Notes on ZCU111 RFSoC Characterisation

Some tests were performed to assess the ZCU111 RFSoC ADCs for suitability for RA applications. This document is an on-going record of some measurements and characterisation that are currently underway to understand details of the RFSoC ADCs

Spectrometer bandpass response with typical noise power input

To assess the presence of any spurious signals in the ADC spectrum from either on chip or on-board sources a sensitive integrating spectrometer design was programmed into the ZCU111. The design consisted of 4 of the 8 ADC cores instantiated, comprising one core in each of the 4 tiles. The DACs were not instantiated. The ADCs were run as close as easily possible to the maximum sample rate of 4.096 GS/s, and the firmware design timing constraints were set to this clock rate as we would require operation at this clock rate for a potential CABB upgrade for instance which has 2.048 GHz bands.

To avoid having to de-solder tiny caps, which is required to allow an external clock source, the closest sample rate easily achievable using the Xilinx delivered clocking resources on the ZCU111 was used. This required reprogramming the LMX2519 synthesiser chips, which have a crystal reference of 122.88MHz to give a 245.76 MHz reference to the FPGA ADC PLL input which is multiplied up to 245.76 * 16 = 3.93216 GHz. This is close enough for characterisation purposes.

The firmware design consists of a 32 point full precision critically sampled coarse filter-bank delivering 16 sub-bands of 122.88 MHz. Any one of these sub-bands can be selected and is passed to a further 4096 channel critically sampled fine filter-bank with a post power detection and integration. A wide-band noise source was applied to the input and filtered to less than the 2GHz sampled bandwidth. Below is an example spectrum obtained in sub-band 8 with 5s integration. The spectrum looks extremely clean with no evidence of contaminating signals.

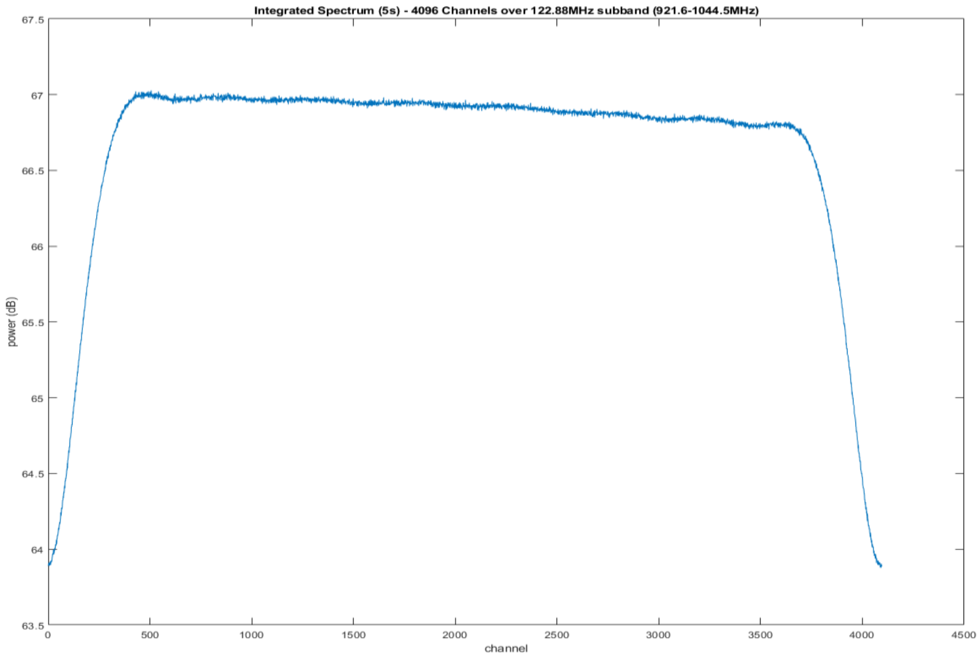

Figure 1 Sub-band 8

This sub-band includes the ¼ sampling frequency which would be expected to show interleaving spurs with an 8 way time interleaved converter such as this if the time or amplitude calibration was off. There is no obvious contamination by any spurious signals. A second sub-band is shown below this time at low frequency which should contain the fundamental frequency of many of the board clocks.

