NORTH ESSEX
ASTRONOMICAL SOCIETY

Meteor Scatter Detection

IAN LAUWERYS

Introduction

- Astronomy come rain or shine, day or night!
- Detect meteors by means of means of radio scatter
- Requires Antenna, Software Defined Radio Dongle, PC or Laptop
- Low cost (£50-£60 not including computer!)
- Details in the Sky At Night Magazine
- Everything you need here including article:
  - https://www.britastro.org/radio/downloads.html

- What if I was to tell you that there was a type of astronomy that you could do at any time of day or night, and that wasn't affected by clouds or rain? Sounds like a dream right, but meteor scatter detection is cheap and easy to get in to.
- Assuming you've got a computer or laptop available, you can put together the rest of the equipment needed for 50 or 60 quid.
- Whilst we'll take a look at some of the science and maths involved in meteor detection, don't let that put you off. If you can work a hand saw, screwdriver and tape measure you've got the skills needed to get this project off the ground.
- I first came across the idea in a series of articles in the Sky at Night magazine. The information you need to get going is available from the British Astronomical Society's web site, including a downloadable copy of the original articles. All of the kit you need can be obtained from B&Q and Amazon.

Basic Premise

- Meteoroids enter the atmosphere at high speed
- At heights between 120km and 80km the bow shock heats and ionises the air creating the glowing trail we see visually
- Electrons in the ionised meteor trail can reflect radio signals

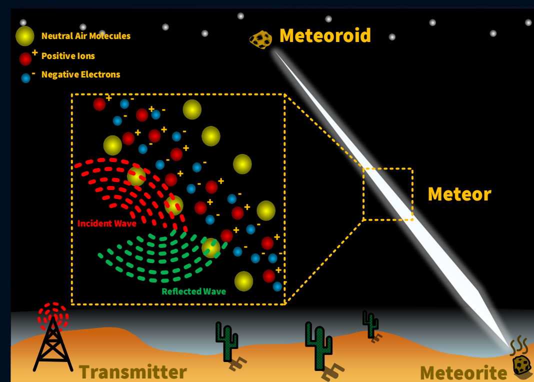

- The basic premise of meteor scatter detection is fairly simple.
- Meteoroids are small particles of space rock or metal that enter the atmosphere. The meteoroids are usually no bigger than a grain of sand, but they can travel at speeds of up to 160,000mph. That’s fast enough to get from Essex to Australia in about four minutes, but you’d end up pretty toasty if you tried it!
- At heights between 120km and 80km up in the atmosphere, the meteoroid slams in to the air molecules and compresses them. Just like pumping up a bicycle tyre, the compression heats the air. The enormous kinetic energy of the meteoroid converts to heat which ionises the air molecules, splitting electrons away from their parent atoms creating an electrically charged plasma.
- The negatively charged electrons are attracted back towards the positively charged ions. In order to recombine in to neutral molecules, the electrons have to get rid of their excess energy in the form of photons of light. This light it what we see as the glowing meteor trail or “shooting star”. By the way, it’s the same process as we see in bolts of lightning – the electrical charge from the cloud heats and ionises the air, creating a bright flash as the electrons recombine with the positive ions.
- Whilst the electrons are free from their parent atoms, they are free to move about in the hot plasma. This mobility enables them to reflect incoming radio waves just like a mirror, and this reflectivity is what we exploit to detect meteors.

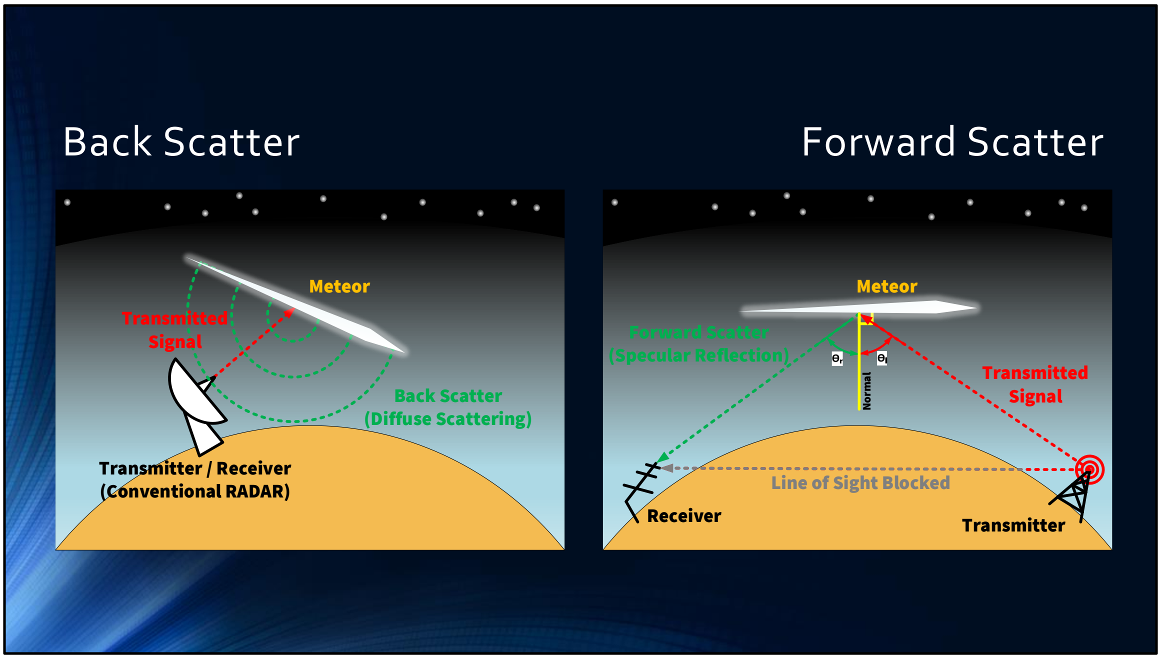

- There are two modes for detecting meteors using radio scatter.
- The first is back scatter. This is undertaken using a conventional RADAR setup. A radio signal is transmitted in to the sky by the RADAR transmitter. When a meteor enters the zone covered by the transmitter, the signal is scattered in all directions by the meteor. A small amount of the original signal ends up back at the transmitter, which also functions as a receiver. When the returning signal is detected we know there is a meteor. This signal reflection is diffuse – imagine shining a torch at a wall covered in matt paint – the light is scattered in all directions.
- The second mode is forward scatter. This is undertaken using a transmitter at one location, and a receiver at a second location some distance away. The signal reflection is specular– imagine shining a torch at a mirror – in this case you'll see that the incoming beam is reflected in a single direction by the mirror.
- In order for forward scatter to detect meteors effectively, we need to block the signal coming directly from the transmitter, and only receive it when it is reflected from the ionised meteor. Conveniently, this can be achieved by locating the transmitter far enough away that it is hidden behind the curve of the Earth, but close enough that both the transmitter and receiver can 'see' the same meteor high in the atmosphere.

- Back scatter has been used for scientific meteor research since the 1950s. A major advantage of back scatter is that by modulating (pulsing) the outgoing beam, we can time the returning radar pulses and calculate the distance from the transmitter to the meteor. We can also measure the Doppler shift of the returning radio frequencies to calculate the speed of the meteor and combine the two pieces of information to determine the direction of travel of the meteor.
- There are some disadvantages for us amateur astronomers though – firstly only a small amount of the diffusely scattered radio energy is returned to the receiver, so you need a relatively large dish to pick it up. Secondly the cost of the equipment is likely to be quite high – I haven't seen any amateur meteor projects of this type, but some people have built doppler rain radars which use the same principles, and these cost several thousand pounds at least. Bearing in mind that rain showers tend to be very large and meteors trails relatively small, I doubt it is practical for an amateur to build a transmitter powerful enough and a receiver sensitive enough to actually work for meteor detection. Finally you'd need a license to transmit on whatever frequency you are going to use.
- Don't be discouraged however – forward scatter detection is a much simpler proposition. Assuming we can find a suitable transmitter (and we certainly can in the UK), we only need a relatively simple receiver setup. The advantage is that a lot more signal is scattered via specular reflection so we don't need a big receiver dish, and since we're not transmitting we don't need any form of radio license either. The downside is that we can only detect doppler frequency shifts of the incoming signal, but not distance to the meteor itself. This in turn means we cannot determine the track of the meteor, but as we'll see, we can still do some very interesting scientific analysis.
- It's also worth bearing in mind that the geometry between transmitter, meteor and receiver has to be right to pick up the forward scatter signal, so you're only going to detect a proportion of meteors passing through the area of interest. More on this later.
- Forward scatter has actually been used by radio Hams for many years as a trick to (briefly) bounce radio transmissions much further than would ordinarily be possible, and indeed the early days of amateur meteor detecting used Ham radio receivers at the detectors. These days we have simpler and cheaper options as we'll see.
- Finally, note that for both setups shown above, there is a specular reflection and a diffuse reflection component – I've only shown one in each diagram for clarity of what we are aiming to detect, but both processes occur in each case.

Forward Scatter Transmitters

- At the average meteor height of 100km, the radio horizon is a circle of 1,150km radius.
- Therefore, the transmitter and receiver should be no more than 2,300km apart in order for an overlapping area of sky to be visible from both transmitter and receiver.
- Typically VHF Analogue TV stations were used for forward scatter detection, but most have migrated to digital broadcast which cannot be used.

- So for our forward scatter detection setup to work, we need a suitable transmitter:

- It should be transmitting continuously on a frequency that we are able to detect.
- It needs to be far enough away that it is over the horizon and not directly detectable by our receiver. At sea level the horizon is 4.7km away, but if the transmitter (or receiver) is at an elevation of only 30m, the horizon is 19.6km away. Many transmitters are much taller than 30m and deliberately located on top of hills so they can be detected over the widest possible area. On top of this one also needs to consider atmospheric refraction ("ducting") – radio waves will bend up over the horizon in the same way as light does, so it would be reasonable to assume a minimum distance of several hundreds of km between transmitter and receiver to avoid direct detection.
- The transmitter also needs to be near enough that both the transmitter and receiver can "see" the same patch of sky containing meteors. If we draw a line out from the receiver straight over the physical horizon and extend it until it reaches 100km above sea level, the average height of meteors, the line will be 1,150km long. Therefore the radio horizon at 100km altitude is a circle of radius 1,150km (again not allowing for refraction). The same is true for the transmitter, and therefore any transmitter and receiver which are less than 2,300km apart will be able to see the same patch of sky.

- For example if our transmitter was located in Gibraltar and receiver in Essex, they would both be able to see the part of the sky where their horizons intersect (like a Venn diagram), as shown in green. The closer the transmitter and receiver, the large the area of overlap. Of course meteors may occur anywhere between 120km and 80km above sea level, so really we're looking at a volume that is wider at the top and tapers towards the bottom.
- Conversely, any pair that are more than 2,300km apart will not work as they do not share any part of the visible sky. So roughly speaking, from Essex we need a transmitter that is located somewhere within the area shown in pink, but far enough away to not be seen directly.
- Historically VHF Analogue TV stations were perfect for this job. One simply tuned the receiver to the frequency of a station far enough away and when a meteor appears, a short burst of reception will occur. Unfortunately pretty much all European stations have now switched to digital broadcasting, which is not suitable for our purposes.

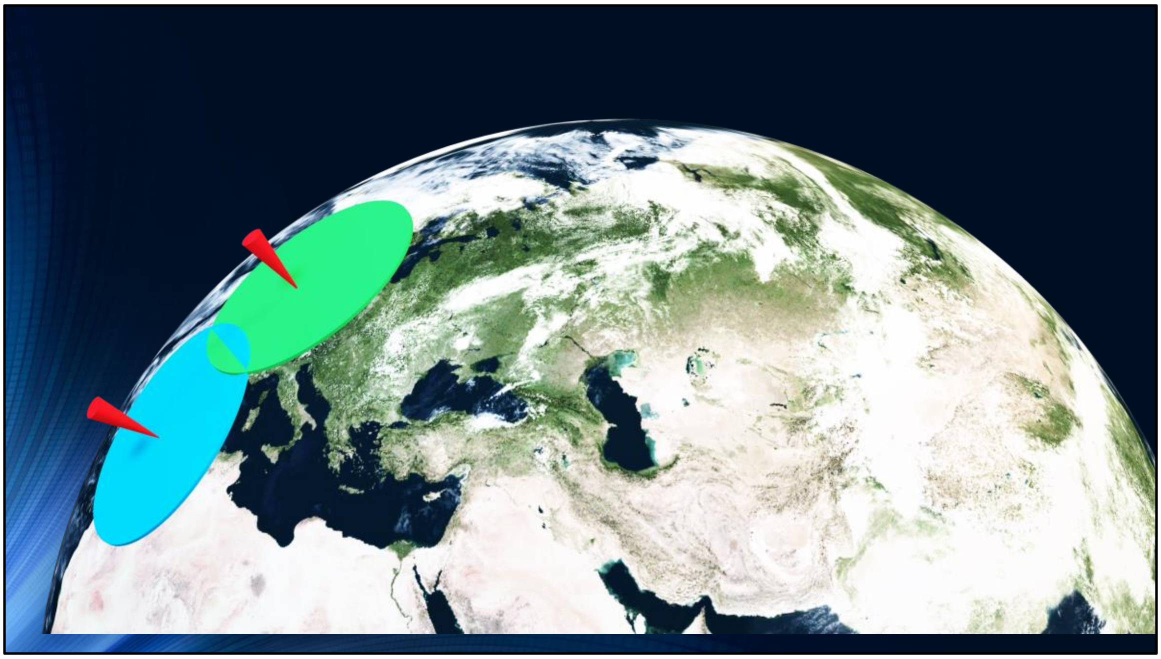

-

- There are several suitable transmitters within range of Essex, including the Brams transmitter and receiver network in Belgium and the VVS beacon also in Belgium, both of which are operated specifically for meteor detection.

- The one many amateurs use however is the GRAVES space surveillance radar which is located near Dijon in central France. This is a high powered continuous wave radar operated by the French military. It's purpose is to track satellites in low Earth orbit at altitudes up to 1,000km which pass over French territory.
- It is a 'bistatic' radar system with the transmitter located on a decommissioned airfield 36km to the East of Dijon, and the receiver at a disused missile site about 365km due south of the transmitter inland from Marseille. Unlike the US Space Surveillance System which has multiple detector sites, GRAVES uses a single detector site providing Doppler and directional information, which when combined using ingenious computer processing can automatically determine orbital elements for about 2,000 objects of interest.
- Details of the transmitter are somewhat sketchy as it is a military installation, but based on the limited public information and empirical evidence it is known that:

- The transmitter sends a single unmodulated Continuous Wave signal on 143.05 MHz, which is conveniently close to the 2m Ham radio band (making receiver and antenna design much easier due to the wealth of experience available).

- It is believed the main transmission covers an arc of 180 degrees, centred on due South, with a beam covering elevations between 15 and 40 degrees above the horizon.
- The transmitter is divided in to four 45 degree sectors, as can be seen from the photograph.
- It is believed that each of the four subsectors is further divided in to six subsectors, each of which is scanned for 3.2 seconds before moving to the next, so the total 180 degree sweep is scanned every 19.2 seconds. The scanning is achieved using a phased array radar, i.e. no moving parts are involved. Instead the many elements of the array each have a phase shifter which causes constructive/destructive interference with waves from other elements, allowing the beam to be steered electronically.

Black Helicopter Territory

- We’re in black helicopter territory here; your mileage may vary, the scanning patterns, transmission power and frequency are known to change for short periods from time to time, including of course shutdowns of the system. Generally though the system operates continuously on the expected frequency.
- E.g. Here’s the Google Street View of the receiver site, I think we can assume visitors are not welcome.

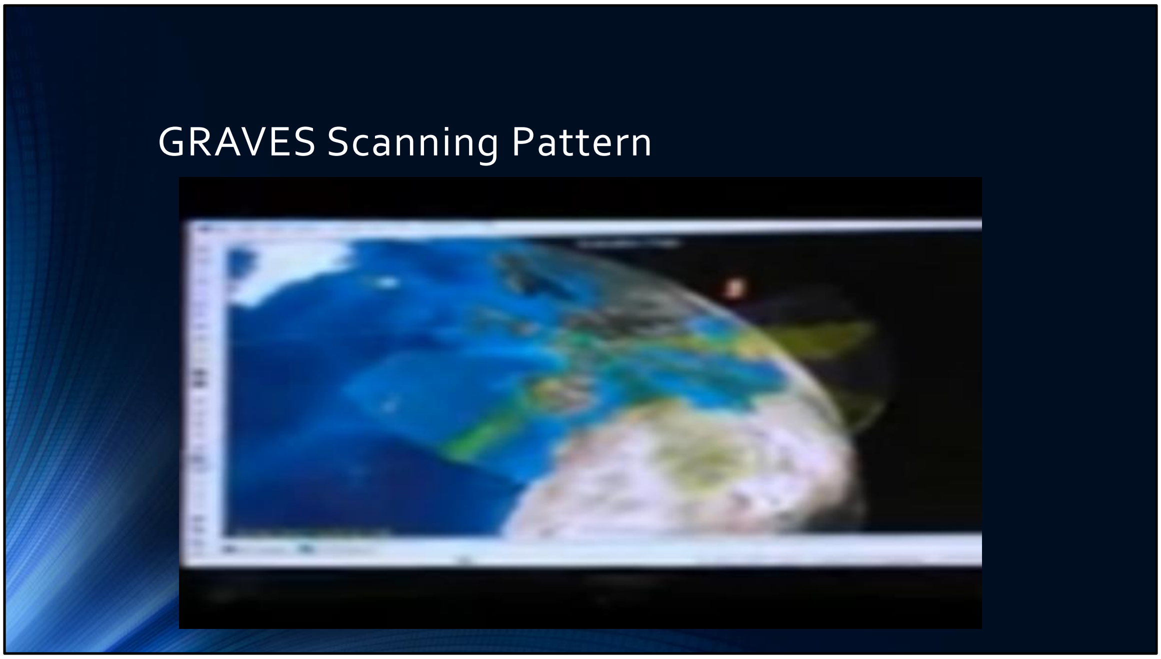

- Here is a speeded up video representation of the scanning pattern – sorry for the poor quality, but at least our guy got out alive.
- You can see the four separate beams scanning their sector of the sky.
- The scanning pattern does make it slightly challenging to interpret meteor results as we'll see later, since the radar signal is not continuous in any one direction.

GRAVES Visibility From Essex

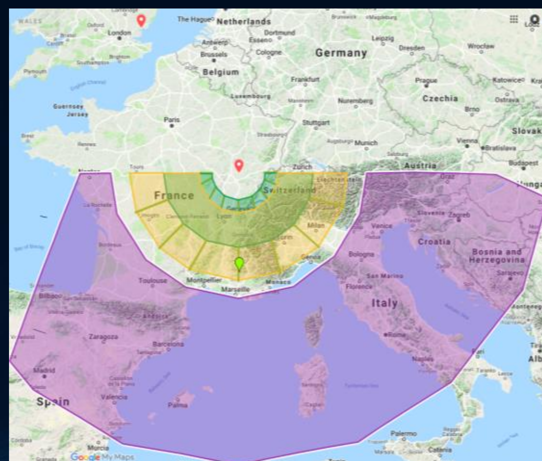

- This map shows the visibility of GRAVES forward scatter signals from Essex (red pin near top). You can see the transmitter site as the other red pin in the middle of France, and the military receiver as the green pin below it (not relevant for our purposes).
- Using what we know about the transmitter beam pattern, I’ve drawn the volume of sky illuminated by the radar. The green area represents the lower (80km) meteor boundary encompassed by the 15 to 40 degree beam elevation. The yellow area is the same for the upper (120km) meteor boundary.

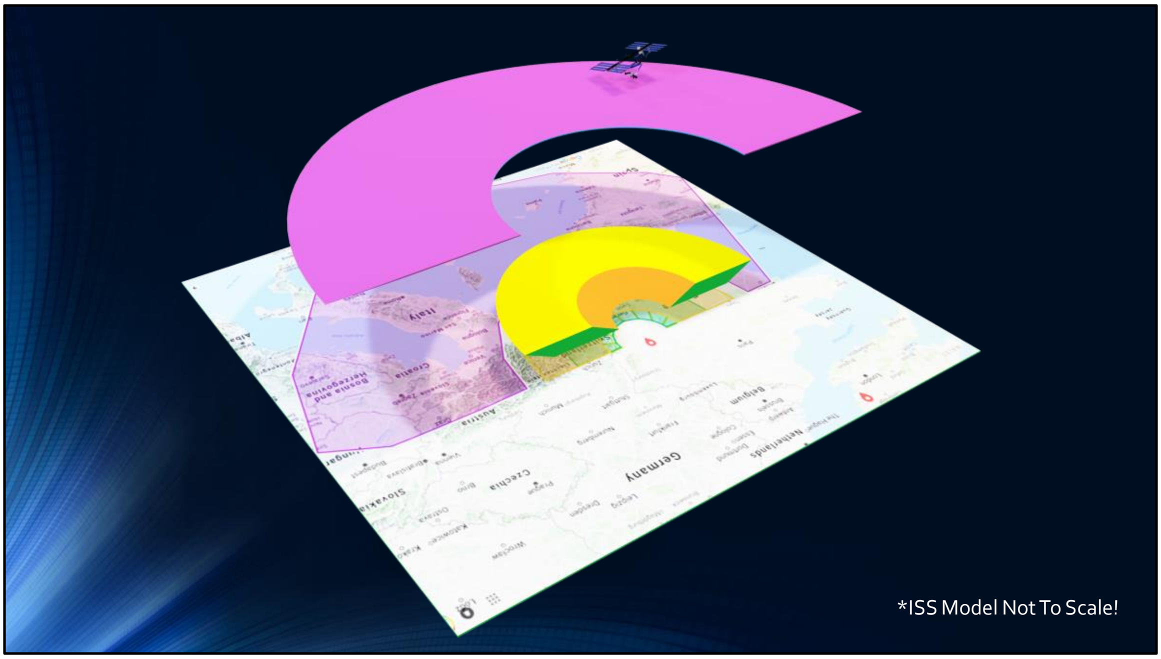

- You can visualise this as a slice of a half cone volume in the sky between the lower green and upper yellow areas. Any meteor passing through this volume could potentially pass through the radar beam and be detected. From our location in Essex, the entire volume is above our radio horizon. As you go further North in Europe and the UK, less and less of the volume is visible (remember the Venn diagram effect from earlier).
- The purple area is where the International Space Station can also be detected by amateurs, it is much higher at 408km and so can be detected successfully much further South of the UK.
- This isn’t the entire story however. Measurements of reflections from the ISS show that the radio beams from GRAVES also extend to the North of the transmitter site. This isn’t surprising, as it’s pretty hard to make a powerful radio transmitter that only sends signals in one direction. It does mean that potentially we can detect meteors over Northern France and the channel, perhaps making it possible to correlate visual/video and radio observations.
- It has certainly been possible to detect meteors via GRAVES as far North as Newcastle and Glasgow, which are well below the radio horizon for the primary transmission pattern. Atmospheric refraction will also help to extend the detection range.

