October 2023 Eclipse: Effects on HF Propagation. As seen using WSJT-X mode FST4W

A talk over Zoom to the Norfolk Amateur Radio Club, 24 January 2024

Gwyn Griffiths G3ZIL

Slide 1

Good evening, it is my great pleasure and honour to give this talk on some of the effects of the 14th October 2023 eclipse over North America. Along the way I'll introduce you to FST4W, a WSPR-like beacon mode within the well-known WSJT-X software suite.

Slide 2

In a 2015 RadCom article NARC member Steve Nichols' reported on several effects on propagation of the April 2015 eclipse over the UK. He highlighted the use of new technologies over the sixteen years since 1999. I'll show that the pace of innovation in technologies and methods for amateurs to study propagation at HF has not slackened.

Slide 3

Here's an outline of my talk. By returning to my own early days with WSPR, I'll be retracing my steps. In that way I hope to carry you with me into perhaps unfamiliar territory - the FST4W mode for one, and graphical ways at looking at measurements and effects on propagation for another.

I'll end with an idea for how UK amateurs might contribute to studies of the 8 April 2024 total eclipse over North America.

Slide 4

Born and bred on the island of Anglesey my initial fascination with radio was piqued by "Charlie George" - transmissions from Holyhead Coastguard at the very top edge of the medium wave on a 1950s Pye set. Most fittingly, one of those Coastguards, the late Paul Lane GW3MQX, helped me become GW3ZIL. But ... perhaps because of living so close to the sea, my curiosity extended to using wireless underwater. I'm the whole person on the left. My school friend Robert Ceen is in the wetsuit with electrodes attached and a transmitter with a throat microphone. He would be underwater and I'd be on the surface with the receiver. This is before the use of the STEM acronym for Science Technology Engineering and Mathematics but that's what we were doing, guided by an incredibly tolerant and supportive physics teacher.

My interest in equipment of that era remains.

Slide 5

It was the 'underwater' interest that edged ahead and became a fulfilling career over forty years. Early years were as an electronics engineer supporting equipment at sea and developing new instruments. I could describe this sonar as a fixed frequency 1 MHz direct conversion receiver with quadrature outputs to an Analogue to Digital Converter with digital signal processing to extract signal level and Doppler shift in two-metre long range intervals away from the instrument. It's a 1987 SDR. It helped that sound underwater only travels at fifteen hundred metres per second.

Later, now a group head, our major effort was designing and building autonomous underwater vehicles. They can go were ships find it difficult, such as under ice here off Greenland, or indeed impossible, such as under the five hundred meter thick ice of floating Antarctic glaciers and ice shelves.

Slide 6

On retirement I returned to amateur radio. Of the myriad of opportunities in the world of amateur radio I discovered WSPR - the Weak Signal Propagation Reporter mode within the ubiquitous WSJT-X package. I think it's because I could use my hands and my head to build simple but very effective receivers and transmitters. I enjoyed the mix of building hardware and writing software, and of incorporating bought-in modules from the Arduino and QRP worlds.

Not all were successful - in the Arduino-based WSPR transceiver I did not pay enough attention to self-generated noise. Having tackled that self-inflicted problem I was left with noise arriving via the antenna.

Slide 7

The problem of noise arriving via the antenna led me to explore noise cancellation using a pair of antennas, altering the phase of one and matching their amplitudes to achieve cancellation of a single dominant radiated noise source. It's a well-known technique. One I set out to automate. This is really a topic in its own right, but the basics are these:

- My prototype was intended for use with WSPR on one band, 40 metres in this case.
- Two dipoles, one with a switched attenuator in one dB steps to7 dB and the other with a switched delay to give the phase shift. The delays were from switching-in lengths of RG174 cable to give roughly 3 nanosecond steps up to 80 nanoseconds, more than enough for 0 to 180 degrees phase shift on 40 metres.
- The two signals are combined, input to the receiver, whose audio output is processed by WSPR in WSJT-X as usual.
- Within the eight second gap at the end of a WSPR transmission the fully automated software goes through attenuator and delay settings to find minimum noise. It applies those settings for the next cycle, and logs them so I could see the benefit.
- This is one of the most impressive noise reduction examples, about 18 dB improvement in signal to noise ratio.

