Contrasting Effects of the 3–5 November 2021 Geomagnetic Storm on Reception in Colorado of WSPR transmissions from North-Eastern North America and those from Australia

Gwyn Griffiths and Glenn Elmore

1. The 3rd to 5th November 2021 geomagnetic storm was a well-documented natural disturbance to the ionosphere. It coincided with some unusual observations of WSPR spot reception by co-author Glenn Elmore in Colorado. Our presentation will try and answer the question, "Was there a causal Space Weather Connection?"

2. Glenn has a remote WSPR receiving site at seven and a half thousand feet in the foothills of the Rockies, it has the benefit of low local noise. Our study concentrates on the seven Megahertz band, where the antenna is a Loop on the Ground. WSPR spots and noise levels are reported using Rob Robinett's WsprDaemon software. To investigate whether the spot count collapse and the increased median distance were connected to the geomagnetic storm we'll look at the WSPR data, the contemporaneous online narrative, and data on the space weather event including excess D region absorption and indicators from ground and space of disturbance to the geomagnetic field.

3. Here's what we mean by unusual patterns. What we have here are heatmaps of noise level, number of spots and median distance. Each column is one day, and each bin in the vertical spans thirty minutes from midnight UTC at the top. The color scale shows increasing values, from blues to reds. November fourth is in the centre.
   You don't need me to point out the spot count collapse on the fourth, from over 200 in thirty minutes to under 20. The increase

in median distance is from the typical two thousand five hundred kilometres or less to over twelve thousand. The spot count and median distance anomalies persist from about oh nine hundred to thirteen hundred UTC.
More subtle, and longer lasting, is a slight reduction in noise level. Glenn's remote site is limited by propagated-in noise. It looks like propagation changed to reduce the incoming noise, coincident with the spot count collapse.

4. Here's a summary of how the online sites spaceweather and spaceweatherlive reported the storm. The precursors were M class solar flares on the first and second November. The speed of the particles was such that the later, faster ones caught up the earlier, slower ones - hence the description of a "Cannibal" Coronal Mass Ejection.
   The result was a G3 class geomagnetic storm spanning the third and the fourth. By the fifth the storm was over, and reports of aurora had come in from California, New Mexico and, of particular interest to us, Colorado.

5. Let's look as some indicators of the space weather. Images of the D region excess absorption are readily available on the NOAA website. I've compiled images one minute apart into these animations. Here's what was happening in the D region during the time of the WSPR spot anomalies on the morning of the fourth. [Run bottom right movie] There was a modest absorption event primarily affecting Africa and the southern hemisphere centred at eleven hundred UTC. No effect over North America. [Run top left movie] In fact the major D region absorption event happened on the day before, at nineteen hundred UTC as the CME hit the earth's magnetic field. At that time, there was no standout event at Glenn's receiver.

6. With excess D region absorption ruled out let's look at indicators of geomagnetic disturbance. Top we have Kp, here the planetary value, which was above storm level from midnight on the third to midday on the fourth.
   Centre we have the magnetic field components from Boulder, Colorado; horizontal (in blue) and vertical (in red). The dip in the vertical component reaches a maximum just after oh nine hundred UTC, returning to normal by about fifteen hundred. This looks a promising space weather connection for our WSPR anomaly.
   The auroral Electrojet index gives a circum-arctic measure of the event from these twelve stations.

7. Some of the orbiting satellites of the US defense Meteorological Satellite Program carry spectrometers covering the short wavelength ultraviolet region where glow from Aurora can be quantified. These images are readily available via a selector at this Johns Hopkins University website.
   Satellite passes are every hundred and one minutes, from north to south and from east to west. Here I've created a composite of three such images for the morning of the fourth November. They are disjointed as the aurora changed faster than the orbit period.
   What's readily visible is that there was auroral glow between grid FN, Eastern North America, and Colorado. And, that the sector of paths to Australia was mostly unaffected - which is what was needed to result in a twelve thousand kilometer median distance. The space weather connection is starting to come together.

8. While the heatmaps are great for quickly spotting anomalies I like to look at time series graphs for detail (and I can control the font size). Here we have the WSPR spot count in 20 minute intervals from grid FN, Eastern North America, to Colorado.

The x axis is time of day UTC. The dots with fine lines are the undisturbed days - showing both day to day variation and an underlying pattern with smooth-sloped, sine-like dips centred on local noon.
The thicker purple line is for the fourth. By oh nine hundred the spot count had dropped well below normal, with few spots received until noon, while by fourteen hundred the pattern was similar to that of a normal day, although with more spots - recall that the noise level was lower during and after the event.
The pattern for the following day is also disturbed until noon UTC.

9. While we've seen this notable event at Glenn's Colorado site, how widespread was it? WSPR gives us an excellent method of answering that question. I've created this animation of WSPR transmitter locations in green, receivers in magenta and red and propagation path midpoints in cyan and blue - admittedly this is a simplification.
   Each frame has spots now in ten-minute intervals, the previous plots having shown the suddenness of the spot count collapse. Time is given in the box, upper right. As a baseline, receivers that hear spots at the indicated time on the first of November are in magenta with the path mid points in cyan. Of course, as the previous plot showed, there is day to day variation, such as here at oh five fifteen, but there is a broad similarity between the path mid points for the first in cyan and the fourth in blue.
   By oh eight hundred we've lost paths, especially those with more northern mid points. At eleven hundred we only have two spots remaining.
   This was clearly not an event just for Colorado; this was a continent-scale event. Recovery is initially patchy, with the northern paths recovering later, and, of course, by this time of day, absorption reducing the overall number of spots.

