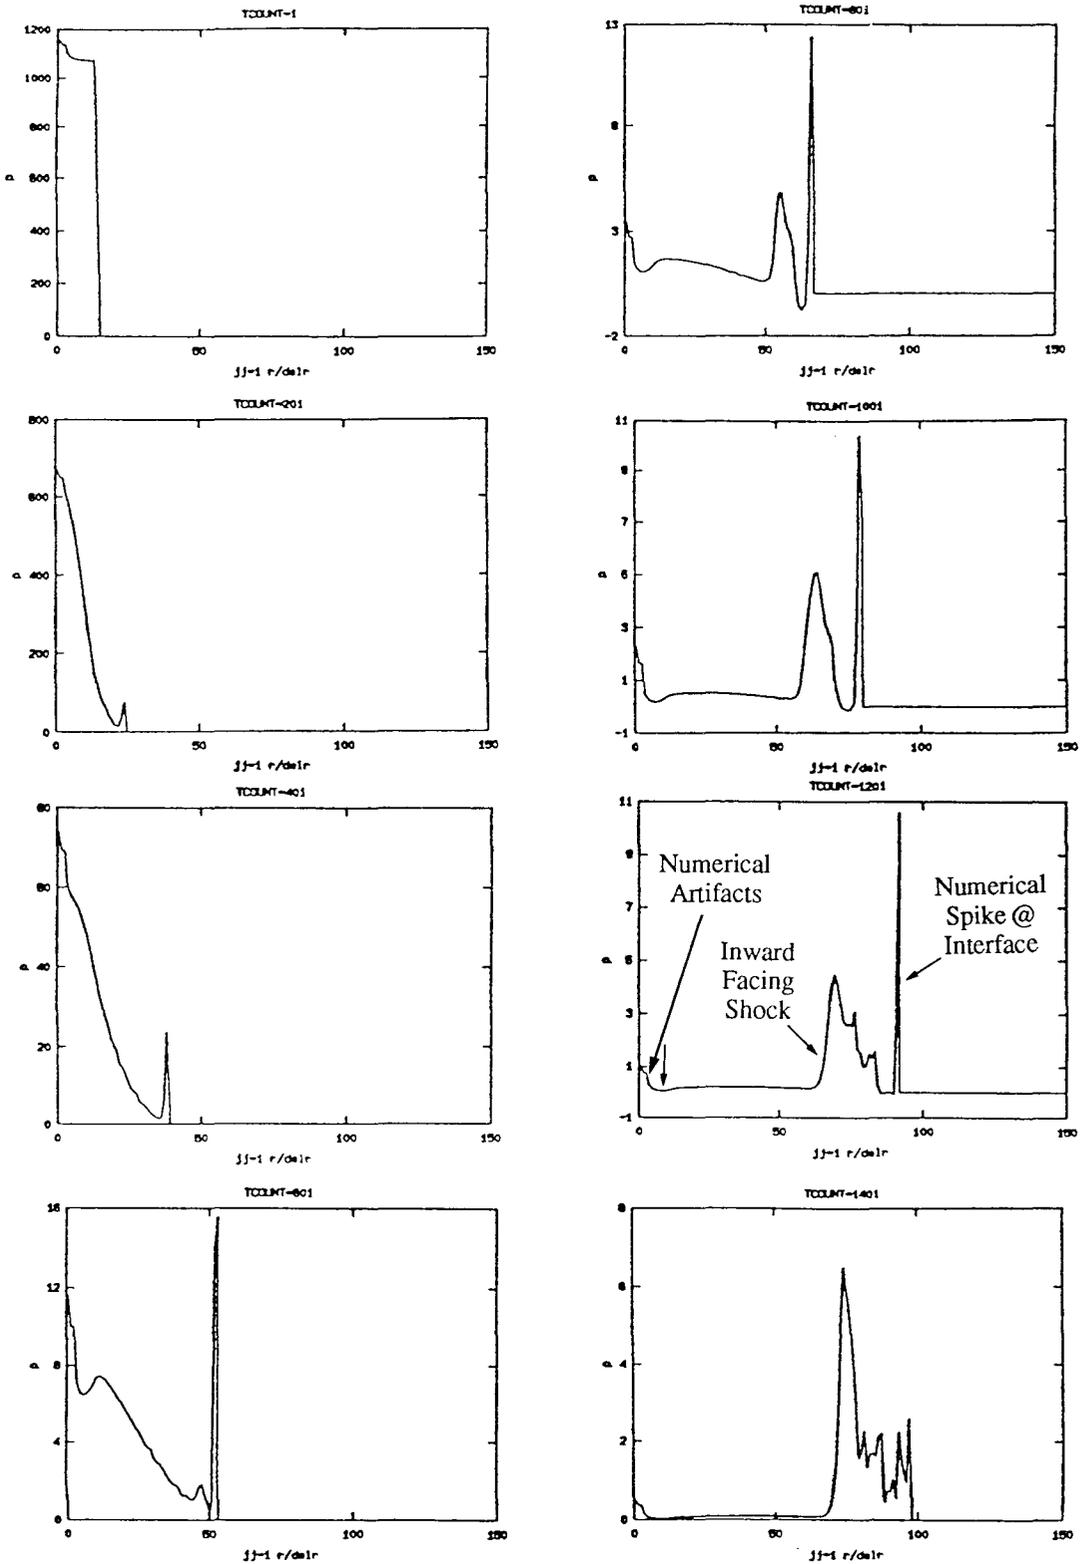


Figure 7.3a

Nondimensionalized pressure (p') (vertical axis)
 v.s. $r / \Delta r$ (horizontal axis) for $\theta = 0.5\Delta\theta$.
 (In this chapter, $\Delta r \approx 3.5$ (cm) and $\Delta\theta = \pi/3.5$ (radians).)

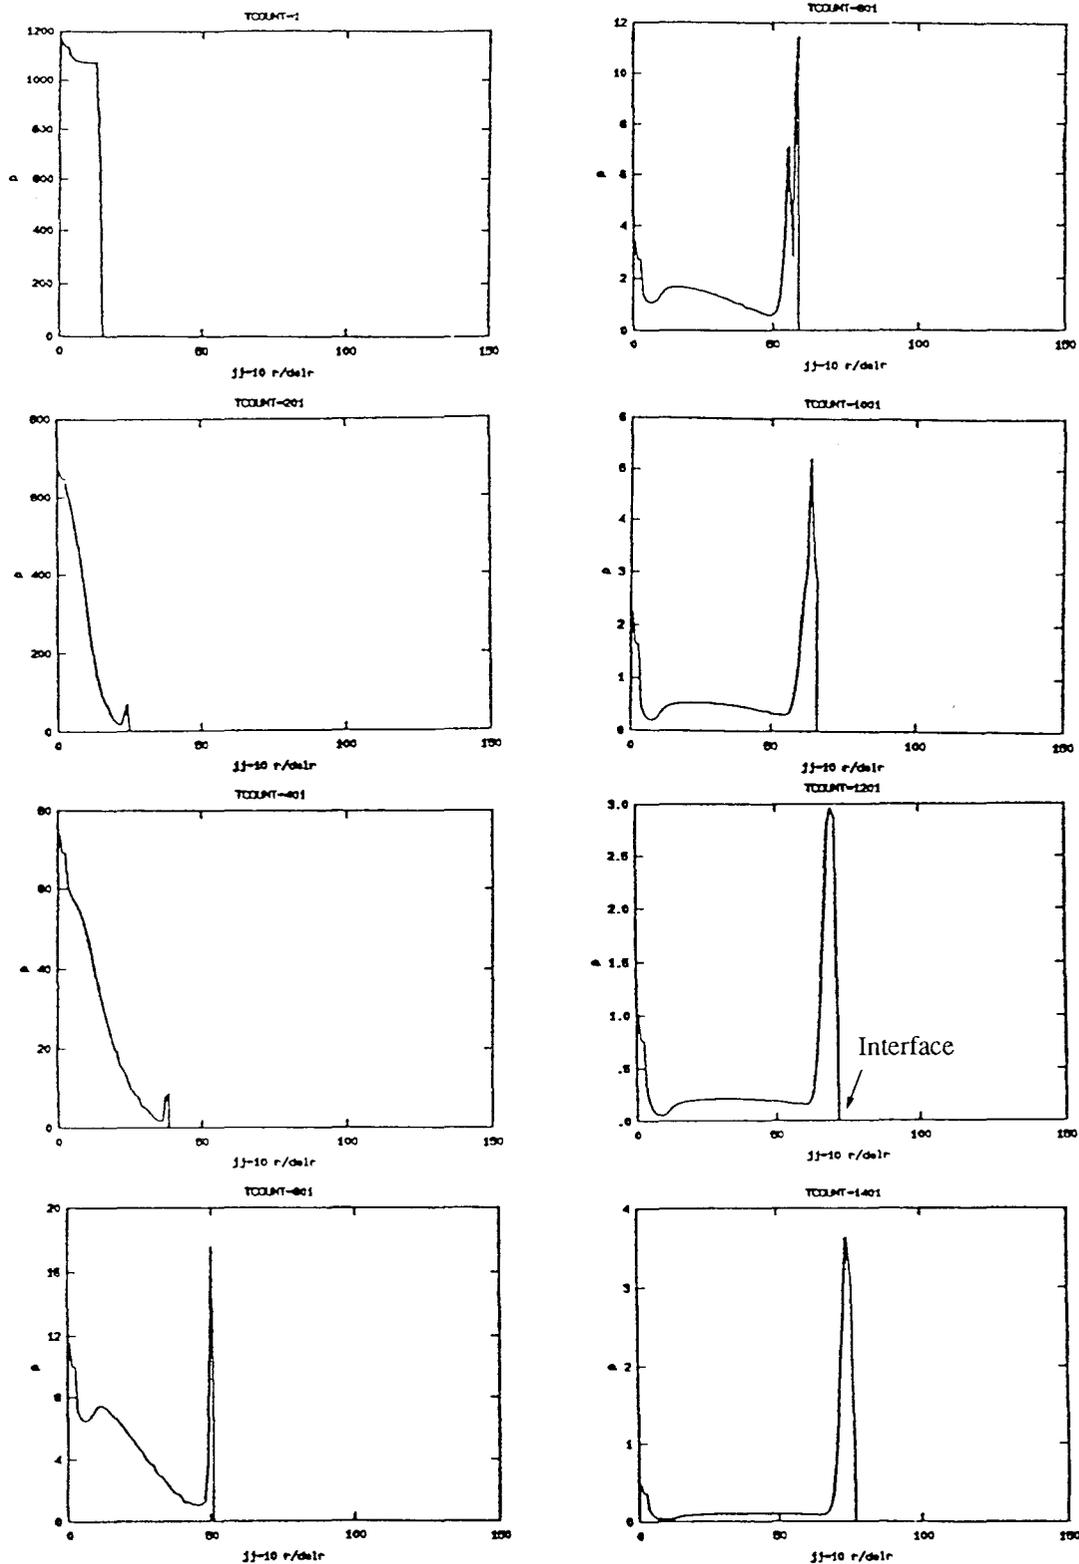


Figure 7.3b

Nondimensionalized pressure (p') (vertical axis)
 v.s. $r / \Delta r$ (horizontal axis) for $\theta = 9.5\Delta\theta$.

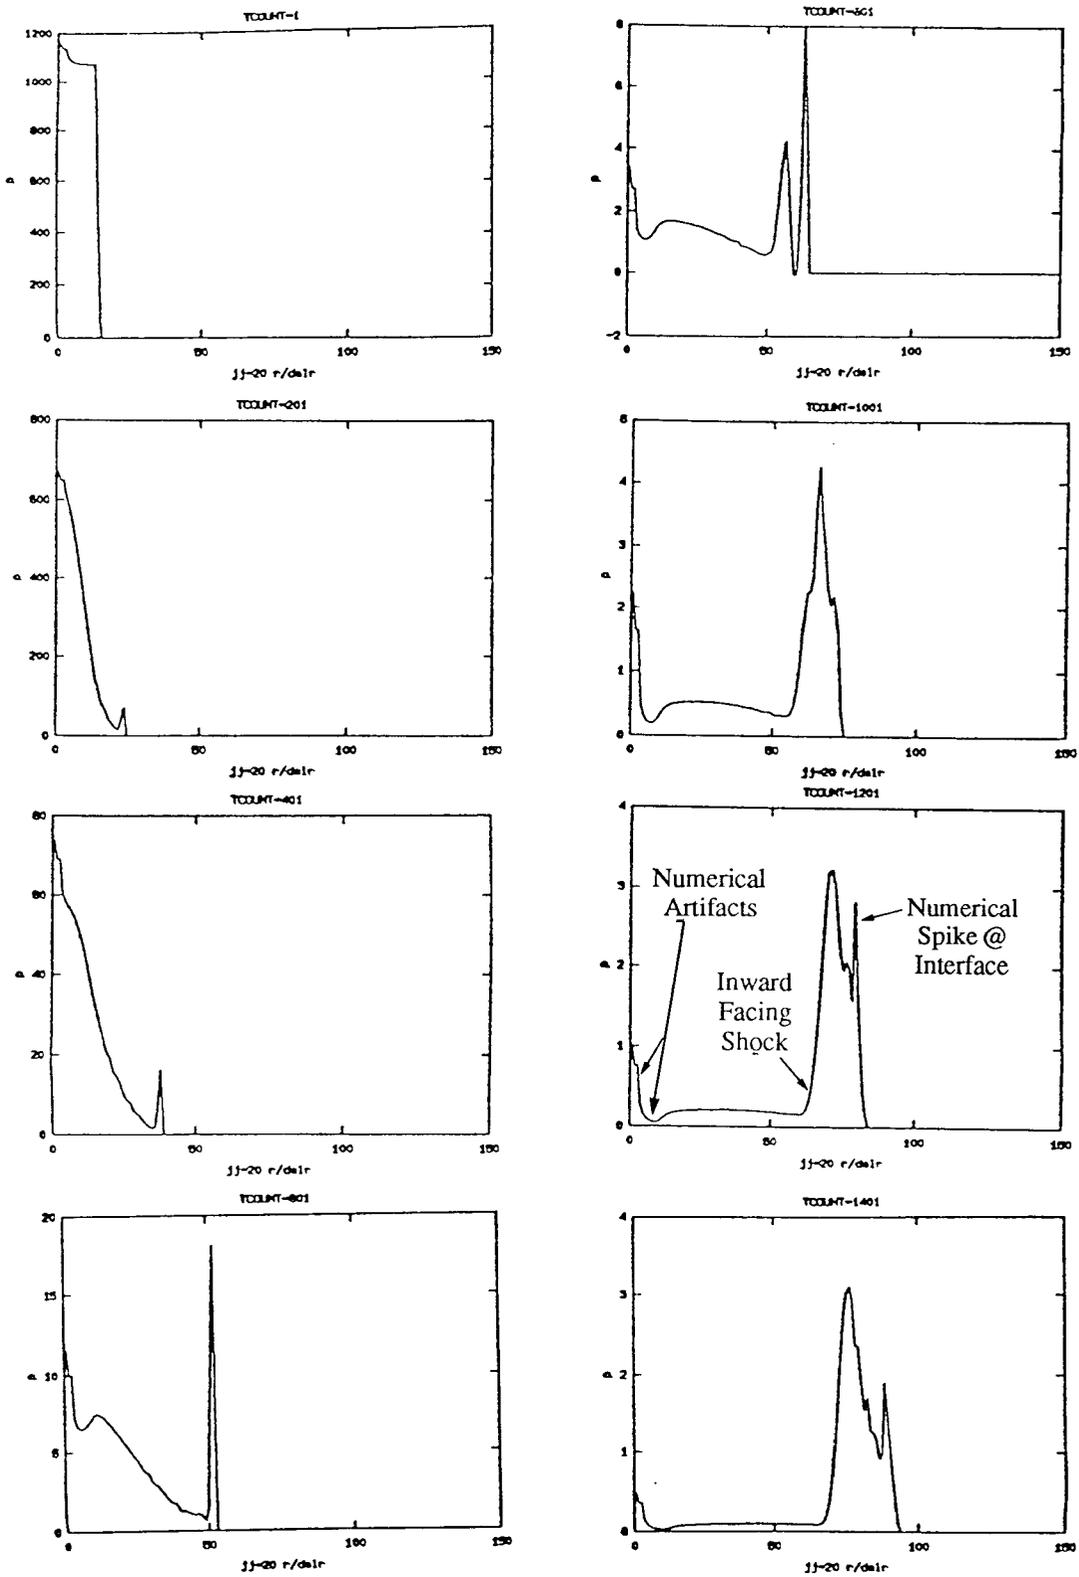


Figure 7.3c

Nondimensionalized pressure (p') (vertical axis)
 v.s. $r/\Delta r$ (horizontal axis) for $\theta = 19.5\Delta\theta$.

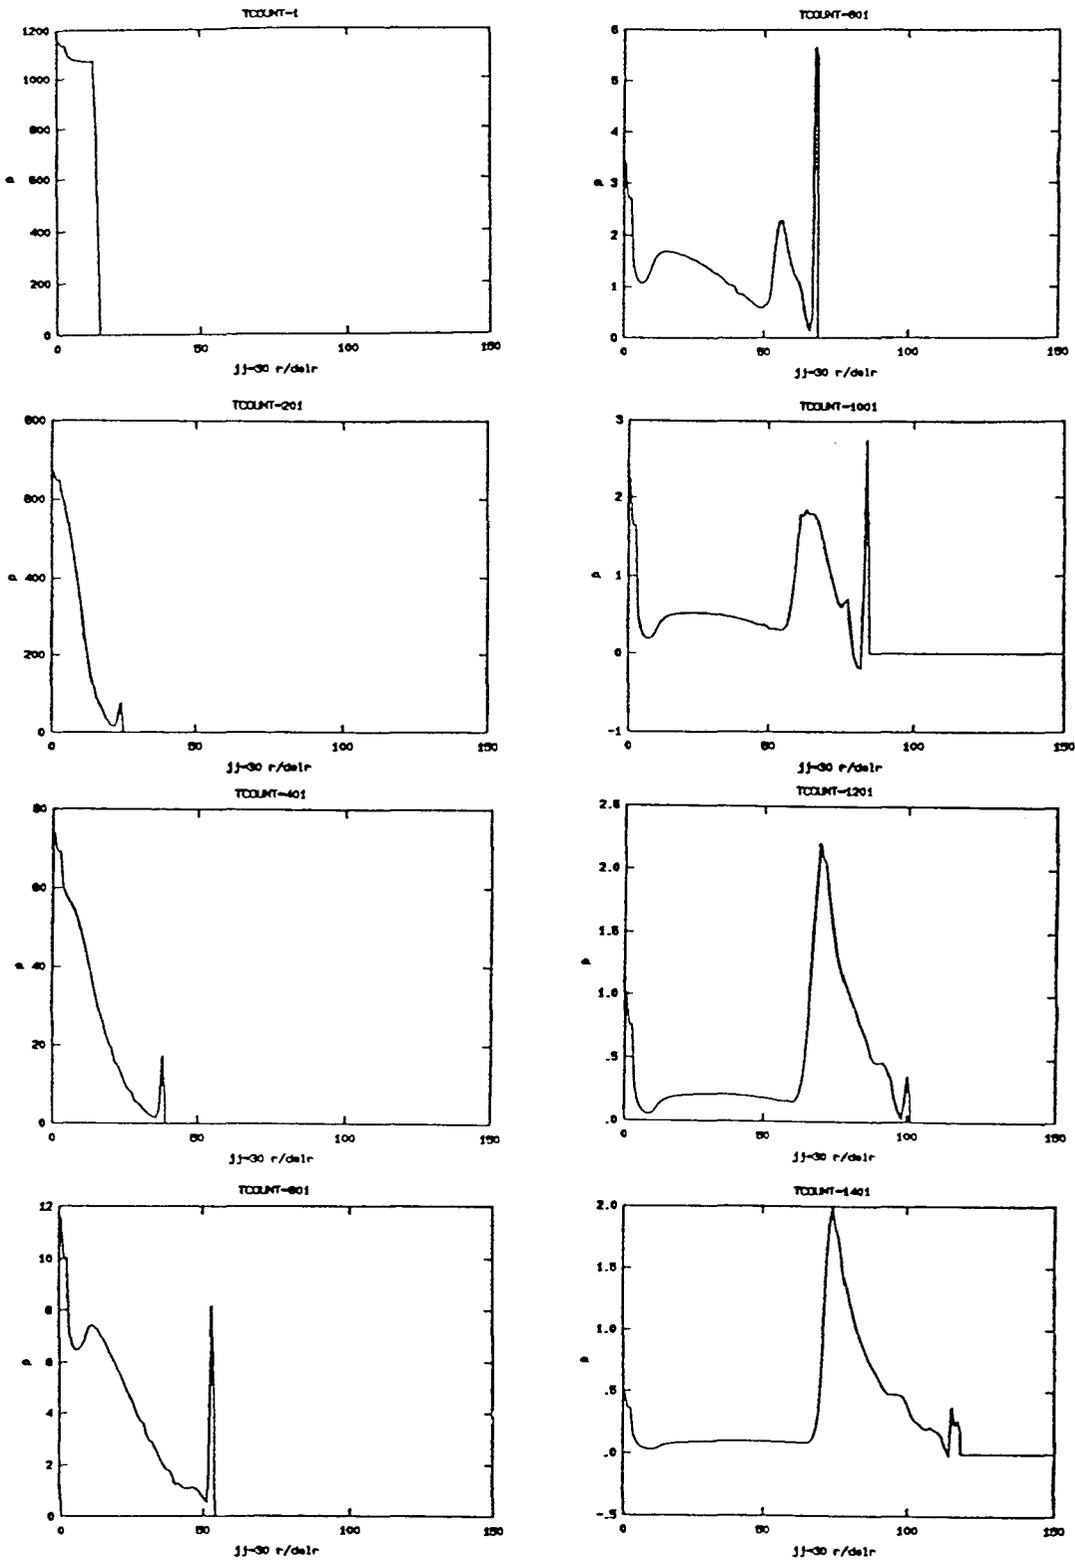


Figure 7.3d

Nondimensionalized pressure (p') (vertical axis)
 v.s. $r / \Delta r$ (horizontal axis) for $\theta = 29.5\Delta\theta$.

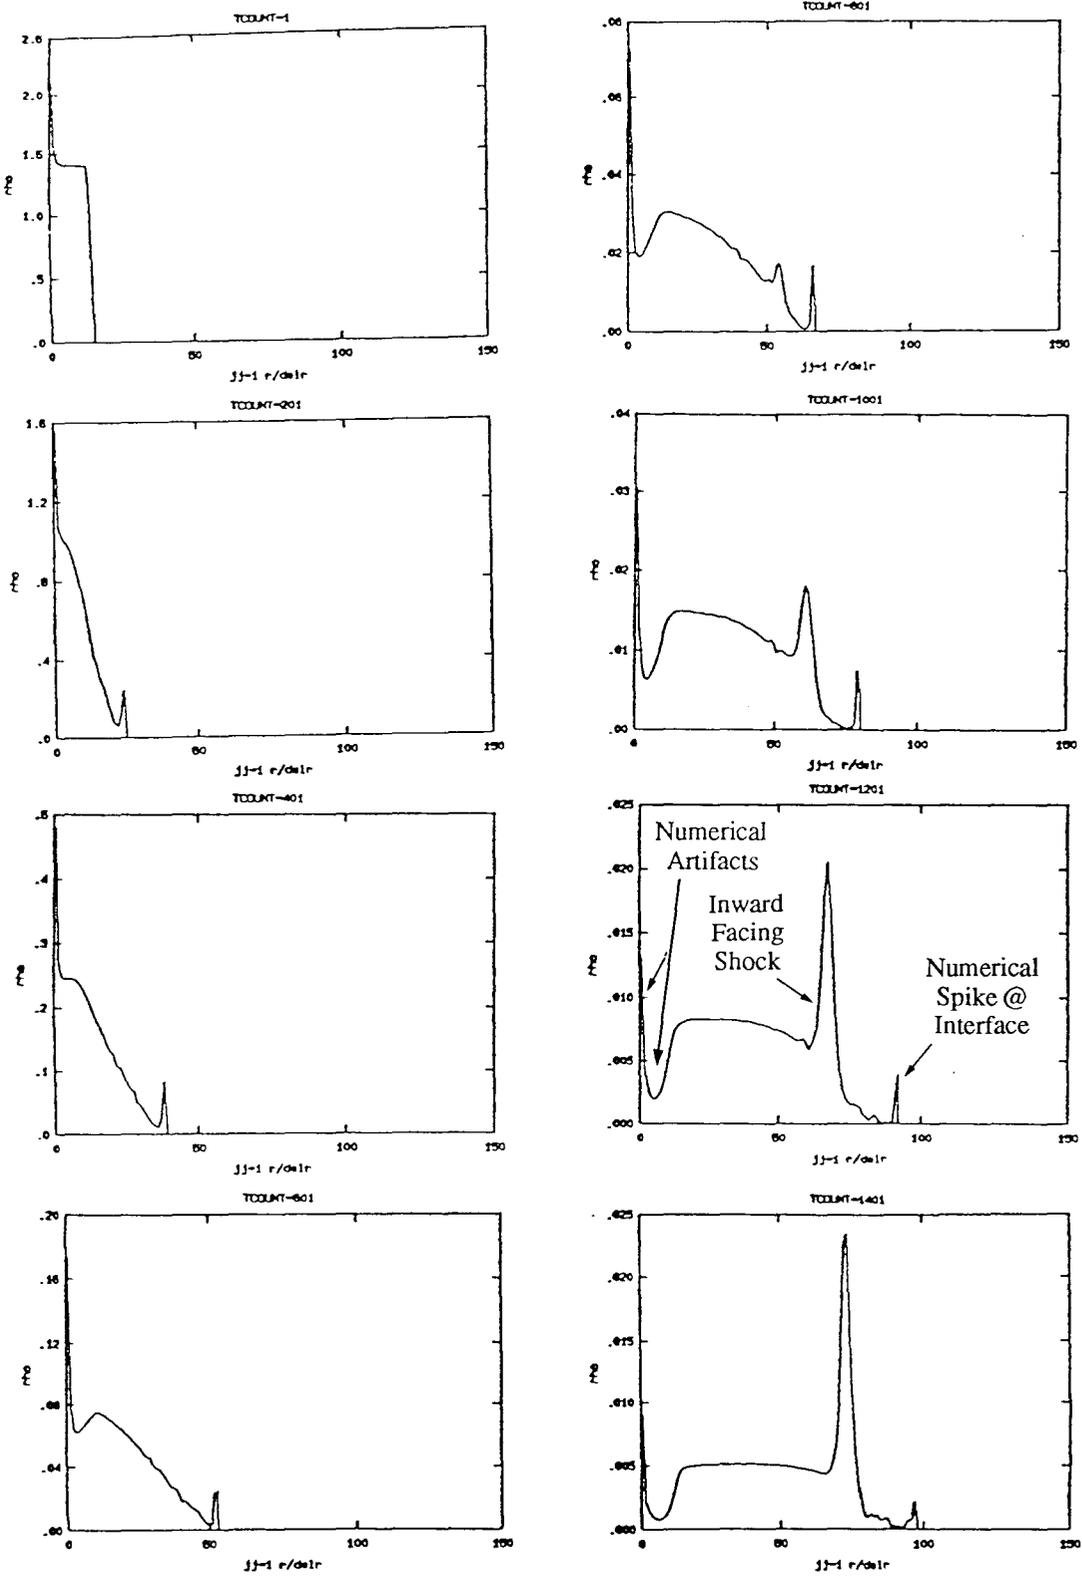


Figure 7.4a

Nondimensional density (ρ') (vertical axis)
 v.s. $r/\Delta r$ (horizontal axis) for $\theta = 0.5\Delta\theta$.

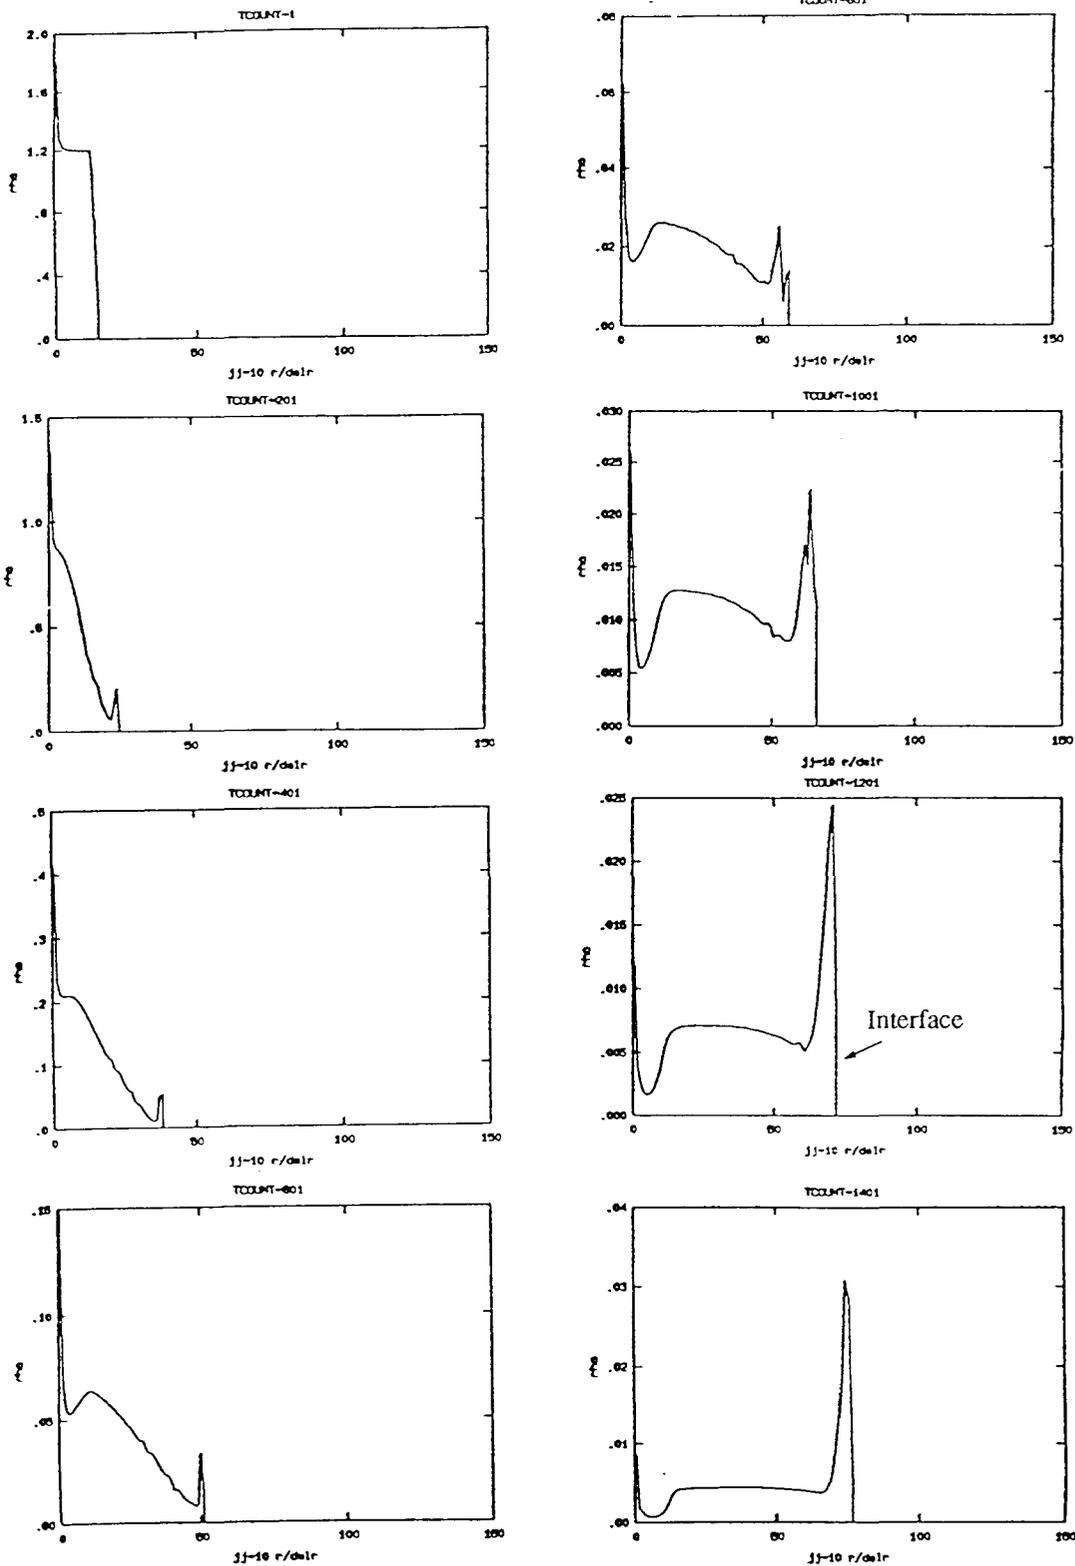


Figure 7.4b

Nondimensional density (ρ') (vertical axis)
 v.s. $r/\Delta r$ (horizontal axis) for $\theta = 9.5\Delta\theta$.

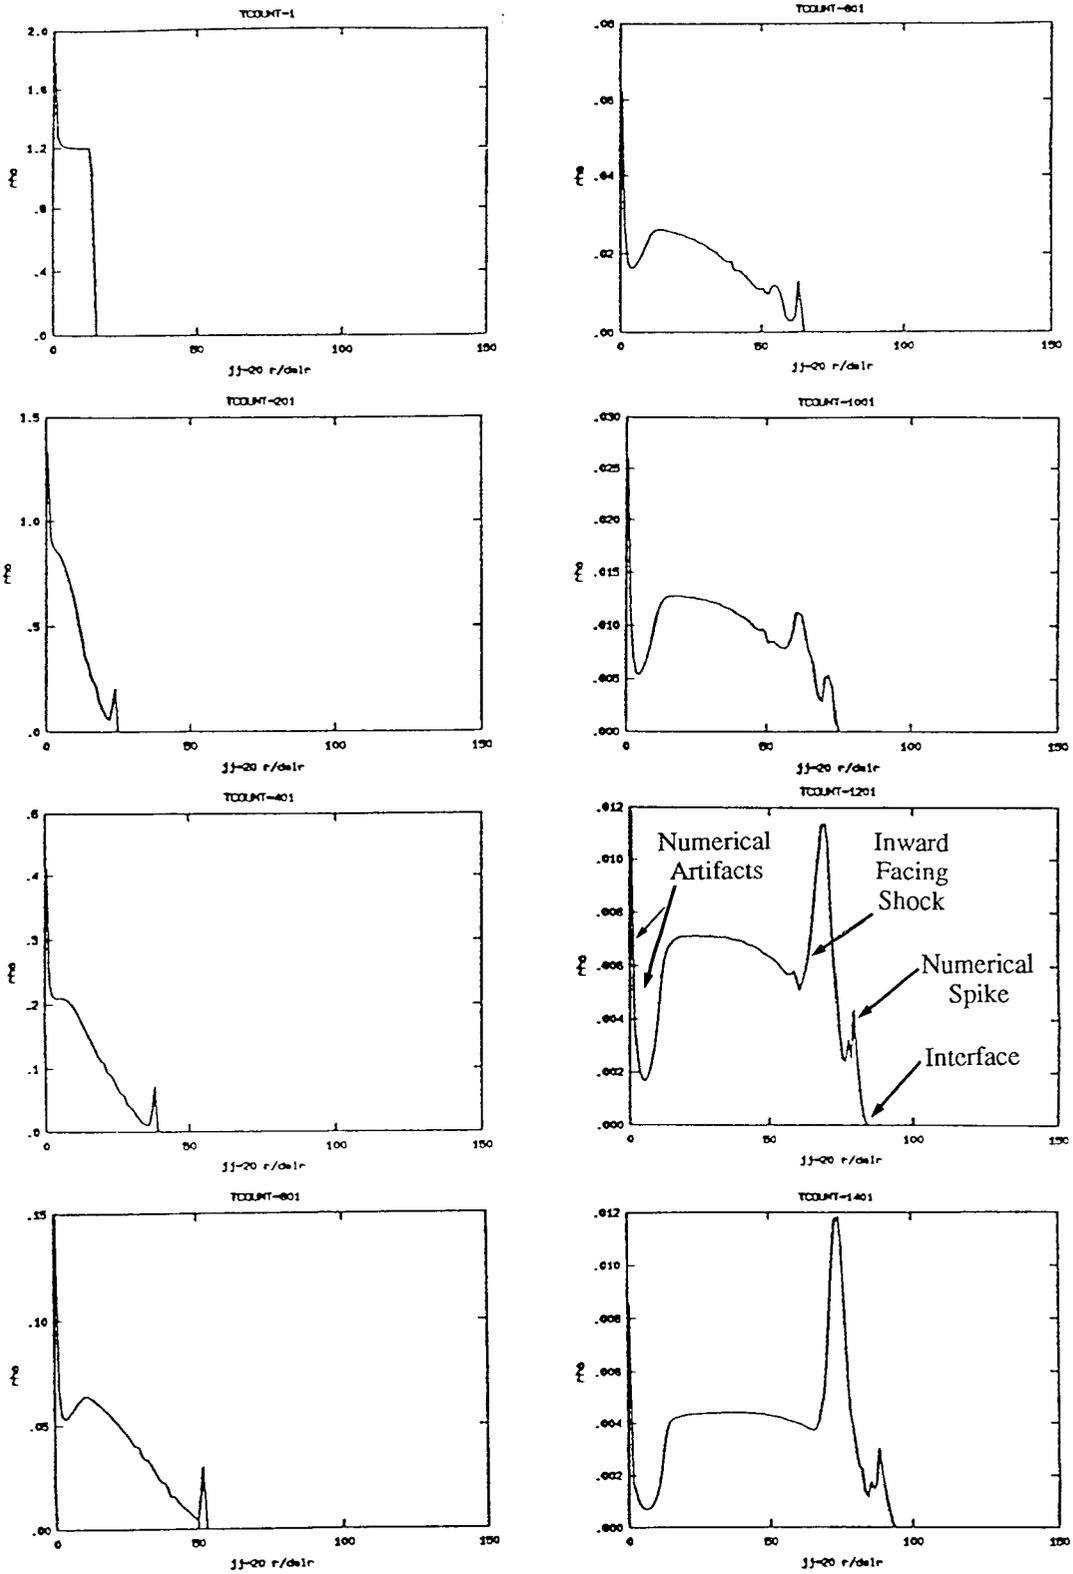


Figure 7.4c

Nondimensional density (ρ') (vertical axis)
 v.s. $r / \Delta r$ (horizontal axis) for $\theta = 19.5\Delta\theta$.

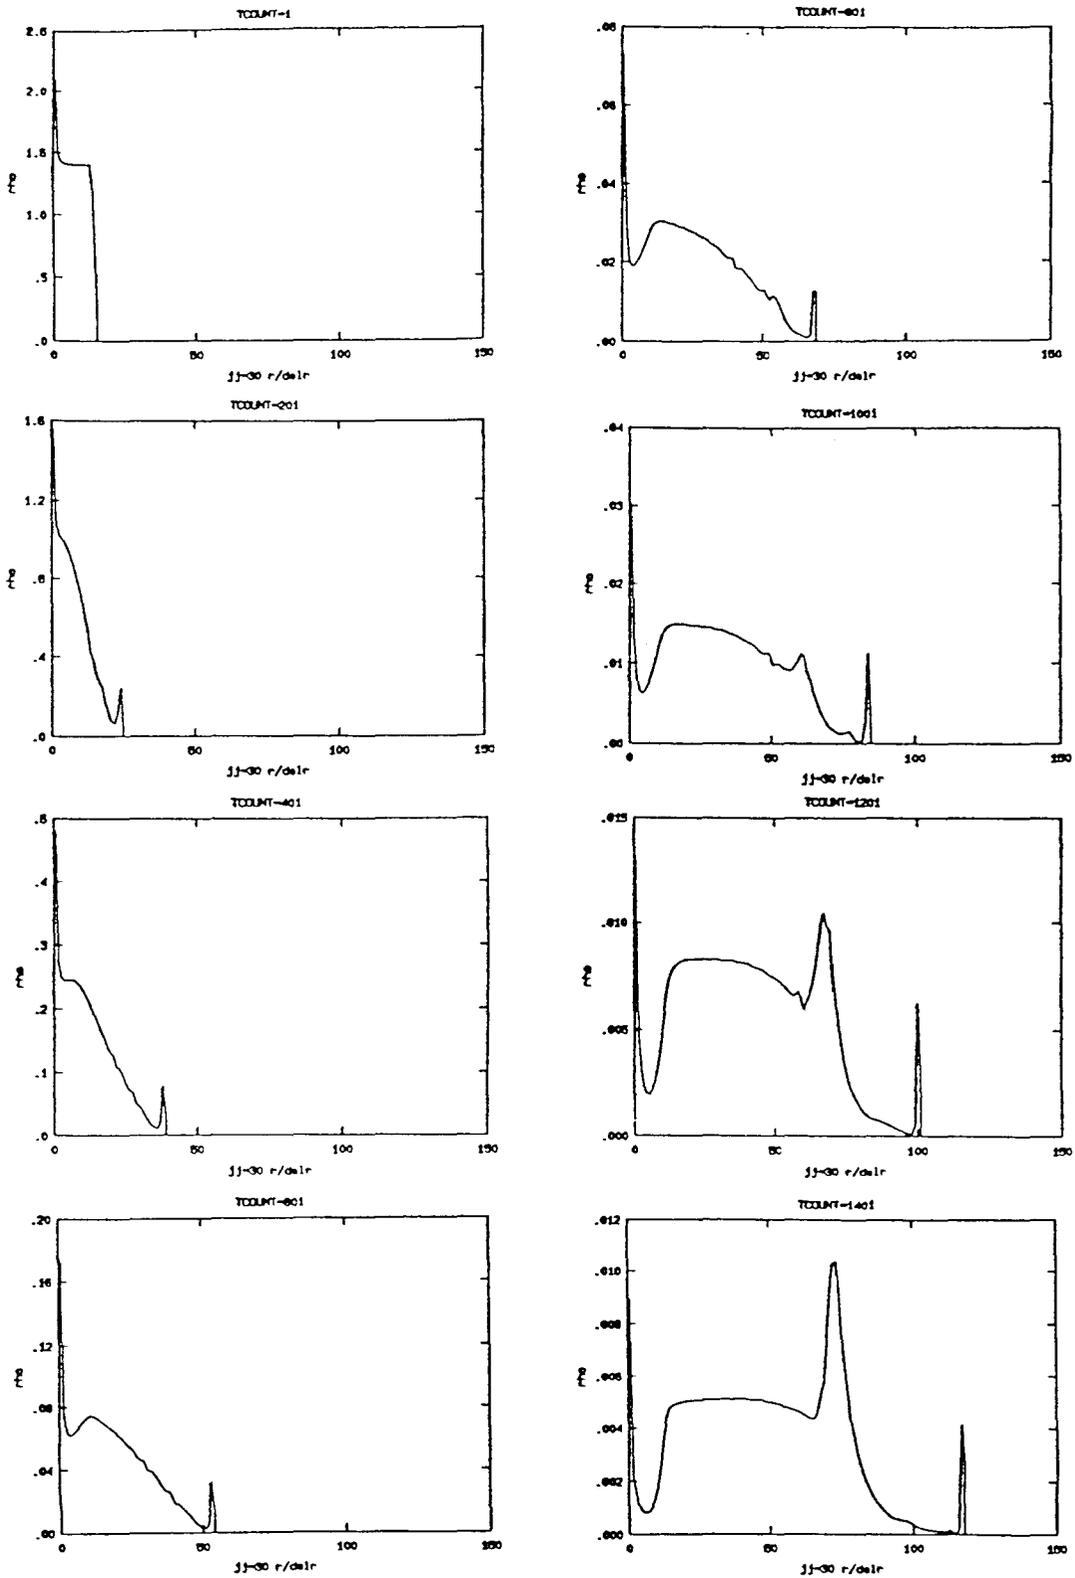


Figure 7.4d

Nondimensionalized density (ρ') (vertical axis)
 v.s. $r/\Delta r$ (horizontal axis) for $\theta = 29.5\Delta\theta$.

