Mon. Not. R. astr. Soc. (1991) **250,** 581–596

Self-similar jets

S. A. E. G. Falle
Department of Applied Mathematical Studies, The University, Leeds LS2 9JT

Accepted 1991 January 21. Received 1991 January 14; in original form 1990 October 2

SUMMARY

General arguments and numerical calculations are used to show that the flow caused by a supersonic gas jet is self-similar under certain conditions. If we assume that the jet has a high initial Mach number and is generated in a region small compared to its length, then the type of similarity solution depends on the density distribution of the gas through which the jet propagates. If this density decreases faster than $1/R^2$, where R is the distance from the source, then the length of the jet increases linearly with time and it may evolve into a classical double if it subsequently encounters a region of higher density. In a more slowly varying external density, the jet is reconfined and the similarity exponent is the same as for an isotropic wind with a constant rate of energy input. At intermediate times this looks like a classical double, but at large times it has many of the characteristics of FRI sources

1 INTRODUCTION

In recent years there has been a great deal of work on hydrodynamical models of extragalactic jets. However, most of this has either concentrated on the working surface at the end of the jet (Norman *et al.* 1982; Kössl & Müller 1988; Chakrabarti 1988) or looked at the flow near the source where one might expect the jet to be steady (Sanders 1983; Falle & Wilson 1985; Wilson & Falle 1985; Wilson 1986, 1987; Daly & Marscher 1988). So far nobody has tried to model the whole of a jet numerically, although there has been some analytic work on the subject (Scheuer 1974; Blandford & Rees 1974).

Since jets generally propagate for distances which are large compared to the scale on which they are generated, it seems natural to look for similarity solutions to describe the large-scale structure. Federenko & Zentsova (1986) have done just this and show that at large times the length of the jet increases like $t^{3/5}$ if the external density is constant. However, they did not consider what would happen in a non-uniform external density, nor did they look at the details of the flow in the jet.

The fact that the jets are long compared to the size of the source galaxy means that the external density varies enormously over the length of the jet and it is important that one takes this into account (Baldwin 1982). We shall see that the nature of the flow depends upon whether the external density decreases more or less rapidly than $1/R^2$, where R is the distance from the source. If it decreases faster than $1/R^2$, then the jet density remains larger than that of the surroundings, if it is so initially and the head of the jet advances at a constant speed. The resulting flow looks very much like a classical double except that it is unlikely to produce much radio emission since only a small fraction of the energy is thermalized. It would, however, become visible if it subsequently enters a region in which the density decreases more slowly.

If, on the other hand, the external density decreases more slowly than $1/R^2$, then the jet blows a low density cavity and is considerably modified as it propagates through the cavity on its way to the working surface. Dimensional arguments suggest that the expansion of this cavity is self-similar with the same similarity exponent as an isotropic wind in a power law external density. Numerical calculations with a constant external density confirm that the global expansion of the cavity is approximately self-similar; although the flow near the working surface is not since it is dominated by a periodic vortex shedding. It turns out that the pressure in the cavity is approximately uniform far from the working surface and this allows a very simple treatment of the way in which the jet is reconfined by the pressure in the cavity.

Although radio jets clearly contain magnetic fields, we do not really know whether or not they are strong enough to be dynamically important. It has been argued that jets can only be confined and stabilized by magnetic fields and there are some recent numerical calculations that explore such effects (Kössl, Müller & Hillebrandt 1990a,b), but in fact laboratory experiments show that pure gas jets are quite capable of propagating for large distances provided their Mach number is high enough (Dash & Wolf 1983). Laboratory jets are eventually destroyed by turbulence, but astrophysical jets can avoid this as long as their opening angle is large enough.

We therefore see no reason why we should complicate matters by trying to include the dynamical effects of the magnetic fields.

2 SIMILARITY SOLUTIONS

The observations tell us that radio jets are produced and collimated on very small scales and their energy is such that they propagate for distances much larger than the size of the galaxy within which the source is embedded. So, as far as the large-scale structure is concerned, we can ignore what happens near the source and simply assume that two symmetric conical jets emerge from their parent galaxy with some power P, mass flux Q and opening angle θ. If θ is non-zero and the jet is initially supersonic, then its internal Mach number will increase with distance from the source and we might as well take it to be infinite. We can also safely neglect the pressure in the gas surrounding the galaxy as it is likely to be small compared to the dynamic pressure of the jets.

Things can be further simplified by supposing that the external gas has a spherically symmetric density distribution centred on the source and that it is at rest. This is clearly not always going to be the case, but it should be appropriate for galaxies at the centre of clusters. It should also be approximately correct for other galaxies while the jet is still short enough not to be significantly affected by the motion of the galaxy through the surrounding gas.

The parameters of the problem are the ratio of specific heats, γ; the density of the surrounding gas, $\rho_e(R)$, where R is the distance from the source; the initial half opening angle of the jet, $\theta (\theta \ll 1)$; the jet power, P; and the jet mass flux, Q.

Since we are interested in similarity solutions, we will suppose that the external density is a power law

$$\rho_e = \frac{a}{R^\alpha}, \tag{2.1}$$

where a and α are constants. From these parameters we can form a characteristic length

$$l_c = \left(\frac{aP^{1/2}}{Q^{3/2}}\right)^{1/(\alpha-2)}, \tag{2.2}$$

which for a uniform external density ($\alpha = 0$) is

$$l_c = \frac{Q^{3/4}}{\rho_e^{1/2} P^{1/4}}, \tag{2.3}$$

where ρ_e is the constant external density.

If we set $\alpha = 0$ and $P = 10^{38}$ W, $Q = P/c^2$, $\rho_e = 10^{-23}$ kg m^{-3}, we get $l_c = 20$ pc, which is very small compared to the typical size of radio jets. The only way in which l_c could be much larger than this is if the initial jet velocity is very small compared with the speed of light. This is unlikely to be the case and so it seems that l_c will always be much smaller than the length of the jet, at least for the constant density case.

Once the jet has become much larger than l_c the swept up ambient material dominates the mass involved in the motion and the flow of this gas ought not to depend on the jet mass flux Q. This is analogous to the situation for a spherically symmetric wind (Dyson, Falle & Perry 1980) and the flow ought to be as sketched in Fig. 1. Here S is a shock bounding the disturbed region and C is a tangential discontinuity separating the jet gas from the swept up ambient gas. This is very similar to Scheuer's model A (Scheuer 1974), except that his relativistic beam has been replaced by a fluid jet.

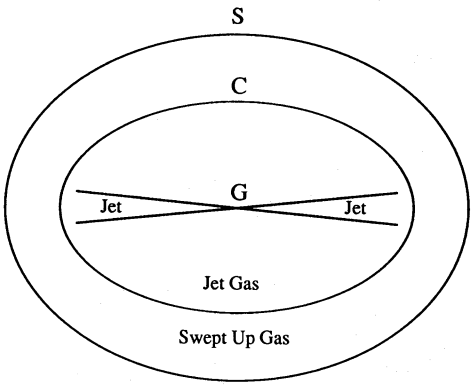

Figure 1. The self-similar flow. Here G is the galaxy producing the jets, S is a shock bounding the disturbed region and C is a tangential discontinuity separating the jet gas from the swept up ambient gas.

The shape of the surfaces S and C will be of the form

$$S(r, z, t, P, Q, a, \alpha, \theta, \gamma) = 0 \tag{2.4}$$

and

$$C(r, z, t, P, Q, a, \alpha, \theta, \gamma) = 0, \tag{2.5}$$

where r and z are cylindrical polar coordinates with the z axis along the jet. Once the flow has become self-similar S and C will no longer depend upon Q and then dimensional arguments demand that their shapes be given by

$$S'(\xi, \eta, \alpha, \theta, \gamma) = 0 \tag{2.6a}$$

and

$$C'(\xi, \eta, \alpha, \theta, \gamma) = 0, \tag{2.6b}$$

where the similarity variables ξ and η are

$$\xi = r\left(\frac{a}{Pt^3}\right)^\beta, \tag{2.7a}$$

$$\eta = z\left(\frac{a}{Pt^3}\right)^\beta, \tag{2.7b}$$

and

$$\beta = \frac{1}{(5-\alpha)}. \tag{2.7c}$$

In a constant external density we get the usual result

$$\xi = r\left(\frac{\rho_e}{Pt^3}\right)^{1/5}, \tag{2.8a}$$

and

$$\eta = z\left(\frac{\rho_e}{Pt^3}\right)^{1/5}, \tag{2.8b}$$

i.e. the disturbed region expands like $t^{3/5}$.

Although the flow in the swept up gas is self-similar in this way, this cannot be true for the flow in the cavity enclosed by the tangential discontinuity C since the self-similar velocity behaves like $t^{3\beta-1}$, but the jet velocity is constant. What must happen is that the jet becomes approximately steady along most of its length and the pressure in the cavity is nearly uniform. That this is true for large times can be seen from the following argument, which is essentially that used by Scheuer (1974).

The speed of expansion behaves like $t^{3\beta-1}$ while the jet velocity is constant, so, provided $3\beta-1\leq 0$, the jet gas must eventually pass through a shock which is almost stationary with respect to the jet. It will therefore acquire a sound speed of the same order as the jet speed and this will be large compared to the expansion speed of the cavity in which the exhaust gas from the jet resides. One would therefore expect the pressure in this cavity to be approximately uniform and to change on a time-scale t. However, if the jet speed is v_j, then the time-scale for the jet, out to a distance L from the source, to adjust to changes in the pressure around it is

$$t_j=\frac{L}{v_j}. \tag{2.9}$$

Hence as time increases the condition $t_j \ll t$ allows a larger and larger fraction of the jet near the source to be approximately steady. The shock structure in such a steady jet has been looked at in some detail by several authors (Sanders 1983; Falle & Wilson 1985; Wilson & Falle 1985; Wilson 1986, 1987; Daly & Marscher 1988).

The question is now what happens at the end of the jet where it hits the end of the cavity. If the jet is not confined between the source and the working surface, then it must have a constant opening angle and since its length increases like $t^{3\beta}$, its radius must also increase like $t^{3\beta}$. If its Mach number is large, then its velocity will remain constant along its length and so the ram pressure p_r at the working surface must behave like

$$p_r \propto (\text{jet radius})^{-2} \propto t^{-6\beta}. \tag{2.10}$$

On the other hand the bow shock advances at a speed proportional to $t^{3\beta-1}$ and it sees an ambient density which behaves like $t^{-3\alpha\beta}$. Its pressure p_b must therefore satisfy

$$p_b \propto t^{6\beta-2-3\alpha\beta}. \tag{2.11}$$

Clearly these two pressures must have the same time dependence and this is only possible if

$$12\beta-3\alpha\beta=2. \tag{2.12}$$

Using (2.7c) we see that this requires $\alpha=2$, i.e. the external density decreases with the square of the distance from the source. We thus have two quite distinct cases depending on whether $\alpha<2$ or $\alpha>2$.

