The RAGazine

Volume 2 Issue 4 May 2015

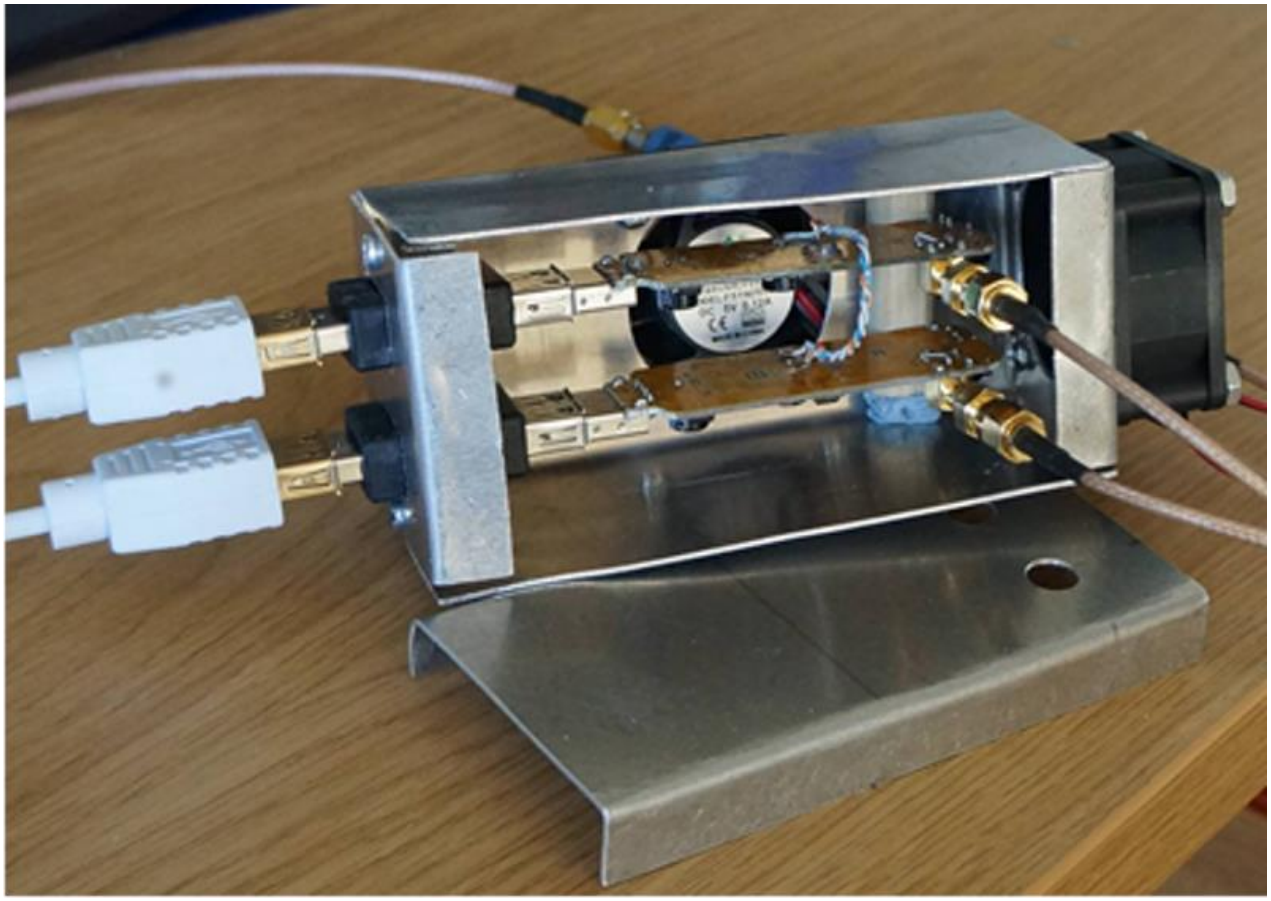

Improving the frequency stability of the RTL2832U SDR

May 2015

BAA, Burlington House,
Picadilly,
London
(+44)207 734 4145
Reg Charity 210769

BAA RAG Coordinator **Paul Hyde** g4csd@yahoo.co.uk
Ragazine Editor **Jeff Lashley** jeffl@spacecentre.co.uk
VLF reports **John Cook** jacook@jacook.plus.com
Web : http://www.britastro.org/radio

Notes for content submission

Content should be emailed to the Ragazine editor a minimum of fourteen days before the next publication date. Content submitted after that may not appear the next issue, but will be held for a later issue.

Observational reports are very welcome on topics where radio techniques are applied to observe astronomical objects, or geophysical events. Articles are welcome on topics of radio astronomy observational techniques, radio hardware and related technology, scientific programming, events, data processing, educational out-reach, book reviews, radio astronomy history etc.

The preferred format for submissions is Microsoft Word (.doc or .docx format). However I know that not everyone has access to Microsoft Office, note that the free office suite LibreOffice is available for Windows, Linux and Mac OS. Note that LibreOffice can save documents in Microsoft Word .docx format. If neither of these applications are available then plain text (.txt) is fine.

Images can be supplied embedded in the document or as separate files. Please include the author credits you require to be included in the article, and indicate whether you want your email address to be shown in the publication.

The BAA is not responsible for the opinions expressed by RAGazine contributors, and such published material does not necessarily express the views of BAA Council or RAG officers. Material without attribution is generally contributed by the editor.

2015 Publication dates and submission deadlines

Release Date	Submission deadline
10th August 2015	27th July 2015
9th November 2015	26th October 2015

RAG Coordinator report

By Paul Hyde

Welcome to the second edition of RAGazine for 2015 and, once again, my thanks to the contributors and to Jeff Lashley for pulling the material together. Please think whether you can offer any material for RAGazine, whether it's from personal work or progress on a club/society project. Your experience helps to inform and inspire others.

Solar eclipse VLF study

We had 17 observers submit reception data for a variety of VLF transmitters during the solar eclipse of March 20th. These reports have been consolidated down into a couple of spreadsheets which are available to anyone who wishes to use the data. There is a short report of the exercise in this edition but the hope is that the information can be used for further analysis of propagation conditions during this period.

Events and talks

As usual, the BAA Winchester Weekend proved very popular with over 150 people attending. The RAG stand had a regular stream of visitors over the Friday and Saturday with particular interest shown in SID and meteor observing activities. The next event for the RAG stand will be the BAA Exhibition Meeting in Cardiff on June 27th.

Future RA talks are scheduled for West Kent ARS (May 11), Radio Society of Harrow (May 15), Wycombe AS (May 20), Heart of England AS (May 28), Basingstoke ARC (June 15) and York AS (June 19). Visitors are welcome to all of these meetings though societies often charge a small entrance fee for non-members. Further information can be found on the BAA website or that of the local society hosting the event. It is always good to meet those with an interest in radio astronomy, whether a beginner or seasoned observer, an amateur astronomer or a radio amateur.

RAG website

We are planning to re-launch the RAG website under a new format, including an extended Resources section containing links to other sites of interest, papers and useful tools. The starting point will be the existing Resources page at http://www.britastro.org/radio/resources.html Please let me know of other links that you feel would be useful to add to this.

Any offers of help in building up the content available through the RAG website would be very welcome.

Best wishes

Paul Hyde

g4csd@yahoo.co.uk

Solar activity first quarter 2015

John Cook

Fig 1 shows activity levels since 2005. Sunspot numbers are courtesy of the BAA Solar Section. SID numbers since the last report are as follows:-

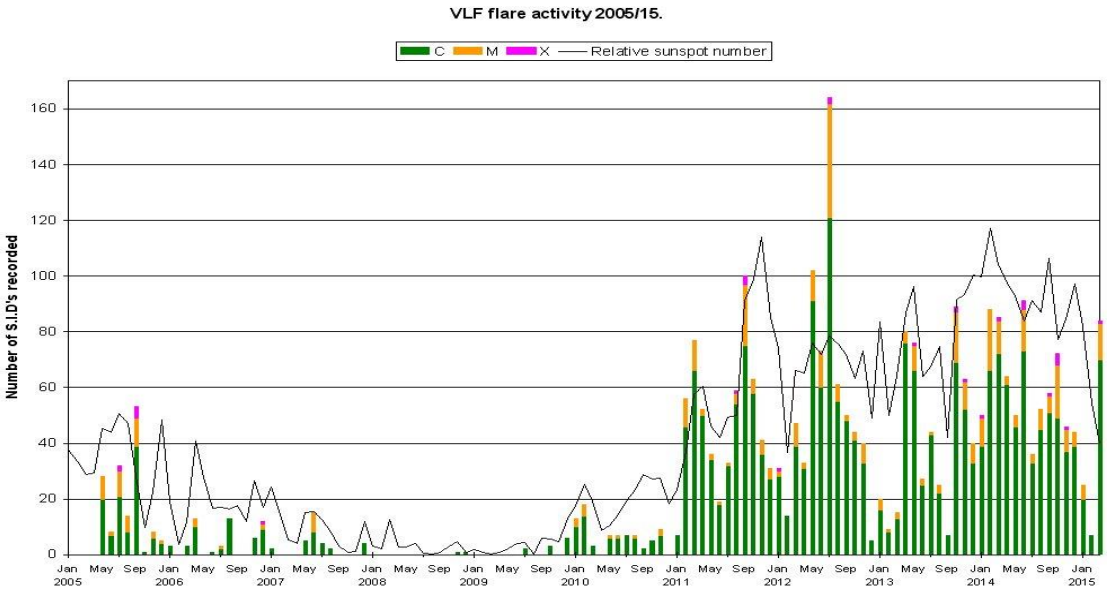

Fig. 1

January = 25 (including 5 M-class)

February = 7 (No M-class)

March = 88 (including 13 M and 1 X-class.)

Very low levels of activity were recorded in January and February, followed by a large increase in March. March also included the partial solar eclipse, widely recorded by group members. While the weather did not co-operate for many visual observers, this was no problem for the radio observations.

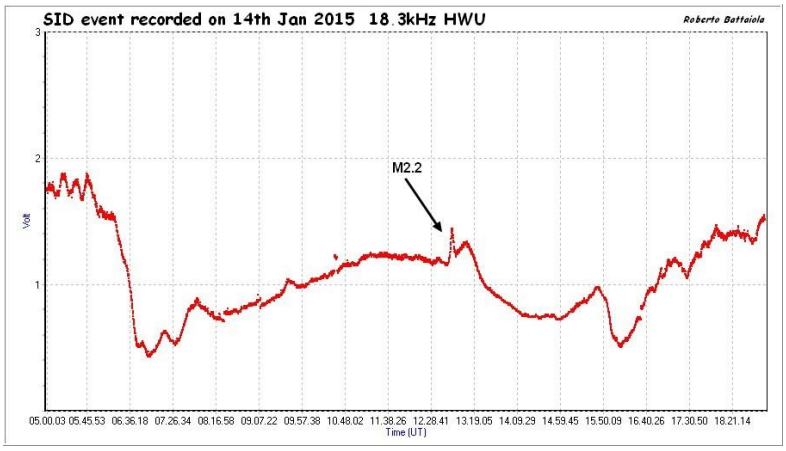

Fig. 2

January 14th included an M2.2 flare producing a SID peaking at 12:58UT, as shown in Fig 2, recorded at 18.3kHz by Roberto Battaiola in Milan. It also had an associated radio burst at 151MHz, shown in Fig 3 by Colin Clements in Lisburn. While VLF propagation had recovered by 13:45, the 151MHz burst lasted until about 14:05. There were no major CME's associated with flares in January, although the effects of a very weak diffuse CME were recorded on the 7th. This appears to have been caused by an eruption partially hidden by a large coronal hole, and not directly related to any flares.