Slide 8

Well, it started in a coffee shop in Sonoma County, Northern California, across the Golden Gate Bridge from San Francisco. Not that I have ever been there... My interest in noise measurement and reduction had led to enjoyable and productive email exchanges with Glenn Elmore, N6GN. He, and many of these good folk, met each week at Coffee Catz. They welcomed my participation over the Internet as if we had known each other for decades. Quite a touching example of the true international spirit of Amateur Radio. The group's interests include installing and maintaining multiple KiwiSDRs at the restored maritime radio receiving station KPH to a substantial software development by Rob Robinett - WsprDaemon. WsprDaemon is a robust decoding and reporting system for WSPR for KiwiSDRs and other multichannel receivers. I contributed to its noise measurement features help look after the database, which accepts many more variables than does wsprnet.org.

Slide 9

Arising from our work with WsprDaemon Rob, Glenn and I started to contribute talks and posters to the annual meetings of HamSci - the Ham Radio Citizen Science Investigation. HamSci is led by an ionospheric physicist, Assistant Professor Nathaniel Frissell of the University of Scranton, Pennsylvania W2NAF. It is a platform to promote projects that bring the amateur community into contact with professional scientists and engineers studying the ionosphere.

Of their many activities it is the HamSci Festival of Eclipse Ionospheric Science for 2023 and 2024 that is central to this talk. We looked at using WSPR, given our deep interest, and how it gave interesting results for the 2017 eclipse over the US and the 2015 eclipse over the UK.

Slide 10

We wanted to move beyond what had been done before. We'd come to realise that SNR is not the only factor that affects the probability of decoding a WSPR signal. Frequency spread has a big effect on this very narrowband mode. Here's a graph from Steve Franke, K9AN, the "F" in FT8 alongside the "T" for Joe Taylor. This graph deserves to be better known. On the X-axis we have signal to noise ratio expressed in a two and a half kilohertz bandwidth. WSPR users may well see the very occasional -32 dB SNR, but low -20s are more common. This graph shows why: as frequency spread increases the SNR you need increases.

Steve Franke labelled this graph "Doppler spread" - but I'll use the wider term "Frequency spread" as it is not only the ionosphere that is spreading the frequency through jittery Doppler shifts it comes from jitter, or phase noise, in the transmitter and receiver.

So how could we get information about frequency spread as well as SNR for our eclipse studies?

Slide 11

The answer was "quite easily" but it was pretty well hidden at the same time. Introduced in the 2.3.0 release of WSJT-X in 2021 FST4W, a mode very much like WSPR, can measure frequency spread. Turning on that option is not a simple click: you have to put an empty file named plotspec in the directory in which wsjt-x is started.

We made detailed studies of equipment frequency jitter, you'll find reports on the wsprdaemon website. More importantly, John Seamons at KiwiSDR and Hans Summers at QRP Labs made significant improvements in their products arising from our findings.

Identifying frequency spread of different propagation modes is the subject of my RSGB Tonight at Eight talk on 5th February. We can clearly identify different propagation modes from SNR and frequency spread. We thought this could be useful during an eclipse. And we weren't wrong.

But there are several challenges ...

Slide 12

So, what equipment could we use? High-end software defined transceivers with external phase locked GPS disciplined oscillators meet the requirements. But so too can lower cost equipment, including kits. The QDX digital modes transceiver from QRP Labs is usable, and even better if clocked by a relatively low cost GPSDO such as the mini-Bodnar. The KiwiSDR with its out-of-the-box GPS aided master oscillator is usable, but again, adding an external GPSDO gives excellent results. The £60 RFZero transmitter module comes with a GPS aided oscillator that is better than that in the KiwiSDR although not quite as good as a phase-locked GPSDO. Here I've added a relay-switched low pass filter board from QRP Labs and a QRP Labs 10 Watt linear amplifier to boost the 20-milliwatt output from the RFZero.