2.1 Case I ($\alpha<2$)

If $\alpha<2$, then the jet cannot have a constant opening angle along its entire length since its ram pressure would be insufficient to drive the bow shock. It must therefore reconfine and it is quite easy to work out where this happens. Since the jet material is expanding adiabatically as it emerges from the source, its pressure will quickly drop to a very low value and will therefore have a negligible effect on the motion. We can therefore assume that the jet emerges ballistically from the source with a uniform density across it. Its density at a distance z is then given by

$$\rho_j \approx \frac{Q^{3/2}}{\pi\theta^2(2P)^{1/2}z^2} \tag{2.13}$$

since θ is small.

Now let $r(z)$ be the radius of the reconfinement shock at distance z. Then if ψ is the angle it makes with the z axis, we have

$$\frac{dr}{dz}=\tan\psi. \tag{2.14}$$

The normal component of the velocity v_n is then

$$v_n=v_j\sin|\phi-\psi|, \tag{2.15}$$

where

$$\phi=\frac{r}{z} \tag{2.16}$$

is the angle the upstream flow makes with the z axis. Since the upstream pressure is very small, the reconfinement shock will be strong and its post shock pressure is

$$p_s=\frac{2}{(\gamma+1)}\rho_j(z)\,v_n^2=\frac{2}{(\gamma+1)}\rho_j(z)\,v_j^2\sin^2(\phi-\psi). \tag{2.17}$$

The jet velocity can be assumed to be constant up to this point, and so we have

$$v_j=\left(\frac{2P}{Q}\right)^{1/2}. \tag{2.18}$$

Putting (2.13) and (2.18) into (2.17) we get

$$p_s=\frac{\mathscr{K}}{z^2}\sin^2(\phi-\psi), \tag{2.19}$$

where

$$\mathscr{K}=\frac{2}{(\gamma+1)}\frac{(2PQ)^{1/2}}{\pi\theta^2}. \tag{2.20}$$

As a first approximation it seems reasonable to assume that the post-shock pressure p_s is equal to the uniform pressure in the cavity p_c. We then have

$$\sin(\phi-\psi)=\sin\left(\frac{r}{z}-\tan^{-1}(dr/dz)\right)=z(p_c/\mathscr{K})^{1/2}. \tag{2.21}$$

Canto, Raga & Binette (1989) used this equation in their analysis of the reconfinement shock in stellar jets. They show that it can be integrated exactly if we transform to polar coordinates, but, since the pressure behind the reconfinement shock will only be uniform if the angles are small, we might as well use this to simplify the analysis. This will be valid provided the jet opening angle θ is small and the post-shock pressure is small compared to the jet ram pressure so

that $\phi - \psi$ is also small. In that case equation (2.21) reduces to

$$\frac{dr}{dz} - \frac{r}{z} = -z(p_c/\mathscr{K})^{1/2}. \tag{2.22}$$

This can easily be integrated to give

$$r = \theta z - z^2(p_c/\mathscr{K})^{1/2}, \tag{2.23}$$

where we have applied the obvious boundary condition

$$\frac{r}{z} \to \theta \quad \text{as} \quad z \to 0.$$

According to equation (2.23) the shock reaches the z axis at

$$z = z_1 = \theta(\mathscr{K}/p_c)^{1/2} \tag{2.24}$$

and it then makes an angle

$$\left(\frac{dr}{dz}\right)_{z=z_1} = -\theta \tag{2.25}$$

with the z axis. Actually in axisymmetry the shock must be normal at the axis, but the resulting Mach disc will only be significant if θ is too large for a regular reflection. For strong shocks this requires

$$\theta > \sin^{-1}(1/\gamma) = 36.9° \qquad \text{for } \gamma = 5/3$$

and so is unlikely to occur for typical jet opening angles.

In all this we have assumed that the cavity pressure p_c is uniform and small compared to the jet ram pressure. Both these assumptions fail near the working surface, since there p_c must be of the order of the ram pressure and the Mach number of the gas in the cavity is not small. However, if the aspect ratio of the cavity is large, then the expansion speed of the sides of the cavity will be small compared to that of its ends. Near the source, the cavity gas will then have a uniform pressure which is small compared to the jet ram pressure. We shall see later that the aspect ratio of the cavity is large for reasonable jet opening angles, and so all we have to do is to check that z_1 is small compared to the length of the jet.

The length L_j of the jet and p_c are given by

$$L_j = \mathscr{A}\left(\frac{P}{a}\right)^{\beta} t^{3\beta} \tag{2.26}$$

and

$$p_c = \mathscr{B} a \left(\frac{P}{a}\right)^{2\beta - \alpha\beta} t^{6\beta - 2 - 3\alpha\beta}, \tag{2.27}$$

where $\mathscr{A}(\theta, \gamma)$ and $\mathscr{B}(\theta, \gamma)$ are dimensionless quantities. Using these and (2.24) we find that

$$z_1/L_j \propto t^{1 + 3\alpha\beta/2 - 6\beta} \propto t^{(\alpha - 2)/(10 - 2\alpha)} \tag{2.28}$$

and this ratio decreases with time as long as $\alpha < 2$. z_1 will therefore eventually become small compared to L_j which means that at large times the reconfinement occurs far from the working surface where the cavity pressure should be uniform.

The fact that the pressure behind the reconfinement shock is small compared to the jet ram pressure means that this shock does not involve much energy dissipation in the jet and should therefore not appear as a particularly prominent feature in radio maps. However, it does have an important effect in that it reduces the opening angle and so ensures that the jet ram pressure at the working surface is large enough to drive the bow shock.

We must now consider what happens downstream of the reconfinement shock. The jet is behaving like a steady underexpanded jet in a uniform external pressure and so there will be another shock cell downstream of z_1. For $z > z_1$ the jet pressure is p_c and since this is small compared to the ram pressure, the jet velocity is still approximately equal to its initial value. The jet Mach number is then

$$M_j(z) = \frac{v_j}{c_j} = v_j \left(\frac{(\gamma+1)\rho_j}{(\gamma-1)\gamma p_c}\right)^{1/2}. \tag{2.29}$$

Substituting for v_j and ρ_j from (2.13) and (2.18) we get

$$M_j(z) = \left(\frac{\gamma+1}{\gamma(\gamma-1)\pi p_c}\right)^{1/2} \frac{(2PQ)^{1/4}}{\theta z}. \tag{2.30}$$

Using (2.24) we can write this in terms of z_1:

$$M_j(z) = \frac{(\gamma+1)}{[2\gamma(\gamma-1)]^{1/2}} \frac{z_1}{\theta z}. \tag{2.31}$$

The jet Mach number at $z = z_1$ is then

$$M_j(z_1) = \frac{(\gamma+1)}{[2\gamma(\gamma-1)]^{1/2}} \frac{1}{\theta}. \tag{2.32}$$

As we are dealing with a small disturbance, the next shock cell will at a distance

$$z_2 = z_1 + 2M_j(z_1)\, r_j(z_1) \tag{2.33}$$

from the source. Substituting for $M_j(z_1)$ from (2.32) and setting $r_j(z_1) = \theta z_1$ we get

$$z_2 = [1 + 2M_j(z_1)\,\theta]\, z_1 = \left[1 + \frac{\sqrt{2}(\gamma+1)}{[\gamma(\gamma-1)]^{1/2}}\right] z_1. \tag{2.34}$$

It is interesting that the ratio z_1/z_2 depends only on γ and not on the jet parameters. It must, of course, also be independent of the inclination of the jet to the line-of-sight, and would therefore seem to be a method of determining the effective γ of the jet fluid. The formula gives

$$\frac{z_2}{z_1} = 3.58 \qquad \text{for } \gamma = 5/3$$

and

$$\frac{z_2}{z_1} = 4.95 \qquad \text{for } \gamma = 4/3$$

and is thus quite sensitive to the value of γ.

Finally, we have to check that the jet ram pressure $p_r(z_1)$ at $z = z_1$ behaves correctly. We have

$$p_r(z_1) = \rho_j(z_1)\, v_j^2 = \frac{(2PQ)^{1/2}}{\theta^2 \pi z_1^2} = \frac{(\gamma+1)}{2\theta^2}\, p_c, \tag{2.35}$$

from equations (2.13), (2.18) and (2.24). As the cavity pressure p_c is proportional to the bow shock pressure, this means that the ram pressure decreases with time in the same way as the bow shock pressure. Furthermore (2.35) shows that our analysis is consistent in the sense that p_c is small compared to the ram pressure provided the jet opening angle θ is small.

Equation (2.35) also gives us an idea of the aspect ratio, $\mathscr{R}$ of the disturbed region. This is simply the ratio of the advance speed of the bounding shock at the end of the jet to that of the sides of the cavity. The shock at the side of the cavity is driven by a pressure p_c, while that at the end of the jet is driven by the jet ram pressure at the working surface $p_r(L_j)$. We expect the jet to expand somewhat between the reconfinement shock and the working surface and therefore

$$p_r(L_j) \leq p_r(z_1).$$

Using (2.35) we then have

$$\mathscr{R} = \left(\frac{p_r(L_j)}{p_c}\right)^{1/2} \leq \left(\frac{p_r(z_1)}{p_c}\right)^{1/2} = \left(\frac{\gamma+1}{2}\right)^{1/2} \frac{1}{\theta}. \tag{2.36}$$

So far we have supposed that the jet boundary is a slip line with zero thickness, but this is clearly not true. Experiments with laboratory jets (Dash & Wolf 1983) show that there exists a turbulent boundary layer which grows more or less linearly with distance down the jet. Canto & Raga (1990) have used a simple momentum balance argument to deduce that

$$\theta_T \propto \frac{1}{M_j} \left(\frac{\rho_c}{\rho_j}\right)^{1/2}. \tag{2.37}$$

Here θ_T is the angle the boundary between the turbulent layer and the laminar jet makes with the z axis and ρ_c is the density in the cavity.