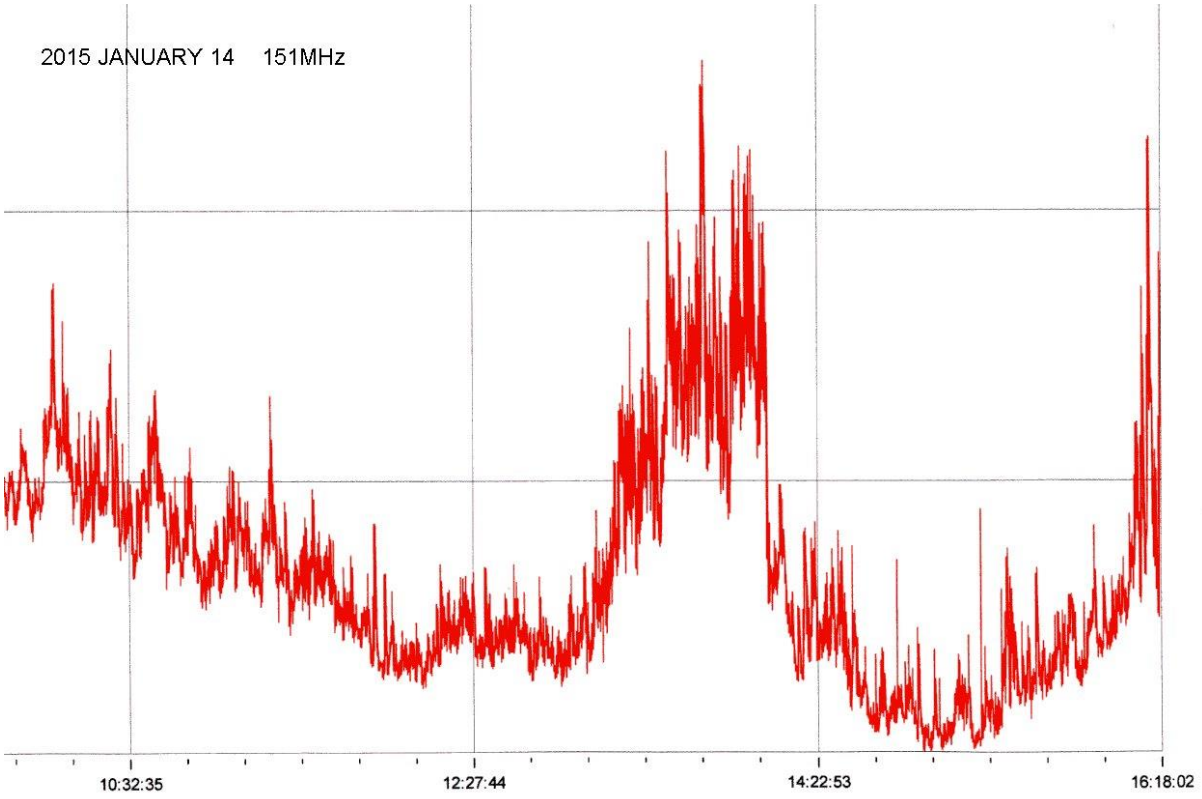

Fig. 3

The SID count in February was the lowest since 2012 December. There were only two M-class flares in the entire month, neither occurring during daylight hours. The background X-ray flux was bellow B9 level all month, dipping below B5 by mid-month. Coronal hole effects dominated the magnetic record, mostly from a long-lived south polar coronal hole. Towards the end of the month an equatorial hole developed, causing some magnetic disturbance on the 23rd and 24th. Fig 4 shows the recording by Roger Blackwell on the Isle of Mull. Note that the magnetometer is reset each day, hence the discontinuity at midnight.

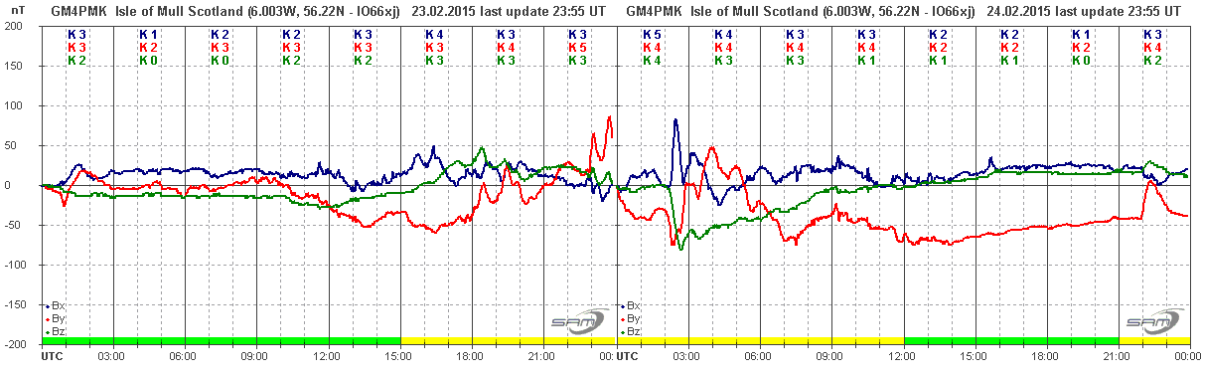

Fig. 4

The major event in March was the solar eclipse on the 20th. While it was partial over the UK, the total phase was visible in the north Atlantic south of Greenland through the arctic to Svalbard and beyond. Maximum phase was at 09:40 in the Faeroes, close to maximum totality. To the D-region of the ionosphere the slow reduction in solar radiation is similar to a sunset, rapidly followed by a sunrise as the sun is again uncovered. Unlike sunset/sunrise, solar radiation remains from above without the illumination from below seen when the sun is on the horizon. Its effect on the various VLF signals is shown in Fig 5, recorded by Mark Edwards. Quite a strong dip in signal level is evident at 22.1kHz (yellow trace) as well as 23.4kHz (brown). Fig 6 shows the 23.4kHz signal recorded by Peter Meadows. The ionosphere reflection point in both cases is well south of the path of totality. A more in-depth analysis of the data received is in progress.

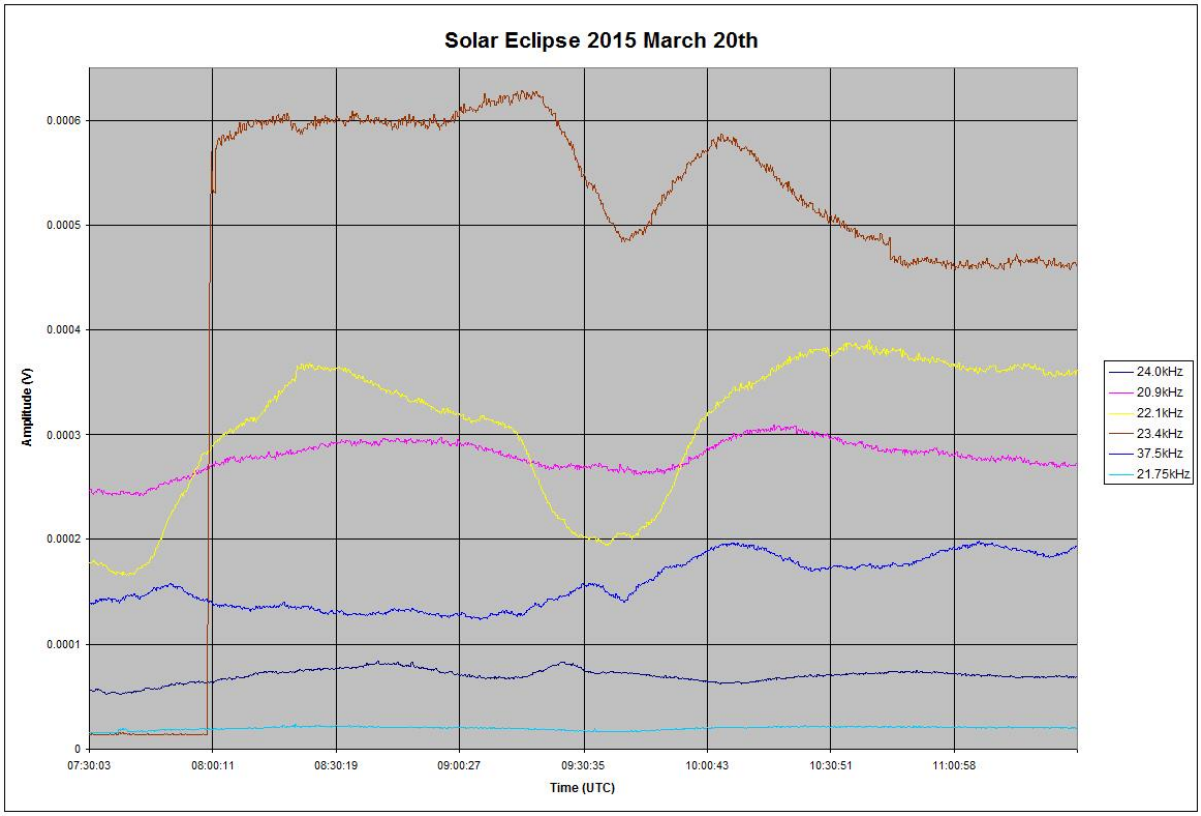

Fig. 5

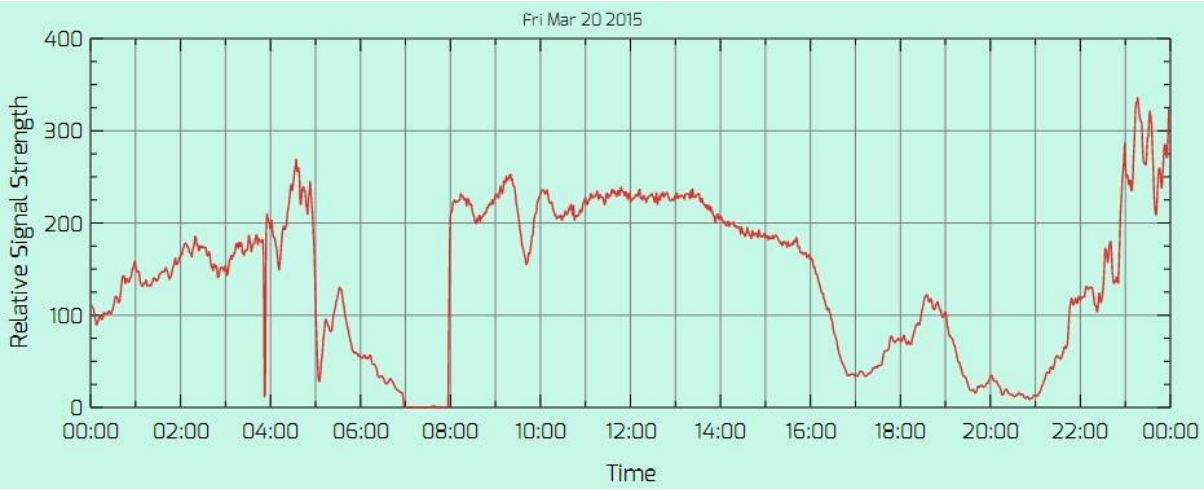

Fig. 6

May 2015

The sun was in a brief quiet period for the eclipse, having been very active indeed earlier in March. As already noted, there was an X2.1 flare late in the afternoon of the 11th. Peaking at around 16:20, the decay of the SID merged with the onset of sunset conditions, and so was difficult to time. Fig 7 shows my own recording from the 11th, showing some of the other activity during the day. The X-ray flux in black is from the GOES15 satellite. The C5.8 flare earlier in the day can be seen to have several sub-peaks in its decay leading to some smaller SIDs seen at 23.4kHz. The X2.1 flare has a very steep rise-time, and has caused a spike-and-wave style SID where the sky-wave has gone through over 180 degrees phase change. Right at the peak of the SID can be seen a very small SFE in the magnetometer trace (green). This flare had an associated CME, although it was not Earth-directed and had only a minor effect on the magnetosphere.

