Use Cases - Separable Modes

Due to the very tight timeline for Atomic COTS (Figure 1, ECP200039) documentation (now complete) no further work has been done on the Separate Filterbank approach (Figure 2, ECP200028). However, both these designs have a degree of commonality in terms of design and both have greater flexibility than the unified design present at CDR. The capabilities of both to implement multiple use cases is explored here.

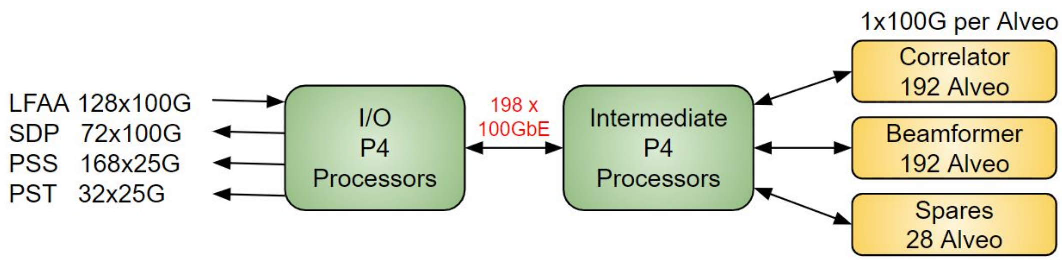

Figure 1: *Atomic COTS approach*

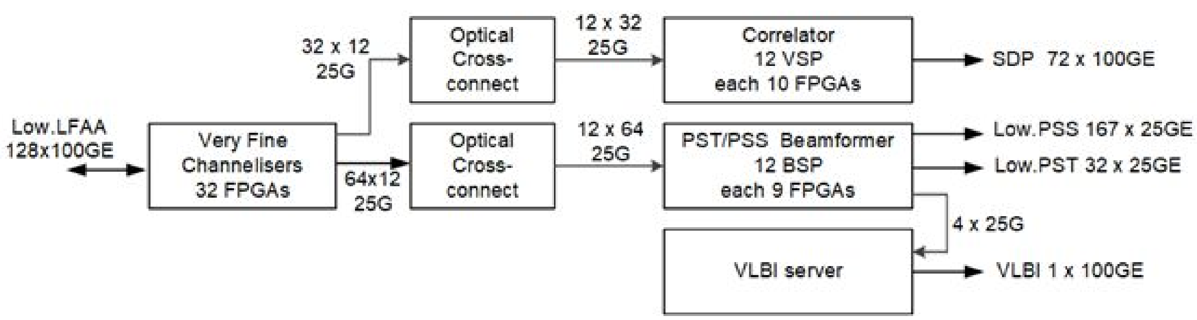

Figure 2: *Separate Filterbank approach*

Both have fixed resources for the beamformer and correlator. The Atomic COTS design has warm spares built in that could be used to provide extra capabilities and the Separate Filterbank approach has spare slots that are fully cabled to allow further Gemini to be installed. Both designs separate out beamforming from correlation. Both designs process a small part of the bandwidth in each correlator and beamformer FPGAs. The main difference is that the Separate Filterbank approach processes a number of fine channels per FPGA and the Atomic COTS coarse channels. But Atomic COTS can then emulate the Separate Filterbank approach by redistributing fine channels (it has the benefit of both approaches).

The number of possible Use Cases that Low.CBF is capable of implementing is huge but the problem is reduced if different modes in a Use Case do not interact. So beamformer modes can be implemented independent of correlatore modes. In the beamformer choice of PSS mode or frequency does not affect PST modes. In the correlator, implementation of zoom modes in one subarray does not affect standard processing in another. Unfortunately the same is not true of substations.

Correlator and Beamformer Separability

If there are N modes in one function and M in the other the total number of modes is NM. If the functions are separable the number of modes is N+M and each can be dealt with separately. In Figure 1 and 2 it is seen that each has separate hardware for correlation and beamforming (i.e., these can be seen as separable modes).