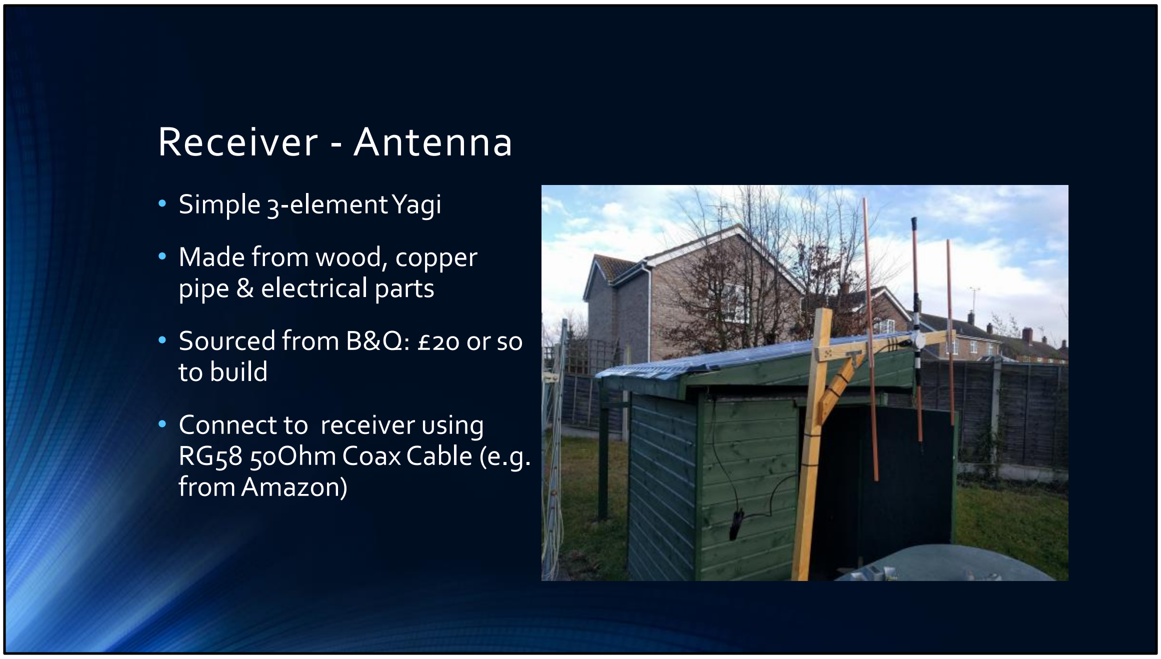

- OK so now we know what we're going to detect and where it is going to be, let's move on to the practicalities of doing so.
- As I said before, we'll take a bit of a look at the theory, but don't worry too much about it, if you can measure and cut, you can get this little project working.
- The first thing we need is an antenna. The design we will use is a simple 3 element Yagi – basically it's the same as an ordinary TV aerial just a little bit bigger as we're going to be picking up longer wavelength signals.
- Most of the parts you need can be picked up in B&Q for about 20 quid – you'll need a 3 metre length of copper pipe from the plumbing section, a few bits of plastic electrical conduit, trunking and a junction box, a bit of wood and some screws.
- The other thing you will need is a length of "RG58" coaxial cable – you can get it in various lengths from Amazon ready made with male BNC connectors, just search for "RG58 BNC Coaxial Cable". A 2 metre length is about £5 – figure out where you're going to put the antenna and how much cable you'll need to get from there to where your computer will be located.
- Coaxial cable is basically a central copper strand, surrounded by an insulator and then surrounded by an outer copper braid and more insulator. It is commonly used for TV aerials and the like, but make sure you get RG58 as this has 50 Ohm

impedance which is necessary for the antenna to work properly.

Radio Wave Basics (1)

- A radio wave is an electromagnetic wave, the same as visible light, microwaves, X-Rays, etc.
- It consists of an oscillation in the electric field (blue) and a corresponding oscillation in the magnetic field (red) at right angles to each other moving through space.
- Radio waves travel at a constant speed in a vacuum, c., the speed of light (299,792,458 metres per second).

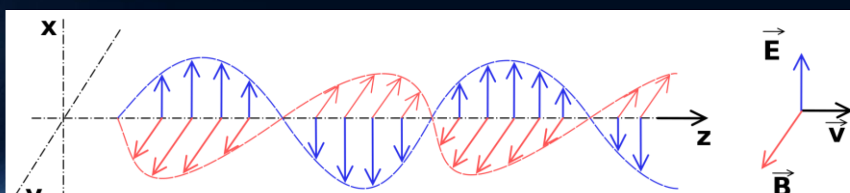

By SuperManu - Own work, CC0

- In order to understand how the antenna works, we'll take a quick look at some radio wave basics. This will also be useful later when looking at meteor detections, so bear with me.
- Firstly, radio waves are just a form of electromagnetic radiation, the same as visible light, microwaves, X-Rays, etc.
- They consist of a oscillations in the electric and magnetic fields at right angles to each other moving through space. The electric oscillation creates the magnetic one, and vice versa so the radio wave propels itself across space of its own accord.
- As you know, light travels at 300 million metres per second in a vacuum, and the usual mathematical symbol for the speed of light is "c".

Radio Wave Basics (2)

- EM waves have three basic characteristics. Firstly their wavelength or frequency.
- Wavelength = distance between successive peaks in the electric (or magnetic) field. Radio wavelengths are metres or more, visible light is in nanometres (thousand millionths of a metre).
- Frequency = number of peaks that pass a fixed point in one second, measured in Hertz (Hz).

$$f = c / \lambda$$

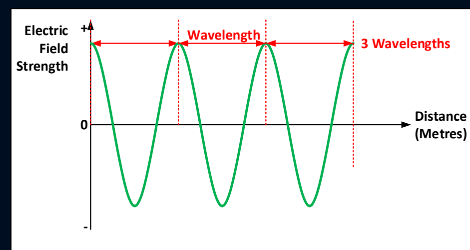

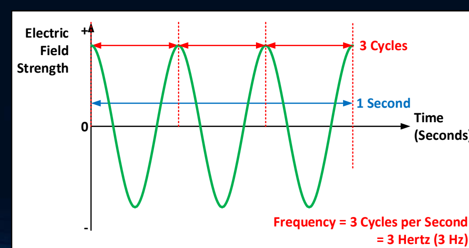

- Each electromagnetic wave has three important characteristics.
- Firstly the wavelength or frequency. These are just ways of measuring the same property:
- We can measure the physical distance between successive peaks in the electric field. The same property applies to the magnetic field, but the electric field is the one we measure using radio equipment The distance between successive peaks is called the wavelength. Radio waves have wavelengths in the range of metres or more. Light has wavelengths in the range of hundreds of nanometres (i.e. thousand millionths of a metre). As we move up to X-Rays and Gamma Rays the wavelengths get shorter and shorter.
- Alternatively, since electromagnetic waves are moving, we can measure the number of peaks that pass a fixed point in one second. That is called the frequency of the wave and is measured in cycles (peaks) per second, otherwise known as Hertz. In our example here three peaks pass by in one second so the frequency is three Hertz (infeasibly low in the EM world but just for illustration!).
- We know the speak of light is fixed, and so the frequency (in Hertz) is the speed of light (in metres per second), divided by the wavelength (in metres). Thus showing that frequency and wavelength are the same quantity just expressed in different units.

Radio Wave Basics (3)

- Secondly we need to know the Amplitude of the radio wave. There are different ways to measure this, but simplistically it is the difference in field strength between successive peaks and troughs of the wave.
- Finally we need to consider the Phase of the wave. This is simply a matter of timing – two waves of the same wavelength where the peaks and troughs arrive at the same time are "in phase", otherwise they are "out of phase".

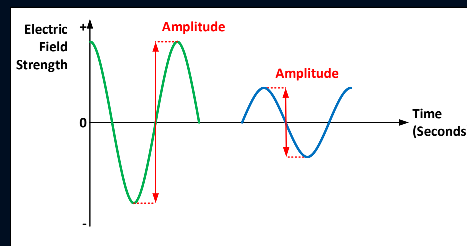

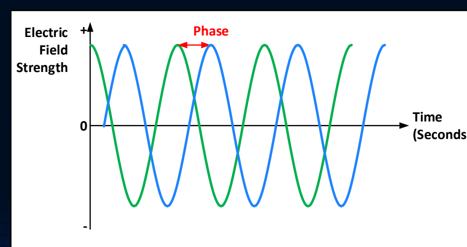

- The second property we need to think about is amplitude. This is simply the strength of wave. For a simple sine wave like this, we can measure the height between a peak and trough to determine the amplitude. For more complex radio waves measuring amplitude needs a bit more thought but we'll skim over that for now.
- The final property to consider is the phase, or more properly the "phase angle" of the wave. In a simple example of two identical sine waves as here, we can think of the phase as the difference in timing between the arrival of the peak of the first wave vs the peak of the second wave, mathematically we express the phase as an angle on a circle via the basic trigonometry you learned at school.
- If both peaks arrive at the same time, the waves are said to be in phase. If they arrive at different times they are out of phase.

Dipole Antenna Principles

- Electric field of incoming radio wave (green)
- Induces an oscillating current in the dipole (black)
- Dipole is half a wavelength, creating standing wave in voltage (red)

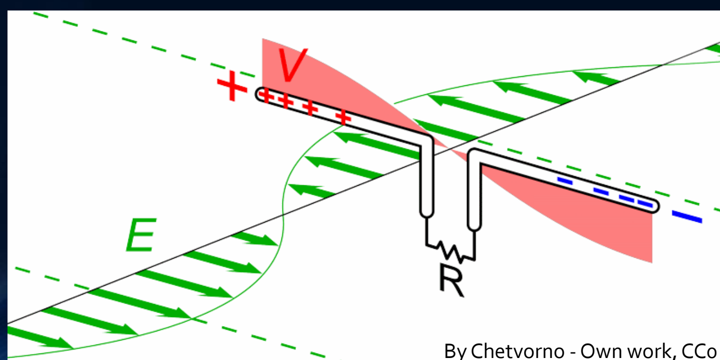

- Now we know a bit about radio waves, let's get back to the antenna. Most antennae are based on a "dipole" design or some variant thereof. This is basically a length of conductive metal (copper, aluminium, etc), or in our case two equal lengths arranged in a line.
- As you now know, a radio wave consists of a changing electrical field and a changing magnetic field travelling through space together as a single electromagnetic wave.
- In order to detect an incoming radio wave, we're interested in the changing electric field – shown here in Green.
- As the electric field passes the metal dipole, it induces an electric current in it. You can see the voltage across the dipole changes in synchronisation with the changing electric field.
- First one half of the dipole becomes positively charged and the other half negatively charged, which causes a current to flow from one side to the other (shown in black). The dipole and cable are arranged so that the current has to pass through our radio equipment on the way.
- As the electric field oscillates, the charges reverse sides on the dipole and flow the other way, backwards and forwards in time with the passing radio wave.

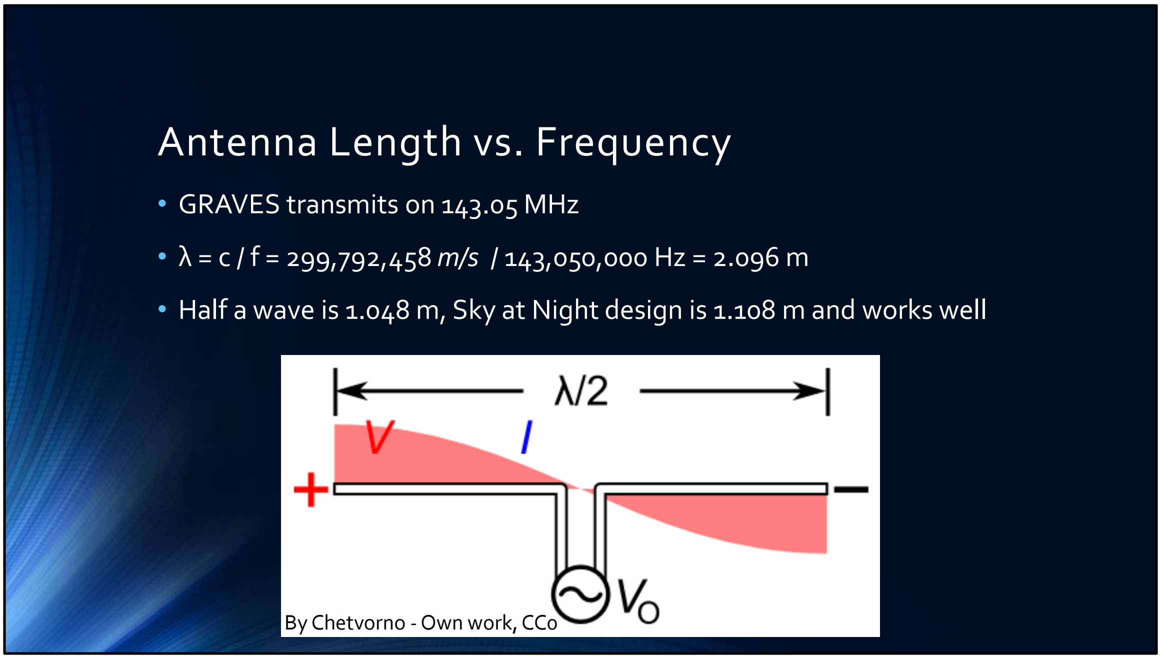

- In order for the dipole to be sensitive to a particular radio frequency, it' overall length needs to be matched to the wavelength of that frequency. If the lengths match, then a standing wave will build up in the dipole, i.e. it will resonate at that radio frequency.
- This is the same principle as sound in a musical instrument, e.g. a guitar string will vibrate at a particular frequency depending on its length. Shorten the string and the resonant frequency increases, making the pitch of the sound increase.
- So we know that the GRAVES transmitter transmits on 143.05 MHz, which means that the radio wave oscillates a bit more that 143 million times per second.
- We know that radio waves travel at the speed of light, which is just less than 300 million metres per second.
- If we divide the speed of light by the frequency, we can work out that the physical length of one wavelength is 2.096 metres, that's the distance between peaks of the electric (and magnetic) fields.
- In order for the dipole to resonate at this frequency, it needs to be half that distance, i.e. 1.048 metres long, i.e the distance from the peak to the trough of the wave. You can see that in action in the diagram here – as the standing voltage wave resonates across the dipole, the half wavelength distance creates the maximum possible difference in voltage between the two ends (in red), creating

the maximum possible current flow through the radio (in blue).

- In fact, any whole number multiple of half a wavelength is resonant and will work; there are reasons for choosing different lengths depending on the application, but for simple cases like this you might as well go for the half wave as it is the most compact version.
- The antenna design in the Sky at Night article we are following is actually a bit longer than half-wave and 1.108m, but it works well – I think it’s just because the author adapted an existing antenna design that was close enough to our needs.

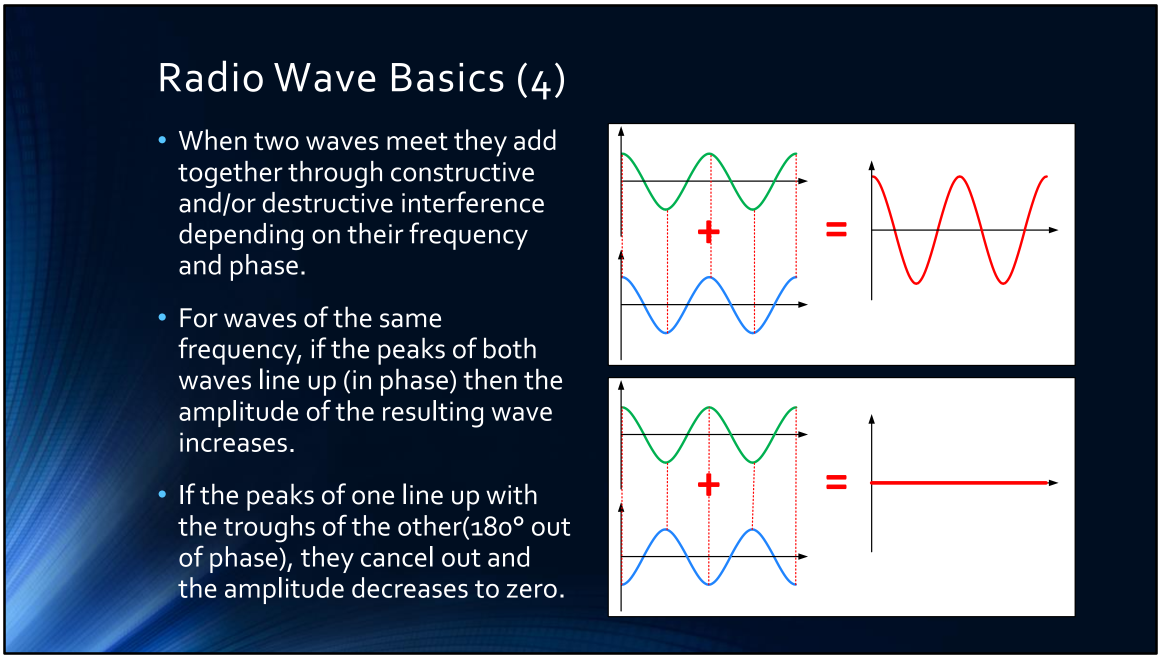

- In order to finish designing our antenna, we need to return to radio wave phases for a moment.
- If two waves meet they will add together to create a new wave through interference.
- In this simple case, we have two identical sine waves. In the top example they are precisely in phase, so the peaks and troughs line up. When they combine we get constructive interference and a new wave is created with the same frequency and phase, but twice the amplitude. (The two peaks add together, the two troughs subtract together).
- In the second case the waves are 180 degrees out of phase, so the peaks line up with the troughs. This is destructive interference, the troughs subtract from the peaks and cancel each other out.

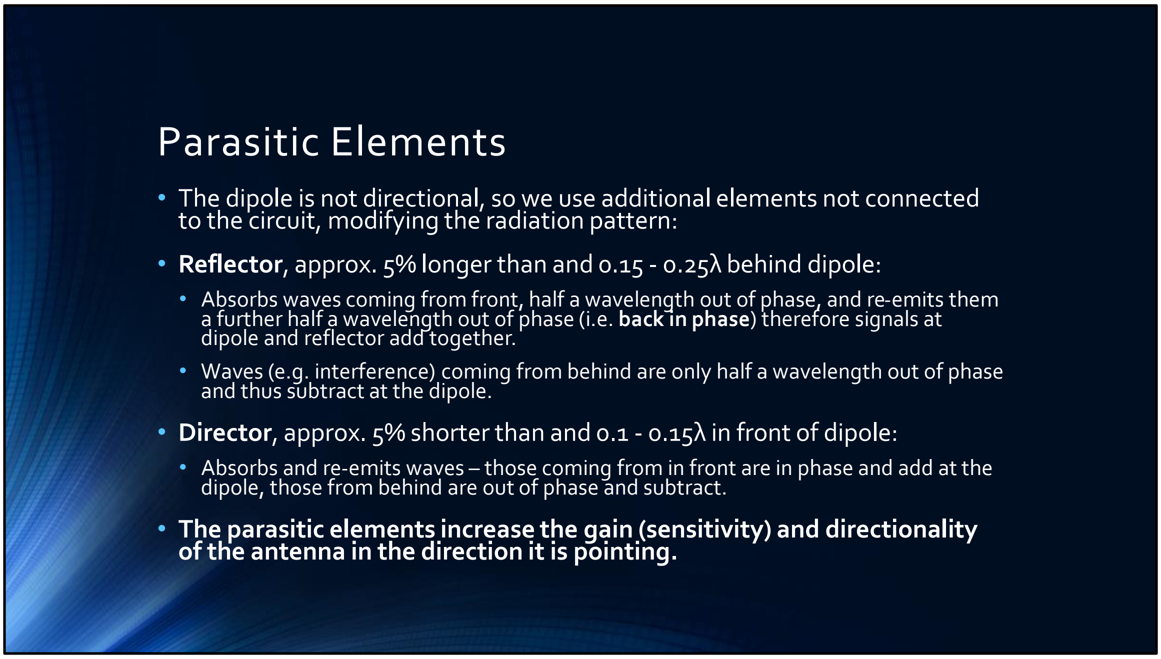

- A bare dipole antenna will pick up signals from any direction – in this case that isn't what we want.
- Instead we add extra 'parasitic' metal elements which are not electrically connected to the dipole, but do affect the way it receives signals.
- At the back of the antenna we add a single Reflector element. This is approximately 5% longer than the dipole itself, and located 0.15 to 0.25 of a wavelength behind it.
- Radio waves of the desired wavelength that come from the 'front' of the antenna are absorbed by the reflector but half a wavelength out of phase, i.e. the peaks of the waves absorbed by the reflector line up with the troughs of the waves absorbed by the dipole. The reflector then re-emits the absorbed waves a further half a wavelength out of phase, so the peaks line up again with the peaks at the dipole. This re-emitted wave is absorbed by the dipole and since the peaks line up, any waves coming from the front of the antenna add together, increasing the strength of the electric field at the dipole.
- Radio waves coming from the back of the antenna are also absorbed by the dipole and the reflector. This time the reflector re-emits them only half a wavelength out of phase so the troughs line up with the peaks of the waves absorbed directly at the dipole. The two out-of-phase waves cancel each other out, effectively blocking

any signals that come from behind the antenna.

- In front of the dipole we add a second parasitic element called a director, which is 5% shorter than the dipole and 0.1 to 0.15 of a wavelength in front. Again this element absorbs and re-emits waves in or out of phase depending on the direction they are coming from.
- The net effect of the director and reflector is to make the antenna directional – it is now much more sensitive to radio waves coming from the front that those coming from behind (which would be interference in this case and therefore not wanted). If you look at a TV aerial you'll realise you can actually add multiple directors in front of the dipole – the more you add the more directional the antenna becomes, and the gain (sensitivity) in that direction increases also. As we'll see, we don't want our meteor antenna to be too directional, so we only use one director.

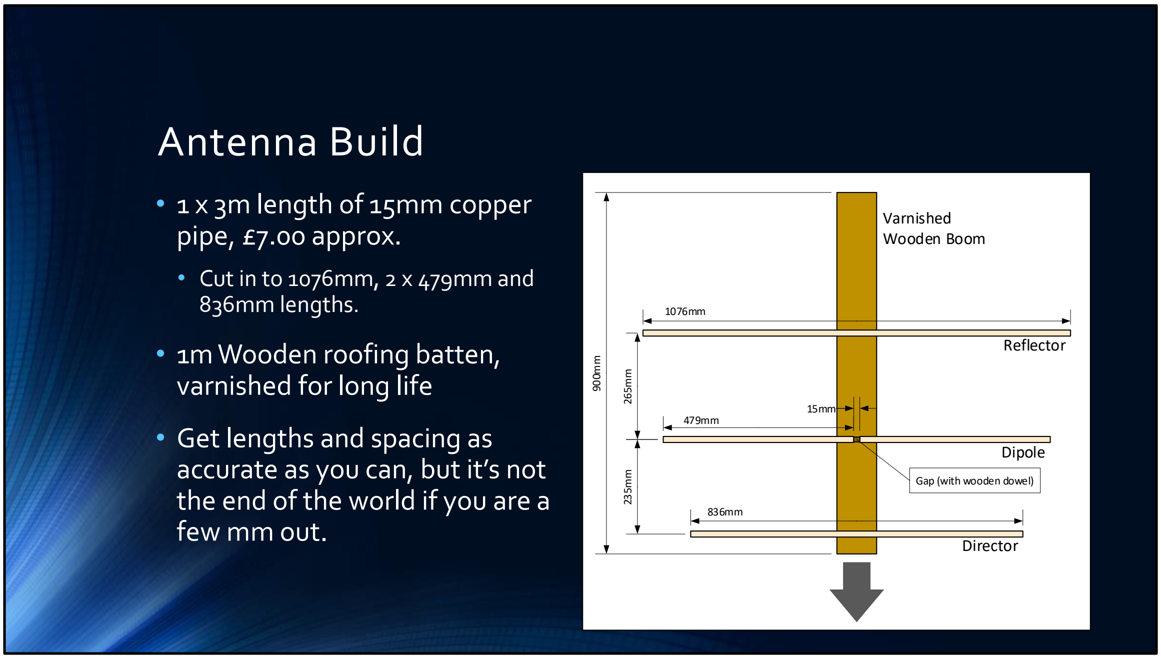

- At this stage I'd refer you to the Sky at Night article which contains full instructions on building the antenna – just download it from the BAA website.
- Tips and tricks for the build follow though.
- A single 3 metre length of copper pipe is enough for all the elements with a few cm left to spare, so measure twice and cut once using a small hacksaw. If anything, it is better to cut the elements a couple of mm longer than the final measurements, since you can file the ends of the pipe down to the precise lengths, whereas making it longer isn't possible!
- Use varnish or wood preserver on the wooden beam since it'll be out in the elements. I just used a piece of cheap roofing batten from B&Q – the boom needs to be wide enough to screw on the three-way electrical junction box, so find that first in the electrical section and hold it up against whatever timber you plan to buy to check.
- You need to get the centres of the three elements lined up down the boom so mark them on the copper pipes and mark a line down the boom as the guide. You need to get the spacing between the elements right as well, so measure and mark them on the boom and then line the centres of the pipes with the marks.
- When attaching the fittings to the boom, check that the pipes will end up parallel both from above (as shown here) and also when looking from the front down the

boom.