In our case we have

$$\rho_c \propto \frac{1}{L_j^{\alpha}} \propto \frac{1}{t^{3\alpha\beta}} \tag{2.38}$$

and

$$\rho_j(z_1) \propto \frac{1}{z_1^2} \propto p_c \propto t^{6\beta - 2 - 3\alpha\beta} \tag{2.39}$$

from equations (2.13), (2.24) and (2.27). Equation (2.32) tells us that $M_j(z_1)$ is a constant and hence we have

$$\theta_T(z_1) \propto t^{2-6\beta} \propto t^{2(2-\alpha)/(5-\alpha)}. \tag{2.40}$$

The entire jet thus becomes turbulent at a distance

$$z_T = z_1 + \frac{r_j(z_1)}{\theta_T(z_1)} = z_1 \left(1 + \frac{\theta}{\theta_T}\right) = z_1 \left[1 + \mathscr{C}\left(\frac{t}{t_c}\right)^{2(2-\alpha)/(5-\alpha)}\right], \tag{2.41}$$

where

$$t_c = l_c \left(\frac{Q}{2P}\right)^{1/2} \tag{2.42}$$

is the characteristic time and $\mathscr{C}(\theta, \gamma)$ is another dimensionless constant. Since $\alpha < 2$ and t_c is short, we can see that $z_T \to z_1$ rather quickly and that the jet thus becomes completely turbulent downstream of the reconfinement shock. Once this happens, it begins to entrain a large amount of gas from the cavity and most of its kinetic energy is dissipated long before it reaches the working surface. What we then have is something that looks very much like an FRI source in that it has no pronounced hot spots near the end of the jet. The M87 jet seems to be a very good example of this if we interpret knot A as being the reconfinement shock (Falle & Wilson 1985).

2.2 Case II ($\alpha \geq 2$)

In the previous case the ratio of the jet density at the working surface to that in the surroundings decreased with time, while if $\alpha > 2$, then the opposite occurs. This means that the bow-shock speed will eventually approach the jet speed and the structure of the working surface should become self-similar with a length scale which increases linearly with time. Note that, as the jet is not reconfined, it cannot become turbulent and is therefore able to transport most of its kinetic energy out to the working surface. Morphologically such a jet would look very much like an FRII source except that the weakness of the terminating shock in the jet means that there would not be much radio emission from the working surface.

However, if such a jet now enters a region in which $\alpha < 2$, then its density would begin to decrease relative to the external density and the jet shock will get stronger. The result would be a very faint unconfined jet ending in an extended radio lobe with a bright hot spot at the working surface, i.e. a classical double. This would, however, be a transient stage since the flow must eventually become like that discussed in Section 2.1.

This case seems to correspond to that first considered by Blandford & Rees (1974) and has subsequently been the subject of a number of numerical simulations (Norman *et al.* 1982; Wilson & Scheuer 1983; Kössl & Müller 1988). As it has already been so intensively studied, there seems to be little point in pursuing it further in this paper.

For $\alpha = 2$, the ratio of densities remains constant and the flow is self-similar with constant advance speed at all times. In that case the terminating shock can be strong enough to generate radio emission and we should get something that looks like an FRII at all times. However, this is a rather special case and is unlikely to be very common.

3 NUMERICAL CALCULATIONS

Although the above analysis is quite useful, it leaves a number of important questions unanswered. It does not tell us much about how the shape of the self-similar cavity is related to the jet opening angle nor does there seem to be any simple way to determine the dimensionless quantities $\mathscr{A}$ and $\mathscr{B}$ that appear in equations (2.26) and (2.27). Apart from that, we

would also like to know what the flow looks like before it becomes self-similar.

As we have already pointed out, previous calculations are only relevant to case II, while we are mainly interested in case I. The difference is that for case I we have to include the entire disturbed region in the computational domain and integrate long enough for the flow to become at least approximately self-similar. This is clearly going to be a very large calculation if performed with adequate numerical resolution and so not only must we use an efficient code, but we will also have to choose our parameters with some care to keep the costs within reasonable bounds.

The details of our numerical algorithm are given in the Appendix, but basically it is a second order Godunov scheme (Godunov 1959) which uses an averaging function to ensure monotonicity. This code has been successfully applied to a number of industrial problems involving the interaction of blast waves and supersonic jets with solid structures and in more than one space dimension seems to be as accurate as PPM (Collela & Woodward 1984) and somewhat faster for the same resolution.

Even with a good code the computational costs are such that we can only consider a single case and even then we have had to make some compromises in our choice of parameters. The time taken for the flow to become self-similar is shortest for a uniform external density and a large initial jet opening angle. We therefore confine ourselves to the uniform density case and use a rather larger opening angle than that found in typical radio jets.

The theory of Section 2 assumes that the flow in the jet is initially conical with some opening angle θ and infinite Mach number. However, since the initial jet radius has to be small compared to the size of the cavity, it is difficult to ensure enough numerical resolution in the jet near the source. The result is that if we tried to input the jet with infinite Mach number and finite opening angle, we would get a very poor representation of a conical flow. It is much better to input the jet with zero opening angle and a finite Mach number. This has much the same effect once the cavity pressure becomes small compared to the initial jet pressure since the jet then expands freely until it begins to reconfine. We shall see that this gives a reasonable approximation to a conical flow at distances greater than a few jet radii from the source.

Ideally we would like the jet mass flux to be as low as possible since equation (2.3) tells us that this leads to a small value of the characteristic length l_c. Unfortunately this is not possible because the jet will not propagate into the domain if its initial density is much smaller than the external density. The simplest thing to do is to set the initial jet density equal to the external density. The Mach number of the jet is then fixed by the choice of opening angle and so the only remaining free parameters are the initial jet radius and the mesh spacing. The initial jet radius has to be smaller than l_c, but we cannot make it much smaller since there must be at least a few mesh points in the jet at input. The mesh spacing should be as small as possible subject to the constraint that the calculation does not become unmanageable before the size of the cavity has become large compared to l_c.

The boundary and initial conditions are fairly obvious. We could consider a one-sided jet, but it is much easier to assume that we have two symmetric jets on either side of the source. If that is so, then we must impose a symmetry condition at the $z=0$ boundary. The other boundary conditions present no difficulty as long as we make sure that no disturbance ever reaches them. The most efficient way to achieve this is to monitor the position of the bounding shock and to add cells to the computational domain as necessary. As we are interested in the behaviour at large times, the initial conditions are of little importance. We therefore started the cal-

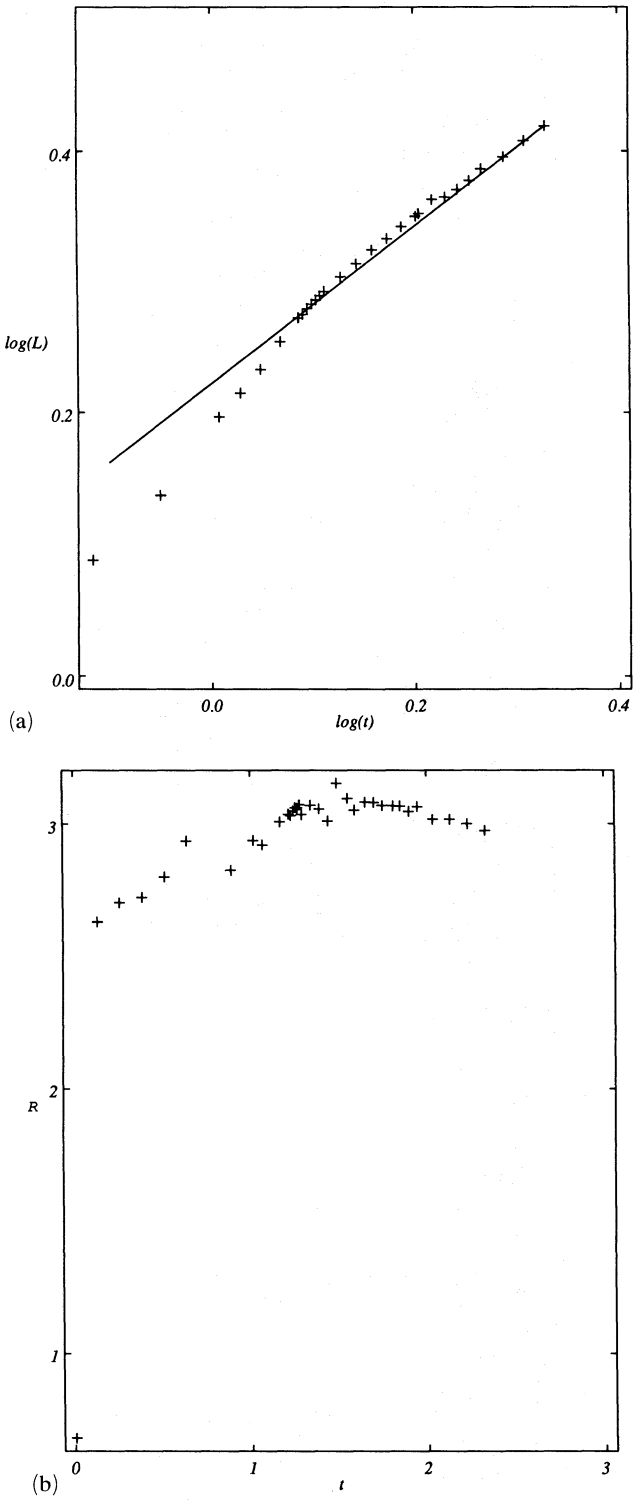

Figure 2. (a) The length of the jet, L, as a function of time. The points (+) are from the simulations and the line has a slope of 3/5. (b) The aspect ratio, $\mathscr{R}$, as a function of time.

culation with the entire domain filled with external gas at rest, switched on the jet and let it do its stuff.

The parameters of our calculation are initial jet radius, 0.02; initial jet Mach number, 5.0; initial jet pressure, 1.0; external density, 1.0; external pressure, 0.001; jet velocity, 6.45; jet mass flux, 8.1×10^{-3}; jet power, 1.87×10^{-1}; and the mesh spacing, 0.005.

These values are a reasonable compromise between what it is possible to compute and the demands of our theory. $l_c=0.04$ and we have been able to continue the calculation until the length of the jet is 2.76, so we ought to have reached the self-similar regime. As there are only four grid points in the jet at input, it is clearly under resolved, but the effective opening angle of about 13° ensures that it has expanded to a decent size before anything interesting happens to it.

The first thing to do is to see whether the disturbed region eventually grows in a self-similar way. Fig. 2(a) shows the length of the jet as a function of time and it can be seen that it does indeed grow like $t^{3/5}$ at large times. The aspect ratio, $\mathscr{R}$, which is plotted in Fig. 2(b), also behaves in the correct way since it approaches a constant value of about 3, which is consistent with the upper limit $\mathscr{R}\leq5.1$ given by (2.36). This is rather encouraging since it tells us that the global behaviour has become self-similar.

However, once we start looking at the details of the flow, we find that things are more complicated than our simple analysis would suggest. Fig. 3 is supposed to correspond to the sketch in Fig. 1. What we have done is to take the solution at the latest time and plot a single density contour at $\rho=1$. Most of the swept up external gas has a density greater than one, so this contour marks the edge of the cavity and the bounding shock. To delineate the edge of the jet, we have also plotted a z velocity contour at half the initial jet velocity. The first thing to notice is that both the edge of the cavity and the outer shock are indented and their shapes are nowhere near as simple as those we have drawn in Fig. 1. This complex structure is caused by vortex rings which are periodically shed from the working surface as it advances. The result is that no part of the flow is really self-similar and it might therefore seem surprising that the overall expansion behaves as it does.