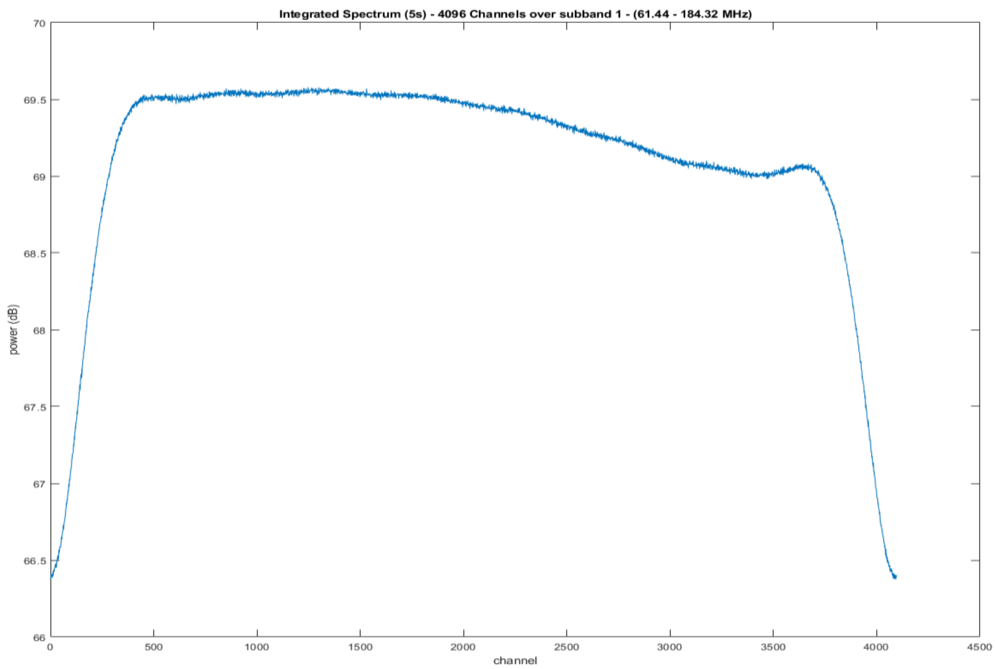

Channel Isolation

To measure the combined adjacent channel isolation of the inputs, board traces and on chip ADCs the following test was performed. Noise at a typical input level was supplied to one channel. At the same time a -2dB FS sine wave input was supplied to the adjacent ADC channel and digitised. The 5s integrated spectrum on the noise channel was recorded to measure the leakage of the tone into the adjacent ADC channel data. The following result shows the spectrum of the noise channel with the signal tone contamination clearly evident. The tone was at 983.0 MHz.