- Don’t worry if things are a mm or two out here and there – this isn’t NASA, but the more accurate you are the more sensitive your antenna will be.
- Don’t forget to leave enough spare boom at the rear to attach the antenna to its mounting point later. 900mm to a metre in total should be enough – much longer and the antenna will be heavy and more likely to sag over time.

Antenna Build

- 2 x short lengths of electrical mini-trunking, screw to batten and clip Director and Reflector in to them
- 1 x 3 way plastic electrical junction box, superglue 2 x lengths of plastic conduit and 1 x plastic compression gland.
- Seal all joints well to keep out water, e.g. silicone sealant, duck tape, etc.

- The article describes how to attach the various elements to the boom – just make sure you know which end is going to be the front of the antenna since the spacing between the three elements is different front to back.
- The electrical conduit and body of the cable gland can be superglued in to the three-way junction box for a solid fit.
- When screwing the junction box to the boom, put some insulating tape on the screw heads afterwards, otherwise they may short out the two halves of the dipole – there's no electrical danger, but the antenna won't work as well!
- Use as much sealant (e.g. bathroom silicone sealant is good) plus electrical or duck tape to seal all the joints. We need to keep water out of the dipole and cable to avoid shorting out the antenna.

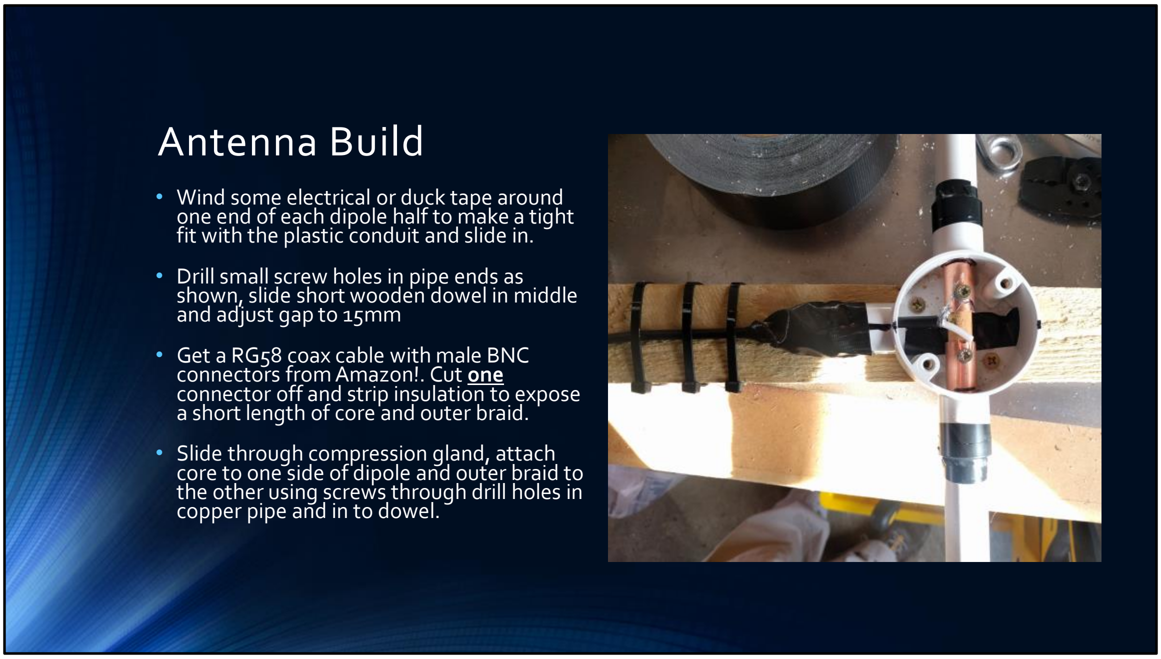

- Before fitting the dipole halves, drill a small pilot hole in the end of each to take the two screws shown – use a small drill bit as the copper is soft and the screws will enlarge the holes and bite in to the copper. You may want to put a bit of wood or similar in the pipe end when drilling to avoid crushing the soft pipe.
- You need to wind a bunch of electrical insulation or duck tape around the dipole at the point where it fits inside the round electrical conduit to make a tight fit. I put a fair amount on and then peeled it off a bit at a time until I could slide it in to the conduit for a tight fit but without forcing it too hard.
- One good trick is to put a length of wooden dowel in the middle of the two pipes as shown. Adjust the gap between the pipe ends to the required 15mm and make sure the whole thing is centred on the boom. Now screw in a couple of small screws most (but not all) of the way in. I used small wood screws for the job. The dowel keeps the gap between the elements at 15mm and helps keep everything rigid.
- To fit the feed cable, get your pre-made coax cable and cut off one of the BNC connectors, keeping the other one intact. If you know where the antenna and computer will go, you might want to trim the cable to length now as the shorter the cable the less signal is lost at the radio end – you can see from my pictures that I coiled up some of the excess just in case – if the cable is 10m or less, don't worry,

but over this length you will want to keep the run as short as possible, or use a higher quality coax cable.

- Next slide the plastic nut and the fat rubber washer from the compression gland over the cable – slide them a way down out of the way, and then insert the cable end through the cable gland body and out of the top of the open junction box. Pull enough through so that you can work on the cable as we'll pull it all back the other way to fit later.
- At the cable trimmed end, you need to strip off a length of outer insulation to expose the metal braid just inside – peel the braid off and tease apart the strands, finally twisting them all together to make a wire.
- Under the braid you may find some metal foil – strip and cut this off back to the same length as the outer insulation.
- Inside this you will find a white dielectric element (a sort of thick plastic 'pipe'). Cut back some of this to expose the inner copper core of the cable. You need to leave sufficient dielectric to keep the outer braid and inner core insulated from each other, and the whole stripped length needs to be as short as possible for best sensitivity.
- Pull the cable back through the gland until you can wind the braid around one of the screws and the core around the other one. Tighten down the screws to grip the wires on to the copper pipes. If you have a multimeter, use the resistance mode to check that there is no connection between the two halves of the dipole.
- Slide the fat rubber washer in to the compression gland body and then screw the plastic nut in to it until tight. This should squash the washer inside the body forming a watertight seal around the cable.

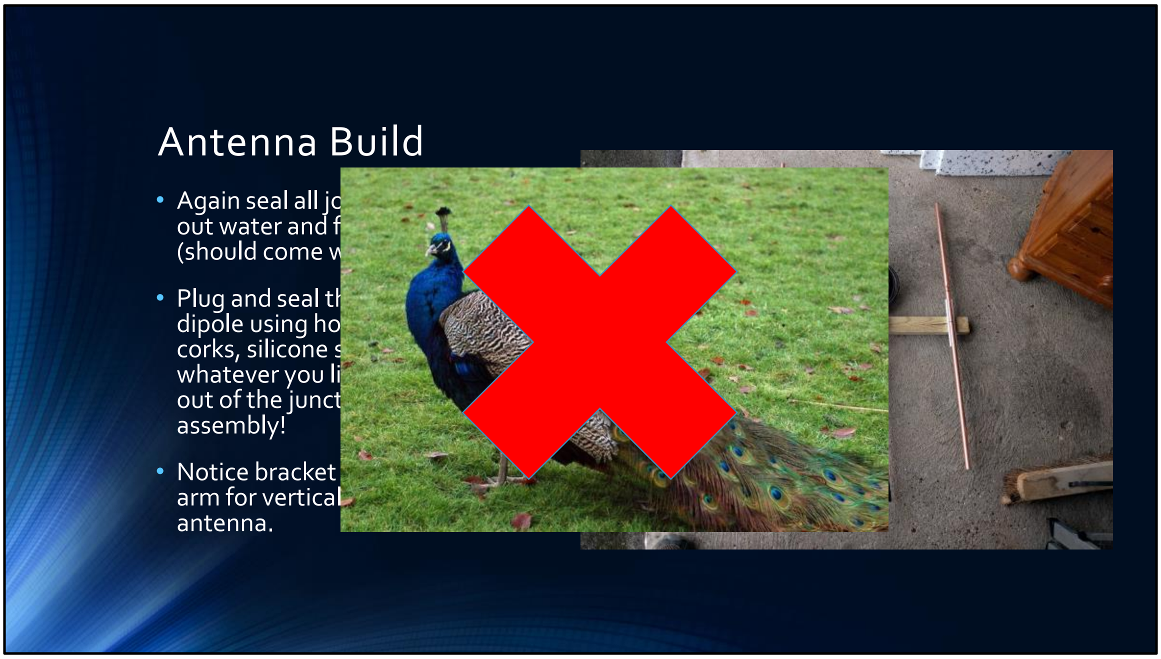

- Next add some more sealant and tape around the cable where it enters the compression gland and then use a few cable ties to attach the cable to the antenna boom so it doesn't pull the electrical connections off the dipole.
- You now need to seal the open ends of the dipole – I used a mass of hot glue and tape, but you could use rubber bungs, artificial cork or similar. We again just need to keep rain out of the junction box and cable.
- Clip the director and reflector in to the electrical mini-trunking – it's a tight fit so push the pipe in to one side first and then push on the other end of the pipe to click it in to place. Ensure the centre marks are aligned on the pipes and boom.
- You don't need to worry about sealing the ends of these two elements. You may want to use some cable ties around the trunking and pipe, to prevent them popping out of place if a really fat pigeon decides to roost on the antenna. At the NEAS observatory, Stuart the peacock decided to roost on our antenna and bent the whole thing until it was pointing at the ground!

Antenna Build

- Simple mounting on a wooden post. Needs to be at least half a wavelength (1m) above the ground.
- Elevation is not critical – anywhere between horizontal and 20 degrees above horizon will make little difference.
- Point roughly South-South-East towards GRAVES transmitter.

- The antenna needs to be at least half a wavelength above ground level, i.e. a minimum of one metre.
- You can align (polarise) the antenna vertically as shown here – I did this for safety as at this height one might walk in to the end of a pipe causing injury. You can see I've just used a wooden post, a wooden brace and some metal brackets to fix the antenna in place.
- The antenna needs to point roughly South-South-East towards GRAVES. It is not that directional so figure out where Polaris is in your garden, stand with your back to it and point the antenna a bit to the left of South.
- The elevation angle is not that important either – horizontal towards the horizon to perhaps 20 degrees above it will be fine.

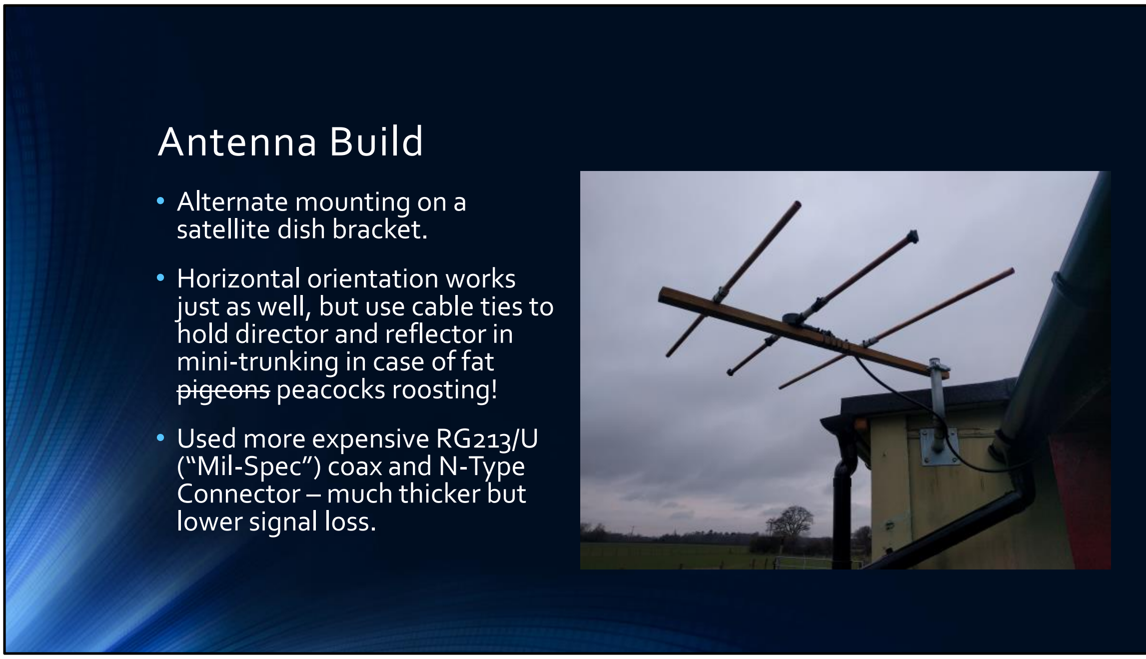

- This is the second antenna that I built for our observatory at NEAS. This time the antenna is a bit higher and a horizontal polarisation was more convenient. I used a satellite dish arm and V-Bracket screwed through the boom this time. Both of these are readily available via Amazon – just make sure the arm is well fixed to the wall (Brian at NEAS added some extra woodwork behind the wall and I screwed the arm in place using large coach screws).
- I used much thicker RG213/U or "Mil-Spec" coax cable for this build. This works well for longer cable runs as the signal loss is far lower, but the cost per metre is much higher.
- The cable is really thick and hard to work with though – it was a struggle to get it in to the junction box and I had to add crimped ring terminals to the end to attach it to the dipole. You'll also need to fit a large 'N-Type" connector to the other end, which again is quite tricky due to the thickness of the cable and metal elements.
- It is worth realising that the cable itself acts as an antenna and that a separate (unwanted) current will flow along the outside surface of the metal braid, in addition to the current we actually want which flows through the cable core and inner surface of the braid.
- To reduce this source of interference, I put a ferrite core around the cable at the point where it exits the junction box. These are available in a variety of sizes for

different cable thicknesses and clip around the cable. The mass of ferrite, which is a mix of ceramic and iron, reduces the unwanted currents flowing up and down the outside of the cable braid. This isn’t completely necessary but is a cheap and easy way to boost the system performance.

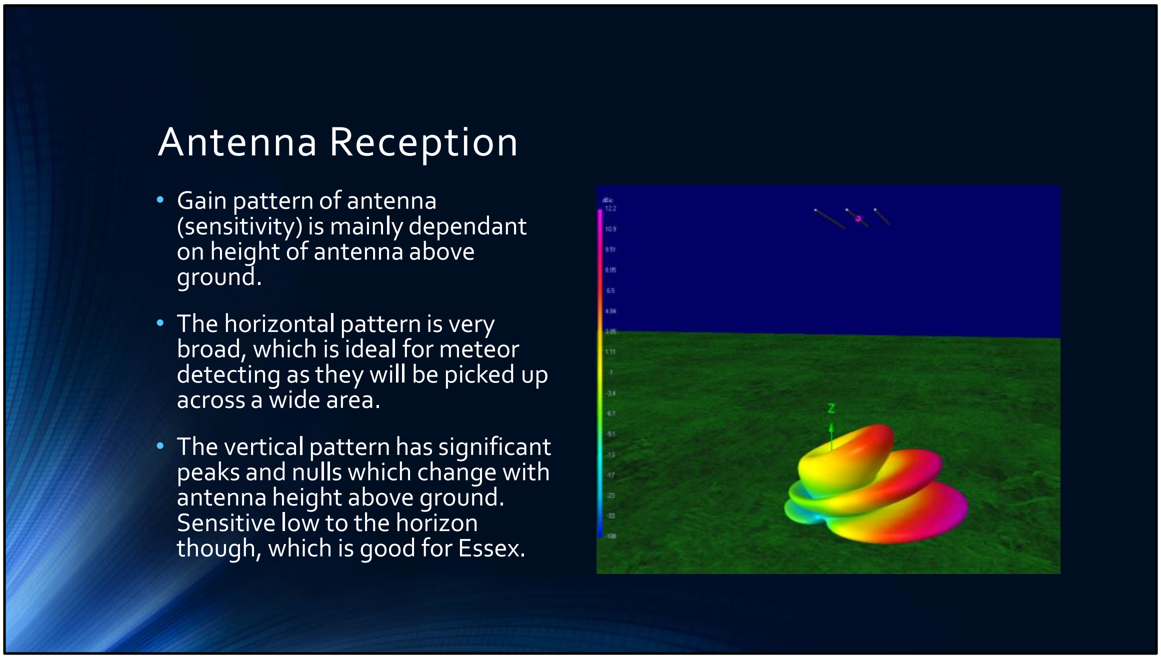

- OK so let’s talk a bit about antenna reception. We already know that the antenna should be much more sensitive to signals coming from the front that from behind.
- The actually pattern of sensitivity is quite complex, but it mainly depends on the height of the antenna above the ground.
- You can see here a 3d representation of our antenna (near the top) which is pointing left to right and located 6 metres above the ground.
- Below it is a computer model of antenna’s sensitivity to 143.05 MHz in every direction – the lobes further from the central (z) axis and coloured pink are most sensitive in that direction, and closest to it and blue/green least sensitive.
- You can see there are three large lobes of high gain (sensitivity) pointing forwards, with large ‘nulls’ at different elevation angles between them – this is unavoidable and we’re never going to cover the whole sky, but importantly you can see the biggest lobe is pointing near the horizon where we expect GRAVES meteor signals to originate.
- You’ll also see that the forward lobes are wide, which means we are covering a broad span of the horizon. This is ideal given that meteors could be detected in quite a big volume of the sky from our perspective. You’ll also notice that there is not much sensitivity to the rear of the antenna which again is what we want to prevent interference.

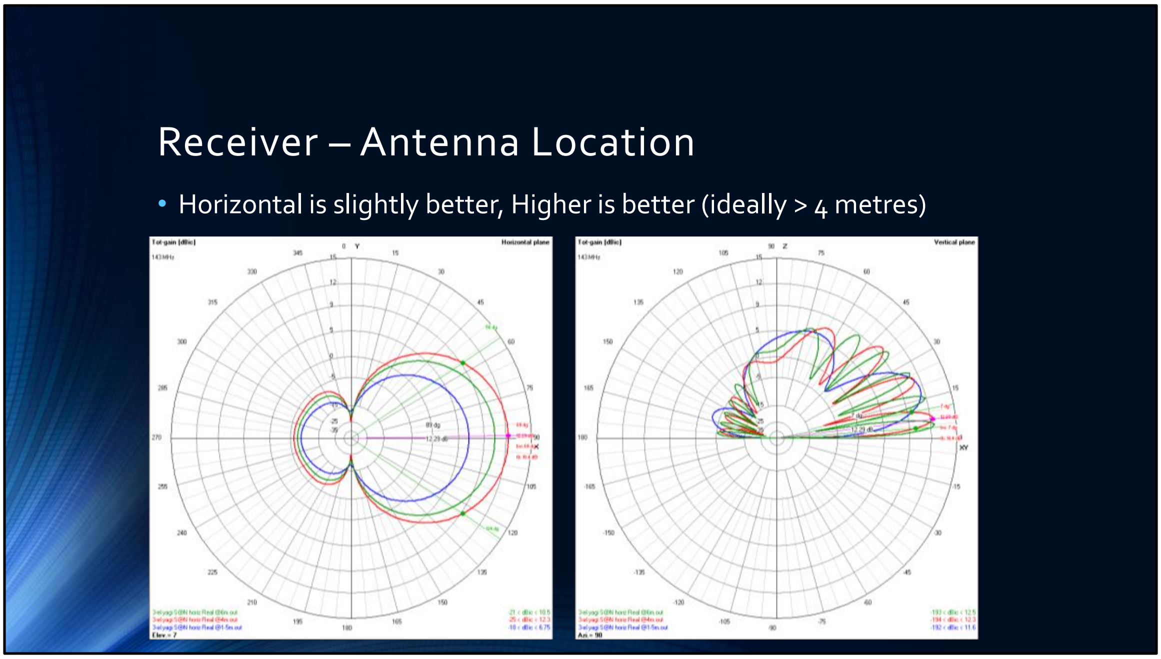

- Here's a different way of viewing the same information.
- Starting on the left, we take a horizontal slice through the reception pattern at the point of greatest sensitivity. The antenna is pointing left to right, so signals would be coming from the right side of the screen.
- I've overlaid three plots for our antenna – the blue plot shows the antenna 1.5 metres above the ground, the red plot is 4 metres above ground and green is six metres. You can see that the forward gain improves significantly as the antenna is raised above the ground.
- On the right side, we have a vertical slice through the pattern. It's a bit harder to read, but you can again see that the lobes to the right get longer as the antenna is raised higher. Interestingly the shape and number of lobes changes with height. There are more lobes and much deeper nulls at 4+ metres, but against that we should consider the improved gain of the lobes. The main reason for the changing pattern is constructive and destructive interference between waves coming directly from the source and waves being reflected off the ground in front of the antenna.
- I've also run simulations for different elevation angles and vertical vs. horizonal orientation and found that it makes very little difference to the reception pattern.
- The upshot is that you should aim to go as high as you can safely manage, but

anything 1.5 metres or more above the ground is workable.

- At this point I should say that you will want a reasonably clear view to the South. If you're right up against a neighbour's house then you're going to struggle to get a signal at ground level. A clear horizon such as we have at both our sites is ideal, and distant buildings (even tall ones) are not an issue as the signal will diffract around them. A normal house across the road is probably workable if you can get the antenna up a few metres.
- If you're really struggling, one option is to locate the antenna inside your loft. The roof will reduce the sensitivity a bit, but long wavelengths like this will penetrate inside buildings reasonably well. Measure your loft hatch diagonally across the corners though. You'll need it to be a metre and a bit so you can fit the assembled antenna through it. If not, you could of course fit the copper pipes to the boom inside the loft if there is space.
- As a last resort you could try a pole attached to the chimney as per a TV aerial, but in that case I'd be tempted to look at buying a proper 2M Ham radio Yagi so that it is robust enough to mount up there.

- If you start doing some reading on antennas and Ham radio, you'll find a lot of stress is put on tuning the antenna properly to get the correct Standing Wave Ratio, aka VSWR or just SWR. This is done either using an antenna tuner device or by adjusting the physical length of the antenna to match the impedance of the antenna and the feed line and is done once the antenna is installed in place.
- Does it actually matter though? If you're transmitting radio signals, it definitely matters since you will be pumping a fair bit of power in to the feed line and antenna. If the SWR is incorrect, then a lot of that power will get reflected back in to the transmitter equipment rather than being radiated from the antenna as planned. This power can damage the radio equipment, and at best you will be transmitting less efficiently than you would like.
- In our case we're only receiving weak signals, so there is zero danger to our equipment. If we can tune the SWR, then the sensitivity of the antenna will be at its best, but you either need an expensive radio transmitter and a relatively cheap SWR meter, or you need an expensive specialist antenna network analyser. Either way you're looking at hundreds of pounds worth of equipment to tune a £20 DIY antenna and it just isn't worth the effort unless you have access to the kit anyway.
- I did manage to get some visting Radio HAMS to test the NEAS antenna and the response was pretty good once I'd sorted out a dodgy connection inside the N-Type connector.