Slide 13

Paul Elliott WB6CXC designed the WSPRSONDE specifically for the eclipse. It can transmit simultaneously on six bands. At 100% duty cycle, and with phase-locked GPS disciplined oscillator stability. During normal conditions this would frowned upon - but as changes in the ionosphere happen quickly during an eclipse a transmission every ten minutes with a 50% chance of being decoded is not going to describe the response of the ionosphere. Five WSPRSONDES were deployed for the October 2023 eclipse and have produced excellent data, as we will see.

Slide 14

This map shows the transmitters in red and the receivers in yellow that were using FST4W during the October 2023 eclipse. I've shown the WSPRSONDE locations with WS. The path of the eclipse is shown in cyan, with the magenta lines bounding the area of the annular eclipse, with up to 95% of the sun obscured. But as you see, the area with the sun 80% obscured was much larger.

I'll be showing results for a selection of transmitter to receiver paths, along and across the path of the eclipse, and at different frequencies, to bring out the variety of effects that we've been able to observe and measure.

Slide 15

We were not able to cover the medium wave or 160 metres, where many useful reports of increased signal levels were reported for the 2015 eclipse over the UK in Steve G0KYA's RadCom article. But we were able to see and quantify the increased signal levels from reduced D layer absorption on 80 metres. Here's just one example, on a 466 km path from northern Utah to Nevada across the path of the eclipse. I've taken the SNR of the FST4W spots and WsprDaemon's noise level measurements from Tom Bunch WO7I to calculate signal level, on the Y-axis. I've marked the start, middle and end of the eclipse at the mid point of this path with the blue lines. Compared with the same time on non-eclipse days we saw a 12 to 15 dB rise in signal level. For the first forty minutes of the sun becoming more and more obscured there was no change in signal level, but then there was a sudden rise. It's these sorts of detail I'm hoping to discuss with ionospheric physicists at the upcoming HamSci workshop in Cleveland in March.

Slide 16

For my next example I'll jump to 10 metres to look at the effect of the eclipse on propagation via the F2 layer. In particular, how reduced ionisation from the sun's shadow dropped the critical frequency, altering the skip zone distance. Here John Clark TI4JWC is operating a WSPRSONDE transmitter. His Costa Rica location is right at the edge of the annular path.

I'm looking at spots received from 4300 to 5000 km distant, by Dennis ND7M, the Northern Utah SDR site KA7OEI-1, either side of the eclipse path, and by KiwiSDRs at KFS and KPH northern California. I'm only looking at whether spots were received or not. With time of day on the X-axis the upper plot shows circuit reliability - that is, what percentage of the ten spots transmitted in 20 minute intervals were received at KFS and KPH. It does not drop to zero, but there was a definite dip.

KFS and KPH also appear on the bottom graph. Here the Y-axis is distance. What I see is first is propagation holding up better to KPH and KFS than to the stations around 4400 km. Second, I see that there were two periods of loss of spots to ND7M and KA7OEI, with a brief recovery in between. Did the eclipse affect one hop then the other? Let's look at that possibility.

Slide 17

Ray tracing is a useful propagation modelling technique that can help our thinking about what was happening during the eclipse on those 4300 to 5000 km paths from TI4JWC. Here I am using the PyLap ray-tracing package, available from HamSci via this Github link. I've specified the date, time and frequency, the latitude and longitude of the transmitter, and a bearing for the path I'm interested in. Those are simple facts. More difficult, yet crucial, is the sunspot number: here it is R12 the value for the chosen day from a twelve-month running mean - which we do not know yet. Instead, I have taken the so-called Effective Sunspot number calculated by Northwest Research Associates available via this link.

Without an eclipse, for R12 as the effective sunspot number for 14 October of 125 we see that the second hop covers ranges from about 3700 km to 5500 km - so we'd expect reception at all four of the receiver sites. That is what we saw before and after the eclipse.