We are by no means the first to come across this vortex shedding. It seems to occur in all high resolution simulations of jets which are less dense than their surroundings and is almost certainly a real effect. As far as the working surface is concerned, the only thing that is different about our calculation is the extent to which the jet has been modified by its passage through the cavity.

One can get some idea of what is happening by looking at the behaviour of isolated vortex rings. Suppose that we have an incompressible flow which consists of a potential vortex ring with radius R and circulation C. Such a vortex ring moves perpendicular to its plane with an induced speed v_i given by

$$v_i=\frac{C}{4\pi R}\left[\ln\left(\frac{R}{R_c}\right)-\frac{1}{4}\right] \tag{3.1}$$

(Fraenkel 1970). Here R_c is characteristic length scale of the vortex core within which all the vorticity is concentrated.

Obviously this formula is not strictly applicable to our case since we do not have a potential vortex and the flow is

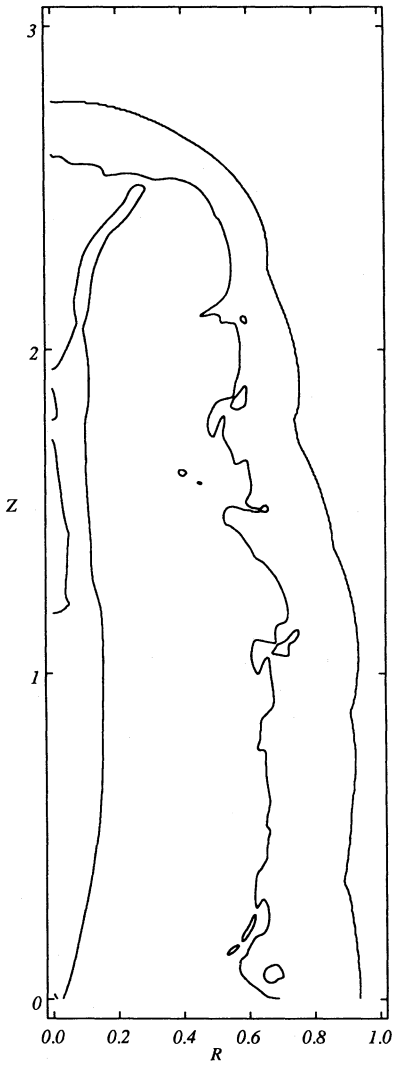

Figure 3. The shape of the cavity and jet at $t=2.3267$ as given by the numerical simulation. We have plotted a single density contour at $\rho=1.0$ and a single z velocity contour at half the initial jet velocity ($v_z=3.2275$).

neither incompressible nor solely due to the presence of the vortex. Despite this it must still be true that the behaviour of the vortex depends crucially on the size of its core and this in turn depends upon how the vortex was formed and the rate at which it is diffused by viscous effects.

The vortex ring at the working surface is formed by the rolling up of the shear layer at the edge of the jet as it impacts with the dense gas at the end of the cavity. If this layer is actually a slip line with zero thickness, then R_c will be zero in the absence of viscosity and v_i should be infinite. In the simulations numerical viscosity ensures that the boundary of the jet is diffused and this keeps v_i finite, but as we increase the numerical resolution we reduce this viscosity and v_i can become very large.

For large enough v_i, the vortex can avoid being swept away by the backflow from the working surface and so it is able to grow as it accumulates more and more vorticity from the jet. The result is an increase in the size of the vortex and a lateral expansion of the end of the cavity. Eventually the

vortex core diffuses to such an extent that the vortex can no longer remain at the end of the jet and it is swept away into the interior of the cavity. Once this happens, a new vortex begins to form and the whole process repeats itself.

If we are right about this, then the time it takes for the vortex to be shed depends upon the effective viscosity. In all the simulations, including ours, the only viscosity is the one introduced by the numerical scheme and this means that it scales with the mesh spacing. We would therefore expect the behaviour of the vortex to depend on both the resolution and the numerical method.

This does in fact seem to be the case. Kössl & Müller (1988) looked at the effect of grid size on the behaviour of a jet whose density at the working surface was less than that of the surroundings. They found that both the intensity of the vortex and the time it spent near the working surface increased with decreasing mesh spacing. This is exactly what we would expect if it is the viscosity that plays the crucial role.

In a real jet it is not likely to be the physical viscosity that matters, but the level of turbulence. The turbulence in the shear layer at the edge of the jet gets swept into the vortex and it is this that determines the rate at which the vortex core diffuses away. Paradoxically this means that the simulations are probably more realistic than they would be if they had incorporated a laminar viscosity and sufficiently high resolution for the numerical viscosity to be unimportant in the vortex.

One might think that this effect means that both our numerical calculations and the analysis of Section 2 are completely useless. Fortunately this is not the case. Although the flow cannot be strictly self-similar. Fig. 2(a) shows that our analysis does describe the global behaviour quite well. We shall also see that the pressure in the cavity is more or less uniform and the jet does actually reconfine in more or less the way that we have predicted.

There seem to be two reasons for this. The first is that, although the vortex shedding dominates the flow near the working surface, it is of little importance near the source. The expansion of the sides of the cavity and the reconfinement of the jet are therefore not too seriously affected by the vortices. We can also show that global effect of the vortex shedding becomes less and less important at large times.

Suppose that the vortex is characterized by a length scale l_v and time-scale t_v. Then if diffusion determines the evolution we must have

$$t_v \propto \frac{l_v^2}{\nu}, \tag{3.2}$$

where ν is the dynamic viscosity. We might expect l_v to behave like the jet radius at the working surface since that is what determines the distribution of vorticity of the gas entering the vortex. Ram pressure balance demands that this increases like $t^{2/5}$ for a constant external density and we therefore have

$$t_v \propto t^{4/5}. \tag{3.3}$$

This means that the time-scale for vortex shedding eventually becomes small compared with the dynamical time t for the flow as a whole. It is therefore not too surprising that the large scale structure is approximately self-similar. The only bad news is that the ratio of t_v to the dynamical time only decreases like $t^{-1/5}$ and it therefore takes a long time before the vortex shedding ceases to affect the overall dynamics.

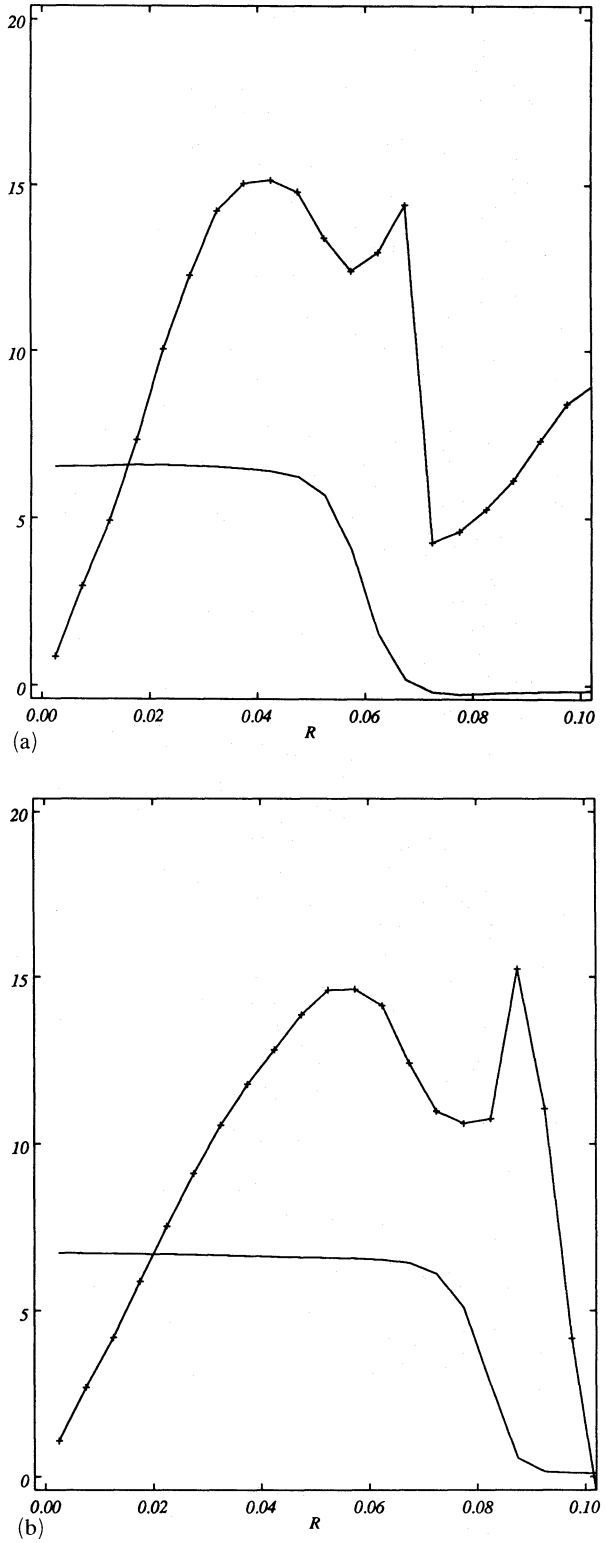

Figure 4. (a) The flow angle ϕ (+ markers) and z velocity (no markers) in the jet at $t=2.3267$ and $z=0.1$ (five times the initial jet radius). (b) The flow angle ϕ (+ markers) and z velocity (no markers) in the jet at $t=2.3267$ and $z=0.2$ (10 times the initial jet radius).

Having established that the overall expansion behaves as expected and that the vortices should be weak near the source, we ought now to look at the reconfinement process to see whether the assumptions of Section 2 are correct. In Fig. 4 we have plotted the z velocity and the angle ϕ that the flow makes with the z axis at 5 and 10 jet radii from the source. The sharp decrease in the z velocity marks the jet radius and we can see that ϕ increases linearly with radius for most of the jet. This means that the flow in the jet is approximately conical before it is affected by the reconfinement shock. Our jet is therefore effectively conical even though it was input with finite radius and Mach number.