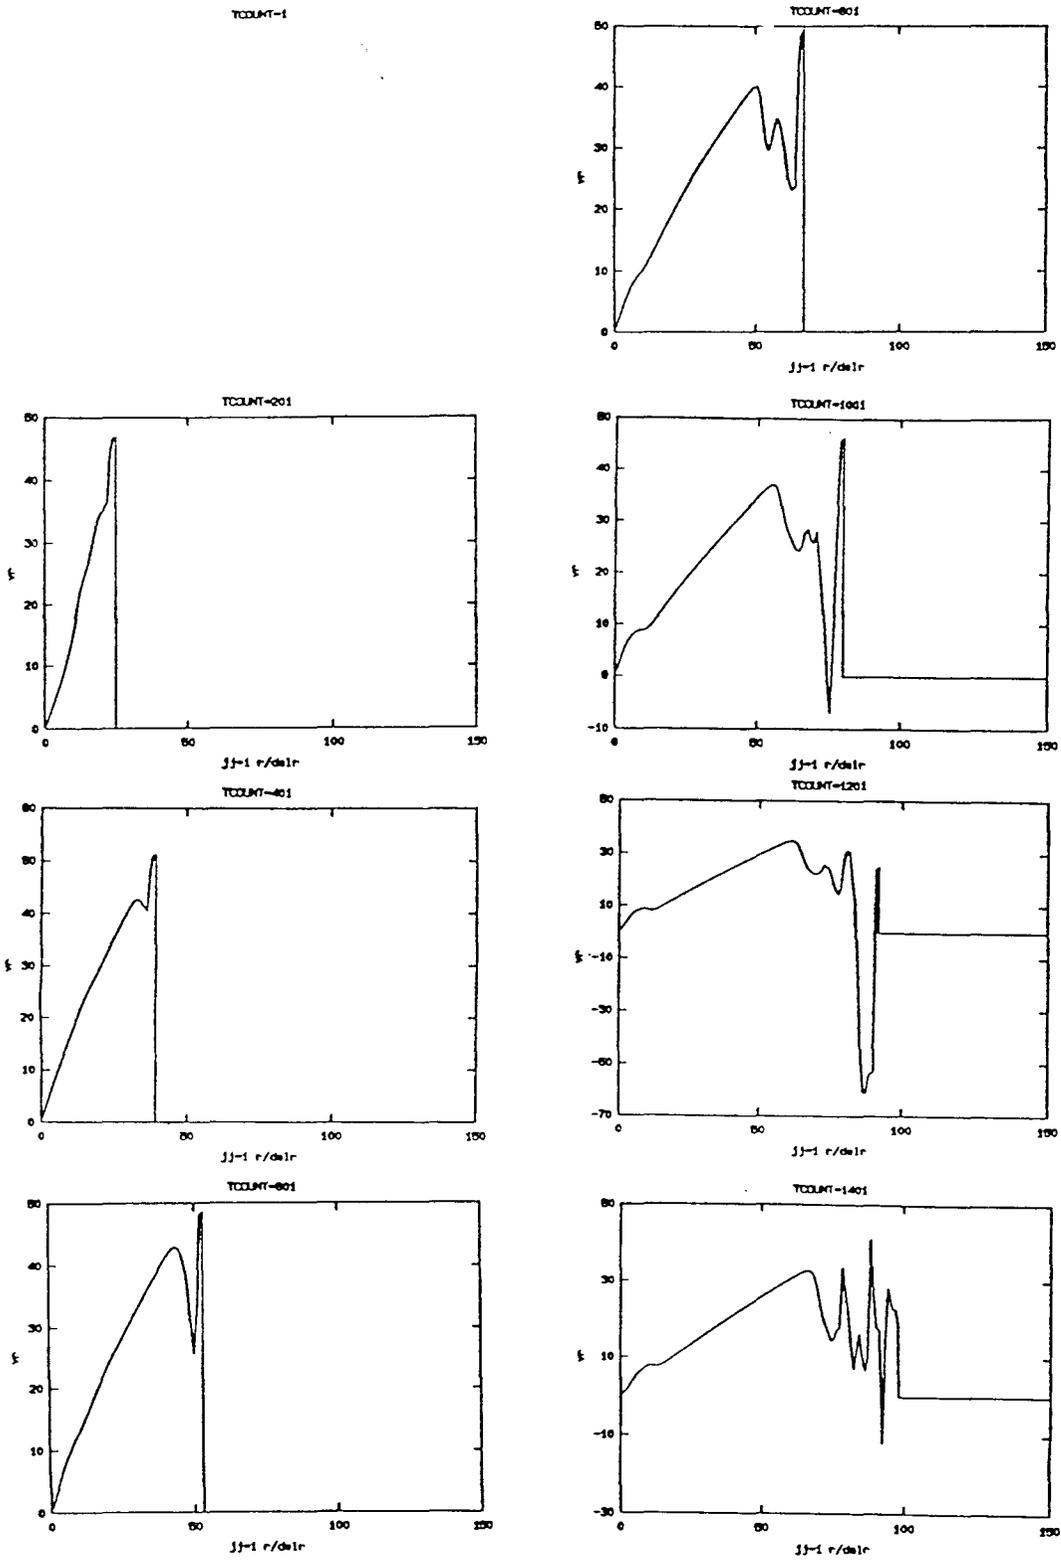


Figure 7.5a Nondimensionalized radial velocity (v_r') (vertical axis) v.s. $r / \Delta r$ (horizontal axis) for $\theta = 0.5\Delta\theta$.

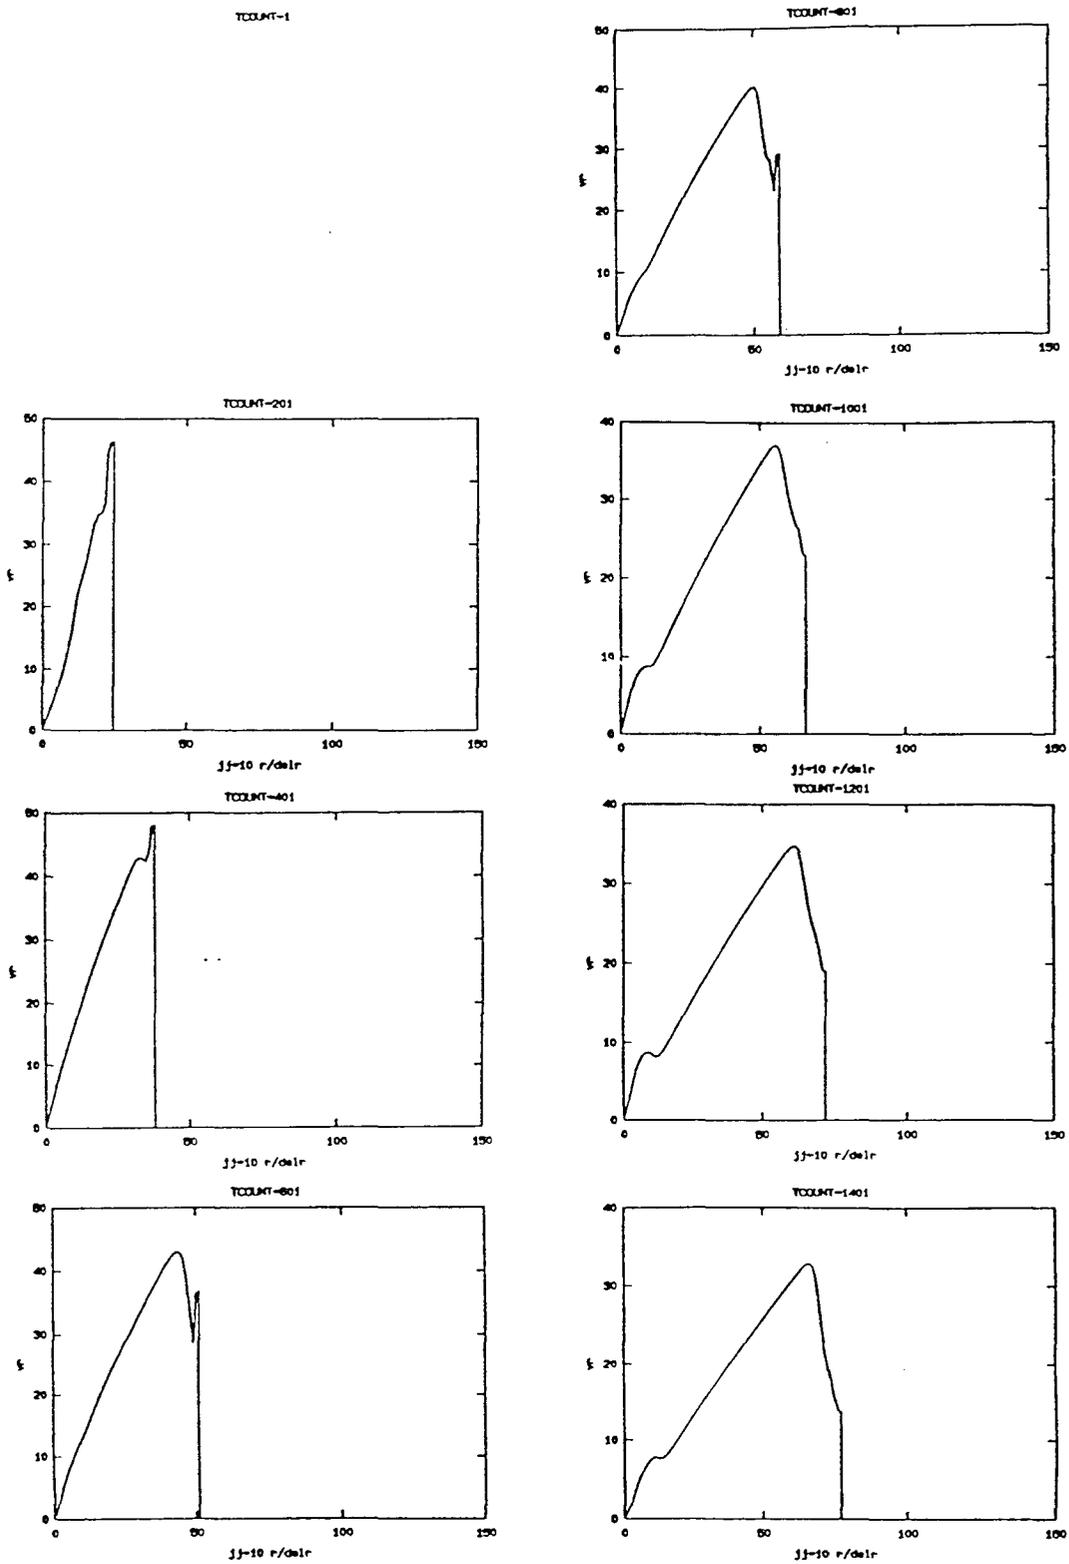


Figure 7.5b Nondimensionalized radial velocity (v_r') (vertical axis) v.s. $r/\Delta r$ (horizontal axis) for $\theta = 9.5\Delta\theta$.

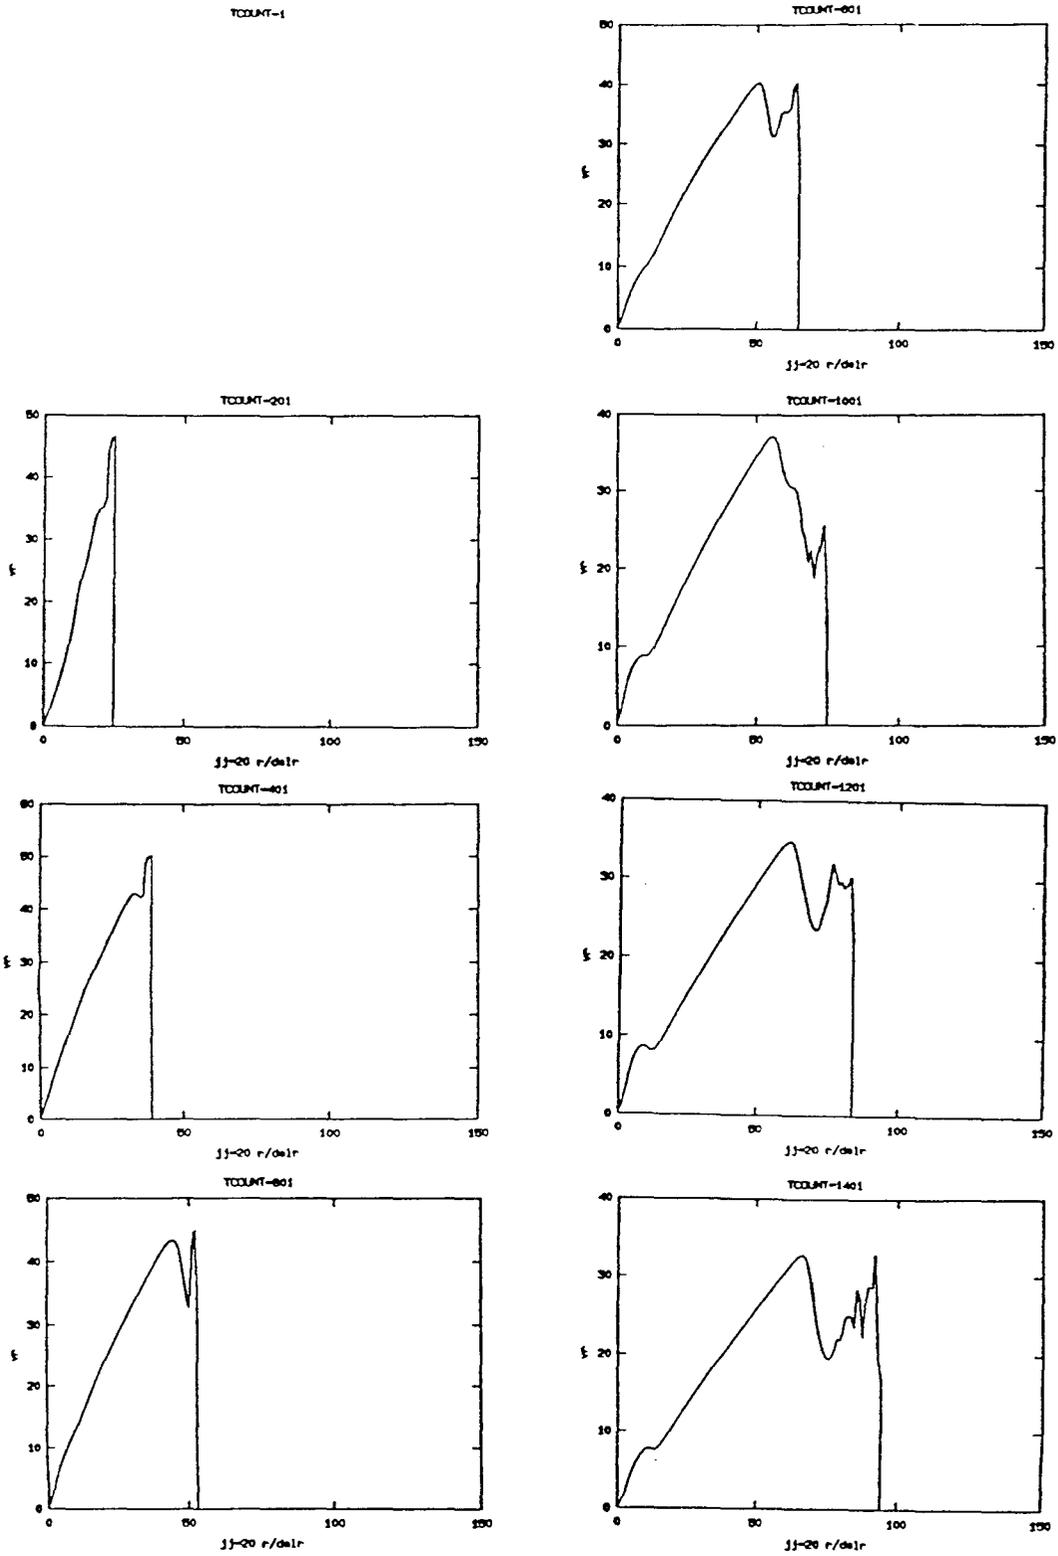


Figure 7.5c Nondimensionalized radial velocity (v_r') (vertical axis) v.s. $r/\Delta r$ (horizontal axis) for $\theta = 19.5\Delta\theta$.

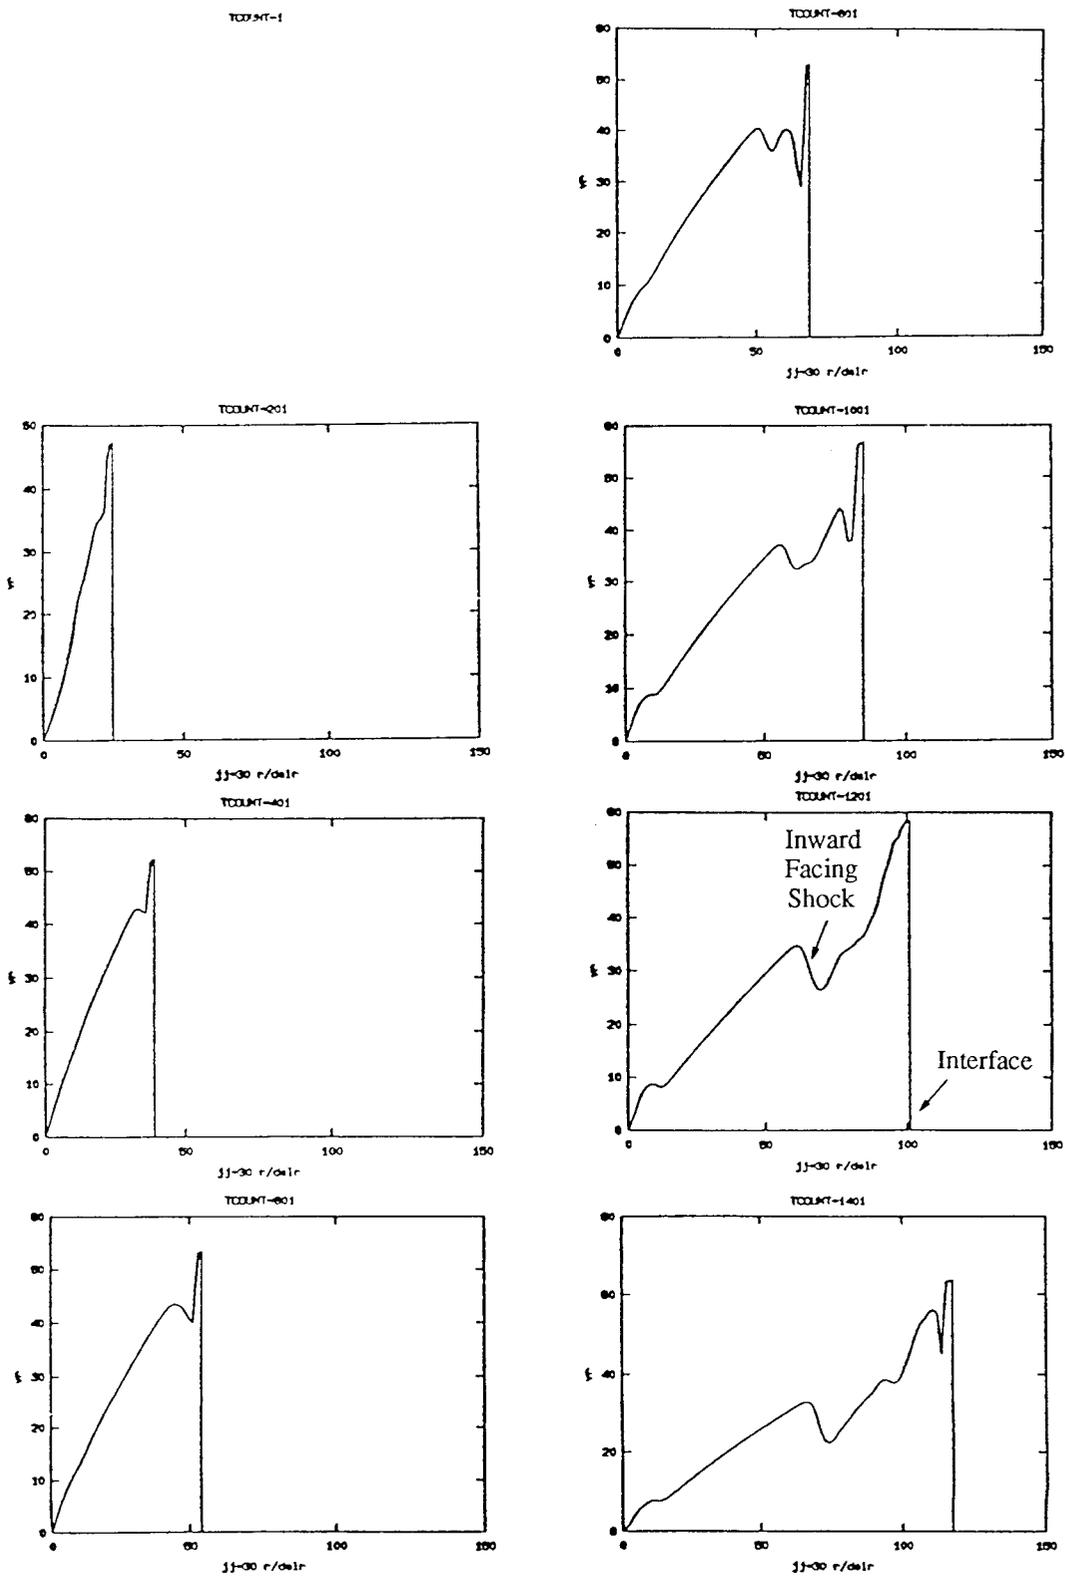


Figure 7.5d Nondimensionalized radial velocity (v_r') (vertical axis) v.s. $r/\Delta r$ (horizontal axis) for $\theta = 29.5\Delta\theta$.

TDCOUNT=1

Features of v_θ are often numerical, especially if $|v_\theta| \ll |v_r|$ (see text).

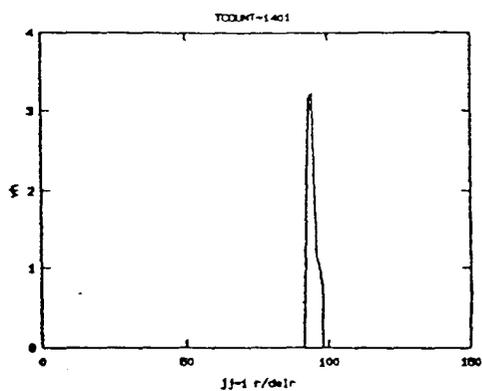
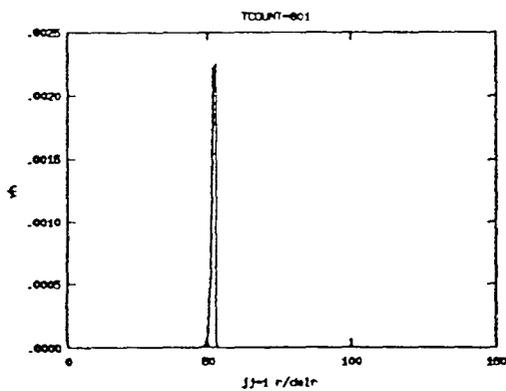
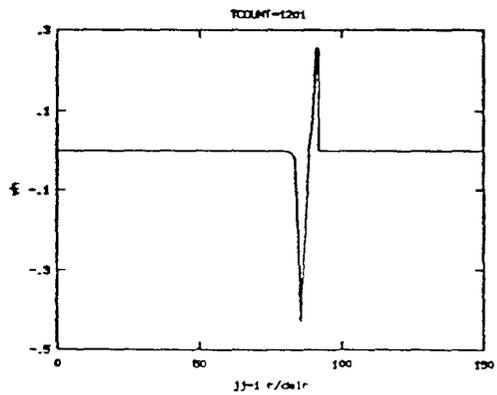
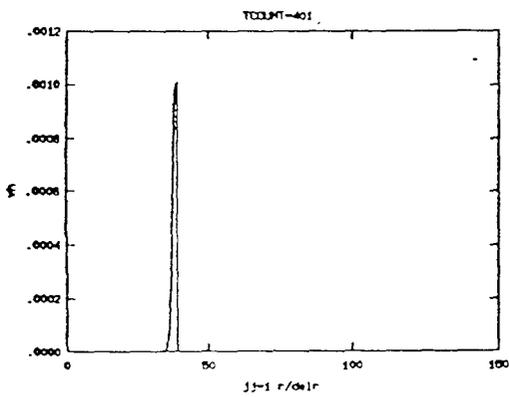
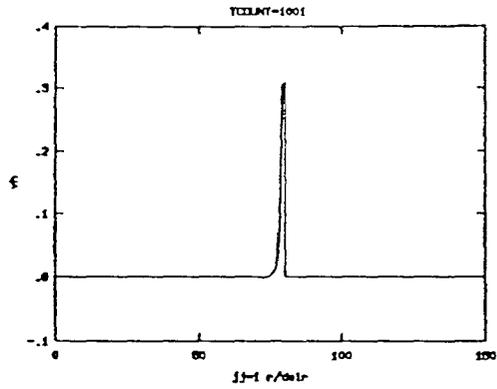
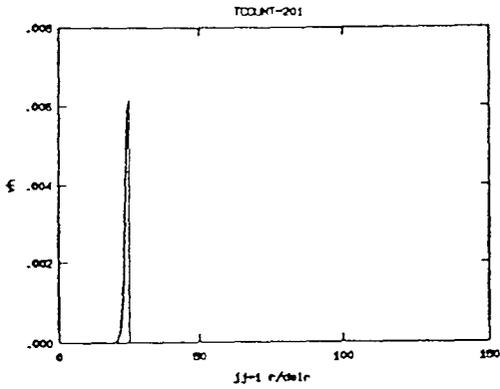
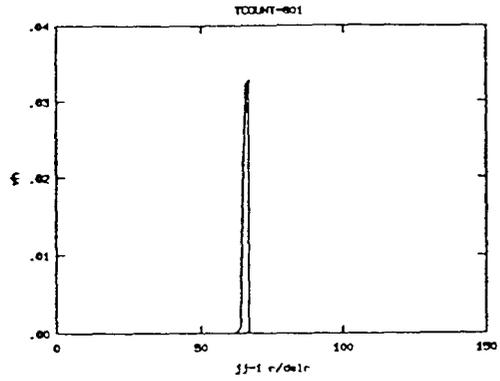


Figure 7.6a Nondimensionalized polar velocity (v_θ') (vertical axis) v.s. $r / \Delta r$ (horizontal axis) for $\theta = 0.5\Delta\theta$.

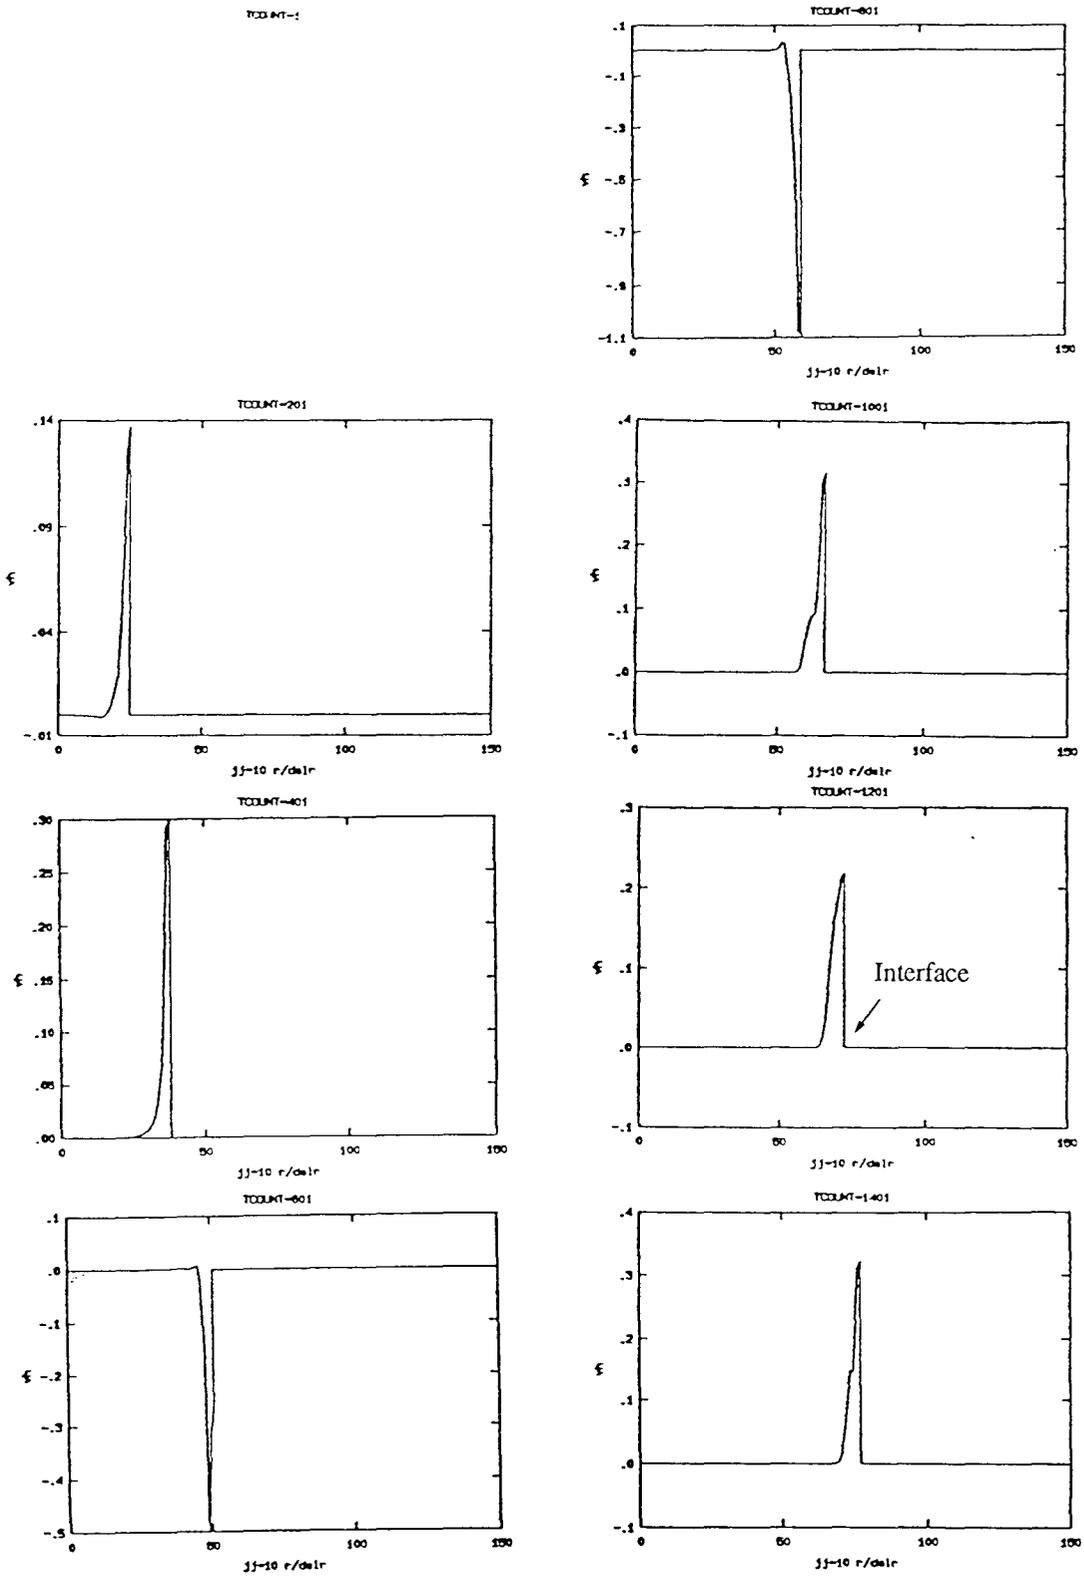


Figure 7.6b Nondimensionalized polar velocity (v_θ') (vertical axis) v.s. $r/\Delta r$ (horizontal axis) for $\theta = 9.5\Delta\theta$.

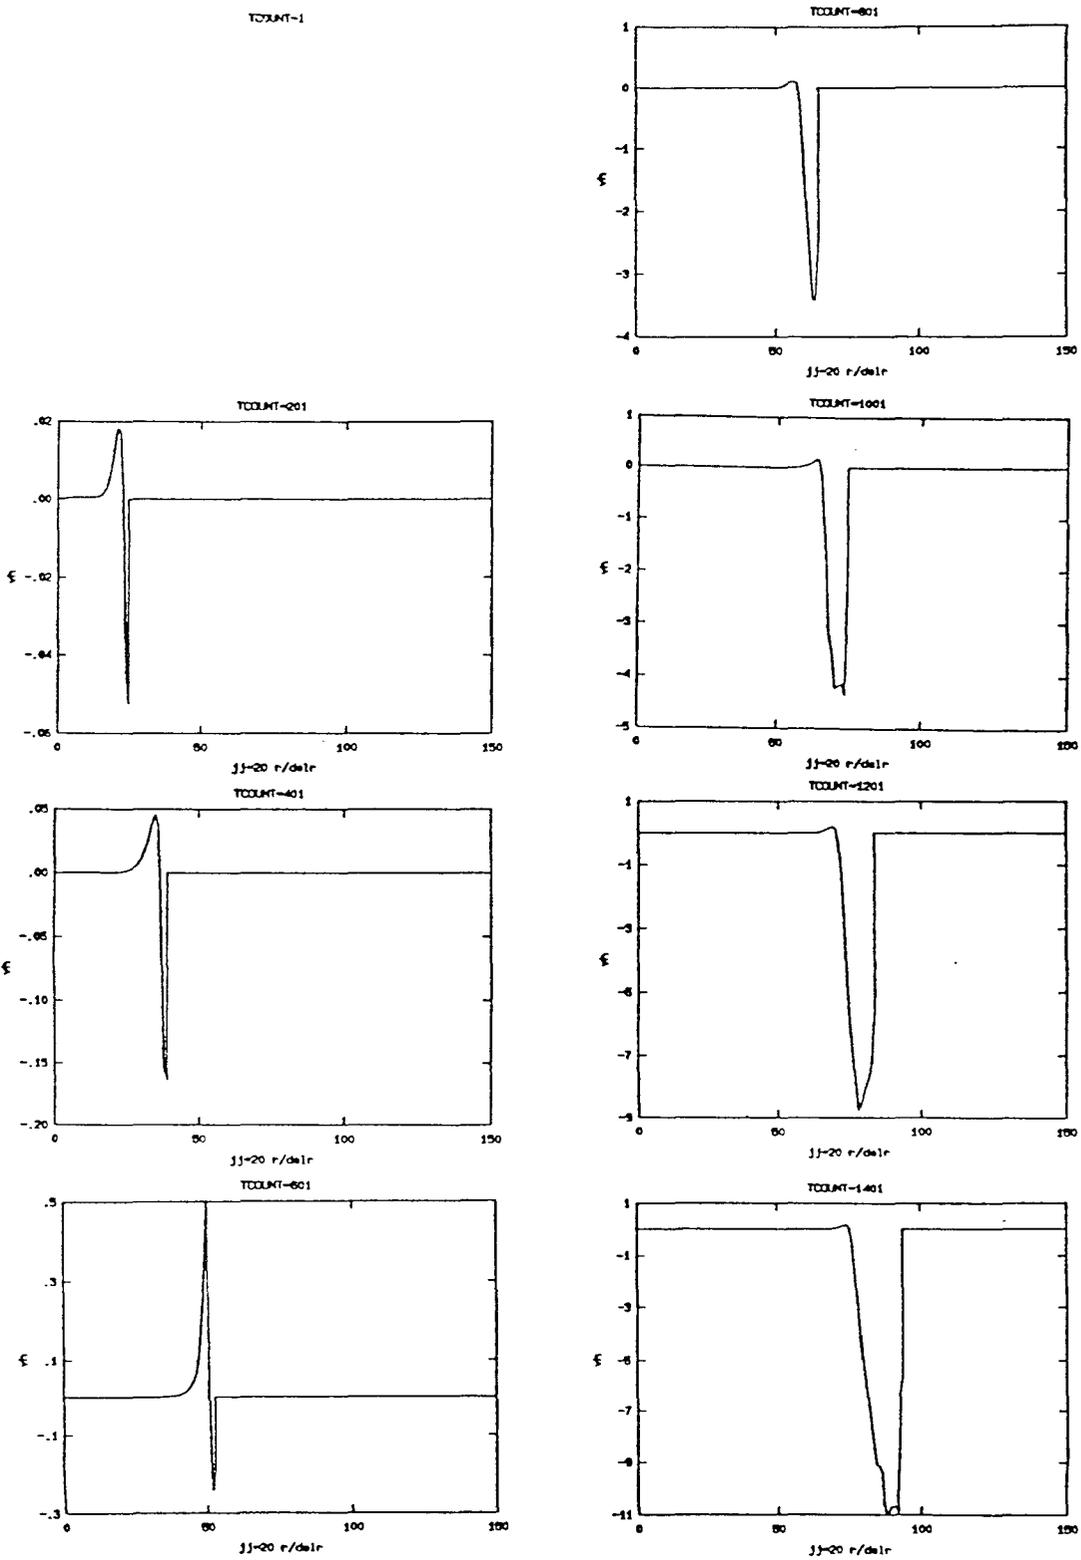


Figure 7.6c Nondimensionalized polar velocity (v_θ') (vertical axis) v.s. $r/\Delta r$ (horizontal axis) for $\theta = 19.5\Delta\theta$.

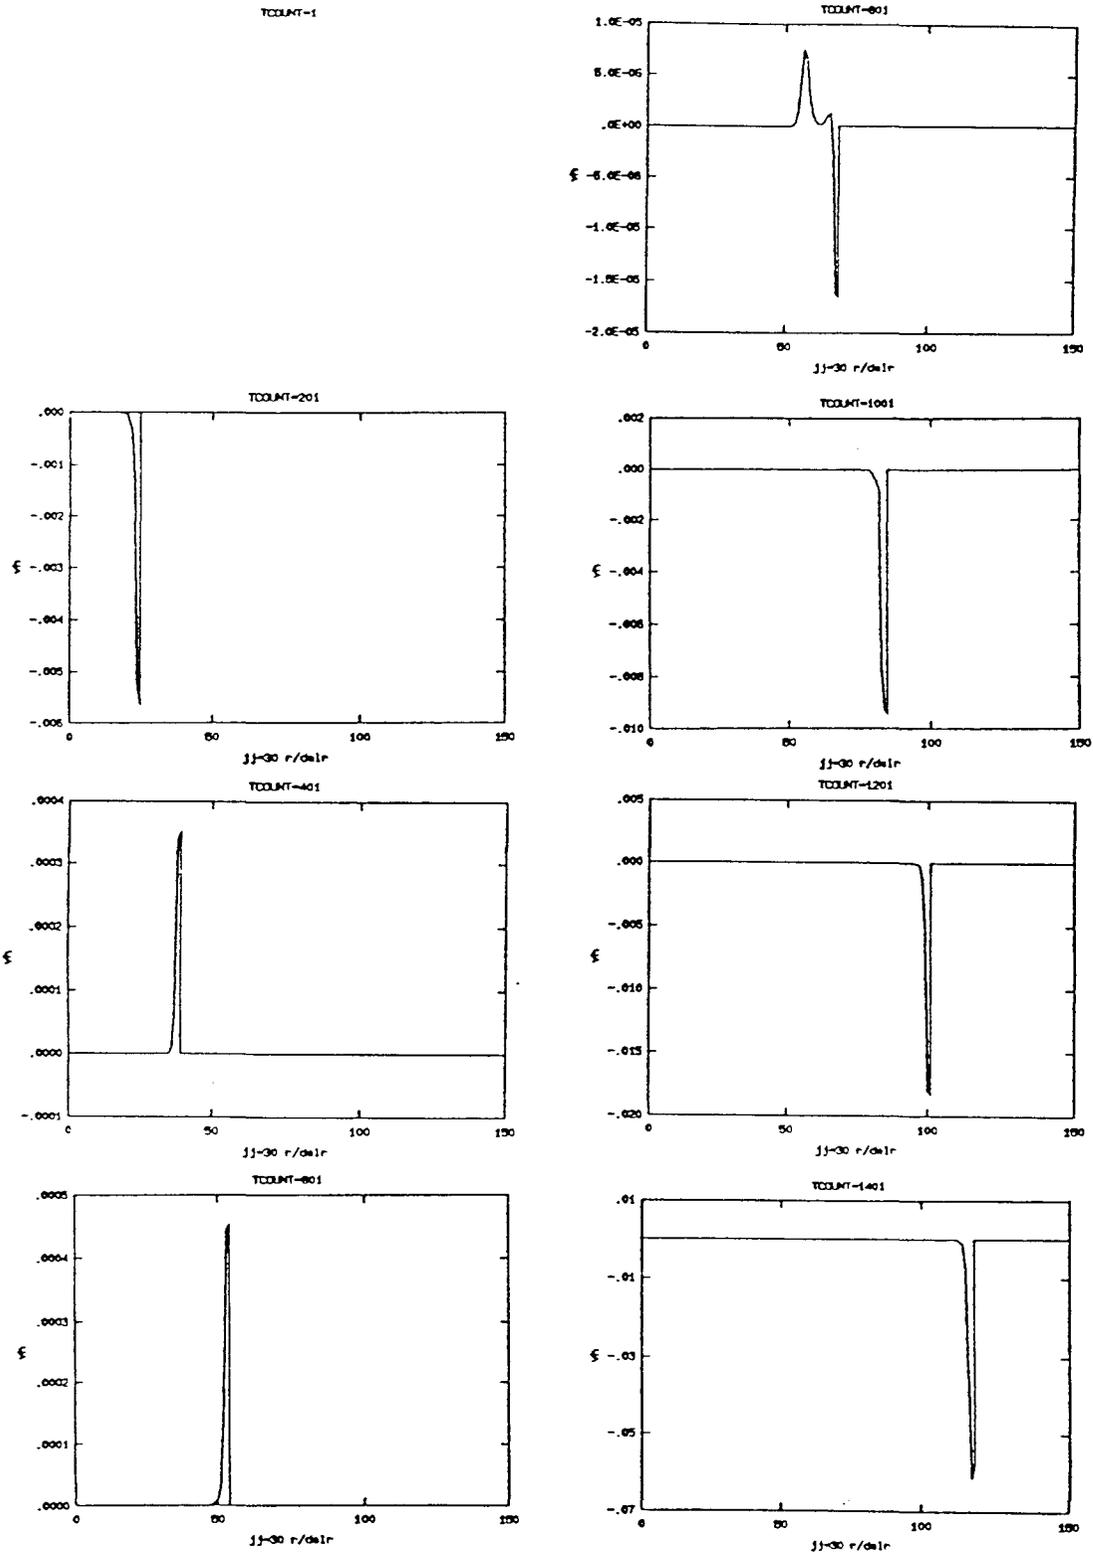


Figure 7.6d Nondimensionalized polar velocity (v_{θ}') (vertical axis)
 v.s. $r/\Delta r$ (horizontal axis) for $\theta = 29.5\Delta\theta$.

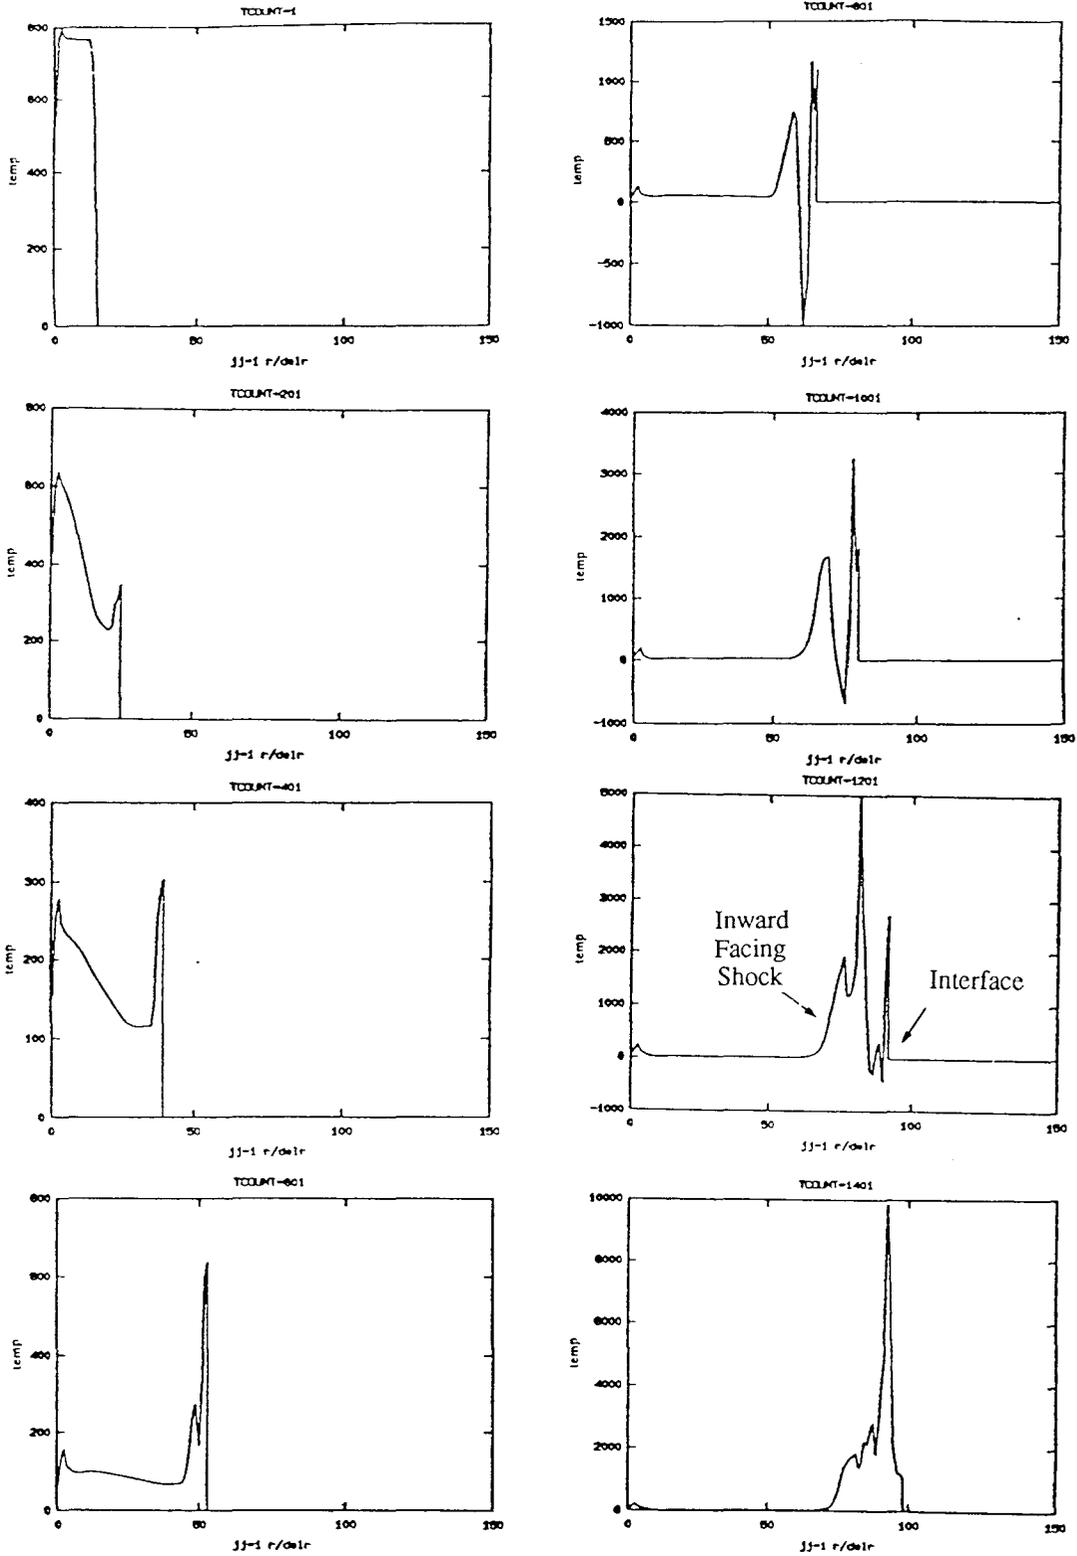


Figure 7.7a

Nondimensionalized temperature (T') (vertical axis)
 v.s. $r/\Delta r$ (horizontal axis) for $\theta = 0.5\Delta\theta$.