Slide 18

By trying lower values for R12 I found that a value of 70 was just low enough to push the start of the second hop skip zone to beyond 4500 km, that is beyond the receivers in Utah and Nevada, while still enabling reception at KPH and KFS. In this simple model experiment we can say that it looks as if the passing of the eclipse was equivalent to the sunspot number dropping from 125 to 70 then back up again over a few hours.

So what about the two gaps with a recovery at the Utah and Nevada receivers? The path of the eclipse was from northwest to southeast. So it first affects the second hop, reducing, as we've seen, the effective sunspot number and pushing the second hop beyond 4300 km, producing a gap. The eclipse shadow travels along this path, while the ionosphere where the second hop takes place recovers, but the shadow has yet to affect the ionosphere where the first hop takes place. We receive a few spots. Then the shadow affects the ionosphere at the first hop, pushing the landing spots to greater distances, and hence the landing spots of the second hop out to greater distances and we get another gap.

Slide 19

Let's now look at a path across the eclipse from Northern California to Colorado with simultaneous transmissions on 20 and 15 metres. Time is on the X-axis, and in the top graph I have circuit reliability. Propagation on 15 metres dropped out then recovered. 20 metres stayed open, although with lower SNR as seen on the lower graph, only to collapse after the eclipse. So the critical frequency on this 1566 km path during the eclipse was high enough to keep 20 metres open but not 15.

Slide 20

What do we see in the ray traces for this path? On the left is 15 metres, the top trace with R12 of 125, as we chose previously, for the non-eclipse case. Paul WB6CXC's transmissions are received by Glenn N6GN. Now I drop R12 to 70 - the value I needed to drop to on the 10 metre path from Costa Rica. It's nice to see that for this different path, and different frequency R12 of 70 is low enough to push the one hop ray landing spots out beyond N6GN.

I have to say I am less convinced about why spots on 20 metres after the eclipse were not decoded. I could argue that N6GN was just too distant for one hop and not distant enough for two hop. If I were to drop the sunspot number N6GN would be within one-hop range - so I can suggest with fair certainty that's what was going on during the eclipse on 20. But still not convinced about afterward.

Slide 21

So far, I've not drawn upon the frequency spread measurement from FST4W. So here goes...

The information here is on a 1808 km path across the eclipse from Dick W7WKR Washington State to Dan KV6X in New Mexico on 20 metres. All three graphs have a common X-axis with time. The top graph is circuit reliability in twenty-minute intervals as we have seen before. The middle one is signal level and the bottom graph is frequency spread. Note that frequency spread is in milliHertz.

What we've found in our studies of frequency spread with FST4W is that one hop propagation usually results in less than 100 milliHertz spread. And, from ray tracing, we know that as twenty metres opens at this distance the propagation starts out as one-hop. Hence what we see on non-eclipse days - the colours other than red - and the eclipse day, are low values of spread. I've ringed these as cluster 'A'.

As the band continues to open frequency spread jumps up with very scattered values above 100 milliHertz. These are values from two-hop and one-hop likely co-existing. On a normal day, those would be the prevailing propagation modes at 1700 UTC on this path.

During the passage of the eclipse, where the blue line shows the time of a dip in circuit reliability and a reduction in signal level we see cluster 'B' of low frequency spread values. There is no equivalent cluster at that time on normal days. What I suggest is happening is that the lower level of ionisation has meant that we see only one-hop propagation, with its lower frequency spread, rather than a mix of one and two-hop.

We'd not have been able to make this deduction from circuit reliability or signal level alone.

Slide 22

Here's another example of how a change in frequency spread tells us about a change in propagation mode. We are now looking at a shorter path, 1055 km, still on 20 metres from Dick W7WKR but this time to KPH, across the eclipse. This path is short enough for only one-hop propagation to be seen on normal days. In the top graph, with time again on the X axis, I'm combining signal level, the cyan dots, on the left hand Y axis, with frequency spread, the orange dots, on the right hand axis. As you see, most frequency spread values in orange are below 100 milliHertz on the right hand scale. Thus confirming the propagation mode as one-hop.

Another way of looking at the data is to plot frequency spread on the Y-axis and signal level on the X-axis. And then add contours, as in a map. We see the tight cluster 'A' during one-hop propagation. The dispersed spots 'B' are not just one-hop but weaker, they represent a completely different mode of propagation.

