Propagation at HF: What can learn using digital modes WSPR and FST4W?

A talk over Zoom in the RSGB Tonight @ 8 series 5 February 2024

Gwyn Griffiths G3ZIL

Slide 1

Good evening, it is my pleasure and honour to give this talk on what we can learn about propagation at HF from reception reports of digital modes.

In particular I'll show how data from WSPR, and its stable-mate FST4W, can tell us whether propagation was via one-hop, two or more hops, a chordal hop or by a mode that was quite a mystery to me.

Slide 2

Here's an outline of my talk. After a few words on my amateur radio and technical background I'll show some early efforts with hardware for WSPR. One thing that intrigued me was that WSPR spots could suddenly disappear, and there were times with hardly any WSPR decodes. I wondered why. I was keen to know more about how signals actually got from here to there. This talk is the story of how I have re-learnt, and seen for myself, HF propagation basics as in, say, the RSGB Handbook. But, I was unable to explain propagation on several paths from descriptions in the Handbook. The explanation came as a surprise and along the way I learnt valuable lessons in just how useful digital modes can be for propagation studies.

Slide 3

As a teenager I was so proud of being an associate member of the RSGB. Pat Hawker's Technical Topics articles in Radio Communication were an important and memorable part of my technical education. Equipment from those years forms part of the 'yesteryear' corner of my shack.

Slide 4

My professional career - the engineering and applied physics side of oceanography - had little to do with radio. But there were parallels. For two decades I worked on sonar systems to study biology and physics beneath the sea. This graphic from the Strait of Gibraltar has depth on the y-axis, from the surface down to 300 metres. Time is on the x-axis, covering seventy-five minutes. The colour scale shows the strength of the acoustic echoes at 38 kHz, a wavelength of four centimetres. What's shown is a packet of internal waves at a depth of 150 metres triggered as the tide changes.

I've drawn on that professional experience of probing the ocean with sound to probing the ionosphere with radio waves as a radio amateur.

Slide 5

On retirement I returned to amateur radio. After a little while I became interested in WSPR as it offered a mix of hands-on design and construction and the challenge of understanding what reception reports were telling me about HF propagation. My little palm-sized direct-conversion receiver shown on the right worked very well. The trickiest part was adjusting the three trimmer capacitors of the front-end crystal filter to position the passband at the WSPR band and the notch at the lower sideband image frequency.

Slide 6

I persuaded an old colleague, Dr Jeremy Wilkinson of the British Antarctic Survey, to take a 30 metre version of my receiver with a Raspberry Pi running WSJT-X for decoding, with wire and coax for an inverted V dipole, on an expedition to the Arctic on the Korean research icebreaker *Araon*. He joined the ship in Alaska, set up the equipment, had it running for some eight days, and received three spots. Three spots in eight days! None of those were from the multitude of stations in North America. What on earth was going on? Were there hardware problems with my receiver? Had the crystal oscillator shifted? Had the trimmers of the crystal filter moved? Was the Raspberry Pi keeping time over the network? Was the antenna actually connected? Was the propagation really that bad to the Arctic?

Slide 7

Those were questions that I could not really answer. But I knew someone that might be able to help.

I'd made four voyages to the Polar Regions on the Royal Research Ship *James Clark Ross*, and knew Mike GM0HCQ by reputation. Through this collaboration we were able to use a Sailor HF receiver on thirty metres with WSJT-X on a PC decoding WSPR and have it running from Southampton to the Arctic and back. This gave us a good starting point. As I live in Southampton I'd be able to compare the number of decodes from the ship at the start and end of the voyage with my own. And Mike would be keeping his eye on the equipment.

say 8 minutes to here

Slide 8

The map on the right shows the ship's track up the North Sea to Svalbard, and onward to 82° North. Black dots show locations of European WSPR transmitters. The circle is the typical skip zone from Southampton that I'd not expect to decode. Cyan and red arcs indicate the auroral region.