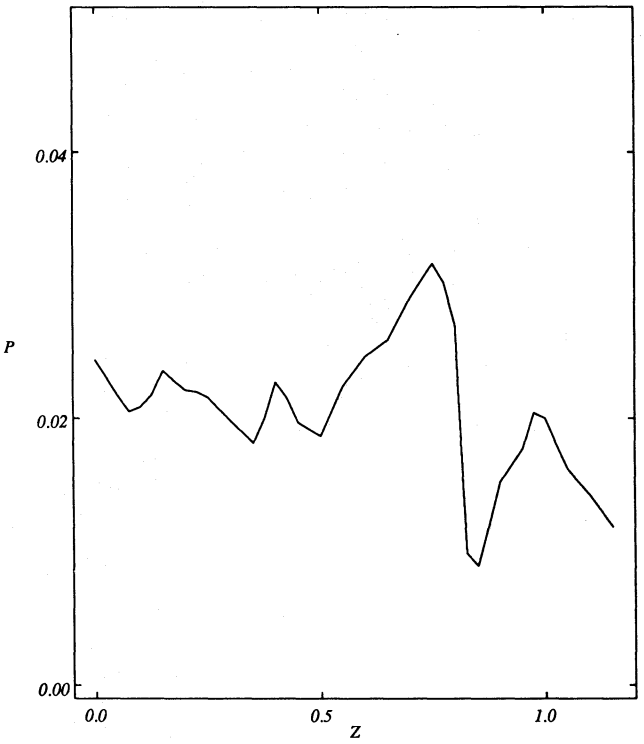

Figure 5. The pressure at the jet boundary as a function of z at $t=2.3267$.

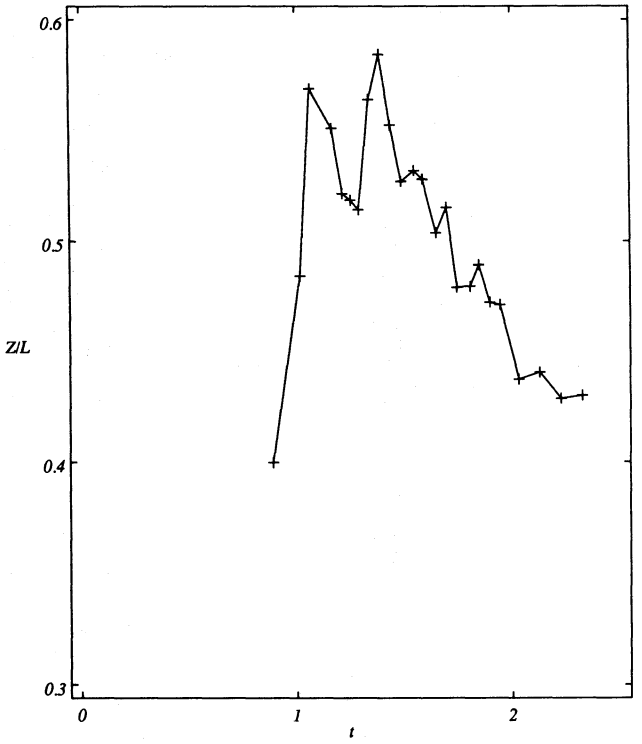

Figure 6. The ratio of the distance to the reconfinement shock, Z, to the length L of the jet as a function of time.

Next we look at the pressure at the jet boundary, which should be equal to the uniform cavity pressure p_c. Fig. 5 shows that, although this pressure is reasonably uniform for $z<0.5$, it varies quite substantially before we reach the reconfinement shock at $z=z_1=1.18$. These pressure variations are due to the vortices in the cavity which have not entirely decayed at this distance. This is simply because we have not continued the calculation for long enough for the distance to the reconfinement shock to become small compared with the jet length L_j. According to our analysis, z_1/L_j should decrease like $1/t^{1/5}$ and we will therefore have to wait a very long time for this ratio to become small. Fig. 6 does, however, show that it is decreasing, albeit very slowly.

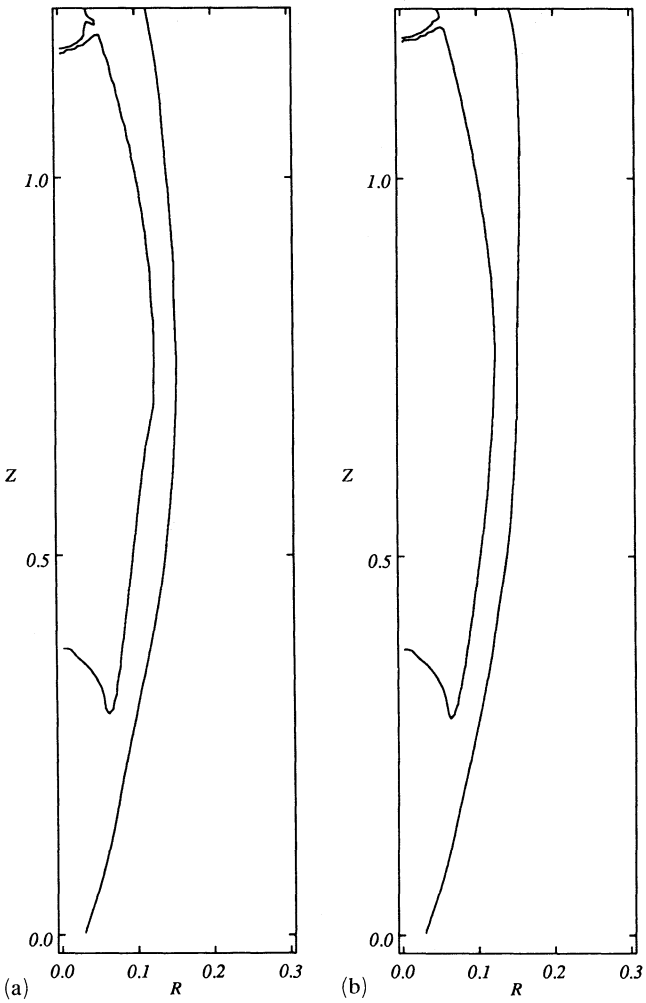

Figure 7. (a) The shape of the jet boundary and reconfinement shock at $t=2.1316$. The jet boundary is delineated by a single contour of z velocity at half the initial jet velocity ($v_z=3.227$). A single pressure contour at $p=0.004$ marks the reconfinement shock for $z>0.3$. (b) The shape of the jet boundary and reconfinement shock at $t=2.3267$. The jet boundary is delineated by a single contour of z velocity at half the initial jet velocity ($v_z=3.227$). A single pressure contour at $p=0.004$ marks the reconfinement shock for $z>0{\cdot}.3$.

Given the size of these pressure variations, we cannot expect the shape of the reconfinement shock to be accurately given by equation (2.23). Fig. 7 shows the shape of the jet boundary and the reconfinement shock at two different times, and we can see that the shock approaches the z axis at a somewhat steeper angle than θ and furthermore there is a Mach disc whose radius is nearly a third that of the jet at this point. This Mach disc is caused by a sharp increase in the cavity pressure downstream of z_1.

It would be nice to get closer agreement between our analytic expressions and the numerical simulation, but this is not really possible. The problem is that the flow approaches self-similarity like $t^{1/5}$, i.e. very slowly and we would therefore have to continue the calculation for a ridiculously long time to get accurate agreement. However, the qualitative features are correct, and the numerical solution does seem to be moving in the right direction.

In a certain sense it does not matter too much that the cavity pressure is not uniform all the way out to z_1, as long as the flow in the jet is quasi-steady up to this point. We can see that this is approximately true by comparing Fig. 7(a) and (b) which show the reconfinement process at $t=2.13$ and $t=2.33$. It takes the jet fluid a time z_1/v_j to reach z_1 and this turns out to be 0.18. In Section 2 we have argued that the reconfinement is quasi-steady if the flow does not change much in a time z_1/v_j and we can see that this is indeed true for $z<1$. So although our simple theory is a bit inadequate, we can always deal with the reconfinement by treating the jet as steady in an external pressure which, although it varies spatially, is aproximately uniform in time. This sort of thing has been done for the M87 jet (Falle & Wilson 1985) and NGC 6251 (Wilson 1986).

So far we have been concerned with the self-similar regime, but the computations also tell us about the flow at earlier times. Fig. 8 shows the pressure, density and velocity at $t=1.75$, at which time it looks quite different from that at $t=2.32$ (Fig. 9). The most important distinction is that the reconfinement shock is much stronger at the earlier time, so strong that it completely disrupts the jet. The highest pressure occurs downstream of the reconfinement shock and the working surface is not particularly prominent. Interestingly, the largest pressure at the end of the jet occurs off axis, so, if the emission correlates with pressure, we would see a ring of enhanced radio emission.

At early times the vortex shedding is much more pronounced and this leads to dramatic variations in the flow on

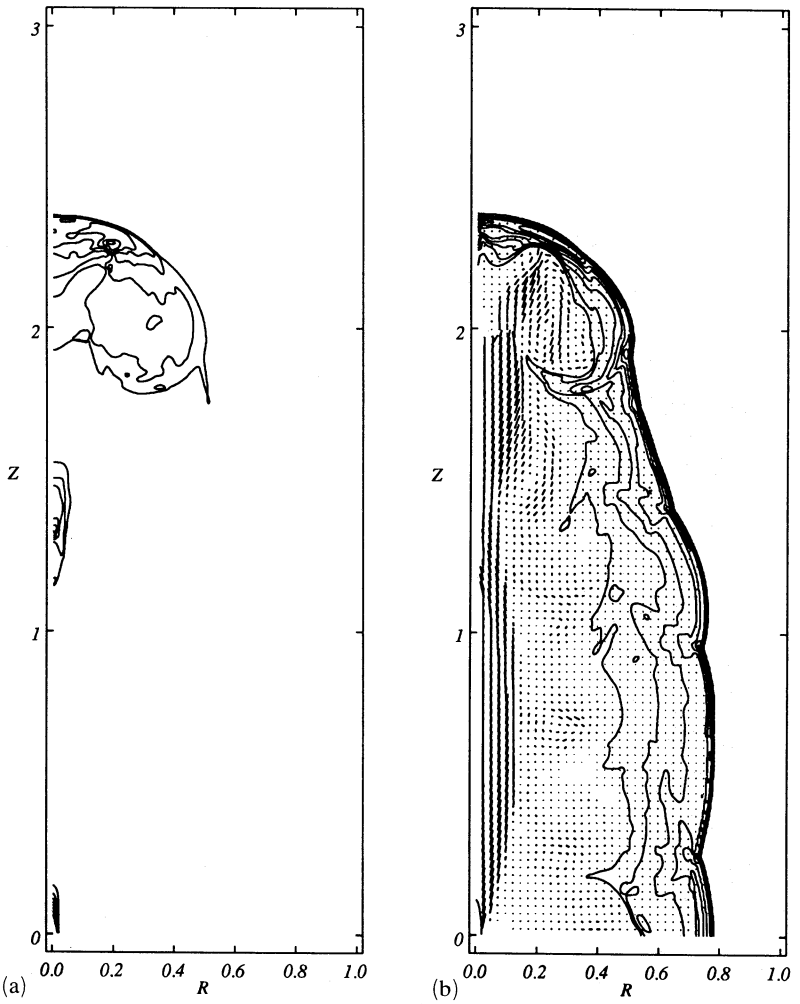

Figure 8. (a) Pressure contours at $t=1.7484$. There are 10 linear contours equally spaced between 0.0 and 1.0. (b) Density contours and velocity vectors at $t=1.7484$. There are 10 linear contours of density equally spaced between 0.0 and 4.645. The maximum velocity is 6.85 and occurs just upstream of the reconfinement shock at $z=1.12$.

the vortex time-scale. As this time-scale is rather short, it makes comparison with observations extremely difficult since small differences in age can lead to large differences in morphology. Fortunately the self-similar flow is much simpler and Fig. 9 shows that it corresponds very well with the morphology of classical doubles. The jet is not too seriously damaged by reconfinement and the structure of the working surface is much the same as in previous numerical simulations.

