Using Doppler DSP to

*by Peter Martinez, G3PLX**

A PREVIOUS *RADCOM* article I wrote, in November 1997, described the use of extremely narrow-band DSP techniques to receive very weak signals on the VLF amateur bands. This was done by using a very narrow-span spectrum analyser to view a small section of the band on a real-time computer display, which showed frequency as the Y-axis, time as the X-axis, and signal-level as brightness. The result was a system capable of receiving signals buried far below the audible noise level, by virtue of its extremely narrow effective bandwidth.

SETTING THE SCENE

SOME READERS of that article may have asked the question: what happens if we use this technique on the higher bands? The answer is that the propagation medium on the higher bands is not as stable as it is on VLF, and instead of seeing a thin straight line on the screen from a transmitted carrier, the signal is spread out over an appreciable band of frequencies. The ionosphere, which reflects HF radio waves, is constantly in motion, causing the path length to vary. As with moving low-orbit Oscar satellites, a changing path length causes a change in the received frequency, known as the Doppler shift. While a moving satellite on VHF can give a Doppler shift of several kHz, the moving ionosphere at HF only gives a Doppler shift of about 1Hz, but this is much more than the 25mHz bandwidths which we were using on VLF and would make communication at these narrow bandwidths impossible on HF.

In this article, I would like to introduce the idea of using these narrowband techniques, not to communicate on HF, but to study the ionosphere itself. The narrow-band spectrum analyser can be used to create 'dopplergrams' of typical HF signals, which can be an effective way of showing many of the features of the ionosphere that have previously only been made visible with equipment well beyond the reach of the average amateur.

AIRCRAFT FLUTTER

TO BEGIN, let us consider perhaps the best-known example of the Doppler effect, familiar to radio amateurs and television viewers alike - aircraft flutter.

If we suppose that the height of the aircraft is small compared to the other distances, then

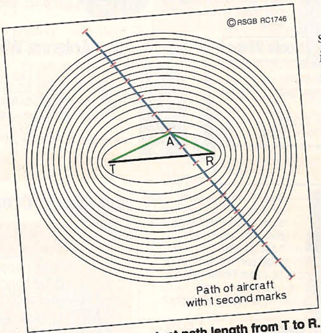

Fig 1: Ellipses of constant path length from T to R.

we can make a two-dimensional model of the reflected path. This is shown in **Fig 1**, and shows what is happening when the signal from T is reflected from a moving aircraft at point A and received at R. As the aircraft moves closer to the line joining T to R the path gets shorter, giving rise to a positive Doppler shift, and as it moves away from the line between T and R, the signal shifts lower in frequency. If we also have a direct path between T and R, the presence of two signals of different frequencies at the same time gives rise to a beat-note, usually at a sub-audible rate, which shows as the well-known flutter fading effect.

The elliptical rings drawn on Fig 1 are lines of constant path-length from T to R via any point on the ellipse. We can imagine that these were drawn by fastening each end of a length of string to T and R, starting with the string one wavelength longer than the distance TR. Drawing the string tight with a pencil at point A, we can then trace out the innermost ellipse. The second ellipse is drawn with two wavelengths of extra string, and so on. An aircraft flying tangential to any ellipse will have zero Doppler shift, and the Doppler shift on an aircraft flying in any other direction, measured in Hz, is numerically equal to the number of wavelength rings that it crosses per second. From this we can see that an aircraft approaching from some distance will appear first with an almost constant positive HF Doppler shift, which will drop, pass through zero as the aircraft either crosses between the transmitter and the receiver or its flightpath becomes tangential to an ellipse, then becomes negative as it flies away. The Doppler shift is proportional to the RF frequency, and with a typical jumbo jet a shift as high as 250Hz is possible on a 144MHz signal. The narrow-band spectrum analyser technique described in November 1997 *RadCom* can be easily used to display these Doppler-shifted signals, even on frequencies much lower than 2m.