On the left is a graph showing spots received. The Y-axis is distance of each spot from the ship. The X-axis is day of the year, day 180 is 29 June and day 230 is 18 August. As the ship left Southampton and travelled north more spots were received. Spots within the skip zone when at Southampton were being received. But as the ship continued north, we saw fewer spots. A glitch meant no data around day 198, but as is crystal clear, very few spots were received with the ship in the far north. Not as low as three spots in eight days as on the *Araon*, but sparse nevertheless. Here was evidence that the poor outcome with my receiver on the *Araon* can't have been entirely down to deficiencies in the hardware - propagation effects must be at play.

Slide 9

This is a graph from Dr Steve Franke, the 'F' in FT8, which deserves to be better known. It shows that it isn't just the signal to noise ratio - SNR - of a WSPR spot that affects decode probability. Doppler frequency spread is also important. Doppler spread is a broadening of the signal spectrum that takes place in the ionosphere due to movement, velocity gradients and turbulent motions. There is also frequency spread due to jitter or phase noise in oscillators within transmitters and receivers. It all adds up, hence I'll use the more general term - frequency spread - in this talk.

Users of WSPR often talk about the threshold of about -30 dB SNR for a decode. We see that it requires frequency spread of less than one Hertz to achieve that performance. As the spread increases one needs an increasingly greater SNR to decode. Being a very narrowband mode makes WSPR more susceptible to frequency spread that wider modes such as FT8.

To avoid confusion I've noted that the spread values on this graph are those between -10 dB points on the signal spectrum. All values I will show later are between -3 dB points, hence smaller - about 40% of the values shown on Dr Franke's graph.

Slide 10

To make progress I needed to measure frequency spread as well as SNR. But how? It was easier than I had feared. Since 2021 there's been a WSPR-like mode in WSJT-X - FST4W - that can be made to measure frequency spread the width between -3 dB points of the spectrum. You turn it on by placing an empty file named plotspec in the directory in which WSJT-X is started.

Before I could carry out an experiment to measure frequency spread on a range of paths there were four hurdles. The first was easily dealt with. Since 2020 I'd made some small contributions to a substantial project of Rob Robinett AI6VN: WsprDaemon. WsprDaemon is a decoding and reporting system for WSPR and FST4W that is more robust, extensible and capable than wsprnet.org. It was easy to add frequency spread measurements into its database.

Slide 11

As for the no users hurdle, I've lucky enough to collaborate with an informal group of amateurs in the US using WsprDaemon. They were willing and able to use FST4W and report frequency spread. As for the equipment-induced frequency spread hurdle, most were using KiwiSDRs, which, as standard, comes with a GPS aided master oscillator. It is not phase locked, or even frequency locked, but it's a decent starting point.

The experiment needed one station to be within line of sight of the transmitter, that station, WW6D did not have a KiwiSDR but an Elad FM Duo with an external OCXO clock. Logging was with standard WSJT-X. After adding the empty plotspec file frequency spread values appear on the console, and Doug kindly copied those manually.

Slide 12

I'm grateful to Lynn WB7ABP, Santa Rosa for providing transmissions of FST4W-120, the two minute long variant. His Apache Labs ANAN SDR transceiver with a GPS disciplined master oscillator assured very low jitter and phase noise and essentially no drift.

To summarise, I was anticipating I'd see frequency spread as on the middle graph on the line of sight path.

In the bottom graph is a simulation of frequency spread for a decent WSPR transmit receive pair - a drift of 0.1 Hz over two minutes, but as we will see, this is not good enough for propagation mode studies.

Slide 13

The purpose of the first experiment in December 2022 on 14 MHz was to gather frequency spread and SNR data over enough paths to be sure of seeing instances of surface wave, one hop and two-hop propagation, as illustrated in this PyLap ray trace diagram.

I'd not expect to see Near Vertical Incidence Skywave on 14 MHz, as the F2 layer critical frequency was unlikely to be high enough. I would expect to see instances of no propagation - when a receiver was within the skip zone from the transmitter. That is, it was closer to the transmitter than the landing spots of one-hop paths.