It seems that we can divide the evolution into three distinct phases. At early times we get strong vortex shedding, a rapidly changing flow and strong hotspots recessed from the end of the jet. Once the flow becomes self-similar, it resembles an FRII source in that the jet is relatively faint with a pronounced hotspot at its end. Later still the jet becomes turbulent downstream of the reconfinement shock and the hotspot at the working surface disappears, leaving us with an FRI source. Although we have only studied the constant density case in detail, the arguments of Section 2 suggest that the flow should be qualitatively similar for any value of α less than 2. The only difference is that the evolution time increases with increasing α.

4 OBSERVATIONAL COMPARISONS

The problem with all comparisons between flow calculations and observations of radio jets is that, despite numerous attempts (e.g. Rayburn 1977; Wilson & Scheuer 1983; Smith *et al.* 1985; Coleman & Bicknell 1987; Heavens & Meisenheimer 1987; Kössl & Hillebrandt 1990c), we do not yet have a reliable way of predicting the radio emission from a given flow. We know that the emission is produced by the synchrotron process and that there must therefore be some mechanism which transfers flow energy to the relativistic electrons and magnetic field, but how this happens is not at all clear. Indeed, we cannot even be sure that the jets are not composed of some exotic material such as relativistic pair plasma which may not behave like a fluid.

Whatever the jets are made of, it seems fairly certain that the relativistic electrons are not accelerated once and for all in the source and then transported to the radio lobes (Begelman, Blandford & Rees 1984). The reasons for this are basically that the source and jet typically have different spectra and the synchrotron lifetimes are too short for the electrons to have retained their energy during their journey to the

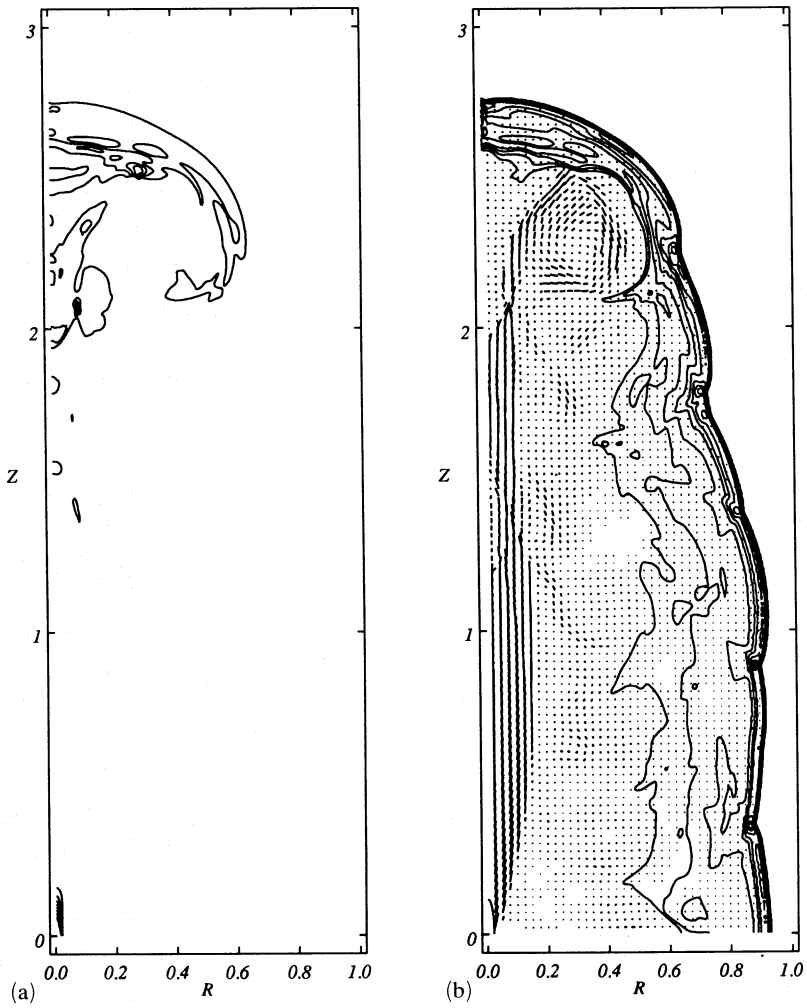

Figure 9. (a) Pressure contours at $t = 2.3267$. There are 10 linear contours equally spaced between 0.0 and 1.0. (b) Density contours and velocity vectors at $t = 2.3267$. There are 10 linear contours of density equally spaced between 0.0 and 4.79. The maximum velocity is 6.86 and occurs just upstream of the reconfinement shock at $z = 1.18$.

radio lobes. It follows that the electrons must be accelerated in the jet flow itself and the most likely candidate for this is some kind of Fermi process. Shocks are a particularly attractive location for this since they can drive the very efficient first-order process that has also been invoked to explain the acceleration of the galactic cosmic rays (Drury 1983). This idea receives some observational support from studies of synchrotron ageing in classical doubles which suggest that the bulk of the acceleration occurs at the working surface where the shocks are dissipating the most energy (Alexander & Leahy 1987).

It is also possible that the turbulence associated with the jet boundary and the vortices could feed a second order Fermi process (Coleman & Bicknell 1987). We have argued in Section 2 that for case I the jet eventually becomes turbulent downstream of the reconfinement shock and subsequent shocks become much weaker. Turbulent acceleration would then become the dominant mechanism and its lower efficiency provides a natural explanation for the lower radio power of FRI sources.

Intriguing as these possibilities are, we do not really want to get too involved with the details of the emission mechanism. Our primary interest is in the large-scale structure, and for that we do not need an elaborate emission model. All we need to assume is that the observed jet boundary more or less coincides with the real one and that shocks and high-pressure regions mean enhanced radio emission.

That given, we should look for radio sources whose morphology is consistent with our calculations. We have already seen that at very large times we expect to get FRI sources, but our detailed calculations only apply to early and intermediate times. The self-similar flow at intermediate times resembles an FRII source and this is rather gratifying since these are extremely common. We would expect sources in the earlier stages of evolution to be rarer, but it is obviously interesting to see if we can find any that fit the bill.

If we look through Leahy & Williams' (1984) list of 39 3C sources, we do find a number with recessed hotspots. The best candidate is 3C134 which has quite prominent recessed hotspots on both sides in addition to the main ones at the ends of the jets. The ratio of the distances to the inner hotspots to the jet length is 0.36 on one side and 0.6 on the other and is thus in the range we would expect at early times. The aspect ratio is hard to estimate, but it seems to be about 0.3. That this agrees almost exactly with our numerical calculation must be pure coincidence since the strength of the recessed hotspots suggests that the flow is not self-similar. It is much more likely that the initial opening angle of the 3C134 jet is much smaller than that in our calculation and that the aspect ratio is still increasing. This is entirely consistent with it still being fairly young.

In the other sources the recessed hotspots are somewhat less pronounced and their positions are harder to estimate, but the trend seems to be what we expect. 3C52 has fairly bright features on each side at 0.66 and 0.4 of the distance to the lobes, while in 3C103 they are much weaker and the ratios are 0.32 and 0.17. In 3C103 there are actually two recessed hotspots on one side and we have taken the distance to the inner weaker one. While this is hardly overwhelming evidence, it is encouraging since our theory says that the reconfinement should be less dramatic the smaller the relative distance to the reconfinement shock.

Of the other sources, 3C139, 3C234 and 3C244 all have recessed hotspots on one or both sides and we could claim that they look like our computations at one stage or another. Whether or not this means anything, only time can tell, but it does at least suggest that this kind of purely gas-dynamical model is still worth pursuing.

We have made much of the presence of strong recessed hotspots in these sources, and so it is worth asking whether there is any other way to produce such features. What we need is something that causes a strong shock in the jet which dissipates a significant fraction of the kinetic energy. Norman, Burns & Sulkanen (1988) have considered a jet propagating in a galactic wind and have shown that the jet can be disrupted if it encounters the terminating shock in the wind. They were particularly interested in flaring of jets and so their shock needed to be strong enough to make the jet flow subsonic. We do not need such an extreme event if all we want is recessed hotspots, but the basic idea is the same.

It seems quite likely that this sort of thing does happen, but it is not clear that it is relevant to the sources we have discussed. Pressure variations in the external medium can only become important at late times when the pressure in the cavity has become comparable to the external pressure and the bounding shock has largely disappeared. None of the sources we have considered look like this. The jets are relatively short and the radio maps do seem to suggest the presence of a distinct cavity, so it does not seem very likely that the recessed hotspots are due to something special in the external medium.

5 CONCLUSIONS

Our intention in this paper has been to study the large-scale structure of supersonic jets in an external medium of varying density. One obvious advantage of this approach is that, for most sources, the only thing the radio maps tell us anything about is the large-scale structure. It also means that we can dispense with the kind of detailed model of the emission that would be necessary if we were to try and understand small-scale features.

What we have found is really quite encouraging. The behaviour of the flow seems to depend crucially on how rapidly the external density decreases with distance from the source. If this decrease is more rapid than $1/R^2$, where R is the distance from the source, then the jet behaves as if it were exhausting into a vacuum and ought to be completely invisible. Only if it at some point enters a region where the density decreases more slowly does it become visible and it will then look like a classical double. However, this case is probably very rare since there seem to be few galaxies in which the density decreases as fast as this over an extended region.

An external density which decreases slower than $1/R^2$ is much more likely and in any case more interesting. The evolution can be divided into three distinct stages with quite different morphologies. At early times we get a very complicated flow dominated by periodic vortex shedding and strong shocks in the jet before it reaches the working surface. In this context strong means that the post-shock pressure is comparable to the jet ram pressure and these shocks can at times disrupt the jet completely. The result is a source with recessed hotspots which can actually be stronger than those

at the working surface. We have been able to find a number of sources which fit this picture and there are no doubt many more, but the flow varies so rapidly with time that it is difficult to make any detailed comparisons with the observations.