We saw two-hop change to one-hop as the critical frequency dropped - what does one-hop change to? It changes to a mode that **appears** to propagate above the maximum usable frequency. That mode is Two-Hop sidescatter. It really is a subject in itself - one that will feature prominently in my forthcoming RSGB Tonight at 8 talk.

And so I ask you to take it on trust that those spots labelled 'B' in the contour diagram and only appearing during the eclipse propagated as follows. During the eclipse KPH became within the skip zone from W7WKR as the effective sunspot number dropped to 70. But normal one-hop propagation could take the transmission to Utah, to where I have the red patch in the map. Most of the energy will have been reflected in the forward direction, to the southeast, but some, a small fraction, will have been scattered to the side, and will have ended up via normal one-hop propagation at KPH. Hence the name: Two-hop sidescatter, with a much weaker signal and much more frequency spread due to scatter from multiple points, than one-hop. It's that combination of weak signal and high spread, measured using FST4W, that enables us to identify this underappreciated propagation mode.

Slide 23

I'm going to change tack away from propagation modes during the eclipse to an attempt to measure how the eclipse affected the height of refraction within the ionosphere on a one-hop path. We can do this because of the ready availability of relatively low cost GPS disciplined and aided master oscillators. Here's another example of how we have moved forward.

There is a very sparse global network of ionosondes: we are lucky to have two in the UK, and Chilton and Fairford. They are scientific instruments that measure several parameters of the ionosphere, including the height of refraction within the different layers. I'm grateful to the GIRO data Centre and the team running the ionosonde at Point Arguello, Van den Berg Air force Base California for these two ionograms. On the Y-axis we have height in kilometers, and frequency in Megahertz on the X-axis. They are one hour apart on the day of the eclipse. Simply as a cartoon I have sketched a one-hop path on both graphs. The apex of my path is the peak of the thin black line. That thin black line represents the amount of ionisation. It's the height at which our signal may well refract, about 334 km at 13:00 and 284 km at 14:00. And subsequently until about local noon, it is as if the F2 layer continues to descend.

Keep in your mind for the next slide that my sketch shows that as the ionisation peak has descended the total path length has decreased - the path is shorter at 14:00 than at 13:00.

Slide 24

When a path length shortens, be it in this radio signal or a sound signal from an approaching siren, there is a positive Doppler frequency shift. This mean Doppler frequency shift is a separate measurement to the Doppler spread we've looked at earlier. The one Hertz frequency resolution of WSPR and FST4W spots on wsprnet.org, as shown here, is not at all adequate for quantifying Doppler shift at HF. However, Rob Robinett's WsprDaemon database holds the data with 0.1 Hz resolution, and this proves just good enough for our purpose. With time on the X-axis we have Doppler shift in Hertz on the Y-axis. These are results on three bands on a 545 km path from Tom WO7I to Dennis ND7M.

80 metres, in brown, was open all night, and so we capture the slow, then quickening descent of the height of refraction as an increasing Doppler shift, and then the descent slows down, the Doppler shift decreases.

40 metres, in orange, opens up, and we capture its Doppler shift for some of the descent. A little later 30 m, in cyan, opens up, while the height of refraction is still descending.

But here's an oddity, at about 14:30 UTC Doppler shift on 80 metres shows zero, quite different to the variations on the other bands. We then lose 80 metres from absorption in the D layer. 40 and 30 metres continue to give us Doppler data through the time of the eclipse.

The Doppler data do not lie on top of each other - that would worry me - as from this equation, for a fixed rate of change of path length Doppler shift is proportional to frequency. Hence a greater shift on 40 metres than on 80 metres and an even larger shift on 30 metres.

Slide 25

Next, we calculate the rate of change of path length from the Doppler shift. This removes the frequency dependence. Now the data do sit pretty well on top of each other, except for 80 metres from 14:30. My reading of this is that refraction of the 80 metre signals had swapped to the E layer from the F2 layer as the critical frequency of the E layer rose high enough, and the E layer was not changing in height.