Slide 14

Here is a map showing the locations of the transmitter, in red, and the receivers in yellow. We'd gather data on mid-latitude paths, mostly around 40 degrees North, but also on a path to the Arctic Circle. We could compare and contrast measurements on the over-water path to Maui to the over-land path to Long Island.

Time does not allow for all the results to be described tonight, but the main findings will become clear.

Slide 15

For each path I'll show a pair of graphs. Both are time series, time in UTC on the X axis. The top graph shows SNR in a 2.5 kHz bandwidth, the bottom graph shows frequency spread in milliHertz - let me repeat, milliHertz. These results are for the 2.4 km path from WB7ABP to WW6D. This shows that this set of equipment, a GPSDO transmitter and an OCXO receiver, did not add measurable spread to that intrinsic to the 120-second length variant of FST4W. I got the same spread at baseband when I connected the audio from one computer running WSJT-X to another as receiver.

some 18 minutes to here

Slide 16

Now we've increased range to 40 km. Recalling that we are using 14 MHz my expectation was for a weak surface wave signal with low frequency spread as it had not been near the ionosphere. But the results were not that simple.

Looking at the SNR time series, I've added curves in cyan for sun's altitude to show local daytime. The two nights A1 and A2, show little variation in SNR, although levels on the two days were different. Daytime SNR was more variable.

On this information you might shrug your shoulders and say no-one expects SNR to stay constant.

But look at frequency spread. During night time, A1 and A2, are much the same, but there's a dramatic increase in frequency spread during the day.

Top right is a scatterplot of frequency spread on the Y-axis with SNR on the X-axis. The contours help pick out distinct clusters.

While clusters A1 and A2 are separated in SNR, both have the same median frequency spread of 59 milliHertz. Cluster B is quite different: its median frequency spread is over ten times greater, at 624 milliHertz. What's clear is that there is no 'middle-ground' - these are two different propagation modes. A1 and A2 are undoubtedly surface wave as I expected, but I had no idea about mode B.

We also see that a GPS aided KiwiSDR adds more frequency spread that we'd like, suggesting we should use GPSDOs for future experiments.

Slide 17

Here we've increased range to 960 km. Propagation must be via the ionosphere. In the SNR time series we have the sun altitude curves. I've labelled three intervals. I1 during the day had to be one-hop. Its frequency spread was reasonably low, its median was 87 milliHertz. The receiver was a GPS aided KiwiSDR, as we'd seen a median spread of 59 milliHertz on the 40 km surface wave path I definitely got the impression that the actual frequency spread on a mid-latitude one-hop path was likely to be lower than we are able to measure with the GPS-aided KiwiSDR.

Mode I2 is still a mystery to me: it had a decent SNR and low spread - it was possibly sporadic E - not unknown in midwinter.

Mode I3 is quite different: it had a low yet pretty consistent SNR and a high frequency spread. The median spread at 623 milliHertz was virtually identical to the 624 milliHertz of our mystery mode B on the 40 km path. But that seemed very strange - implying that during the day there was a propagation mode on 14 MHz via the ionosphere on a 40 km path!

Slide 18

Lets push the path length out to 3762 km over the Pacific to Maui where AI6VN has an unattended GPS-aided KiwiSDR. During the day we had mode I4, high SNR but with significant frequency spread, a median of 266 milliHertz. At night, mode I5 had a lower SNR and lower frequency spread, a median of 83 milliHertz. Recalling that one-hop on the 960 km path had a median spread of 87 milliHertz, I was pretty sure mode I5 was one-hop. PyLap ray tracing showed one-hop was possible. Mode I5 was therefore very likely two-hop, again confirmed using ray tracing.

So we start to build a picture of frequency spread: one-hop as less than 100 milliHertz, two hop at about 260 milliHertz. With mystery modes B and I3 showing frequency spread over twice that of two-hop.

Slide 19

Next let's look at a 679 km path, which I'd expect to be within the skip zone at 14 MHz. Around the middle of both days I've labelled mode I1 on the SNR and scatterplot graphs. Because of its high SNR and low frequency spread this was one-hop propagation. But far more common, happening before and after one-hop, was mode I3, which we now recognise through its fingerprint of high spread and low SNR.