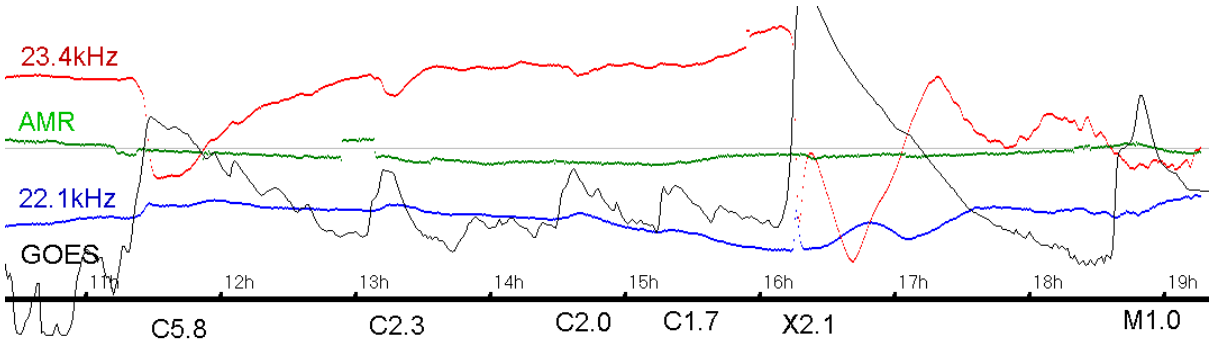

Fig. 7

The BIG magnetic storm on the 17th was due to the combined effects of a filament eruption and some smaller flares early on the 15th. Fig 8 shows the magnetic storm recorded by Colin Clements. The SSC is timed at 04:45UT, and produced a very sharp spike of about +/-40nT. The storm really got underway in mid afternoon, and lasted into the morning of the 19th. My own magnetometer recorded a peak deviation of 213nT just before midnight on the 17th.

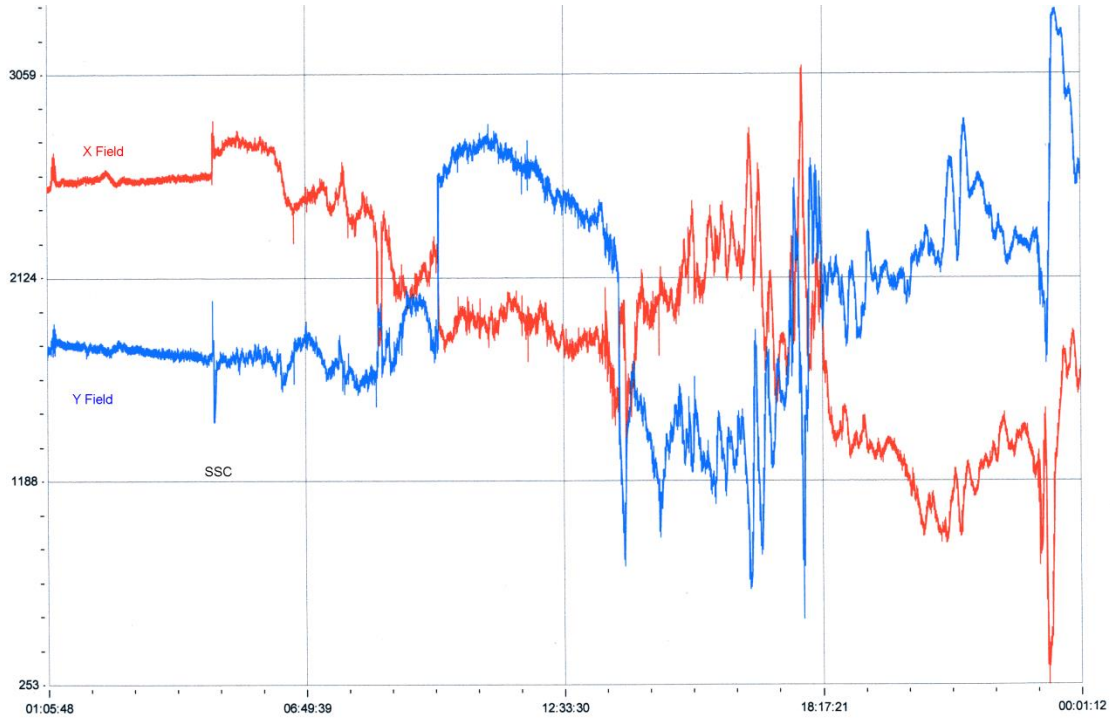

Fig.8

May **2015**

A second SFE was recorded at 14:30 on March 9th. Associated with a fast M4.5 flare, it does not seem to have caused a CME, and no further magnetic disturbance was reported. Coronal hole effects were responsible for most of the remaining activity shown in the Bartels diagram, Fig 9.

Observers: Roberto Battaiola, Roger Blackwell, Colin Clements, Mark Edwards, John Elliott, Paul Hyde, Peter Meadows, Steve Parkinson, Gonzalo Vargas, John Wardle, and John Cook.

My thanks to all contributors. If you would like to add your own observations, please contact jacook@jacook.plus.com.

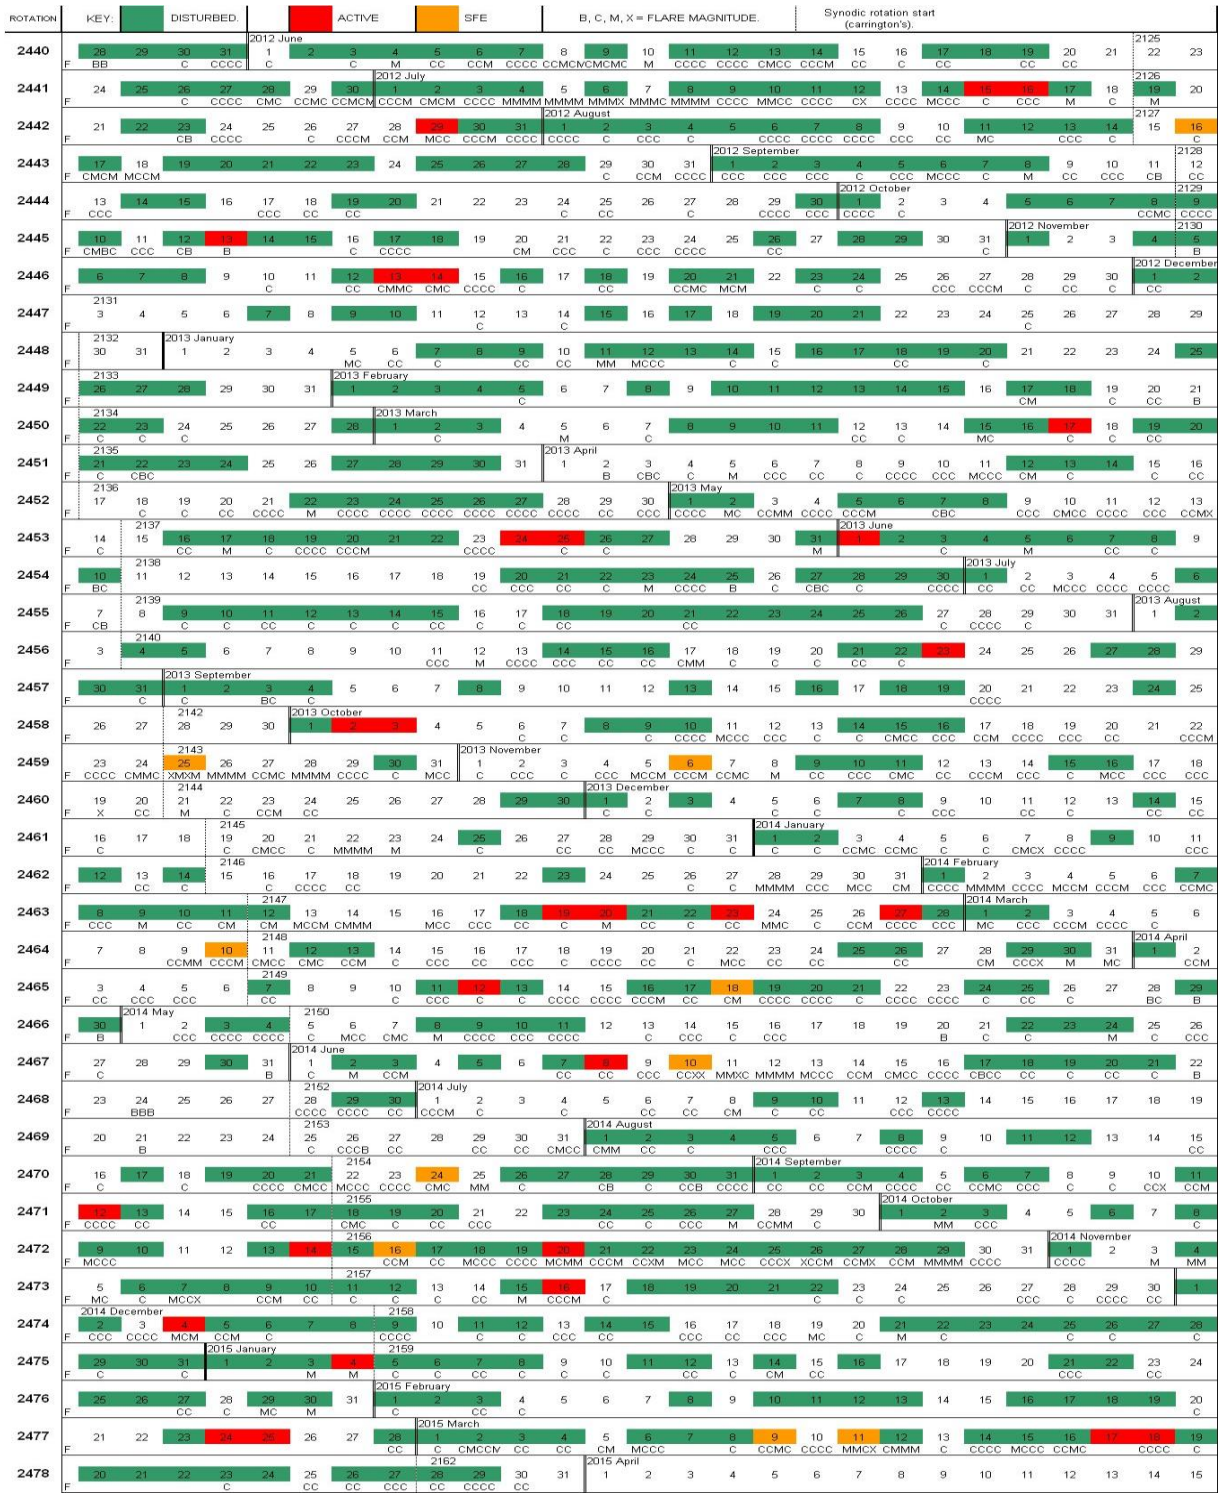

Solar Eclipse 20th March, 2015.

By John A. McKay

A Solar eclipse occurred on 20th March, 2015 with an approximately 90 per cent coverage of the solar disc at visible wavelengths from the location of the 3peaks amateur radio telescope.

The 3peaks amateur radio telescope was configured to track the Sun and to record eclipse data to determine characteristics of Solar radio emissions at a wavelength of 21 centimetres.

It was found that the observed Solar diameter was approximately 1.773 degrees and that the Solar corona became visible at maximum coverage.

Recording data began on the 20th March, 2015 at 07:39 and ended at 11:30.

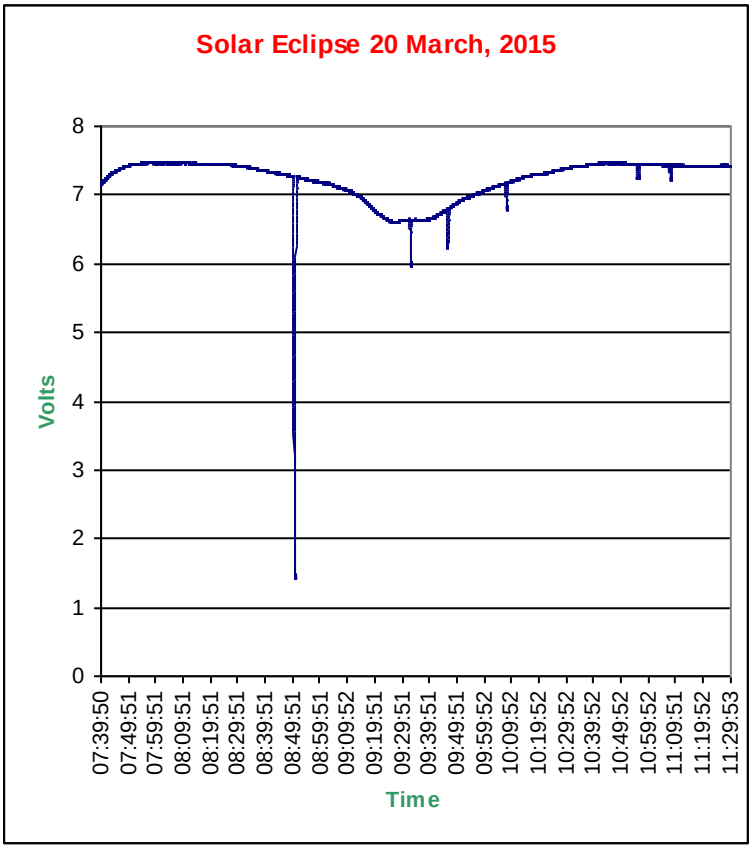

Fig.1

Fig. 1 is a plot of the raw data, volts against time. The dropouts are a result of, as yet unknown, errors in the dish tracking algorithms which causes the dish to "go walkabout" at times but generally the tracking is very good and remains accurate to within a few arc minutes.