- OK so now we have radio signals coming down our feed line, what do we do next?
- In the early days of amateur meteor detecting, you would use a traditional radio receiver. In this case a Ham radio unit designed for the 2 metre band should work OK. The receiver would be set to Single Side Band mode and tuned near to 143.05 MHz. As meteors appear, an audible chirp will be produced by the radio. This sound can be fed in to a PC via the microphone or line-in socket and analysed using the software described later. A suitable unit would cost a minimum of £1500 and you would need to study for and obtain a Ham radio license to be able to use one legally. If you have such a unit and a license, it’s definitely an option to explore but I don’t know enough about the specifics to comment further on this approach and you probably know more than I do about this whole subject area!
- Instead we are going to use a Software Defined Radio to capture and analyse the signals. I’ll explain how they work shortly, but the upshot is that this is a small dongle that connects to the antenna at one end and to your PC via USB at the other end. The Sky at Night article recommends using the FunCube Pro+ dongle. This costs approx. £170 and is quite a high-end item and unnecessarily advanced for our purposes.
- Instead we are going to use an RTL2832U based dongle. There are lots of these available very cheaply online, Amazon, eBay, etc. Basically they use the chips from

a standard digital TV receiver dongle, so something you'd plug in to your PC in order to watch Freeview TV. The chips are pretty much standard, but the driver software is modified so that instead of being limited to receiving digital TV pictures, the computer can access the raw radio data coming straight from the dongle.

- You could use pretty much any RTL 2832U SDR dongle, just Google it, but in order to work with our chosen software you'd need to ensure it comes with an "ExtIO.DLL" driver which isn't the case for all of them. If I'm talking Greek here, don't worry. Just buy the NooElec NESDR SMArt dongle. This is £20-25 from Amazon and will work really well. It also comes in a nice aluminium enclosure which is good since the thing will get pretty hot when it is in operation. Note that the NESDR's are sold in a variety of kits with extra antennas and stuff – good if you want to play around with it, but actually you only need the cheapest option which just contains the actual dongle and nothing else for this project.
- The other thing you will need are a SMA Male to BNC Female adaptor to connect the dongle to the feed line – about £3. (If you use the thick cable and N-Type connector you'll need an SMA to N-Type adaptor instead).
- You can plug the dongle directly in to the PC but actually I'd recommend getting a USB extension cable as you want to locate the dongle away from the computer to reduce interference as much as you can.

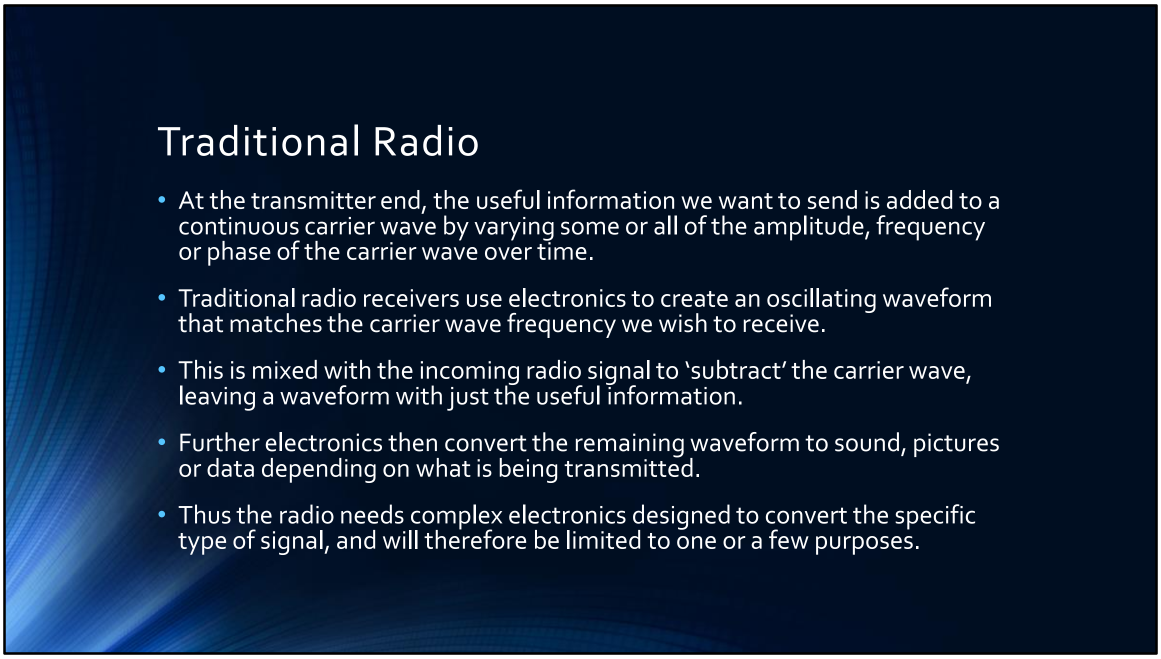

- Traditional radio receivers use electronics to create an oscillating waveform at the radio frequency we wish to receive.
- This is mixed with the incoming radio signal to 'subtract' the carrier wave, leaving a waveform with just the useful information.
- Further electronics then convert the remaining waveform to sound, pictures or data depending on what is being transmitted.
- The useful information may be added to the carrier by varying some or all of the amplitude, frequency or phase of the carrier wave over time.
- Thus the radio needs complex electronics designed to convert the specific type of signal, and will therefore be limited to one or a few purposes.

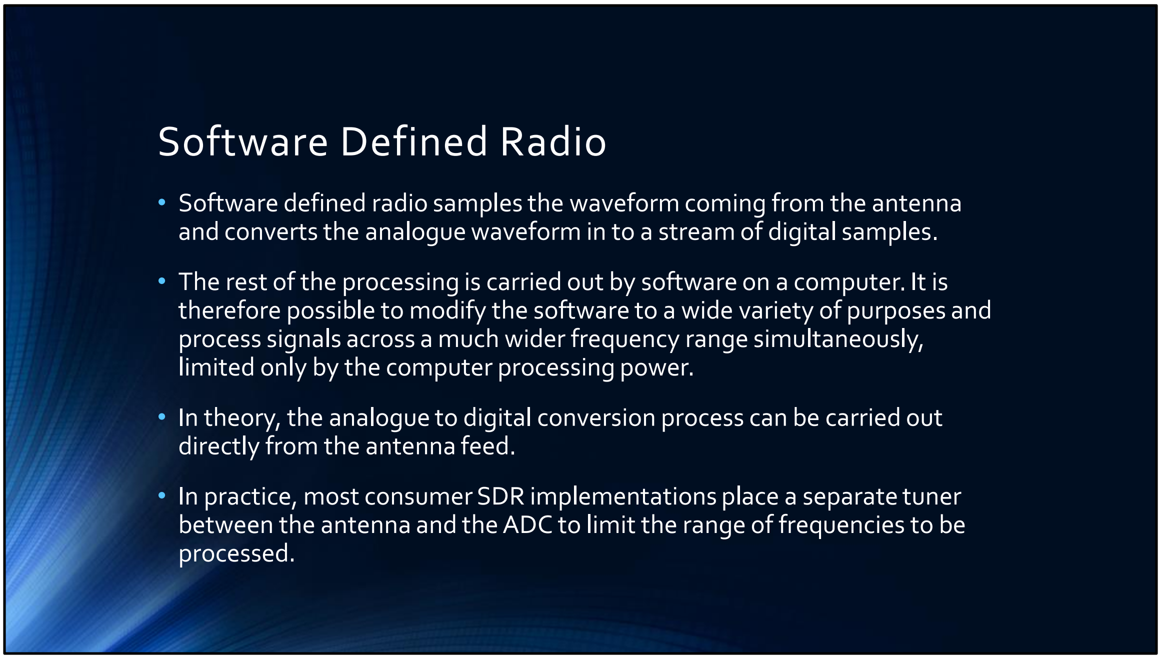

- Software defined radio samples the waveform coming from the antenna and converts the analogue waveform in to a stream of digital samples.
- The rest of the processing is carried out by software on a computer. It is therefore possible to modify the software to a wide variety of purposes and process signals across a much wider frequency range simultaneously, limited only by the computer processing power.
- In theory, the analogue to digital conversion process can be carried out directly from the antenna feed, but in practice there are limits on the sampling rate that restrict the range of frequencies that can be measured.
- So in most consumer SDR implementations, a separate tuner is placed between the antenna and the ADC to limit the range of frequencies to be processed centred around the frequency tuned to.

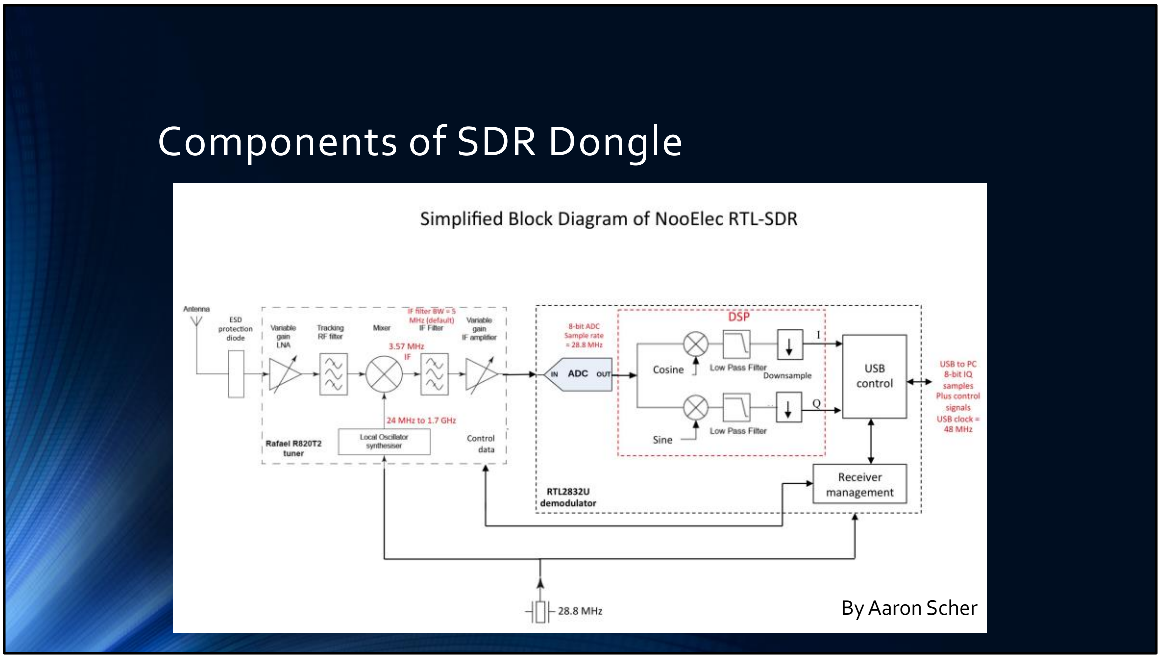

- Again, if some of this goes over your head don't worry because its over mine too. You don't need to understand how things work to use the SDR Dongle, but out of interest:
- On the left we have the feed coming from the antenna. This is passed in to an R820T2 tuner chip. The chip amplifies the signal, passes it though a filter to clean it up and then mixes the signal with a waveform coming from an oscillator. I'll explain why in a minute. The mixed signal is filtered again and amplified again.
- It then passes to the RTL2832U chip, where the analogue to digital converter circuitry measures the incoming voltage and converts it in to a stream of numbers between 0 and 255 (8 bit samples) with 0 being the lowest voltage and 255 being the highest voltage.
- These ADC samples are taken at a frequency of 28.8 MHz, i.e. nearly 29 million samples (numbers) every second. Nyquist sampling theorem dictates that the highest frequency we can measure is half the sampling rate, so this chip can measure radio signals across a bandwidth of 0MHz to 14.4 MHz. Now we know that the GRAVES frequency is 143.05 MHz, which is ten times higher, so what gives here?
- The frequency of the oscillator can be set by the computer anywhere between 24 MHz and 1.7 GHz. Mixing the oscillator frequency with the incoming radio signal

allows us to tune to a frequency anywhere between the minimum and maximum oscillator frequencies, but the output signal will just be 0 - 14.4MHz centred on the chosen oscillator frequency. which the ADC can sample successfully. Our SDR software knows what frequency we are tuned when we're displaying the results.

- In practice, the RTL-SDR has to send the samples via the USB2 connection. The USB2 clock runs at 48 MHz and each pair of I/Q samples is 16 bits, so the sample rate that can be sent at is 48 MHz / 16 = 3 MHz at most. In practice this will be reduced due to some of the bandwidth being used by USB control signals, etc. Thus the RTL2382 downsamples the IQ stream to 2.4 MHz, and Nyquist again shows that we can monitor a range of radio frequencies 1.2MHz wide.
- In order to reduce the workload for the computer, Spectrum Lab can then 'decimate' the samples. In the case of my configuration, we ignore three out of four samples. So the sample rate is 2.4 MHz / 4 = 600 KHz, i.e. a radio frequency range 300 KHz wide.
- We're really only interested in a 1 KHz to 6 KHz wide frequency range at most, so this is more than adequate for our purposes.
- Phew!

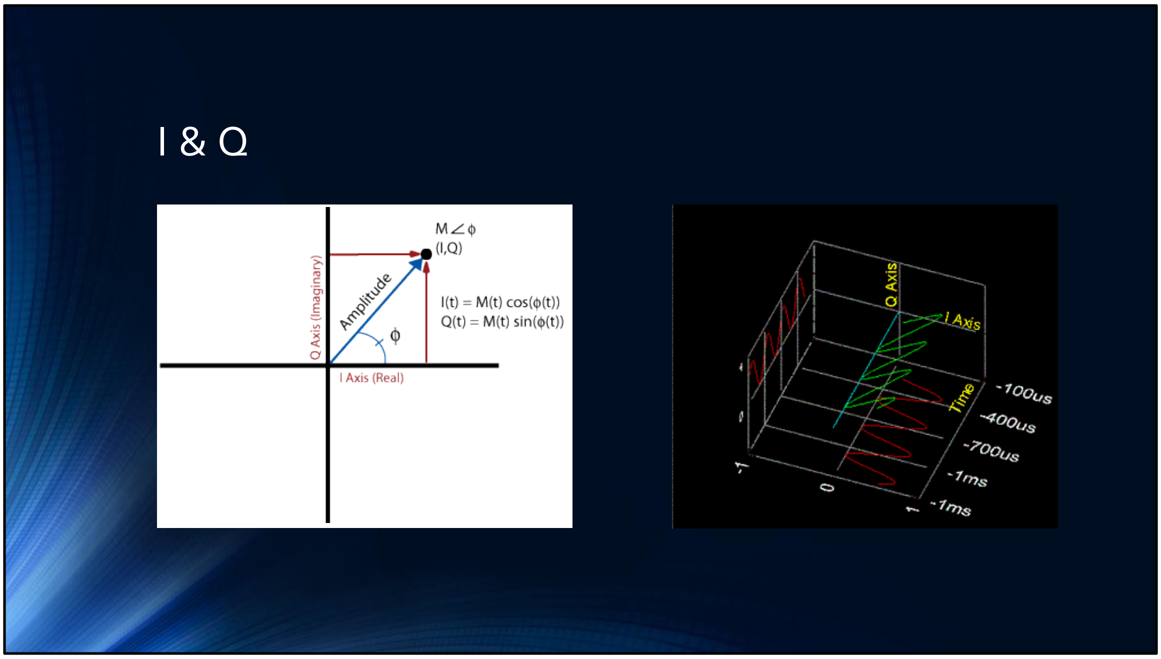

- Delving slightly deeper, the dongle's Analog to Digital converter measures the amplitude of the incoming radio wave and converts it to an 8 bit (0-255) sample. There is a bit of a problem in that the samples could be interpreted as the correct original waveform or a mirror image of it on the opposite side of the central frequency and there is no way to know which is right.
- So microcontroller circuitry processes that stream of numbers to split it in to two data streams, called I and Q.
- Basically each pair of I/Q samples are created by splitting the incoming signal in to two identical copies and adding a sine wave or cosine wave to the two copies. In ye olden days this would have been done in circuitry, but in the SDR world it is all done mathematically by the microcontroller in the dongle. This produces a set of samples that are in phase with the original signal – I, and a set of samples that are 90 degrees out of phase, called quadrature or Q.
- An I/Q pair can be visualised as the coordinates of a point of a graph. The graph on the left shows an I/Q sample pair at a single point in time. The distance from the origin to the point represents the amplitude of the waveform at that time, and the angle of the line from the origin to the point the phase angle of the wave. If we measure the rate of change of the phase angle over time, we also discover the frequency of the waveform.

- On the right the red lines represent the changing I & Q values being plotted over time. If we convert these Cartesian (X/Y) coordinates in to polar form (distance from origin and angle - green line) you can see that we get back the original waveform that was sampled, in this case a simple sine wave.
- The I/Q stream unambiguously identifies which of the two possible waveforms identified by the original sample stream is the right version. Additionally we can now process and extract the original information easily regardless of whether it was added using amplitude, frequency or phase modulation, since the I/Q values contain all that information conveniently split out and readily accessible for mathematical processing..
- Given the purpose of the RTL2832U chip was to watch digital TV, there is further processing that is normally carried between the dongle and the software driver to extract the TV pictures from the I/Q data streams ready for display.
- The clever bit of using these cheap and easily available chips for SDR is that we use different driver software that bypasses the TV processing part and just sends the RAW I/Q data straight to the computer allowing us to process it for any kind of signal we like, not just digital TV.

Software

- Many free SDR applications available, e.g. SDR#, GNU Radio
- Use Spectrum Lab – complex to set up but config files available from BAA site
- Automates detection 24/7
- Images, “Sound”, Data Logging

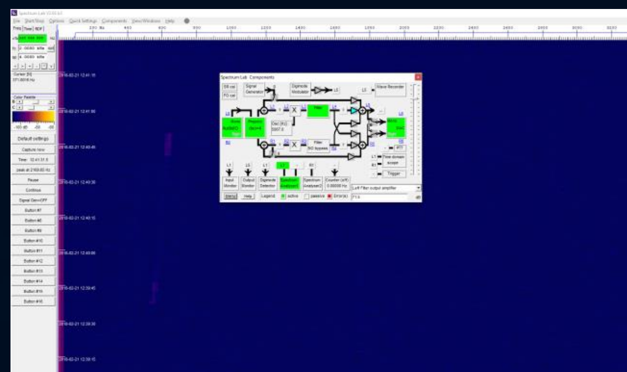

- The final step of the process is to process the I/Q data stream to extract and analyse the signals that we are looking for. There are many free SDR applications available for Windows PCs. One of the easiest to get going with is SDR#. You can download it, hook it up to your dongle and listen to the radio quite easily. There are lots of instructions online to doing so and it is a good way to verify that your dongle is working.
- The most advanced piece of free software is called GNU Radio, but I’m still experimenting with it so watch this space.
- We’re going to use an intermediate piece of software called Spectrum Lab as it allows us to automate the detection and logging of meteor data, capturing images of the waveforms, converting them to sounds and logging events to data files for later analysis. I’ll run through how to set up the software in a bit, but first lets look at the principles of what the software needs to do.

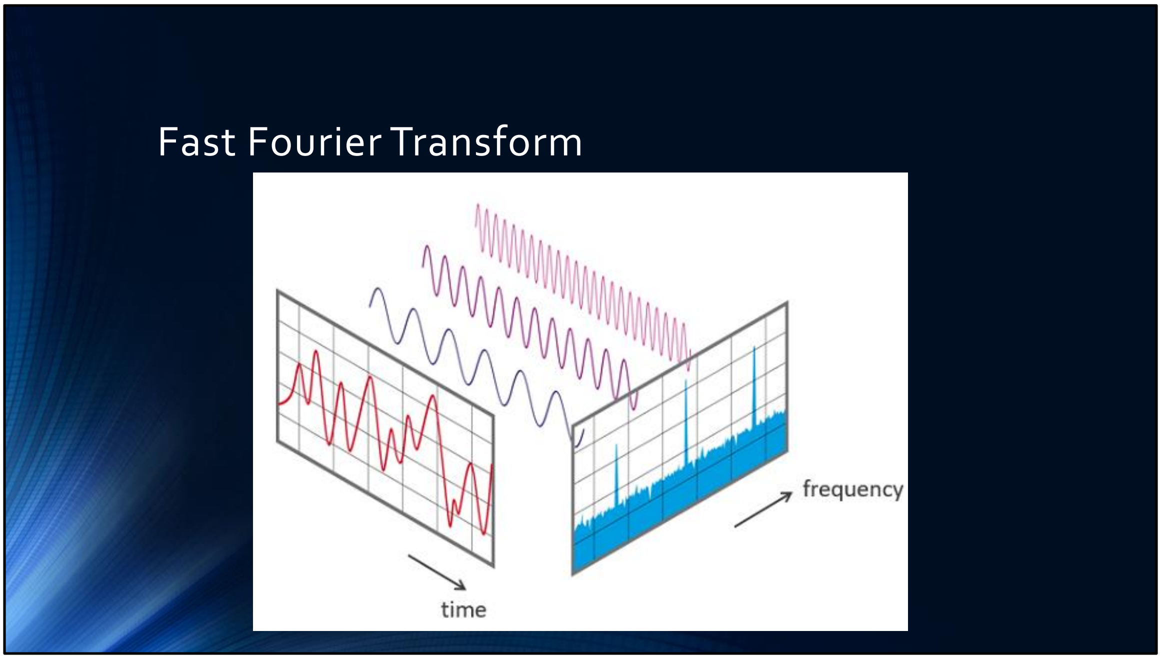

- The basic principle is that the SDR software reads the incoming stream of I/Q sample pairs from the SDR dongle. As you'll recall, the I/Q data contains all the information we need to reconstruct the original waveform received at the antenna.
- In the case of most a radio signal, that waveform will be complex as it contains amplitude data that varies over time and that is spread across a range of frequencies; in this case we'll be looking at a frequency range about 2 kHz wide close to the GRAVES transmitter frequency. In order to extract the individual frequencies contained in the signal, we need to use a Fourier Transform.
- The Fourier Transform is a mathematical process that can break any complex waveform down in to a set of simple sinusoidal waveforms. In our case we're talking about a radio wave, but Fourier Transforms can be applied to any type of information that varies across time, or indeed across space or some other variable. Fourier transforms are used in a huge variety of signal processing applications including sound, mobile phones, image processing and of course radio.
- If you look at the example, you can see that we have an apparently complex waveform varying over time. By applying a Fourier Transform, we can see that it is just the combination of three simple sine waves, with different frequencies. If we combine the three different sine waves, constructive and destructive interference

between the peaks and troughs of the waves will recreate the original waveform.

- In our case, many more waves will exist, and they will vary in frequency, amplitude and phase angle, not just in frequency as shown in the simple example here.
- Typically the Discrete Fourier Transform (set of sine waves) is calculated using a Fast Fourier Transform algorithm which makes the process computationally feasible. In order for the FFT to work, we must gather a sufficiently large set of I/Q samples to extract the constituent waveforms. As we'll see later, this means we have to make some trade-offs in the real world use of FFTs.

