

Visibility file times and 'Quacktime'

Background:

The RRI receivers have three elements that need to change state at the start of an observation, depending on the contents of the schedule database for that observation. These elements are:

- A beamformer on each tile. The desired set of hex delays need to be sent to the beamformers. This takes around 0.2 seconds, and will produce RFI on the tile output during this time.
- A variable attenuation input stage for each of the 16 inputs (8 tiles in X and Y). This is currently set one input at a time using serial commands over an I2C bus, taking around 2.2 seconds in total. During this process, inputs will be changing between the old and new attenuation settings - if these are the same, there shouldn't be any effect on the output signal. New software (yet to be deployed) sets 8 inputs (four tiles in X and Y) at a time, dropping the total change time to around 0.3 seconds.
- The desired set of 24 coarse channel outputs need to be selected for transmission over fibre to the correlator building, by commanding the AGFO module. This takes around 1.5 seconds, and it's not known what the output data looks like during this time, even if the initial and final channel sets are the same.

The receivers don't start changing state until the instant of the given observation ID (a time in GPS seconds). After that instant, if all is going well, pointing the beamformers takes 0.2 seconds, changing the coarse channels takes 1.5 seconds, and changing the analogue attenuation takes 2.2 seconds (all three changes happen in parallel). All hardware (pointing, AGFO and ASC attenuation) is always sent change commands for every observation.

Sometimes, there are additional delays in some or all receivers, typically due to high load on the main M&C PostgreSQL server, or the PostgreSQL servers inside the individual receivers. That happened often in the past, less so recently.

Discarding visibility data:

For visibility files, the amount of data that needs to be discarded will depend on:

- How long the receiver hardware takes to produce a stable data stream, given tile pointing, coarse channel selection, and attenuation setting.
- How many seconds passed after the obsid before the visibility data actually started being recorded. That number is exactly two seconds now, but for very old data it's ranged from being 1 second *before* the obsid to 4 seconds after, and not necessarily the same for each GPU box.
- What the correlator averaging is in the current mode - for example, a fraction of a second of bad data will contaminate an entire two second dump.

Starting in mid 2017, I added two fields to the metafits files:

- The QUACKTIM header contains the amount of time (after the obsid, not after the start of recorded data) contaminated by pointing/freq/attenuation changes, rounded up to the next correlator dump time. It depends on whether the channel set and attenuation in this observation are the same as the previous observation, or different.
- The GOODTIME header is the unix timestamp of the first instant of 'reliable' data (obsid + QUACKTIM, converted to a unix timestamp).

At the `_end_` of an observation, a few seconds of data are missing (any files in the queue are discard), but you shouldn't have to discard any additional data that's been recorded. It's possible that for very old observations, the data might have kept on being recorded into the start of the next observation, but it certainly isn't doing that now.

Summary - a chronology of MWA data:

2012 to 2014(ish): *The Dark Ages* - chaos and uncertainty, before the civilised era. Code changed day-to-day, there were no dipole tests or flags, and the QUACKTIME and GOODTIME cards were a distant dream. Leap second offsets were hard-coded in dozens of places. Don't trust any timestamps to better than a few seconds. The data in a visibility file starts anywhere from -1 to 4 seconds after the obsid, and the first 0-4 seconds should be discarded, depending on correlator averaging time. This era is why Cotter discards the first four seconds of data, by default. Here be dragons in the data.

2014(ish) to mid 2017: *The Renaissance* - the code was stabilising, dipole tests and flags exist. From here, visibility files should always start 2 seconds after the start of an observation. There is no QUACKTIM or GOODTIME card in the original metafits files, but newly created metafits files will have those cards.

mid 2017 to now: *The Modern Era* - stable correlator code, no more leap seconds. QUACKTIM and GOODTIME headers, and data files that `_always_` start exactly 2 seconds after the obsid. Because of that 2 seconds without data, only attenuation settings should ever contaminate recorded data, only if the initial and final attenuations are different, and only for one correlator dump time.

2021 onwards: *The Future* - new (already tested) receiver code drops attenuation change times to < 0.5 seconds. The new MWAX correlator solves all of the problems of the old correlator, and has no bugs of its own. The promise of new receivers heralds the dawn of a new age.