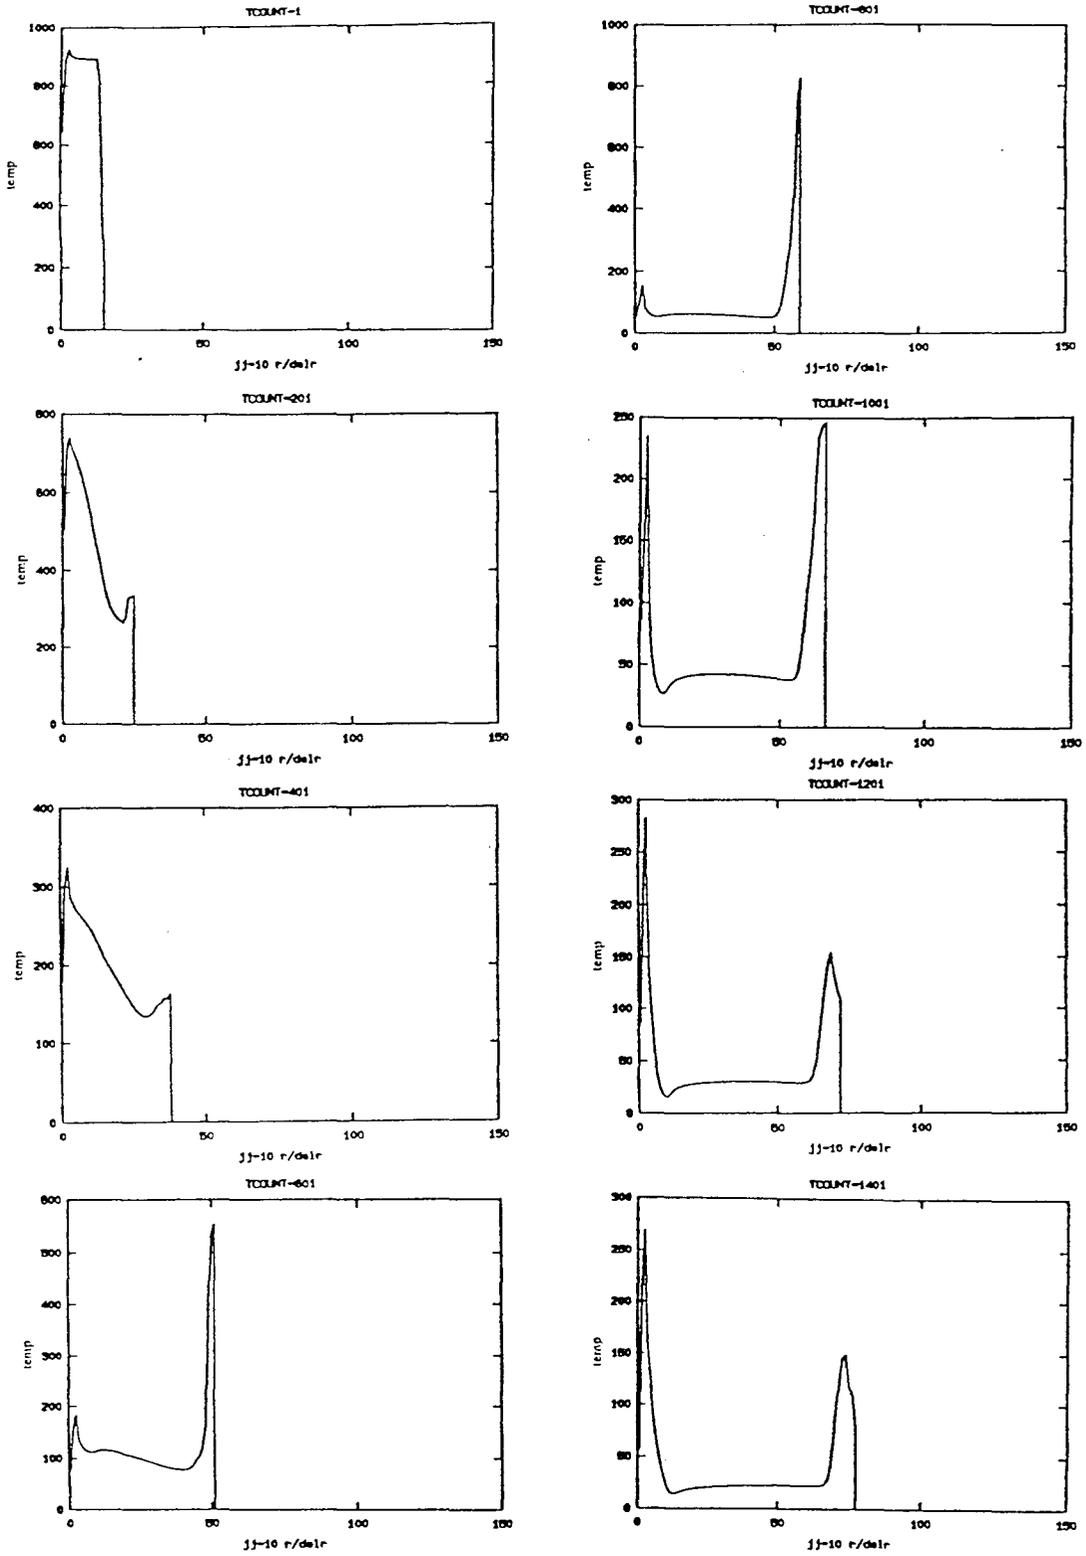


Figure 7.7b Nondimensionalized temperature (T) (vertical axis) v.s. $r / \Delta r$ (horizontal axis) for $\theta = 9.5\Delta\theta$.

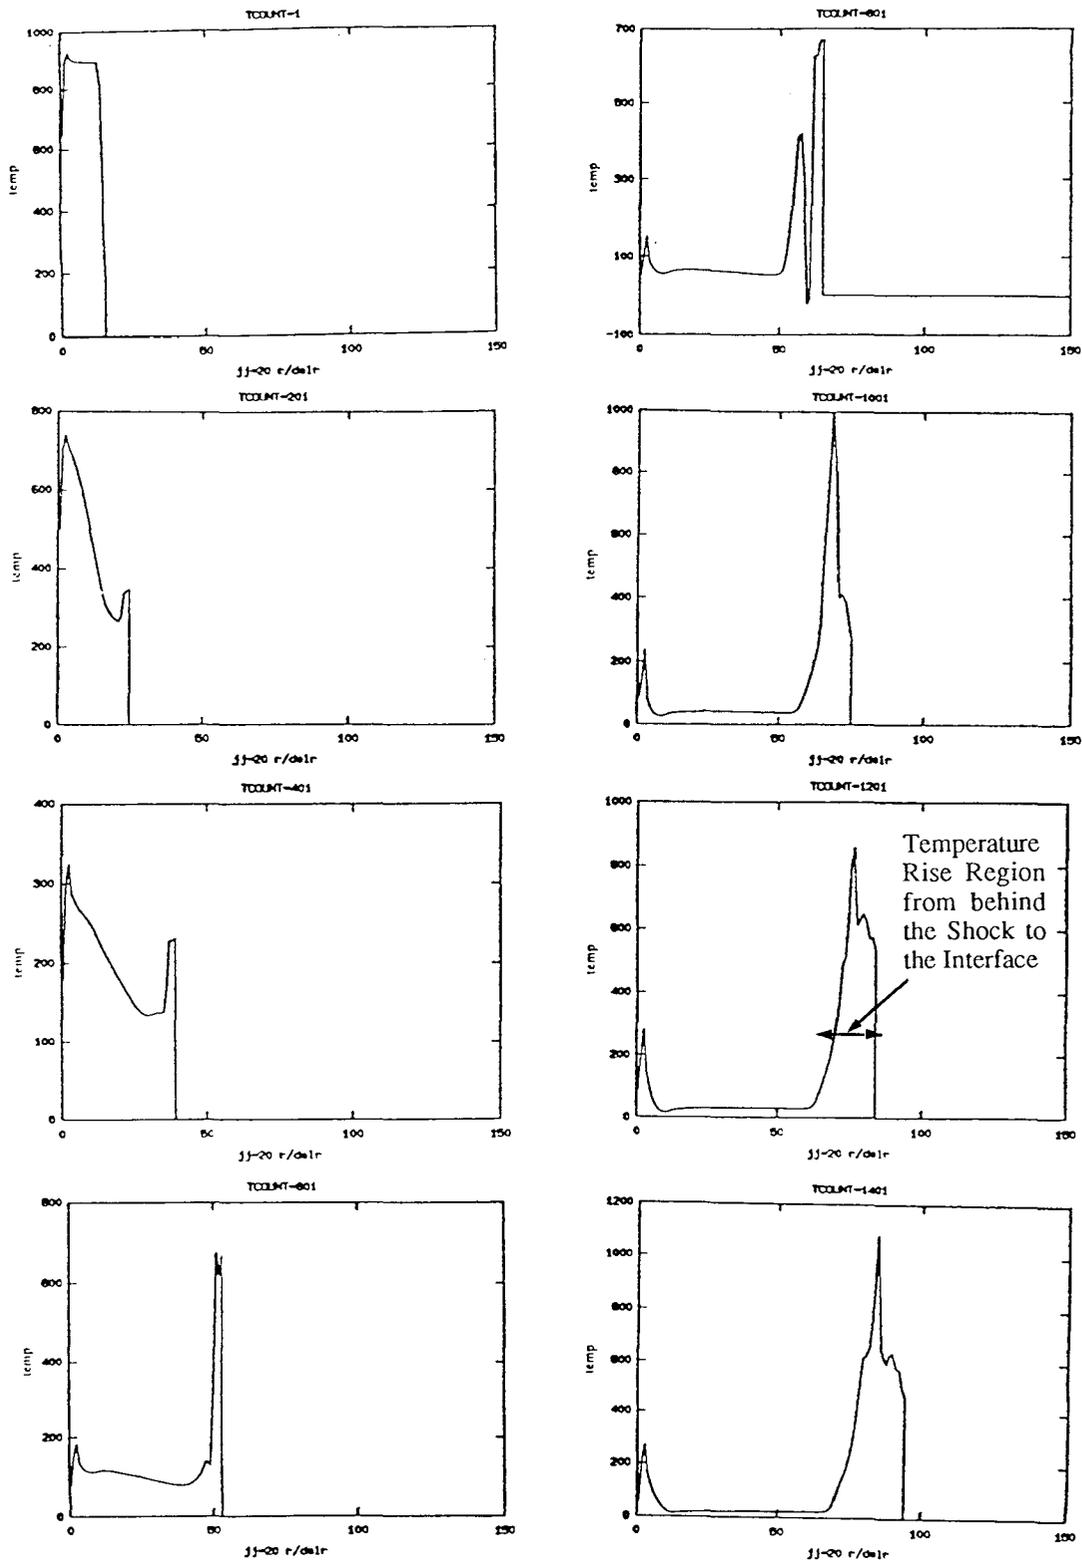


Figure 7.7c

Nondimensionalized temperature (T') (vertical axis)
 v.s. $r/\Delta r$ (horizontal axis) for $\theta = 19.5\Delta\theta$.

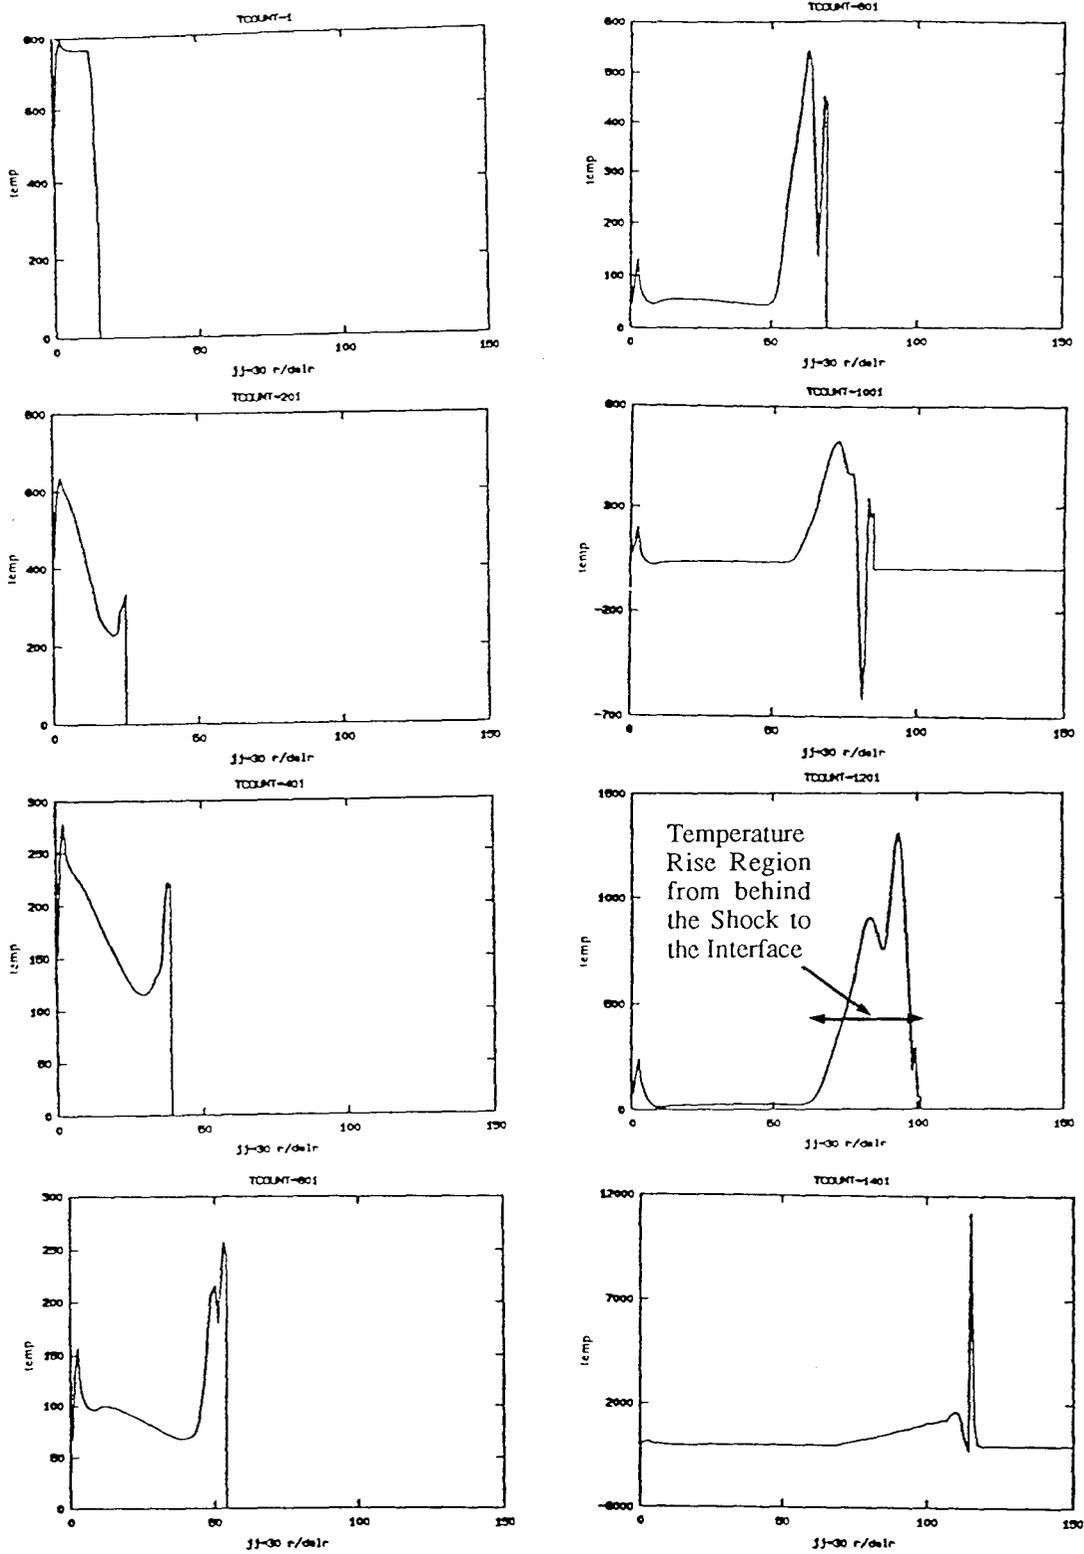


Figure 7.7d

Nondimensionalized temperature (T') (vertical axis)
 v.s. $r/\Delta r$ (horizontal axis) for $\theta = 29.5\Delta\theta$.

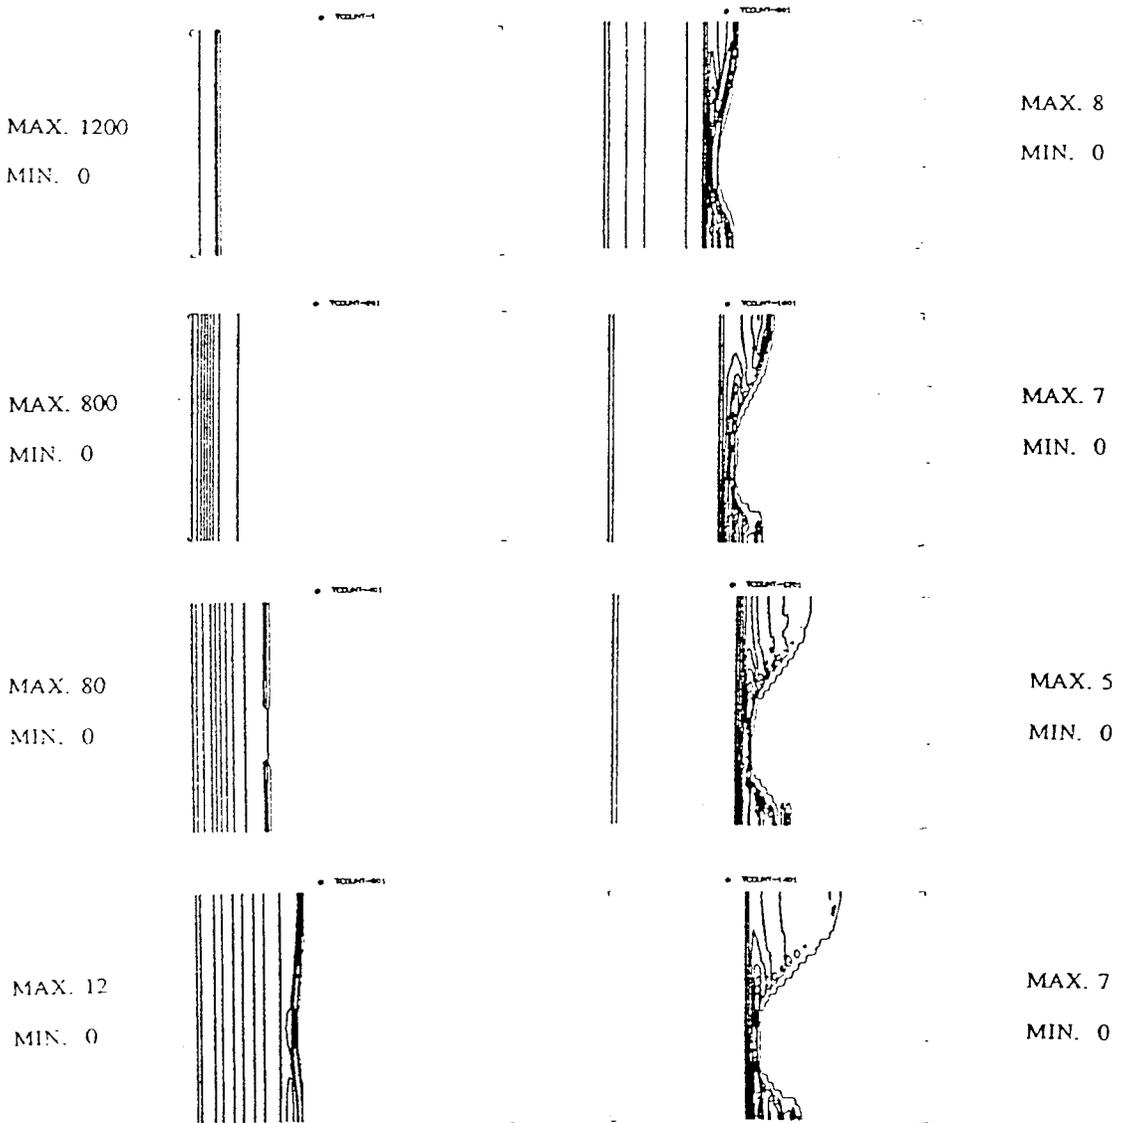
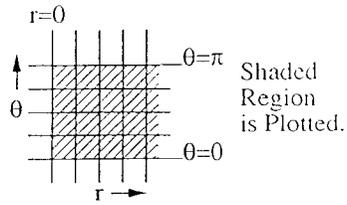


Figure 7.8

Contour plots of nondimensionalized pressure (p') in the rectangular r - θ plane.

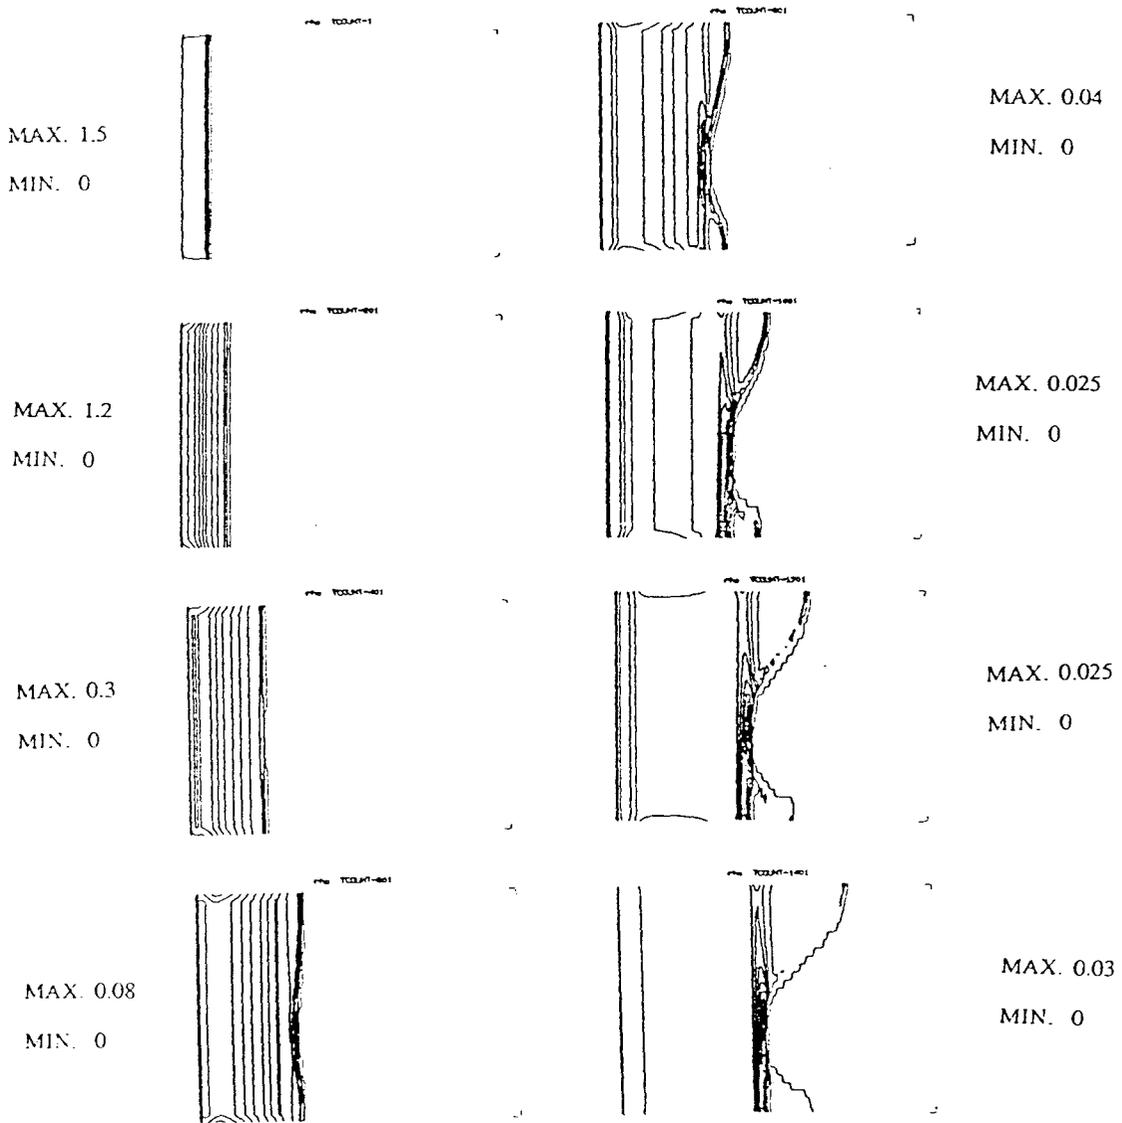


Figure 7.9

Contour plots of nondimensionalized density (ρ')
in the rectangular r - θ plane.

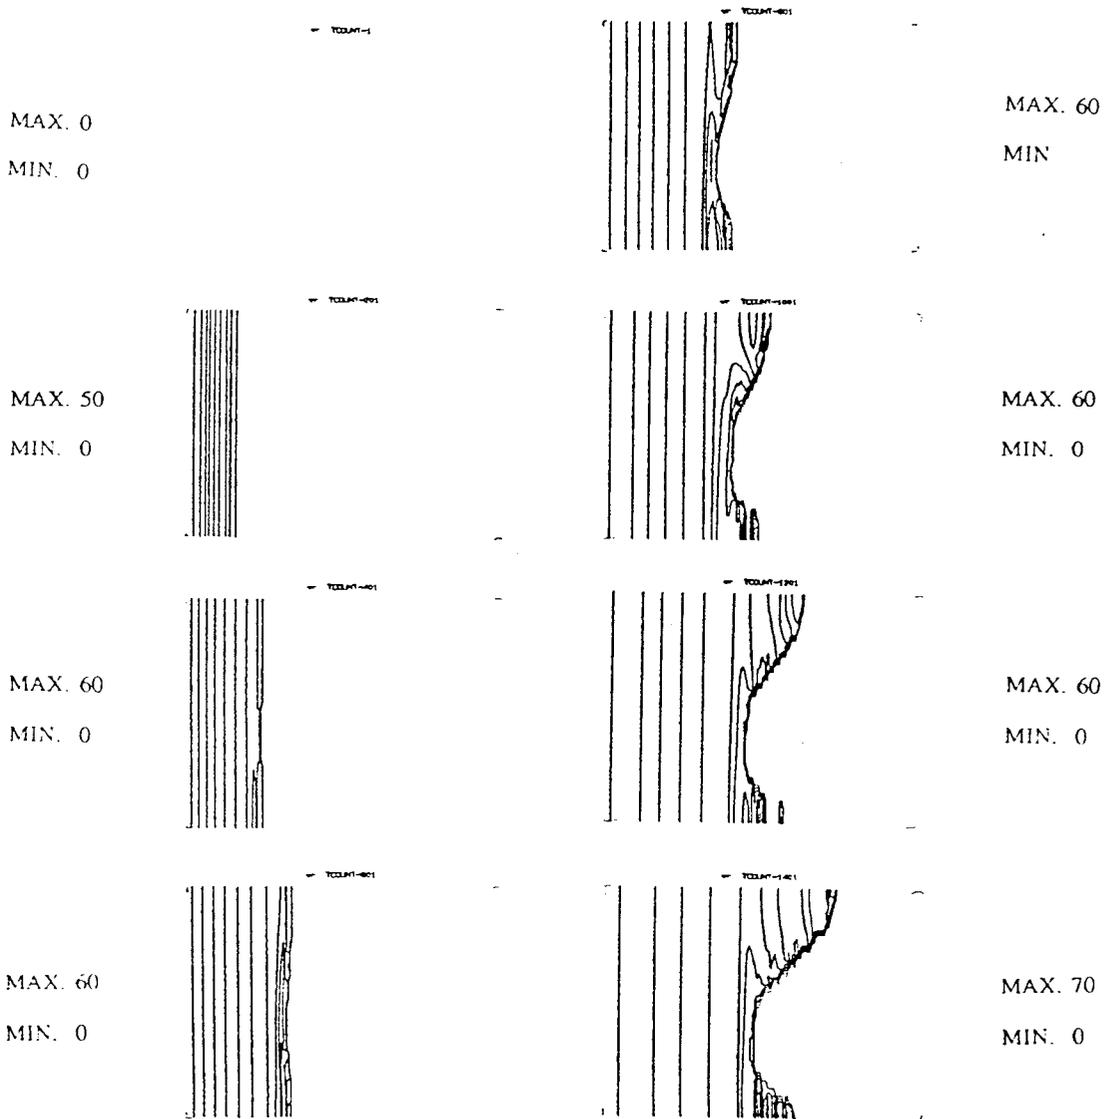


Figure 7.10 Contour plots of nondimensionalized radial velocity (v_r') in the rectangular r - θ plane.

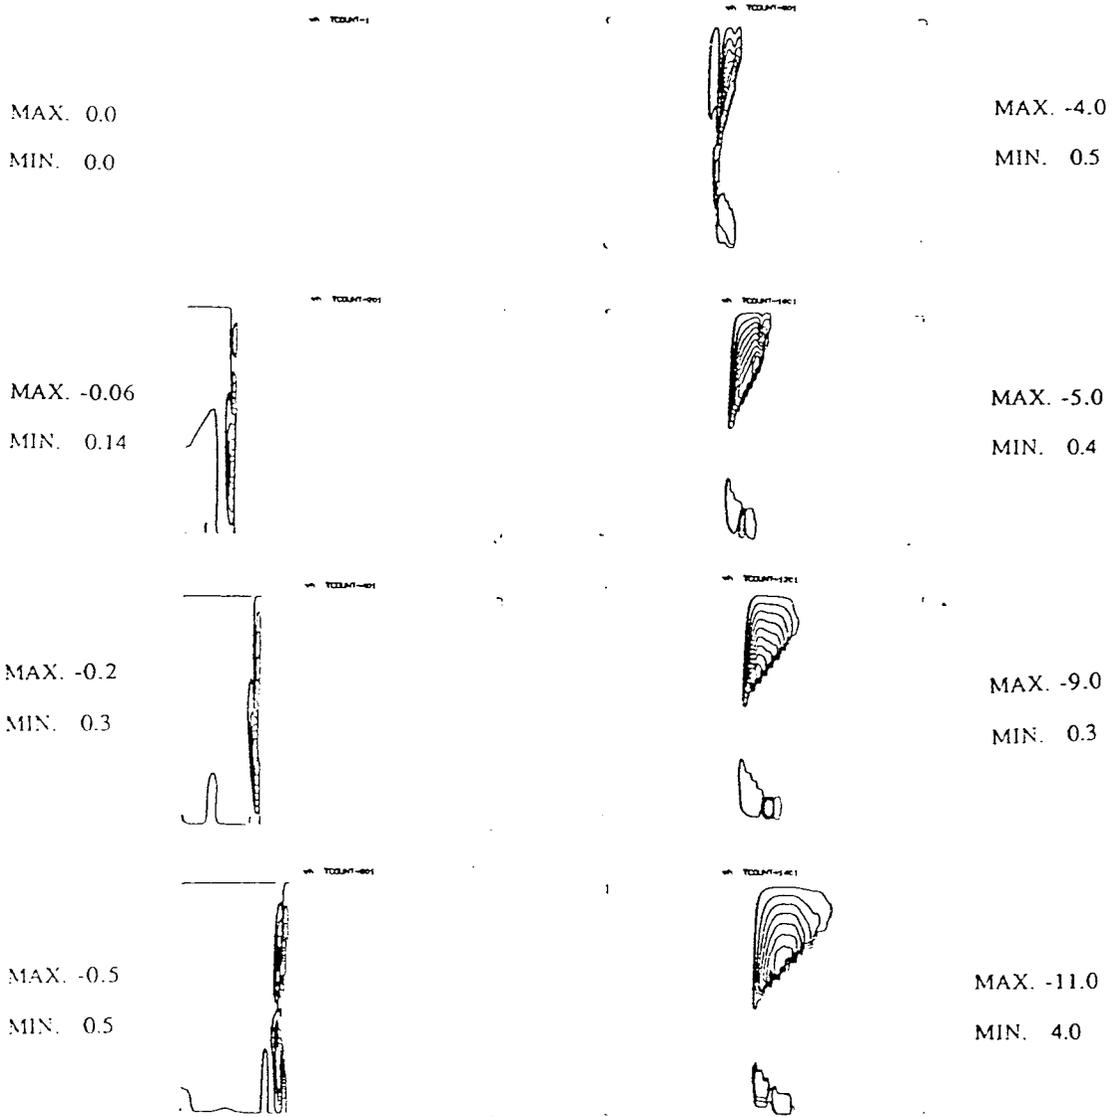


Figure 7.11 Contour plots of nondimensionalized polar velocity (v_θ') in the rectangular r - θ plane.

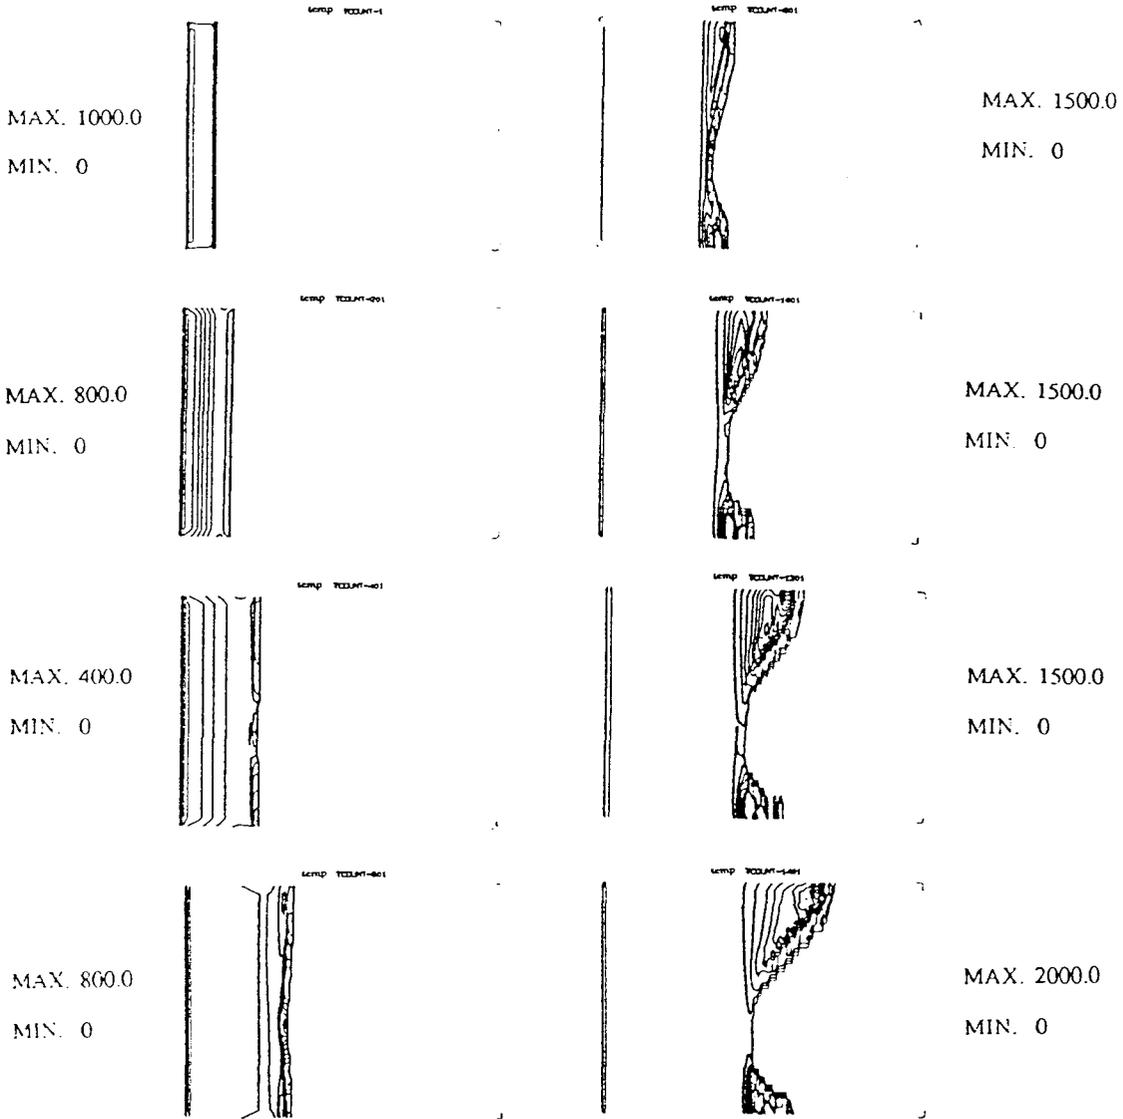


Figure 7.12 Contour plots of nondimensionalized temperature (T') in the rectangular $r-\theta$ plane.

