C.P. No. 1226

C.P. No. 1226

MINISTRY OF DEFENCE (PROCUREMENT EXECUTIVE)

AERONAUTICAL RESEARCH COUNCIL

CURRENT PAPERS

Transonic Fan Noise

By

D. Hawkings,

Loughborough University of Technology

LONDON: HER MAJESTY'S STATIONERY OFFICE

1972

PRICE 26p NET

CP No. 1226*
March, 1972

TRANSONIC FAN NOISE

by

D. Hawkings,
Loughborough University of Technology

1. Introduction

Although transonic compressors were first developed in the 1950's, it has been their recent introduction into large commercial turbofan engines which has aroused widespread interest in their distinctive acoustic radiation. A transonic fan is basically an axial flow compressor which operates with a supersonic relative velocity at the blade tips ($M_{rel} \sim 1.4$ say), but a subsonic velocity at the blade roots ($M_{rel} \sim 0.7$). The axial velocity is everywhere subsonic. Such a device has good aerodynamic performance, with acceptable losses, and has now found widespread application in the aero-engine field. To the uninitiated, it is surprising that such a machine works at all, let alone works well! The fully supersonic fan, on the other hand, because of its high losses, is unlikely to find practical application in the near future.

The acoustic radiation from transonic fans has a unique character. Spectral analysis of the far field noise shows that it is dominated by a large sequence of discrete tones at all multiples of the basic fan rotation frequency ω. These tones are variously called 'multiple tones', 'combination tones', or 'buzz saw tones', and are most unexpected since arguments of symmetry suggest that only harmonics of the blade passing frequency (BPF = $\omega \times$ No. of fan blades) could be generated. Furthermore, only BPF harmonics are generated by subsonic fans. Other features of multiple tone radiation have been demonstrated experimentally. The acoustic signal is very steady, repeating almost perfectly with every complete rotation of the fan. However, every fan appears to have a unique acoustic 'signature' which is impossible to predict. Other experimental evidence will be discussed later.

In this report, we first present in physical terms the currently accepted explanation of this phenomenon. This is followed by an explanation of the various mathematical theories which have been used to give quantitative substance to these physical ideas. However, the deficiencies of these theories are then discussed, partly just to highlight them, but mainly to provoke ideas as to how the mathematical models might be improved. Finally, their implications regarding the design of quiet transonic fans are reviewed.

The effects of acoustic liners upon multiple tone radiation are not considered here. However, in practice these do make substantial reductions in the noise, and a judicious combination of liners and the source reduction techniques suggested in § 5 may well be the best short term solution to this problem.

*Replaces A.R.C.33 543.

2. Physical Explanation

Much experimental evidence on multiple tone generation suggests that linearised acoustics is inadequate to describe this phenomenon. Linear theory (Ref.1) suggests that for a fan operating at supersonic tip speeds, strong discrete tones would be generated at the rotor and propagate with constant amplitude along the duct. However, if all the blades were identical these tones would be confined to harmonics of the blade passing frequency, and not include other multiples of the rotation frequency ω. Even if small differences in the blades were admitted, so that all multiples of ω were generated, the BPF harmonics would still dominate, and again the amplitude of all the tones would remain constant along the duct.

Experimental observations made within the inlet ducts of transonic fans flatly contradict these predictions of linear theory (Ref.2). First, the overall amplitude of the acoustic signal decays steadily along the duct, usually as $x^{-\frac{1}{2}}$ near the rotor, and as x^{-1} at larger distances (x is distance from rotor face). Second, the spectral content of the acoustic signal changes with x. Near the rotor, the BPF harmonics are dominant, and there is only a relatively small amount of energy in the other rotation order frequencies. However, as x increases, there is a steady decline in the importance of the BPF harmonics, and an increasing prominence of the other tones. Ultimately, the spectrum is filled with a fairly uniform distribution of tones at all multiples of ω. Clearly, this simultaneous decay and spectral evolution of the signal cannot satisfactorily be explained on linear acoustic theory.

A much better physical understanding is obtained if the problem is regarded as one of weak shock wave generation and propagation, rather than sound wave generation. Since the supersonic parts of the blades must generate shock waves, clearly some account must be taken of these. In fact in-duct observations show that the 'acoustic' waveform ahead of the rotor has the characteristic sawtooth form that is associated with the passage of a sequence of shock waves, so that the forward acoustic field is just the shock field of the fan blades. It is easily verified that the shock waves are indeed weak in the aerodynamic sense; typically $\Delta p/p_o < 0.1$.