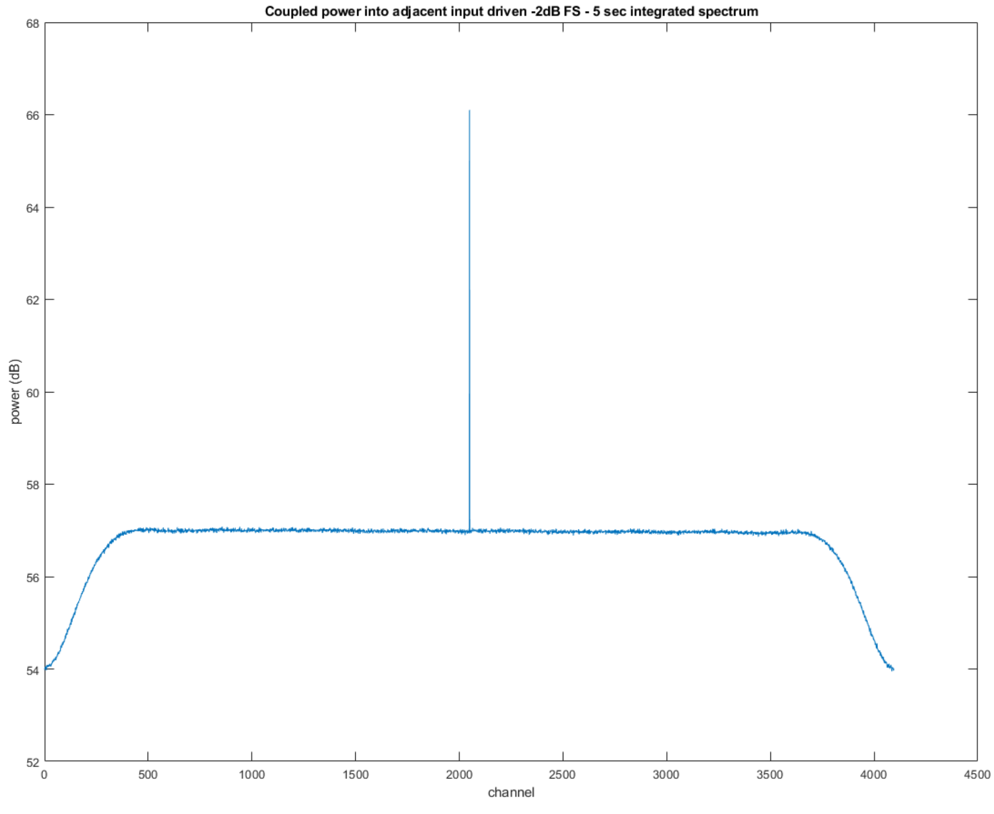

From the spectral density of the input noise, the amplitude of the signal tone and the size of the coupled component the adjacent channel isolation was measured to be -65 dB.

System Stability

An important consideration for using the RFSoC for RA astronomy is system stability. That is if the amplitude and bandpass response of the ADC remains stable over time to allow signals to be integrated for improved SNR. This is especially important for single-dish spectrometers and though less so, also important for use in an interferometer. There is also the question of whether unknown background calibration mechanisms in the RFSoC ADCs might cause periodic shifts in the response. To try and answer some of these the output of the integrated spectrometer was measured over a period of hours and the reduction in variance measured over time. An Allan Variance analysis was performed to estimate the typical system stability of the bench setup and the output time series examined for evidence of abrupt shifts. The noise source used was insulated as best I could but I did not have access to a temperature stabilised noise source so it is highly likely a lot of the long term drift could be due to noise source temperature dependence.

Below is the Allan Variance and time series for the entire averaged power across ~100 MHz of the full sub-band encompassing 1GHz with 1 second integrations (i.e. channels were added and divided by number of channels used so is a sensitive total power measurement).