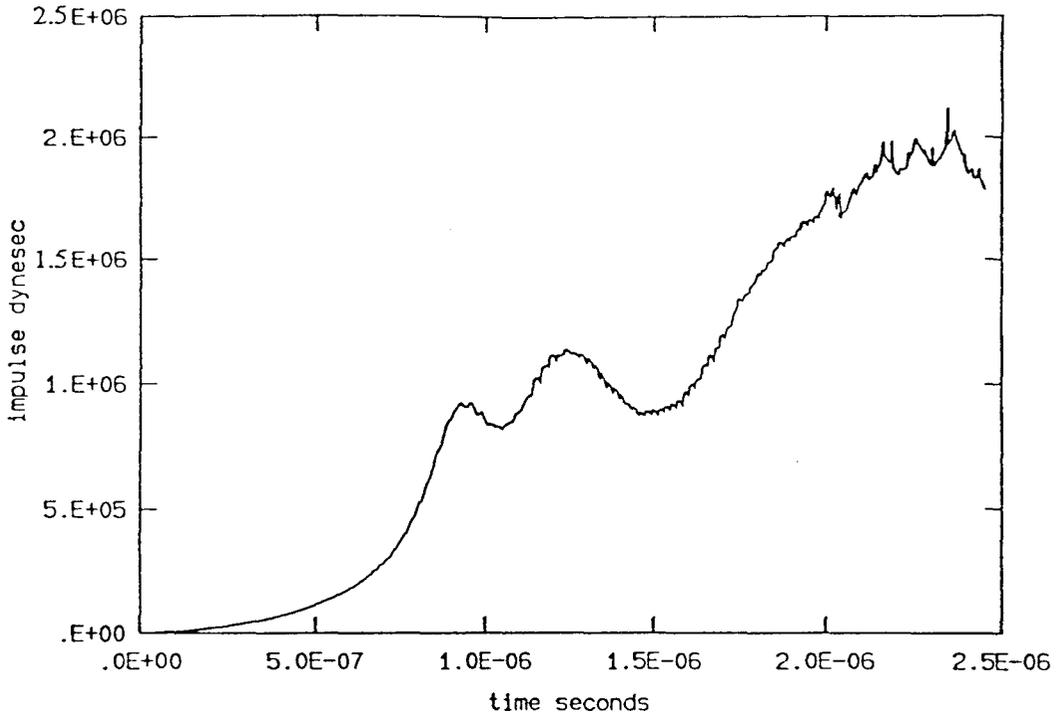


Figure 7.13a

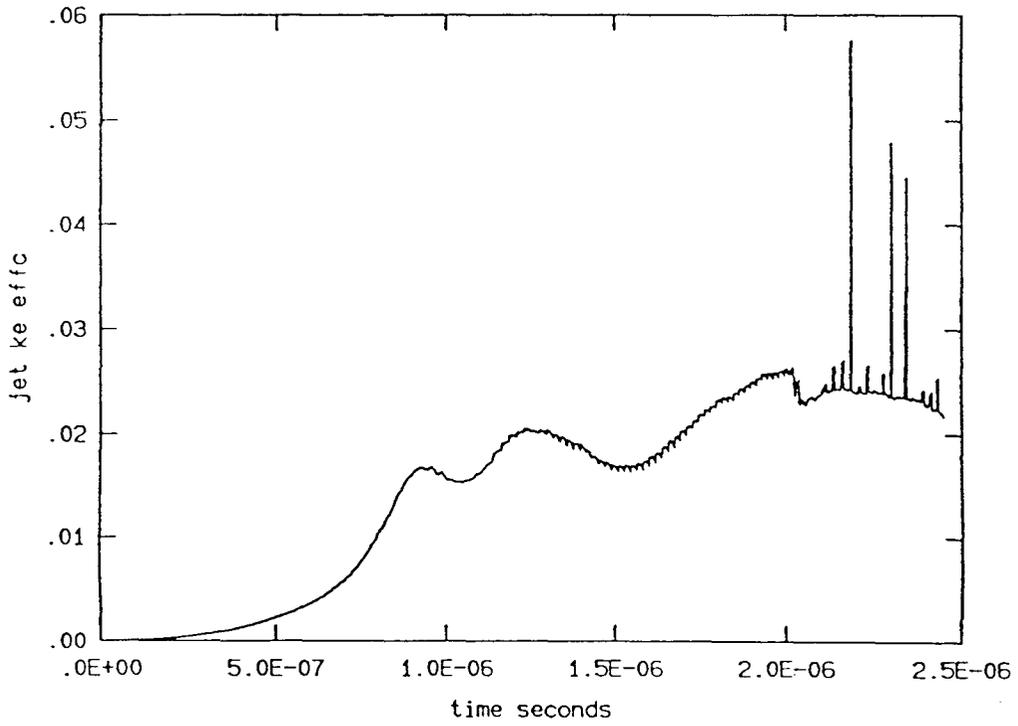


Figure 7.13b

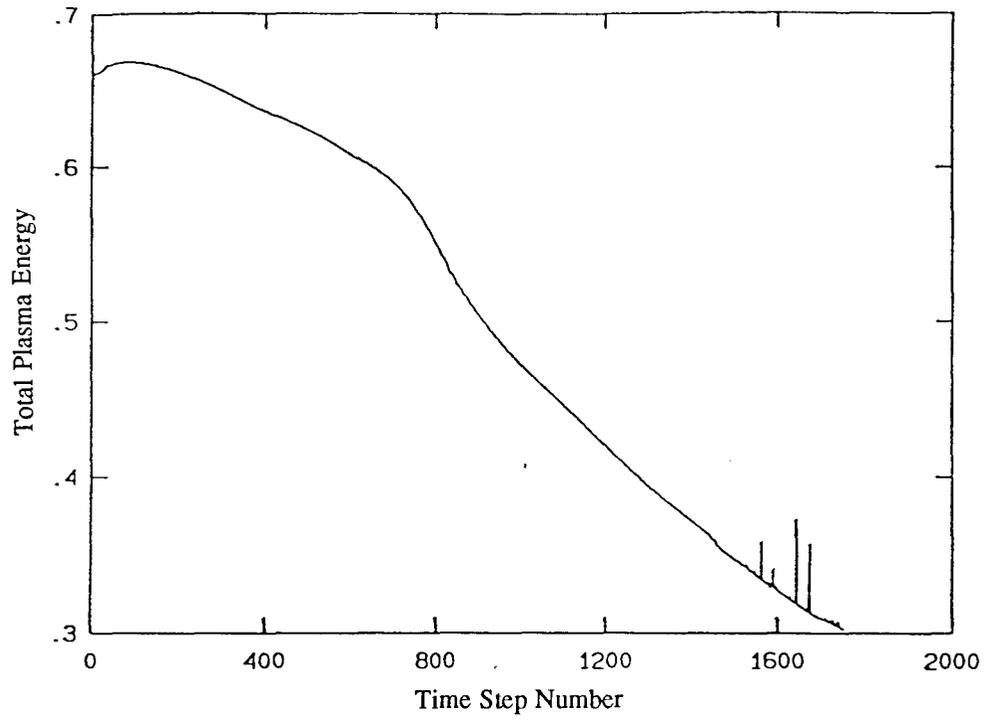


Figure 7.14

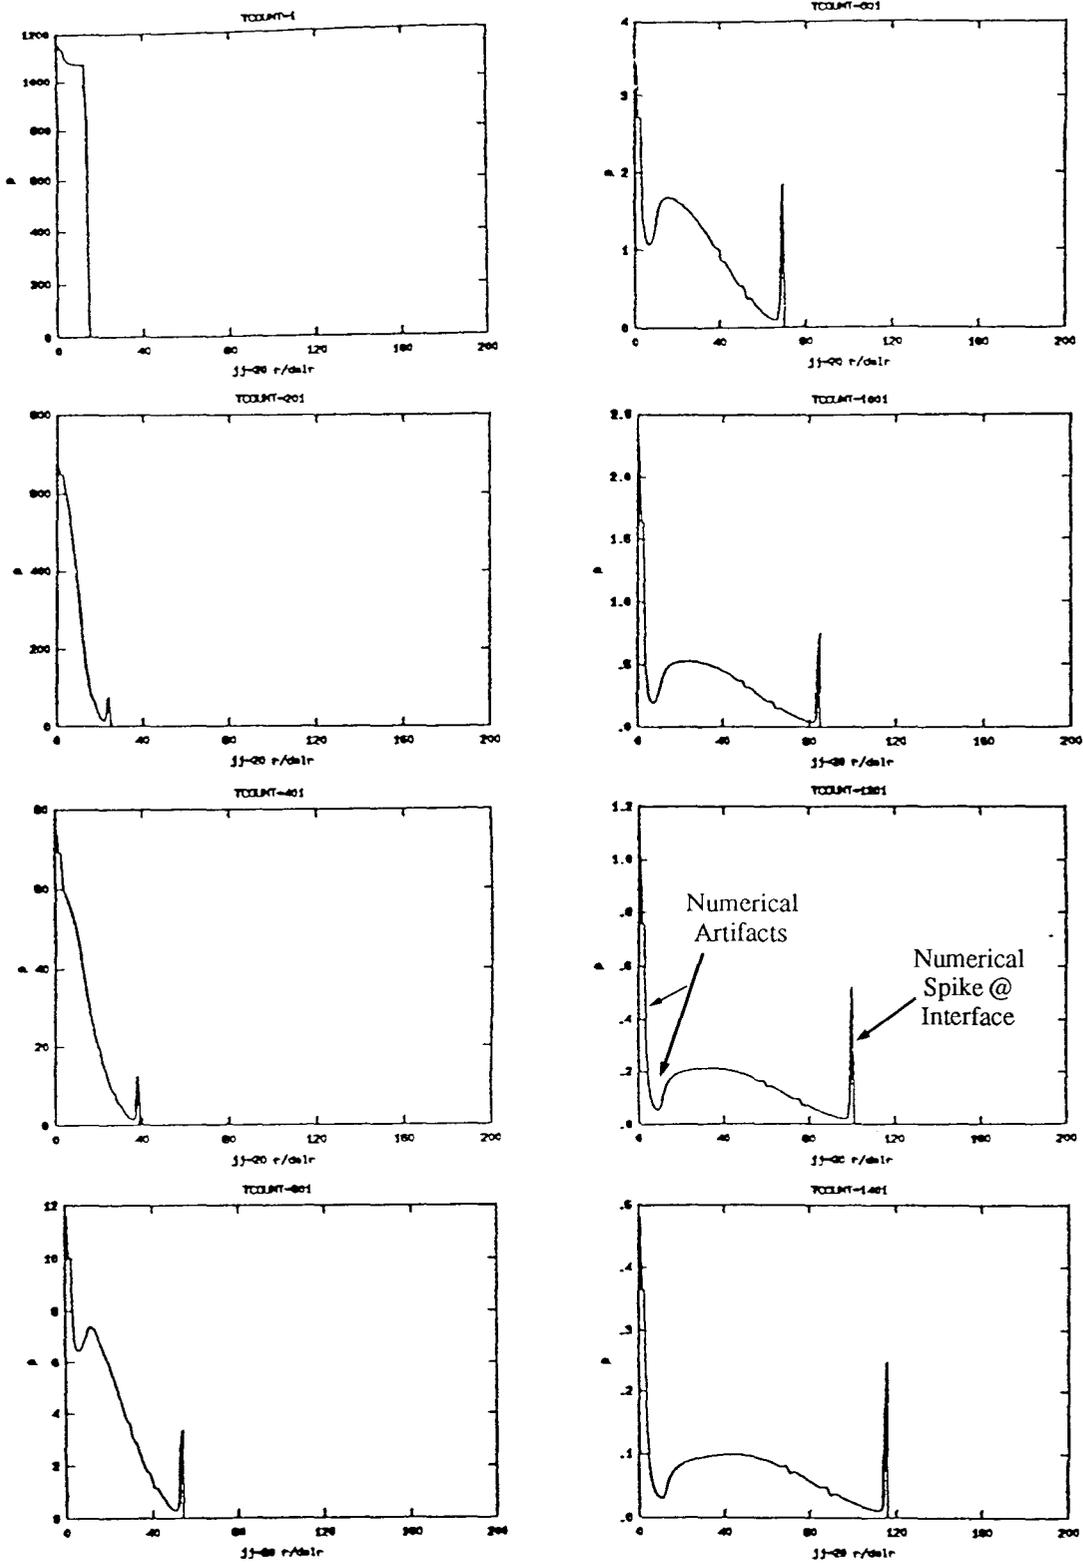


Figure 7.15

Nondimensionalized pressure (p') (vertical axis)
 v.s. $r / \Delta r$ (horizontal axis) for Case 2
 (Free expansion).

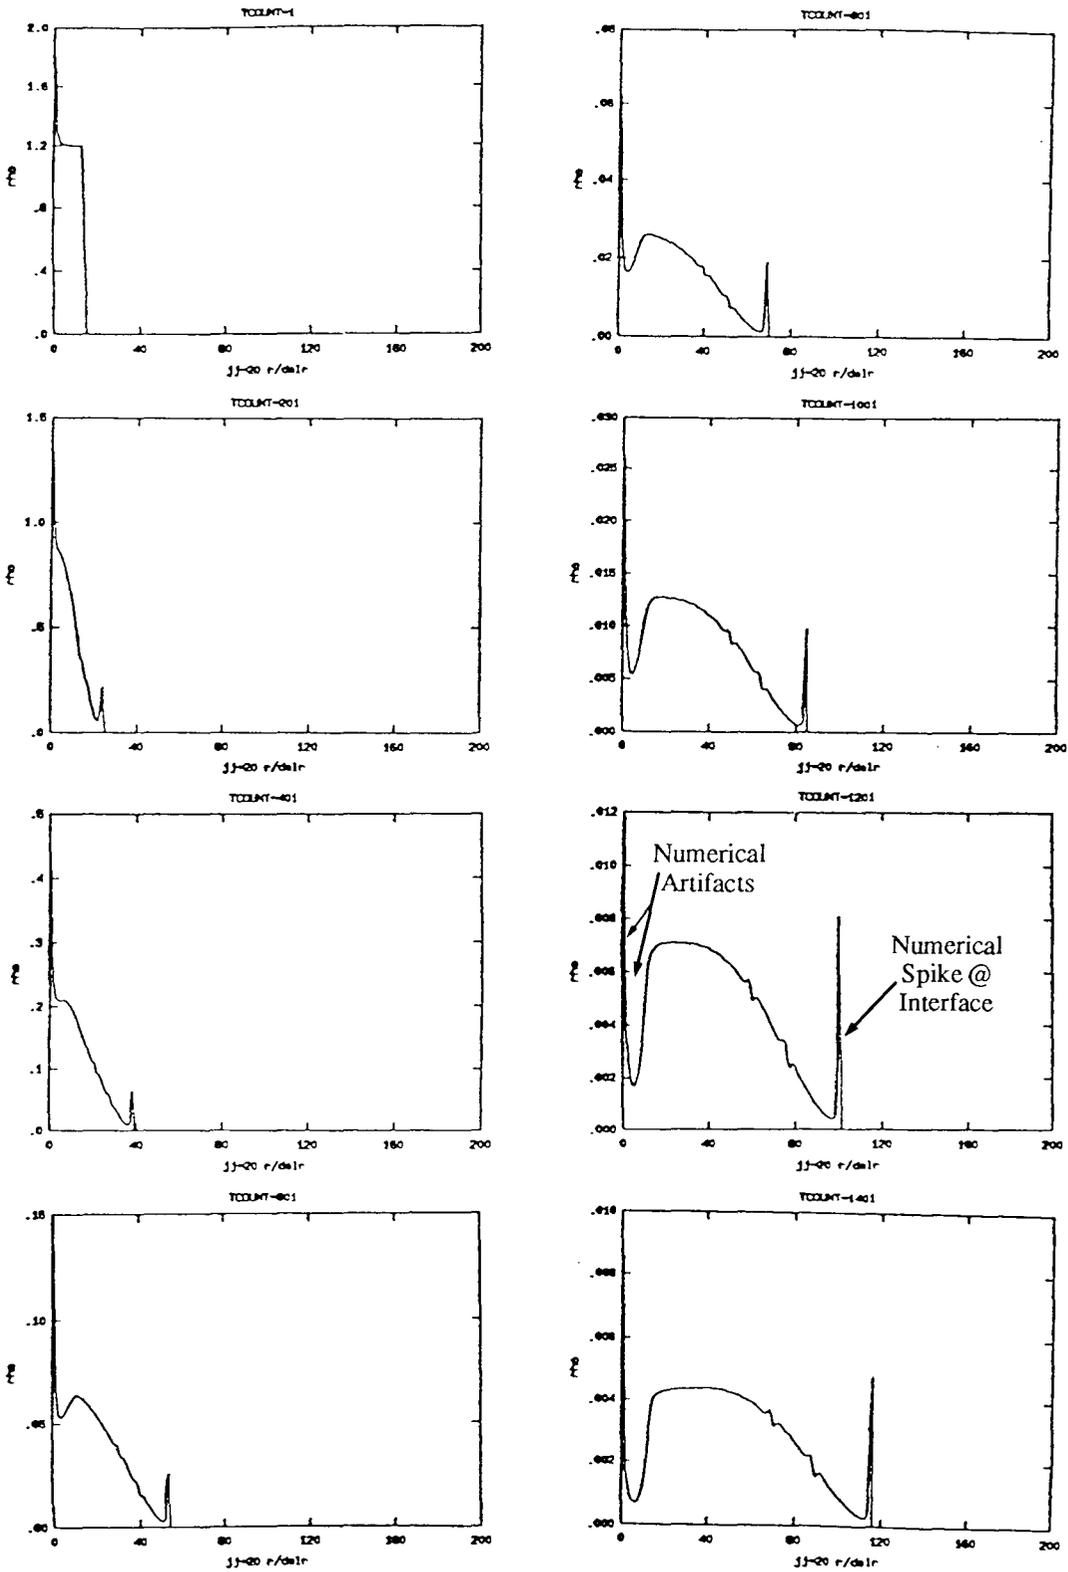


Figure 7.16

Nondimensional density (ρ') (vertical axis)
 v.s. $r / \Delta r$ (horizontal axis) for Case 2.

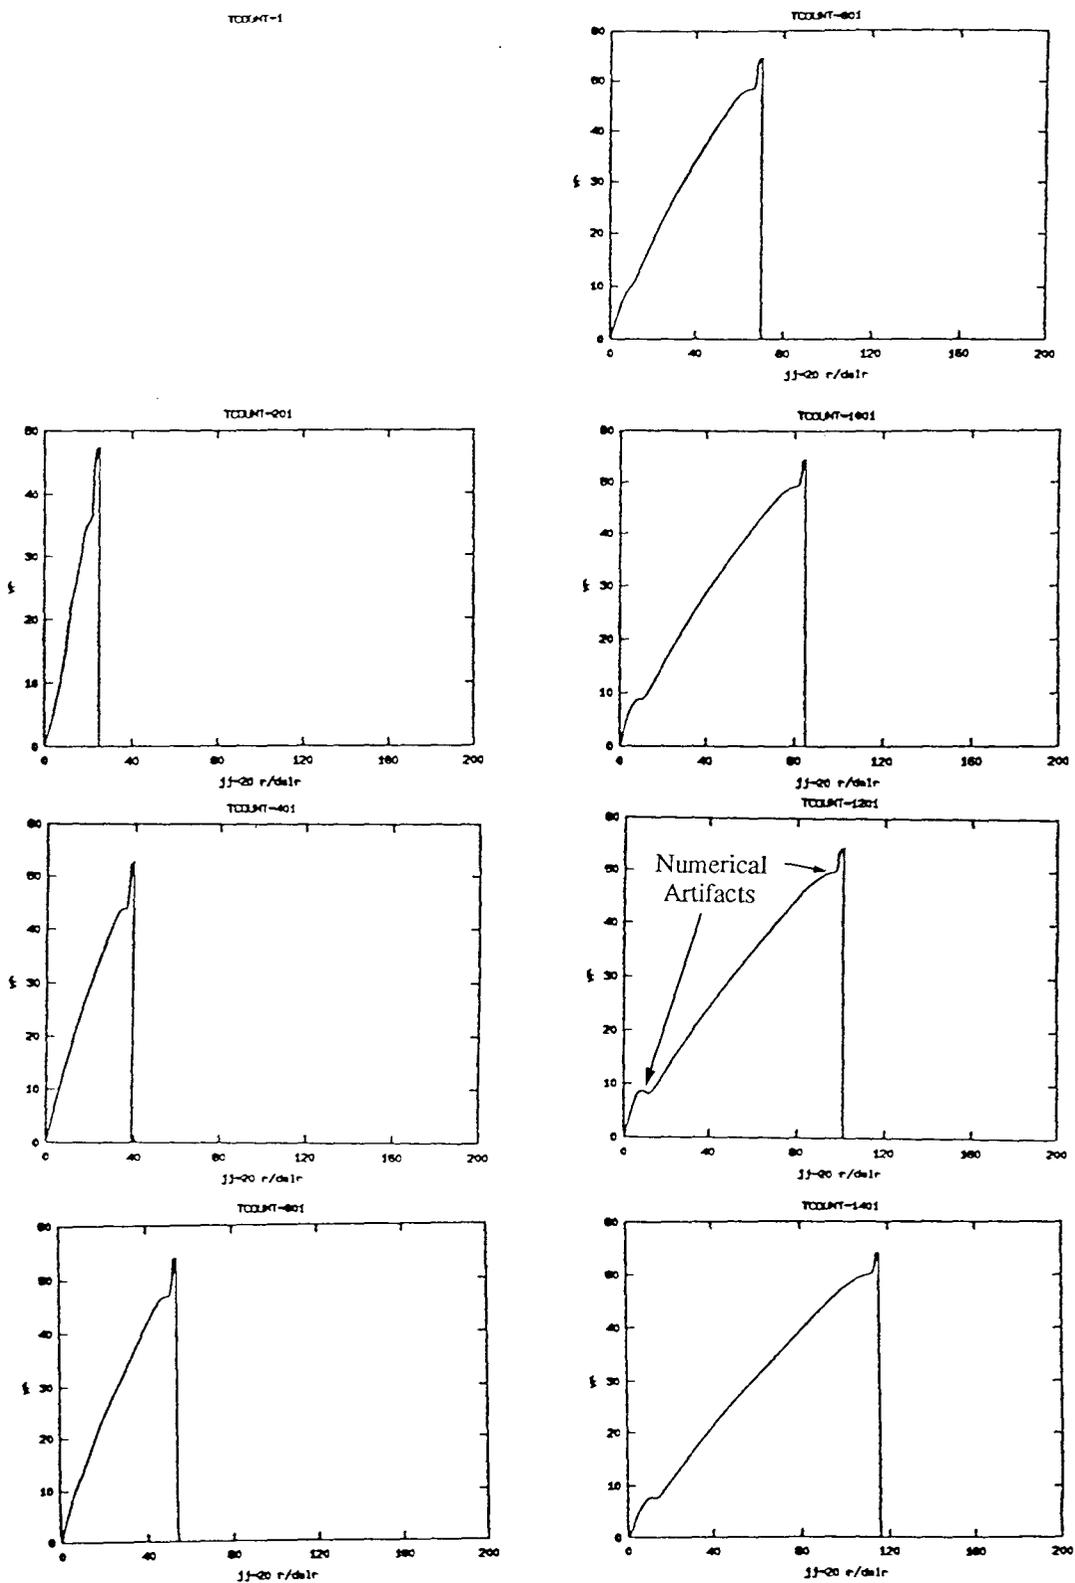


Figure 7.17 Nondimensionalized radial velocity (v_r') (vertical axis) v.s. $r/\Delta r$ (horizontal axis) for Case 2.

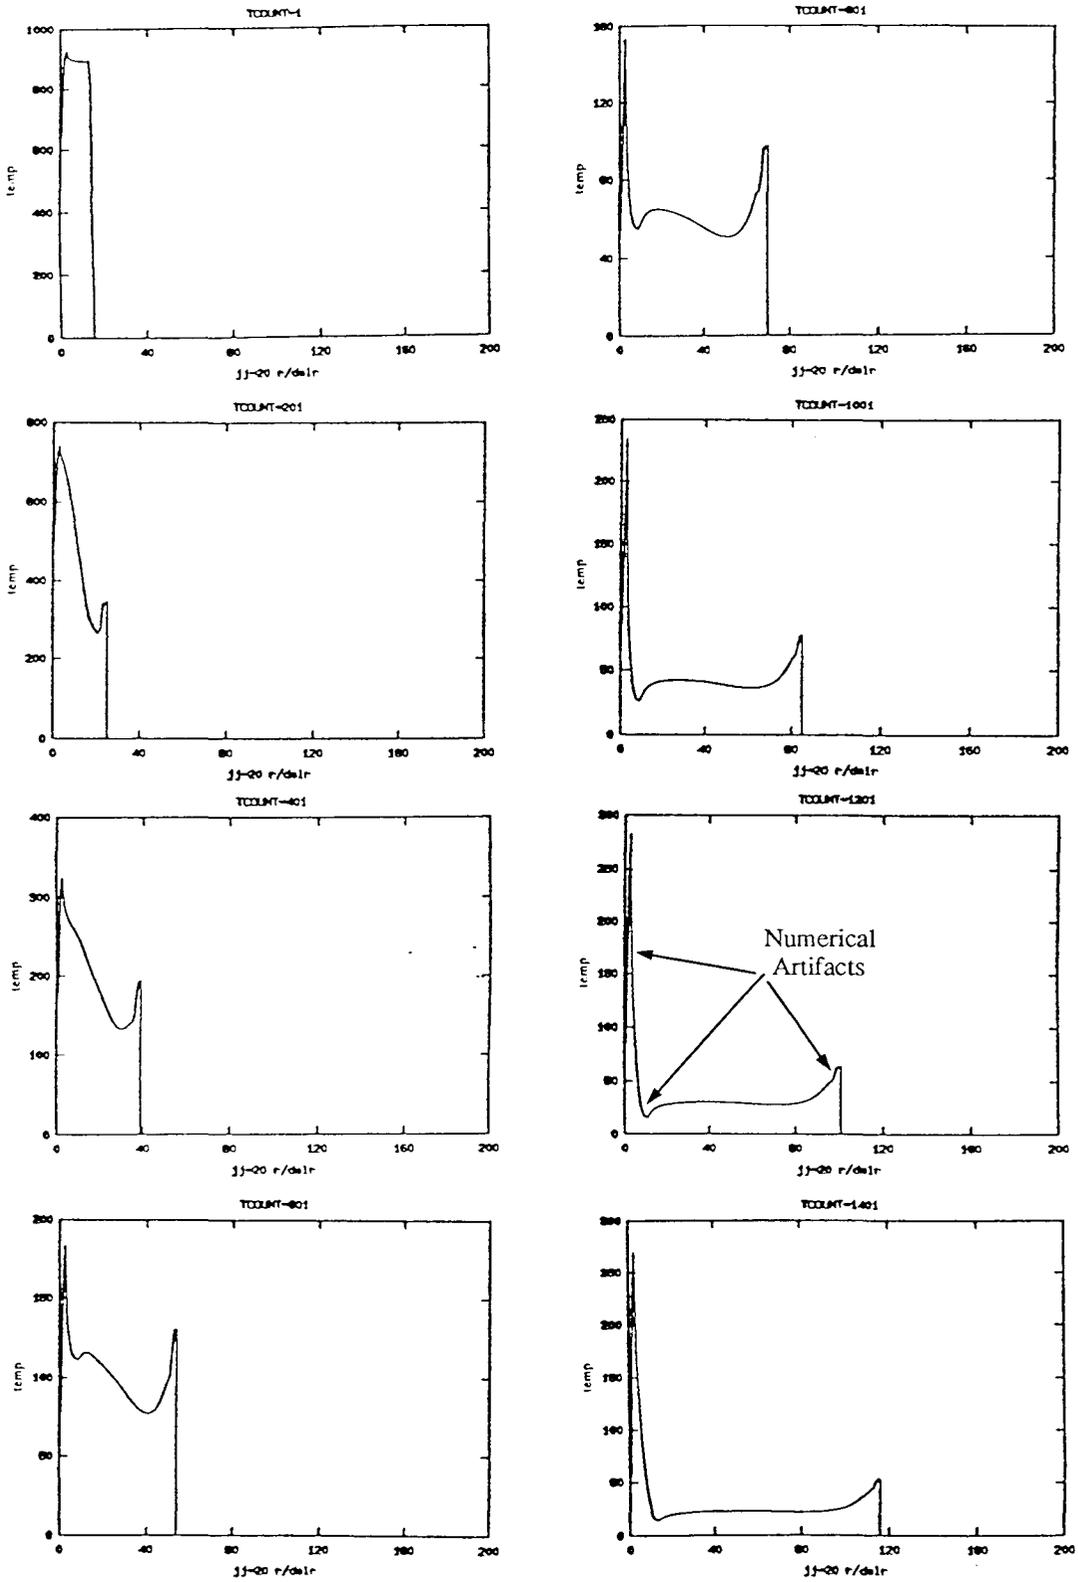


Figure 7.18

Nondimensional temperature (T) (vertical axis)
v.s. $r / \Delta r$ (horizontal axis) for Case 2.

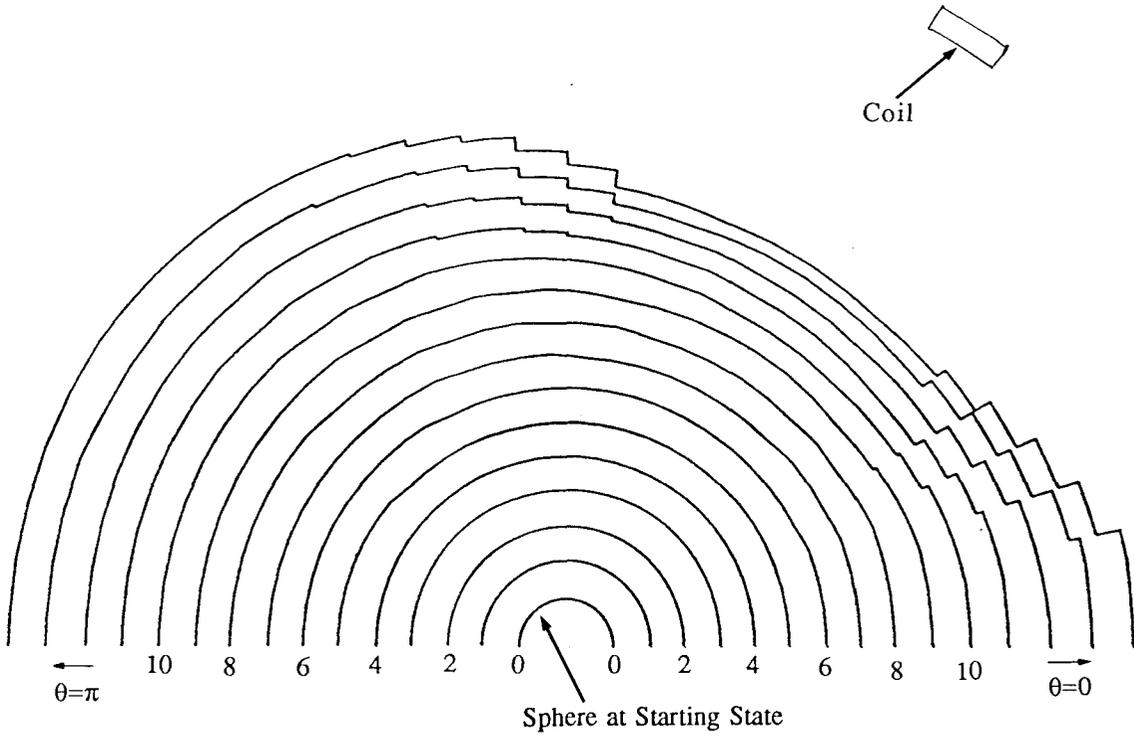


Figure 7.19 Locations of the interface at evenly spaced time intervals for Case 3.
($ER \approx 4.2$, $ER_i \approx 0.56$, $PR \approx 1.5$)

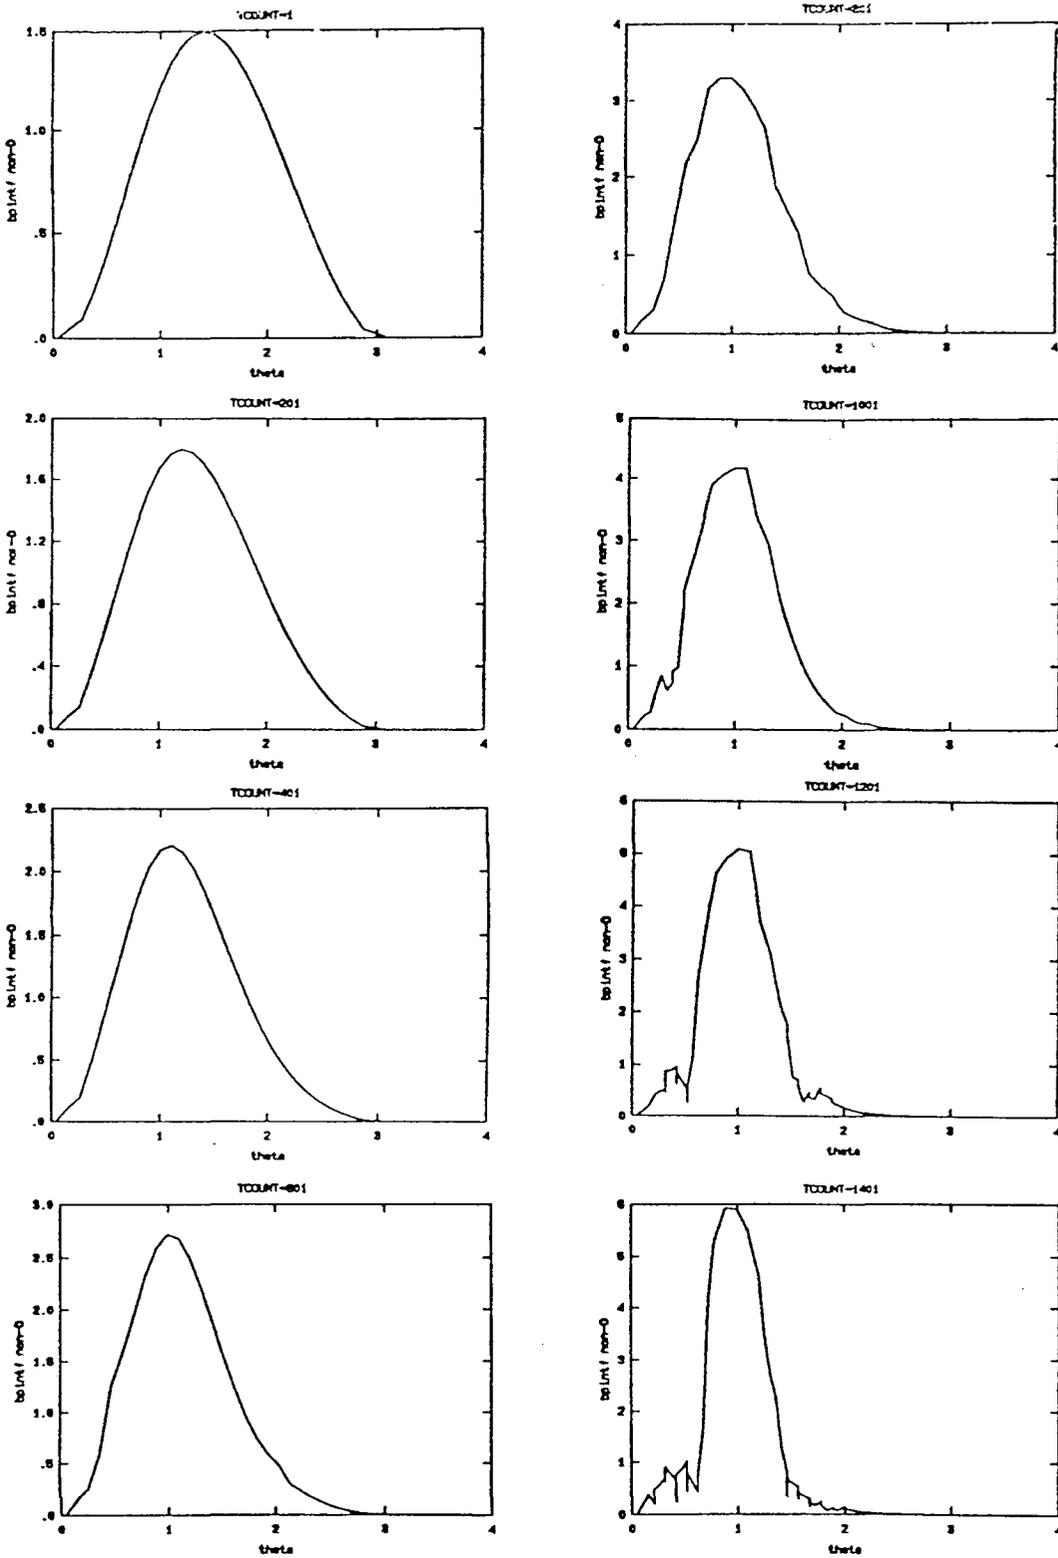


Figure 7.20

Nondimensionalized interfacial magnetic pressure (p_B'')
 (vertical axis) v.s. polar angle (θ) in radians (horizontal axis).

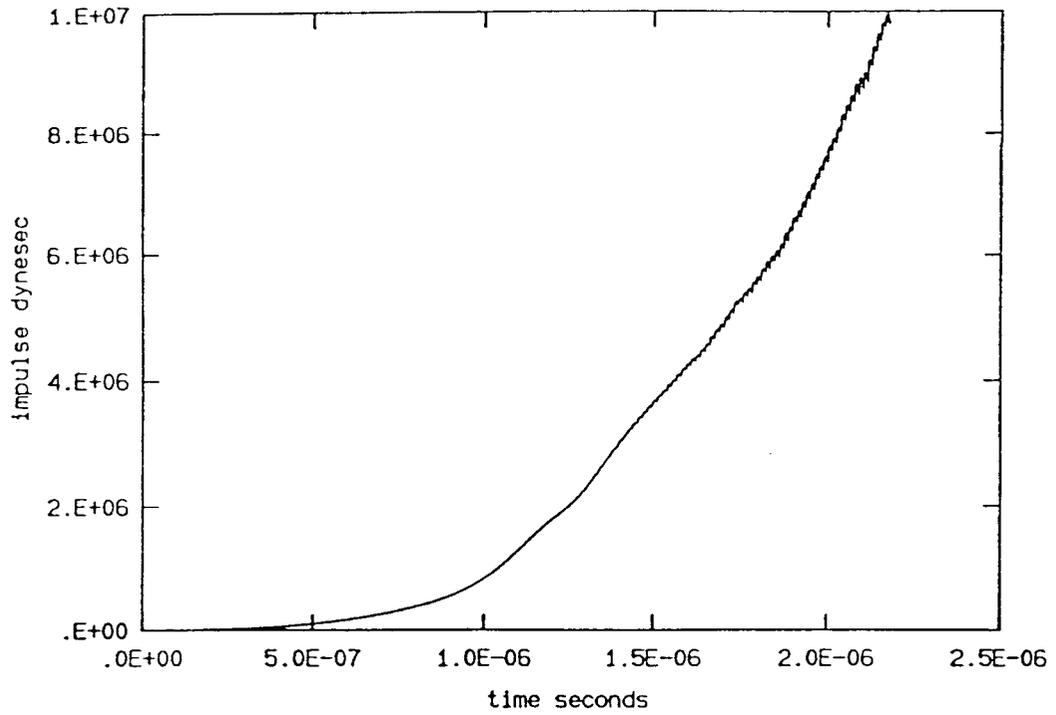


Figure 7.21a

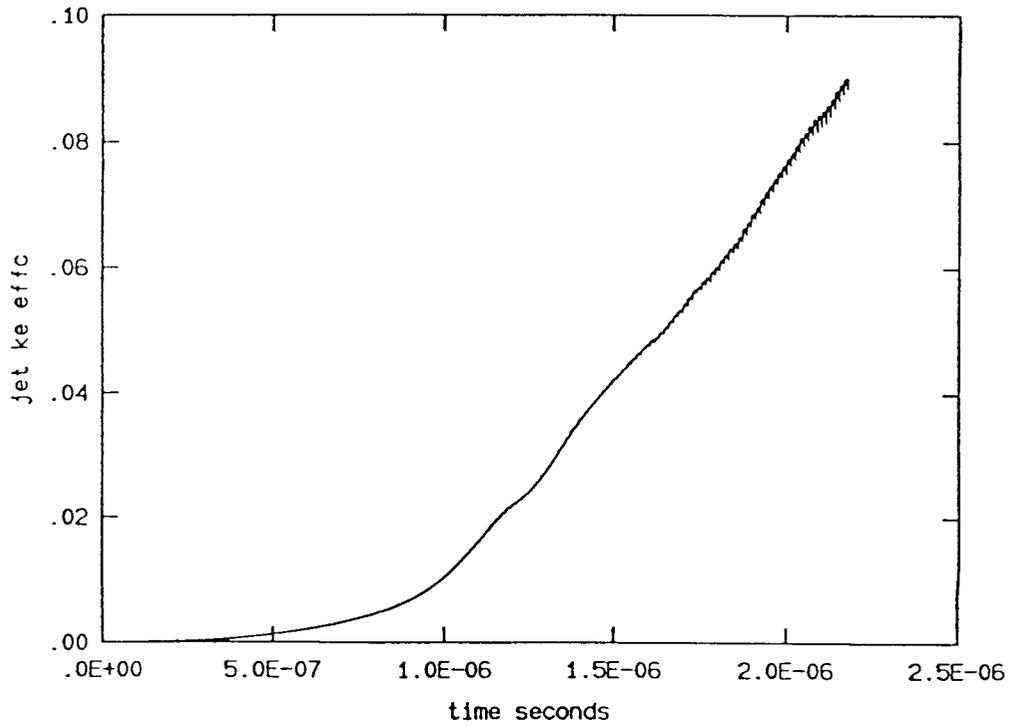


Figure 7.21b

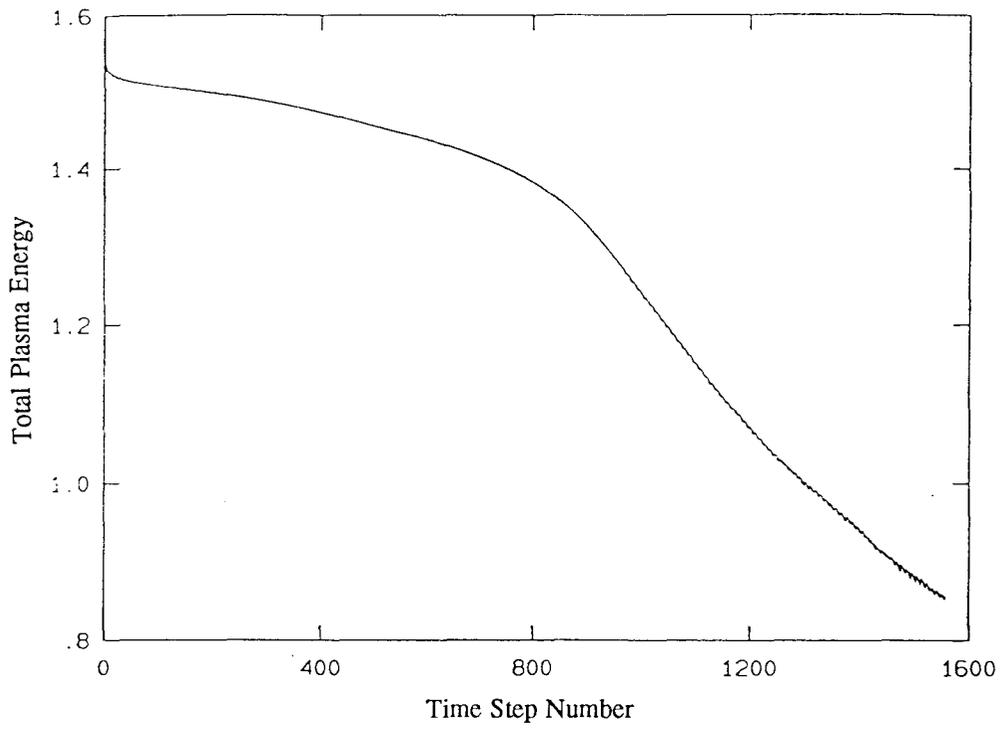


Figure 7.22

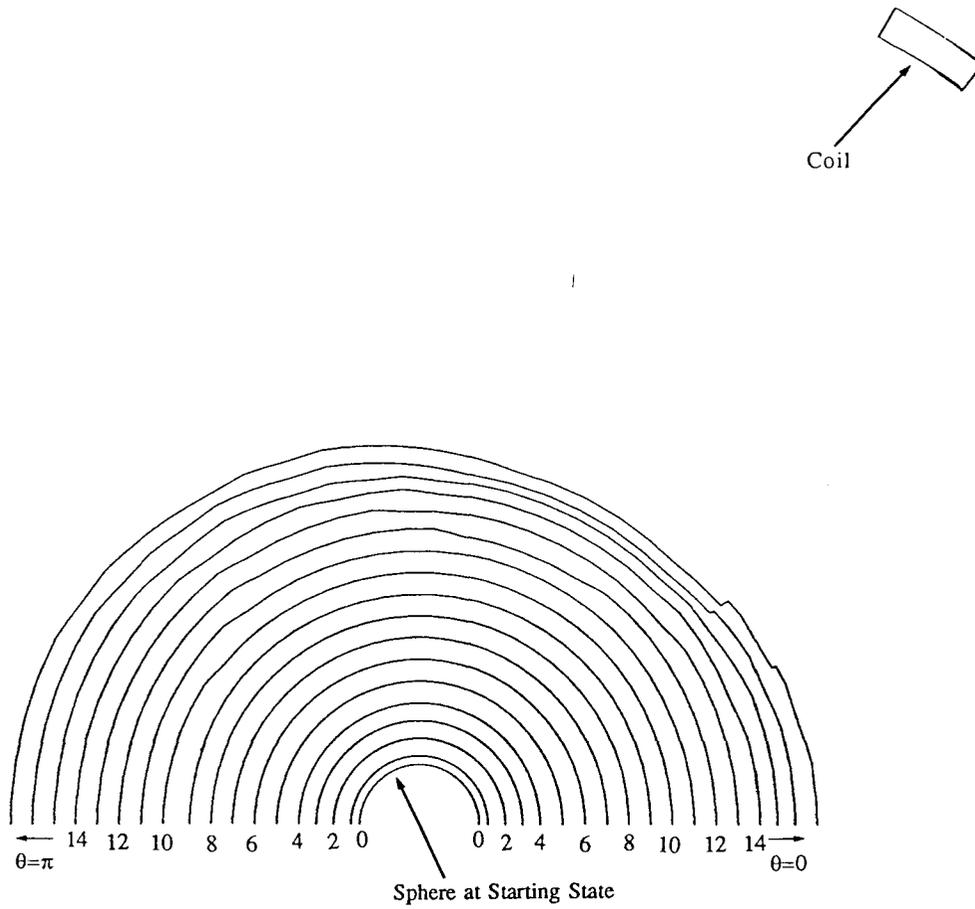


Figure 7.23 Locations of the interface at evenly spaced time intervals for Case 4.
($ER \approx 5$, $ER_i \approx 0.67$, $PR \approx 0.0007$)

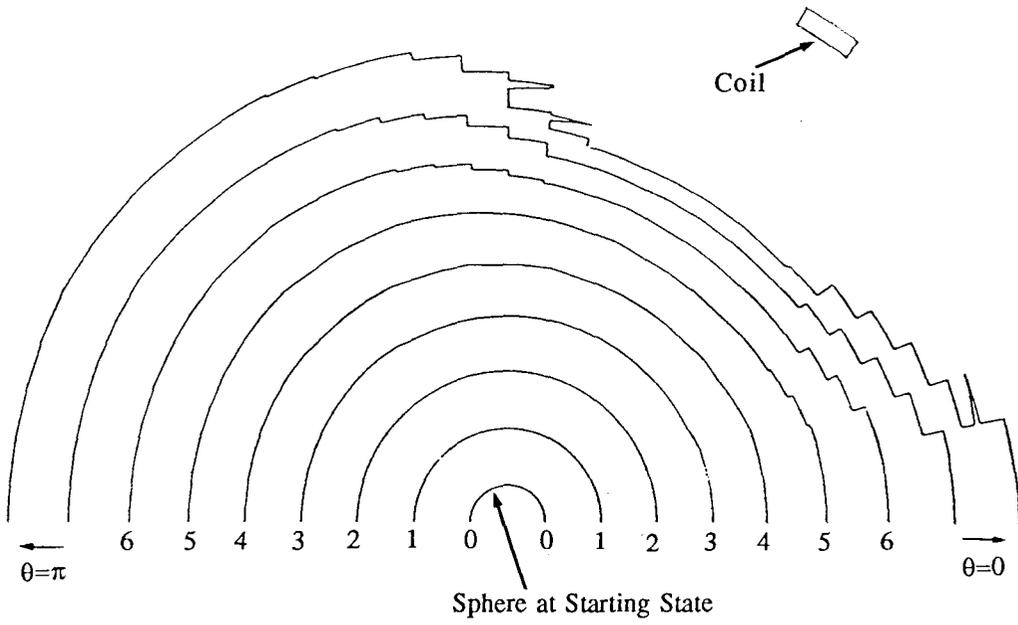


Figure 7.24 Locations of the interface at evenly spaced time intervals for Case 5.
 ($ER \approx 0.19$, $ER_1 \approx 0.55$, $PR \approx 1.5$)