Fig 2 shows a dopplergram taken on a BBC World Service AM transmitter on 15.485 MHz at Penrith, Cumbria, about 40km north of my home, at about 4pm. This signal is at the limit of surface-wave propagation, but well inside the skip zone. Listening in USB mode, with the carrier tuned to give a 1000Hz tone, it shows signs of aircraft flutter from time to time. The horizontal centreline of the dopplergram represents the carrier frequency, where the direct signal is visible. The top and bottom of the dopplergram are respectively 25Hz above and below the carrier frequency. The marks along the bottom edge represent 10 minute intervals.

The striking feature of Fig 2 is the very large number of aircraft Doppler trails visible; far more than can be detected by ear. The narrowband analyser is doing a far better job of separating out the faint, closely-spaced signals than the human ear can. To get some idea how far away some of these aircraft are, we can note that the area swept out between an aircraft trace

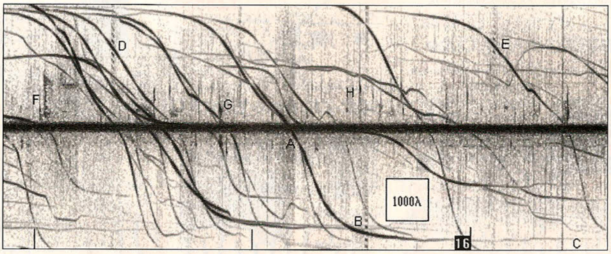

Fig 2: The Doppler trails from numerous aircraft.

*High Blakebank Farm, Underbarrow, Kendal, Cumbria LA8 8HP.

Study HF Propagation

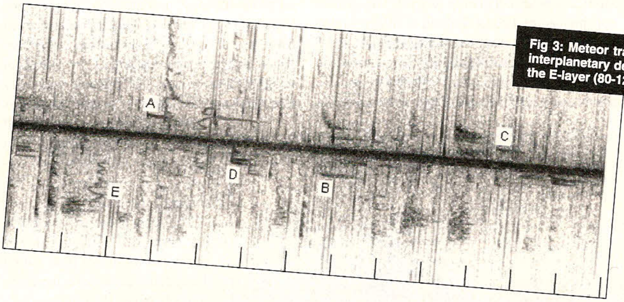

Fig 3: Meteor trails, caused by the burning-up of interplanetary debris at about the same height as the E-layer (80-120km).

and the centre-line over a given time interval is numerically equal to the number of wavelength ellipse rings crossed by the aircraft in that time. We could, for example, estimate this area by counting the number of 1Hz by 1sec squares in the area swept out on the dopplergram. Since we know that the aircraft is at its closest point when the trace crosses the centre-line of the dopplergram, an estimate of the area of the triangular section between the point where the trace crosses the centre-line and the faintest detectable signal will give us a figure for the most-distant path length minus the closest. We can then say that the most distant path length is certainly longer than this figure. Since the path is from the transmitter to the aircraft and back to the receiver, we can thus get a lower limit for the distance to the aircraft by halving this path-length. The square area marked 1000λ in Fig 2 represents a distance of 1000 wavelengths, that is a path-length of 19.3km at this frequency. Using this method, the aircraft which crossed the centre-line at point A, and produced the trail to B, and was just detectable at point C, was at least 150km away at point C.

Some of the traces in Fig 2 do not follow smooth curves, and these can be attributed to aircraft which are changing speed or direction. Some traces run along the centre-line for some time, and these result from aircraft that are flying along the line between transmitter and receiver. It is interesting to speculate on the possibility of using several receivers spaced out around a single broadcast transmitter and combining all the Doppler signals to give an almost completely 'passive' radar picture of the aircraft activity in the area. The sensitivity is actually limited in Fig 2 not by receiver background noise, but by close-in sideband noise in the BBC transmitter, resulting from low-level second-order distortion in the modulated signal.

There are other strange things in Fig 2, some of which can be explained and some which cannot. At some points, for example the points marked B, D, and E, some aircraft seem to aquire sidebands at multiples of 1.3Hz for a few minutes, which seems to suggest that the reflected signal is being modulated in amplitude or phase at the 1.3Hz rate. This does not seem likely and these sidebands may not be what they seem.