With the aid of weak shock theory, we can now explain some of the observed features. The supersonic part of each blade generates a shock wave at its leading edge. These shocks are 'locked' to the fan, and rotate with it; a stationary observer is swept by the same sequence of shocks every fan revolution, and consequently observes a very repetitive signal. In a stationary frame, the shocks appear to spiral forward towards the duct inlet, and thence into free space. Because of non-linearity and the dissipation inside the shock waves (due to viscosity and heat conduction) their strengths diminish as they travel forward, and in fact the decay rate at large distance is roughly like x^{-1}. This tallies with the observed decay rate.

Weak shock theory can also explain the spectral evolution process, if it is realised that the velocity of a shock wave increases with its strength. If all the blades were identical, their shocks would initially be of the same strength and relative location. Consequently they would all travel at the same speed and decay at the same rate, and so remain in a perfectly regular array. However, if one of the shocks was initially slightly stronger than the rest, it would travel faster than the others, and catch up the one ahead of it.

These/

These two would merge, and the new shock would then continue to chase and catch the others. And so on. Ultimately if this process were allowed to continue, only one shock in the system would remain. This type of behaviour is familiar in sonic boom generation, where all the minor shocks generated by the aircraft ultimately run into the head or tail shock. Thus provided we assume there are some initial differences in the fan blades - such as those tolerated in manufacture - then a steady degeneration from an initially regular to a completely irregular shock pattern can be explained. This accounts for the spectral energy transfer described previously. It also accounts for the uniqueness of each fan's acoustic signature; they all have a slightly different pattern of initial irregularities.

Thus, the currently accepted view of multiple tone generation is that it is a natural consequence of the non-linear shock behaviour associated with the supersonic aerodynamic flow. We now proceed to consider the mathematical analyses that have been used to give quantitative support to these ideas.

3. Mathematical Theories

Mathematical theories of multiple tone generation used to date have relied heavily upon either simple supersonic aerodynamic ideas or on shock propagation concepts. It is thus convenient to classify the theories as either 'aerodynamic' or 'non-linear acoustic' depending upon their approach to the problem. Ultimately they all rest upon the same physical ideas and give essentially the same answers.

The 'aerodynamic' theories (Refs.3,4,5,6,7) replace the rotor by a two-dimensional cascade representing a typical supersonic section of the blading. The forward flow field for a set of identical blades with curved suction surfaces is depicted in Fig.1. It consists of alternate shock waves attached at the leading edge and expansion waves emanating from the suction surface. Simple aerodynamic arguments show that all flow variables - such as pressure, velocity - are constant along the expansion rays. The Mach number and flow angle in each expansion wave family are related by the Prandtl-Meyer function. The shock waves separate adjacent families of waves and are determined by the property that they bisect, the angle of the two expansion waves intersecting at that point. These conditions are sufficient to determine the whole of the inlet flow field.

For a perfectly symmetrical cascade, it is found that sufficiently far upstream ($x >>$ radius of curvature of suction surface) the shocks are separated circumferentially by the blade pitch s, the pressure falls linearly between the shocks, and the shock strength $\Delta p/p_o$ is given by

$$\frac{\Delta p}{p_o} = \frac{s}{x}\,\frac{2\gamma}{(\gamma + 1)}\,\frac{\sqrt{M^2 - 1}}{M^2}\,(\sqrt{M^2 - 1}\,\cos\theta - \sin\theta)^2. \qquad \ldots.\ (1)$$

Here M is the free stream relative Mach number and θ the mean stagger angle of the blades. Nearer the cascade, the flow field is more complex, and decays with distance like $x^{-\frac{1}{2}}$. It should be noted that the far field result (1) does not depend upon either the incidence or curvature of the blades but only upon the free stream conditions. The precise blading of the cascade does not influence the far field pattern.

The/

The flow field of an imperfect cascade can be inferred from the above model. The inclination of the expansion waves on each blade can still be related to the surface direction via the Prandtl-Meyer function, and the values of pressure, etc., on each ray deduced. However, they now vary from blade to blade. Likewise the shock waves are still located by the 'angle property' mentioned above, but again vary from blade to blade. In principle the whole flow field can then be calculated, although certain complications arise where two shock waves merge. In practice it is necessary to resort to numerical techniques to calculate the upstream pressure profile for cascades with more than a few non-uniformities.

By means of a series of examples, a picture can be built up of the behaviour of irregular cascades, and this picture confirms the previous ideas. The stronger shocks overtake the weaker ones, and the regularity of the shock pattern degenerates as it moves away from the cascade, transferring spectral energy from the BPF harmonics to the other rotation order tones. It also shows the important point that during this transition phase, the decay of the overall level of the signal is slower than if the cascade had no imperfections (Ref.6).