Whatever this mode was, it's not rare, we've seen in on paths from 40 to 960 km, at night and during the day, day-in and day-out. It appears to be an ionospheric propagation mode that seems to work 'above the maximum usable frequency'.

Slide 20

To show that this mode does not just occur over North America, my own 14 MHz FST4W signals from Southampton received by Nigel G4HZX Beckenham, South London, a path of 108 km during the day had a median spread of 614 milliHertz - it's the mystery mode again!

Slide 21

After much reading, I began to think the mystery mode was two-hop sidescatter. Let me explain, in this diagram we have the 960 km one-hop path from WB7ABP to KA7OEI in white. As evening progresses, and the F2 layer critical frequency drops, the ray landing spots from California move beyond KA7OEI, and the receiver falls into the skip zone.

Let's say the landing spots now start at 1300 km, a one-hop path might take the signal to the ocean off Baja California. The large majority of the signal power will be reflected from the ground in the forward direction, the dashed line. But some small fraction will have been scattered at all angles, one ray could follow a path to KA7OEI via one hop propagation. Hence the name for this mode - two hope sidescatter.

Slide 22

I've tested this theory for several paths. Tonight I'll look at the path where we've seen the mystery mode to dominate, from WB7ABP California to Rick, KK6PR in Oregon. When we look at a propagation prediction output from VOACAP there's no likelihood of an open path at 14 MHz. But one clearly exists...

Slide 23

Simple ray tracing is no help - it just confirms KK6PR well within the skip zone. But what is useful is that the ray trace shows a focus of landing spots at the edge of the zone where one-hop reception becomes possible. This is because higher angle rays from the transmitter are refracted at higher altitudes in the ionosphere than those at lower angles. We have an aggregation of ray landing spots near the edge of the reception zone. This will prove of interest later, as indeed will be the idea of 'ray landing spots'.

Slide 24

I've shown two-dimensional ray traces from PyLap - and I've introduced the idea of 'ray landing spots'. I'll now use PyLap in 3D mode. Think of it as conventional 2D ray traces stepped the bearing from the transmitter in one-degree increments. When we do that for the transmitter at WB7ABP we get the ray landing spots in black shown on the first map. KK6PR is well within the skip zone. In an academic paper on two-hop sidescatter the authors placed a transmitter at the site of each ray landing spot, and mapped its landing spots. That is a tedious approach and needs a lot of computing time.

Instead, I assume that reciprocity holds well enough for my illustrative purposes. And so I place a transmitter at the receiver location and calculate its landing spots, in magenta.

For step three I multiply the number of landing spots from the transmitter and from the receiver in each one-degree by one-degree latitude longitude box - for example, in this tiny orange box. This gives me a measure, a metric, of the likelihood of backscatter from that area, everything else being equal. From this map we can get an impression that there's much overlap - hence high values of the likelihood metric - to the west and to the south-east.

Slide 25

This becomes clearer in a contour plot of the likelihood metric - the highest values being the regions of most intense red. This is where the model suggests sidescatter happens at 0400 UTC. The thin-ring form is because of the ray path geometry and from the focussing effect of high and low rays near the edge of the skip zone

But is this model result at all credible?

Slide 26

Lynn, WB7ABP was kind enough to perform a rotating antenna experiment. He'd rotate his five element 20 metre yagi in ten degree steps every four minutes. If the path was over the great circle to KK6PR the maximum signal level would be on a heading of 010 degrees. It wasn't. Instead, the maximum was with a beam heading of 320 degrees. That's forty degrees clockwise of the model prediction, which I consider a fair result. Plotting the Yagi's beam pattern via surface wave at a distance of 2.4 km gives the asymmetric plot in green, but we do not really know what the beam pattern actually looks like at elevation angles for a range of 1500 km.

Slide 27

We now have some sense of where sidescatter may occur at 0400 UTC - but how might it vary through the day. Here is an animation over 24 hours, starting at 0000 UTC, which is afternoon, 1600 UTC, local time. The afternoon peak moves to the west during the evening and weakens before becoming an arc to the south. In the early hours local time a peak forms to the east and around noon flips to and fro to the west, before settling to the west in the afternoon.