VLF observations during the 2015 Total Eclipse

Alan Melia G3NYK

The opportunity to observe the ionosphere during a close total eclipse of the sun is too good to miss for anyone interested in radio propagation. The chemistry of the ionosphere is driven almost completely by the Sun. The altitude of the D-region is slightly difficult for a lot of formal observations: it is too low for rockets and too high for balloons. Thus a lot of the knowledge has been gained by observing its effects on radio signals.

It was quoted by the author of the 1999 eclipse report that an eclipse is similar to the sunrise/sunset effect. At sunset the daytime D-region decays due to recombination rates exceeding the ionisation rate of the setting Sun, with the reverse happening at sunrise. However in the first minutes of sunrise or the last minutes of sunset the D-region is also illuminated by rays which graze the ground before hitting the D-region from underneath. It is believed this produces a strongly absorbing region that virtually destroys the sky-wave. Thus at the bottom of the "dips" seen at these times the signal received is mainly due to the ground wave. I have monitored the signal from HGA near Budapest on 136 kHz for several days, and I find that, whilst I get slightly varying daytime levels at noon, the level at the bottom of the morning dip is very consistent and is very close to that predicted by the ITU recommended model for calculating ground-wave signal strength.

In an eclipse the mechanism is slightly different. The ionosphere is only being illuminated from above, then the sunlight is obscured by the moon, just as though the illumination had been faded down to dark and then up again. Ideally we would like to know what happens in the D-region of the ionosphere when this occurs.

The possible results we could aim to achieve are: the altitude of the "apparent reflection height"; the amount this moves upwards during the eclipse; the time delay between the onset of darkness and the disappearance of the ionisation. Then we could look at how the percentage of darkness on the penumbral region affects the change in signal level. Knowing the change in height we may be able to allow for the phase change in the signal strength observations and thus use lots of the data collected by the observers outside the main shadow of totality.

For the 2015 March eclipse, 17 observers monitored up to 17 VLF services, of which six were off-air during the period of the eclipse. Figure 1 gives the approximate position of the location of both observers and transmitters to show the range of signal paths covered. All of the paths fall within the 70% eclipse profile, but unfortunately only one transmitter (Grindavik, Iceland) lies north of the track of totality.

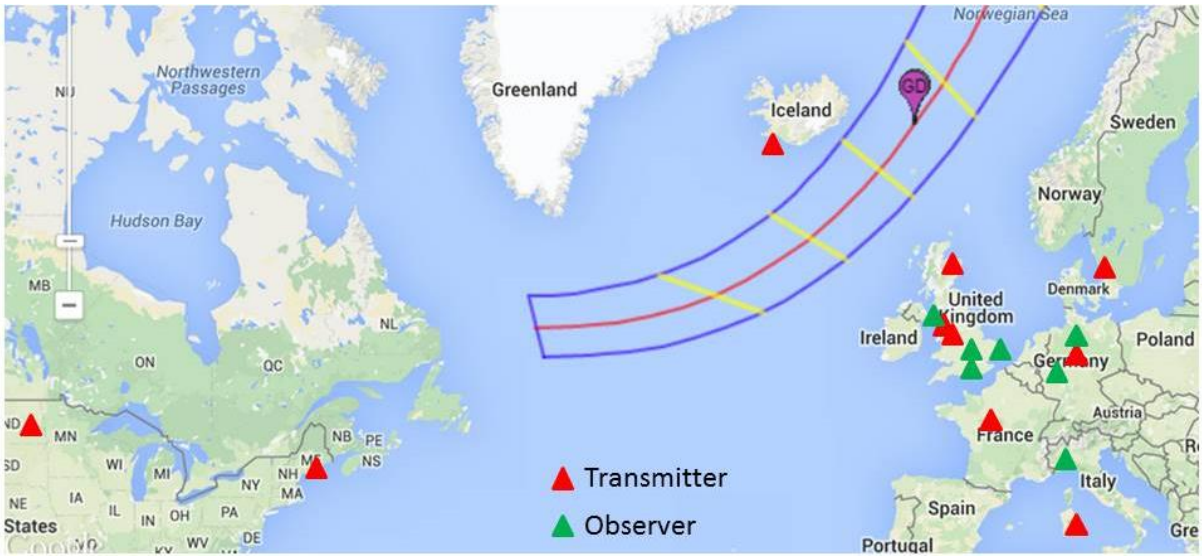

Figure 1 Approximate locations of Observers and the VLF transmitters used in the study.

Map data ©2015 Google, ENGI. Eclipse Predictions by Fred Espenak, NASA's GSFC

A total of 63 sets of observations were provided for the analysis, each covering between 0800 and 1200 UT. Many of these were amateur radio astronomers who monitor VLF station signal amplitude as a method of detecting Solar X-ray flares via Sudden Ionospheric Disturbance events (SIDs). These observers have provided data on the change in received signal level and I have also acquired data from two stations which have slightly more sophisticated equipment able to monitor the received signal phase.

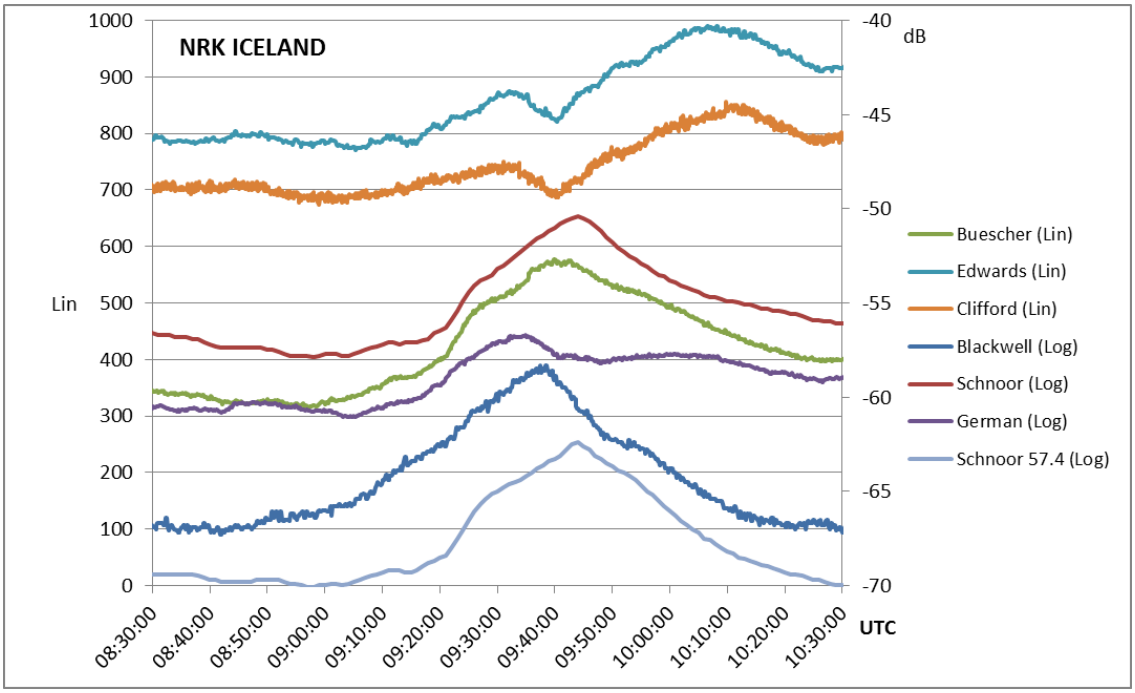

Fig. 2a

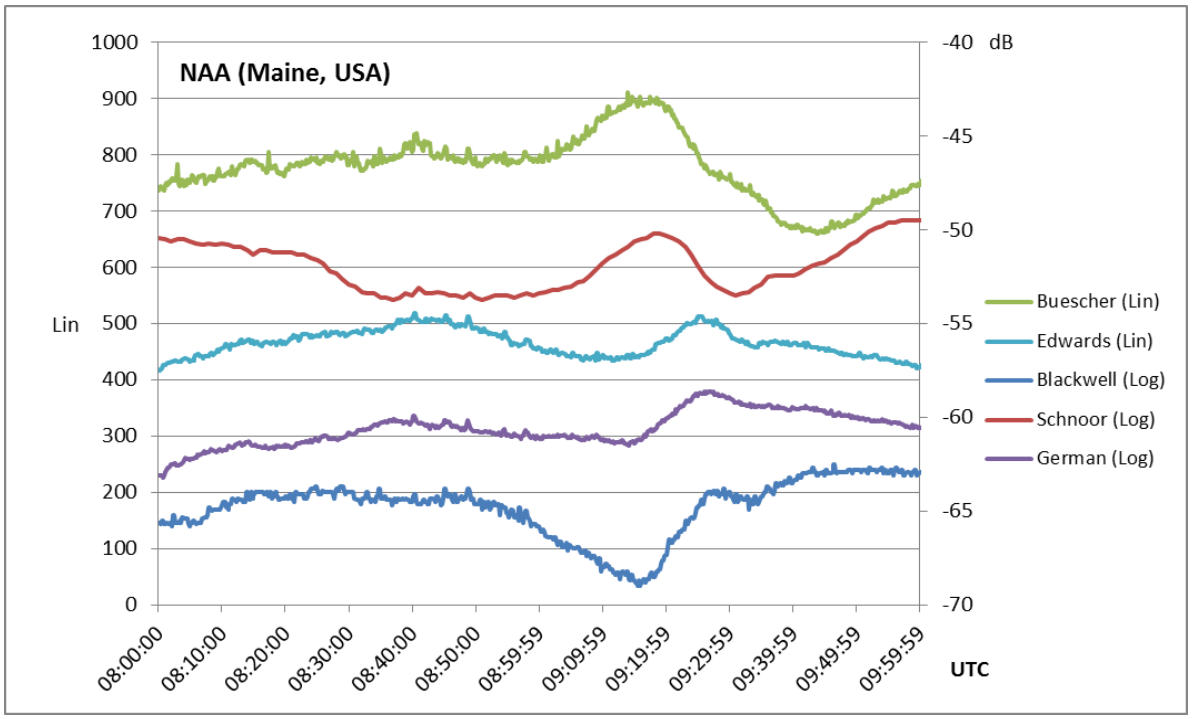
NAA (Maine, USA)
Lin
dB
UTC
Buescher (Lin)
Edwards (Lin)
Blackwell (Log)
Schnoor (Log)
German (Log)