A variant of this approach has been proposed in Ref.5, where a statistical analysis has been attached to the model. The 'expected' acoustic field is then related to a statistical description of the cascade irregularities (i.e., their means and variances) rather than their detailed values. This has the obvious advantage of eliminating the need for a precise description of the cascade, but the converse disadvantage that a particular measured acoustic field may be quite different from the predicted one, since it is only a single realisation of the field, and not the 'expected' one.

The 'non-linear acoustic' theory (Refs.7,8), on the other hand, regards the flow as unsteady one-dimensional, rather than steady two-dimensional. It assumes as known an initial pressure profile just ahead of the cascade and analyses its subsequent progress along the duct. Thus this theory cannot take account of the geometric details of the blades, nor predict the $x^{-\frac{1}{2}}$ near field decay, but it gives more insight into the ultimate shock propagation process. The analysis proceeds in two stages. First, the development in time of the given initial profile is calculated. Second, the axial distance x covered by the shocks in this time is then deduced, bearing in mind that they spiral obliquely into the incoming flow. This allows an effective separation of the decay process - due to non-linearity and dissipation - from the propagation process - due to compressibility and convection by the axial inflow.

For a uniform cascade, this approach leads to the following result. The shock strength at large time t is related solely to the time and the initial spacing λ, and is independent of the initial amplitude;

$$\frac{\Delta p}{p_o} = \frac{2\gamma}{(\gamma + 1)} \frac{\lambda}{a_o t} . \qquad \text{.... (2)}$$

For a flow at relative Mach number M and angle θ, the non-dimensional time $a_o t/\lambda$ is related to x/s by

$$\frac{a_o t}{\lambda} = \frac{x}{s} \frac{M^2}{\sqrt{M^2 - 1}} \left\{\sqrt{M^2 - 1} \cos\theta - \sin\theta\right\}^{-2} . \qquad \text{.... (3)}$$

Clearly,/

Clearly, combining (2) and (3) gives exactly the same answer as before (1). For an irregular cascade, an initially asymmetric profile is assumed and its subsequent development computed numerically. The relationship between t and x can still be taken as (3), since this accounts for the propagation of the shocks at a mean speed a_o; the effects of excess velocity are automatically accounted for in stage one. Again, computed results based on this approach confirm the tendency of the shock pattern to degenerate and transfer its spectral energy to the rotation frequency harmonics (Ref.8).

A third possible method of analysis would be to use the Whitham weak shock theory (Ref.9), which systematically improves a linear theory by taking account of the cumulative non-linear effects. This theory is used in sonic boom computations. This method also gives the result (1), but has certain extra advantages, see § 4.

It should be mentioned in passing that a 'special case' exists in the aerodynamic theories. If the blade suction surfaces are all flat and parallel, then it can be shown that the incoming gas must also flow in this direction, and there are no shock or expansion waves present. The upstream gas is entirely undisturbed by the blades. At first sight, this would appear to be an ideal solution, as there is clearly no forward noise. However, if it is admitted that all cascades have some slight imperfections, then a pattern of shock and expansion waves returns to the system, and these are of similar magnitude to those normally found on curved blades. Thus although nominally flat blades may make some differences to the multiple tone generation, they will not completely eliminate it; they are an ideal solution, not a practical one!

Thus, to summarise, all the present theories are effectively the same, and confirm and quantify our physical understanding of the problem. Clearly some implications can be drawn from these theories regarding the possibility of low noise transonic fans, but this is postponed to § 5. In the next section, we discuss some of the deficiencies of these theories, in the hope that this may provoke some ideas on possible improvements.

4. Theoretical Deficiencies

The above theories are attractive both for their analytical and conceptual simplicity, and for their apparent agreement with experimental experience. Admittedly, a certain amount of 'adjustment' is required to correlate theory and experiment (e.g., careful selection of the right M and θ to insert in (1)), but none the less these theories are often regarded as adequate. However, because they are two-dimensional (2D), they should strictly only apply in narrow annular ducts of constant cross-sectional area. Their application to real fans - where there is significant spanwise variation in conditions, as well as duct area changes - is clearly suspect, and merits a more critical appraisal.