VII-A41

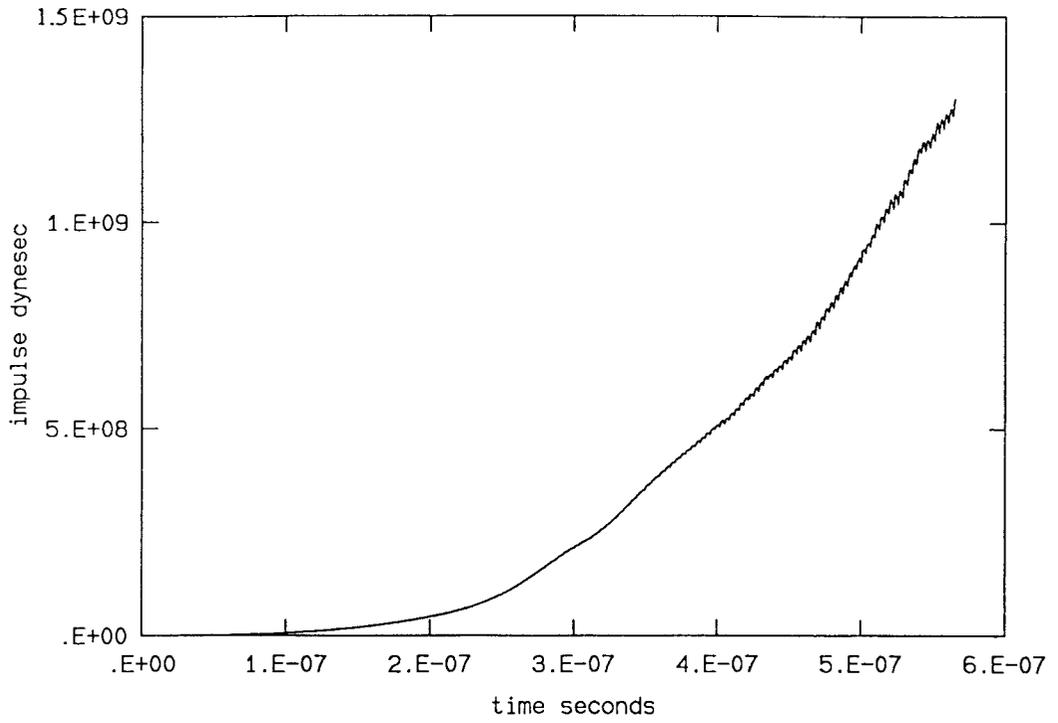


Figure 7.25a

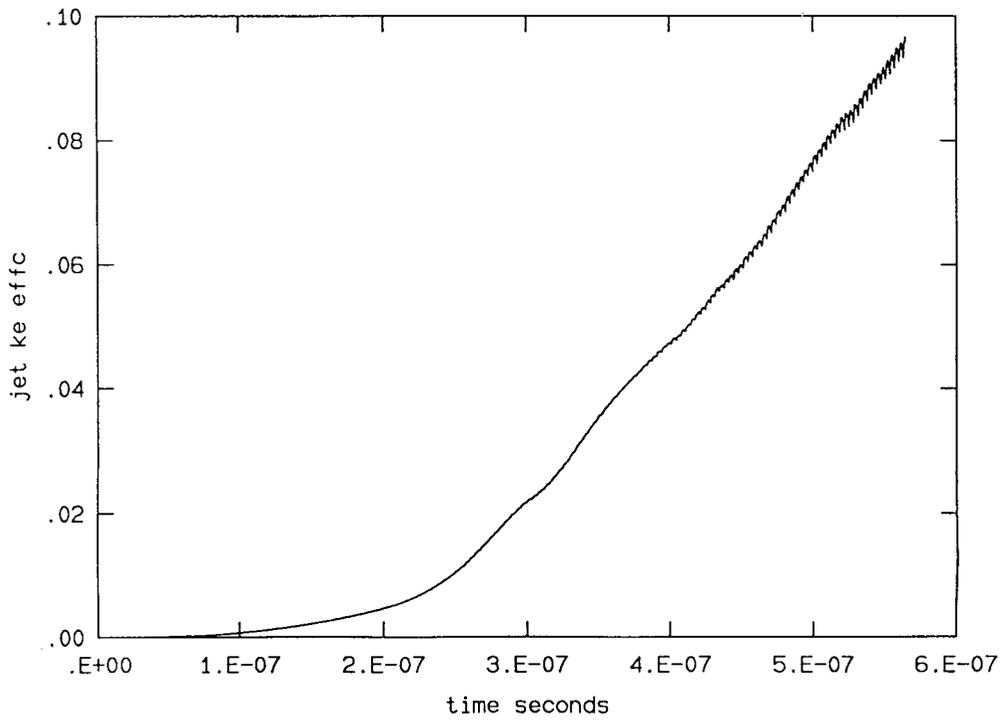


Figure 7.25b

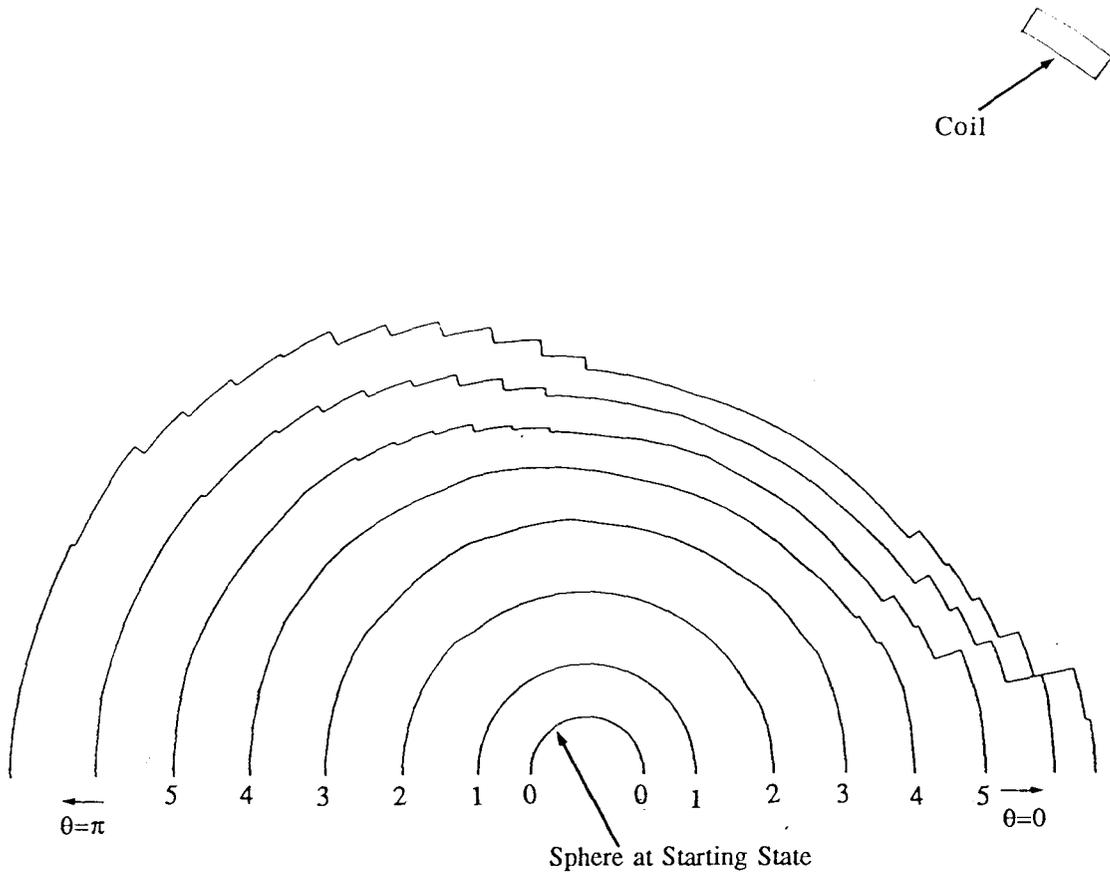


Figure 7.26 Locations of the interface at evenly spaced time intervals for Case 6.
($ER \approx 10.5$, $ER_i \approx 1.36$, $PR \approx 0.0014$)

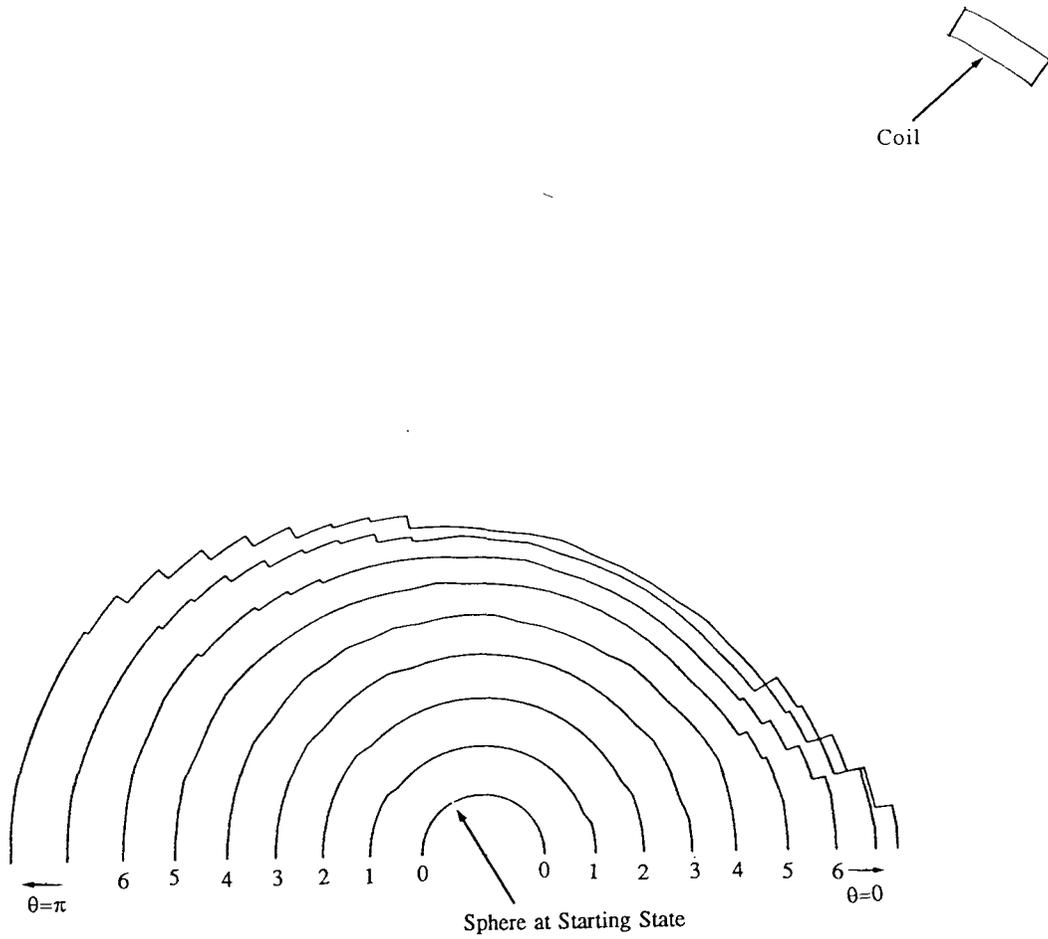


Figure 7.27 Locations of the interface at evenly spaced time intervals for Case 7.
($ER \approx 5$, $ER_1 \approx 6.6$, $PR \approx 0.3$)

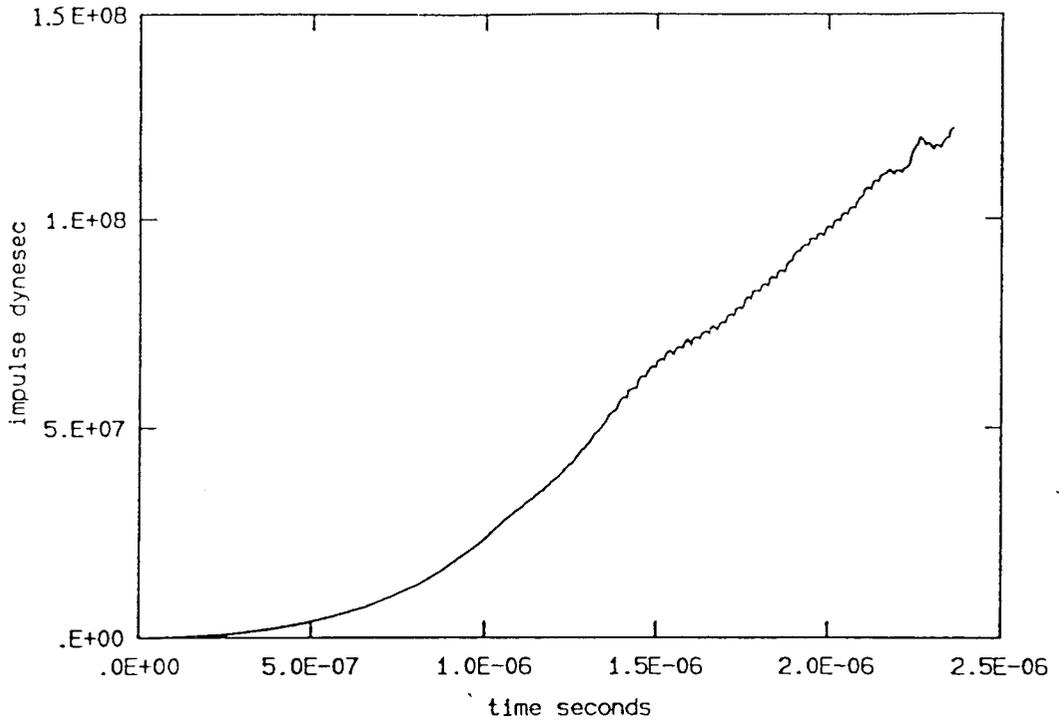


Figure 7.28a

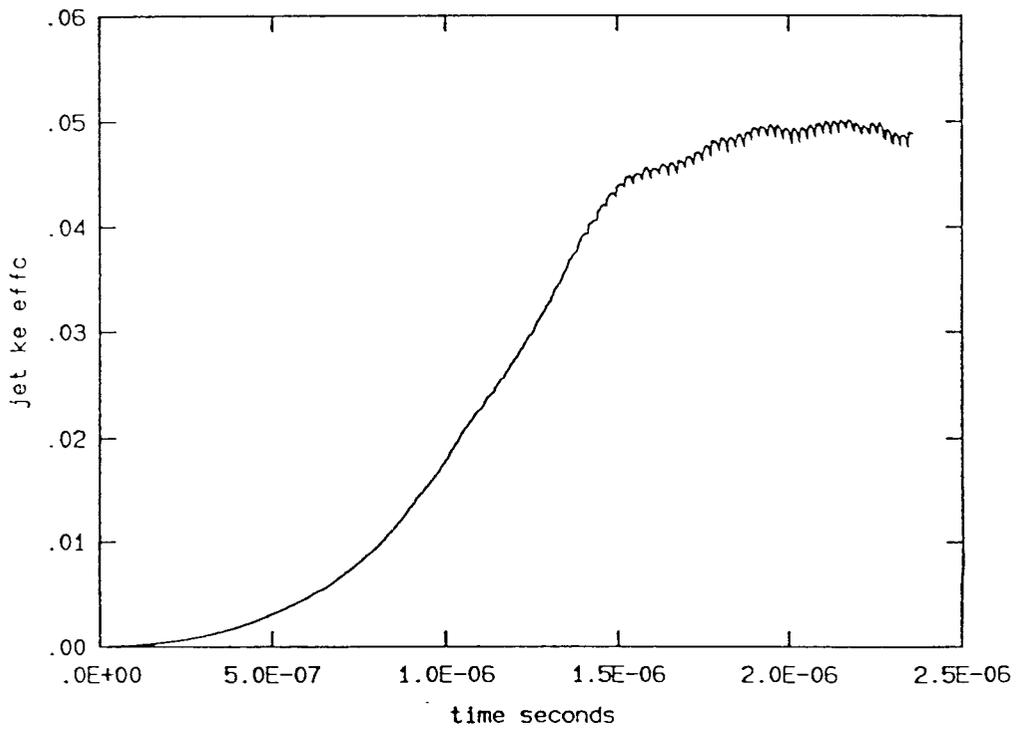


Figure 7.28b

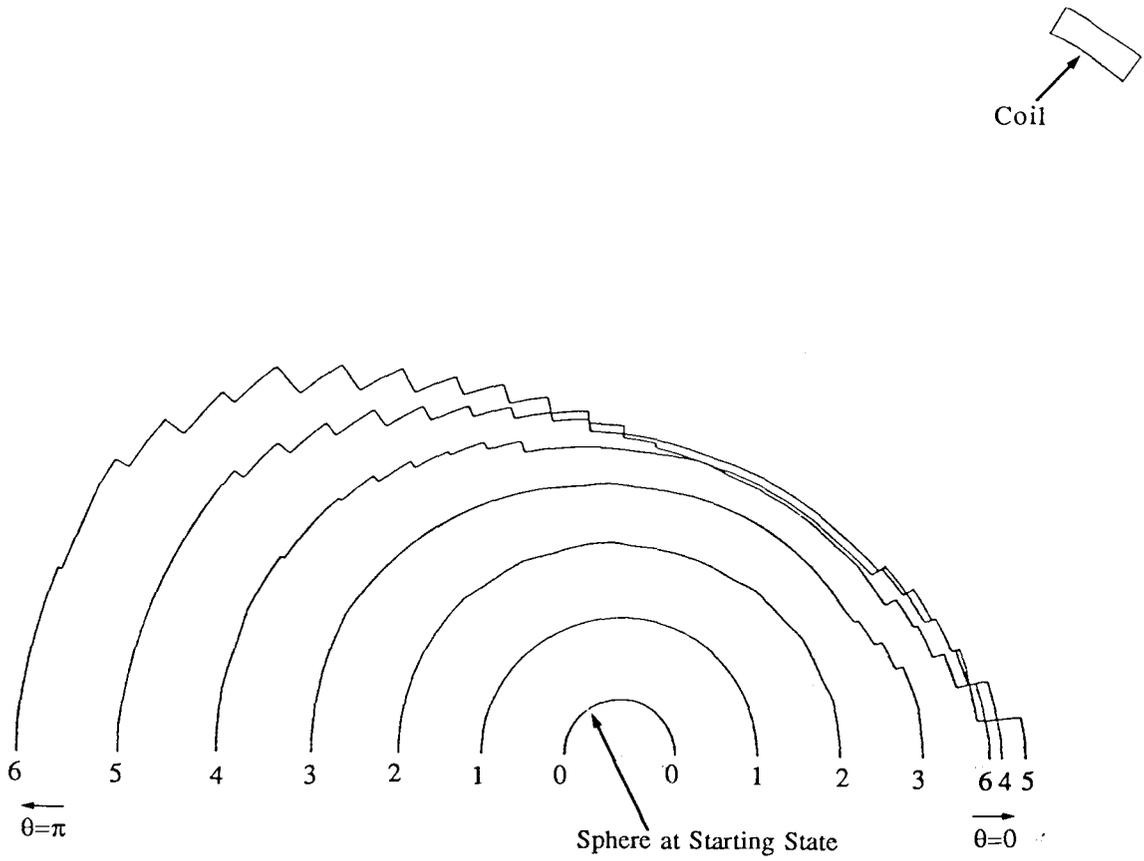


Figure 7.29

Locations of the interface at evenly spaced time intervals for Case 8.
 ($ER \approx 0.58$, $ER_i \approx 5.2$, $PR \approx 14$)

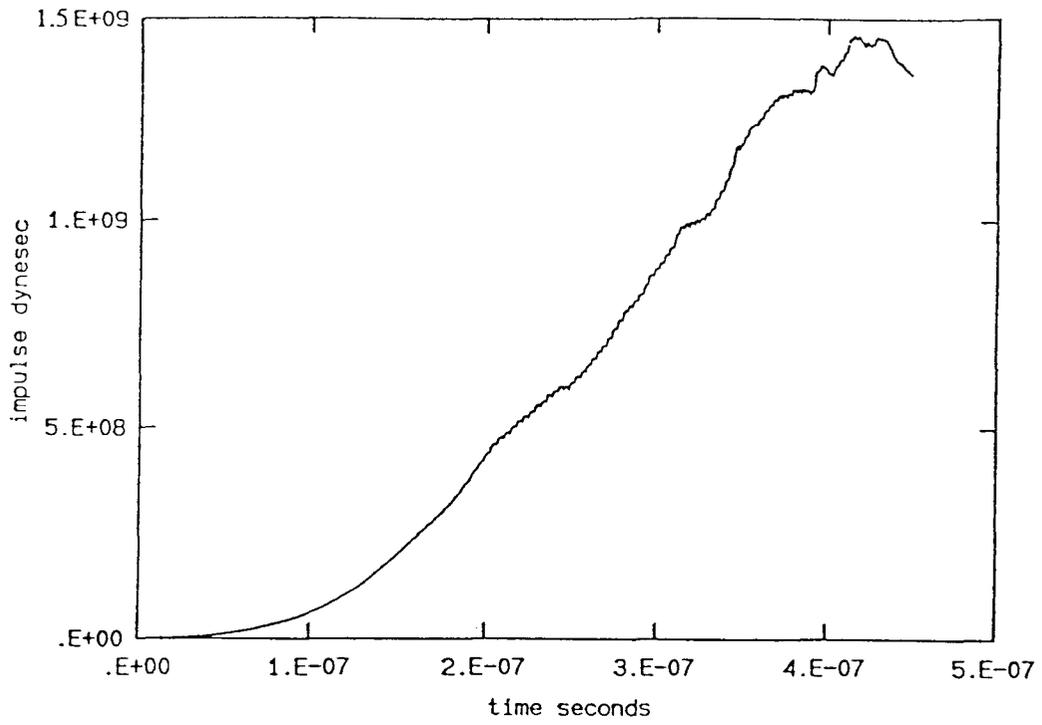


Figure 7.30a

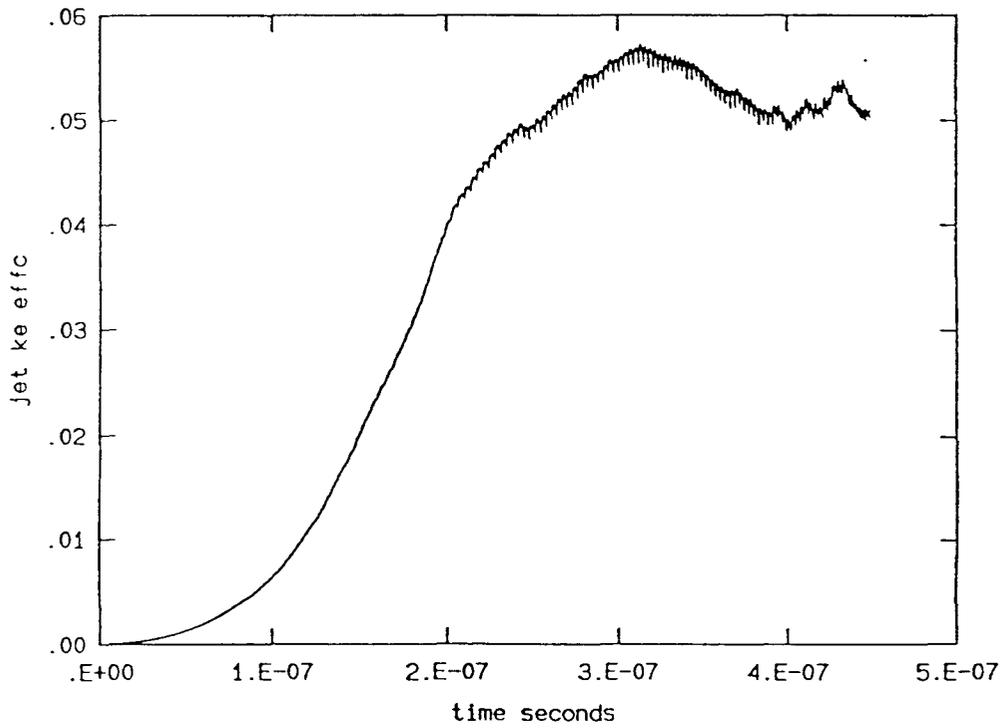


Figure 7.30b

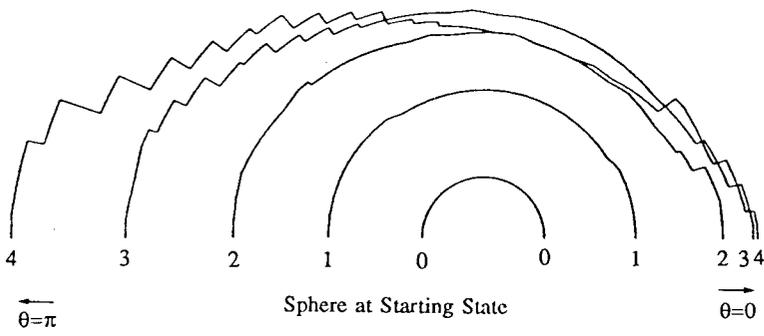
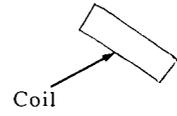


Figure 7.31 Locations of the interface at evenly spaced time intervals for Case 9.
 ($ER \approx 0.6$, $ER_i \approx 24$, $PR \approx 70$)

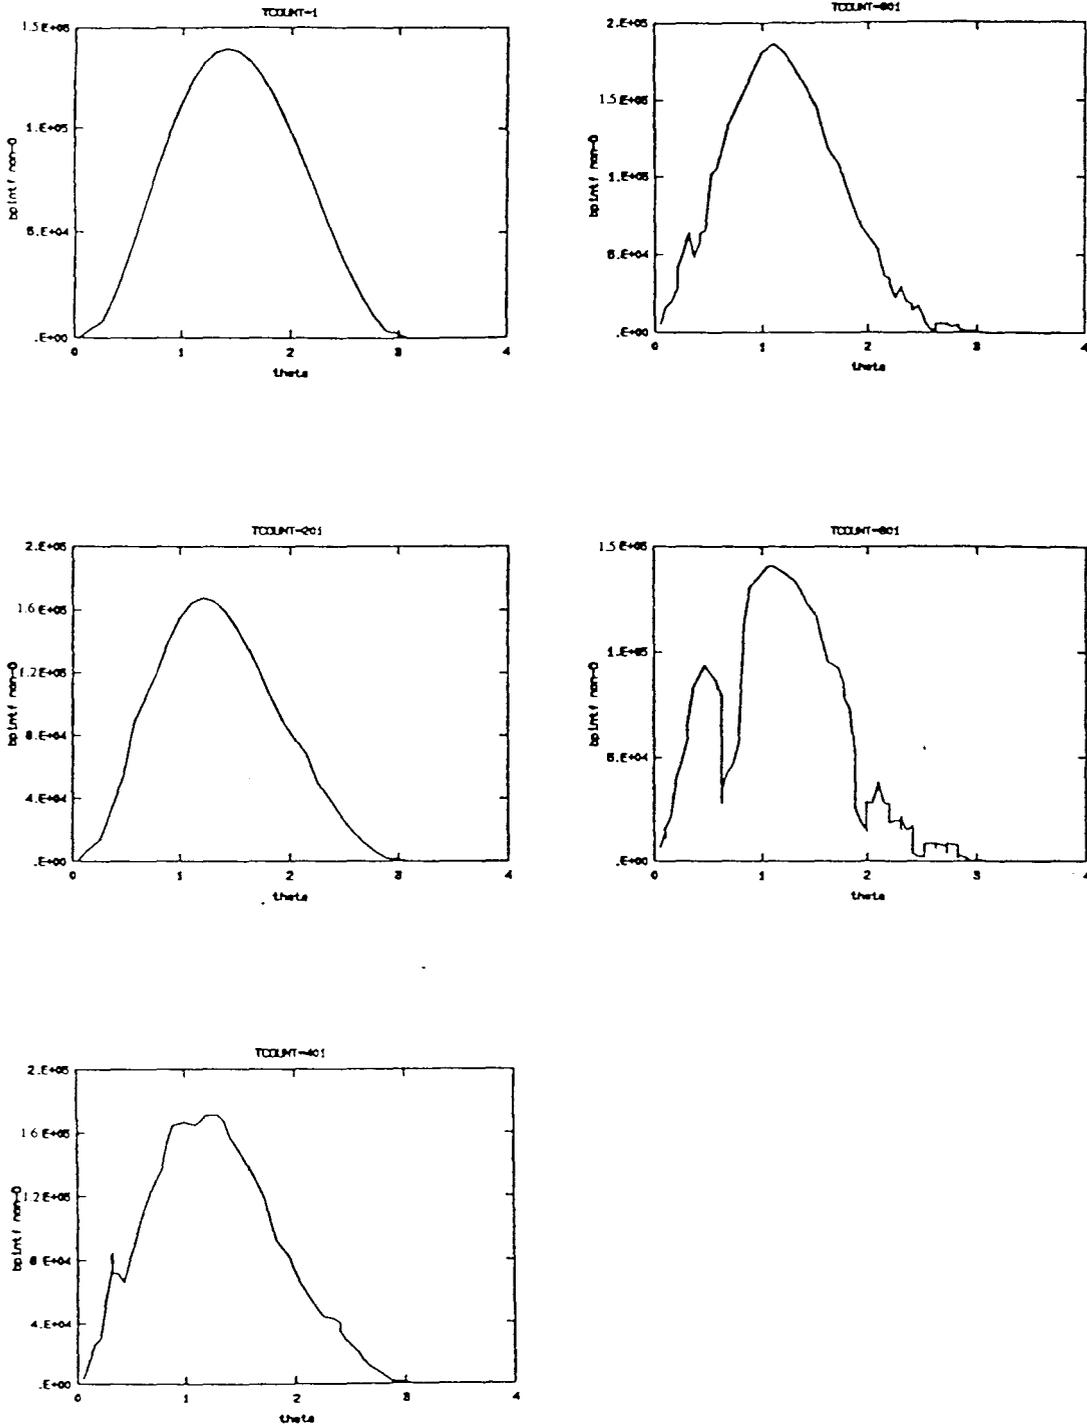


Figure 7.32

Nondimensionalized interfacial magnetic pressure (p_B') (vertical axis) v.s. polar angle (θ) in radians (horizontal axis).

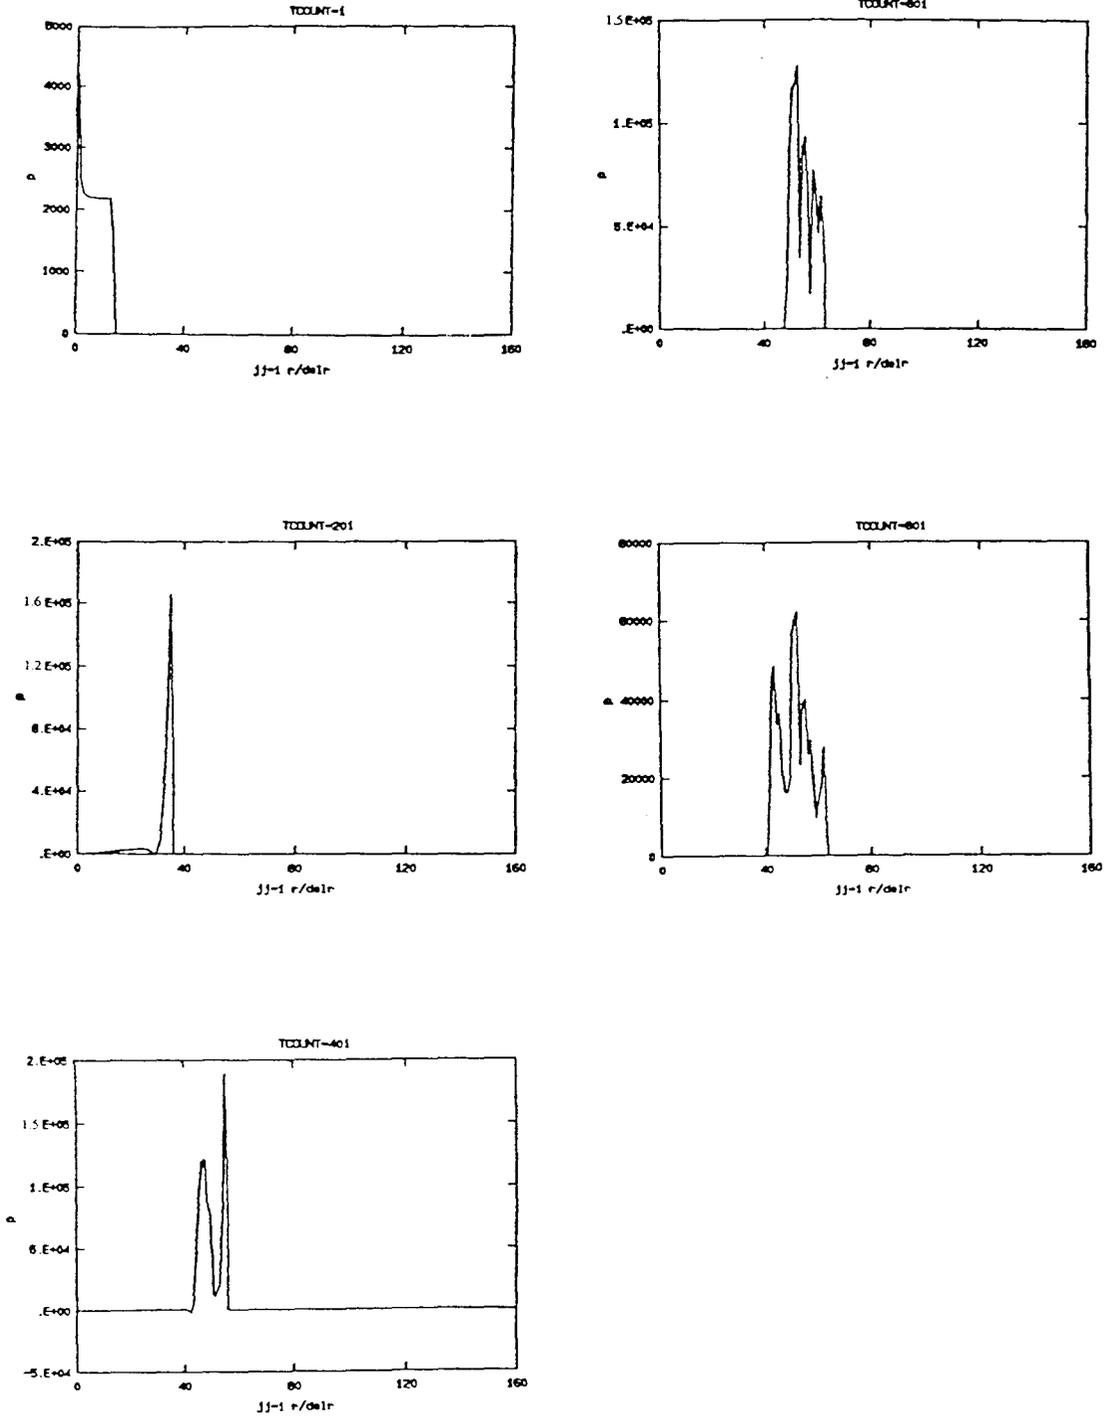


Figure 7.33a

Nondimensional pressure (p') (vertical axis)
 v.s. $r / \Delta r$ (horizontal axis) for $\theta = 0.5\Delta\theta$.

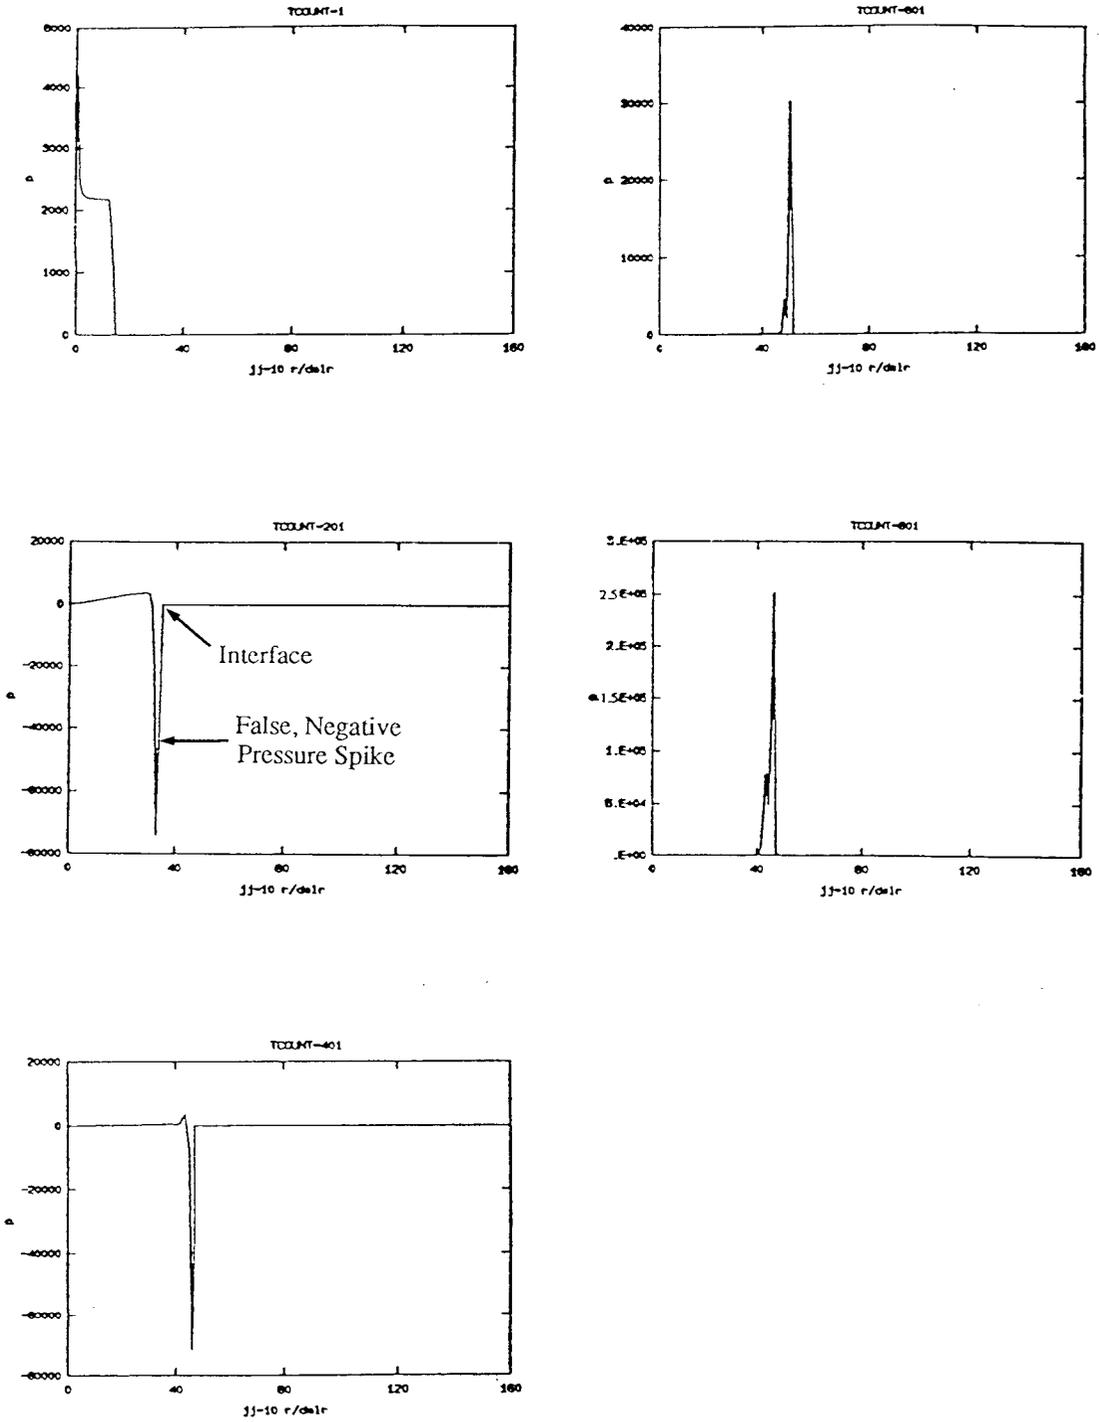


Figure 7.33b

Nondimensionalized pressure (p') (vertical axis)
 v.s. $r/\Delta r$ (horizontal axis) for $\theta = 9.5\Delta\theta$.

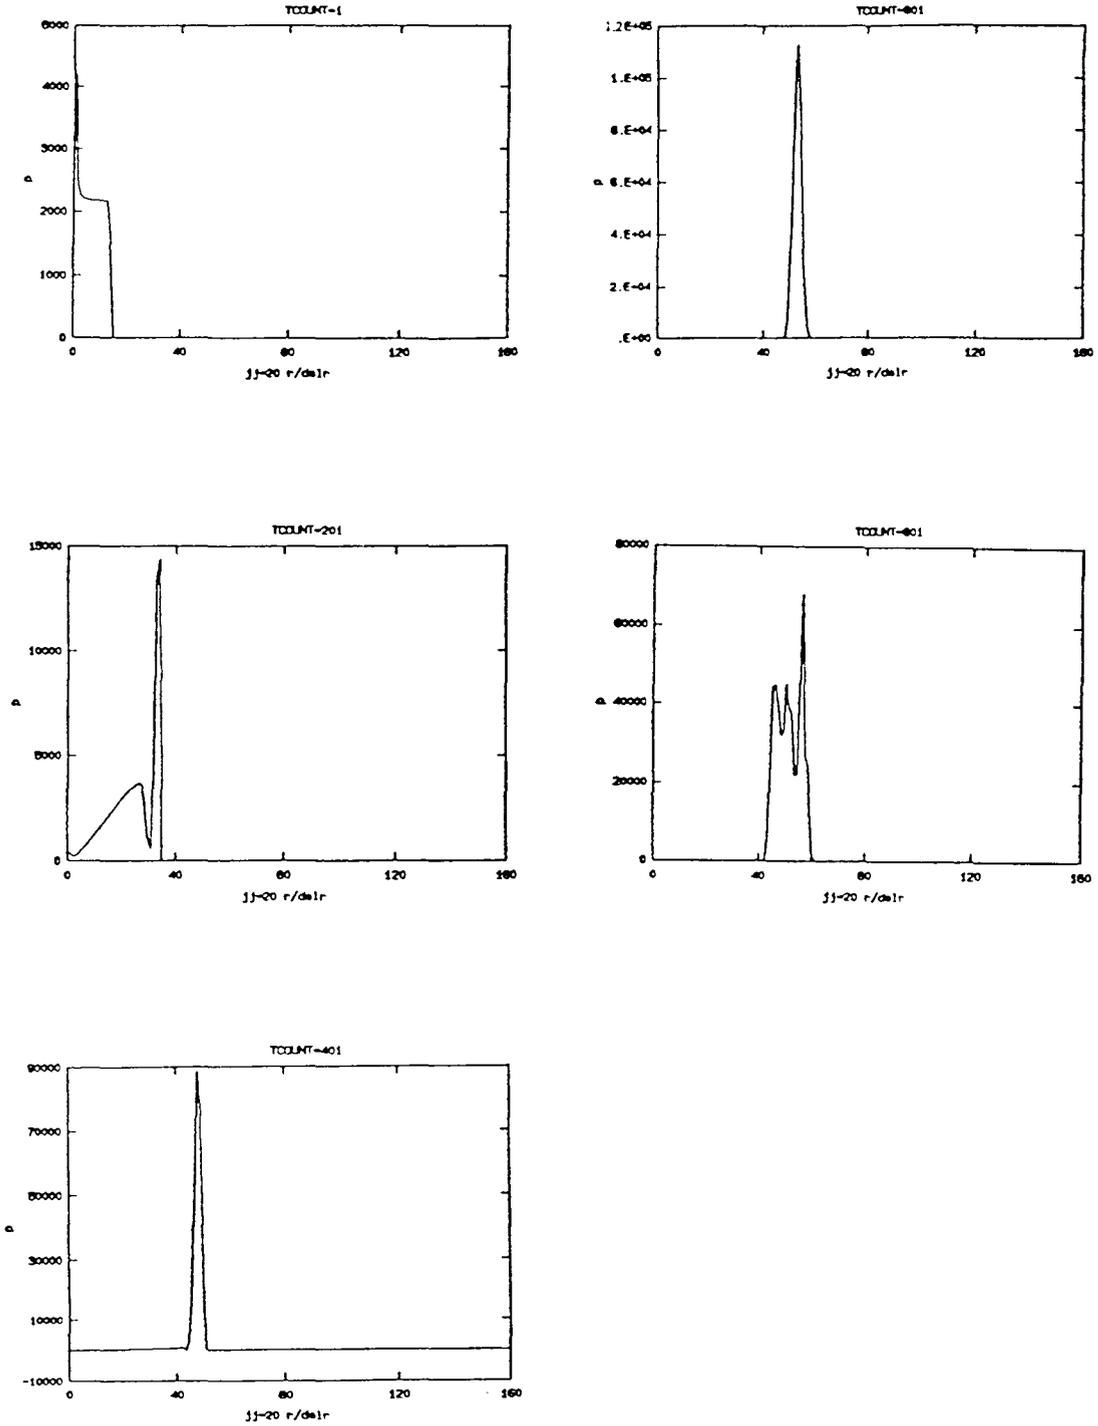


Figure 7.33c

Nondimensionalized pressure (p') (vertical axis)
 v.s. $r / \Delta r$ (horizontal axis) for $\theta = 19.5\Delta\theta$.