Fig. 2b

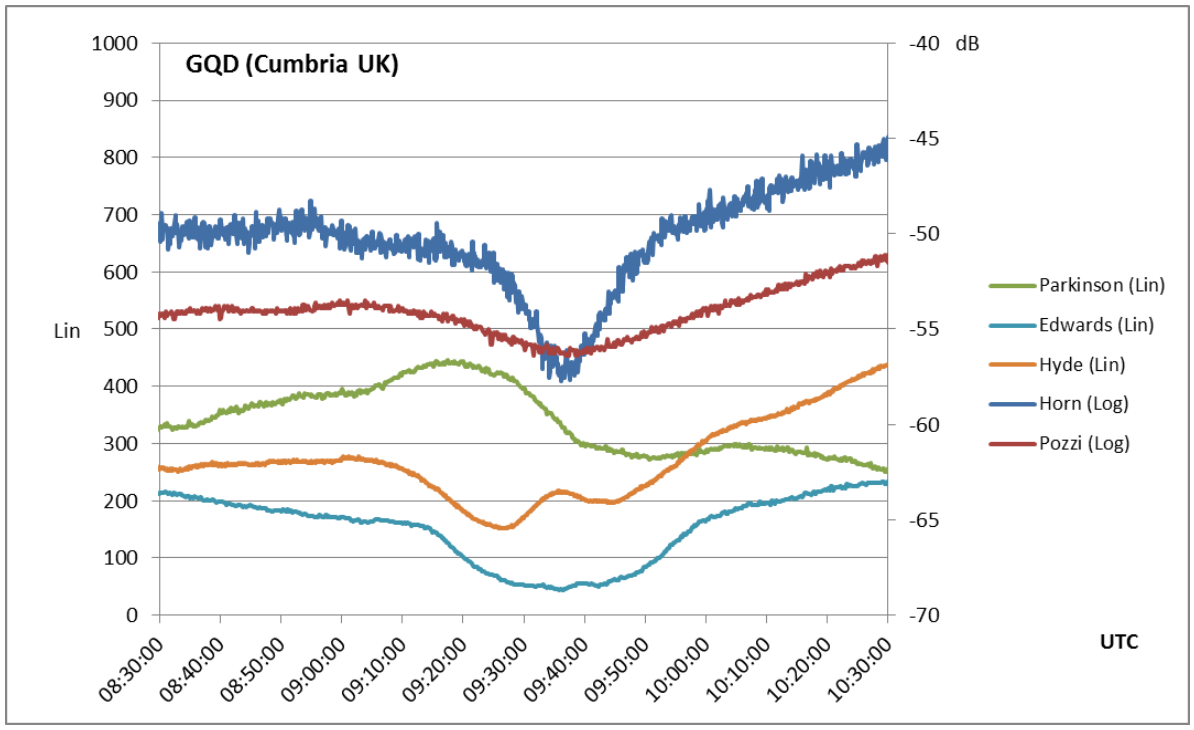
GQD (Cumbria UK)
Lin
dB
UTC
Parkinson (Lin)
Edwards (Lin)
Hyde (Lin)
Horn (Log)
Pozzi (Log)

Fig. 2c

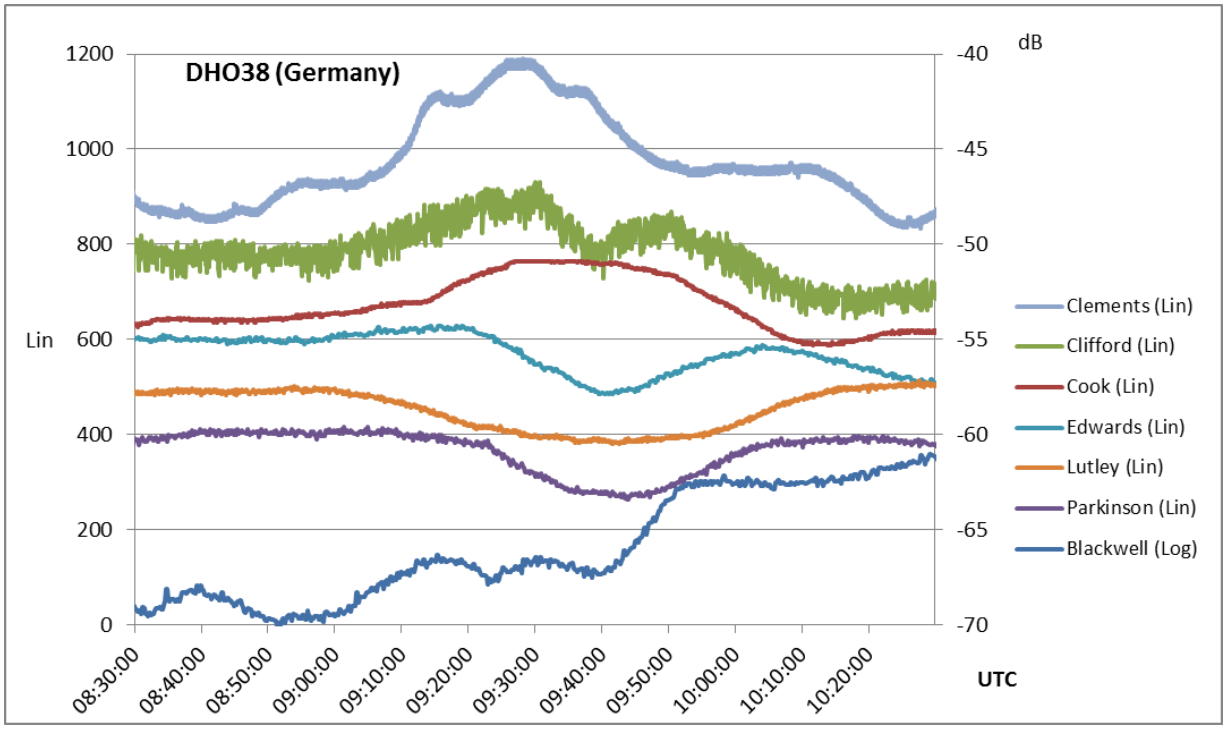

Fig. 2d

Figure 2a Received signal strengths during the eclipse for NRK (37.5 kHz)

Figure 2b NAA (24.0 kHz)

Figure 2c GQD (22.1 kHz)

and Fig 2d DHO39 (23.4 kHz)

Note that some traces are against the right hand logarithmic scale, and that traces have been offset vertically to aid comparison.

Initial analysis

Whilst changes in the ionosphere, in particular the D-region, will cause changes in the amplitude of the received signal the effect is complicated by changes in absorption as well as changes in path length. Changes in absorption below the apparent reflection height cause changes in the strength of the ionospheric wave component. The resultant amplitude seen by the receiver is due to the addition of signals over the two paths, where the phase of the ionospheric wave may cause it to reinforce or to partially cancel that due to the ground-wave. The observation of the received phase is independent of any absorption effects and provides a direct measure of changes in the ionospheric wave path length.

The first analysis covers the height change using observations of phase change made by the two stations able to measure this. These observed the Icelandic military transmitter NRK at Grindavik on the south-west corner of the island, at a great circle distance of 1600km. Although their paths are slightly too far

south for the "reflection" to be in the middle of the predicted path of the totality, we must remember we are interested in the illumination at an altitude of 50km. With the sun at about 30 degrees elevation the shadow at 50km will be around 70 to 100 km south of the ground shadow, which just touched the Faroes.

The data I currently have suggests that the "apparent reflection height" changed by rising by about 5 to 6 km during the event. The paths are not sufficiently different to resolve the ambiguity as the phase rotates through 360 degrees. The actual value depends also on the starting value. The actual phase change recorded was 80 degrees by one station and a similar amount (though contaminated by some observing effect) at the other site. This is probably accurate enough to start looking at amplitude data and making some estimations of how much of the amplitude change was due to the phase shift.

This initial analysis suggests that the apparent reflection height increased by about 5.5 to 6 km. Further work is needed to refine this. It is hoped that this work will yield some information on the ionisation decay times and the time to re-ionise, but there is some discussion on the accuracy of timestamps from some observers which has not yet been resolved.

My thanks go to the various observers who have made data available for this work: Roger Blackwell (UK), Wolfgang Buescher (DE), Phil Busby (UK), Colin Clements (UK), Chris Clifford (UK), John Cook (UK), Mark Edwards (UK), Mike German (UK), Mark Horn (UK), Paul Hyde (UK), Terry Jeacock (UK), Bob Kerswill (UK), Andrew Lutley (UK), Paul Nicholson (UK), Steve Parkinson (UK), Claudio Pozzi (IT), Peter Schnoor (DE), and Ian Williams (UK)

Special thanks should also go to Paul Hyde for his work in consolidating the submitted data prior to further analysis. The resulting spreadsheets are available to anyone wishing to do their own analysis work using the results.

Comparing the RTL2832 and FCD Pro Plus receivers for meteor scatter applications

Paul Hyde

Meteor scatter using the GRAVES space surveillance radar system as the illuminating source provides a simple way into radio astronomy. The reflected signal levels are comparatively high, making local noise less of a problem, whilst the frequency of 143.05 MHz makes for a relatively small antenna. Combining the FUNcube Dongle Pro Plus receiver with Spectrum Lab software provides a simple way of displaying meteor echoes and of automatically detecting and recording individual events.

Whilst the FCDPP is much cheaper than a communications receiver, it is still £150 (inc. p&p) which is a disincentive for the beginner. By comparison, the RTL2832-based TV Dongles can be purchased for less than £20, in some cases substantially less, though there is a risk of ending up with a sub-standard device if you go for the cheapest in the market. The performance of these devices has been questioned as they do not include the front-end filtering or the additional preamplifier stage of the FCDPP. This makes them less sensitive and prone to overloading from strong signals in adjacent bands, in particular the commonplace Pager transmitters at 138 and 153 MHz. In addition, the RTL2832 devices are limited to 8-bit operation, reducing the dynamic range. Finally, there is no simple way of using them with Spectrum Lab, unlike the 'Plug and Play' nature of the FCDPP. Spectrum Lab is currently the only application that has the functionality needed for detecting and logging meteor events automatically.

Nevertheless, the low cost of the RTL2832 devices makes them very attractive and David Morgan has produced a paper showing how they can be used with Spectrum Lab in conjunction with the SDR# application.[1] Here I try to build on this work by comparing the resulting performance of the RTL2832 and FCDPP. In fact, I was pleasantly surprised at just how good these cheap devices can be.

Installing the RTL2832

The first step was to install the SDR# application and get the RTL2832 receiver working with the Zadig drivers. I used the 'Quick-Start' installation procedure detailed on the RTL-SDR website.[2]

[1] www.britastro.org/radio/downloads/TECHNIQUES_FOR_USING_THE_RTL_Dongle.pdf

[2] See www.rtl-sdr.com/rtl-sdr-quick-start-guide/

NB. Several problems have been noted with conflicts between the Zadig driver and the FUNcube Dongle receiver.[3] *The FCDPP was designed to present itself as a standard external audio device and hence does not need any special drivers. Zadig is ONLY for use with the RTL device and you must avoid associating it with the FCDPP.*

It is essential to be sure of your frequency calibration before searching for meteor pings and both receivers exhibit their own manufacturing offsets. The RTL2832 is significantly poorer in this regard and at initial install mine was nearly 10 kHz out at the GRAVES frequency, compared to only 470 Hz for the FCDPP. You can correct this offset within SDR# using the Configuration screen (under the 'cogwheel' icon) but finding a suitable reference frequency is not easy. The ideal solution is to use one of the amateur radio beacons such as GB3VHF, but these do not provide national coverage. Another possibility is one of the VOR aircraft navigation beacons located around the country with frequencies between 112 and 118 MHz. Information on these can be found on the internet[4] and you are likely to find two or three that show a signal when tuned to the listed frequency. These beacons are a useful resource and more information can be found in another paper by David Morgan.[5]

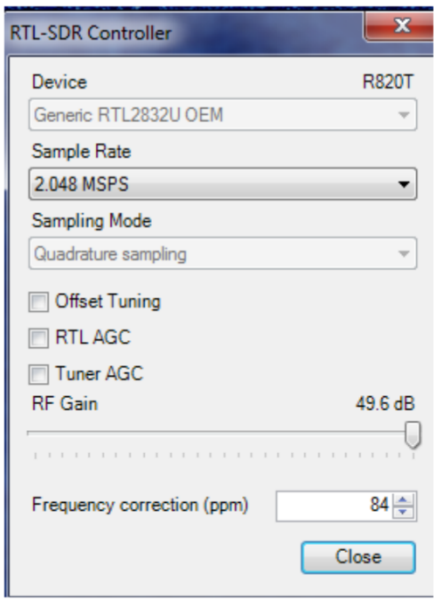

The SDR# Configuration Screen for setting the gain of the RTL2832 TV Dongle and compensating for the frequency offset. The correction should be increased until the frequency displayed for the beacon is as close as possible to the published figure.

