FST4W on the HF Bands: Why - What to expect - Equipment - Results.

Gwyn Griffiths G3ZIL, Glenn Elmore N6GN, Rob Robinett AI6VN, Lynn Rhymes WB7ABP, John Watrous K6PZB

Corresponding author: Gwyn Griffiths gwyn@autonomousanalytics.com

Slide 1

Good day. My aim today is to show you that FST4W - a digital protocol akin to WSPR - is of real value on the HF bands and isn't just useful at LF and MF. All co-authors have made substantive contributions to gathering the information to show you today and we are grateful to our collaborators acknowledged here for their enthusiastic assistance.

Slide 2

Why bother with FST4W? Well, there's the nice to haves - a decode threshold about 1.4 dB lower than WSPR for the 120 second variant. And the longer variants have even lower decode detection thresholds. FST4W also has greater tolerance to Doppler spread, decoding at 7 dB lower SNR than WSPR at 2 Hz Doppler.

But there are downsides. If using WSJT-X you have to set one sequence length to listen for. Admittedly, there is a technical challenge - there is no drift compensation as in WSPR - and as sequence lengths increase the requirements on equipment spectral spread become more challenging at HF. The waterfall from N6GN on the right shows a real FST4W transmission at its 33rd harmonic - smearing from one tone frequency to another because of equipment or ionospheric spread can set a limit to the mode's usefulness. It is understandable that these downsides mean few people are using the mode.

But, as we will show, each of these downsides provides a challenge that we are already overcoming.

Slide 3

Here's the Game-Changer for why FST4W at HF deserves more attention - it measures spectral width, although not by default, and it's a little tricky to turn on.

This means FST4W gives you the critical tool to counter the downsides of spectral spread whether from equipment or the ionosphere. In the graph of decode threshold SNR against spectral spread for FST4W and WSPR at right I've showed two regions where spots are unlikely to be decoded, SNR too low and spread too great. At times with WSPR it is not clear which is at play. With FST4W we can be.

On first thought it seems almost unbelievable that this 900-second FST4W transmission from Fort Collins to Santa Rosa was not decoded - but it wasn't - too much spectral spreading despite the very high SNR.

And - the final why - once we have backed off equipment spectral spread the remaining spread from the ionosphere may be of interest to the science community.

Slide 4

As to what to expect here's an example time series of SNR in the top panel of FST4W-120 transmissions from Fort Collins to Maui. As in all the examples today the transmitter uses a phase-locked GPSDO - keeping one variable fixed is of great help. Here the receiver is a standard KiwiSDR with its out-of-the-box frequency aiding, before we'd uncovered a bug.

The bottom panel is the game-changer - a histogram of spectral width for the decoded spots. We'll show later how to split this out into contributions from the transmitter, receiver and the ionosphere.

Slide 5

Getting a feel for what to expect is made easier using the fst4sim program that comes bundled in WSJT-X. Given this set of parameters fst4sim generates the specified number of wav files that one decodes with the jt9 executable. Here we've kept al variables fixed except for fdop, generating 50 wav files at each fdop, and plotting the results.

Importantly we keep in mind that fdop is the total spread fro the transmitter, receiver and the ionosphere, so we label as spectral spread. From 50 runs at each fdop we get the decode probability with spectral spread in red; what FST4W considers the spectral width in green and the output SNR in purple. The output SNR decreases with spectral spread as energy from one tone band leaks to another. Decode probability reducing is because we are not decoding those with greater spectral spread, so we get a biased spectral width that becomes increasingly biased low.

For reference the results from the last slide were all within the grey box, comfortably within the usable range.

Slide 6

We can also use fst4sim to estimate performance at one sequence length, say 900 seconds, where we have spot count and spectral width for a shorter sequence, say 300 seconds.

In orange is the histogram of spectral widths from 14 MHz FST4W-300 transmissions from Fort Collins to Maui. Its right hand tail was already extending into the region of reduced probability of decode, the solid red curve. But look at the decode probability curve from fst4sim for dash 900, dotted red. When we multiply that probability with the spot count we get a predicted spot count for dash 900 as the brown curve - just 98 out of 431 spots for dash 300. In fact it was worse than that, Rob decoded only 8 dash 900 spots in two days.

Slide 7

Let's move to equipment. Given what we've just seen it is not surprising that early tests with FST4W on the HF bands were disappointing. Kenji JJ1BDX has kindly allowed me to share his graph and experience. He could only get his own FST4W-300 transmissions to decode at no higher than 3.5 MHz. The drift on his FT-891 was just too large on transmit - it's fine on WSPR where linear drift is removed, but not on FST4W. The receive side of commercial transceivers can also be problematic - temperature rise on a transmit cycle may result in excessive frequency drift in the following receive cycle - it may take several cycles for the TCXO to become sufficiently stable to decode FST4W.

Slide 8

So what equipment does work? All three transmit stations in this study used variants of ANAN software derived transceivers with phase-locked GPSDOs. These performed excellently. We have made an initial study of the remarkable $69 QRP Labs QDX digital modes transceiver, but its novel single signal generation method results in too much spectral spread on transmit in its current implementation. It is fine on receive.

KiwiSDRs have provided the vast majority of spot data for this study, about 50/50 between those with out-of-the-box GPS aiding and external phase-locked GPSDOs. And an Elad FM Duo with GPSDO works excellently.