There are other strange marks, as at F, G, and H in Fig 2, and numerous short vertical streaks both above and below the centre-line, which can be explained, and take us into the main uses of dopplergrams for propagation study. To explain these we need to look much higher in the sky than the aircraft. It helps to find a signal coming from much further away, so that aircraft reflections are below the horizon and are not visible. **Fig 3** shows just such a trace, from a French AM broadcast transmitter on 21.580MHz at about 1000km distance, recorded at about 1000UTC in April. There is no surface wave, but this station is still within the skip zone so there is no skywave; the stripe across the centre of the trace is probably a scatter signal. Again, the vertical scale is 50Hz top-to-bottom, but the horizontal timebase is twice as fast as Fig 2 and the marks along the bottom now represent one minute intervals.

METEORS

THE VARIOUS streaks and squiggles in Fig 3 are in fact meteor trail reflections. Meteors are small particles of interplanetary debris, typically the size of a grain of sand, which burn up on entering the top of the earth's atmosphere. Occurring at heights between 80 and 120km, they leave an ionised trail rather like a vapour trail which gradually fades away, and a radio signal can be reflected from this trail. Unlike an aircraft, which can be considered as a point reflector and scatters RF in all directions, a meteor trail is a line, and gives a strong mirror-like reflection when the trail is at the correct angle to the path.

This is the textbook explanation for meteor-scatter, but it does not yet explain the strange shapes we see on the dopplergram. If the meteor trail was stationary in the sky, the reflection would be exactly on the same frequency as the transmitter and would appear exactly on the dopplergram centre-line. Meteor pings have Doppler shifts because the trails are moving, blown along by high-altitude winds and will give a positive or negative Doppler shift depending on whether the trail is upwind or downwind of the path between the transmitter and the receiver. This accounts for some of the blobs in Fig 3, namely the horizontal streaks at A, B and C, but there are several places where the reflection seems to be spread out over a range of Doppler shifts, with a tendency to have a leading 'nose' and two or more tails. Once again, we can explore what is happening by reference to a diagram.

Fig 4 is a simplfied diagram of a vertical section through the atmosphere at a height of about 100km where a typical trail is being formed, and shows at (a) a meteor which is falling vertically downwards, with the wind blowing from left to right across the diagram. In general, the wind speed will not be the same at all levels. Let's suppose that the wind is stronger at the higher levels, as shown by the arrows, and that the transmitter is in front, to

Fig 4: How the initially straight trail from a vertically falling meteor can be blown out of shape by high altitude winds.

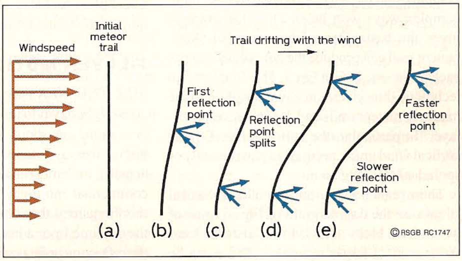

the right of, and below the diagram and the receiver is behind it, so that the diagram is a slice across the path which includes the meteor trail. At the instant at which the trail is formed, there is no reflection, since the trail is vertical and the transmitter and receiver are on the ground 100km below and not lined up with the trail. After a few seconds maybe, the wind gradient along the trail has blown it into the shape at Fig 4b and the middle of the trail becomes perpendicular to the transmitter/receiver plane and a signal is reflected. Note that different parts of the reflecting part of the trail are moving with different speeds, or in other words it is rotating, so that the reflected signal has a range of Doppler shifts simultaneously. It is this which gives the Doppler trace its flat 'nose'. As the trail shape blows further away from the vertical, Fig 4c, the middle of the trail rotates out of line and ceases to reflect, but the sections of trail either side now come into line and reflect. The upper reflecting section is travelling faster than the lower one, and so the nose of the Doppler trace splits into two tails, one higher in frequency than the other, just like the trace at D in Fig 3.

Fig 4 shows only one possible simple path geometry. In general, the alignment of the trail and the transmitter/receiver paths will be more complex, and the wind speed may vary in more complex ways with height, but this example gives the basic idea of how a wind-blown meteor trail can produce the convoluted shapes such as the one at E in Fig 3. The dopplergram technique thus gives some real insight into the nature of meteor trails and the physics of the E-layer. In particular, the ability to see E-layer vertical wind gradients may help in the study of sporadic-E propagation.