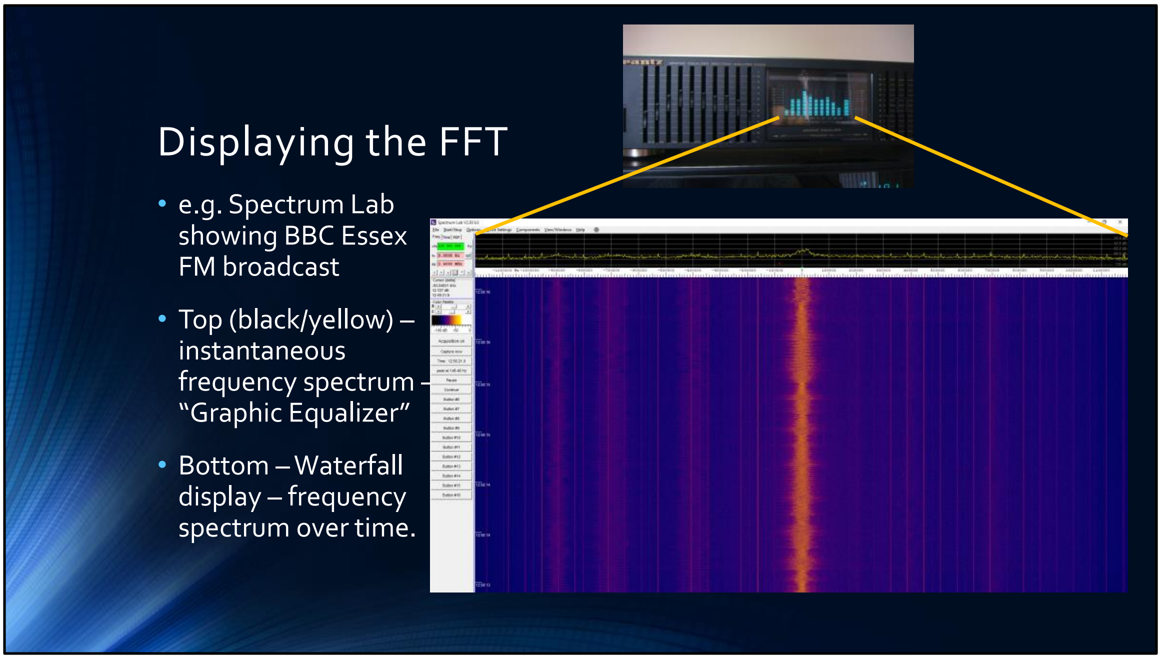

- Having performed an FFT on our signal, we can now display the results. In Spectrum Lab we can display the instantaneous spectrum, shown in black and yellow at the top here. This behaves exactly like one of those graphic equalizer displays that we all had on our HiFi in the ’80’s. The amplitude (strength) of the signal is shown as a point on the graph at each frequency ‘bin’.
- In order to visualise how the signal changes over time each instantaneous plot is also added to the “waterfall’ display below. The strength of each frequency bin is shown as a darker or brighter colour on one ‘strip’. Each time the instantaneous spectrum is recalculated and displayed, a new strip is added to the waterfall and it scrolls down the screen.
- So you can see that the FM broadcast consists of a constantly varying set of frequencies (i.e. Frequency Modulation). The frequencies run left to right in this display, and time runs from top (now) to bottom (older).
- Spectrum lab can be configured so that the waterfall scrolls from top to bottom, as shown here, or from right to left (as we’ll see later). Typically I use the right to left display as the screen is wider than it is tall, and so we can fit a longer time range on screen in that format.

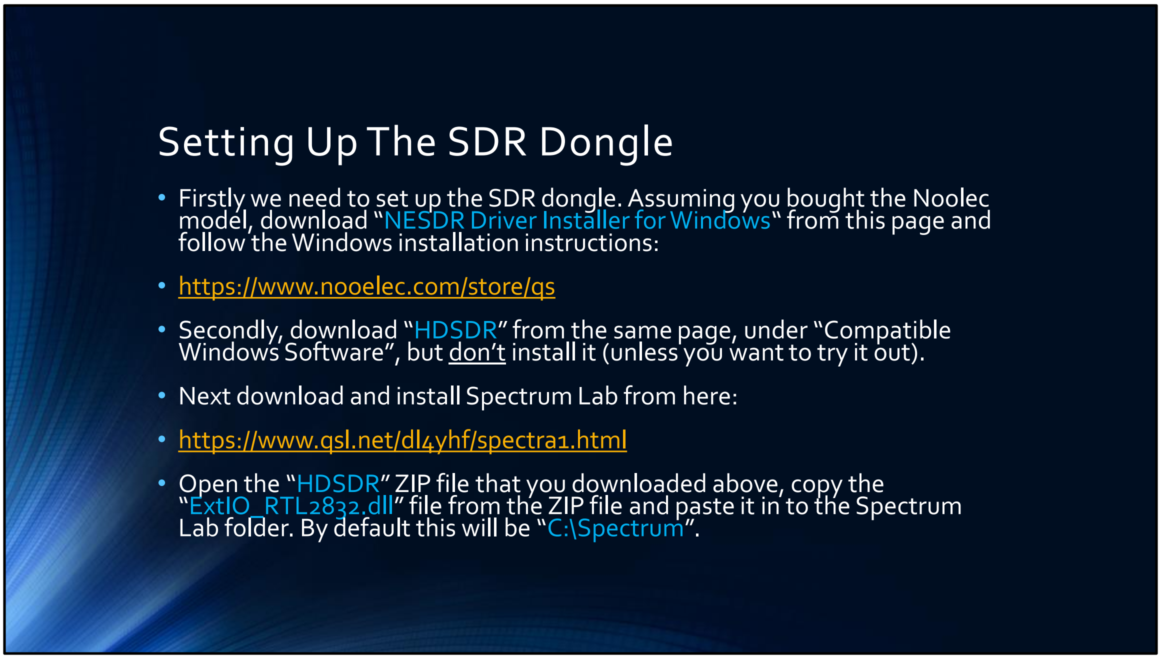

- Setting up the SDR dongle to work with Spectrum Lab requires a few steps.
- Firstly download the NESDR Driver Installer for Windows" from the page linked here. You need to follow the instructions to replace the default driver that installs when you first plug in the dongle with the special driver that allows direct access to the I/Q data stream. The 'Zadig' software does this but you have to take care to pick the right device and driver as per the instructions on the page.
- Whilst you're on that page, also download the "HDSDR" software listed under "Compatible Windows Software". You don't need to install this, but there is no harm in doing so if you're keen to try out different SDR software. If so just open the zip file and follow the instructions in the "Readme" file.
- Next download Spectrum Lab from this page and install it.
- The final step is to give Spectrum Lab access to the SDR dongle. To do this, you need to open the HDSDR zip file you previously downloaded, and copy the "ExtIO_RTL2832.dll" file from the ZIP folder and paste it in to your Spectrum Lab folder. This will be "C:\Spectrum" by default but could be different if you decided to install it elsewhere.

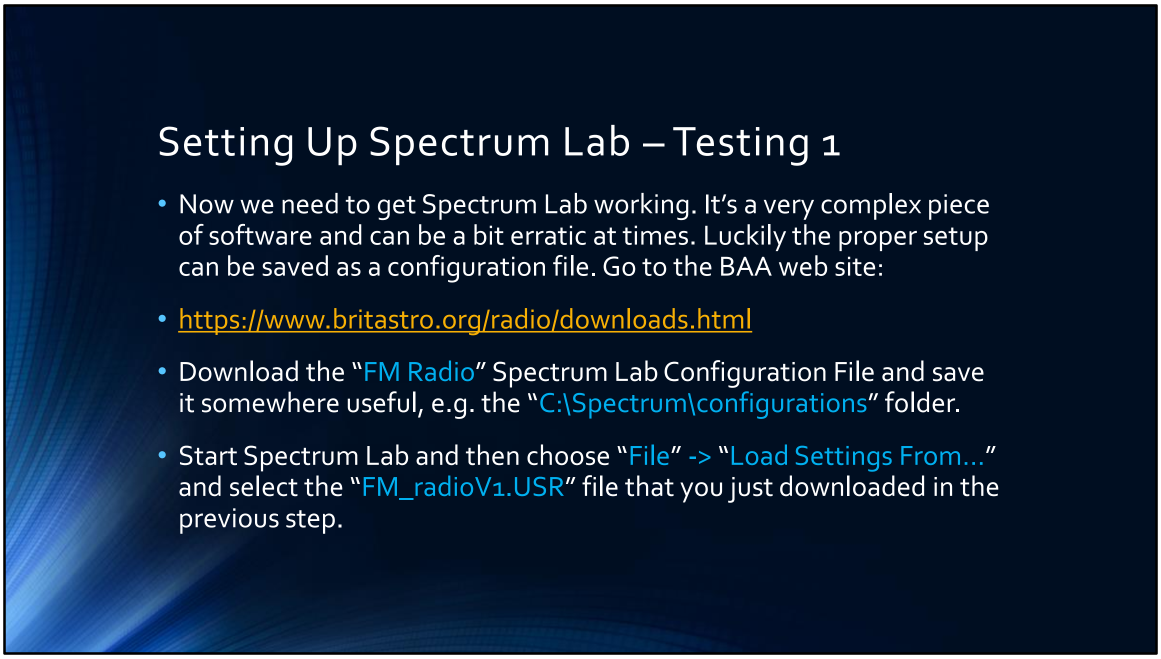

- Spectrum Lab is a tough bit of software to get your head around. It has evolved over the years and there are lots of complex screens and buttons the function of which are not always that obvious. Luckily the entire configuration of Spectrum lab can be saved in a file allowing you to get it set up with little or no knowledge to start with.
- If you go to the BAA Web site we referenced earlier, you can download the FM Radio configuration file and save it in to the Spectrum Lab configurations folder.
- Now start Spectrum Lab and load FM_radioV1.USR using the File, Load Settings From option. We’re going to start with this as it will help you easily verify that your setup is working before we go hunting for rather more elusive meteors.

Setting Up Spectrum Lab – Testing 2

- Finally we need to get Spectrum Lab talking to our SDR dongle:
- In Spectrum Lab choose "Options" -> "Audio settings, I/O device selection".
- Click the "..." button next to "Input Device", and choose the "ExtIO_RTL2832.dll" file you pasted in to "C:\Spectrum" and click "Open".
- Set "decimate input SR by" to "4" and then click "Apply".

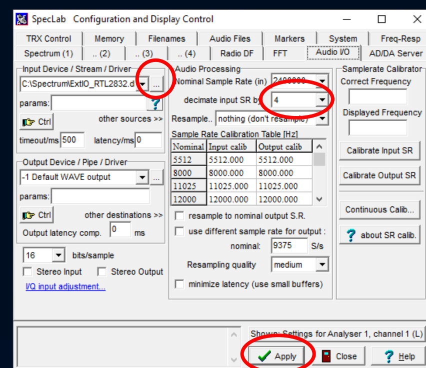

- The configuration file won't be set up for out SDR dongle, so the next step is to go in to Options, Audio settings, I/O device selection.
- Amongst the gazillion controls, find the three dots on the button near "Input Device" and click it. Now you need to navigate the to the C:\Spectrum folder, find the ExtIO_RTL2832.dll file you previously put there and Open it. The ExtIO file simply links Spectrum Lab to the incoming I/Q data stream from the SDR dongle's driver. If you want to us a different dongle, it should work provided it is supplied with its own EXTIO dll file.
- Whilst we're here, find decimate input SR by and set it to 4 to reduce the workload for the computer by only processing one in every four samples received. That reduces the width of the frequency range we can monitor, but it is still much wider than we need.
- Now click the "Apply" button to return to the main screen.

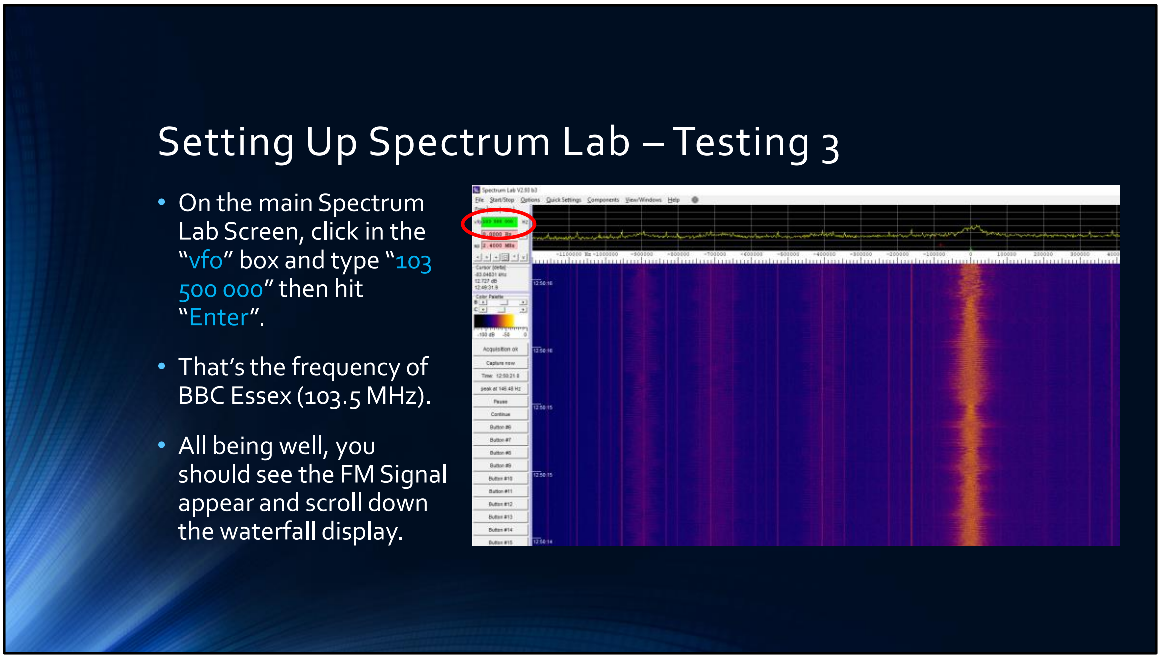

- Finally click in the "vfo" box near the top left and type 103 500 000 (including spaces) and hit the enter key to tune the SDR dongle to BBC Essex. That's the channel's frequency in Hz, i.e. 103.5 MHz.
- If things are working you should see the FM signal running down the middle of the waterfall display. Spectrum lab can't decode FM back to the audio, so if you want to listen to the radio you'll need to use SDR# or HDSDR instead.
- We now know our dongle is working. If it isn't, you're going to have to troubleshoot – check the dongle is plugged in and that you installed the driver properly, and make sure the antenna is connected. It's not optimised for this frequency but the local broadcast is strong and you should pick it up regardless.

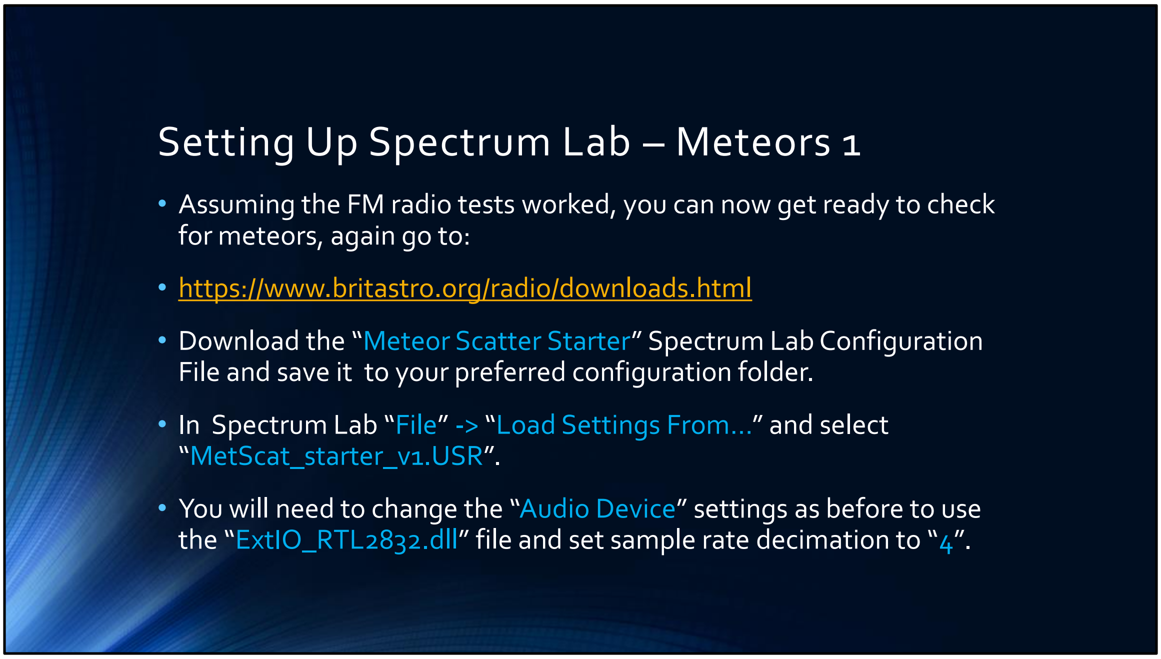

- If all has gone well so far, you can now start looking for meteors. Go back to the BAA web site and download the Meteor Scatter Starter configuration, then load those settings in to Spectrum Lab.
- You'll then need to repeat the steps as before to change your Audio Device to the ExtIO dll for your dongle and set the sample rate decimation to 4.
- You may want to Save the configuration via the File menu after applying it, so that you can just reload it later with the correct driver already selected.

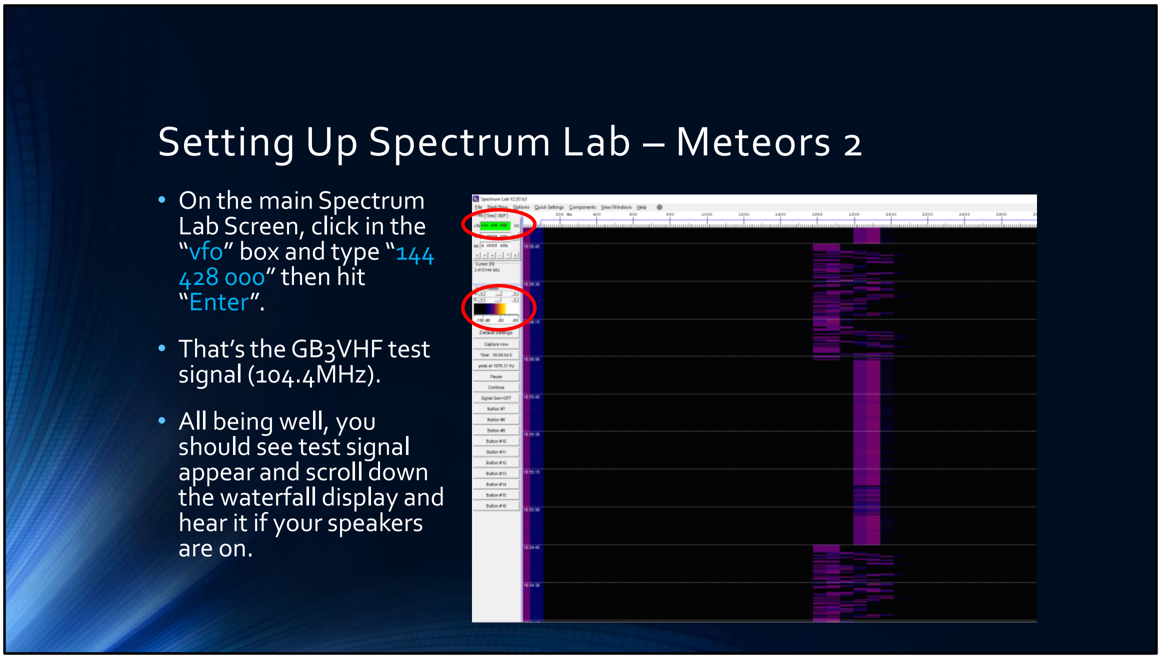

- Again click in the "vfo" box near the top left and type 144 428 000 (including spaces) and hit the enter key to tune the SDR dongle to the GB3VHF test beacon. This transmits on 144.4 MHz and is a useful test since it is much lower power than a local FM radio station.
- If things are working you should see the test beacon signal running down the middle of the waterfall display. If you turn on the speakers on your PC you may be able to hear the beacon as well. It alternates between a continuous tone, a series of Morse-code dashes and a series of rising and falling tones on a repeating cycle.
- If you can detect the beacon, you're good to go as it is a similar strength to fainter meteor echoes from here in Essex. If you're elsewhere in the country, you may not be able to detect the beacon as it is quite directional and not very powerful.
- Whilst you're on the test beacon, it might be a good time to adjust the brightness and contrast of the waterfall display. First move the contrast "C" slider left until the range of colours in the little colour display chart is spread out more. Then move the brightness "B" slider to the left to make the background of the waterfall a dark blue colour, not black as you might miss faint signals. You'll need to tweak the B and C sliders repeatedly to get the best contrast so that the strongest part of the beacon signal is nearly white and the background is a dark blue colour.

Setting Up Spectrum Lab – Meteors 3

- On the main Spectrum Lab Screen, click in the "vfo" box and type "143 048 000" then hit "Enter".
- That's just below the GRAVES transmitter on 143.05MHz.
- Now wait for a meteor to appear!

- Finally we are ready for meteors. Click in the "vfo" box near the top left and type 143 048 000 (including spaces) and hit the Enter key to tune the SDR dongle to the GRAVES radar transmitter. It transmits on 143.05 MHz so we tune slightly below that. We have 143.048 on the left of the display and this configuration shows a frequency span of 4 kHz, placing GRAVES right in the middle of the screen.
- Now we wait for a meteor. This might take some time, so be patient. You'll have more luck in the early hours of the morning through to lunchtime. Afternoons and evenings will usually be quieter so you will have to wait longer to see a meteor. We'll find out why later.

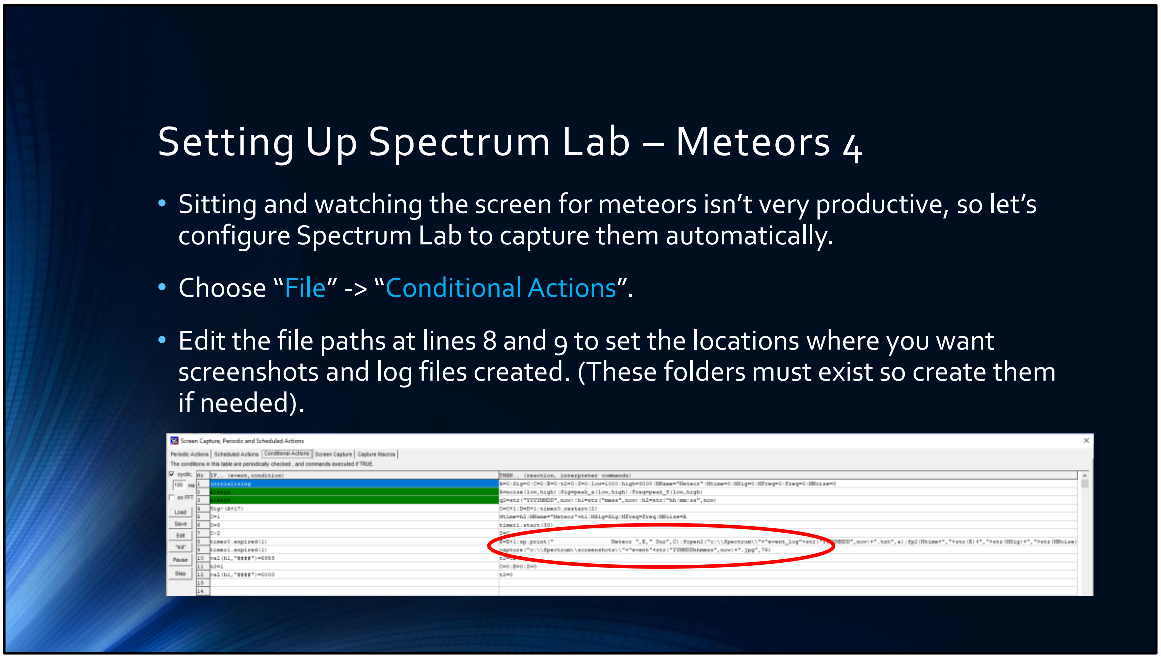

- Sitting for hours watching the screen isn't very productive, but the reason we use Spectrum Lab is so that we don't have to.
- Choose File -> Conditional Actions to see the conditional action script that is constantly running in the background. The basic meteor setup runs this script ten times per second, and for any line where the test in the left (If) column is true, it executes the commands in the right (Then) column.
- You don't need to be a software developer, but if you are you can read all about Spectrum Lab's command interpreter in the help documentation. For the rest of us, just look at lines 8 and 9.
- Every time a meteor is detected, line 8 marks it on the screen and records it in a log file in C:\Spectrum\event_log_date.txt. If you open that file with notepad you'll see a list of all the meteors recorded so far on the given date. If you want to change the name or location of the file, just edit the path where it starts C:\\Spectrum… Note the use of double backslashes (as backslash is the escape character in the interpreter).
- At line 9, Spectrum Lab creates a screenshot in C:\Spectrum\screenshots. Again you can edit the path as needed and then hit apply. Any folder you choose must already exist so create it in Windows explorer if needed.
- So all you need to do is start Spectrum Lab, load the starter configuration and tune

to the GRAVES frequency. You can walk away and when you come back later hopefully you'll find a bunch of meteor screenshots in the screenshots folder.