There are two key predictions of the 2D theory, namely the approximately x^{-1} decay of the signal level, and the degeneration of the waveform with increasing axial distance. Taking the latter point first, although this is an undoubted prediction of the 2D theory, to my knowledge it does not tally accurately with real experimental results. For example in Ref.10, a test compressor was inspected and its non-uniformities measured. These were used in a computer run to predict the forward pressure signature and spectral content. However, the agreement between the predicted and measured spectra was poor; the theory over-predicted the importance of the first few

shaft/

shaft order tones, and under-predicted the mid-range tones (around $\frac{1}{2}$ BPF). Admittedly, the rotor was inspected statically, whereas it is its condition at speed which is important, but unless it was significantly more uniform at speed than it was at rest, this is not likely to make much difference. Again, in my own investigations I have tried using pressure profiles measured at one station to predict those at another, with very little success. Here also the theory over-predicted the low order tones. Thus, it is suggested that the 2D theory does not accurately predict the evolution of the waveform, but only in a qualitative manner.

What can be said about the x^{-1} law? Let us digress and return briefly to the Whitham theory approach. This has successfully been extended to cases of weak shock propagation where the area of the shock front expands, e.g., cylindrical or spherical waves (Ref.11). This theory considers the propagation of the shocks along 'ray tubes' whose cross-sectional area varies in a known way $A(x)$. For periodic waves, the shock strength varies as

$$\frac{\Delta p}{p_o} \propto \left(\sqrt{A}\int_o^x \frac{dx}{\sqrt{A}}\right)^{-1} \qquad \ldots . \ (4)$$

so that if A varies as x^{β},

$$\frac{\Delta p}{p_o} \propto \frac{1 - \frac{\beta}{2}}{x} . \qquad \ldots . \ (5)$$

Thus with the singular exception of spherical waves $(\beta = 2)$ provided the area expands as a simple power of x, the shock strength always decays as x^{-1}. Physically this is because the area expansion attenuates the excess velocity behind the shock, and hence the rate at which energy is fed into the shock and dissipated. This reduces the dissipative decay of the waveform, but is exactly offset by the attenuation due to area expansion.

These comments suggest that the x^{-1} variation of shock strength is not a feature unique to 2D theories, but could be found in any theory that took account of the changing duct area. Thus, this experimental observation is not sufficient justification for the 2D theories, and it must be concluded that these models provide at best only a qualitative description of real fan situations. In fact other predictions of the Whitham theory highlight the danger of placing too much direct quantitative reliance on these models. For example, because the 'acoustic' energy flowing along the ray tubes is proportional to $(\Delta p/p_o)^2 A$, (this energy would be conserved in linear theory, but decreases due to dissipation in shock theories) it varies as

$$E = \left(\frac{\Delta p}{p_o}\right)^2 A \propto x^{\beta - 2}\,(1 - \tfrac{\beta}{2})^2, \quad (\beta < 2). \qquad \ldots . \ (6)$$

Thus although the shock strength always decays as in the 2D model, the local acoustic energy flux does not; it must be corrected by multiplying by the duct area and the algebraic factor $(1 - \frac{\beta}{2})^2$.

Although the 2D theories are inadequate, there do not appear to be many easy improvements to be made. One possibility is a detailed application of the 'ray tube area' theory to the shock propagation. It would be desirable

to/

to also incorporate the effects of the varying main stream axial velocity, a feature which was ignored in the previous discussion but which might well be significant. This model would be a modest improvement on the 2D theories, as it would now apply to narrow annular ducts but of varying cross-sectional area, and it does appear to be a feasible problem. Although it would expose the precise effects of area change, it is unlikely that any major new features will come to light. It would however, improve the prediction of the waveform degradation process, since increasing area slows down the non-linear effects of excess velocity, and so will inhibit the rapid build up of the low order tones.

Beyond this, one has to face the fact that the real flow is truly three-dimensional, with significant radial variation in conditions, and that many more physical effects must be considered. The relative blade speed varies with radius, and passes through a sonic condition at some intermediate point. (As the blade rotates, this point sweeps out the 'sonic circle', and the cylinder formed on this circle by generators parallel to the axis is termed the 'sonic cylinder'.) The blade aerodynamic flow near the sonic point is likely to exhibit all the complexities of transonic airfoil flow, and obtaining even a partial understanding of this region may well be difficult. Presumably there are no shock waves generated inside the sonic circle, but the transition to a shocked flow outside the circle is likely to depend critically upon local conditions such as leading edge bluntness, boundary layers, etc. Consequently this region may well be a further cause of non-uniformities in the initial shock pattern which will be exaggerated during propagation. If this is so, some means of stabilising this sonic transition may be desirable.

The sonic cylinder also plays a role in the forward propagation of the shocks. It is well known that the shock front generated in free space by an accelerating body contains a cusp. For a body rotating in a circle the shock front never penetrates inside the sonic cylinder, instead the inward arm is cusped and then propagates outwards. Fig.2 shows a typical example, it is a contour plot of the shock front generated by a point disturbance rotating at Mach number $M = \pi/2$. A similar feature still applies to a ducted rotor, although the outer wall now reflects the shocks inward again and alters the picture from Fig.2. However, it still remains true that no shocks enter the sonic cylinder, so what happens if the forward ducting is narrowed to a diameter less than that of the cylinder? Do the shocks get trapped and fail to propagate out of the duct?