There remain a large number of faint vertical streaks on the dopplergram of Fig 3, some of which have blobs at their lower ends or occur at the start of larger squiggles. These can be explained by the Doppler shift of a signal reflected from the expanding tip of the ionised trail as it is actually being formed. The surface of the leading end of the trail will be rounded rather than flat, so the scattered signal will be fainter and not so highly directional, and the Doppler shift will not be just due to the windspeed but contain a component due to the much higher velocity of the meteor itself. Each meteor ping thus starts with a faint wide-angle 'chirp' from the HF side. Some of these faint chirps are received even though the resulting trail is never in the right alignment to produce a real ping, probably from meteors which fall quite close to the transmitter or the receiver.

F-LAYER MOTION

FOR THE NEXT dopplergram, **Fig 5**, we come right down to the lower HF bands. Fig 5 is a signal on about 5MHz, received at my home over a distance of about 80km. It is actually an unmodulated spare channel of a commercial multi-channel RTTY signal. On this frequency the ionosphere reflects during the daytime from a height of about 300km, so the reflection point is almost overhead. For this dopplergram the vertical scale is 3Hz top-to-bottom, and the time marks along the horizontal scale are at 10 minute intervals with the UTC hours also marked. The signal is reflected from the F-layer and is seen to be wandering about 0.5Hz up and down in frequency at random. This is not just due to frequency instability in the transmitter or receiver. If this transmitter had been close enough to show the direct (ground wave) signal, this would have been visible as a straight line across the centre of the chart.

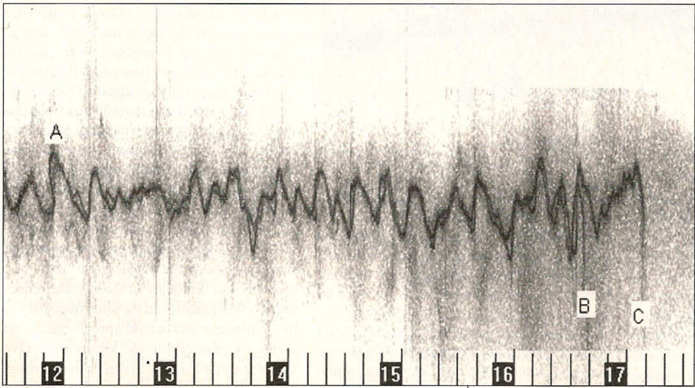

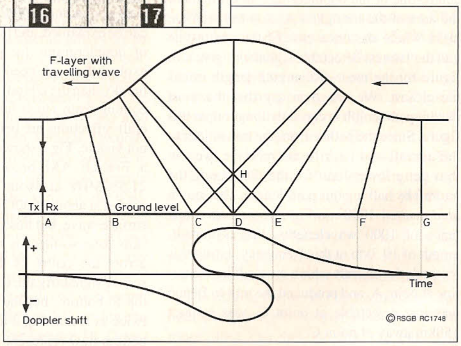

Fig 5 (above): Vertical motion of the F-layer leads to Doppler shift, including a 'switch-back' effect at point A. Left and right hand polarised signals from the same transmitter fade out at different times (points B and C). Fig 6 (right): Multiple reflection points across an upward bulge in the F-layer, illustrating the 'switch-back' effect at point A of Fig 5.

The ionosphere is basically a flat horizontal reflecting layer, and it can be shown that, unlike the meteor-scatter case, there can be no Doppler-shift due to horizontal motion of a horizontal layer. The observed Doppler shift is actually caused in this case by vertical motion of the F layer. Some of this motion can arise from weather effects in the atmosphere below, but in the same way that waves can propagate long distances across the surface of a pond, such disturbances can propagate horizontally as waves at the height of the F-layer.

Even quite small disturbances, due to thunderstorms in the atmosphere below, become magnified in size as they expand into the lower densities of the higher altitudes. If we could monitor the Doppler shift at two or more receiving sites, it might be possible to see the direction of motion of these waves, but there is another phenomenon which demonstrates the travelling-wave effect very clearly, and incidently might also help to convince some readers that this wobbly trace is not just transmitter or receiver frequency drift. At several points in time the wandering Doppler trace seems to have an almost vertical rising edge, and indeed at point A in Fig 5 the trace seems to defy the laws of physics and lean backwards in time. This can certainly not be receiver drift, but we can explain how this strange effect occurs by reference to **Fig 6**.