- If you don't get any captures, try adjusting the condition on the left side of line 4. By default meteors must have a signal to noise ratio of 17dB. Reduce the number 17 to maybe 15 to increase the sensitivity. Conversely if you are getting lots of false captures (i.e. background noise) try increasing the 17 to a slightly bigger number.

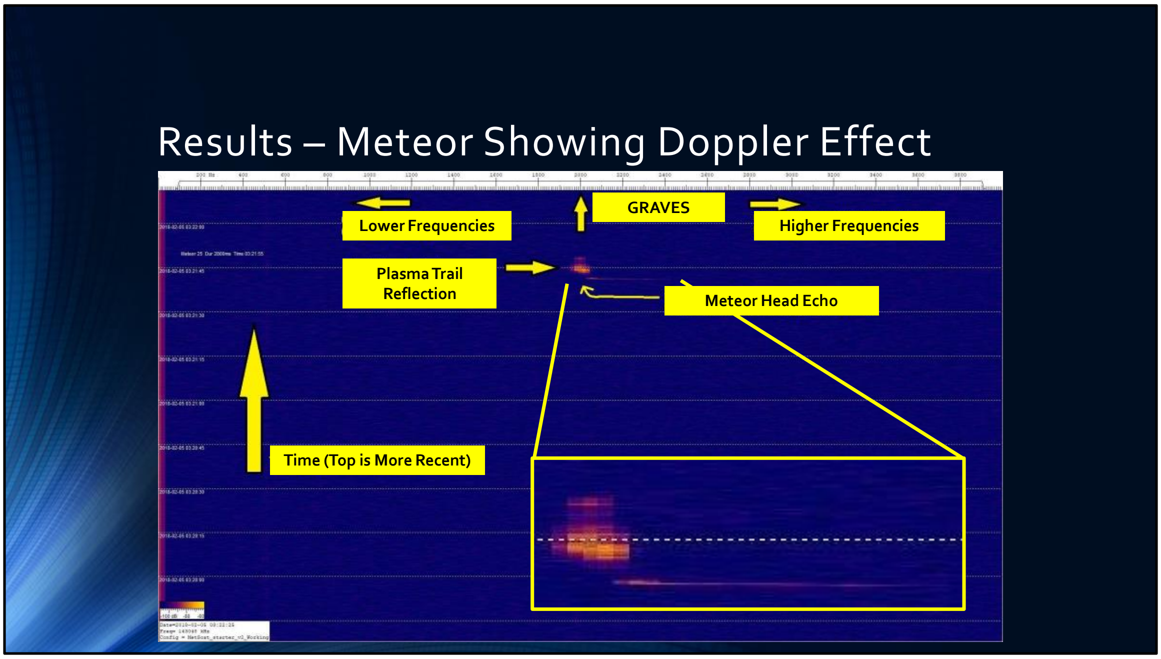

- Here's an example of a fairly typical meteor capture using the starter setup. Now it's really important to bear in mind that this is not a 'picture'' of the meteor, this is basically a graph showing the how the frequency of the reflected radar wave changes over time.
- So across the top you can see the range of frequencies (in Hz) relative to vfo frequency, which is at the left side (i.e. 0Hz relative to 143.048 MHz). The GRAVES radar frequency is 2,000 Hz above that at 143.05 MHz, so appears in the middle of the display and then we have higher frequencies to the right up to 143.52 MHz.
- The time axis is at the left here, with the dashed lines 15 seconds apart, more recent times to the top going back in to the past down the screen.
- In the middle you can see a meteor reflection – I've blown it up to the bottom right here so that you can see it more clearly.
- There are two key features that you will often see with meteor reflections. First we have the meteor head echo. That's this long streak that is sloping slightly. That tells us that the radio frequency we are receiving is changing rapidly – in this case dropping quickly from a higher frequency to a lower one in a fraction of a second. The cause of this frequency change is the Doppler effect. You'll be familiar with this if you've heard a siren on an emergency vehicle. As the vehicle comes towards you the pitch of the siren is higher. Basically the source of the sound is catching up

with the sound waves already emitted, causing the wavelengths to appear shorter from our point of view and thus increasing the pitch of the sound.

- As the vehicle passes us and starts moving away, the pitch of the siren drops noticeably, since the sound source is moving away from us making the wavelengths appear longer from our perspective.
- In the case of the meteor, the fast moving meteoroid creates a small ball of plasma around itself that reflects the radar waves. The same Doppler effect changes the frequency of the reflected radar waves as you can see here. Now it's tempting to think that what is going on is exactly the same as for the vehicle siren, i.e. the meteor is moving towards or away from us. That's true but it is a little bit more complicated as we'll see later.
- Once the meteor has flashed through the sky, it will leave an ionised trail in its wake as previously explained. This broad, strong reflection here is the radar reflecting off the wake. Again we'll take a more detailed look at what is going on here, but for now you can see that the trail is broadly centred on the GRAVES frequency, because it's fairly static and not moving much relative to us there is little or no Doppler shift. It also persists for about seven seconds in this case. Most often you'll get head echoes with no persistent trail, and sometimes a trail with no head echo but those are the two basic components of each meteor detection to look out for.

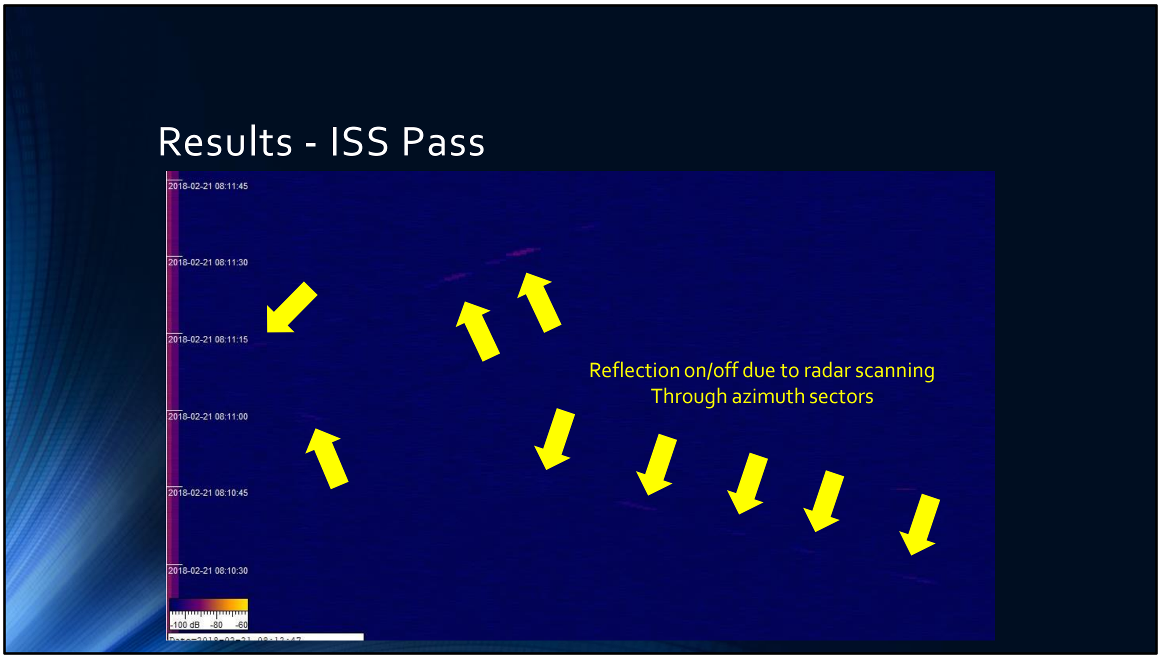

As previously mentioned, as well as detecting meteors, it is relatively easy to pick up the International Space Station at certain times.

- Here we see a capture using the BAA basic meteor setup.
- The ISS shows as a series of dashed lines as the radar scans across the sector of the sky, repeatedly illuminating the ISS as it scans across it.
- It is worth bearing this scanning behaviour in mind when analysing meteor results, since:
- Firstly some short duration meteor events will always be missed if the radar beam isn't pointing at the right part of the sky at that moment.
- Longer duration events such as persistent trails will often show peaks and troughs in reflectivity as the beam scans – this isn't the only cause of variations as we'll see, but it is a significant one.
- Most importantly, you can clearly see the Doppler shift in frequency as the station first moves towards us, and then moves away from us.
- The ISS is extremely useful to understand what is going on with these shifts since, unlike meteors, we can correlate the location and movement of the ISS with the signals we receive.

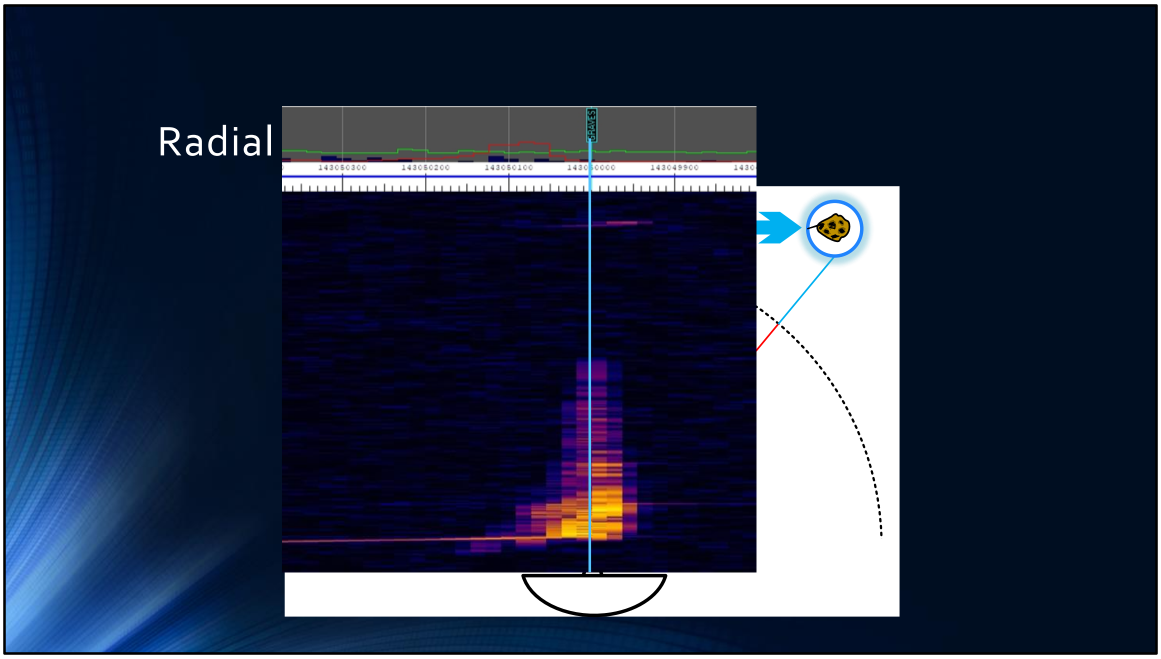

We’ve talked briefly about the Doppler shift which occurs due to the motion of the meteor, but let’s look in a bit more detail at what is going on:

- Simplistically one might assume that if the meteor is moving directly towards the observer, that the transmitted radio frequency will be shifted to a higher frequency at the receiver, or to a lower frequency if the meteor is moving away.
- That’s certainly true, but if the meteor was moving at a constant velocity towards us, we’d expect the received frequency to be shifted to some fixed higher frequency. Instead, when we look at the head echoes in Spectrum Lab, the received frequency changes rapidly over time rather than remaining at some fixed higher frequency. So what’s going on?
- One way that a rapid change in Doppler shift might occur is if the meteor changes velocity over time, so might we assume that the meteor slows rapidly as it enters the atmosphere?
- As it turns out, that isn’t the case. The slowing of the meteor is not significant enough during the short duration of the head echo to produce such a large change in Doppler shift, so something else must be going on.
- If we look at the case where a meteor is moving directly across the line of sight of the receiver we can see the cause.

- Although common sense might lead us to think that the meteor is neither moving towards us or away from us, this diagram shows that not to be true. You can see that as the meteor moves from left to right, the distance to the receiver changes. In fact the only way that the meteor could not be moving towards or away from us is if it follows a circular path with us at the centre (the black dotted line).
- Meteors travel in straight lines, not curved or circular ones, so by definition it must either be moving towards or away from us at (nearly) all times.
- In reality the path of the meteor will rarely be directly across our line of sight, and so there is also an additional intrinsic motion towards us or away from us. If we add that intrinsic motion to the line of sight effect, we can determine the meteor's Radial Velocity.
- It is the Radial Velocity that determines the Doppler shift of the radar frequency.
- In this simple case, you can think of the black dotted line as a graph showing the rate of change of the meteor's Radial Velocity. At the point of closest approach, the velocity is zero, before closest approach the radial velocity decreases at a decreasing rate, and after it the radial velocity increases at an increasing rate. At the extremes, the radial velocity is almost (but never exactly) the same as the meteor's intrinsic velocity, as our line of sight converges with its actual motion.
- You can observe exactly this effect if you watch a high-flying airliner leaving a contrail across the sky. As the airliner passes directly over head, it appears to be moving relatively quickly, but as it then recedes away from us, the apparent speed reduces due to our changing line of sight.
- If we overlay a couple of meteor head echoes on the diagram, you can see how the radial velocity correlates with the radar plot. Frequency decreases from left to right here, and time is most recent at the top and older at the bottom.
- You can see that the first (larger) head echo has been created before the meteor's point of closest approach, as it's left of the transmitter frequency (blue line). In this setup, the higher frequencies are on the left (opposite of the previous examples, sorry!) thus indicating motion towards us. It's a bit hard to see, but the head echo is ever so slightly curved, which shows that the rate of change of frequency is not constant. This is as we'd expect for a radial velocity effect, since the rate of frequency change reduces as the meteor nears the point of closest approach.
- The persistent ionised trail is (mostly) centred on the transmitter frequency, which shows it is not moving much, again as we'd expect.
- If you look at the top, you can see a second meteor head echo, this time it starts to the left of the transmitter frequency and finishes to the right of it, indicating that it starts by approaching us, passes through the point of closest approach and ends by moving away from us.

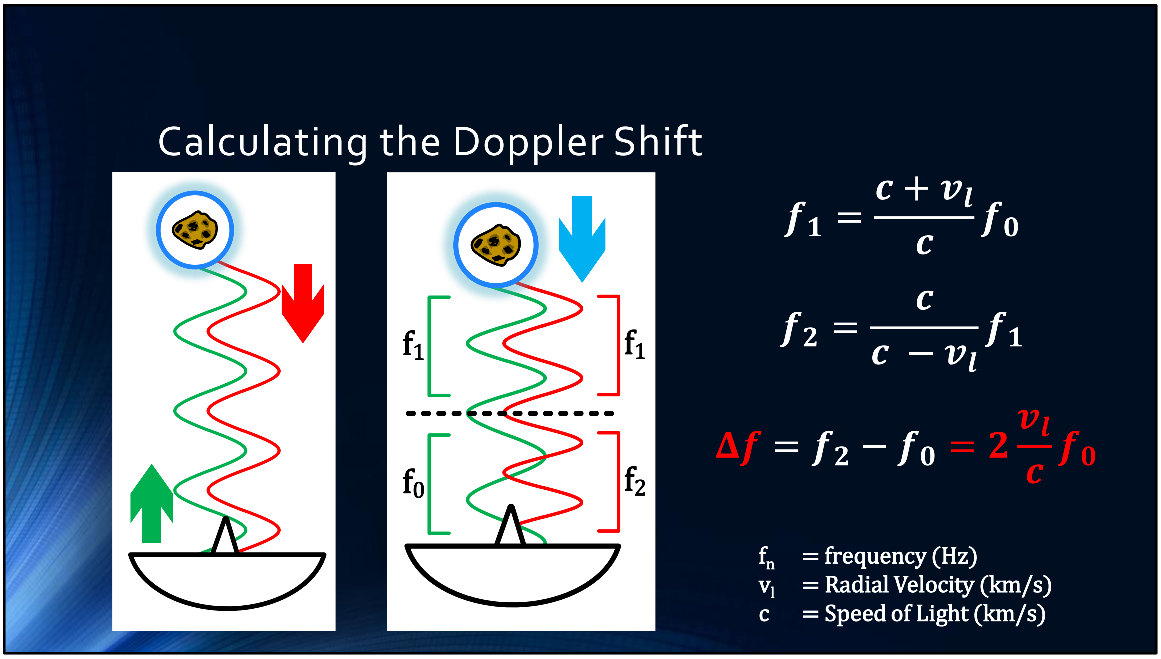

Calculating the Doppler Frequency Shift requires a bit of maths, so bear with me here.

- If an object is stationary relative to the transmitter/receiver, then of course there is no Doppler shift. The outgoing radio wave (green) is reflected by the meteor head (blue) and arrives back at the receiver at the same frequency (red).
- Now let's assume the meteor is moving towards the transmitter/receiver; I've shown it directly approaching for simplicity, but in reality we use the radial velocity, i.e. combination of intrinsic motion and line of sight effect:
- The outgoing radio wave is sent at some frequency (f0). Because the meteor is moving towards the transmitter it sees the radio wave arrive at a higher frequency, (f1). Everything is relative here – it doesn't matter if the transmitter is moving towards a stationary meteor, or a meteor is moving towards a stationary transmitter – the effect and the maths is the same.
- We can calculate the received frequency (f1) as shown in the first equation, i.e. the speed of light plus the radial velocity divided by the speed of light gives us the shift factor, which is multiplied by the transmitted frequency to give us the received frequency at the meteor.
- Because the meteor is reflecting the incoming signal, it is effectively behaving as a second transmitter, so therefore we have to calculate a second Doppler shift. So

the frequency reflected by the meteor (f1) is seen as f2 at the receiver using the second equation here.

- To calculate the change in frequency we just combine the equations for f1 and f2 and simplify them, so delta f (i.e. the difference between f2 and f0) is 2 times the transmitted frequency (f0) times the radial velocity over the speed of light.
- You will notice that the maths assumes the transmitter and receiver are in the same location, i.e. a backscatter setup. Of course with our home made receiver we are using forward scatter, so the maths would be slightly different as, for example, the meteor could be moving towards the Graves transmitter and away from our receiver. In practice these equations are good enough for our purposes as any difference in motion towards the transmitter and away from the receiver tend to approximately cancel each other out.

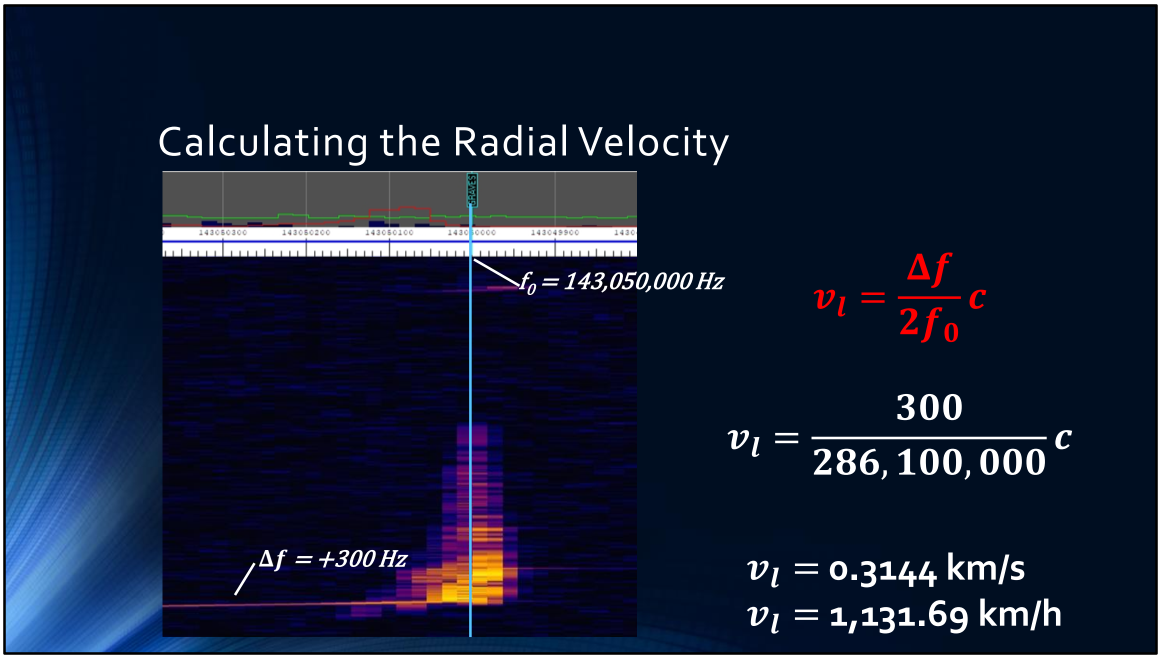

So can we calculate the approximate radial velocity of one of our detections? Let's give it a try.

- I've just rearranged the previous equation to calculate the radial velocity from a known frequency shift as shown here.
- Looking at our example, f0 is 143.05MHz, and if we pick a point somewhere on the head echo here, we can see that it's 300Hz higher.
- Plug those two numbers in the equation and we can calculate that the radial velocity at that point is 0.3144km/s, or a bit over 1,100km/h.
- The intrinsic velocity could be in the region of 200,000km/h but what we're seeing here is its relative motion towards us, not it's intrinsic velocity.
- In fact there is generally a large difference between typical meteor velocities and the radial velocities that this setup will detect:
- Bear in mind that we are looking at a relatively narrow range of frequencies here. Typically for my setup, I am looking at a bandwidth of 1KHz centred on the Graves frequency.
- Even if we widen that to say 6KHz, the maximum radial velocity we can detect is about 11,000km/h towards or away from us, so maybe a 20th of a typical meteor's intrinsic velocity. So we're only going to detect meteors moving quite shallowly across our line of sight. Anything that is coming too directly towards us or away

from us will exhibit a Doppler shift that puts it well outside the range we are monitoring.

- This may explain why we sometimes get static meteor trails appearing with no accompanying head echo – the trail isn't moving much relative to us, so can be detected, but the geometry of the head's motion is too extreme to allow us to detect it.