Slide 26

The next step, getting from path velocity to the height of refraction is the trickiest. It requires a little geometry as set out here. We are missing one vital measurement - the path length itself. Doppler shift tells us the rate of change of path length but tells us nothing about the actual path length. So what I have done is to calculate an estimate of the path length from WO7I to ND7M from the height of the F2 layer from the Point Arguello ionosonde once. At 15:00 UTC the height *h* was 230 kilometres from which I calculated the path length as 721 kilometres. At that time the Doppler shift on 40 metres was 0.3 Hz, giving a path velocity of minus 6.7 metres per second. We have measurements every 120 seconds, so in the time between measurements we calculate that the change in path length as 120 seconds times minus 6.7 metres per second which is minus 0.8 kilometres. We now have a path length at 15:02 UTC as 721 kilometers minus 0.8 kilometers which is 720.2 kilometers. We put that path length in the bottom equation and arrive at our height estimate for 15:02 UTC of 229.3 kilometers. We now have a new starting value, and apply the same maths for the next two-minute interval.

Slide 27

Here is the result. Time is on the X axis, and height of refraction in kilometers on the y-axis. The vertical blue line at 15:00 UTC is where I took the height from the ionosonde. The ionosonde estimates several heights; here I've shown two of them, the F2 minimum virtual height in black and the peak height in grey. Our heights for 80, 40 and 30 metres are in brown, orange and cyan. Our measurements track each other pretty well. Comparing with the ionosonde, it looks as if our measurements track the peak height in the morning, local time, and swap to tracking the minimum height in the afternoon.

As to whether this makes sense - I need to talk with an ionospheric physicist.

How has the eclipse affected the height of refraction? Not easy to tell from this plot, so let us zoom in, and compare our estimated height with that on the following day.

Slide 28

To the graph of the previous slide I have added refraction height on the 15th October as the average for the three bands, this is the magenta line with dots. Immediately we see that on a 'normal' day our method shows a smooth change in height, quite different to the eclipse day. In the bottom graph I have subtracted the height on the 15th from that on the 14th - to give a 'height anomaly' - a measure of how different the height of refraction was on the eclipse day.

To be honest I was astonished with this result. The times line up well, and the result is smooth and not noisy. In fact, it is a less noisy record than could be obtained from professional ionosondes.

I think this is a tremendous tribute to what can be done with relatively inexpensive amateur equipment, recalling that here the transmitter was a homebrew WSPRSONDE from Paul Elliott WB6CXC and the receiver a commercial KiwiSDR. Furthermore, there's been no need for custom digital signal processing: the developers of WSJT-X have done it all for us.

Having mentioned Steve Franke for the 'F' and Joe Taylor for the 'T' it is only fitting, as I draw to a close, to acknowledge the 'S' in FST4W - the late Bill Sommerville G4WJS.

Slide 29

Looking forward to the total eclipse of 8th April 2024 - I suggest there is an interesting opportunity for UK amateurs to contribute. The eclipse itself will have ended with dusk over the mid Atlantic but you can see that great circle paths would carry our signals well into the eclipse zone. Here's an example path from me to Peter K6RFT, with some other WSPR top spotters shown in yellow. On 15 metres the second and third hops would likely be affected by the eclipse. This is akin to the fascinating path from Costa Rica to the western US that I showed. I'd like to encourage UK WSPR users to transmit and receive on 15, 12 and 10 metres on the 7th through 9th April and in return I'll work on the resulting data.

Slide 30

To end ... I hope I've captured tonight some of the sense of fellowship, adventure, technical challenge, and quiet achievement that I've encountered on my return to amateur radio. This has been my niche - to use the tools developed by others - be it hardware or digital communications protocols - and to extract information beyond that envisaged by the originators. And although there is no science breakthrough here, I've found it fascinating to work on this simple data to bring to life details of propagation on the HF bands. This is high quality information that has been gathered by the amateur community, and is safely within the triplicated WsprDaemon database.

Slide 31

I'll leave this slide up - here's where you can find out more.

Thank you