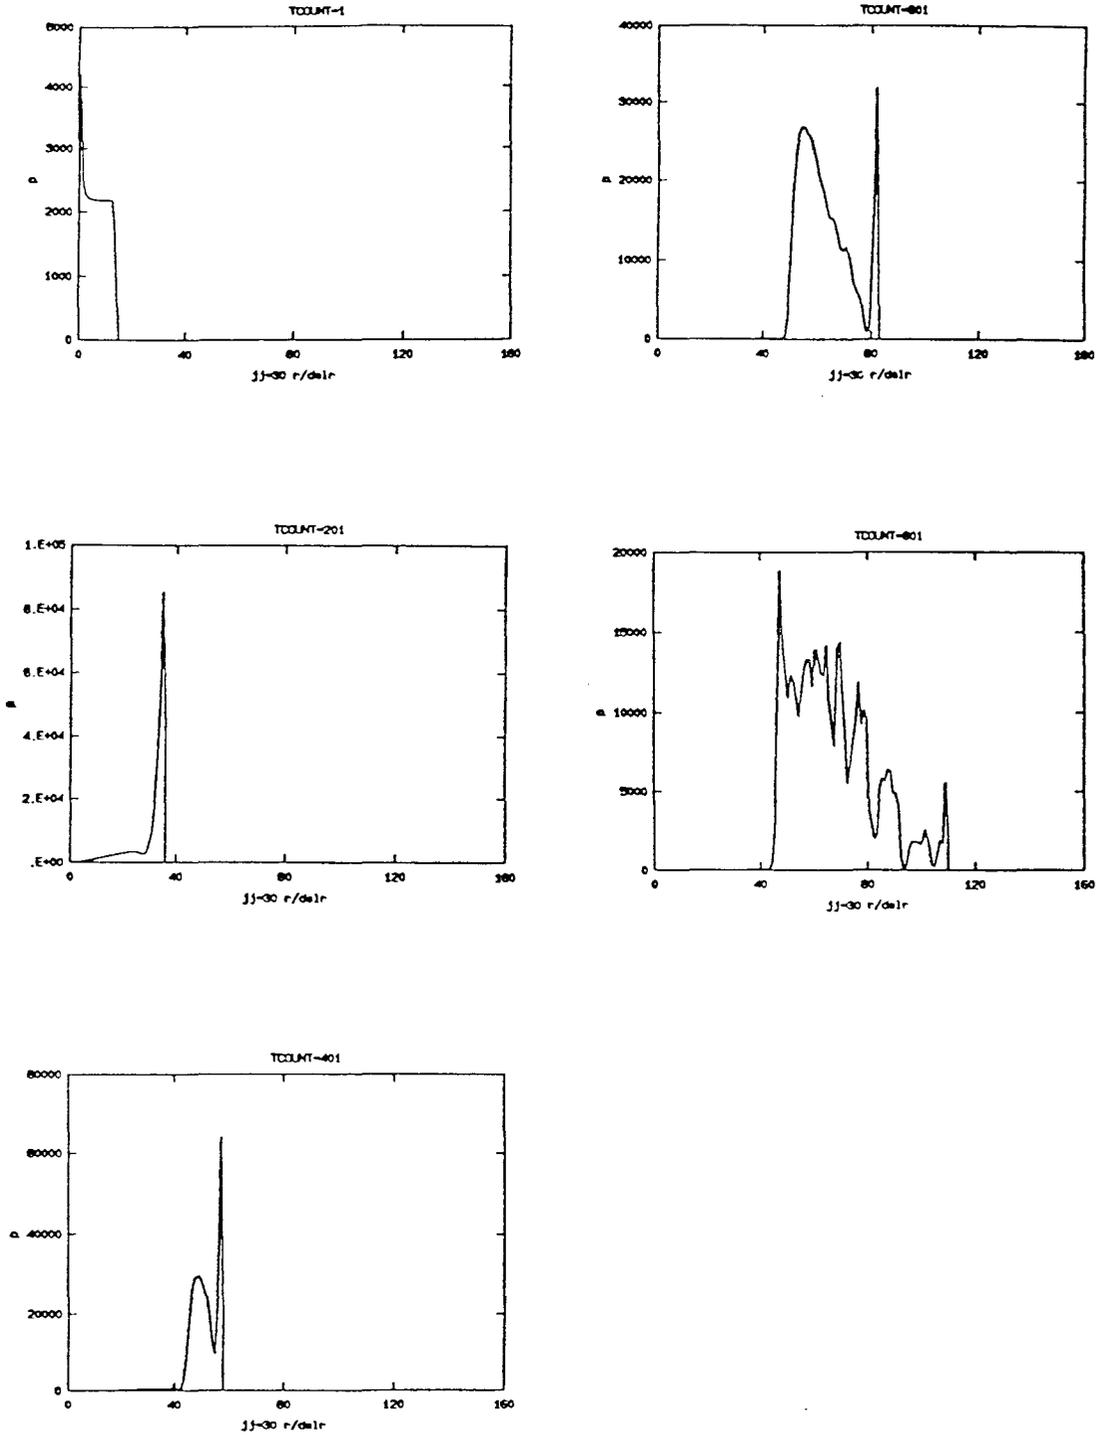


Figure 7.33d

Nondimensionalized pressure (p') (vertical axis)
 v.s. $r/\Delta r$ (horizontal axis) for $\theta = 29.5\Delta\theta$.

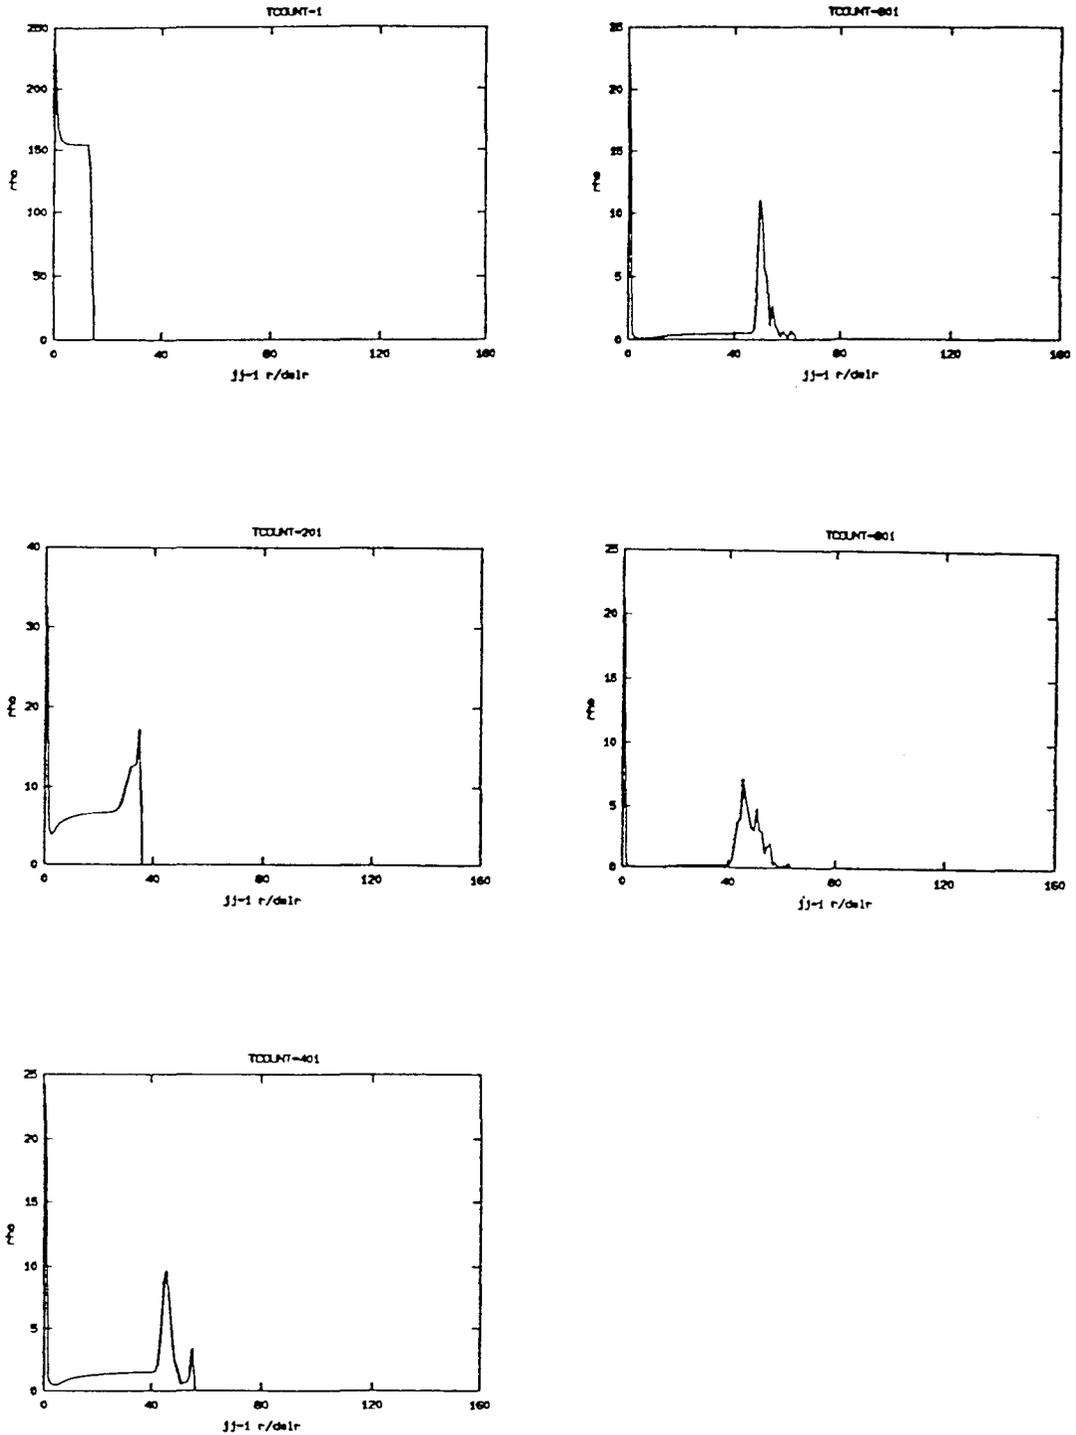


Figure 7.34a. Nondimensionalized density (ρ') (vertical axis) v.s. $r/\Delta r$ (horizontal axis) for $\theta = 0.5\Delta\theta$.

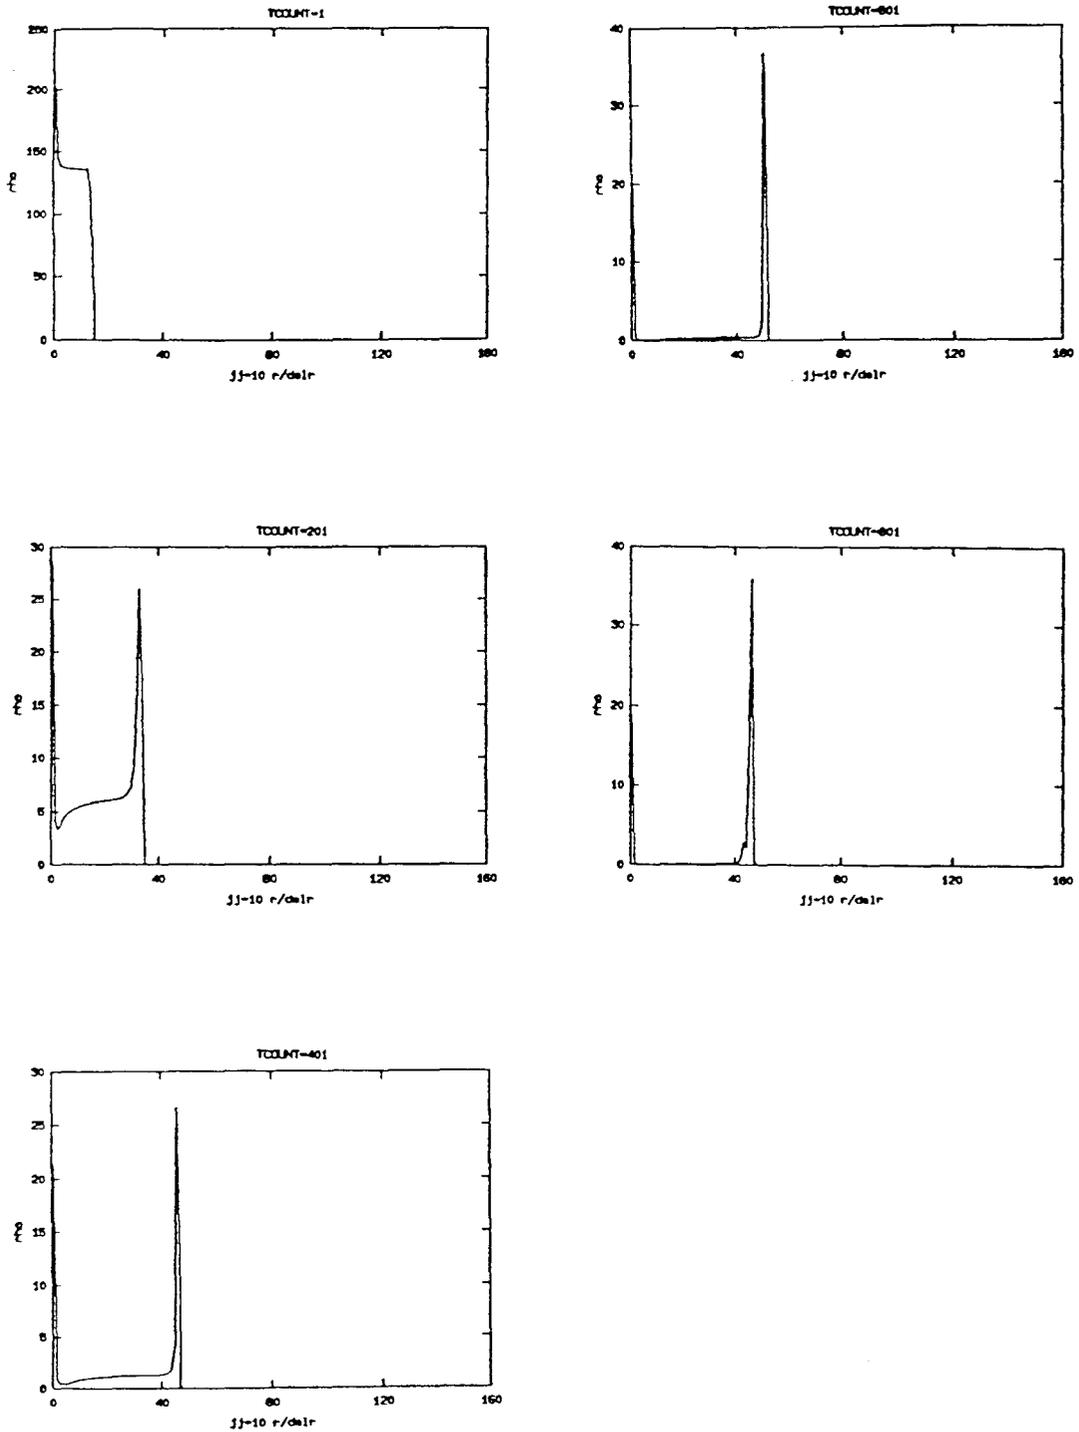


Figure 7.34b

Nondimensional density (ρ') (vertical axis)
 v.s. $r/\Delta r$ (horizontal axis) for $\theta = 9.5\Delta\theta$.

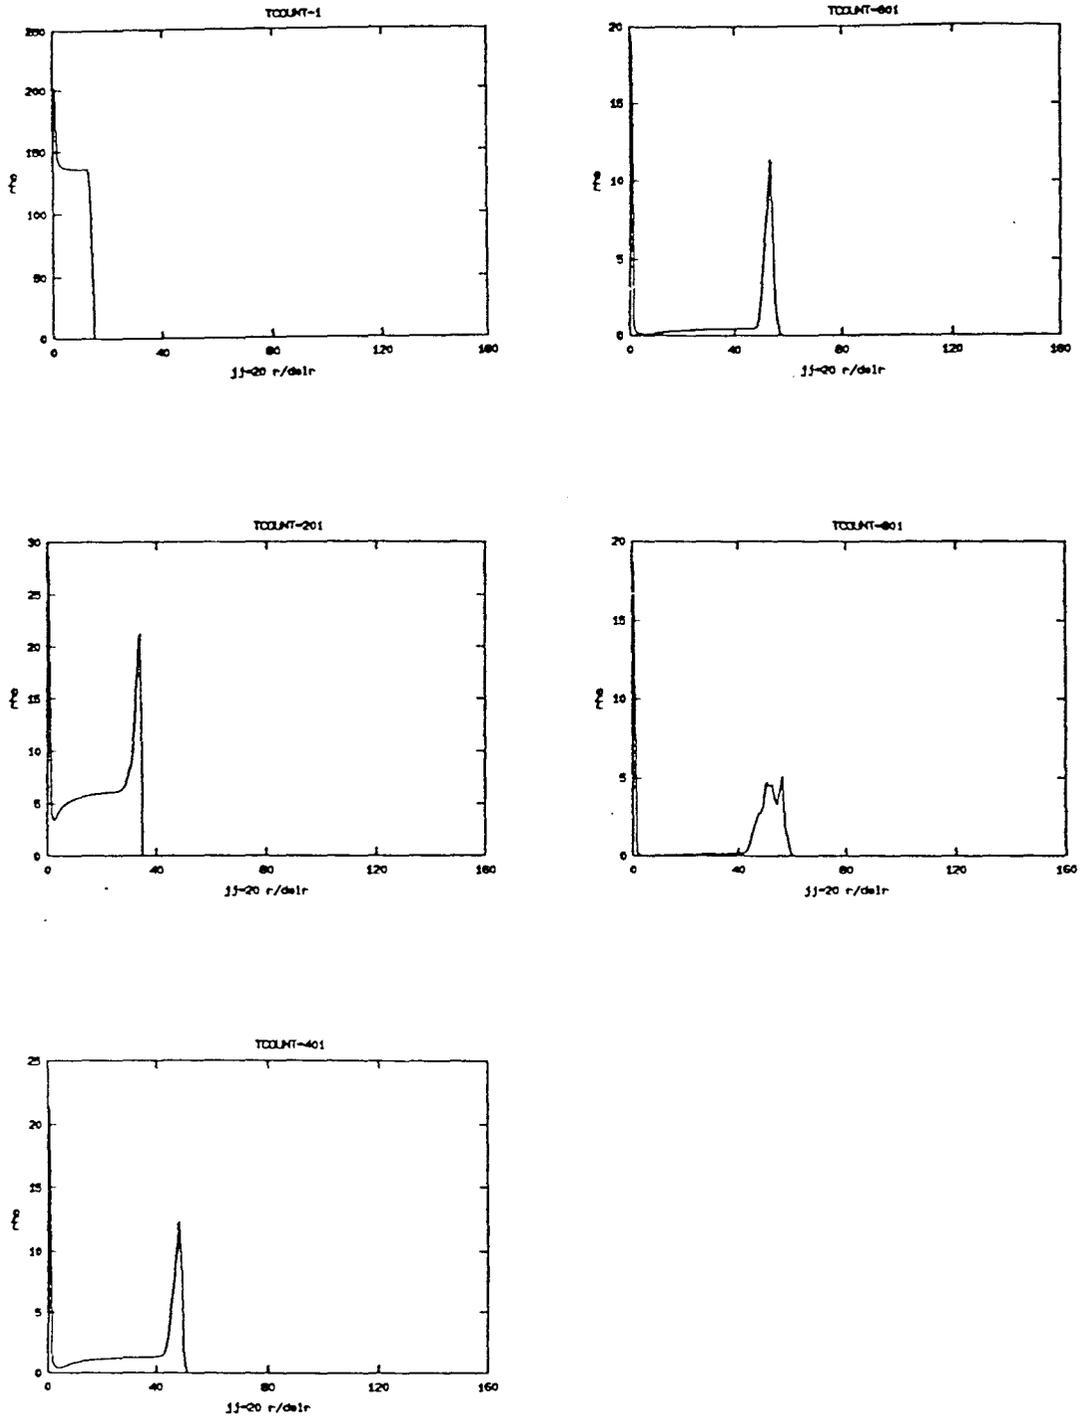


Figure 7.34c

Nondimensionalized density (ρ') (vertical axis)
 v.s. $r / \Delta r$ (horizontal axis) for $\theta = 19.5\Delta\theta$.

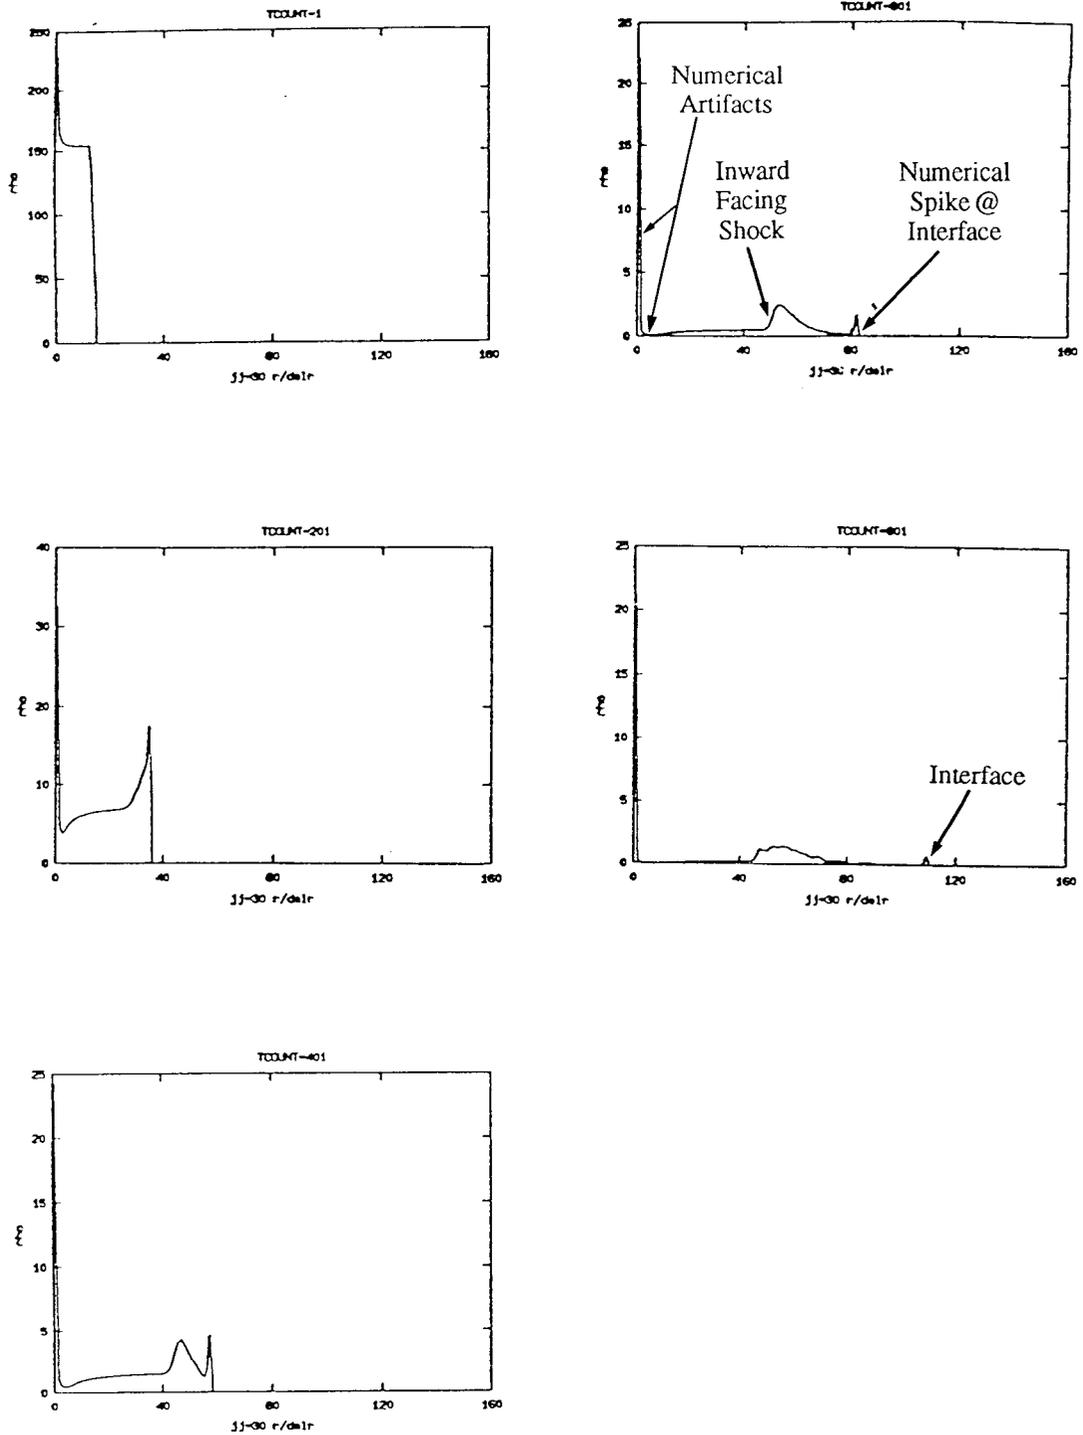


Figure 7.34d

Nondimensionalized density (ρ') (vertical axis)
 v.s. $r / \Delta r$ (horizontal axis) for $\theta = 29.5\Delta\theta$.

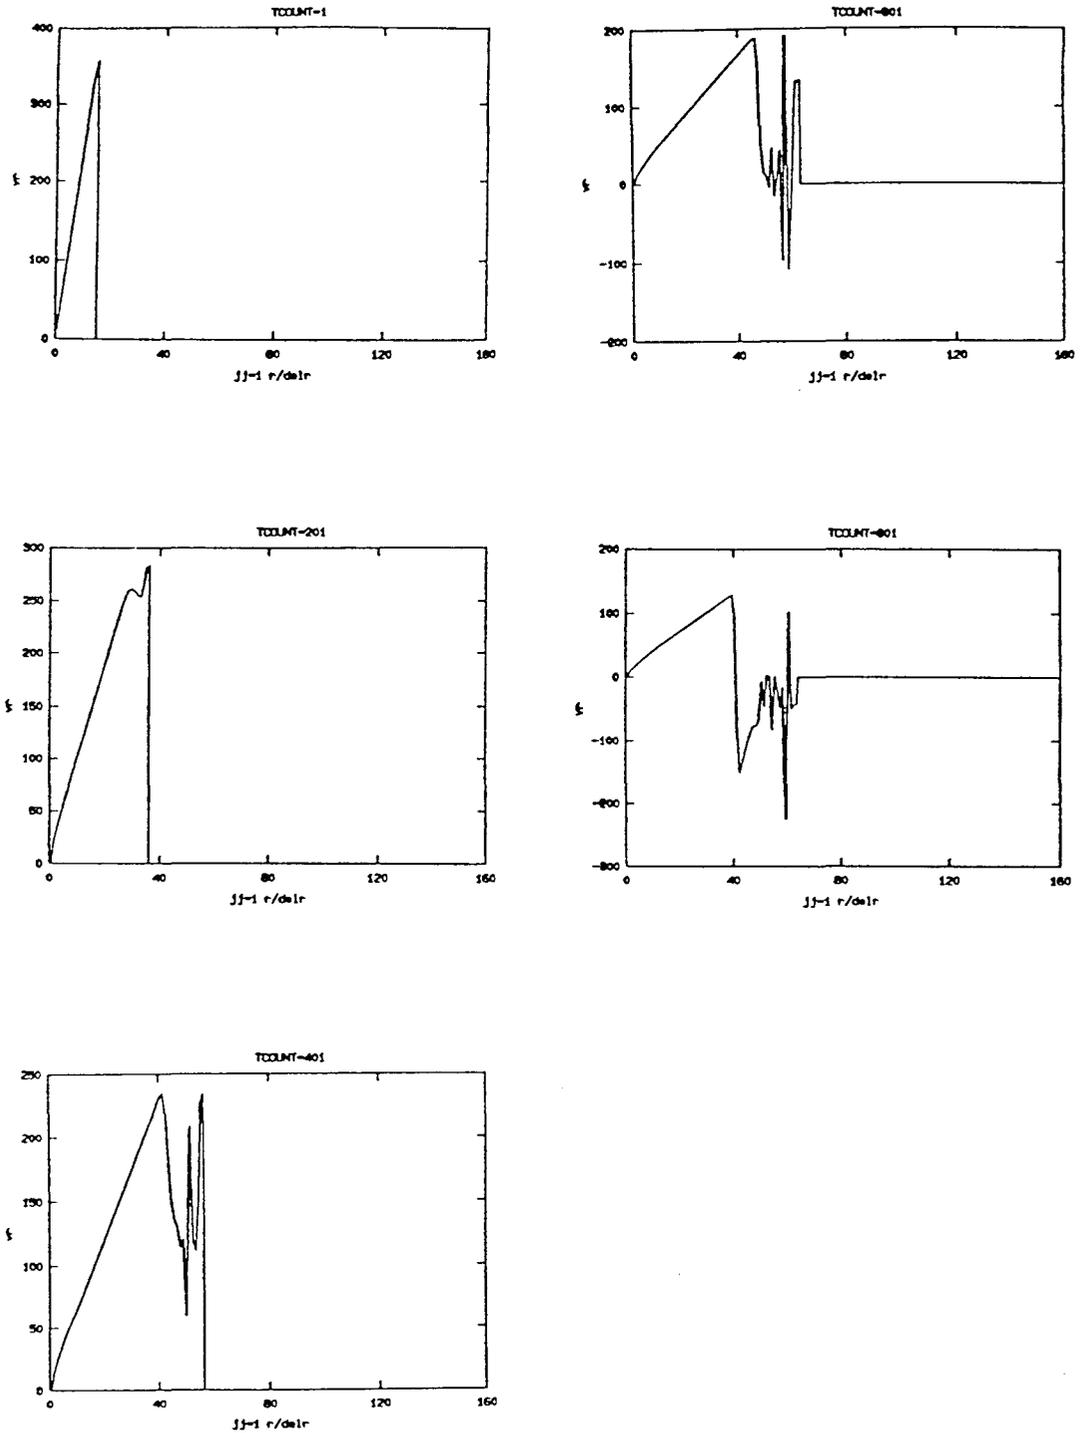


Figure 7.35a Nondimensional radial velocity (v_r') (vertical axis)
 v.s. $r/\Delta r$ (horizontal axis) for $\theta = 0.5\Delta\theta$.

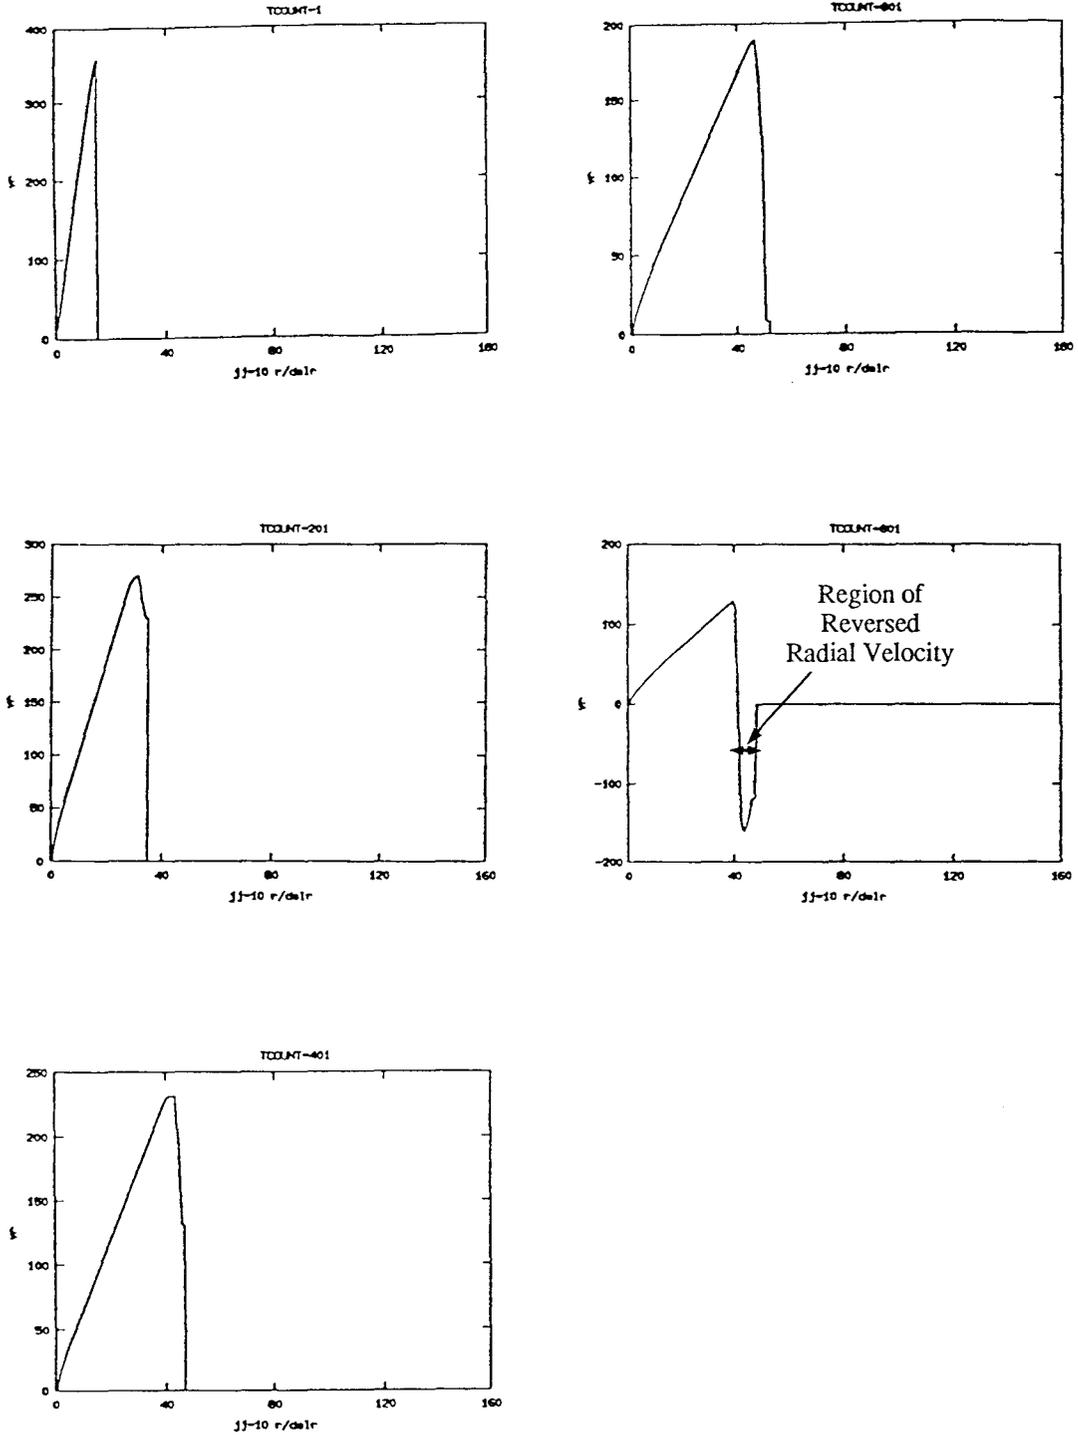


Figure 7.35b Nondimensionalized radial velocity (v_r') (vertical axis) v.s. $r/\Delta r$ (horizontal axis) for $\theta = 9.5\Delta\theta$.

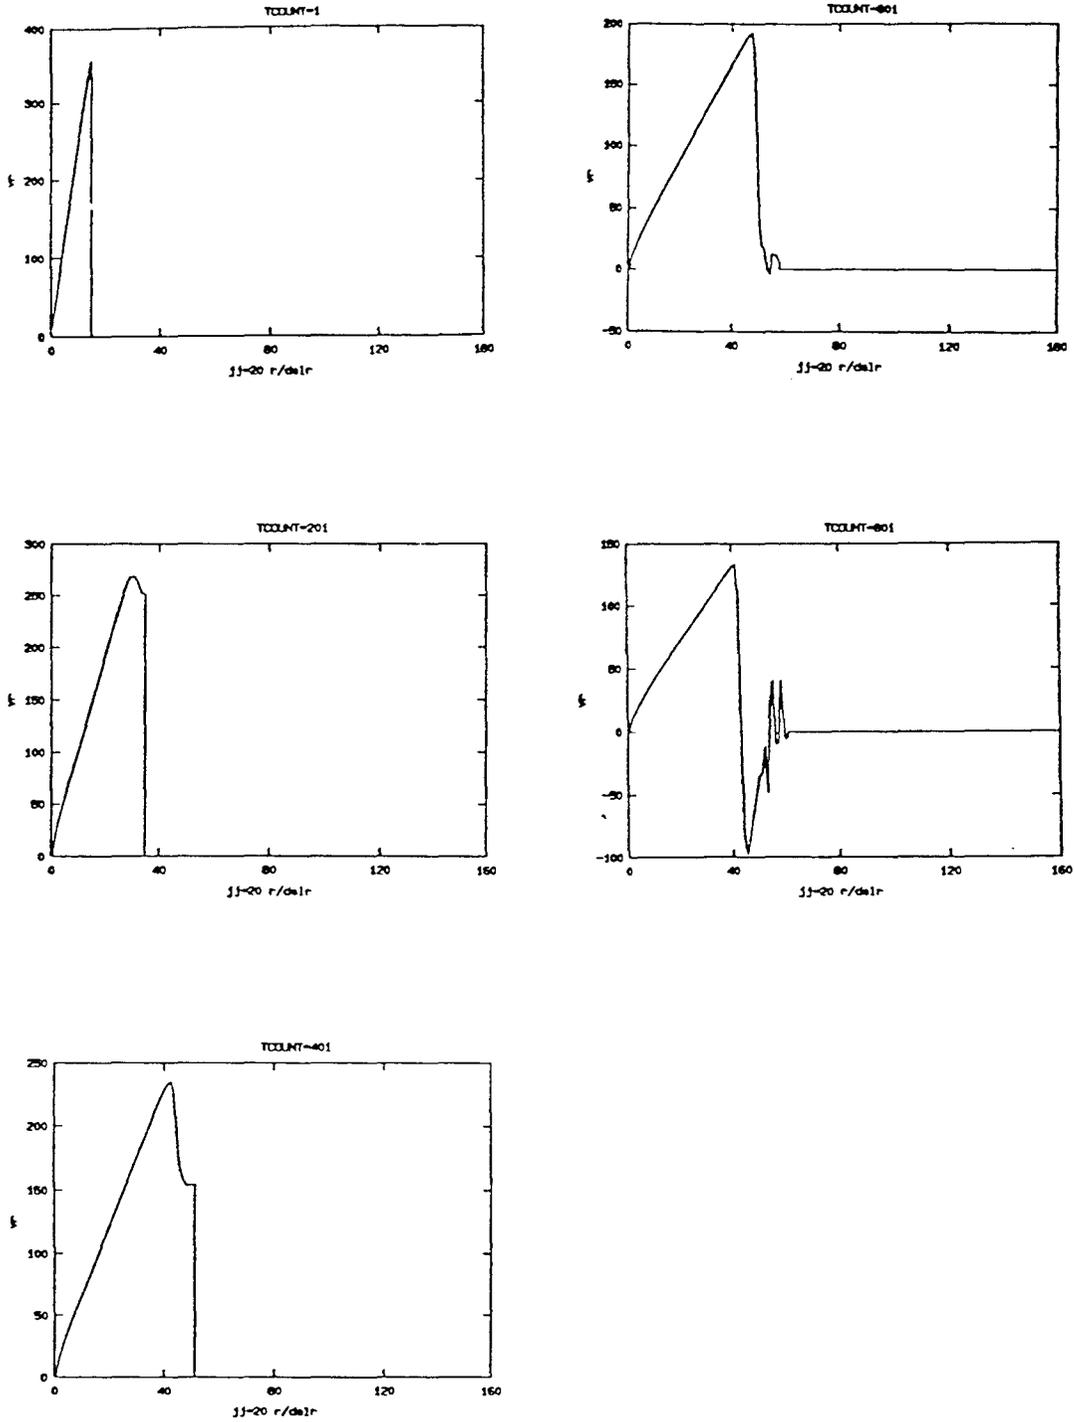


Figure 7.35c Nondimensionalized radial velocity (v_r') (vertical axis) v.s. $r / \Delta r$ (horizontal axis) for $\theta = 19.5\Delta\theta$.

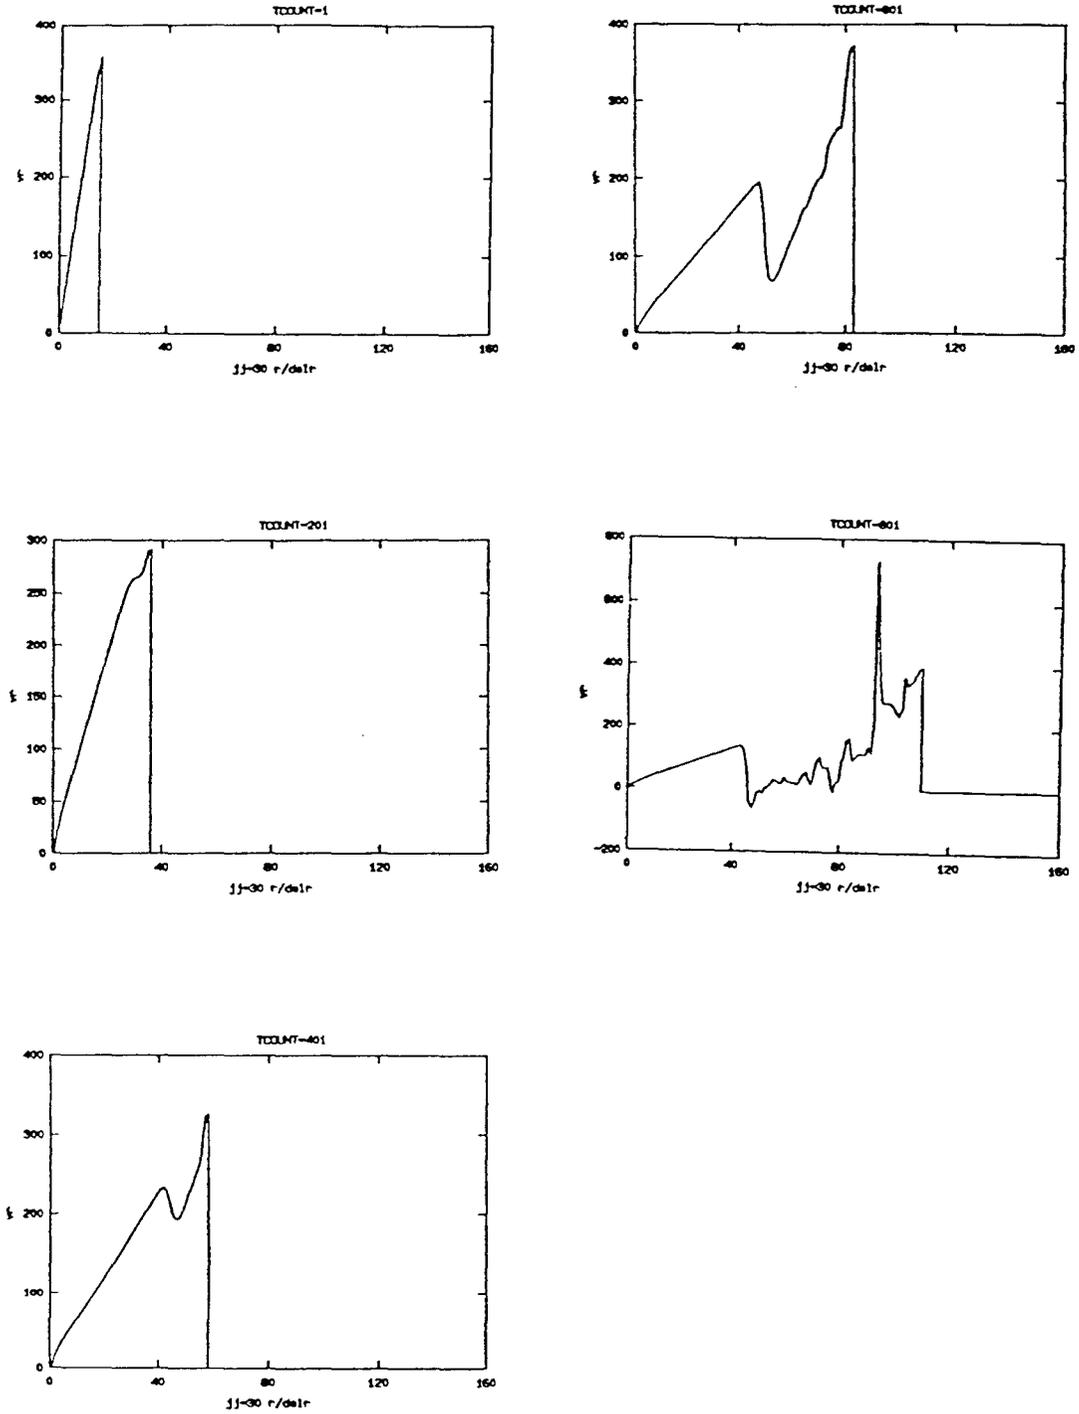


Figure 7.35d Nondimensional radial velocity (v_r') (vertical axis) v.s. $r / \Delta r$ (horizontal axis) for $\theta = 29.5\Delta\theta$.