This violent stage eventually gives way to a sort of self-similar solution in which the size of the disturbed region grows as a power of time in a power-law external density. The jet is reconfined at some distance form the source, but this does not lead to much energy dissipation and there should therefore be a prominent hotspot at the working surface. Such a flow looks like a clasical double with a faint jet and edge-brightened radio lobes. There is still some vortex shedding to give an interesting structure to the lobes, but its importance is on the wane.

At very large times the reduction in opening angle and Mach number brought about by the reconfinement shock allows the jet to become turbulent. This reduces the ram pressure in the jet and hence the strength of the hotspots at the end of the jet. Although the cavity blown by the jet still expands according to the same formula as before, the morphology is completely different and resembles that of an FRI source. The lower radio power of such sources would then be a consequence of the fact that particle acceleration by turbulence is a much less efficient process than shock acceleration.

The point we are trying to make is that the jets propagate in an environment which is largely of their own making. This is an old idea and was first explored in detail by Scheuer (1974), but more recent numerical work has tended to somewhat neglect this aspect. What we have shown is that quite reasonable assumptions give us a picture which seems to be consistent with the observations.

Gratifying as this is, there is still much to be done. We have not really been able to follow the flow far enough into the self-similar regime, nor have we looked in detail at the dynamical effects of turbulence. This is not due to lack of computing resources, but to the fact that the flow becomes self-similar so slowly that any computation on a fixed mesh would exhaust the power of current machines without getting much further than we have. What is needed is an adaptive algorithm which automatically refines the mesh in complicated regions and uses lower resolution in the bland parts of the flow. Such a code is currently under development for industrial applications and we will soon be able to apply it to this problem. However, this is not enough by itself. We have argued that turbulence plays an important role even in the early stages and we must therefore include some kind of turbulence model in our calculations. There are a number of possibilities, none of them entirely satisfactory, but it does seem that a modified version of the $k-\varepsilon$ model works quite well for the shear layer at the edge of supersonic jets (Dash & Wolf 1983). It is therefore a good idea to incorporate this into future calculations.

ACKNOWLEDGMENTS

This work benefited from numerous conversations with J. R. Giddings and M. J. Wilson. I would also like to thank the referee for several useful suggestions. The computations were carried out on a Sun 4/110 at Mantis Numerics Ltd.

REFERENCES

Alexander, P. & Leahy, J. P., 1987. *Mon. Not. R. astr. Soc.*, **225,** 1.
Baldwin, J. E., 1982. *IAU Symp. No. 97,* Albuquerque, USA, p. 21, eds Heeschen, D. S. & Wade, C. M., D. Reidel, Dordrecht.
Begelman, M. C., Blandford, R. D. & Rees, M. J., 1984. *Rev. mod. Phys.*, **56,** 255.
Blandford, R. D. & Rees, M. J., 1974. *Mon. Not. R. astr. Soc.*, **169,** 395.
Boris, J. P. & Book, D. L., 1973. *J. Comp. Phys.*, **11,** 38.
Canto, J. & Raga, A. C., 1990. *Astrophys. J.*, submitted.
Canto, J., Raga, A. C. & Binette, L., 1989. *Rev. Mex. Astr. Astrofis.*, **17,** 65.
Chakrabarti, S. K., 1988. *Mon. Not. R. astr. Soc.*, **235,** 33.
Collela, P. & Woodward, P. R., 1984. *J. Comput. Phys.*, **54,** 174.
Coleman, C. S. & Bicknell, G. V., 1987. *Mon. Not. R. astr. Soc.*, **230,** 497.
Daly, R. A. & Marscher, A. P., 1988. *Astrophys. J.*, **334,** 539.
Dash, S. M. & Wolf, D. E., 1983. *AIAA Paper* 83-0704.
Drury, L. O'C., 1983. *Rep. Prog. Phys.*, **46,** 973.
Dyson, J. E., Falle, S. A. E. G. & Perry, J. J., 1980. *Mon. Not. R. astr. Soc.*, **191,** 785.
Falle, S. A. E. G., 1988. *Supernova Shells and their Birth Events,* p. 63, ed. Kundt, W., Springer.
Falle, S. A. E. G. & Wilson, M. J., 1985. *Mon. Not. R. astr. Soc.*, **216,** 79.
Falle, S. A. E. G. & Wilson, M. J., 1988. *Numerical Methods for Fluid Dynamics III,* p. 418, eds Morton, K. W. & Baines, M. J., Clarendon Press, Oxford.
Federenko, V. N. & Zentsova, A. S., 1986. *Soviet Astr.*, **30,** 24.
Fraenkel, L. E., 1970. *Proc. R. Soc. London A,* **316,** 29.
Godunov, S. K., 1959. *Mat. Sb.*, **47,** 357.
Heavens, A. F. & Meisenheimer, K., 1987. *Mon. Not. R. astr. Soc.*, **225,** 335.
Kössl, D., Müller, E. & Hillebrandt, W., 1988. *Astr. Astrophys.*, **206,** 204.
Kössl, D., Müller, E. & Hillebrandt, W., 1990a. *Astr. Astrophys.*, **229,** 378.
Kössl, D., Müller, E. & Hillebrandt, W., 1990b. *Astr. Astrophys.*, **229,** 401.
Kössl, D., Müller, E. & Hillebrandt, W., 1990c. *Astr. Astrophys.*, **229,** 397.
Leahy, J. P. & Williams, A. G., 1984. *Mon. Not. R. astr. Soc.*, **210,** 929.
Norman, M. L., Burns, J. O. & Sulkanen, M. E., 1988. *Nature,* **335,** 146.
Norman, M. L., Smarr, L., Winkler, K. H. A. & Smith, M. D., 1982. *Astr. Astrophys.*, **113,** 285.
Osher, S., 1982. *Numerical Methods for Fluid Dynamics,* p. 179, eds Morton, K. W. & Baines, M. J., Academic Press, Oxford.
Rayburn, D. R., 1977. *Mon. Not. R. astr. Soc.*, **179,** 603.
Roe, P. L., 1981. *J. Comp. Phys.*, **43,** 357.
Roe, P. L., 1985. *Lectures in Applied Mathematics,* **22,** 163.
Roe, P. L., 1986. *Ann. Rev. Fluid Mech.*, **18,** 337.
Sanders, R. H., 1983. *Astrophys. J.*, **266,** 73.
Scheuer, P. A. G., 1974. *Mon. Not. R. astr. Soc.*, **166,** 513.
Smith, M. D., Norman, M. L., Winkler, K. H. A. & Smarr, L., 1985. *Mon. Not. R. astr. Soc.*, **214,** 67.
van Leer, B., 1982. *Lect. Notes in Phys.*, **170,** 507.
van Leer, B., 1977. *J. Comp. Phys.*, **23,** 276.
Wilson, M. J., 1986. *Mon. Not. R. astr. Soc.*, **224,** 155.
Wilson, M. J., 1987. *Mon. Not. R. astr. Soc.*, **226,** 447.
Wilson, M. J. & Falle, S. A. E. G., 1985. *Mon. Not. R. astr. Soc.*, **216,** 971.
Wilson, M. J. & Scheuer, P. A. G., 1983. *Mon. Not. R. astr. Soc.*, **205,** 449.

APPENDIX

Most modern shock capturing schemes for the Euler equations are both upwind and conservative. By upwind we mean that it takes account of all the characteristics, not just the streamlines, and by conservative we mean that conservation is satisfied to machine accuracy. These two properties are important because only for an upwind scheme is it possible to impose the correct boundary conditions, and conservation is necessary to guarantee the correct jump conditions at shocks.

In axisymmetry strict conservation is impossible since then we have to work with the radial momentum and this is not a conserved quantity. However, by using a finite-volume formulation, it is possible to construct an algorithm which becomes conservative at shocks as the mesh spacing tends to zero. Of course, we cannot actually go to the limit of zero mesh spacing, but this only causes problems near the axis of symmetry and these can be minimized by ensuring that the flow is well resolved there.

A scheme can be described as upwind if it is equivalent to a characteristic scheme for the linear problem. This can be achieved in a number of ways, but the simplest, and the one that we shall adopt, is that due to Godunov (Godunov 1959). This involves computing the flux of the conserved quantities from the solution to Riemann problems at the interfaces between the computational cells. It has subsequently been shown that the same effect can be obtained in other ways and this has led to the development of a number of schemes all of which can be described as characteristic based. Examples are flux-vector splitting (van Leer 1982), Roe's scheme (Roe 1981) and Osher's scheme (Osher 1982). These all use an approximate solution to the Riemann problem which is exact for the linear problem and certain other cases. Those interested in this subject should consult the excellent reviews by Roe (Roe 1985, 1986).

In axisymmetry the Euler equations can be written in the form

$$\frac{\partial \boldsymbol{U}}{\partial t}+\frac{1}{r}\frac{\partial r\boldsymbol{F}}{\partial r}+\frac{\partial \boldsymbol{G}}{\partial z}=\boldsymbol{S}, \tag{A1}$$

where

$$\boldsymbol{U}=(\rho, \rho v_r, \rho v_z, e) \tag{A2}$$

is a vector of conserved quantities and

$$\boldsymbol{F}=[\rho v_r, p+\rho v_r^2, \rho v_r v_z, v_r(e+p)], \tag{A3}$$

$$\boldsymbol{G}=[\rho v_z, \rho v_r v_z, p+\rho v_z^2, v_z(e+p)], \tag{A4}$$

are the fluxes in the r and z directions. Here p, ρ are the pressure and density and e is the total energy per unit volume

$$e=\frac{p}{(\gamma-1)}+\frac{1}{2}\rho(v_r^2+v_z^2). \tag{A5}$$

The source term is given by

$$\boldsymbol{S}=\left[0, \frac{p}{r}, 0, 0\right] \tag{A6}$$

and accounts for the fact that the radial momentum is not conserved.