10. What would be needed to cause such a widespread collapse in WSPR spot count over North America?
Three possibilities are: a much reduced signal to noise ratio, a much increased Doppler spread, or a combination. Let's quantify those factors. Looking at signals at N6GN from one of the stations with highest SNR, KC2TER New Jersey, in this time series from twenty hundred on the third to twenty hundred on the fourth, the "cliff edge" is clear - there is no spot with an intermediate signal level between -119 dBm and the limit of detection at about -142 dBm. We'd need the signal level to drop by over twenty dB.
We're grateful to Steve Franke for providing this graph of WSPR decode probability against SNR for a range of Doppler shifts. Recall that WSPR tones are one point four six Hertz apart. There's a graceful degradation out to two Hertz, but by three Hertz the probability of decode is less than fifty percent whatever the SNR. Can we see evidence for either or both of these causes?

11. Here we have heatmaps of the number of spots in SNR bins of two dB on the y axis and ten minute intervals on the x axis for the same set of western North America receivers as in my earlier animation. If Doppler spread over two Hertz dominated we'd expect to see a transition to few spots with low SNR, spots with higher SNR would dominate. While, for the fourth, there appears to be some loss of low SNR spots compared with other days, the transition to Doppler over two Hertz may have happened so quickly we don't see much of a signature.
If loss of SNR dominated we'd expect to see a transition to lower SNR - but I don't see it - the "cliff edge" dominates. It's difficult to draw a robust conclusion.
Incidentally we see that recovery took some three days.

12. Perhaps we can gain a better insight from ionosonde records. Unfortunately we only have data at either end of our path - I've chosen to show that from Boulder. What we see in a rather typical graph at oh seven hundred is a maximum usable frequency of six point seven Megahertz at fifteen hundred kilometers followed by a drop to an MUF of 4.8 Megahertz at fifteen hundred kilometers at oh nine thirty. Note that the electron count, the black line, also dropped.
By ten hundred the echoes are faint, and there is no estimate from the automatic scaler for MUF and no electron count trace. I'd welcome insights from those of you with more experience interpreting ionograms as to what these indicate and the likely affect on seven Megahertz propagation.

13. Some ionosondes, but not Boulder, provide information on Doppler shift from movement of the ionosphere plasma. These are from the Idaho National Laboratory.
Imagine you are looking straight up, so that the centre of these images is overhead, and I've marked North, South, East, and West. The edges are at an angle of depression of about forty degrees, at a range of three to five hundred kilometers. These are at a frequency of two to three Megahertz.
There is significant variability between images fifteen minutes apart, so there is a real risk of selection bias. But ... here goes ... The images at oh three fifty three and thirteen thirty-eight show minimal Doppler, ranging to half a Hertz either side of zero. During the event, at oh nine oh eight, the Doppler is one-sided at plus nought point five to plus one point six Hertz, that is, a movement of the plasma toward the ionosonde location. Fifteen minutes later the echoes are from the southwest, and negative nought point six to negative one point six Hertz, that is, moving away from the ionosonde. The direction of movement is the same at both times, north to south. And, with a Doppler shift of one point six Hertz at three Megahertz equivalent to three point

seven Hertz at seven Megahertz there's more than sufficient Doppler to prevent WSPR signals being decoded.
This is useful circumstantial evidence.

14. Can we say anything about Doppler shift directly from WSPR data? Perhaps. In some circumstances. WSPR spot data in the wsprnet.org database is inadequate; the frequencies have been truncated to one Hertz resolution. Rob Robinett's WsprDaemon logs the one tenth of a Hertz resolution data provided by the standard WSPR decoder. Glenn has a GPS disciplined oscillator for the KiwiSDRs at his remote site, and some transmitters have adequate stability for us to attempt a Doppler analysis.
Here we have data for grid FN stations AC2ZR as yellow dots and N8VIM as purple dots. The time series run from the twenty eighth of October to the November sixth. The upper panel is the variance of each received frequency over a one-hour interval. The lower panel is the frequency anomaly - the difference between a spot frequency and the average frequency of the previous thirty spots.
Visually the fourth of November is different for both metrics and stations. If we have stationary statistics, the distribution of the variance should follow a Chi-squared distribution - a special case of a Gamma distribution. Here I've fitted a Gamma distribution to the AC2ZR frequency variance. The probability that the variance of just over nought point five Hertz squared at oh six hundred was from that Gamma distribution was one point five times ten to the minus five per cent. It was clearly from a different distribution of greater-than-normal Doppler spread. And the lower panel shows values of over two Hertz prior to and after the gap.

15. Finally, let's get to why the median distance leapt to over twelve thousand kilometers. The paths from Australia to

Colorado likely had their last encounter with the ionosphere well to the south of where we saw auroral glow on the satellite images. In the upper graph the black line and dots are the spot counts in thirty minute intervals averaged over twenty eighth October to November second. On the third the pattern was quite different, and the uplift persisted to the fourth.
While the noise at Glenn's site dropped by two dB on the fourth this was not the cause of the much higher number of spots. Looking at histograms of the SNR between the twenty eighth and the second and the fourth there are many more spots with SNR over minus sixteen dB - it's not simply plus two dB due to the lower noise. I've no explanation for this uplift in the number of spots, beyond a speculation that it may be due to the higher X-ray flux from earlier CMEs.

16. In summary, we've shown how a sudden precipitous drop in WSPR spot count at seven Megahertz across North America was associated with a G3 geomagnetic storm. The timing and duration of the drop coincided with a dip in the vertical component of the magnetic field at Boulder and with auroral glow observed by a satellite. Ionosonde records showed reduction in electron count and MUF and increased Doppler - consistent with a lower probability of decoding WSPR spots. In contrast the last ionospheric refractions for the signals from Australia were to the south of the auroral disturbance, and it is possible that earlier higher X-ray flux enhanced rather than diminished propagation on those paths.
And, on that point, we welcome comments, and thank you for your attention.

12 March 2022