MWA beamformer temperatures

Each MWA tile has an associated beamformer in the desert next to it, which combines the signals from the 16 dipoles using a specified set of delays to 'point' the tile in the desired direction. The desired set of delays is sent to the beamformer over coaxial cable (also used to carry 48 VDC to the beamformer, and RF signal from the beamformer). The data is sent as a 253 bit packet using a clocked serial protocol, and after each new pointing is sent, sending 24 more clock pulses results in an 8-bit status byte and a 12- bit temperature value being sent back from the beamformer. Because the beamformer can only send temperature data after a new pointing packet is received, beamformer temperatures are only ever known at the instant each observation starts.

The beamformer temperatures were saved to a PostgreSQL table in the 'telemetry' database, and unlike the main schedule database, these tables were not replicated off site.

Most MWA data analysis doesn't use the beamformer temperatures, but the EoR team has noticed gain variations that means the MWA temperature is a required input to their data reduction.

Lost data

In February 2018, a combination of hard disk failures and human error meant that the database table containing all beamformer temperatures (before Feb 6, 2018) was lost. Since then, many measured beamformer temperatures were recovered from various locations (observation metafits files in various places, the Carbon/Whisper 'round robin' database used by our system health visualisation tool, and a few ad-hoc database dumps made over the years). That was still only a tiny fraction of the total number of observations.

The recovered beamformer temperature measurements, combined with measurements taken after Feb 6th, were used to generate fit parameters that could be used to estimate each of the beamformer temperatures using the current air temperature, and the current rate of change of air temperature in deg C per second. This fit was only valid at night time (to avoid the problem of variable amounts of solar heating), but since only the EoR team is interested in the temperature data, restricting it to night-time observations was acceptable.

The fitting model is a simple linear relation:

Tpred = Tair + offset + (m * dTair/dt)

Where Tpred is the predicted beamformer temperature, dTair/dt is the rate of change of beamformer temperature in deg C per second, and 'offset' and 'm' are the fit parameters. The fit parameters were chosen using the 'curve_fit' function in the scipy.optimize library, using offset=3.0 and m=-1.0 as an initial guess.

Firstly, the parameters were fit individually for each tile, for all measurements. A second fitting stage was carried out for some tiles after the splitting the measurements into two or more date ranges. The split into date ranges was done by visually inspecting the results of the first fitting stage, and locating dates at which the mean errors changed substantially. Where it was possible to verify using site trip

maintenance reports, these dates corresponded with beamformer substitutions (to repair damaged units) on that tile.

Some tiles have three or four sets of air temperature fit parameters, each valid over a different range of dates. The final parameters used give a typical standard error of +/- 1.2 deg C between measured and estimated temperatures. A sample plot of measured versus estimated beamformer temperature (below) for typical tile shows the errors. The different coloured dots represent time ranges with different fit parameters for that tile.

The air temperature data to calculate the predicted beamformer temperatures came from averaging measurements of the outside air temperature at each of the 16 receivers, and from a weather station operated by CSIRO only a few kilometres from the MWA. Where gaps in both of these data sets overlap, predicted beamformer temperatures could not be calculated.

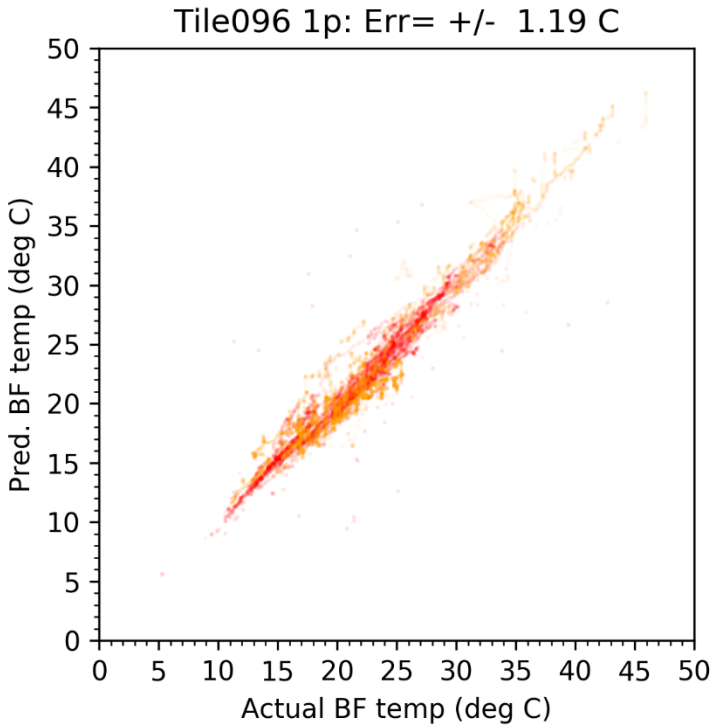

Figure 1: Predicted versus actual beamformer temperatures for Tile096, 2013-2018. Red dots are for measurements before 2017/06/15, orange dots for after that date. Fit parameters used to generate the predicted temperatures are different for the two periods.

Results

The calculated fit parameters were used to generate predicted temperatures for all observations between 2012/6/1 and 2018/11/08 where air temperature measurements were available, and where valid data files were recorded by the correlator. This includes daytime observations, but solar heating during the day means these predictions are a lower limit only.

The results were written to the 'bf_temperatures' table in the main schedule database, mirrored to Pawsey in real time, along with air temperature measurements at the start of each observation in the 'air_temperatures' table.

In total, between the first G0009 (EoR) observation on 2013/07/20 and now (2018/11/08) there have been 117644 valid observations for project ID G0009, of which 15281 (13.0%) have measured beamformer temperatures saved, 100885 only have predicted temperatures saved (85.8%) and 1478 have neither measured nor predicted temperature data available (1.26%).

New measured beamformer and air temperatures are being saved in real-time to the same bf_temperatures and air_temperatures tables used to store the historical data.

Because the fit parameters for each tile will change as beamformers are swapped for maintenance, new observations will NOT have predicted beamformer temperatures calculated and saved.

Access to beamformer temperature data

Beamformer temperatures are present in the metafits files created through the metadata web service (http://mwa-metadata01.pawsey.org.au/metadata/getfits?obsid=NNNNNNNNNN) in the same data table that contains tile locations, cable lengths, etc. The 'BFTEMPS' column contains a single value for each input – the measured beamformer temperature if it is available, the predicted beamformer temperature if it can be calculated, or zero if neither is available. Note that metafits files generated before 2018/11/08 (using the web service, or downloaded with the data files) will not contain the predicted beamformer temperatures.

There is also a dedicated beamformer temperature web service, that returns a JSON structure. It is called as http://mwa-metadata01.pawsey.org.au/metadata/temps?obsid=NNNNNNNNNN and returns a dictionary with tileid as a key and tuples of (TileName, recieverID, slotNumber, Temperature) where the temperature value is either the measured or predicted value, or null, as above. It can also be called with the 'dictformat' argument, to return a dictionary for each tile instead of a tuple, containing more information (the tile name, receiver ID, slot number, air temperature, rate of change of air temperature in deg C per second, the measured beamformer temperature, and the predicted beamformer temperature). For this format, add '&dictformat=1' to the end of the above URL.

If you need more information about the beamformer temperature fitting (the actual fit parameters, error plots like the above for all tiles, etc), please contact me directly (Andrew.Williams@curtin.edu.au).