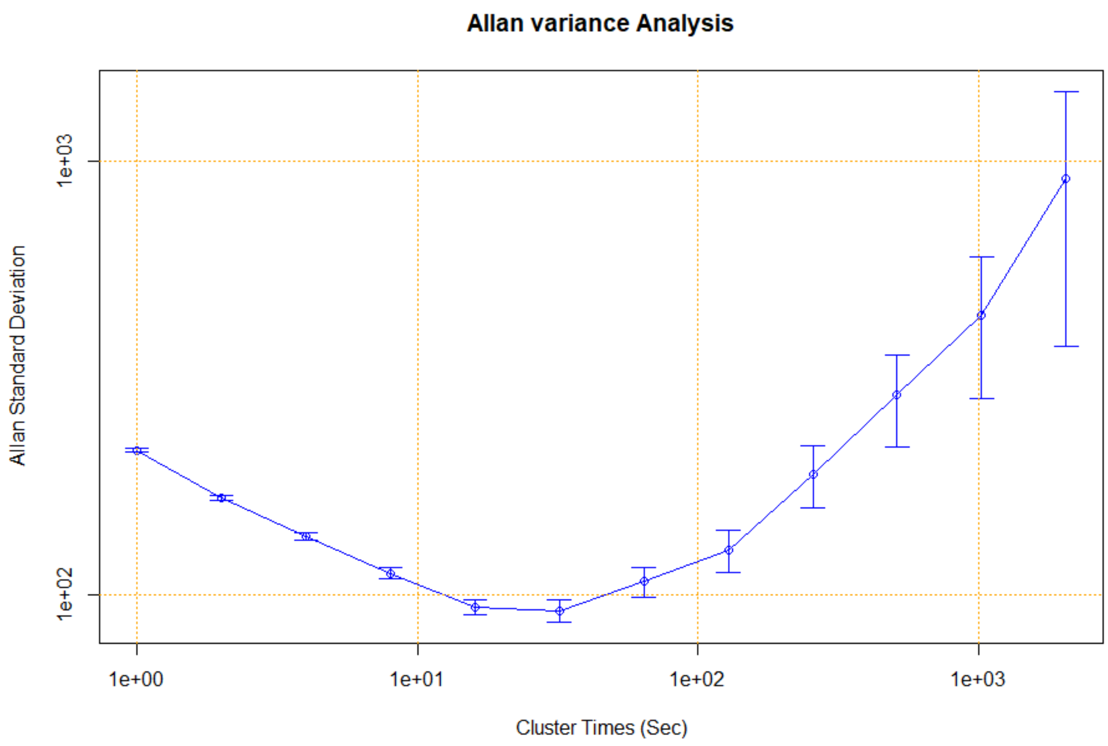

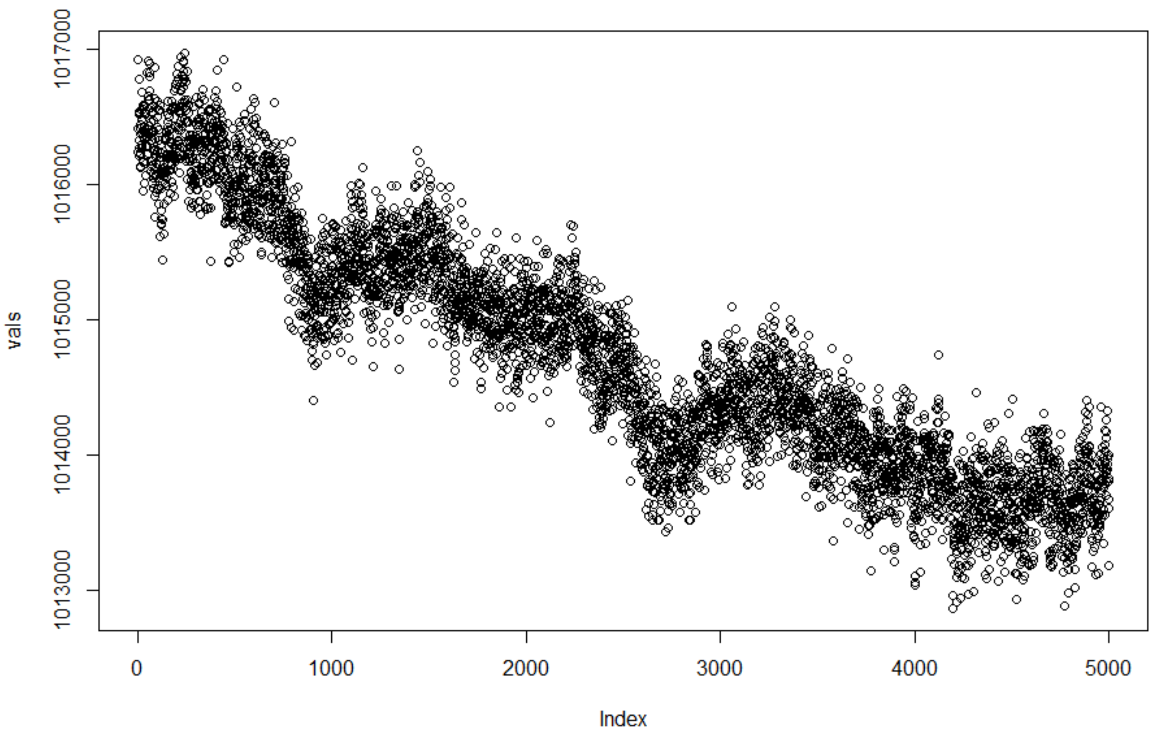

Fig 5. Channel power vs sequential 1s integration number.

At this level of sensitivity the AV indicates a stability period and useful integration time of around 10s.

A second measurement was taken this time considering just one of 4096 sub-channels, as might be relevant for example for a spectral line observation. At this bandwidth the statistical counting noise is naturally much greater and is dominant over the total power drift leading to a much longer useful integration time in the 100s of seconds.