I then installed the VB Cable application[6] and selected the (MME) CABLE Input (VB-Audio Virtual C) option from the SDR# Audio Output drop-down list. You may need to restart SDR# before this option becomes visible in the selection list.

Finally I started Spectrum Lab and in the Audio I/O screen set the Input Device as CABLE Output (VB-Audio Virtual C)

[3] www.funcubedongle.com/?p=1425
[4] See the global map at http://ourairports.com/big-map.html#lat=0,lon=0,zoom=2,type=Hybrid and check the Navaids box to show the VOR beacons near you. Select any within 50 miles and note their frequencies.
[5] www.britastro.org/radio/downloads/A-PRELIMINARY-NOTE-ON-DETECTION-OF-AIRCRAFT-VOR-NAVIGATION-BEACONS.pdf
[6] http://vb-audio.pagesperso-orange.fr/Cable/index.htm

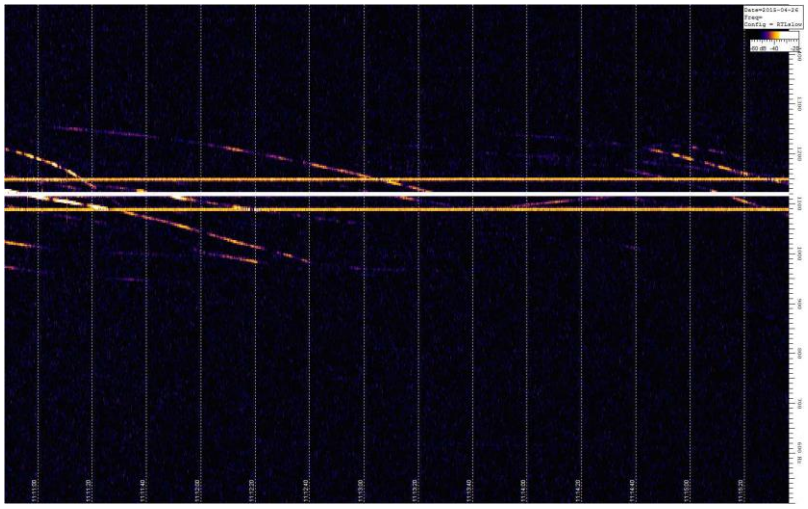

If the system is set up correctly you should see something like this, showing the beacon and its two 30 Hz sidebands running horizontally across the screen. The other traces are the beacon signal after reflection from moving aircraft.

Intermodulation performance

My first, and greatest, surprise was the absence of noticeable intermodulation products. I have a 138/153 MHz Pager site less than a mile away and on the same bearing as GRAVES. This created severe problems with the original FUNcube Dongle receiver and I had to add in a filter at the antenna to make the system work. The problem was solved with the new FCD Pro Plus which includes its own filter network. I was expecting to encounter the same problems with the RTL2832 device, but so far I have not seen any sign of breakthrough even with the RF Gain set at maximum. I'm sure that problems will occur when using these devices for weak signal applications, but I have not found them to be a problem for meteor scatter work.

Catching meteors

I used a cheap CATV splitter to run both FCDPP and RT 2832 receivers from a common antenna. This makes for a mismatch of 50 and 75 ohm impedances but the losses will be insignificant.

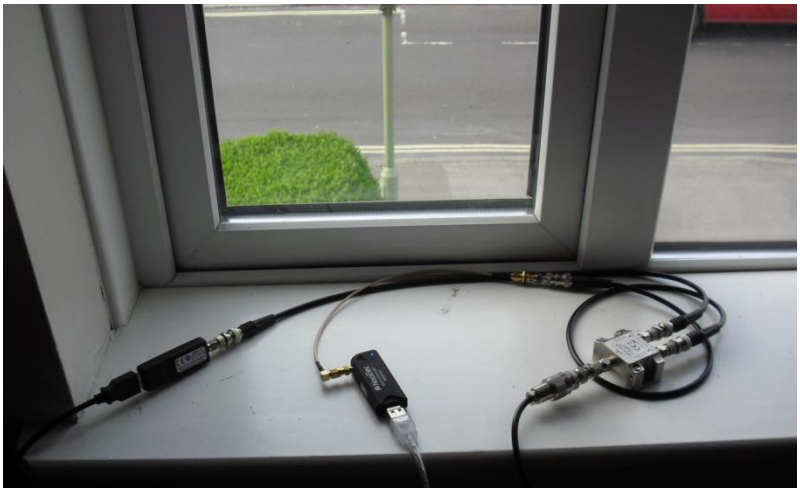

I then set both systems to a high screen scroll rate to capture the detail and obtained the following comparative images, from the FCDPP system on the left and the RTL2832 system on the right..

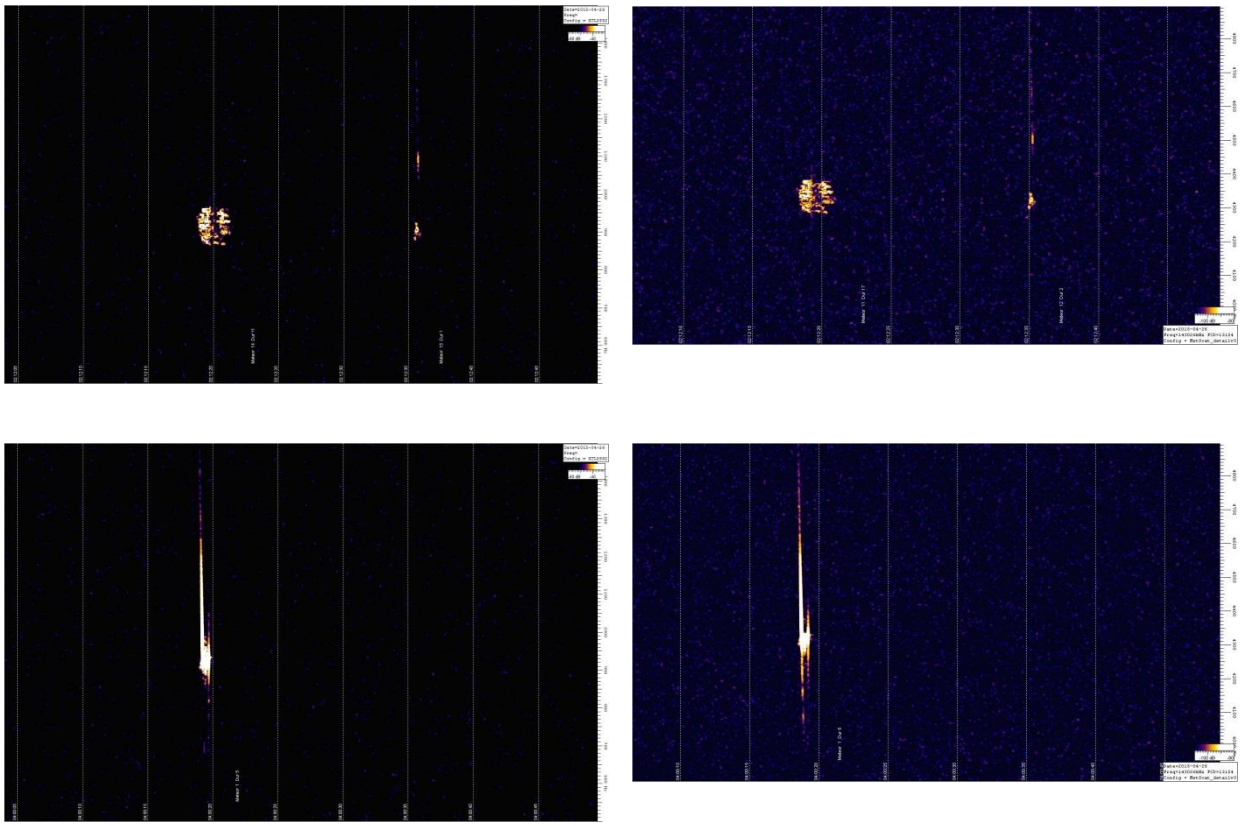

I did find a difference between the RTL2832 and FCD systems in the number of meteor events recorded in the daily log. This seems to be down to the setting of the RF Gain of the RTL2832 receiver but further work is needed to understand what's going on. However, provided you do not change the gain setting the activity profile should remain comparable with previous days.

Frequency stability

This is an area where the performance of the RTL2832 is poorer than that of the FCDPP but not to the extent of making it unsuitable for meteor scatter work. I used a Conditional Actions script running in Spectrum Lab to log the frequency and amplitude of a strong VOR beacon over a day.

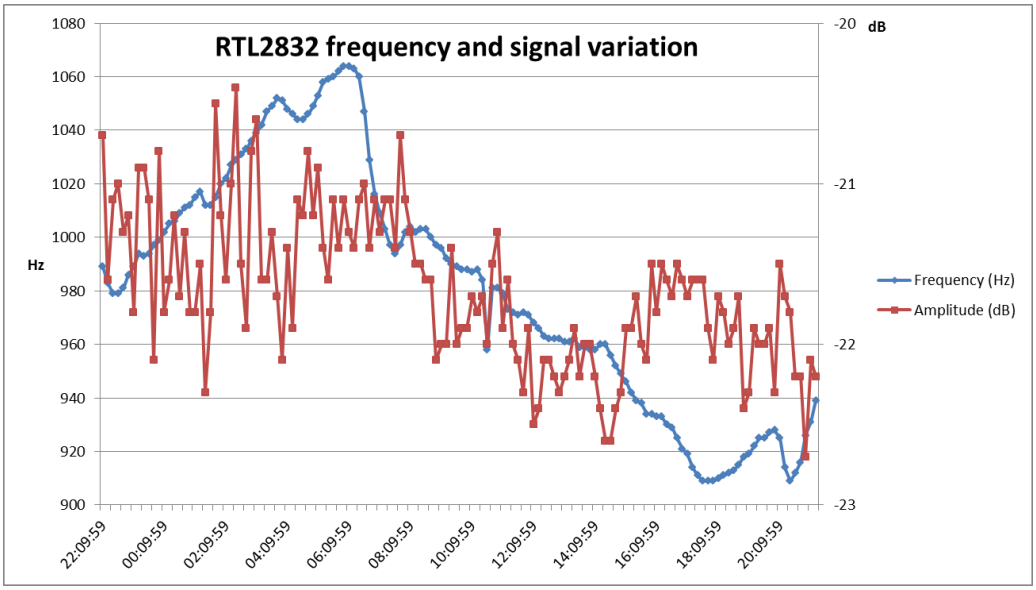

I measured around 160 Hz of frequency drift over a day, compared to around 10 Hz for the FCDPP. The amount of drift poses a problem if you are trying to use Doppler shift to measure the Line of sight velocity of the meteor trail but it does not affect the ability to capture timestamps and durations of individual events.

Conclusions

The RTL2832 provides acceptable performance for meteor scatter work at a fraction of the cost of the FUNcube Dongle. This makes it possible to build a meteor scatter system for around £50, not including the computer and display. This makes it an ideal entry-level project for those wanting to use radio techniques in astronomy.

Using SDR# and VB Cable to interface between the RTL2832 and Spectrum Lab is wasteful of computer CPU resource. It may be possible to avoid this by using the new ExtIO_USRP package from Balint Seeber[7] which is the next step to explore!

Paul Hyde

Improving Frequency Stability of the RTL2832U SDR

for dual-Channel H-line Receivers or Interferometry

Peter East

Abstract

The RTL2832U DVB-T dongle used as a cheap software defined radio (SDR) has found many applications in radio astronomy. Whilst the basic stability performance is limited, it has mostly been adequate used with available GUI software. Some novel progress has been made in frequency locking pairs of dongles, offering the possibility of tracking coherency good enough for bistatic radar and interferometry.

In this note, cheap and simple methods of improving the short and long-term frequency stability of the RTL SDR and aligning data clocks are explored.