To construct a scheme for equations (A1), we divide the computational domain into cells with mesh spacing h such that the i, j cell occupies the region $(i-1)h \le r \le ih$, $(j-1)h \le z \le jh$, whose volume is

$$V_{ij}=\pi(2i-1)h^3. \tag{A7}$$

Now suppose that we know the solution at some time $t=t_k$ and we want to calculate that at a later time t_{k+1}. We integrate (A1) over the i, j cell and from $t=t_k$ to $t=t_{k+1}$ to get

$$\frac{\boldsymbol{U}_{ijk+1}-\boldsymbol{U}_{ijk}}{t_{k+1}-t_k}+\frac{2}{(2i-1)h}(i\boldsymbol{F}_{i+1/2jk+1/2}-(i-1)\boldsymbol{F}_{i-1/2jk+1/2})$$

$$+\frac{1}{h}(\boldsymbol{G}_{ij+1/2k+1/2}-\boldsymbol{G}_{ij-1/2k+1/2})=\boldsymbol{S}_{ijk+1/2}. \tag{A8}$$

Here

$$\boldsymbol{U}_{ijk}=\frac{2\pi}{V_{ij}}\int_{(j-1)h}^{jh}\int_{(i-1)h}^{ih}\boldsymbol{U}(r, z, t_k)\, r\, dr\, dz, \tag{A9}$$

is the mean value of $\boldsymbol{U}$ in the i, j cell at time t.

$$\boldsymbol{F}_{i+1/2jk+1/2}=\frac{1}{(t_{k+1}-t_k)h}\int_{t_k}^{t_{k+1}}\int_{(j-1)h}^{jh}\boldsymbol{F}(ih, z, t)\, dz\, dt, \tag{A10}$$

$$\boldsymbol{G}_{ijj+1/2k+1/2}=\frac{2}{(t_{k+1}-t_k)(2i-1)h}\int_{t_k}^{t_{k+1}}\int_{(i-1)h}^{ih}\boldsymbol{G}(r, jh, t)\, r\, dr\, dt, \tag{A11}$$

are the fluxes averaged over time and the cell interfaces and

$$\boldsymbol{S}_{ijk+1/2}=\frac{2\pi}{(t_{k+1}-t_k)V_{ij}}\int_{t_k}^{t_{k+1}}\int_{(j-1)h}^{jh}\int_{(i-1)h}^{ih}\boldsymbol{S}(r, z, t)\, r\, dr\, dz\, dt, \tag{A12}$$

is the source term averaged over time and the volume of the cell.

Equation (A8) is exact and forms the basis of all conservative schemes for the axisymmetric Euler equations. It is the way in which the quantities with half-integral suffices are approximated that distinguishes one such scheme from another.

In a first order Godunov type scheme the approximations to the fluxes and source term are derived by assuming that the solution is uniform within each cell and constant over a time-step. The fluxes are determined by ignoring the source terms and solving a one-dimensional Riemann problem at the cell interfaces. A Riemann problem for a hyperbolic system of non-linear conservation laws is defined to be an initial value problem with discontinuous initial data of the form

$$\boldsymbol{U}(x)=\boldsymbol{U}_l=\text{const} \quad \text{for } x\le 0,$$

$$\boldsymbol{U}(x)=\boldsymbol{U}_r=\text{const} \quad \text{for } x>0 \text{ at } t=0.$$

For the Euler equations, the solution for such initial data will in general involve shocks, centred rarefaction waves and contact discontinuities, but it will always be true that the value of $\boldsymbol{U}$ (the resolved state) at $x=0$ is constant for $t>0$, although not in general equal to either $\boldsymbol{U}_l$ or $\boldsymbol{U}_r$. The resolved state is given by a set of non-linear algebraic equations which have to be solved iteratively, but this can be done very efficiently (van Leer 1977). In any case the discontinuity in $\boldsymbol{U}$ is

$O(h)$ in smooth regions so that most of the Riemann problems that occur can be linearized.

If we let $\boldsymbol{U}_*(\boldsymbol{U}_l, \boldsymbol{U}_r)$ be the resolved state, then in a first-order Godunov scheme the fluxes are given by

$$\boldsymbol{F}_{i+1/2jk+1/2} = \boldsymbol{F}[\boldsymbol{U}_*(\boldsymbol{U}_{ijk}, \boldsymbol{U}_{i+1jk})],$$
$$\boldsymbol{G}_{ij+1/2k+1/2} = \boldsymbol{G}[\boldsymbol{U}_*(\boldsymbol{U}_{ijk}, \boldsymbol{U}_{ij+1k})], \tag{A13}$$

i.e. the left and right states are those in the cells on either side of the interface and the flux is that in the corresponding resolved state.

The source term has to be treated with some care. It is no good assuming that $\boldsymbol{S}$ is uniform in each cell since it depends upon the radial coordinate. Instead we have to evaluate $\boldsymbol{S}$ by setting the pressure constant and then integrating over the cell. The result for the non-zero component of $\boldsymbol{S}$ is

$$\left(\frac{p}{r}\right)_{ijk+1/2} = \frac{2\pi}{(t_{k+1}-t_k)\,V_{ij}}\, p_{ijk} \int_{t_k}^{t_{k+1}} \int_{(j-1)h}^{jh} \int_{(i-1)h}^{ih} dr\, dz\, dt$$
$$= \frac{2p_{ijk}}{(2i-1)\,h}. \tag{A14}$$

This form of $\boldsymbol{S}$ ensures that a state with uniform pressure and zero velocity is an exact steady solution of the difference equations.

The above is a straightforward extension of Godunov's scheme to axisymmetry. To make this second order in time and space we have to introduce some structure inside the cells and allow for the variation in $\boldsymbol{U}$ over the time step. Second order time accuracy can be achieved by using the first order scheme to obtain an intermediate solution $\boldsymbol{U}_{ijk+1/2}$ at the half time $t=(t_{k+1}+t_k)/2$, which is then used to determine the second-order fluxes.

Second-order space accuracy requires a knowledge of the gradients within each cell, but we cannot just use central differences for this since Godunov's theorem (Godunov 1959) warns us that such a scheme will not be monotonic in the neighbourhood of shocks and tangential discontinuities. Instead we have to use a non-linear switch which has the effect of reducing the order of the scheme in such regions. Broadly speaking there are two approaches to this. The first is flux corrected transport, which uses a flux limiter to modify the second-order flux so that it does not cause oscillations in regions where the second derivatives are large (Boris & Book 1973). Although this is a sound idea, it is very complicated in more than one dimension and has a number of unpleasant side effects. The alternative is to apply a non-linear averaging function to the gradients which has the effect of taking the smallest of the two gradients on either side of the cell (van Leer 1977).

We shall adopt van Leer's approach except that it is better to work with the vector of primitive variables.

$$\boldsymbol{P}=(\rho, v_r, v_z, p) \tag{A15}$$

The gradients in these variables is calculated from

$$\left(\frac{\partial \boldsymbol{P}}{\partial r}\right)_{ijk+1/2} = \mathrm{av}\left(\frac{\boldsymbol{P}_{ijk+1/2}-\boldsymbol{P}_{i-1jk+1/2}}{r_g(i)-r_g(i-1)}, \frac{\boldsymbol{P}_{i+1jk+1/2}-\boldsymbol{P}_{ijk+1/2}}{r_g(i+1)-r_g(i)}\right),$$
$$\left(\frac{\partial \boldsymbol{P}}{\partial z}\right)_{ijk+1/2} = \mathrm{av}\left(\frac{\boldsymbol{P}_{ijk+1/2}-\boldsymbol{P}_{ij-1k+1/2}}{h}, \frac{\boldsymbol{P}_{ij+1k+1/2}-\boldsymbol{P}_{ijk+1/2}}{h}\right), \tag{A16}$$

where $\mathrm{av}(a, b)$ is the averaging function and $r_g(i)$ is the radial coordinate of centre of gravity of the cells with r-index i,

$$r_g(i) = \frac{2}{3}\frac{(3i^2-3i+1)}{(2i-1)}\,\mathrm{h}. \tag{A17}$$

The averaging function has to have the following properties

$$\mathrm{av}(a, b) \to \tfrac{1}{2}(a+b) \quad \text{as } a \to b,$$
$$= 0 \quad \text{if } ab < 0,$$
$$\to a \quad \text{as } |b|/|a| \to \infty,$$
$$\to b \quad \text{as } |a|/|b| \to \infty.$$

There are obviously an infinite number of functions with these properties, but our attitude is that the simplest is the best. We therefore use

$$\mathrm{av}(a, b) = \frac{(a^2 b + ab^2)}{(a^2+b^2)} \quad \text{if } a^2+b^2 > 0,$$
$$= 0 \quad \text{if } ab < 0, \tag{A18}$$

which clearly has the desired behaviour.

Once we have got suitable gradients in each cell, we can set up the left and right states for the Riemann problems. For the radial fluxes we have

$$\boldsymbol{P}_l = \boldsymbol{P}_{ijk+1/2} + [ih - r_g(i)]\left(\frac{\partial \boldsymbol{P}}{\partial r}\right)_{ijk+1/2},$$
$$\boldsymbol{P}_r = \boldsymbol{P}_{i+1jk+1/2} + [ih - r_g(i+1)]\left(\frac{\partial \boldsymbol{P}}{\partial r}\right)_{i+1jk+1/2}, \tag{A19}$$
$$\boldsymbol{F}_{i+1/2jk+1/2} = \boldsymbol{F}[\boldsymbol{U}_*\{\boldsymbol{U}(\boldsymbol{P}_l), \boldsymbol{U}(\boldsymbol{P}_r)\}],$$

and similarly for the axial fluxes.

Having dealt with the second-order fluxes, we must now derive an expression for the source term. All we have to do is to use our linear pressure distribution to evaluate the integral in equation (A12). The result is

$$\left(\frac{p}{r}\right)_{ijk+1/2} = \frac{2}{(2i-1)\,h}\left\{p_{ijk+1/2} + \left[h\left(i-\frac{1}{2}\right) - r_g(i)\right]\left(\frac{\partial \boldsymbol{P}}{\partial r}\right)_{ijk+1/2}\right\}. \tag{A20}$$

Once we have calculated the second-order fluxes and source term in this way, we can use them in equation (A8) to advance the solution through the complete time-step. The resulting explicit upwind scheme is as conservative as it is possible to be in axisymmetry and is second order in smooth regions. Being explicit, it is subject to the usual Courant stability condition

$$t_{k+1} - t_k \le \min_{ij}\left(\frac{h}{\lambda_{ijk}}\right), \tag{A21}$$

where

$$\lambda_{ijk} = [(\gamma p/\rho)^{1/2} + (v_r^2 + v_z^2)^{1/2}]_{ijk} \tag{A22}$$

is the maximum wave speed in the i, j cell at time $t = t_k$.

As we have already remarked in Section 3, this scheme has been applied to a number of industrial problems and although much of this work is not available in the open literature, some of the results are given in Falle (1988) and Falle & Wilson (1988). Although it is much simpler than the Piecewise Parabolic Method (PPM), it seems to perform just as well in most multidimensional applications. The reason for this seems to be that although PPM is superior in one dimension, most of these advantages are lost when it is operator split to cope with several space dimensions. Our scheme is only partially operator split in that, although we solve one-dimensional Riemann problems, we do not factorize the scheme into r and z operators.