Slide 9

Here's how we've got rid of a major impediment to uptake on receive. Rob Robinett's WsprDaemon V3 allows the KiwiSDR user to declare which modes from WSPR to all FST4W lengths to decode for each receiver and band. By recording one minute wav files and concatenating as needed to present to the wsprd and jt9 decoders only one Kiwi receive channel is needed to decode all these modes on one band simultaneously. That's a big step forward to the one-at-a-time approach in WSJT-X.

Moreover, Rob switches on measurement of spectral width and logs the data, and is working on extracting the tenth-of-a-Hertz frequency resolution that FST4W can also, optionally, provide.

Slide 10

Moving to results - let's compare how many spots with SNR at or lower than -30 dB are decoded by KiwiSDRs with standard frequency aiding and those with a phase-locked GPSDO. Remembering that, for decoding at the lowest SNRs we need low spectral spread. Or should we not bother? Might not ionospheric spread dominate? Not on the 1500 km path from Fort Collins to the Bay area for FST4W-300. A test for no difference was rejected - KiwiSDRs with GPSDOs won hands down in this test on 7 MHz. It wasn't down to antennas or location, the GPSDO receivers were in urban or suburban locations while KFS and KPH with their log periodic antennas and quiet locations just had standard aiding.

Slide 11

This is how we've gone about ascribing spectral spread to receiver, transmitter, and path. In outline, we started with a GPSDO transmitter and receiver pair 21 km apart over a line of sight path. We gathered FST4W-300 spectral widths, histogram top right, and fitted a Gamma distribution. Gamma because it is positive only and copes with extended right tails, which is what we have seen in virtually all cases. The fit is in red.

If we assume no spread over the line of sight path, and as we have identical GPSDOs at both ends, we ascribe measured spread 50/50 to receiver and transmitter.

We can do this by halving the Gamma distribution Shape factor, keeping the Scale factor constant. Now we have the parameters to generate numerical simulations of the spectral spread of GPSDO transmitters and receivers individually. That's the middle histogram.

Next was an experiment over a 38 km path between a GPSDO transmitter and a standard KiwiSDR receiver. We again gathered spectral width data for a histogram and fitted a Gamma distribution. From the parameters we generate a numerical simulation. We make the perhaps contestable assumption of no spread from the path and subtract spot by spot the numerical simulation of the GPSDO transmitter spread from the total spread to arrive at a Gamma distribution for the KiwiSDR spread alone, bottom right.

Slide 12

Having estimates of spectral width from a GPSDO transmitter and KiwiSDR receivers we can subtract their contributions from the measured spectral widths over example propagation paths to get an estimate of the ionosphere's contribution. This is where FST4W perhaps becomes of interest to science.

All three examples here are on 14 MHz FST4W-120 from Fort Collins, to Santa Rosa, Maui, and New Zealand. Not only is FST4W-120 capable of global paths on 14 MHz it adds this entirely new set of measurements compared with WSPR. Over the path to New Zealand there's not much to choose between the Gamma fit in red and a Gaussian in blue.

One of the few equivalent Doppler spread histograms we found in the open literature is on the right - across the Auroral Oval - note the long tail well suited to a Gamma distribution.

Slide 13

Moving from histograms to time series and to comparing two bands we see a great deal that we, as radio amateurs, just do not understand. The three panels on the left are for 7 MHz FST4W-120, those on the right for 14 MHz FST4W-300. The top panel is SNR - much as WSPR - the middle panel is spectral width and the bottom panel is the rate of change of spectral width with time.

First, the spectral width behaviour on 7 MHz was quite different to 14 MHz - why, we have no idea. If anything, the geomagnetic conditions were somewhat more disturbed during the 14 MHz trials.

Second, note that the spectral width on 14 MHz on 6 September was essentially constant over a 30 dB change in SNR - that surprised us.

Third, equally surprising, were the large changes in spectral width on 7 MHz between measurements just four minutes apart, except for around 0700 UTC, that is, about local midnight at the mid point of the path. But even then, the magnitude and variation of the spectral spreading was greater than on 14 MHz.

Slide 14

Recognising that there is still a great deal for us to understand, here is a graphic that tries to give a quick-look assessment of FST4W for 14 MHz based on our results.

There's a plot for each sequence length and a colour code for likelihood of decoding spots. Both X and Y axes are spectral widths - with example paths on the Y axis and different transmitter/receiver combinations on the X.

FST4W-1800 is only usable on 14 MHz over line-of-sight paths and with GPSDOs each end. FST4W-900 over one-hop paths is feasible but only likely with a GPSDO transmitter and a receiver with a narrower spectral spread than a KiwiSDR.

For FST4W-300, as we saw from Fort Collins to Maui and with a KiwiSDR, decodes were possible, but much reduced. With a receiver with lower spread, the decode probability would improve. But Global paths are unlikely whatever the receiver/transmitter combination.

Slide 15

Finally, FST4W is not just useful at LF and MF, it is entirely practical on the HF bands with readily available equipment. However, it does need care to minimise frequency drift and spectral spread of transmitters and receivers. We've shown tangible benefits of GPSDO receivers in decoding spots with low SNRs.

Most importantly, FST4W adds a new variable to measure - spectral width. We've suggested a method to separate out the spectral widths of transmitters, receivers and the propagation paths. After doing so we've seen intriguing patterns in spectral widths worthy of further study.

We've further work to do on characterising spectral width following changes to the KiwiSDR, and are in conversation with QRP Labs about the spectral width of the QDX transceiver on transmit. Personally I'd like to understand the mode's curious SNR algorithm and there are endless possibilities to study SNR and spectral width over interesting propagation paths.

Do give FST4W at try.

Thank you.