Introduction

The RTL2832U SDR has low thermal mass so the operating temperature and the quartz crystal's frequency stability is at the mercy of the ambient temperature and the effects of internal power dissipation. Typical warm-up temperature variation is of the order 1 part in 10^5 but this is usually masked by the frequency calibration correction of up to 1 part in 10^4. Calibration offsets can of course

[7] See www.qsl.net/dl4yhf/speclab/settings.htm#RTL_SDR

be measured and corrected. Once the device has reached thermal equilibrium, however, frequency stability can improve to 1 part in 10^6.

Juha Vierinen has posted a method of running at least three RTL units from the quartz crystal of one SDR, so locking them in frequency. The procedure involves removing crystals from slave dongles and linking to the crystal buffer amplifier inputs as described in his link.... http://kaira.sgo.fi/2013/09/16-dual-channel-coherent-digital.html

The basic frequency stability of the RTL2832U does not appear good enough for coherent applications, but this note describes a simple method of improving stability, linking a pair of SDRs and assesses their performance for application to a twin channel H-line receivers or interferometers. For this application, it is important that the data from a pair of SDRs is also accurately aligned in time and this problem is also addressed.

Basic Performance

Measurements on a standard RTL SDR show that in the USB plugged-in quiescent state, power dissipated is approximately 0.4W.

When operating, collecting data or visible with SDR# software, the dissipated/heating power increases to about 1.3W. At 1420MHz, from operating switch-on, the tuning of a standard unit drifts several kHz before reaching thermal equilibrium after maybe 10 to 15 minutes.

Longer term stability then appears to fluctuate by a few hundred Hertz.

Operated with SDR# software, this performance when thermally stable, is adequate for most applications and any calibration error can be corrected by suitable offsets.

However, when the SDR is used intermittently and data collected with rtl_sdr software tools, the few kHz tuning drift over the observation period can colour Hydrogen line spectrum analysis detail for example.

Not all RTL2832U dongles are equal. More than half the dongles bought from various internet sources exhibit a 30dB gain reduction at frequencies above about 1380MHz after several minutes of continuous operation. This appears to be a thermal effect and seems to be cured by keeping them cooled.

Tests running two dongles on one or two PC/laptops using the rtl_sdr software tool tasking two dongles simultaneously has shown a very wide variation between actual dongle start-up times of up to several tens of milliseconds; almost certainly due to operating system preferential housekeeping tasks.

Improving Frequency Stability

The method adopted here takes advantage of the low SDR thermal mass, by blowing copious ambient air across bare SDR pec's to extract dissipated heat above ambient, so minimising the SDR quartz crystal temperature rise. To improve the cooling efficiency mini heatsinks were thermo-plastered onto the integrated circuits and crystal. With these modifications, thermal equilibrium is achieved more quickly, now in a few tens of seconds. Figure 1 shows some switch-on temperature drift comparisons.

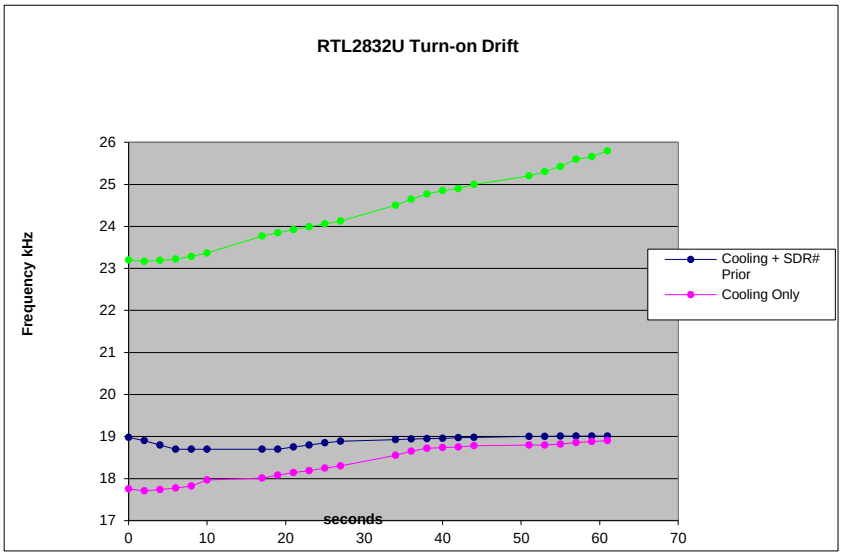

Figure 1 RTL SDR Turn-on Tuning Drift

Referring to Figure 1, the green curve shows the basic switch-on drift of a standard, unmodified unit with time as the dissipated power warms it up. The drift rate is approximately 3kHz in 60 seconds. It is clear from this curve that thermal equilibrium is not yet reached. It is obvious that were data collected in this period the source signal spectrum would be spread accordingly.

The magenta curve shows the frequency drift with fan cooling. Already it is observed that the initial frequency offset, due to the quiescent 0.4W dissipation (~5kHz) has been compensated. Also under full working power, thermal equilibrium is approached after about 40 seconds. The overall drift in 60 seconds is reduced to about a third of the uncooled case.

The blue curve results from preheating the SDR using SDR# to run the dongle prior to switching to rtl_sdr software to collect the data. In this case the drift deviation is less than 0.2kHz and thermal equilibrium is reached within 30 seconds. Once thermal equilibrium is reached, indicated frequency variation of less than 10Hz is observed. Separating this data in time produces data with an indicated stability of less than 1 part in 10^8.

Calibration offsets still need correction however, and will need adjusting for the current ambient temperature.

Mechanical Construction

Figure 2 shows the mechanical layout. The two bare dongles are aligned side-by-side and mounted in a folded aluminium box with two 5V, 0.6W fans, one blowing in and the rear fan sucking out. These are driven by a 5V power supply unit. The dongles, of course, being powered from the laptop USB ports.

150Ω chip resistors were soldered across the RTL RF inputs to reduce the dongle nominal 75Ω input impedance to match the 50Ω rf amplifier/filter characteristic impedance.

Twisted pair rather than coax is used for interconnecting the crystal inputs to keep the connection convenient and short.

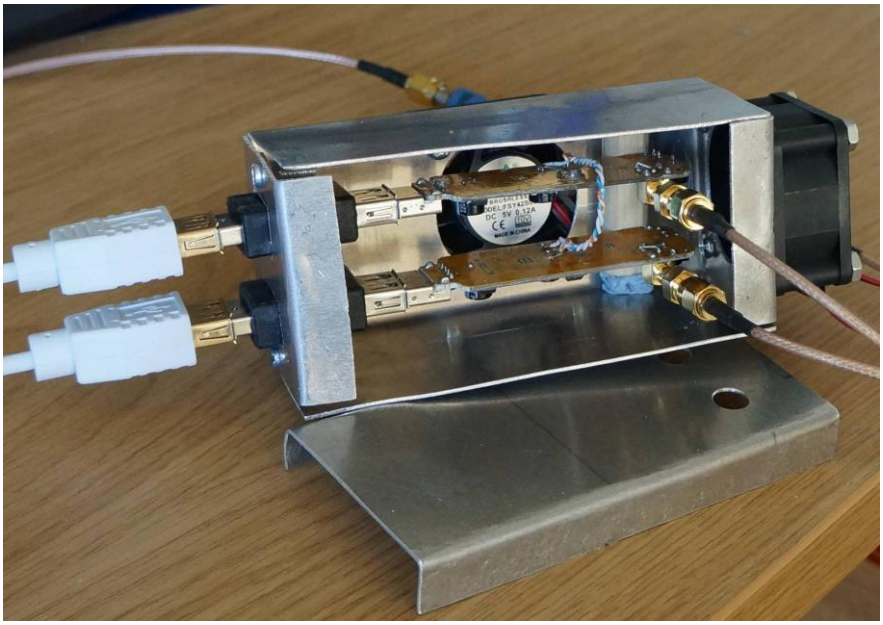

Figure 2 Duo-SDR Cooling Construction

Operational Tests

Figure 3 shows the test set-up. A stable 1420MHz pure-tone signal is fed to two RTL SDRs via an RF power splitter. The two SDR USB outputs are fed to two laptops which are data-synchronised by the data capture software running the rtl_sdr tool. True synchronism depends on the laptop clocks running in synchronism. This is optimised in linux using the 'chrony' system tool and in Windows by synchronising both laptops to the same internet time server. Experience has shown that different laptop clocks do not stay in synchronism for long, so time synchronising should always be carried out prior to serious measurement.

Data collected is sectioned and sections analysed for indicated frequency, phase tracking and relative delay.

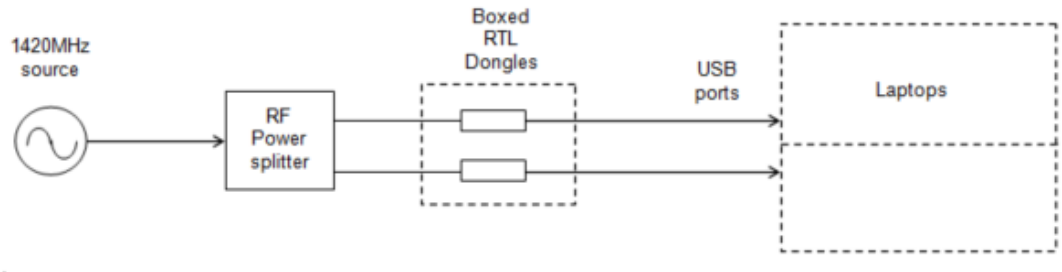

Figure 3 Test Schematic

Although it has shown to be feasible to lock two dongles in frequency, it is unlikely that the data clocks will be aligned. For perfect frequency alignment, clock misalignment will be exhibited as a constant phase shift.

Correlation of data sets has shown that laptop data delays, even with time synchronised laptops can be a few tens of milliseconds. These delay differences can go either way and are probably due to operating system housekeeping. Both Windows and Linux exhibit time variations.

Correlation delays are determined using a wideband noise source replacing the oscillator shown in Figure 2 and special purpose software (see Appendix). The injected noise power is switched on for a short period and should be sufficient to dominate over system noise.

Tracking accuracy better than half the clock period is observed.

Figure 4 shows an implementation for obtaining both calibration and observation data.

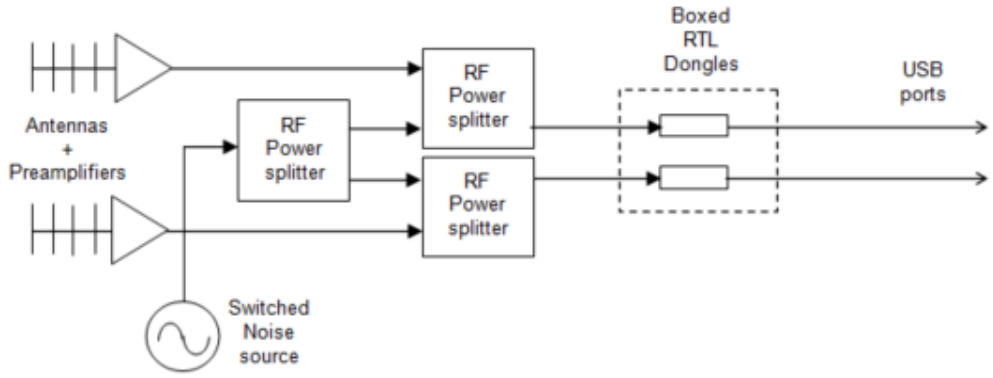

Figure 4 Time Tracking Primer

Simultaneous data are recorded, first with the noise source switched on at a level which swamps the antenna signal for a second or so before the wanted signal data.

Data files are then processed in software suitable for separating calibration and source data, determining the timing alignment, and then aligning source data records prior to examining the source signals for interferometer phase variation, for example.

Time Correlation Alignment

As well as the operation system timing uncertainty, due to the separate dongle clocks, the antenna waveforms are not sampled at precisely the same time, so after system delay compensation, there will be a systematic phase error in the two data records. It is possible to estimate this phase error from the correlation results, but as it is initially unknown, the traditional complex correlation formulae cannot be applied.

However, the signal amplitudes in the two channels do still correlate in time and by allowing for the mean amplitude offsets, the following simple amplitude correlation procedure can be adopted.