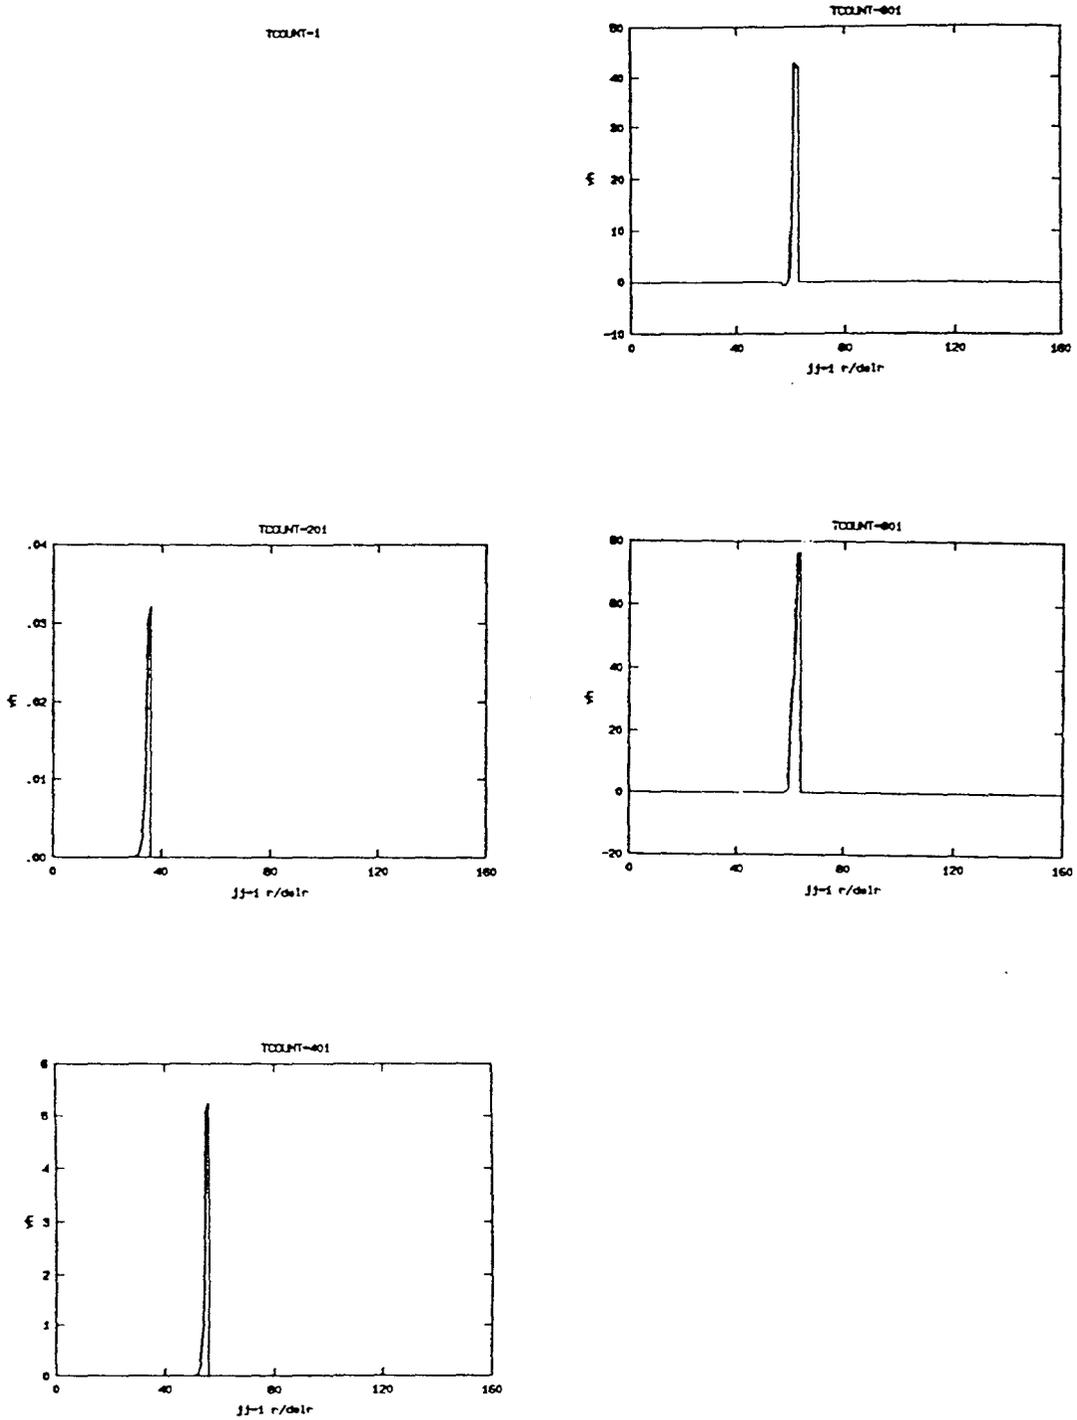


Figure 7.36a Nondimensionalized polar velocity (v_θ') (vertical axis) v.s. $r / \Delta r$ (horizontal axis) for $\theta = 0.5\Delta\theta$.

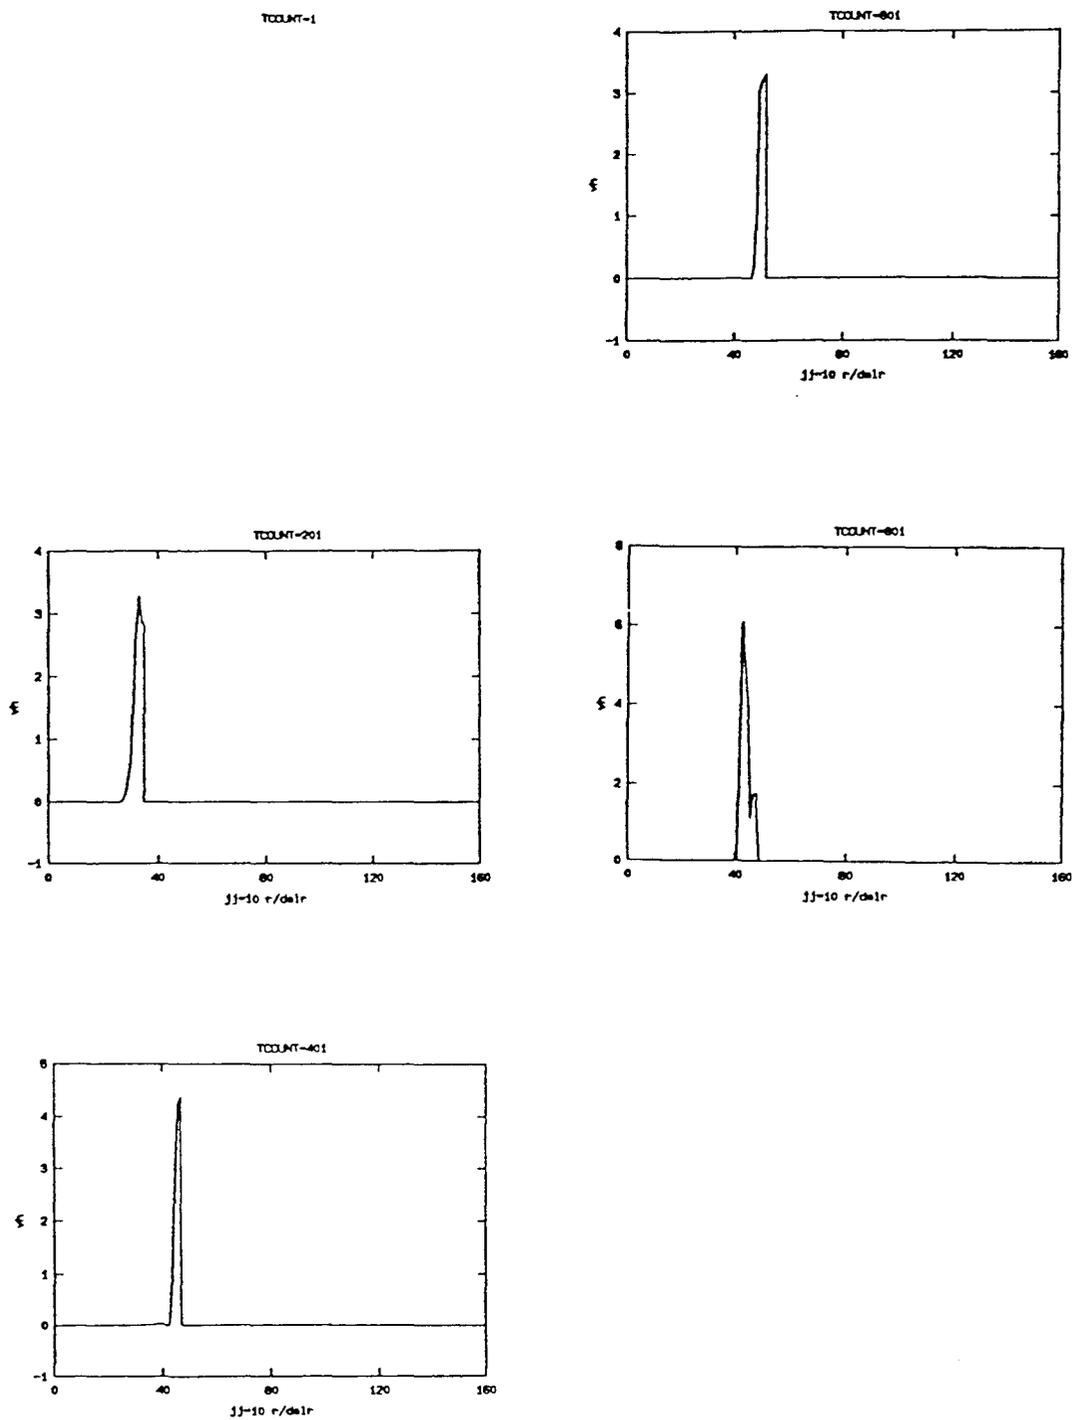


Figure 7.36b Nondimensionalized polar velocity (v_θ') (vertical axis) v.s. $r/\Delta r$ (horizontal axis) for $\theta = 9.5\Delta\theta$.

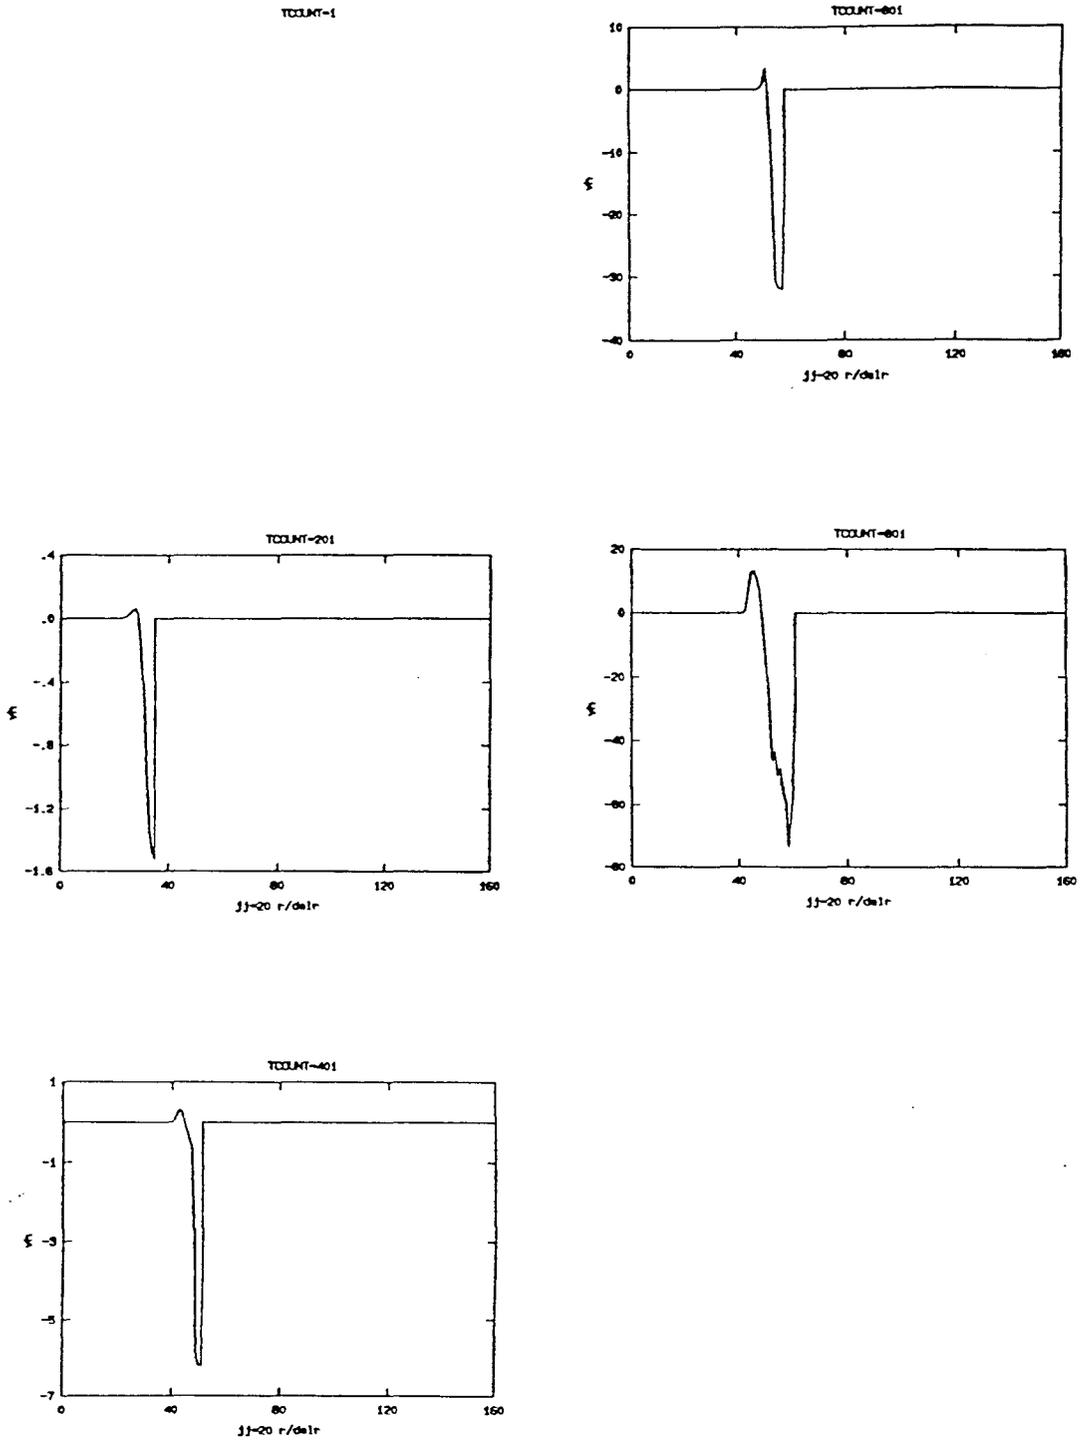


Figure 7.36c Nondimensional polar velocity (v_θ') (vertical axis)
 v.s. $r / \Delta r$ (horizontal axis) for $\theta = 19.5\Delta\theta$.

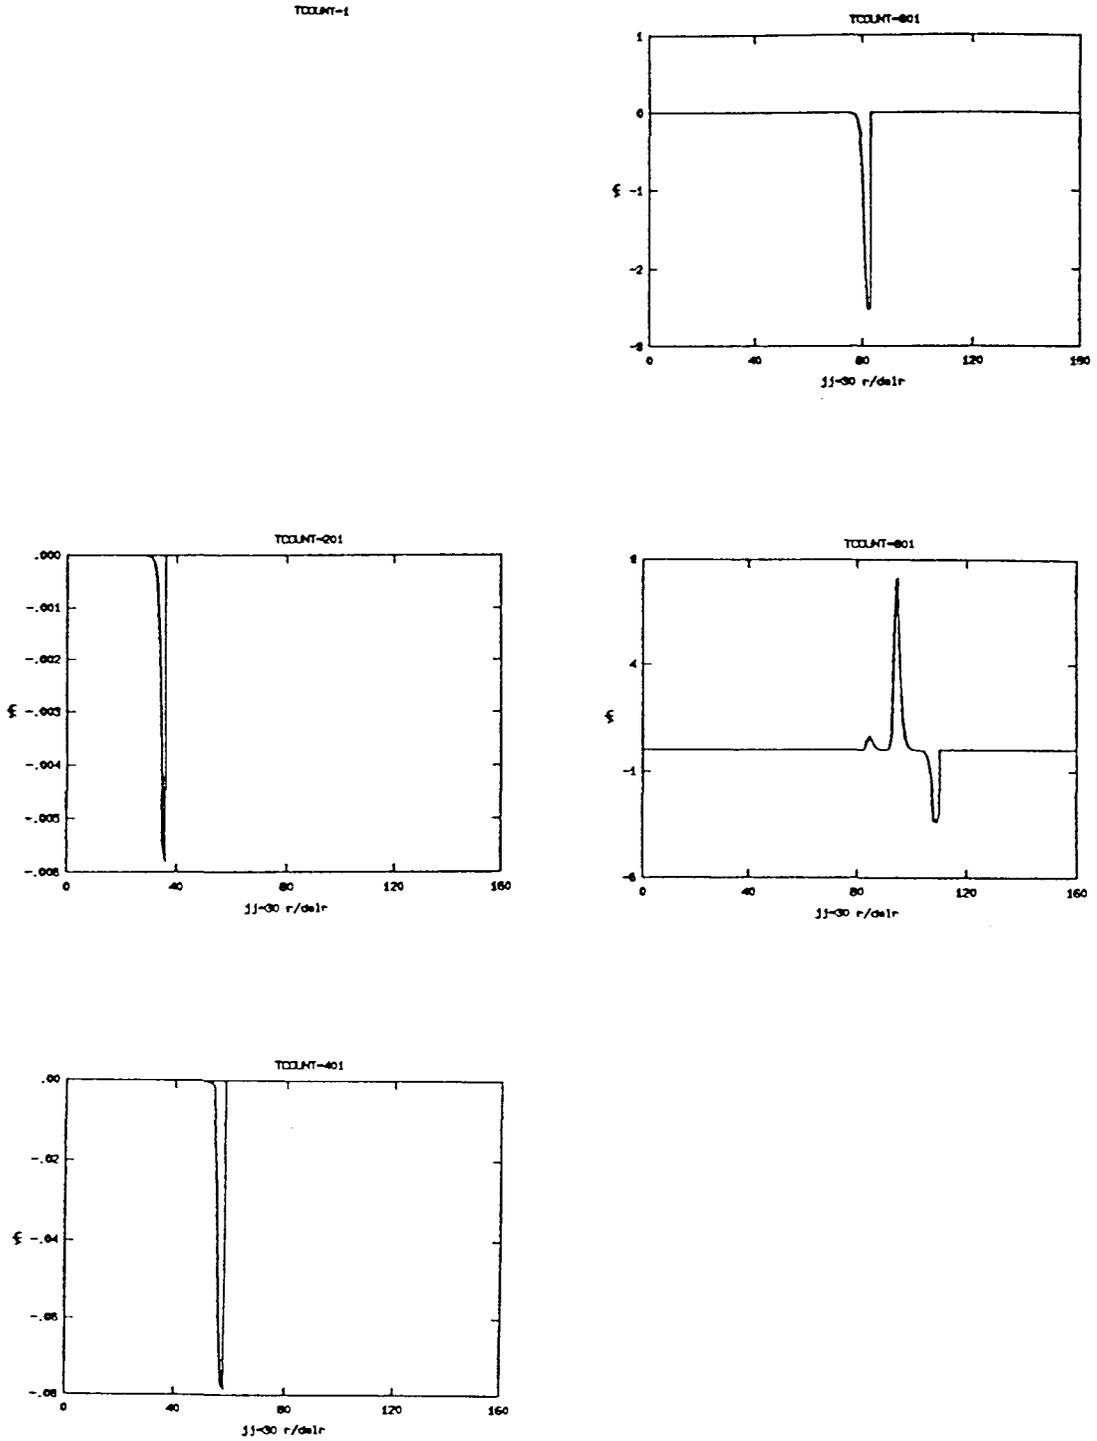


Figure 7.36d Nondimensionalized polar velocity (v_θ') (vertical axis) v.s. $r/\Delta r$ (horizontal axis) for $\theta = 29.5\Delta\theta$.

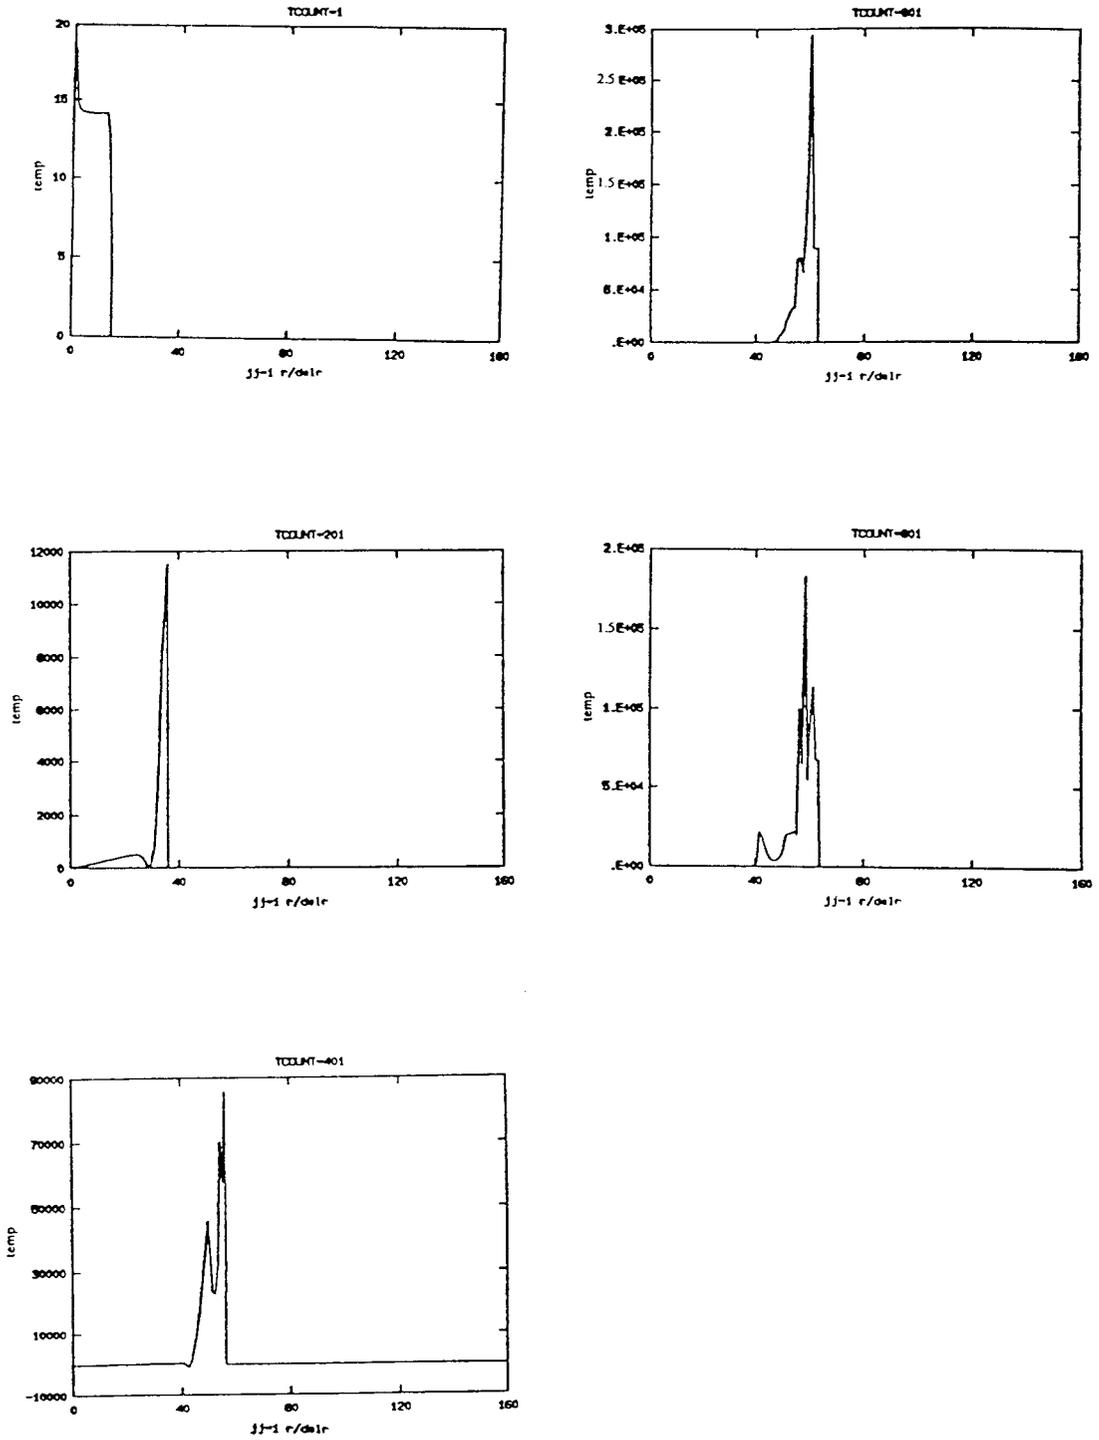


Figure 7.37a

Nondimensionalized temperature (T^*) (vertical axis)
 v.s. $r / \Delta r$ (horizontal axis) for $\theta = 0.5\Delta\theta$.

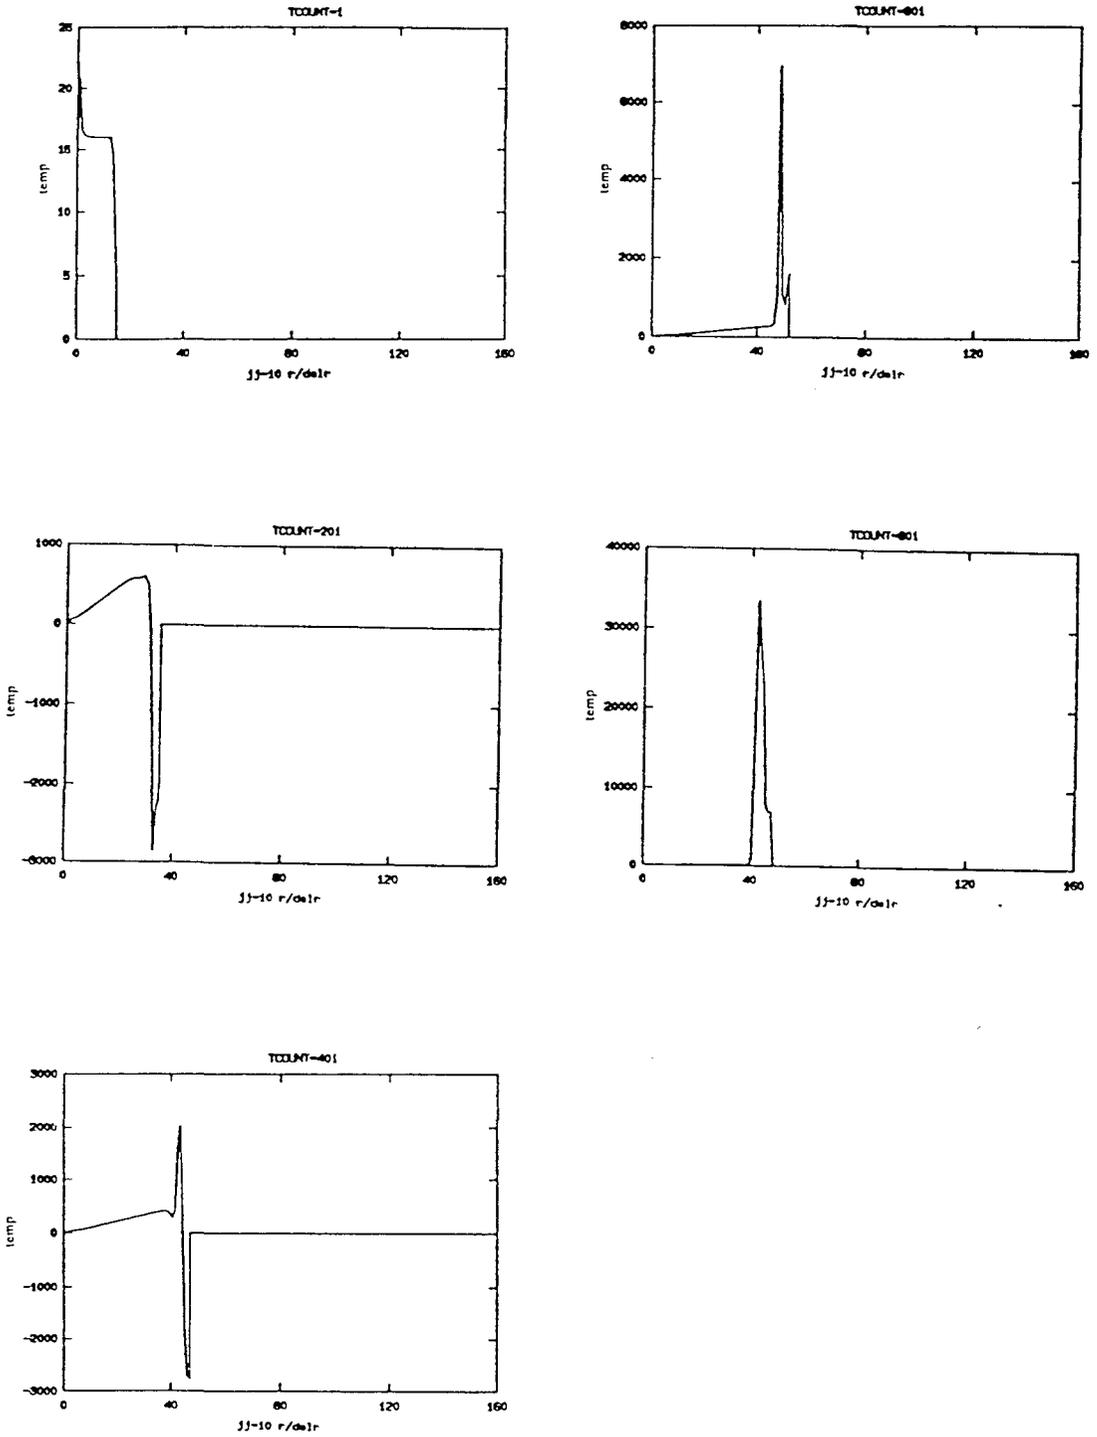


Figure 7.37b Nondimensionalized temperature (T') (vertical axis) v.s. $r/\Delta r$ (horizontal axis) for $\theta = 9.5\Delta\theta$.

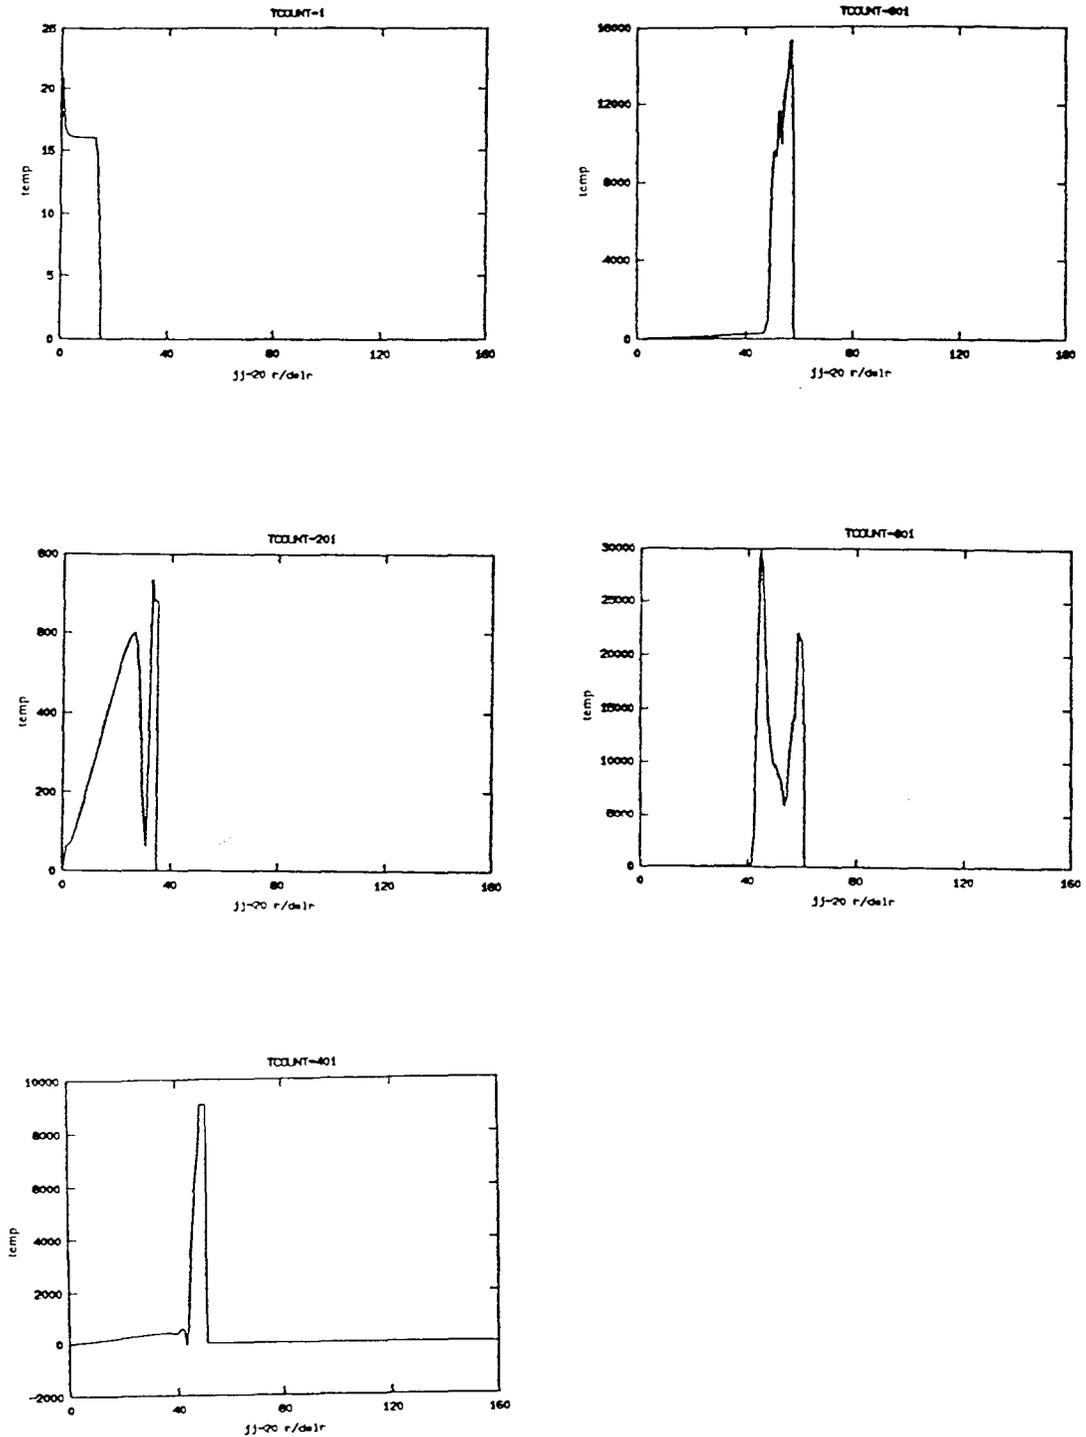


Figure 7.37c Nondimensionalized temperature (T') (vertical axis) v.s. $r / \Delta r$ (horizontal axis) for $\theta = 19.5\Delta\theta$.

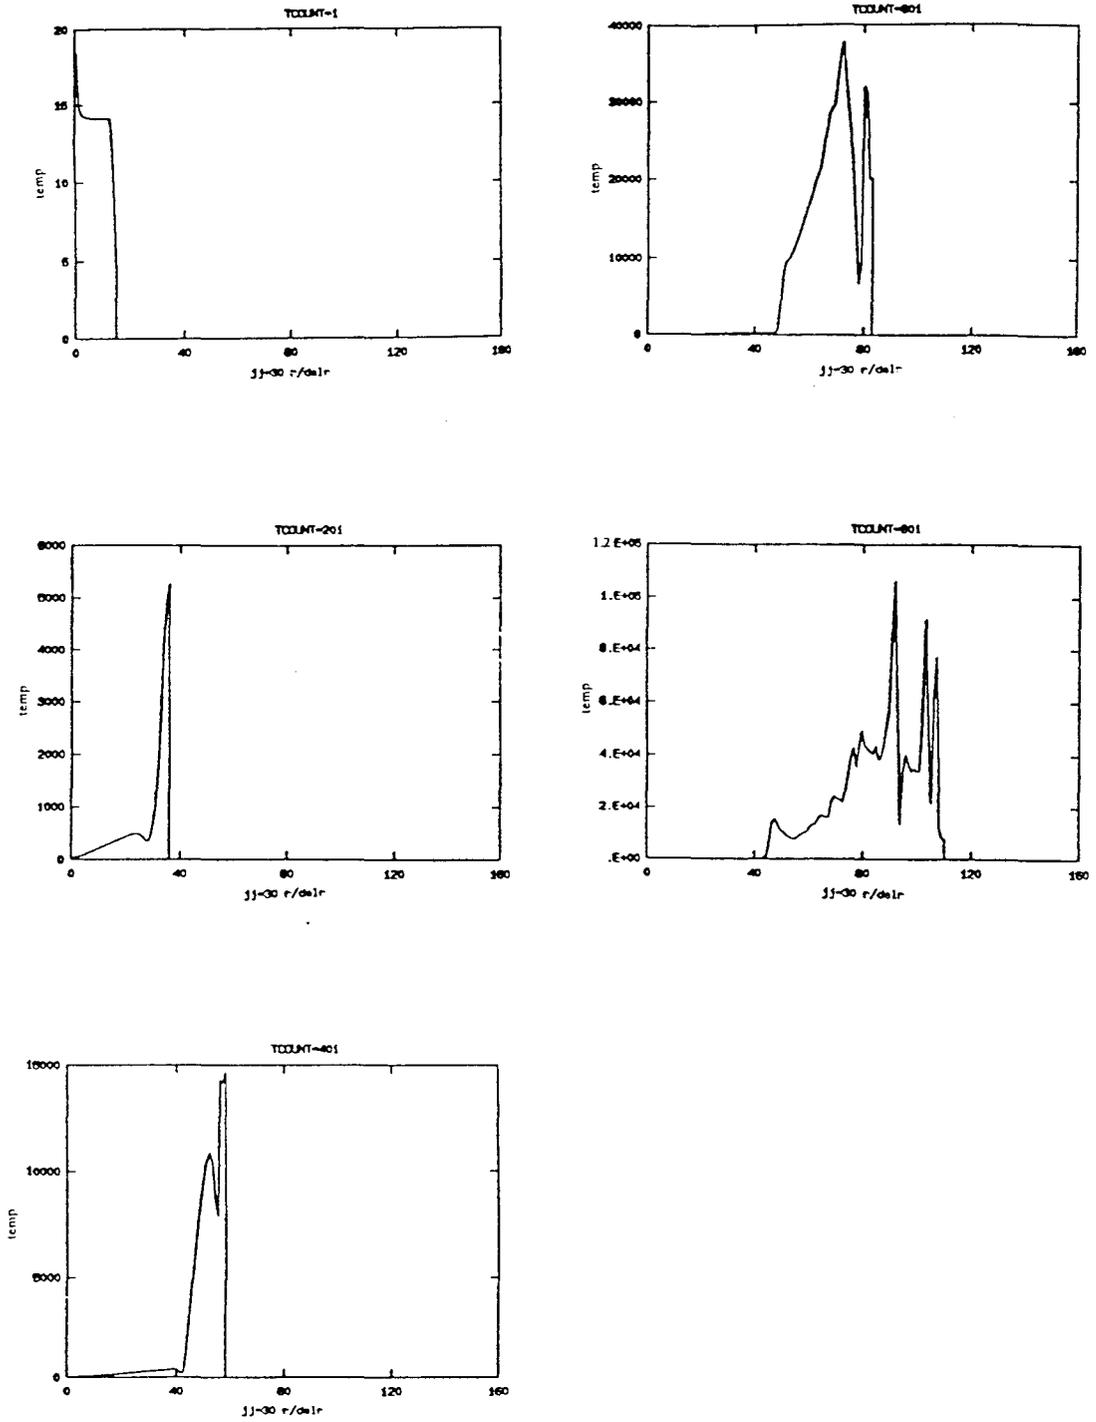


Figure 7.37d Nondimensionalized temperature (T') (vertical axis) v.s. $r / \Delta r$ (horizontal axis) for $\theta = 29.5\Delta\theta$.