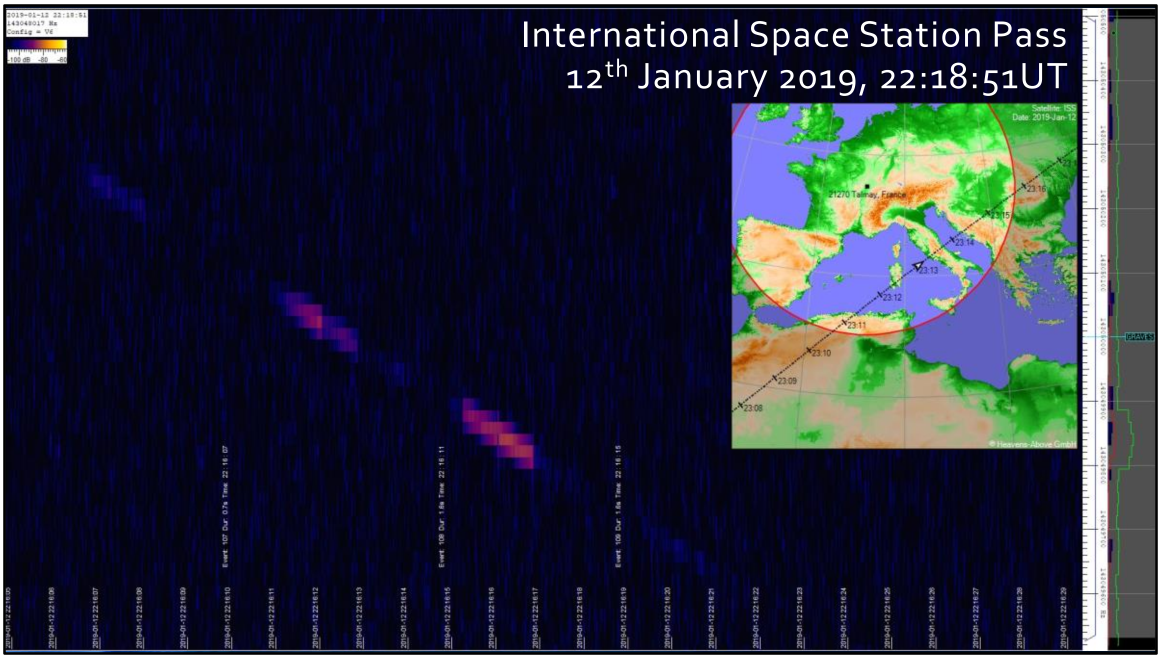

- As I hinted before, using the ISS we can get some idea of how radial velocity relates to the rate of detected frequency change.
- Since the ISS moves at a constant 7.66 km/s, any differences in the rate of frequency change are purely down to line of sight effects due to its track relative to the transmitter and receiver.
- Here we can see a pass from January 2019; I've use Heavens Above to plot the ISS's known track on a map.
- The observant will note that ground track timings are Dijon local time, i.e. 1 hour ahead of UT. The really observant will also note that there is a discrepancy of Spectrum lab clock being 8 minutes ahead of ISS predictions. This is due to Spectrum Lab timing issues, as the recorded times tend to drift from the PC clock the longer you leave SL running.
- In any event I'd think it reasonable to assume the forward (specular) scatter is centred near the arrow at 23:13 on the map. Note the angle of the changing Doppler shift as the ISS passes through this point.

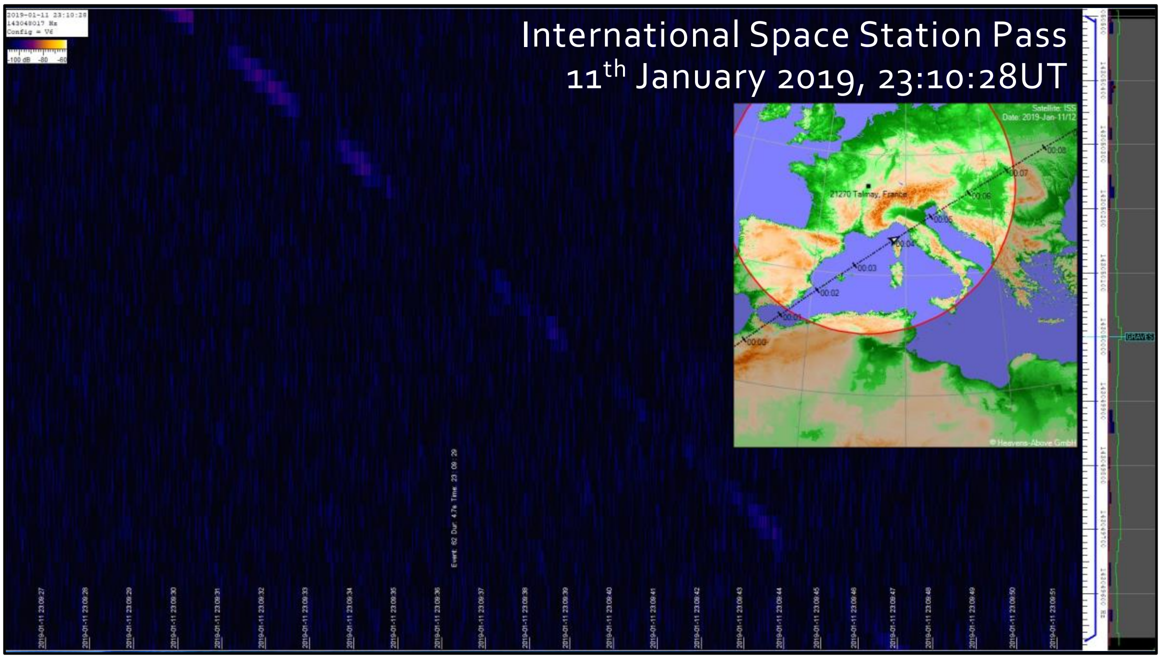

- Here is a second plot from a day earlier in January 2019. You'll note that the track is somewhat nearer to both the transmitter and receiver sites.

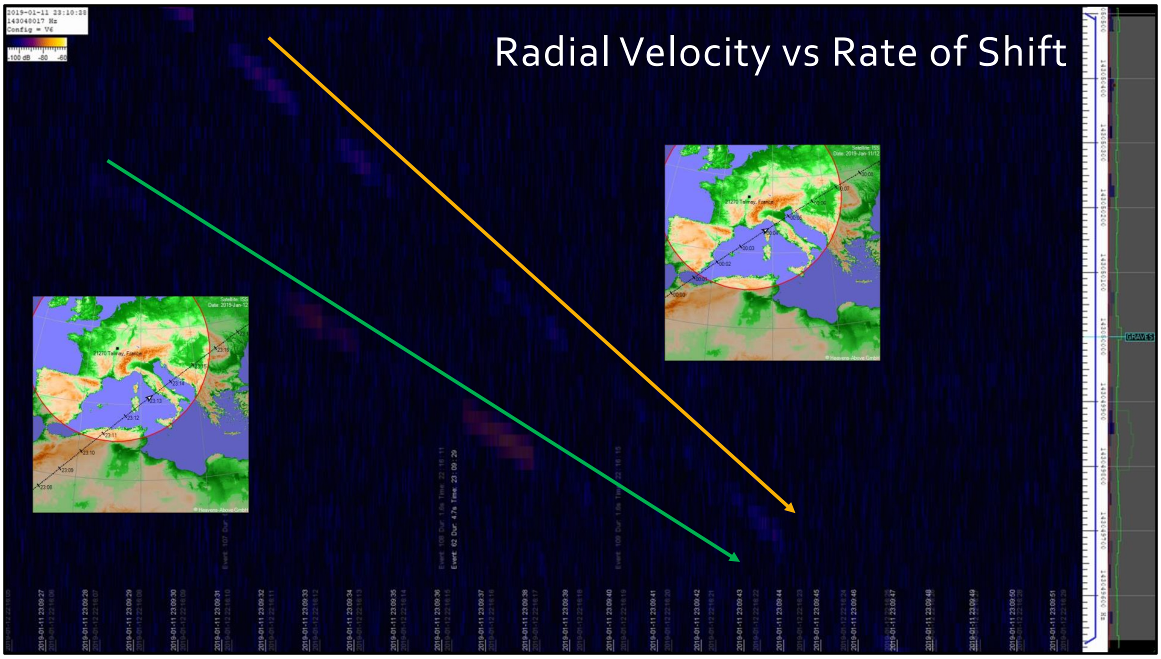

- Overlaying the two plots, you can see that the angle of the radar plot for the nearer pass (in orange) is steeper than the angle for the pass that was further away (in green).
- This is exactly what we'd expect to see. If the object passes across our line of sight nearer to us and the transmitter, then its apparent distance will change more rapidly.
- Thus the radial velocity and rate of frequency change will be correspondingly faster.
- This is fine if we know the intrinsic speed of an object such as the ISS. By looking at the angles of the plots we can at least estimate which passes are nearer or further away, and indeed we can have a stab at calculating the distance to the point of specular reflection.
- There are papers explaining the calculations for back scatter radars, but it is worth a try for forward scatter if we use the simplifying assumption that differences between the outgoing and reflected signal paths approximately cancel out.
- It is much harder to do similar calculations for meteors using forward scatter detection. We know neither the intrinsic velocity nor the distance to the specular reflection point.
- If we use meteors from showers with known average velocities (determined by

back-scatter detection), we can have a go at estimating the distances, but it is nowhere near as reliable.

- I leave it to you to look in to this further if you're so inclined.

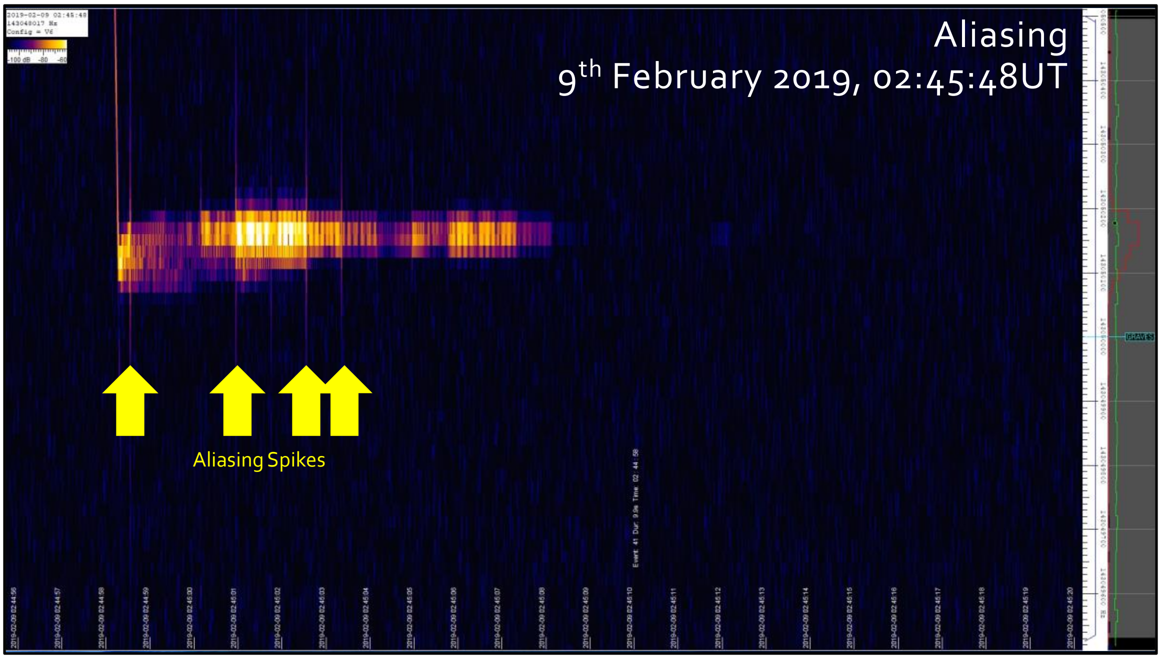

- Having taken a look at what we can determine from head echoes, next we'll examine the static tail reflections.
- The first thing to be aware of is aliasing artefacts.
- In this plot from February, you can see the head echo of the meteor at the left, but along the static trail you can see further 'spikes'.
- These spikes are not real signals, but a product of the sampling and FFT processes we are using to analyse the radar signal.
- Aliasing spikes generally protrude equidistantly above and below the static trail, and are vertical, whereas a head echo will generally be more prominent on one side of the static trail and at a slight angle due to the radial velocity change over time.

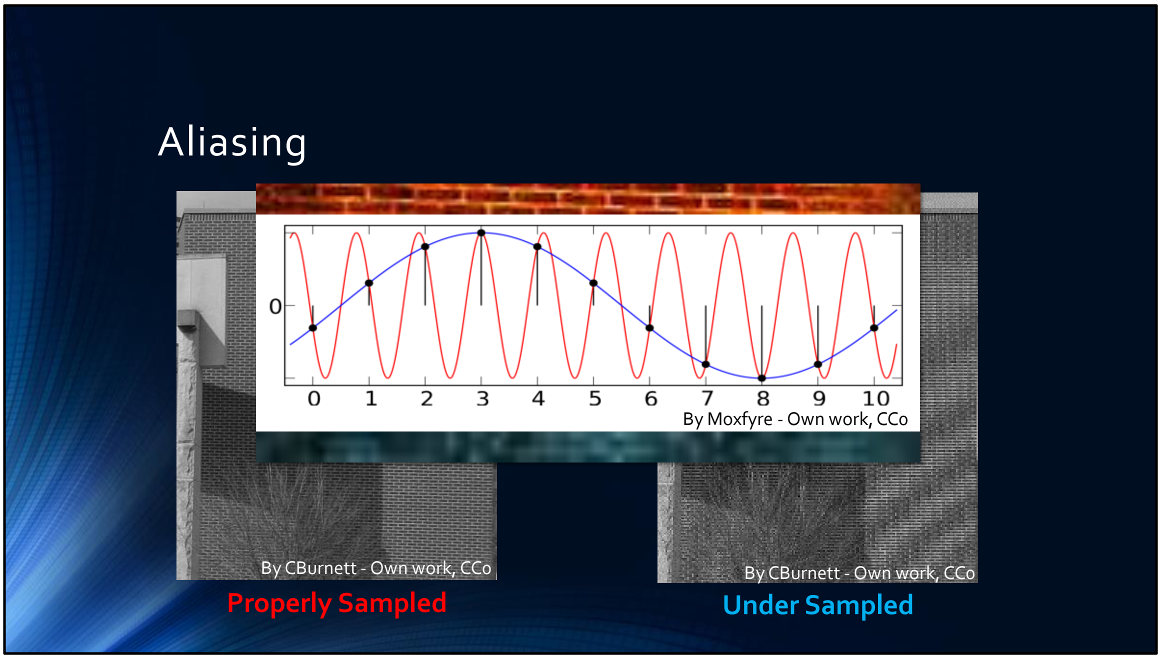

- Aliasing is cause by under-sampling. It affects all sorts of signals that we process digitally, including sound, radio waves, and in this example a digital image.
- Nyquist-Shannon Sampling Theorem dictates that in order to convert an analogue waveform in to an accurate digital representation, we have to sample the waveform at twice the rate of the highest frequency component that we wish to represent.

- Starting simply, if we wanted to sample a simple sine wave with a frequency of 2kHz, we would need to take samples at a rate of 4kHz, i.e. we need to measure the amplitude of the waveform at least 4,000 times per second. If we do so, and then use those measurements to drive a speaker in our TV or radio, then we'll get back that original 2kHz sine wave as a continuous tone.
- Now bear in mind that we already discussed that any waveform can be broken down in to a series of simple sine waves using the FFT process, so for any complex waveform we simply need to identify the highest frequency component it contains using the FFT, and then sample at twice that frequency to capture and accurately recreate the original waveform.

- What happens if we sample at less than the Nyquist rate though? Well any signal

(sound, image, radio, etc.) is a form of energy. Your physics teacher will have told you that energy can neither be created nor destroyed, only change its form.
- Even if we turn an actual signal (say a sound) from energy to numbers, the same thing is true. The energy that we don't fully sample has to go somewhere, and it shows up in the numbers.

- So where does it go? If you look at this digital photo of a brick wall on the left, it has been properly sampled, i.e. for every little feature such as the lines of mortar between the bricks, there are at least two pixels (samples). On the right is the same picture, but deliberately under sampled, i.e. some finer features are only represented by one pixel.
- The energy (photons of light hitting the camera sensor), still has to go somewhere in the digital world, and what you see is a whole series of waves and patterns (aliasing artefacts). These waves are the under sampled energy appearing in the picture.

- So what's going on? Well this graph shows one of the higher frequency waveforms in the picture in red. The red waveform represents part of the changes in brightness across a strip of the image caused by (say) the mortar lines between the brickwork.
- If we under-sample this waveform (i.e. measure is at less than twice per wavelength – the black dots), and then try to re-display the waveform on our screen, we find that the blue wave appears instead – i.e. a much coarser amount of detail.

- So when we have aliasing due to under-sampling (in space as here or over time for a sound or radio wave), the under-sampled high frequency components 'fold back' and appear as lower frequency components.

- That is part of the reason for the aliasing spikes in our meteor trail – the under-sampled frequency components have come out somewhere in the plot. Also, because each FFT calculation (single strip of pixels) in the plot uses a block of samples several seconds long (the window in FFT speak), the aliased energy concentrates in certain points at the start or end of a particular window creating these spikes.

- As a quick aside, why don't our digital smartphone and DSLR images suffer from aliasing effects all the time? Well the camera manufacturer puts an anti-aliasing filter in front of the sensor. This is a very slightly frosted layer of glass that blurs the frequencies that are too high to be properly sampled by the camera's sensor pixels. Instead of the excess energy being aliased in to ugly waveforms, it is physically spread uniformly over the entire image as a very low frequency set of waves before it is sampled. This reduces the contrast in the image a bit, but is far better than aliasing. Similar filters are used in audio sampling to block frequencies above the chosen Nyquist rate from entering the sampling system at all.

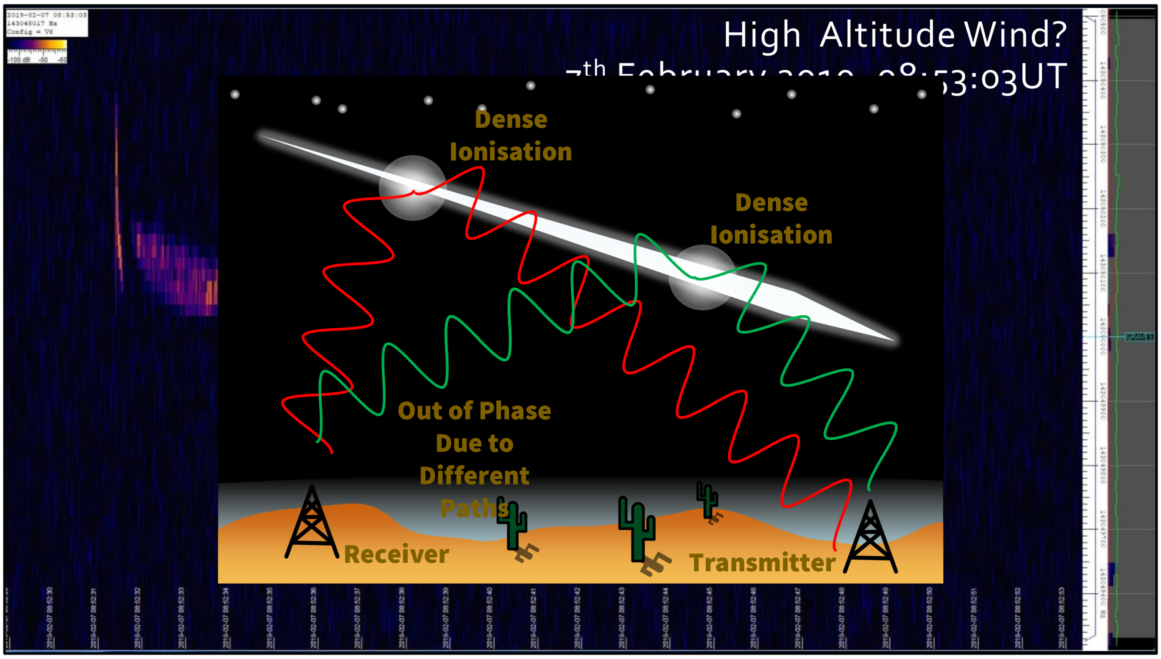

- Here's an interesting plot of a static tail reflection. We we'd expect the reflection to be centred on the transmitter frequency if it is not moving.
- In this case it is some way above it, but that's just an artefact of the cheap SDR dongle – the frequency can drift with temperature and tuning doesn't appear to be particularly linear.
- Nonetheless you can see the reflection varies in frequency over the course of a few seconds.
- In many cases this is due to the ionised meteor trail physically moving – blown by high altitude winds it can shift, distort and break in to sections before it decays, creating a small and varying Doppler shift.
- The other two things to consider with a trail are the radar scanning pattern and interference:
- Firstly you've seen that the GRAVES radar scans each subsector once every 3.2 seconds. The beam is not likely to be that tight, so depending on the geometry and nature of the ionised trail, you may get a continuous return as here, or the static tail make be chopped up in to chunks with gaps between them as the beam scans on and off the trail.
- Secondly different parts of the trail mail be more or less reflective depending on the density of the ionisation. This creates 'Fresnel Zones' where the return signals

from two or more parts of the trail constructively or destructively interfere with each other, making the total signal at the receiver vary in strength as the trail decays at different rates along its length and/or breaks up due to high level winds.

- You can see the cause here – strongly reflecting parts of the trail cause different radar returns along longer or shorter paths, which causes the signals to arrive out of phase at the receiver and thus constructively or destructively interfere.
- As far as I know it is impossible to disentangle the varying signal strengths due to the scanning beam and due to interference, plus the windowing/aliasing effects of the FFT process will tend to amplify stronger and weaker returns over time. Upshot is that the trail will often vary in strength over time rather than just fading uniformly as the ionised particles recombine.

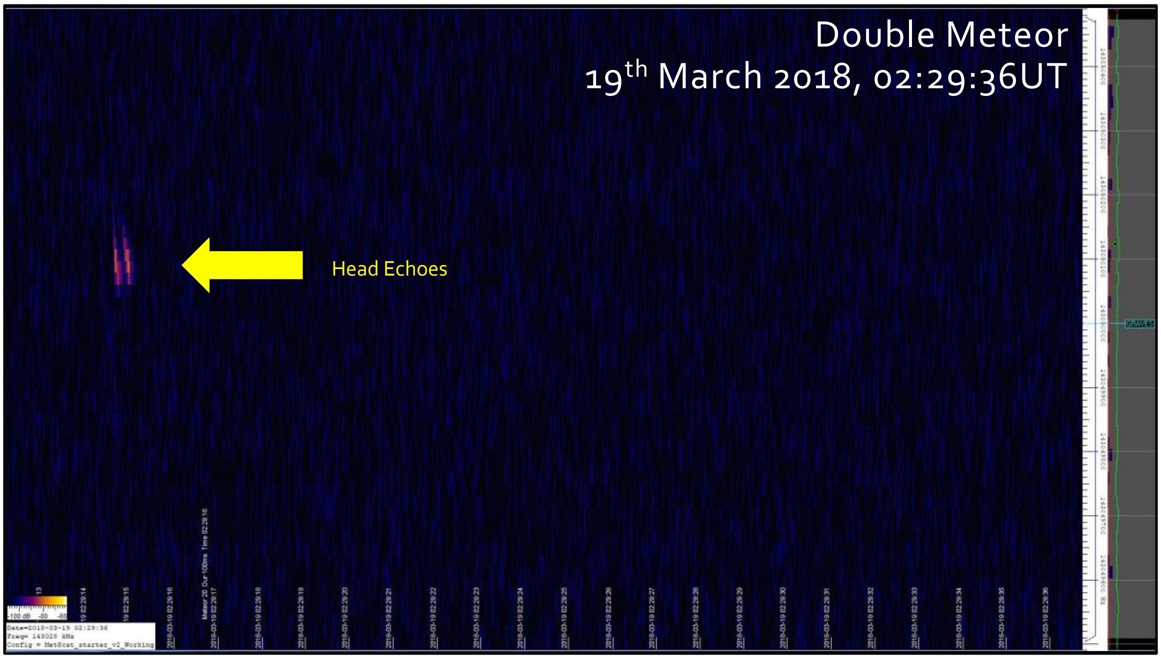

- I have frequently captured ‘double meteors’ as shown here.
- I believe these are genuinely two meteoroids travelling along roughly the same path in quick succession; indeed on a couple of occasions I have visually observed this ‘double meteor’ effect where one meteor is quickly followed by a second along the same path.
- It’s not that surprising given that meteoroids generally form from small particles detached from large comets, so it would be reasonable to expect them to end up in ‘trains’ on the same orbital paths.