If, datA and datB are the two data series derived using the Osmocom rtl_sdr, .bin files, the sample amplitudes are derived from (sequential real and imaginary components); for n = 0,1.....etc:

$$ampA(n) = \sqrt{(datA(2n))^2 + (datA(2n+1))^2}$$
$$ampB(n) = \sqrt{(datB(2n))^2 + (datB(2n+1))^2}$$

The correlation process applies the following formula, which produces a single peak only when the two records align as shown in Figure 5

$$Corr(s) = \frac{4}{N}\sum_{n=0}^{N}\frac{\left(ampA(n) - \overline{ampA}\right)\left(ampB(n+s) - \overline{ampB}\right)}{\overline{ampA}.\overline{ampB}}$$

where,

$$\overline{ampA} = \frac{1}{N}\sum_{n=0}^{N} ampA(n) \quad \text{and} \quad \overline{ampB} = \frac{1}{N}\sum_{n=0}^{N} ampB(n)$$

are the group sample amplitude averages.

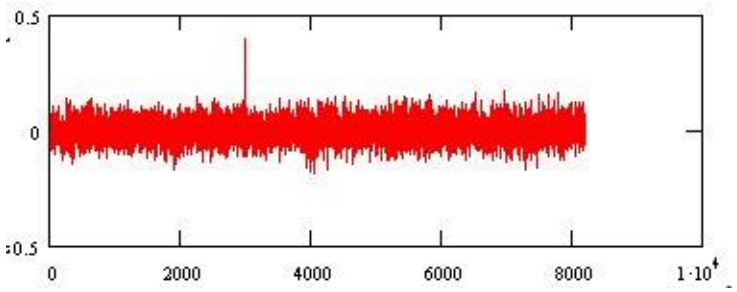

Figure 5 Correlation Result

Figure 5 was based on a sample group *N*, of 512 and demonstrated alignment at sample number 2979, peak = 0.37 and sample number 2980, peak = 0.4. The sample rate was 250kHz resulting in data *B* arriving some 11.92ms before data *A*. Improved signal-to-noise ratio is possible with increasing group sizes but the software correlation time increases accordingly.

Software

Two software tools have been written. The first carrying out correlation, (COR_TIM.exe) to find the time slip and the second (F_ALIGN.exe) to trim the data files so that they align to the optimum data sample and are then useful for phase measurement.

Phase Measurement

Once the data streams have been aligned phase comparison can be effected either globally for narrow band signals or selectively in active frequency (filtered or FFT) channels using the I/Q data by applying the trigonometrical identities,

$$sin(a-b) = sin(a)cos(b) - cos(a)sin(b), \qquad cos(a-b) = cos(a)cos(b) + sin(a)sin(b)$$

or in data terms,

$$Phase(A\text{-}B) = tan^{-1}\left(\frac{AQ(f)BI(f) - AI(f)BQ(f)}{AI(f)BI(f) + AQ(f)BQ(f)}\right)$$

where, *AI* and *BQ etc.* represent the FFT bin I and Q components.

The conventional interferometer sum and difference outputs can also be synthesized.

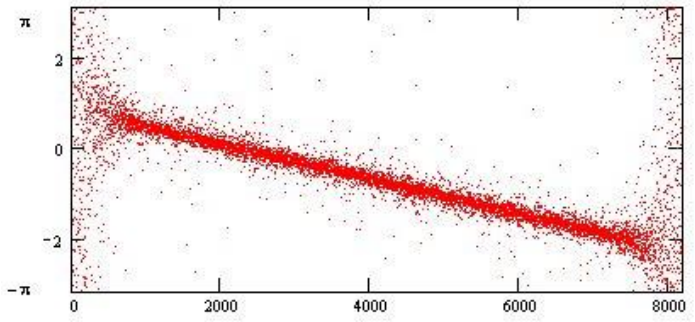

Figure 6 Residual Phase v Frequency

Figure 6 is the result of applying the phase equation to the bins of an FFT covering the whole of the SDR band (250kHz). Note that the phase changes linearly with frequency as expected from the residual delay of the two SDR data clocks not aligning. The residual delay is calculated from dt = dφ/df or π/(2π.250kHz) = 2μs. This is confirmed by the two adjacent correlation peaks being roughly equal. The plot shows that this clock timing phase error can be calibrated by an initial noise injection for each measurement so that interferometer phase data can be correctly identified.

Discussion

The techniques described above have been proven on noise calibration data but there are attractive implications for wider economical application in the general radio astronomy field. For the signal processing, very powerful PC's are available but work on a suitable software Graphic User Interface to run the software tools would be of benefit. For data collection, credit card PC's like the Raspberry Pi controlled over a network with software such as PUTTY would maintain the economical mantra.

A receiver channel could comprise, an antenna, broadband rf amplification plus a suitable filter, an SDR dongle and a Raspberry PI.

For the hydrogen line work, the cost with cabling and couplers, is about £250-£300.

Two application possibilities are,

a) Interferometry

The complex FFT outputs containing the wanted signal can be sum-difference combined for conventional analysis. With low level noise signals, integration may be necessary to improve the signal-to noise ratio. This can be carried out by averaging the amplitude processed outputs of successive short FFT data blocks to produce the √n improvement required. Drift scan interferometry can produce very large files.

b) Phased Arrays

The techniques described here can be applied to multiple antennas. A central temperature-compensated crystal oscillator at 28.8MHz can be amplified and split to feed multiple SDR dongles to lock them all in frequency. Similarly a suitably amplified noise source can be divided for the initial calibration sequence. To produce a single broadside beam, the complex FFT outputs are combined directly prior to determining the amplitude for sensitivity-improvement summation. It is feasible to correct and adjust relative phase for beam steering in the post processing.

Conclusions

The paper has shown that with some simple modifications, the frequency stability of very cheap DVB-T dongles can be improved, at modest cost, to a few Hz at 1400MHz matching professional equipment performance. Coupling the crystal circuits between a pair of dongles and mounted as described, serious vector interferometry is possible. The main problem is the amount of post data collection processing for data alignment and phase extraction required, but with software written in 'C' language the delay between data collection and results viewing can be minimised.

Very accurate and stable measurement of the Hydrogen line should add significantly to the quality of data and derived results.

Appendix - Software Processes

1. Data Collection - RN_RTLAT.exe (run_rtl_at - linux)

This uses the PC clock to run the rtl_sdr[6] program at the command-line set time, generates a bin file which is then processed with the FFT algorithm to output spectrum averaged text file.

Separate PC clocks should be updated at the same time to the same internet server. PC clocks drift in time accuracy and should be re-set prior to serious measurement.

To achieve near simultaneous recording, two terminal windows, one for each SDR, run a copy of the control software that is set with the desired recording parameters.

The command format is:-

RN_RTLAT <"rtl_sdr data.bin + command line">< Number of FFT points><hr min sec>

Typical control software commands to run two SDRs (-d 0 and –d 1) on a single computer for MS WINDOWS are,

>RN_RTLAT "./rtl_sdr dat0.bin –f 1420e6 –d 0 –g 49 –n 100e6" 256 14 24 00

>RN_RTLAT "./rtl_sdr dat1.bin –f 1420e6 –d 1 –g 49 –n 100e6" 256 14 24 00

These are run in two Command windows set to the working directory, which also contains the Osmocom rtl_sdr tools.

2. Correlating Data Sets COR_TIM.exe (cor_tim - linux)

Although multiple SDR dongles can be programmed to sample data at the same time, PC housekeeping and internal clock monitoring is unsynchronised so that data collection start times can vary by several tens of milliseconds. This correlating software identifies the time slip for correction. With this information, data sets can be aligned to within a data clock period.

The command format is:-

COR_TIM <infile1> <infile2> <outfile> <sample length> <No 8192 blocks>

The number of data bocks chosen for this command only needs to be large enough to ensure the correlation peak becomes evident.

A typical command line and response is shown in Figure A1.

```
COR_TIM aa01.bin aa02.bin ff.txt 512 8

Number of Input Bytes = 4000000   No. of IQ Samples=65536

Block   Sample  Corr    Offset  Corr    Offset
Size    Width   Peak1   1 on 2  Peak2   2 on 1

8192    512     0.210   -255    0.183   -4675
No Blocks cross-correlated: = 7679
8192    512     0.225   3216    0.682   8013
No Blocks cross-correlated: = 15871
8192    512     0.225   3216    0.682   8013
No Blocks cross-correlated: = 24063
8192    512     0.225   3216    0.682   8013
No Blocks cross-correlated: = 32255
8192    512     0.225   3216    0.682   8013
No Blocks cross-correlated: = 40447
8192    512     0.225   3216    0.682   8013
No Blocks cross-correlated: = 48639
8192    512     0.225   3216    0.682   8013
No Blocks cross-correlated: = 56831
8192    512     0.225   3216    0.682   8013
No Blocks cross-correlated: = 65023

Infile1=aa01.bin    Infile2=aa02.bin    Outfile=ff.txt   No Samples=512
```

Figure A1 Correlation Tool Response for Data Files

Perfect correlation with data clock alignment is unity. This example shows a correlation of peak of 0.682, comparing a sample block of 512 of file 2 located on file 1, 8013 samples later. Random noise peaks around the 0.2 are visible, but the true correlation peak or adjacent peaks are normally more than double this level. A requirement is for the calibration noise source to dominate over system noise. To avoid possible distortion at the file start, the software ignores the first 7000 IQ pairs.

3. Data Set Alignment F_ALIGN.exe (f_align - linux)

The command format is:-

F_ALIGN <file_In> <file_Out> <No of 8192 IQ blocks> <Align Address>

For the files in Figure A1, we need to determine the alignment address and file size (number of 8192 IQ blocks).

The number of useable 8192 IQ data blocks is given by, (4000000-7000-8013)/2/8192 ~ 486 complete blocks. The alignment start address for files 1 and 2 are (7000+8013) and 7000 respectively.

The example relevant commands are,

F_ALIGN aa01.bin aa01a.bin 486 15013

F_ALIGN aa02.bin aa02a.bin 486 7000

This will produce two files of the same length and virtually data clock synchronised. This is demonstrated re-running COR_TIM with the new data files as shown in Figure A2.

```
COR_TIM aa01a.bin aa02a.bin ffa.txt 512 8

Number of Input Bytes = 19660800   No. of IQ Samples=65536

Block   Sample  Corr    Offset  Corr    Offset
Size    Width   Peak1   1 on 2  Peak2   2 on 1

8192    512     0.692   0       0.692   0
No Blocks cross-correlated: = 7679
8192    512     0.692   0       0.692   0
No Blocks cross-correlated: = 15871
8192    512     0.692   0       0.692   0
No Blocks cross-correlated: = 24063
8192    512     0.692   0       0.692   0
No Blocks cross-correlated: = 32255
8192    512     0.692   0       0.692   0
No Blocks cross-correlated: = 40447
8192    512     0.692   0       0.692   0
No Blocks cross-correlated: = 48639
8192    512     0.692   0       0.692   0
No Blocks cross-correlated: = 56831
8192    512     0.692   0       0.692   0
No Blocks cross-correlated: = 65023

Infile1=aa01a.bin    Infile2=aa02a.bin    Outfile=ffa.txt   No Samples=512
```

Figure A2 Correlation Tool Response for Aligned Files

4. Data Analysis BIN_TXTR.exe (bin_txtr - linux)

This tool converts binary data to text format so that data samples can be processed using Excel or Math CAD software.

The format is,

BIN_TXTR <binfile_In> <Textfile_Out> <No. of 8192 IQ blocks> <Start block> <End Block>

Text files can be very large but a typical example is,

BIN_TXTR aa01a.bin aa01aa.txt 0 512

This produces text files formatted with sample number, *Q* data, *I* data; about 94MB size.

5. Software Link

http://www.y1pwe.co.uk/RAProgs/NewSW.zip

Peter W East. March 2015

Coming up in the next issue:

How to display real time audio FFT using Python.

Editors note: I ran out of time to complete the above article for this issue, so I will include it in the next one. Currently the example code works on a PC but not on the Raspberry Pi. I hope to be able to show it working on the Pi by the next issue.

Jeff Lashley