The final, and perhaps most important feature of the shock propagation process is the fact that acoustic waves inside a duct are dispersive. Due to the constraint imposed by the wall, only certain duct modes can exist, and these all propagate at different speeds. Hence any initial wave shape is distorted, and this combined with the fact that the flow is non-linear means that the shock propagation is essentially a problem in a non-linear dispersive system. Although some recent investigations have been made in this area they are specifically restricted to unshocked modes. However, they do show an important new effect termed 'subsonic choking' (Ref.12) in which a rotating mode is prevented from propagating forward by an axial velocity <u>less</u> than sonic due to the combination of non-linear and dispersive effects. Maybe there is a similar effect on shocked modes.

Thus to conclude, it appears that the 2D theories are severely limited in their relevance to real transonic fans, in that many important physical effects are not incorporated. Unfortunately, the desirable theoretical improvements will be difficult, involving concepts from both transonic flows and non-linear dispersion theory, and so may not yield to theoretical attack for a long time.

5./

5. Conclusions

The physical and mathematical models described in §§ 2 & 3 obviously have various implications regarding the possibility of designing quiet transonic fans. These will now be described in more detail.

(1) Multiple tone noise is an inevitable consequence of the steady aerodynamic flow around the blades, and results from the special behaviour of the blade shock waves. Its detailed form depends greatly upon the pattern of dissimilarities between the blades, but does not depend strongly on the nominal blade design. Consequently it is neither accurately predictable nor repeatable from one fan to another, nor is it likely to be seriously altered by changes in fan blading.

(2) The propagation of the shock waves is a process of simultaneous decay and degradation of the original ordered waveform into a smaller, disordered one. Thus it is not possible to have low overall levels without widespread distributions of tones.

(3) If low overall level is the primary aim, then the design objective is simply to maximise the non-dimensional time $a_o t/\lambda$ spent by the shocks in the duct. This can be achieved by either reducing the blade spacing (increasing the number of blades) or increasing the duct length. Alternatively, the velocity triangles can be chosen to increase the factor in Eqn.(3) (see Fig.7 or Ref.8); in general an axial velocity as high as possible is required, as this makes it harder for the shocks to propagate forward. This implies an increase in relative Mach number or stagger angle, but requires major aerodynamic changes to the rotor and may be impossible due to other aerodynamic constraints.

(4) The effects of increasing the duct area can be described qualitatively, and are unhelpful for achieving low overall levels. This is because in expanding ducts, the rate of destruction of acoustic energy is reduced. Also the axial velocity will decrease and so allow the shocks to travel down the duct more quickly.

The above points appear to cover most of the low noise implications in our present understanding. Certainly, much more detail could be extracted from the 2D theories, but in view of the comments in § 4, these are likely to be of limited application to real fans. These implications lead to a rather pessimistic view of the possibility of quiet high speed fans. A 20 dB reduction in overall level at the duct exit would require at least a tenfold increase in $a_o t/\lambda$, and although this is not completely impossible it does appear to be at the limit of what could be reached using present aerodynamic technology.

There then remains the possibility that the present theories miss some vital physical mechanism which could be turned to advantage, such as those mentioned in § 4. It is clear that the present theories are far from exhaustive, and it is quite probable that some new features will be revealed theoretically in due course. However, there appears to be an urgent need for some experimental work on this problem, in particular looking at the strengths, locations, and behaviour of the shock waves in the duct near the rotor and

especially/

especially their three-dimensional features. This could be done both on current fan configurations and also on some of the more radical possibilities that might be suggested. Such experimental work would both stimulate the theory, and perhaps expose some of these other mechanisms that hopefully will ultimately allow the operation of quiet transonic rotors.

Acknowledgement

This report was prepared whilst the author was supported by an SRC research grant.