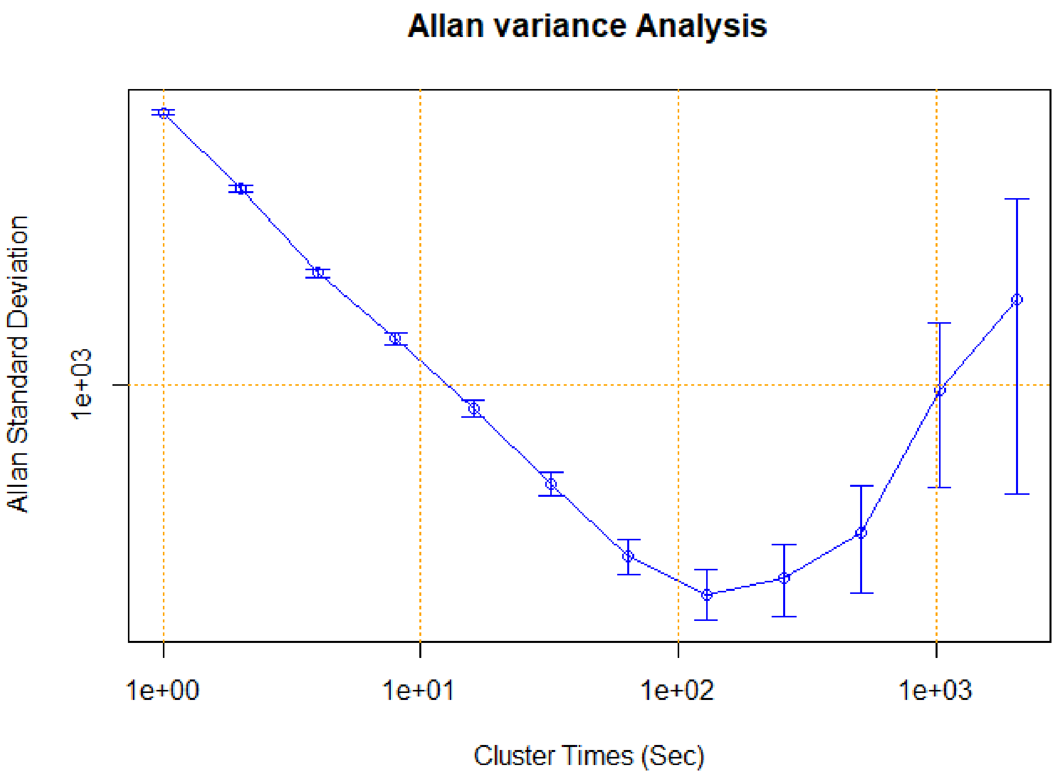

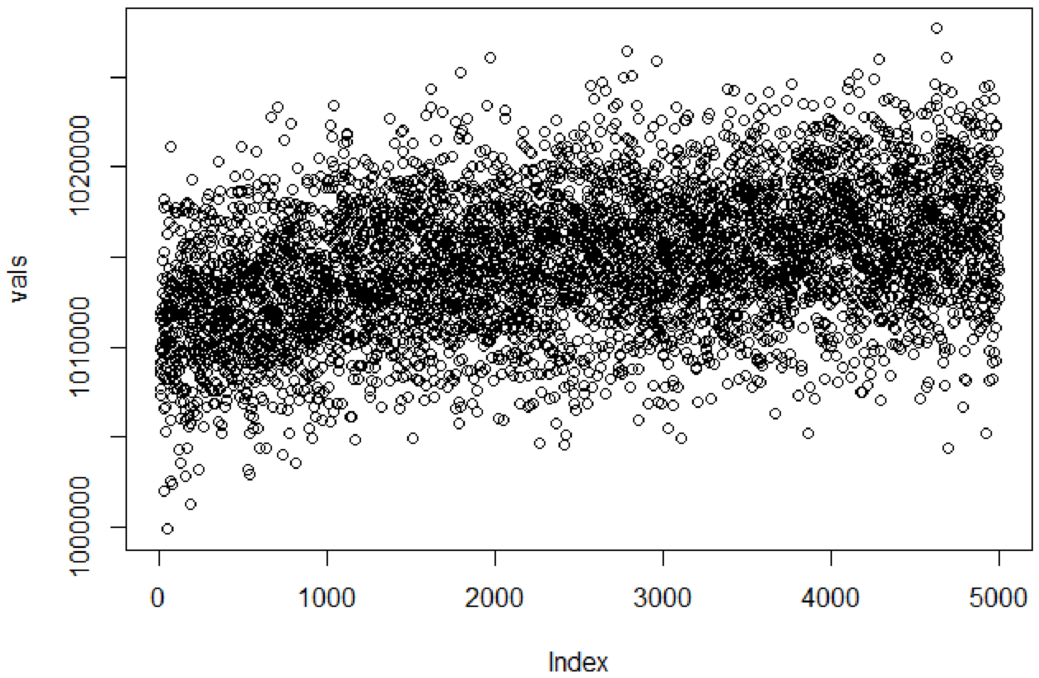

Fig 7. Time series of 1 individual channel output power sequential 1s integrations.

