CAG *I-LOFAR* Realtime Event Triggering Proposal May-2016

On 9-Jun-2016 a substantial portion of the CAG infrastructure was transferred to the SCSS School Cloud. The remaining bones of the infrastructure were restructured as a *Server Pool*, a *Blade Realtime Cluster* for processing *I-LOFAR* data, a *Dell2950 Realtime Cluster* for event triggering, and *CAG Storage* for data and results. Here we outline a proposal for realtime event triggering.

I-LOFAR data processing

There will be continuous data traffic at 3Gbps on paths from Birr Castle to the CAG storage. The Birr Castle LOFAR station has 96 high-band (120-240MHz) and 96 low-band (10-80MHz) antennae, both dual-polarised, i.e. effectively there are 192 x HB and 192 x LB antennae. Depending on the science targets, an antenna/receiver chain is selected and configured to prefilter and amplify a specific antenna output signal, effectively selecting a 100MHz slice of the spectrum, either 0-100MHz, 100-200MHz or 200-300MHz. The result is then sampled at 160MHz or 200MHz to 12-bit resolution and fed to four FPGA-based Remote Station Processing (*RSP*) modules and two FPGA-based Transient Buffer Boards (*TBB*) with memory banks for each RSP. The TBBs contain buffers that are intended to be used for storing transient and cosmic ray events. These buffers store a 1.3sec snapshot of the raw 100MHz data (can be upgraded to 10.4sec), triggered either externally or by an internal detection algorithm in each signal path. The contents of these buffers will be extracted from the Birr Castle station and streamed to CAG storage for realtime processing.

The snapshot intervals and buffer data transfer rates (max.6:1 duty cycle) will depend upon how fast the processing is. It is suspected this will be intractable for frequent realtime execution on the existing hardware, but this can be upgraded as needed. LOFAR digitizes 96 dual-polarized antennas at 200MHz x 12-bits (in all ~4.8 x 10^{11} bps), so for example 1sec of raw data will take ~2.5mins to stream at ~3Gbps to anywhere (whether in Birr or TCD). If an event occurs at time *t*, event triggering will not see it until too late.

Realtime event triggering

Here we propose instead that a small quantum of each snapshot is *sampled* very frequently for event triggering. This will allow events to be recognised before it is too late. In addition the sampled data is then a valuable fine-scale sampled time series of raw data.

Let us assume that at time *t* only a small quantum of each snapshot, the leading epoch T_Q, is streamed to CAG storage. If the event processing takes time T_P, and if $(T_Q + T_P) < T_S$, (where is the sampling interval) then if an event occurs at time *t*, event triggering can hopefully detect it quickly enough to request streaming of the entire snapshot. If not, then it can request a new quantum.

In the limit, if each TBB is upgraded to 10.4sec and raw data is streamed at ~3Gbps, and event processing takes time $T_P = 0$, the max.quantum of data streamed before 10.4sec elapsed would be ~3.10^{10} bits (~4GB), and therefore max.T_Q = 625mS, i.e. if 625mS of data were streamed at ~3Gbps then infinitely fast event triggering could detect it quickly enough to request streaming of the entire snapshot. Unfortunately, if the event occurred later in the snapshot it would not be recognised.

Instead let us propose that a much smaller quantum is streamed, say 6.25mS (i.e. 3 x 10^9 bits). At ~3Gbps this would take 1sec to stream. Then if in an event is detected the whole 10.4sec snapshot can be streamed, otherwise a new quantum can be requested. In effect 6mS of the raw data is sampled every second for event triggering until an event is detected, whereupon it is captured. Of course the triggering cannot be infinitely fast, but a shorter quantum, say 1mS, may be tractable in realtime.

Maximising stream concurrency

Each LOFAR stream is rate limited to 140Mbps, so in theory 21 streams are needed for an aggregate 3Gbps of streaming. There are eight compute nodes with 10Ge interfaces, so let us propose 24 streams, three per node. There are six storage RAIDs, each able to stream concurrently, so each RAID can accommodate streams from 2 nodes. This arrangement will maximize stream concurrency.