References/

References

No.	Author(s)	Title, etc.
1	Tyler and Sofrin	Axial flow compressor noise studies. SAE Trans. 70, pp.309-332. 1962.
2	Sofrin and Pickett	Multiple pure tone noise generated by fans at supersonic tip speeds. International symposium on fluid mechanics and design of turbomachinery. Penn. State University. 1970.
3	Yamaguchi	On the inlet flow field for a two-dimensional cascade with curved entrance region. Bulletin of JSME 7, 25, pp.91-95. 1964.
4	Kurosaka	A note on multiple pure tone noise. JSV 19, 4, pp.453-462. 1971.
5	Pickett	The prediction of the spectral content of combination tone noise. AIAA Paper No. 71-730. 1971.
6	Fink	Shock wave behaviour in transonic compressor noise generation. ASME Paper No. 71-GT-7. 1971.
7	Morfey and Fisher	Shock wave radiation from a supersonic ducted rotor. Aeronaut. J. 74, pp.579-585. 1970.
8	Hawkings	Multiple tone generation by transonic compressors. JSV 17, 2, pp.241-250. 1971.
9	Whitham	The flow pattern of a supersonic projectile. Com. Pure & App. Maths. 5, pp.301-348. 1952.
10	Kantola and Kurosaka	The theoretical and experimental investigations on multiple pure tone noise - Part 1. NASA CR-1831.

No.	Author(s)	Title, etc.
11	Whitham	On the propagation of weak shock waves. JFM 1, 3, pp.290-318. 1956.
12	Tam	On finite amplitude spinning acoustic modes and subsonic choking. JSV 16, 3, pp.393-405. 1971.