Imagine an upward bulge in the F-layer moving, like a tidal wave, from right to left across the top of the diagram above a stationary transmitter with the receiver next to it, shown at point A in Fig 6. In fact, to make the diagram less cluttered, Fig 6 has been drawn showing a stationary wave and with the transmitter moving along the line A to G, which will give the same effect. At point A, before the transmitter gets under the bulge, the path length to the layer is constant and there is no Doppler shift. At B the layer is starting to move upwards and away from the transmitter, giving an LF Doppler shift at the receiver, and this gets larger at C, but we also just start to get some signal from the opposite side of the bulge which is moving and rotating downwards and towards the transmitter, so another signal appears, with a 'nose' and a positive Doppler shift. At D, with the transmitter and receiver right under the centre of the wave, there are now three points of reflection, one from the upward-moving leading edge of the bulge, one from the top, and another from the downward-moving trailing edge.

At E we just lose the reflection from the leading edge as it rotates out of line, and at F and G we have only the upward Doppler shift from the trailing edge as the travelling wave passes away.

Only travelling waves which have a small enough radius of curvature will cause this 'switch-back' effect. The audible effect is that an otherwise steady signal develops deep fading for a few minutes as the three doppler-shifted paths beat with other.

DOUBLE VISION

IF WE LOOK at the rest of Fig 5 carefully, we can see many places where there are two traces close together, and it's not difficult to imagine that virtually the whole dopplergram actually consists of two traces which are weaving in and out of each other. Listening to the signal on the speaker while watching the dopplergram build up, it is clear that the fading rate is fast when the two traces are widely separated and slow when they are close together.

To explain this double trace, we need to go a bit deeper into the theory of HF propagation than the average amateur radio textbook. Such books tells us that the electrons in the ionosphere oscillate in sympathy with the electromagnetic wave, and in the case of a signal below the critical frequency fired vertically upwards, slow it down to a standstill and return it back along its original path. What these books do not tell us is that the earth's magnetic field has a subtle effect on this process. The physics is complex, but the result of the interaction between the radio wave, the electrons in the ionosphere, and the magnetic field, is that a right-hand circularly polarised signal propagating through the ionosphere has a slightly different propagation speed compared with a left-hand circularly polarised signal. This means that two waves polarised in opposite senses will take two slightly different ray paths through the ionosphere. One of these paths is known as the Ordinary ray and is independent of the magnetic field, and the other, which is slowed down slightly by the magnetic field, is known as the eXtra-ordinary ray. The amount of slowing down depends on the strength of the component of the earth's magnetic field along the path.

Since a linearly polarised signal can be resolved into two contra-rotating circularly-polarised signals, it follows that even linearly-polarised signals passing through the ionosphere are split into two separate ray paths, except for certain special cases at the magnetic poles and equator. The dopplergram shows where the two rays separate in frequency due to the motion of the ionosphere itself.

For most of the day the O and X signals stay close to each other, but at sunset when the ionisation level drops below the level needed to cause reflection, the O signal vanishes first. This can be seen at point B at 1638 UTC in Fig 5. As it goes, the reflection point moves rapidly upwards as the ionisation drops, giving rise to a characteristic LF 'chirp'. But the extra-ordinary ray, boosted by the extra spin that the electrons get from the magnetic field, keeps going until point C at 1709 UTC, when it too vanishes with a downward chirp. Interestingly, although this particular signal shows deep fading throughout the day as the O and X rays beat with each other, for the 30 minute period where only the X trace is present, the signal is quite constant in level.

Behind the main trace there is a fuzzy signal with a greater Doppler shift, which does not track that of the main trace. This is probably due to small-scale irregularities in the F-layer, caused by solar wind particles trapped along the earth's magnetic field. This fuzzy trace can sometimes get much stronger, swamping the main trace during periods of solar activity. There is an audible rapid flutter on the signal when this occurs. It is interesting to note that there is still a fuzzy background trace after the main signal has faded out at sunset.