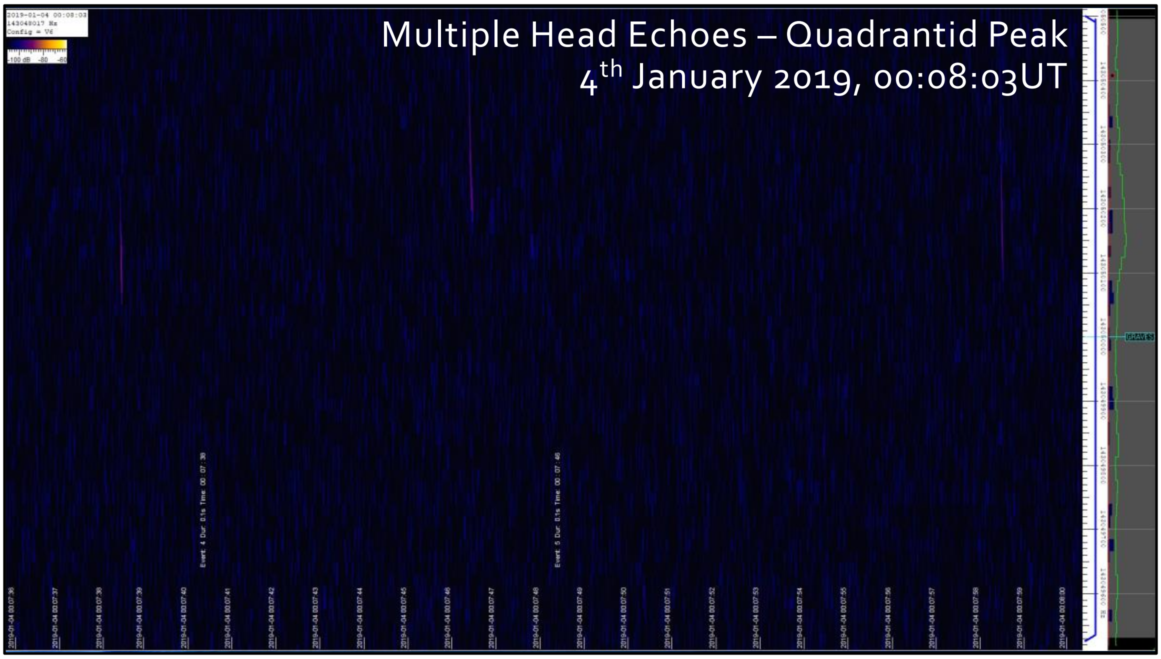

- During meteor showers, multiple meteors may be detected per minute depending on the hourly rate.
- This is the peak of the Quadrantid shower on 4th January 2019, with a predicted rate of c. 100 per hour.

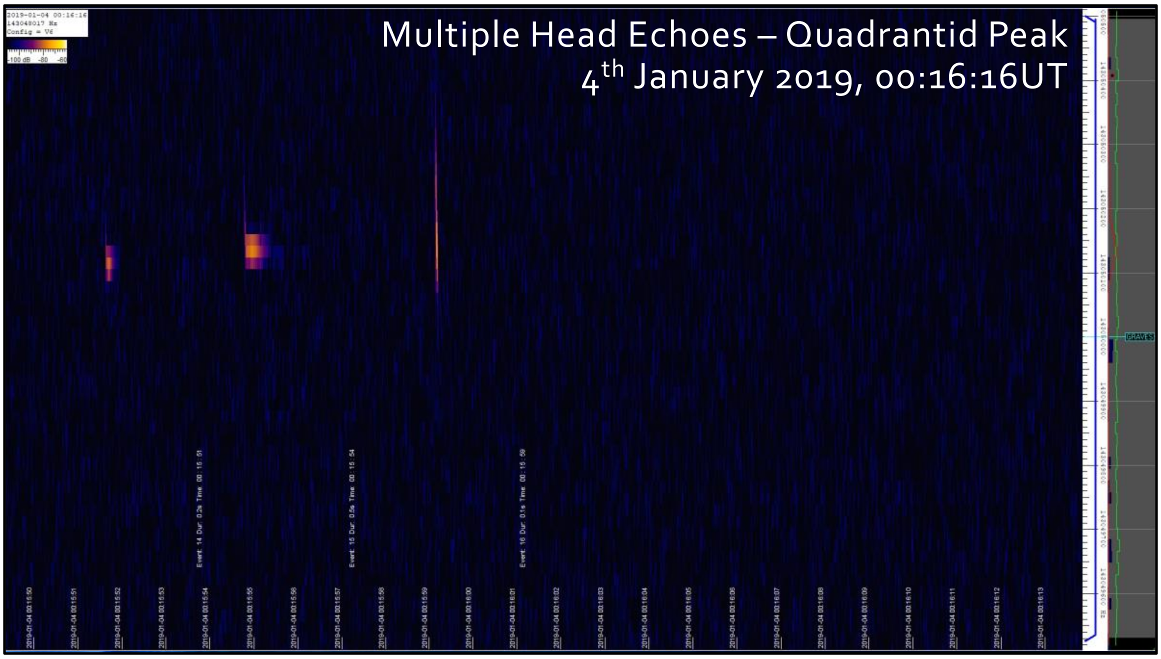

Here are more from the same shower about 8 minutes later.

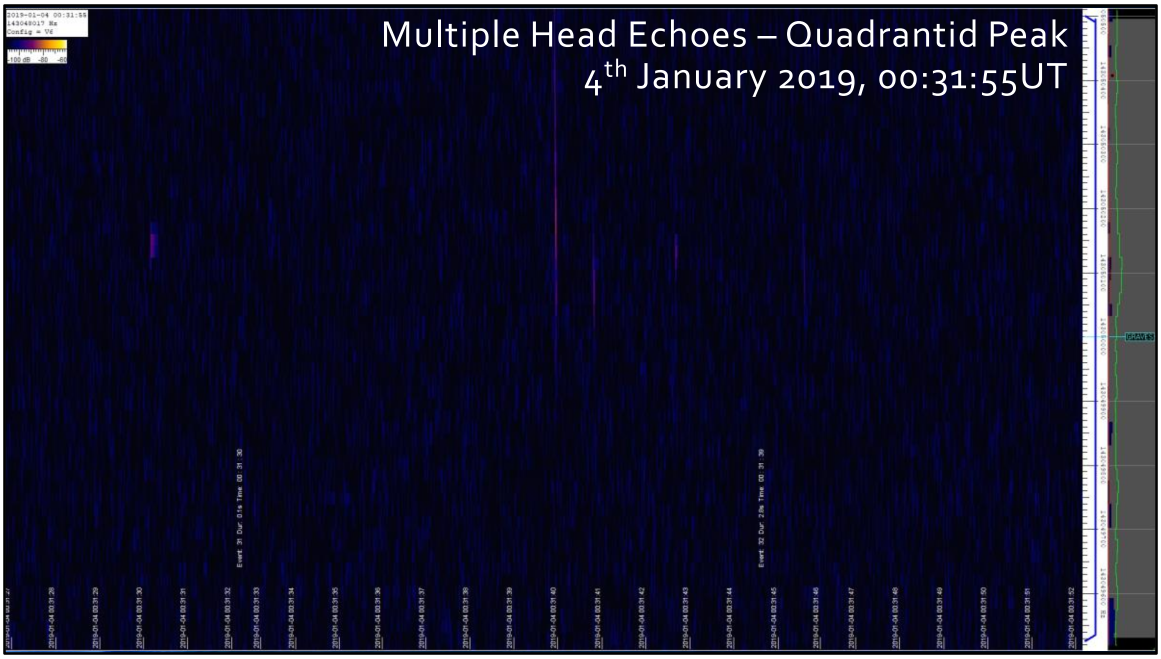

- And even more, all within 30 minutes on the peak evening.

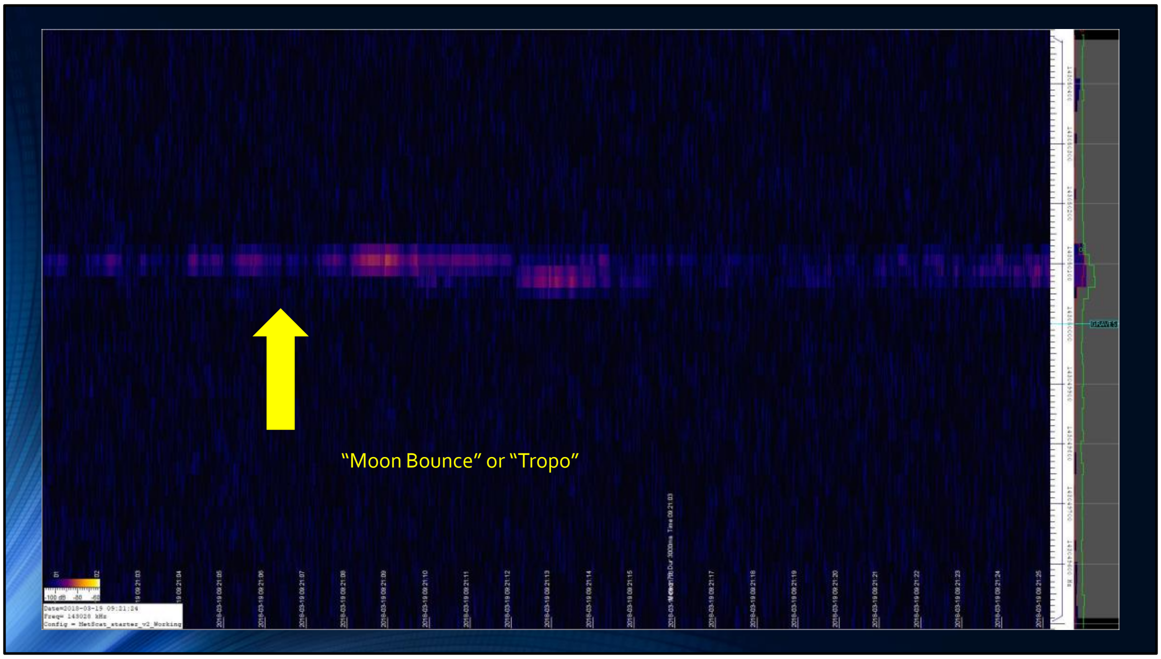

- One common phenomenon is to have a period of continuous radar returns over several minutes or longer. This can be caused by one of two effects:
- Firstly 'Moon Bounce' where the Moon passes through the specular reflection point between transmitter and receiver. These signals are generally quite weak and I don't have any clear evidence of having received one myself.
- Other observers with more advanced setups definitely do get them, e.g. using filters and amplifiers at the antenna to increase SNR. The geometry can be calculated and it is possible to predict and filter out false returns due to Moon Bounce with more sophisticated software. I've not done this myself though.
- More commonly one will receive the radar signal directly due to tropospheric ducting, a 'tropo'.
- During periods of stable high pressure weather, a temperature inversion may occur. Normally the temperature of the troposphere (lower part of the atmosphere) falls by between 6 and 10 degrees Celsius for each 1,000 meters gained in altitude. During an inversion, above a certain height the temperature actually increases with height for some distance.
- This increases the refractive index of the air and allows radio waves to be bent through the atmosphere by refraction and travel much further than normal.
- Thus during a 'Tropo' we are able to detect the outgoing radar transmission which

is normally hidden behind the curve of the Earth.

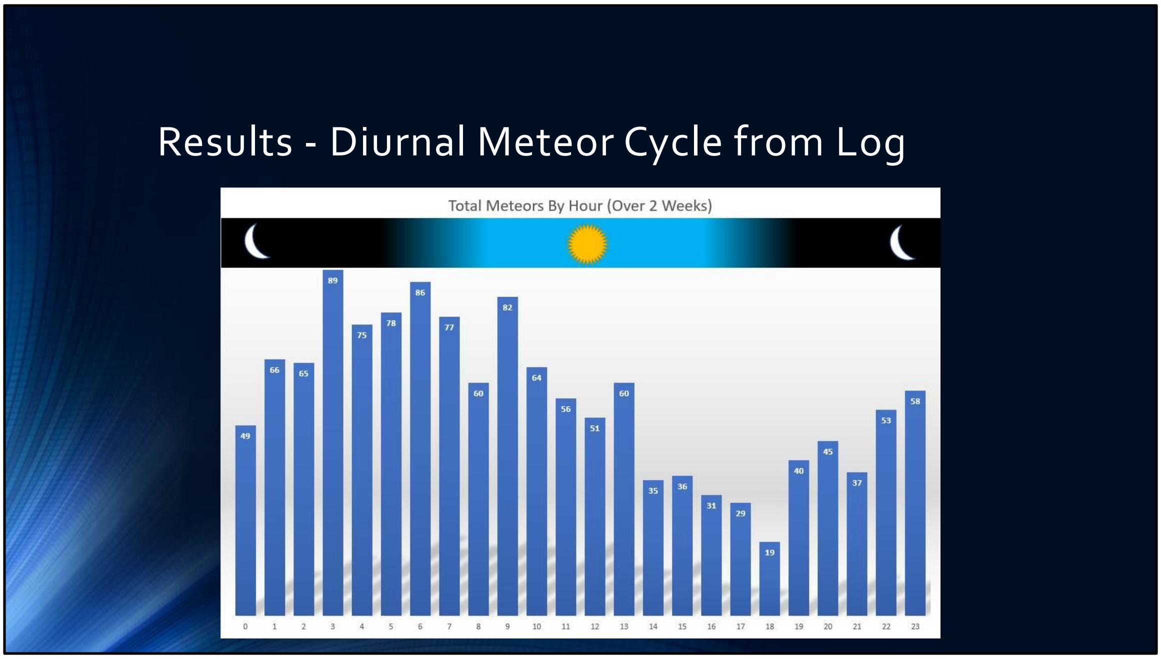

- One fun and simple experiment is to measure the average rate of sporadic meteors by hour of day (i.e. outside a meteor shower period.
- Here you can see I've just counted the number of meteors arriving in each hour of the day and average the result over a period of a couple of weeks.
- A clear pattern emerges quickly – there are fewer meteors during the afternoon and early evening, with the numbers picking up during the early hours of the morning.

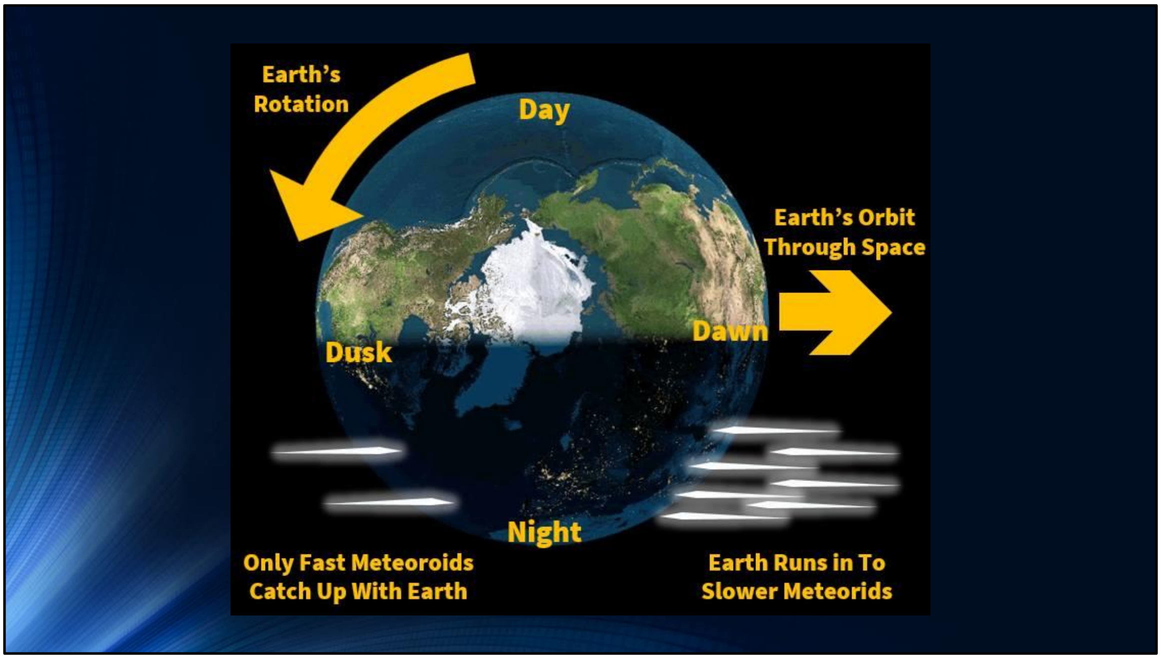

- This diurnal meteor cycle is easy to explain.
- The Earth travels through space on its orbit at about 67,000 miles per hour.
- If you imagine driving your car during a summer evening, lots of bugs will splat in to your windscreen because that's the direction you are going, whereas they don't splat the rear windscreen.
- During the afternoon and early evening, we are effectively looking out of the Earth's 'back window', so only faster meteoroids travelling in the right direction will enter the atmosphere.
- During the early hours and the morning, we're looking out of the front window, so far more meteoroids are swept up as the Earth runs in to them, just like the bugs hitting the car.

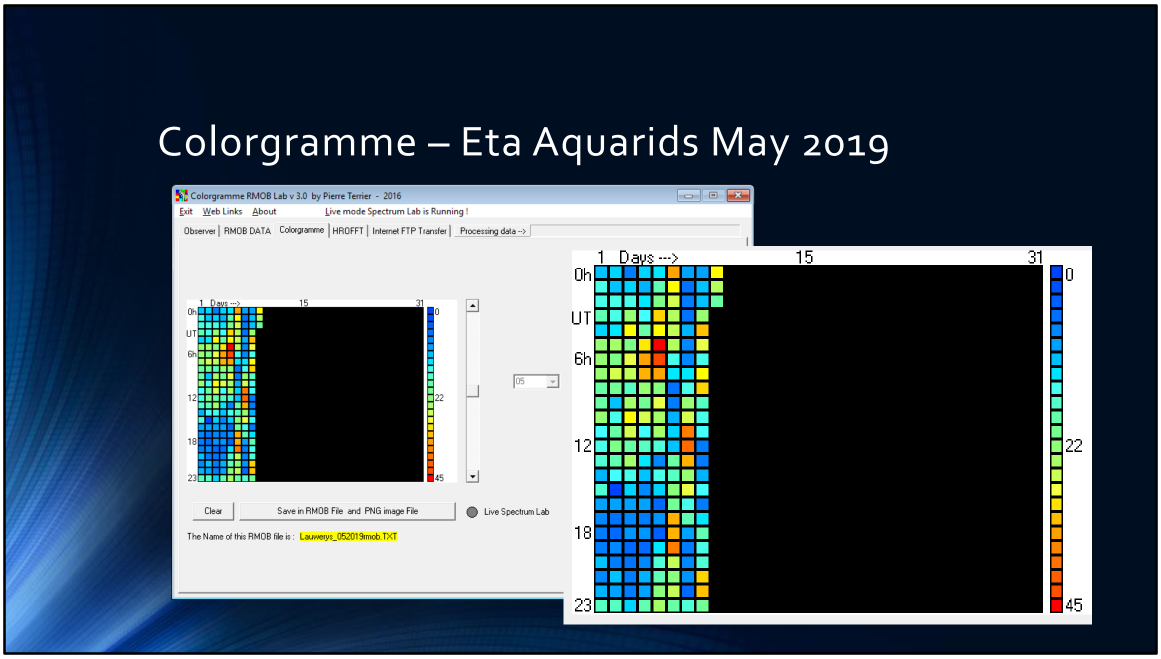

- We can use the same simple counting techniques to identify meteor shower peaks as well.
- This is the same sort of data analysed from the log files generated by Spectrum lab, and plotted using a free application called 'Colorgramme'.
- What you can see is the number of meteors per hour (midnight at the top through the day to 11:00pm at the bottom), and day of the month, one day per column from left to right.
- The colour scale represents the number of meteors logged each hour.
- You can see that there is a very bright area on the 5th of May between about 3:00am and 7:00am. This is the peak of the 2019 Eta Aquarids shower. You can see the peak building up in the previous two days around the same time.

Collaborative Data Sharing - RMOB

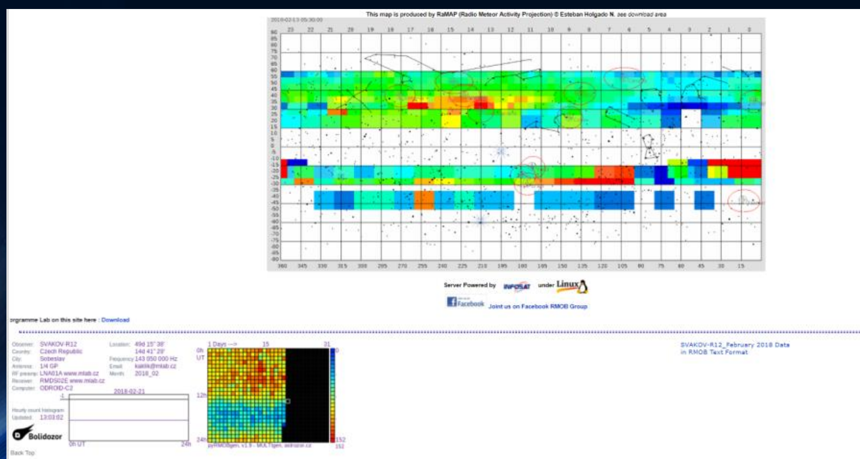

- Colorgramme can be used to upload your data to the Radio Meteor Observers Bulletin web site (RMOB).
- Lots of users around the world upload data regularly and the site has some features that try to correlate data from all observers, correcting for location on the Earth to try to identify meteor shower peaks and approximate locations of their radiants on the sky.
- I'm not sure how successful this is, but the idea is to try and pick up much lower rate / unknown showers by collecting data over many years.
- There are other collaborative meteor detecting sites and organisations out there and this is a fascinating subject to delve in to, but one that is relatively easy to get started with.

Improved Spectrum Lab Actions

- Finally, if you do decide to give this a try you may want to take a look at some improved conditional action scripts that I've made available for Spectrum Lab.
- I won't go in to the details, but suffice it to say that the starter actions from the BAA web site have a number of problems for accurate data collection and logging.
- The script I've made available addresses these and you'll probably want to give something like this a try if you get serious about meteor detecting.

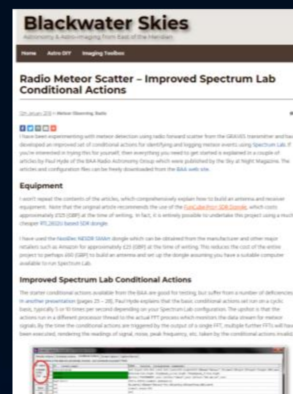

Questions?