PAC

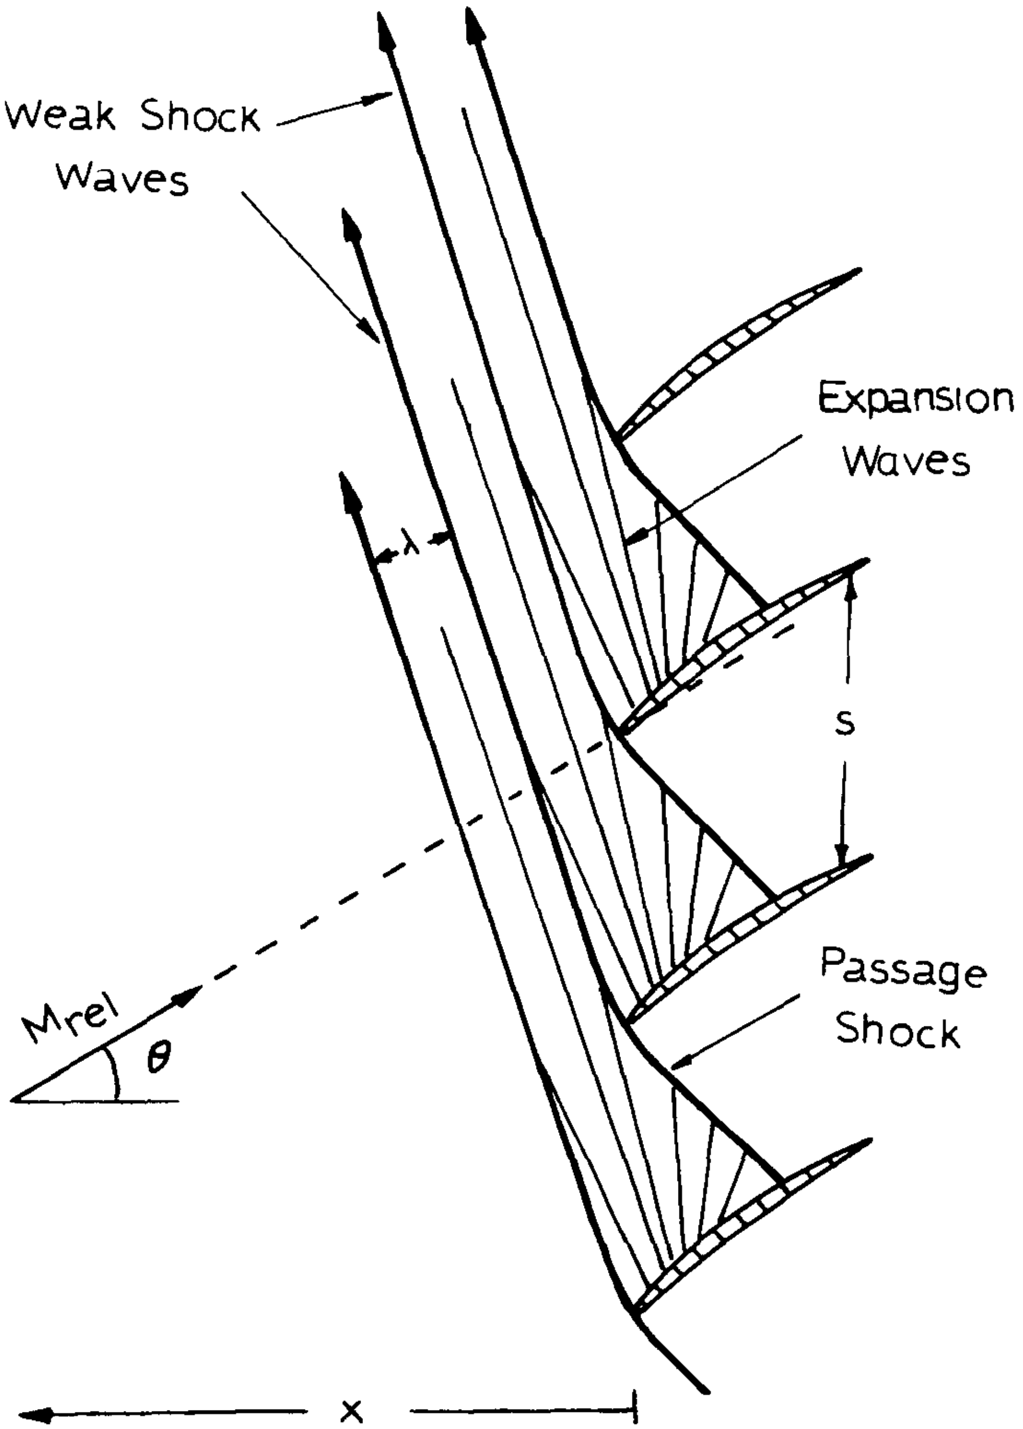

Fig 1 INLET FLOW OF TWO-DIMENSIONAL CASCADE

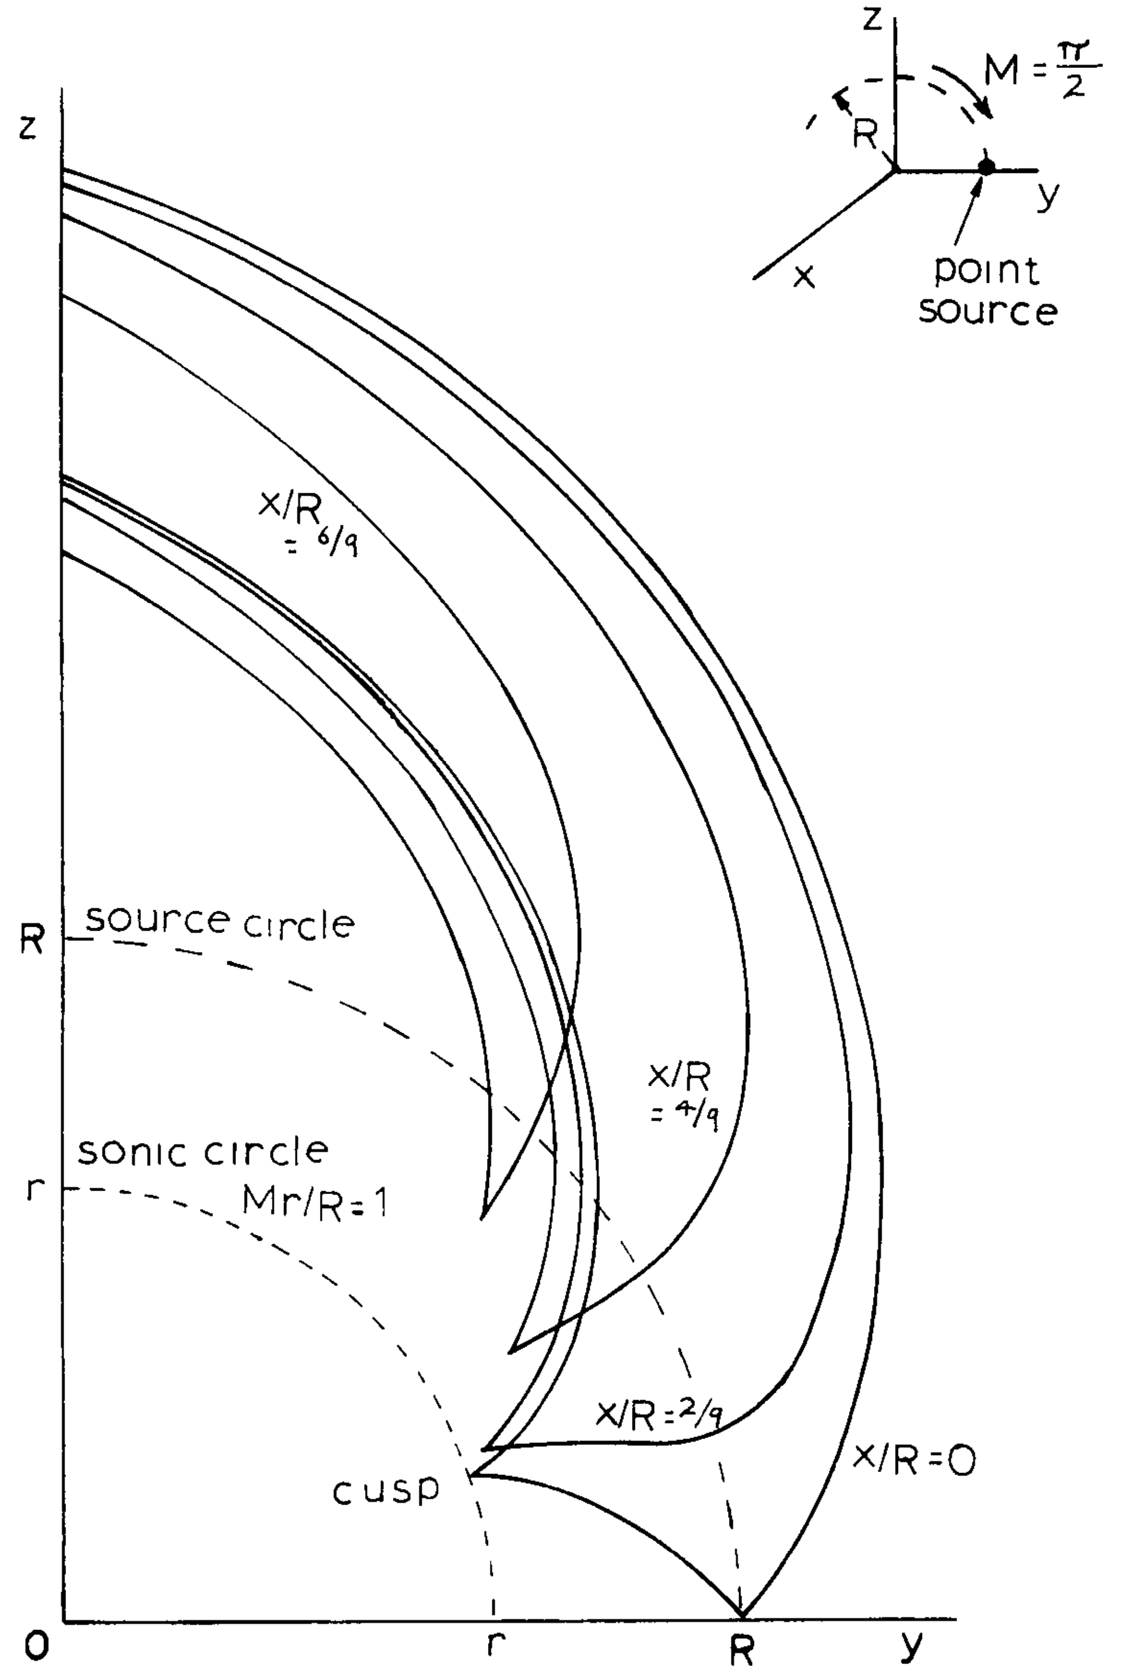

Fig 2 Shock front caused by point source in circular motion

R13162/1/502109 K4 8/72 P

DETACHABLE ABSTRACT CARDS

ARC CP No.1226
March, 1972

Hawkings, D.

TRANSONIC FAN NOISE

Multiple pure tone generation by transonic fans is discussed. This is explained in terms of weak shock theory, and the various related mathematical models are outlined. The implications of these models for the design of quieter transonic fans are emphasised. Some important physical features not yet included in these theories are also highlighted.

ARC CP No.1226
March, 1972

Hawkings, D.

TRANSONIC FAN NOISE

Multiple pure tone generation by transonic fans is discussed. This is explained in terms of weak shock theory, and the various related mathematical models are outlined. The implications of these models for the design of quieter transonic fans are emphasised. Some important physical features not yet included in these theories are also highlighted.

ARC CP No.1226
March, 1972

Hawkings, D.

TRANSONIC FAN NOISE

Multiple pure tone generation by transonic fans is discussed. This is explained in terms of weak shock theory, and the various related mathematical models are outlined. The implications of these models for the design of quieter transonic fans are emphasised. Some important physical features not yet included in these theories are also highlighted.

C.P. No. 1226

HER MAJESTY'S STATIONERY OFFICE

Government Bookshops

49 High Holborn, London WC1V 6HB
13a Castle Street, Edinburgh EH2 3AR
109 St Mary Street, Cardiff CF1 1JW
Brazennose Street, Manchester M60 8AS
50 Fairfax Street, Bristol BS1 3DE
258 Broad Street, Birmingham B1 2HE
80 Chichester Street, Belfast BT1 4JY

Government publications are also available through booksellers

C.P. No. 1226

SBN 11 470784 7