Fig 5 showed a near-vertical path in order to demonstrate the vertical motion of the F-layer. For the next dopplergram we move to a 400km path and come down to 3.5MHz, in order to show the effects of low-angle propagation, and to introduce the E-layer which does not reflect at 5MHz. **Fig 7** uses the same vertical and horizontal scales as Fig 5, but is taken at sunrise

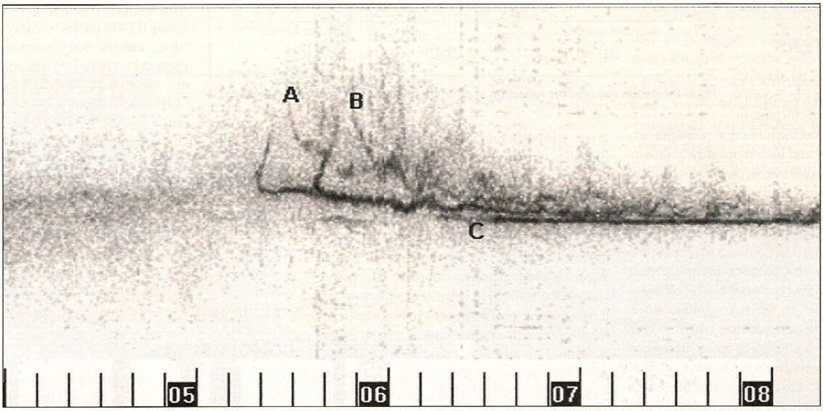

Fig 7: F-layer traces at dawn, showing the 'chirp' from the HF side as the reflection height drops, and the onset of the E-layer.

over the path from Great Yarmouth, on the east coast of England, to the writer's home at Kendal in north-west England. The transmitter is a navigational data broadcast on 3572kHz, which radiates 24 hours per day and has a nice, clean carrier in the centre of it's digital PSK modulation.

THICK LAYERS

IT IS PERHAPS no surprise that the sunrise in Fig 7 shows the F-layer traces chirping in from the HF side as the reflection point moves, first very rapidly and then more slowly downwards. There are two traces, the extra-ordinary ray appearing first. But, like the travelling wave in the F-layer, these traces seem to defy the laws of physics and move *earlier in time* after their first appearance. To explain this, we again need to go a bit deeper into ionospheric theory than the average textbook. The textbooks tell us that if we fire a radio wave obliquely upwards at a shallow angle it will reflect from the ionosphere, but at a steeper angle it will go through and out into space. This simple model assumes that the ionosphere is a thin layer. To explain the sunrise fold-back effect we need to take into account that the F-layer is thick, with the ionisation level varying smoothly from a low intensity above and below the layer to a high intensity in the middle, rather like the ionisation profile shown in **Fig 8a**, which shows the ionisation level along the horizontal axis and the corresponding height vertically.

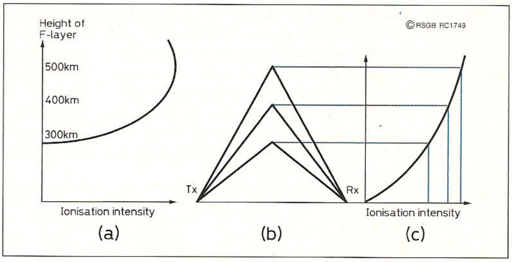

Fig 8: Ionisation required for reflection at different heights for a given path distance.

Consider what happens to a ray which is fired at an angle upwards, and suppose that it is only just reflected from a thin ionospheric layer at a height of 300km, and received back on the ground 400km from the transmitter. A ray at a higher angle will skip out into space and a ray at a lower angle will reflect back to the ground further away. The receiver is thus right on the edge of the skip zone. But let's follow the higher ray for a moment. It passes upwards through the 300km level to a greater height, say 400km. If the ionisation at 400km was high enough (it would have to be higher than at 300km, because of the steeper angle), this ray too could be reflected to our receiver at 400km distance. Three such rays at different take-off angles are shown in Fig 8b. With the help of some mathematics we can draw a smooth curve, showing the ionisation intensity that would be needed to reflect a signal from a thin layer at *any* height down to our receiver 400km away. This curve is shown in Fig 8c, and is known as the transmission curve for that distance.