I don't think the stability of the noise source is good enough to get near the inherent stability of the ADC. Perhaps thhe best method would be to use one of the cryogenic receivers in the lab which are probably the best stabilised noise sources available.

To look for evidence of any periodic calibration effects the FFT of the high sensitivity time series fig 5 was taken and is below.

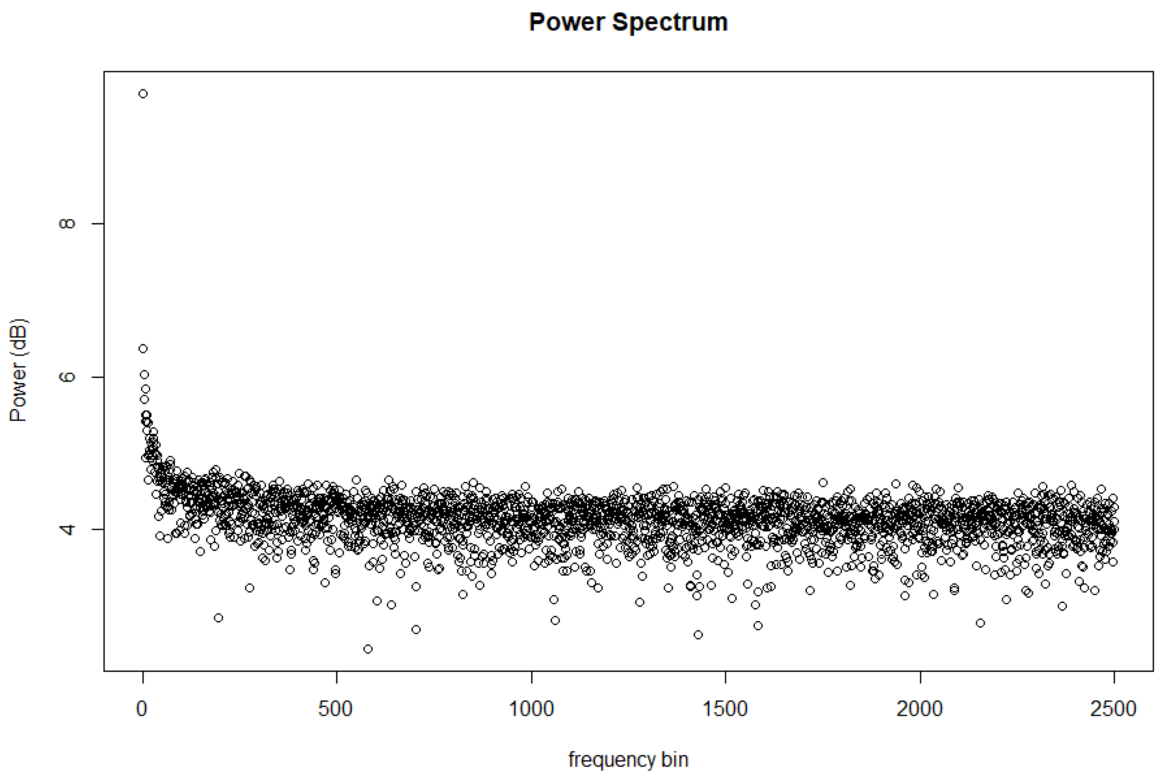

Fig 8. FFT of fig 5 time series

There is no obvious peak corresponding to a periodic calibration interval that shifts the amplitude. The spectrum illustrates the typical drift 1/f noise spectrum expected. This is not conclusive of course just suggestive. Calibrations may occur on random intervals or when triggered by a parameter drifting beyond a certain range.

That said, I'm not convinced yet that some calibration effects are not occurring. The time series plot figure 5 does not look random to me in places. There seem to be several places where there are consistently high or low consecutive power values that then return rapidly to the baseline. For example look at the spikes of samples near 5000 and 4900, and the negative going set near 1700. This requires some further examination and seeing if it is possible to freeze the calibration and what effect this has …

Paul Roberts CSIRO Nov 2018