Again, we can ask, is this result at all credible?

Slide 28

Unfortunately I don't have a twenty-four hour rotating antenna experiment to show. Instead, let's look at the variation of signal level through the day, as a scatterplot with contours and as a histogram of spots received in each hour. On the small map I've plotted the locations of the maximum sidescatter metric for each hour. I'll ask for your indulgence in not expecting exact correspondence in time between model and actual data, concentrating more on the key features. From about 09:00 to 14:00 UTC there was a drop in the number of spots decoded and their signal level - I suggest this was when we saw that very thin arc to the west, moving to the south, and then to the east, at the greatest round-trip distances from transmitter to receiver. From 14:00 UTC the signal level increased somewhat, and the number of spots rose.

But there was a dip in the number of spots decoded even though signal level was height between 21:00 and 01:00. I suggest that this is because of excess frequency spread, with signals at the receiver arriving from both the land to the west and the ocean to the west. We see from the map that the maximum metric location flipped back and forth.

about 36 minutes to here

Slide 29

Leaving two-hop sidescatter let's look briefly at Chordal Hop - a propagation mode where there is an ionosphere to ionosphere hop that avoids a ground reflection as here, on a 6920 km path on 14 MHz. I've come clean on the slide - I can't *quite* get the rays to touch ground at K6RTF. I tell myself that the model of the ionosphere within PyLap is more about climate that day-to-day ionospheric 'weather' - and chordal hop is not be seen every day. Greater than modelled ionisation at the second hop near 5000 km would be needed for the rays to reach ground at K6RFT.

This graph is of frequency spread against time of day for ten days. On just two of those days, only in the morning, there were clusters of spots with low frequency spread - 100 milliHertz or less. From all I've learnt to date - given the low spread, these spots had not reflected off ground or ocean - it's quite possible they were chordal hop.

Slide 30

Having explored frequency spread on several paths I return to where I started - the Arctic - on a path from Southampton to VY0ERC on Ellesmere Island on 14 MHz. The scatterplot on the left is for my path to Tom, WA2TP Long Island. This provides a mid-latitude reference, with a median spread of 296 milliHertz. The shorter path to VY0ERC has over double that spread - and only a quarter of the spots decoded. I carry over the peaks from each graph to the other - to show they are completely different clusters. I am pretty sure that there were spots with frequency spread of 800 milliHertz and above that were not decoded. Excess spread was likely the problem - not simply lack of SNR.

Slide 31

I applied many of the techniques I've shown toning to study effects on HF band propagation of the 13 October eclipse over North America. You can find my talk on the Norfolk Amateur Radio Club YouTube site. There were fascinating results on the 5000 km path from Costa Rica to the US west coast. Propagation was interrupted on the second and third hop on 28 MHz at distances of around 4400 km and circuit reliability, top graph, reduced at 5000 km. Those results set me thinking - For the 8th April 2024 eclipse the second and third hops on 21 MHz and up from the UK to Eastern North America will fall within the zone with the sun obscured by 80% or more. I encourage UK WSPR users to try 21 MHz and up on the 7th through 9th as a contribution to the more local studies in the US of this eclipse. I look forward to working up the results.

Slide 32

To close, here is a summary of what I have learnt over the last eighteen months of studying frequency spread at HF. Please appreciate this is work in progress and may be revised and expanded.

Indicative equipment-derived spreads are in magenta. I was surprised just how low frequency spread could be via the ionosphere, and it's only through using GPS disciplined oscillators both ends that we can properly measure such low spread. A plain crystal oscillator is just not good enough for these studies.

Two-hop sidescatter has been a revelation - and deserves further study and appreciation.

It may seem that the information within a WSPR or FST4W reception report is minimal and ephemeral. But, as I've shown tonight, there's a great deal we can learn about HF band propagation from these measurements, made carefully and systematically.

Thank you.

Slide 33

Leave this more info one up.

4204 words

Slide 28

Slide 29

Slide 30

Slide 31

Thank you