If we now superimpose this curve on top of the actual F-layer ionisation profile of Fig 8a, we can predict what will happen if the F-layer is thick, that is, if it consists of a stack of thin layers with different ionisation intensities. We find that the two curves cross at two places. This means that there will be reflection from *two* different heights in our new thick layer simultaneously. The high path is sometimes called the Pedersen ray, and is almost always ignored in the average textbook. Now we can begin to understand the dopplergram: before sunrise the ionisation intensity is too low and the ionisation profile never crosses the transmission curve, but there comes a point where the ionisation profile just touches the transmission curve at a tangent. A signal just starts to reflect over a range of heights and with a range of Doppler shifts, and this corresponds to the leading 'nose' of the dopplergram of Fig 7. As the ionisation intensity increases still further, the reflection height splits, just like the splitting of the reflection point on a drifting meteor trail, and this gives rise to the separate upper and lower arms of the dopplergram trace. The Pedersen ray eventually vanishes when the skip zone disappears completely, and we can go back to the textbook thin layer model, at least until sunset.

The amount of fold-back increases on longer paths, where the take-off angle is lower, but at vertical incidence there is no fold-back and the signal just 'chirps' vertically down from the HF side. In Fig 7 it is possible to see two more faint traces, marked A and B, which show no fold-back. They would therefore seem to be at a higher angle than the main traces. They could perhaps be reflections from a higher layer, but are more likely to be double-hop paths.

There is a similar effect at sunset, effectively a folded-back version of the fade-out shown in Fig 5. The dopplergram of a sunset fold-back event is just another way of looking at the phenomenon known as skip fading, where the heterodyne between the high path and the low path gives an audible rhythmic fading pattern, which becomes slower and deeper, getting to the zero-beat just as the signal drops out. Using the same 'area swept out' technique that we used to estimate aircraft distances, we can estimate the height change between the high ray and the low ray. In other words, we can measure the thickness of the layer.

E-LAYER

FINALLY IN Fig 7, at the point marked C, a third trace starts at 0620 UTC. This does not start with the same characteristic chirp from the HF side, and has only a very small Doppler wobble right through the day. This is the E-layer, which appears at a height of about 100km during the daytime, and is only ionised enough to reflect 3.5MHz and not frequencies as high as 5MHz, so it doesn't show in Fig 5. Unlike the F-layer, the thickness of which explains the fold-back effect, the E-layer really *is* thin. The travelling waves that were visible on the F-layer trace in Fig 5 are hardly visible on the E-layer trace, because the vertical wave movement is much smaller in the higher density E-layer lower down. Fig 7 confirms what the textbook tells us, namely that the E-layer progressively blankets the F-layer, and this gets weaker due to D-layer absorption during the middle of the day.

CONCLUSIONS

THERE ARE MANY more interesting propagation effects which the dopplergram technique can show, but which cannot be shown in this article because of space limitations. For example, a solar flare shows up as a spike of positive Doppler shift of several Hertz on an F-layer trace, with a rise-time of a minute or two and a slower negative recovery, and the effects of aurora which are audible as flutter on high frequencies can be clearly seen on the LF bands. Aircraft reflections and meteor pings can likewise still be seen on dopplergrams right down to 3.5MHz. A local beacon on 28 or 50MHz will show signs of back-scatter on a dopplergram when the band is open for long-distance propagation even though it is well inside the skip-zone, and the fascinating subject of sporadic-E propagation has yet to be fully explored with Doppler techniques.

Although frequency stabilities better than 1Hz are needed to produce dopplergrams like those shown in this article and the writer has achieved this by locking the reference oscillator on his TS930s transceiver to a standard frequency broadcast, even without this high stability the basic dopplergram shapes can be recognised. Unlike professional ionospheric research, which is done with high-power pulse transmitters that we as radio amateurs would not be able to use, dopplergrams can be produced with a conventional amateur SSB receiver, a readily available low-cost DSP starter kit, and an unmodulated carrier, which can often be provided by an existing broadcast or beacon transmitter in a suitable location. The dopplergram technique could therefore become a very useful tool for amateur propagation study. ♦