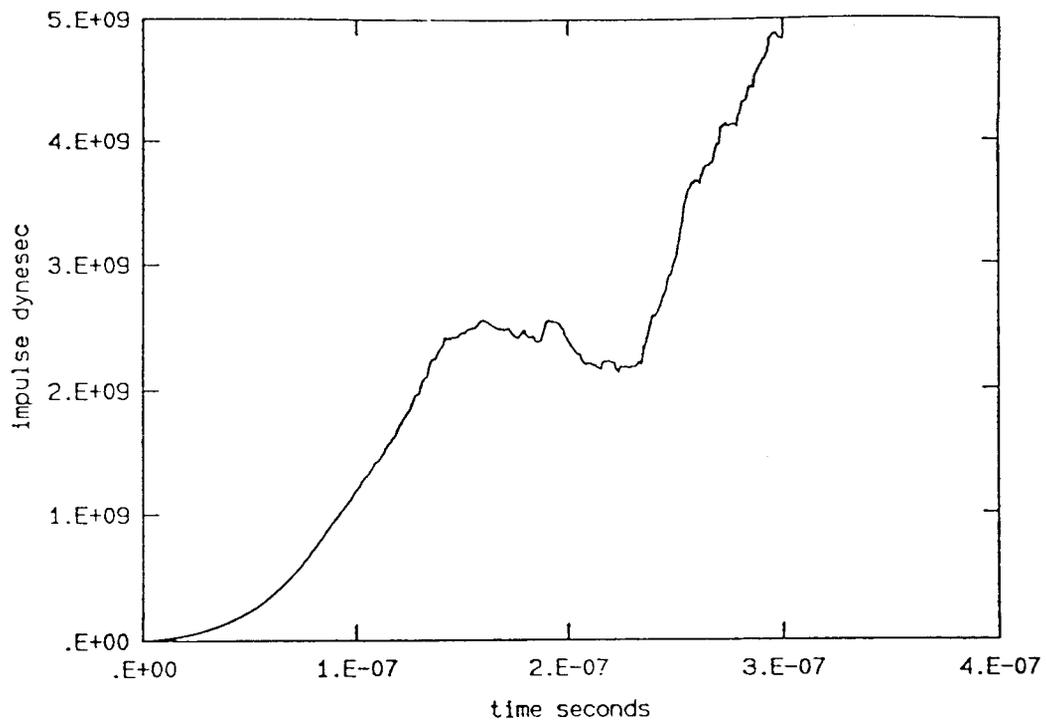


Figure 7.38a

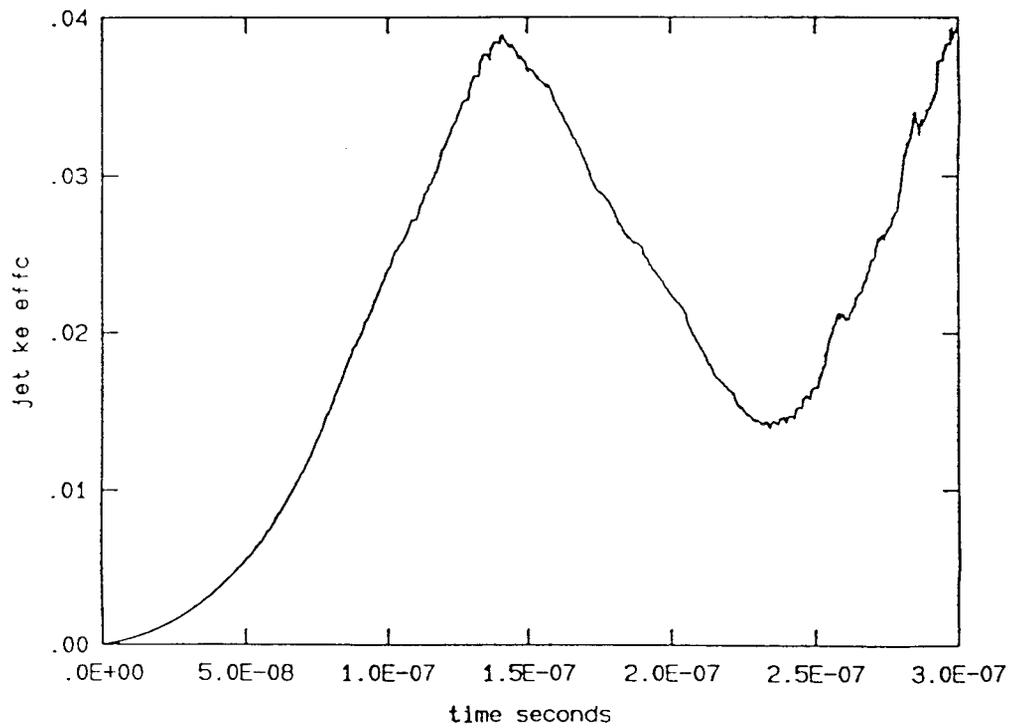


Figure 7.38b

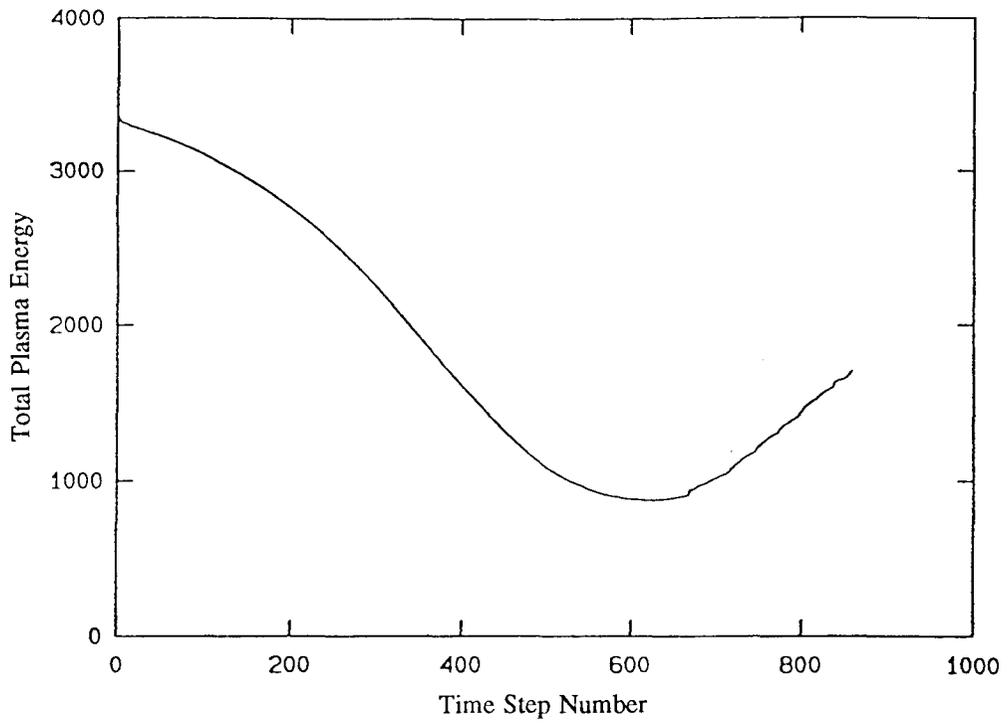


Figure 7.39

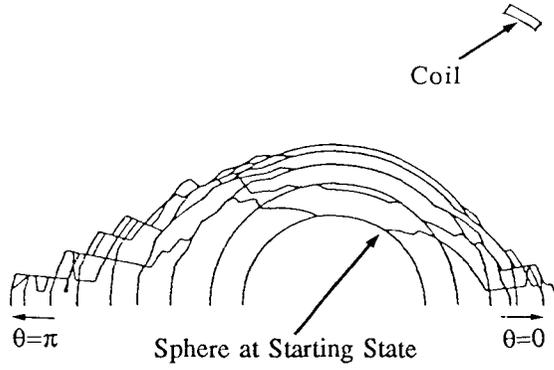


Figure 7.40a Locations of the interface at evenly spaced time intervals for Case 10. ($ER=4.2$, $ER_i \approx 266$, $PR \approx 720$)

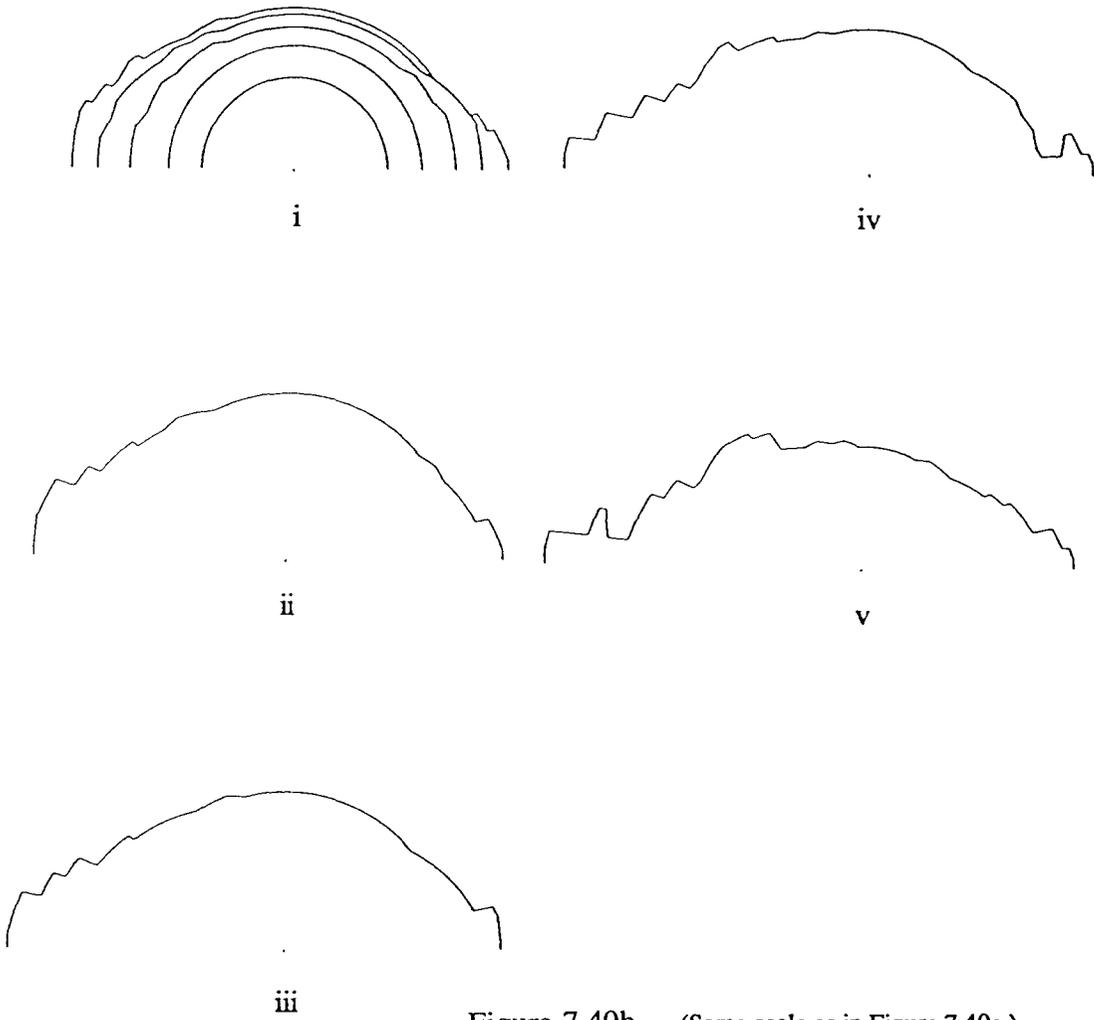


Figure 7.40b (Same scale as in Figure 7.40a.)

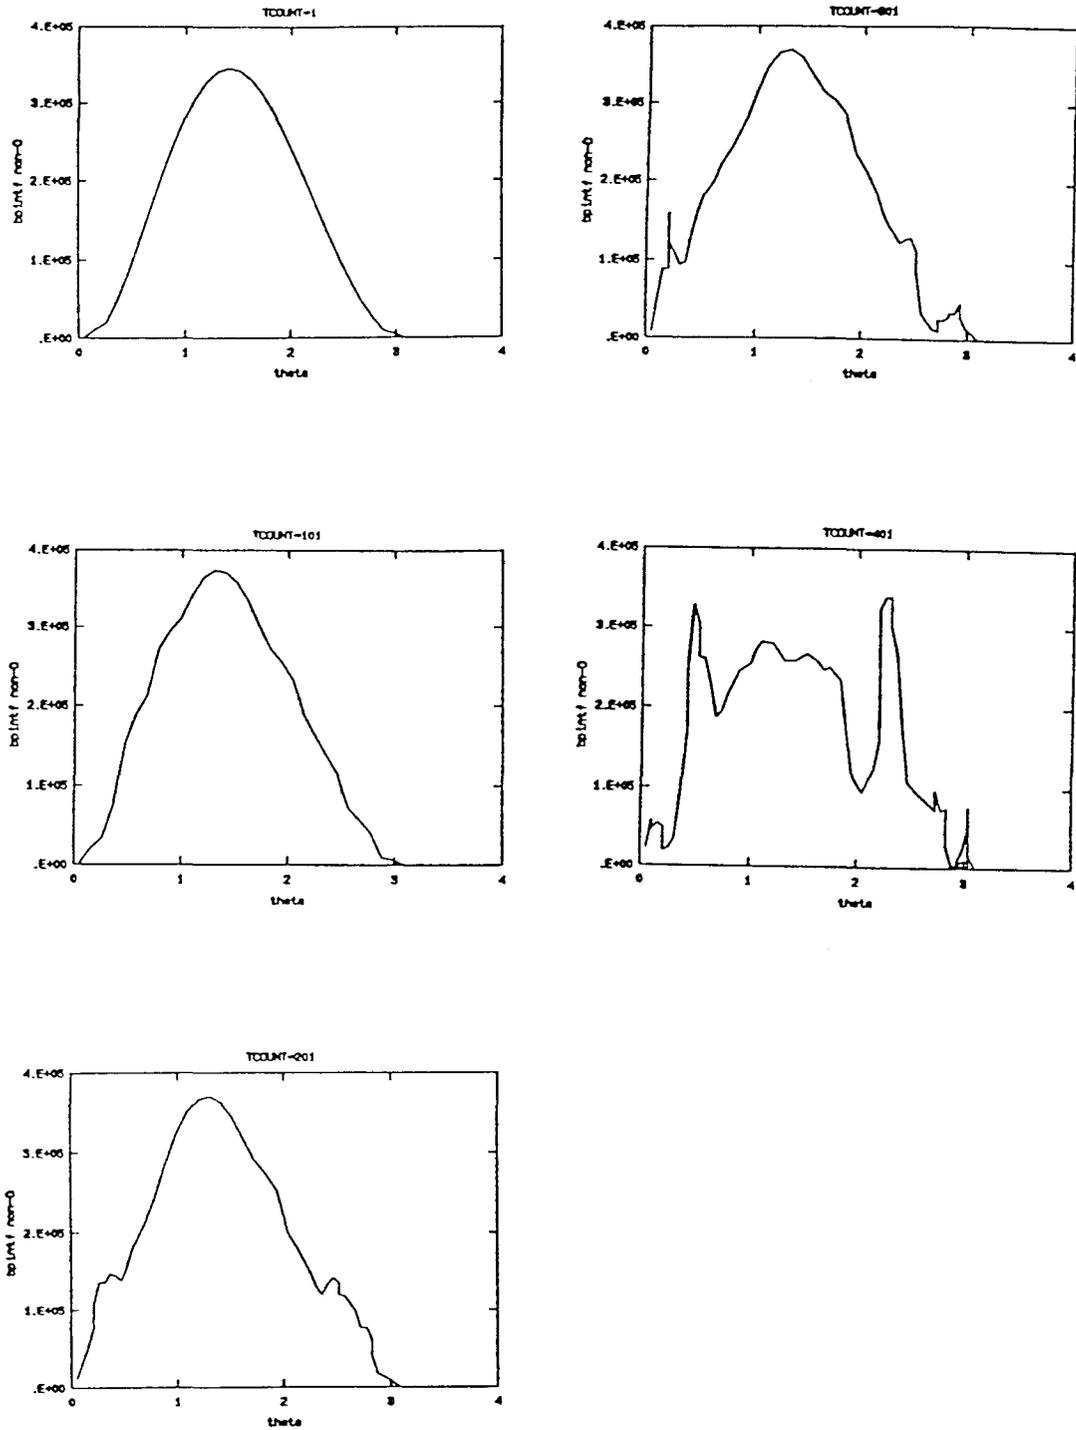


Figure 7.41

Nondimensionalized interfacial magnetic pressure (p_B') (vertical axis) v.s. polar angle (θ) in radians (horizontal axis).

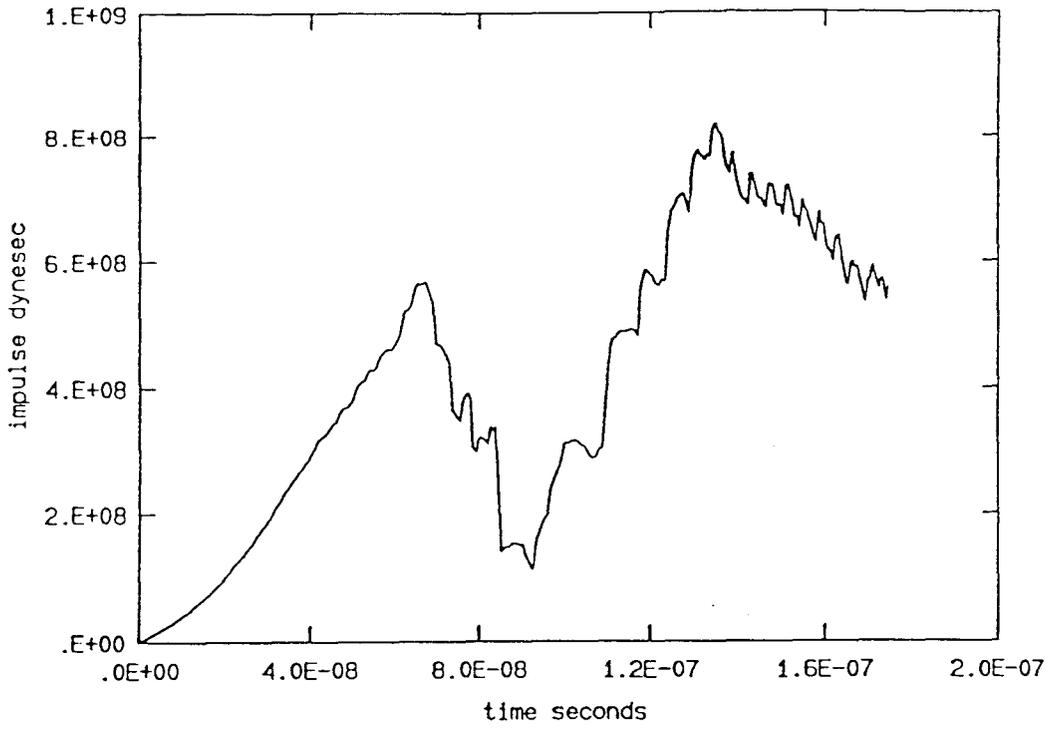


Figure 7.42a

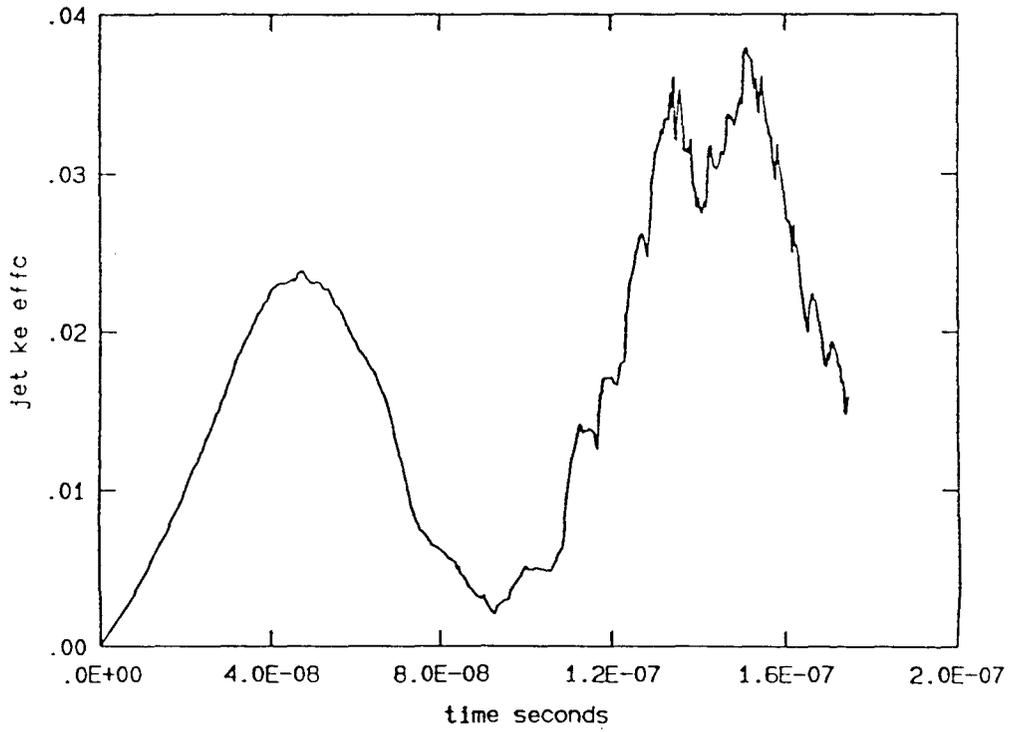


Figure 7.42b

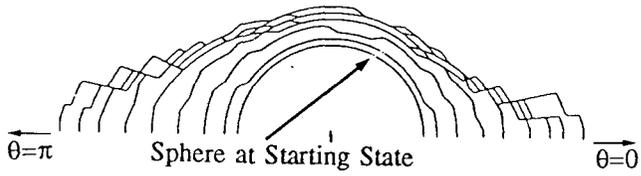
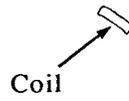


Figure 7.43a Locations of the interface at evenly spaced time intervals for Case 11.
($ER \approx 155$, $ER_1 \approx 320$, $PR = 0.33$)

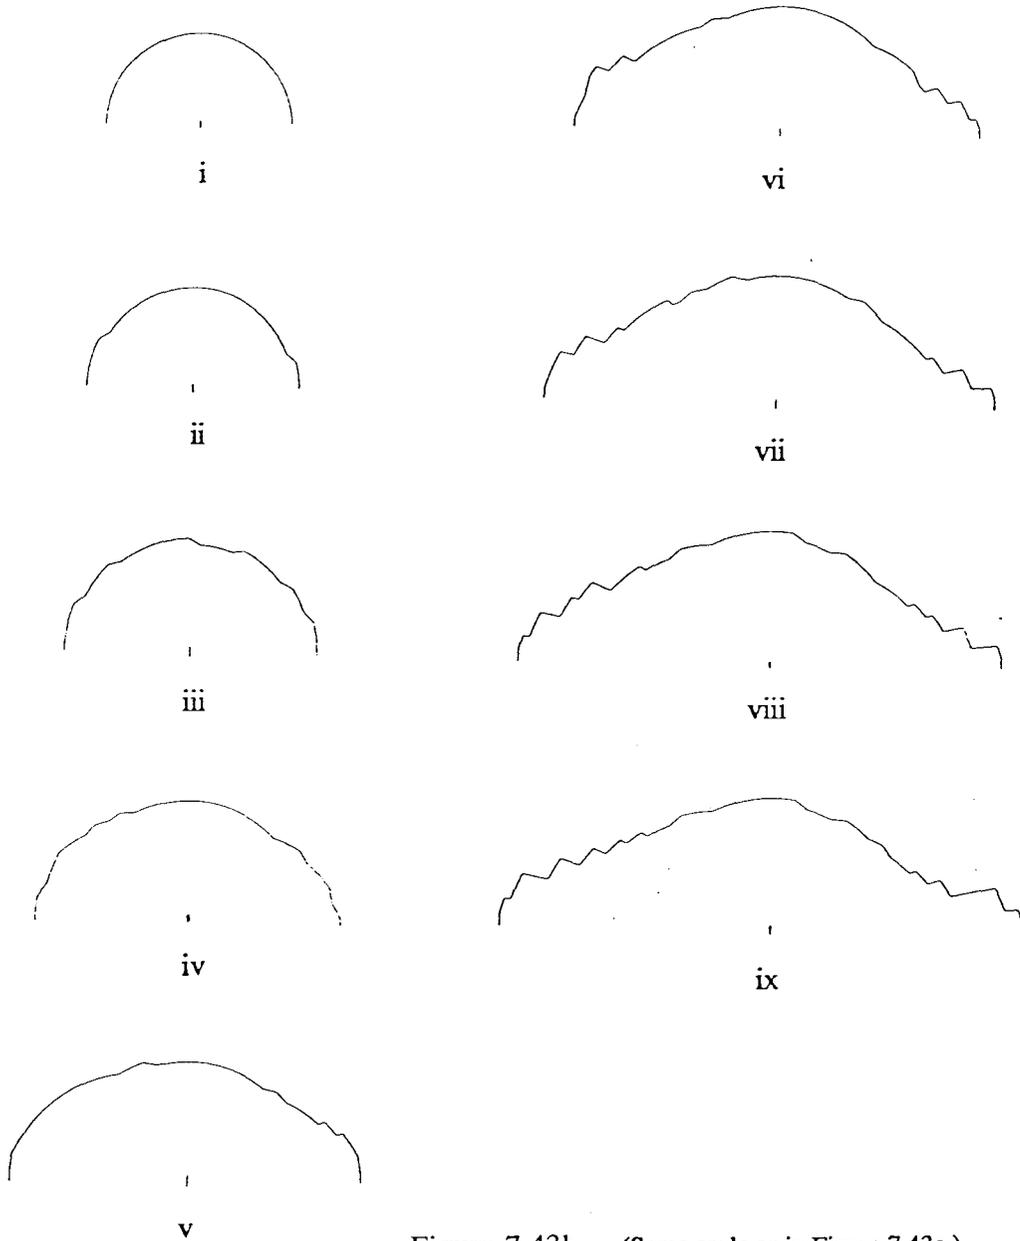


Figure 7.43b (Same scale as in Figure 7.43a.)

Chapter 8

Summary and Conclusions

The ICF pulse rocket is a potential space propulsion system, which, if feasible, will be capable of orders-of-magnitude higher performance than those of most other space propulsion systems that are within the grasps of near-future technology. Possessed of the dual advantage of both extremely high specific impulses and high thrust-to-weight ratios, the ICF pulse rocket should make possible high payload fraction missions carrying several hundred tons of payload to the outer planets with transit times on the order of months. The ICF pulse rocket is also one of the most likely candidates for "near-term" interstellar propulsion. Although, in the light of presently envisioned ICF pellet structures, it seems a little too optimistic to expect specific impulses as high as 10^6 seconds, specific impulses within an order of magnitude of this value can still be attained if the initial plasma debris energy can be converted successfully into a primarily unidirectional axial flow.

The standard ICF pulse rocket, as envisioned today, derives its thrust by redirecting an isotropically expanding ICF debris plasma with a suitably configured vacuum magnetic field, against which the plasma expands. The field forming the thruster is produced by current coils attached to the vehicle. The successful operation of the thruster therefore relies on the plasma's behaving as a highly conducting fluid, which repels external fields.

The first step in a study of the behavior of an ICF debris plasma in a magnetic thruster must be an investigation of its bulk flow, and the work in this thesis concentrated in numerical simulations of this aspect of the problem. The performance of the thruster was assessed assuming that idealized conditions hold, most notably, that the plasma will behave as a perfectly conducting medium. How closely a real plasma will behave in such a

way should be determined later, based on the results obtained from simplified analyses of bulk expansion.

Analyses that could be carried out for the bulk flow without relying on numerical simulations were generally limited to gross estimations of the variation of bulk-averaged plasma parameters with time (see Part 2 of the appendix for Chapter 3). When attempts were made to model realistic geometries and parameters, the calculations tended to become manually intractable. Simple calculations indicated that the assumption of MHD (single-fluid approximation, valid for low-frequency plasma processes) is, for the most part, acceptable, but the enormous expansion ratios experienced by the fluid in the proposed designs left something to be desired. Namely, the plasma will attain very low densities and temperatures early in the expansion, and the low temperatures can be detrimental to a successful operation of magnetic thrusters because of the accompanying low electrical conductivities. In fact, the bulk-averaged plasma temperature will fall well below the ionization temperature before the plasma has expanded to radii at which dominant thrust production takes place. The very rapid cooling resulting from the three-dimensional expansion will not be easy to counter (such as by physical processes that have sometimes been suggested, as recombination), and technological advancements to make possible stronger (although not by orders of magnitude) coil currents (using multicoil high field designs should help here) and hotter ICF debris plasmas (such as higher fuel ratios in the pellet and higher burnup fractions; these improvements will also make possible higher specific impulses) may be necessary. And if the plasma becomes low in temperature and resistive enough for resistive heating to be important, that is not a good indication, because, as has been stated many times, high conductivities are necessary for magnetic thrusters to function properly. Besides, raising the temperature of a plasma (or gas) that is not fully ionized does not necessarily lead to lower resistivities. Preliminary estimates (not discussed) suggested that high current multiple coil designs may make possible expansion

ratios that are small enough to result in plasma parameters that are more suitable for redirection by the fields. However, because of the very nonuniform structure that develops in the plasma blob, and in particular, the formation of a narrow shock heated shell-like region, a final conclusion regarding the acceptability of the currently proposed thruster parameters is pending. In addition, the growth of instabilities such as the Rayleigh-Taylor instability could possibly destroy the propulsion concept as it now stand.

Following previous work by other authors, simulations based on the thin-shell approximation were run for proposed thruster configurations and variants thereof; the interface motion was tracked until the "plasma" was expelled a large distance away from the thruster's region of influence. The hydrodynamic simulations of Chapter 7 showed that the thin-shell approximation is a crude but acceptable way to model plasma expansion against a vacuum field as long as the expansion ratio is large enough to be of relevance to flows in the thrusters of ICF pulse rockets. Promising propulsive efficiencies were obtained for a range of thruster parameters and geometries. For the single-coil design, the efficiencies clearly peaked for a specific choice of the two design parameters, the initial field-to-plasma energy ratio and the aspect ratio. If the initial field-to-plasma energy ratio is too low, the interface is affected substantially by the fields only in a narrow range of polar angles about the line of sight of a field coil. If the ratio is too high, even the downstream side of the interface is strongly deflected downwards against the central axis. The interface then developed a rather flattened appearance, betraying, in this case, the lack of a high degree of preference for rearward ejection over forward ejection.

The production of thrust was found to peak around the time that the interface experiences a severe deceleration and reverses its direction of flow. In the low-field device, the interface expanded with very little deceleration until it was very close to its turn-around position, at which point it experienced a very sudden deceleration. In contrast, with the high-field device, the deceleration, and thus thrusting, proceeded more steadily

throughout the expansion process, and velocity reversal of the interface occurred while the interface was still a large distance away from the field coils. However, even with the high-field device, significant redirection of the interface occurred only when the interface was in a very localized region about its halting position.

Generally, the reflection proceeded as if the particles were being reflected off a divergent rigid-walled nozzle, but not necessarily a paraboloid, as the discrepancy between the effective paraboloidal (nozzle) angle of subtention for the thrust efficiency and the jet kinetic energy efficiency for each thruster showed. In a low-field device, the deceleration/redirection occurs only very close to the coils, so it may be optimal to arrange the coils in the form of a contour for a most desirable rigid-walled thruster. The coils in such a device may be best configured in the form of a paraboloid (with the explosion point located at its focus), rather than in the form of a hemisphere (with the explosion site at its center of curvature) as in the Daedalus design. This is especially so because in designs such as the Daedalus, the radially expanding particles receive little polar component momentum upon reflection, and the reflected particles may rebound outwards. Compared to the importance of the geometry of the thruster, the initial field-to-plasma energy ratio does not have much significance for the flow dynamics in a very low field thruster. Unfortunately, low-field thrusters are difficult to simulate because the interface approaches the coil surfaces very closely, and numerical calculations much beyond reflection were abandoned in this work.

The favorable efficiencies obtained by the thin-shell calculations do overestimate the actual performance, even with plasma physics issues set aside, chiefly for two reasons. First, the method assumes all of the debris material to be collected into a perfectly conducting thin shell (by virtue of the deceleration arising from the finite applied pressure at the interface) located exactly at the interface, where the external force is directly applied. Second, internal fluid processes such as the conversion of kinetic energy into internal energy are not taken into account. Both of these simplifications are not precisely accurate.

For example, in the thin-shell calculations, the shell elements, for high fields, are forced down onto the central axis, when in reality, some of the reflected elements may actually rebound outwards without exiting the thruster. Also, with a substantial fraction of the particles falling onto the axis, an apparent good collimation of the shell elements representing the interface does not necessarily imply high efficiency.

After the thin-shell approximation, the next step in analyzing the bulk flow is to approximate the plasma by an unmagnetized perfectly conducting fluid that obeys the classical laws of hydrodynamics, not assuming any shell formation *a priori*. In such a simulation, the code must be able to capture automatically the unknown interface motion, unless at each time step, the correct location of the interface is determined by the less novel method of iteration, a difficult procedure. In a successful capture of the interface, the motion and shape of the interface is found self-consistently (without employing cumbersome iterations) so that the externally applied pressure at the interface and the interfacial fluid pressure are matched at every time step. This is because pressure discontinuities cannot exist across infinitesimally thin interfaces across which mass does not flow. In the case of an ICF pulse rocket, the field pressure acting on the moving boundary is a function of the geometry and location of the boundary itself. The code should also have the capability of shock capturing because of the relative difficulty of shock-fitting techniques.

In both the thin-shell calculations and the two-dimensional hydrodynamic calculations, the externally applied pressure at the plasma interface was calculated (using Maxwell's equations) as that which would actually apply there as a magnetic pressure, if both the plasma and field coil structures, situated as in proposed magnetic thruster designs, behaved as ideal perfect conductors. It will be stressed that in this work, the fields were assumed to exist solely in the vacuum region exterior to the perfectly conducting plasma, and the magnetic fields thus entered the dynamics of the plasma only through applying a

magnetic pressure force normally at the plasma-vacuum interface. By virtue of the surface currents that prevent penetration of the vacuum fields into the perfect conductors, the magnetic flux between the plasma surface and the surfaces of the field coil structures was prescribed to stay fixed at the same value as that which passed through the field coils in the absence of the plasma. Also, as the plasma interface moved, the adjustment of the vacuum magnetic field to the new location of the interface was approximated as being instantaneous on the characteristic time scale of the plasma expansion; i.e., a quasi-static approximation was used for the field calculations.

For handling the hydrodynamics, the PIC algorithm was chosen with the anticipation of its possessing a good interface-capturing capability (along with its other advantages mentioned in Chapter 5), and this proved to be the case. The automatic tracking of an applied pressure vacuum interface whose motion is primarily normal to itself is believed to be rather novel to this work, and was also an important issue in the development of the codes. Interface capturing was achieved mainly through leaving the outermost fluid cell out of the calculations and by specifying appropriate boundary conditions at the effective boundary. Especially in the slab-geometry codes, interface capturing proceeded very successfully even under adverse circumstances.

In the two-dimensional codes, limitations imposed by the computing hardware on the finite difference calculations caused degradations in the quality of the observed results. Satisfaction of the pressure matching condition at the interface was sometimes not immediately apparent, because of such problems as the limit on the grid resolution available. Although the response of the fluid pressure to the changing magnetic pressure was rather slow in the two-dimensional case, the interfacial fluid pressure showed all signs of converging upon the interfacial magnetic pressure. The bulk profile (less the numerical spikes and dips) was believed to be given correctly in all cases.

As expected, the finite applied pressure at the interface and the resultant interface deceleration, even if not persistent, caused an inward-facing shock to form from the

interfacial regions, and this was also captured nicely. The outflowing fluid experiences its primary deceleration (and redirection, depending upon the polar angle) across this shock, generally a finite distance behind the interface. Sometimes, high viscosities had to be incorporated in the numerical algorithm to suppress exaggerated numerical oscillations at discontinuities and the ringing instability. The ringing instability, which is a numerical instability often encountered in PIC simulations in regions of low speed flows, was especially troublesome when higher-order planar-geometry simulations were attempted for extended time steps in the work of Chapter 6.

The planar-geometry PIC code, when used to simulate free expansions into a vacuum or shock-tube problems, gave results that matched well with results of known theory. Also, it was found that generally, with small expansion ratios and a high amount of mass contained in the interior regions, the interface could remain almost immobile for a long time (on the time scales of the expansion) around its maximum expansion position at velocity reversal. With a fluid filled blob, interior fluid dynamics (wave propagation and heating phenomena included) could play an important role in determining the interface dynamics, and thus the performance of the device.

In the two-dimensional hydrodynamic PIC simulations, the geometries and field variations of a realistic magnetic thruster were employed, as in the thin-shell calculations, although only the single-coil design was investigated. It was found that due to the deceleration of the interface and the rapid decrease of the bulk density (and pressure) with the expansion, a relatively narrow region, in which the density and pressure rise high above the bulk values, indeed does form. This was so, even for weak interfacial pressures, as long as the expansion ratio was not too small. This narrow and highly peaked shell-like region (difficult to represent completely accurately in a numerical simulation, especially with a coarse grid), bounded on its inner side by an inward-facing shock, generally traveled away from the interface, if the interface experienced acceleration after the “shell” had formed. This was observed commonly at polar angles where the

applied interfacial magnetic pressure decreased with the expansion. The shell stayed against the interface at other angles; i.e., the sharp pressure and density rise region remained confined to a narrow region situated against the interface, as the interface continued to decelerate at such angles. Whether the "shell" stayed at the interface or not depended upon the deceleration/acceleration history of the interface, and thus also the pressure profile near the interface and the history of the applied pressure at the interface. At polar angles for which the shell-like structure travels away from the interface, the density and pressure drops low in a rarefaction region between the shell-like structure and the interface, and the radial velocity rises again towards the interface. However, the temperature, which also rises at the shock, remains high (at a relatively constant value) all the way to the interface. The raising of temperature in the interfacial regions (due to rethermalizations), which occurs as part of the deceleration/redirection process, and in particular, due to shocks, is a welcoming phenomenon for propulsion applications because of the necessity to maintain sufficient conductivities in that region as the bulk flow cools dramatically from the expansion.

The qualitative features of the interface motion and deformation obtained from the two-dimensional hydrodynamic runs are similar to those obtained from calculations employing the thin-shell approximation, as are the trends of the interface dynamics with varying field-to-plasma energy ratio. The flow initially expanding into the forward hemisphere is found to be redirected primarily in the negative radial direction, back towards the explosion site, not only implying a rather inefficient propulsive flow, but also suggesting the possibility of some of the fluid's not leaving the thruster in one reflection. Under the approximations used, confinement of the plasma in the forward axial direction appears to be quite good for reasonably high field thrusters.

With a fluid filling the blob and with conversions between kinetic and internal energies allowed for, more parameters are required to describe the flow process than in the thin-shell case. One of the primary parameters will still be the field-to-plasma energy ratio.

However, the quantity that should be used in place of the field energy here is really the work done by the plasma against the magnetic pressure applying at the interface (this may be termed as the energy of the displaced fields), and not the field energy contained in the entire thruster.

The hydrodynamic simulation of two-dimensional expansion flows often failed before all interesting stages of the propulsive flow could be treated, as a result of cells devoid of, or nearly devoid of, computational particles forming, where in reality, the density should be nonzero. More computational particles must be used unless the fixed grid approach is abandoned.

In the two-dimensional hydrodynamic simulations, memory limitations of computers forced the use of plasma parameters rather far removed from those that would describe flows in the thrusters of ICF pulse rockets, if all important phases of the expansion flow process were to be treated. In other words, if realistic parameters were chosen, the hydrodynamic simulations could only be carried out part of the way through the important phases of the expansion process. In particular, to treat more than a portion of the important phases of the expansion flow process, the expansion ratio to be experienced by a fluid during a simulation had to be reduced quite significantly from that typical of real situations. This is because, in PIC, the lowest density that can be treated depends upon the number of particles per cell used to represent the initial density. Thus the maximum expansion ratio treatable is limited by the total number of computational particles employed.

Because the thruster size was chosen to be the same as that of a proposed design, and a maximum expansion distance on the same order as the characteristic thruster dimension was considered desirable, the two-dimensional hydrodynamic simulations were begun with a uniform density sphere with a radius as large as 50 cm when the field-coil radius of the thruster was only 650 cm. For a real ICF plasma that is born with a radius that is very small compared to the dimensions of the thruster, the state, at say a 50 cm

radius, will be very far removed from those of the uniform blobs assumed in Chapter 7. The use of this type of an unrealistic starting state possessing uniform density and pressure at a very large radius, resulted in a nonnegligible fraction of the debris mass's being contained in a large region not influenced by the deceleration of the interface during the simulation of the expansion process.

Primarily for this reason, the propulsive efficiencies obtained from the simulations were notably lower than those obtained by simulations that assumed all the mass to be collected into a thin region that always stays precisely at the interface at which the external pressure is applied. However, taking into account the unrealistic starting state selected, the efficiency values were acceptable. The conversions between kinetic and internal energies, and in particular, the incomplete conversion back from thermal energy to directed kinetic energy, played only a secondary role in lowering the propulsive efficiencies below the thin-shell case, even when these unrealistic starting states were employed. It should be realized though, that in a realistic hydrodynamic analysis, the conversion between kinetic and internal energies (commonly involving shocks) forms an integral part of the deceleration/redirection process of the fluid flow.

Although the thin-shell model cannot take into account many of the losses that will be present in a fully hydrodynamic model, hydrodynamic simulations initiated from a realistic starting state are expected to give efficiencies much closer to those obtained by the thin-shell simulations than those obtained by the hydrodynamic simulations of Chapter 7 because of the very unrealistic starting states employed in the latter.

In order to execute two-dimensional hydrodynamic simulations more closely approximating, in terms of parameters, flows in thrusters that are being proposed for ICF pulse rockets, while at the same time treating the propulsive flow well after the first reflection due to the recoil of the fields, a code with nonuniform, moving grids should be written so that large expansion ratios may be handled without requiring a prohibitive

number of computational particles. Although the peak of the thrusting process may be over after the first inward reflection of the interface, it will still be important to follow the fluid flow after that stage because the behavior of a thin-shell simulation with no internal fluid pressures deviates strongly from reality thereafter. The code should allow starting the simulations from a small radius of the interface at which a uniform state* may be considered acceptable for the plasma blob. If the number of computational particles used in a simulation can be reduced, a code will also be able to run a larger number of time steps, under the constraints of limited CPU time, because the lower the number of particles, the less time it takes for the computer to execute calculations for one time step.

Once bulk flow simulations using realistic numbers are performed, the results can be used to calculate various plasma parameters and characteristic time scales for relevant phenomena such as field diffusions and instabilities. This will be crucial in estimating the feasibility of a thruster design.

However, the results of a classical hydrodynamic simulation are valid as an approximation to the real behavior of a plasma only if the idealizing assumptions used are legitimate. It cannot be stressed too much that a real plasma, no matter how close to "ideal," and especially at low temperatures, may behave quite differently from a perfectly conducting fluid obeying just the laws of classical hydrodynamics. This will no doubt be a source of difficulty in getting the magnetic thrusters of ICF pulse rockets to operate successfully, much less attain the efficiencies quoted in this work for ideal situations. For example, while possessing a tendency to follow field lines, which itself will work against the efficient ejection of a jet once the field diffuses into the plasma (although plasma-field separation will probably occur decently), a real plasma will diffuse across fields, and may also experience cross field drifts due to polarizations. A perfect conductor does not allow external fields to diffuse into its interior. But a real plasma expanding in the magnetic

* The choice of a uniform state to serve as the starting state for simulations is of course arbitrary, but it simplifies matters especially when good "ICF debris data" is not available.

thruster of an ICF pulse rocket will have finite resistivity values dependent upon the state of the local plasma. Because, as mentioned above, the principle of operation of a magnetic thruster for an ICF pulse rocket relies on the plasma in the thruster to behave as a very good conductor, repelling external fields, the time scale of magnetic field diffusion into the plasma, relative to the characteristic time scale of plasma expansion will be a parameter of great importance to the problem. If the resistivity is low enough for the characteristic time scale of field diffusion into the plasma to be much longer than the expansion time scale of the plasma in the thruster, field diffusion into the plasma will not be a problem.

Therefore, after the classical hydrodynamic stage of investigations is completed, the next major step will be the incorporation of more detailed and realistic plasma physics, and most notably finite resistivities, into the numerical calculations. Namely, this will be an upgrade of the classical hydrodynamic code to an MHD code.** Although hybrid codes taking into account some aspects of a two-fluid approach may possibly be attempted, time-scale and length-scale considerations make a two-fluid numerical treatment of the bulk flow impractical. However, two-fluid simulations treating flow phenomena locally in space and time should be possible, and will be necessary for obtaining a more complete understanding of the real behavior of a plasma in a magnetic thruster. Models incorporating realistic velocity distributions of plasma particles will also be important in determining accurately, the plasma expansion and deceleration/redirection process. Equally important will be the addition of realistic physics, such as radiation and recombination, to the calculations and the incorporation of self-consistent transport coefficients. Also, a real

** By "an upgrade of a classical hydrodynamic code to an MHD code," primarily the inclusion of finite resistivities is implied here. This will mean dealing with a magnetized plasma. Also, even if we have a perfectly conducting problem, we will no longer be associated with classical hydrodynamics if the fluid is magnetized. However, it is also true that even if we are dealing with an initially unmagnetized plasma, and set the resistivity term appearing in the equations to zero, an MHD analysis will differ from a classical hydrodynamic analysis of an unmagnetized perfectly conducting fluid, unless the simplest form of the MHD equations is used. Therefore, a distinction is made in this work between a strictly classical hydrodynamic analysis and a perfectly conducting MHD analysis of an (initially) unmagnetized fluid.

ICF plasma may be born permeated by a rather strong magnetic field and may experience rotational motion as well, and these and other details need considerations too. Finally, as the expanding plasma system will be vulnerable to physical instabilities, instability analyses will be a very important consideration in future work. For example, the bulk flow is of the kind especially susceptible to the Rayleigh-Taylor type instability (flute instabilities). If the time scales of development of potential instabilities are shorter than the expansion time scale of the plasma, instabilities may well render the device inoperable.

In conclusion, according to the approximate analyses carried out in this work, the present designs of magnetic thrusters for ICF pulse rockets, if the highly idealized conditions assumed hold, do appear to be capable of thrust efficiencies at least as high as about 70 %, and jet kinetic energy efficiencies (generally not equal to the square of the thrust efficiency) well above 50 % (these values were obtained through application of the thin-shell approximation, with particles falling onto the central axis). Although it can be shown through a simple calculation that the effective exhaust velocity is not equal to the initial isotropic expansion velocity of the debris multiplied by the thrust efficiency, this product should give a rough estimate for the effective exhaust velocity. This implies that effective specific impulses on the order of the value corresponding to the initial isotropic expansion velocity may be achieved. It is true, though, that the flow was not always being redirected in the most favorable fashion for propulsive purposes in the proposed designs, especially with the low-field cases such as that of the Daedalus.

For the single-coil high-field thruster, performances peaked for rather high field-to-plasma energy ratios (between about 15 and 30, based on the thin-shell approximation). With the low-field thruster, a multicoil design will be necessary, but as long as the coils are arranged in an optimal geometry, the low-field design should be able to achieve propulsive efficiencies as high as those of high-field thrusters. However, because the impulse is

transferred to the vehicle structure over a much more spread over period of time during the expansion process in a high-field thruster, high-field thrusters possess a far better "shock-absorbing" property than their low-field counterparts. The need for multiple coils in the hazardous environment of neutrons and x/γ rays, and the close approach of plasma to the coil structures in the low-field device also seem to render the high-field design more attractive, although coil current (and thus coil conductivity) requirements are more demanding with the latter.

Unfortunately, the favorable performances predicted in this work do not directly translate into the behavior of real magnetic thrusters. As stated above, the extremely large expansion ratios experienced by the plasma in the proposed designs, while the plasma is still in the thruster, will cause the plasma not to have the very high conductivities necessary for favorable operation of the device. The plasma will not even be fully ionized, but the nonuniform distributions that develop within the plasma blob with shock heating near the interface does supply some hope in maintaining high enough conductivities, at least at the important interfacial regions. Even at very high temperatures and conductivities with little field diffusion, getting the thruster to work will not be easy because of such phenomena as instabilities. Furthermore, hydrodynamic simulations of extremely large expansion ratio flows with parameters very closely resembling those of situations to be found in proposed thrusters are yet to be carried out. Thus again, repeating an earlier remark, no conclusions as to the feasibility and performance (e.g., efficiency) of the proposed designs can be drawn at this point. And even the two-dimensional fully hydrodynamic simulations utilizing realistic parameters, by assuming not only perfect conductivity, but also classical hydrodynamic behavior instead of real plasma behavior, are highly idealized, and their results will be over optimistic, not truthfully bringing out all the unfavorable phenomena that may actually occur.

Regarding future work, it will be important first to complete the bulk flow analysis by carrying out the classical hydrodynamic simulation, using realistic parameters closely

resembling those proposed for ICF pulse rocket thrusters, all the way through all the important phases of the flow process with a very large expansion ratio, in order to obtain data on the temporal and spatial variations of fluid variables under the exact conditions of proposed thrusters. The thruster design parameters may then have to be modified as necessary. After that stage, in-depth studies involving detailed plasma physics should be conducted in an attempt to come up with a working engine. Much work remains to be done in this field, but as the fundamental operating principle of a magnetic thruster for an ICF pulse rocket is physically sound, and the thruster performance under idealized conditions, attractive, it is well worth the effort to continue the study.