Looking at the Separate Filterbank approach the input data is common in the Very Fine Channelisers. Within the Very Fine Channeliser the only common components are the data ingest and corner turn operation and the readouts are independent. All signal processing and subsequent data flows are independent and the correlator and beamformer functions are separable.

In the Atomic COTS design an FPGA is either a beamformer or correlator (i.e. separable modes). However, the correlator and beamformer share the P4 switches and each has different needs in terms of routing and data duplication. Inputs are defined by LFAA and there is always the capacity to accept 384 channels from a station. Inputs to the FPGAs have sufficient bandwidth to accept the data from the switches. On the inter-P4 Switch link there is no data duplication, for standard processing, and the data rate on a 100GbE link is 33Gbps.

For both designs correlator and beamformer processing can be considered separately from each other. For the correlator, as will be shown, there are further levels of separability.

Beamformer

Both designs are capable of forming

- 16 PST beams on 300MHz of bandwidth
- 500 PSS beams on $3/8^{ths}$ of 300MHz (112.5MHz) or 250 PSS beams on 3/4 of 300MHz (225MHz)

All on 512 stations. The factors that affect the beamformer are subarrays and station beams. Subarrays reduce the number of inputs to a beam. Station beams may reduce the bandwidth available for a beam. Both multiple station beams and/or subarrays reduce the compute requirements on the hardware. What remains constant is the 384 channels from each of 512 stations and the distribution of that data across all FPGAs producing beams. So all beamformer modes for any configuration of subarrays and station beams are possible. And importantly this processing does not affect the correlator processing.

Correlator

For both designs the correlator can implement

- A “standard” correlator 5.4kHz channels on 300MHz for 512 stations, OR
- 4 zoom bands (arbitrary bandwidth per band) on 300MHz for 512 stations, OR
- A 5.4kHz resolution correlator on 75MHz for 1024 substations, or 18.75MHz for 2048 substations, which were the original requirements. Atomic COTS documentation has been expanded to describe multiple substation modes. This has yet to be done for the Separate Filterbank approach.

It will be shown that for (512) stations that station beams, subarrays, and zoom vs “standard” processing are separable.

Station Beams

A station beam is a group of channels for a subarray that has a different delay polynomial to channels in another station beam. The implementation of delay correction occurs at the filterbank which processes all data from a station. After the delay correction the processing is the same for all channels assigned to a substation, independent of the delay polynomial attached to any particular channels. As far as modes of operation are concerned the correlator can process any combination of station beams assigned to a subarray. This is also true for subarrays of substations. In the correlator the configuration of station beams is independent of any other mode.

Subarrays of Stations

When processing stations in each FPGA both approaches receive data from all stations for a subset of the bandwidth. For each time sample this is a data plane in memory as illustrated in Figure 3.

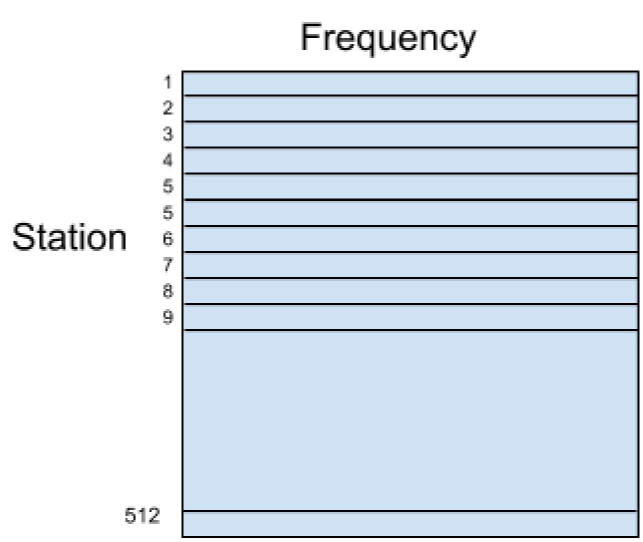

Figure 3: *Station-Frequency Data Plane*

Subarrays are formed by assigning stations to it. Figure 4 illustrates a case where the telescope is divided into three subarrays: blue, yellow and green

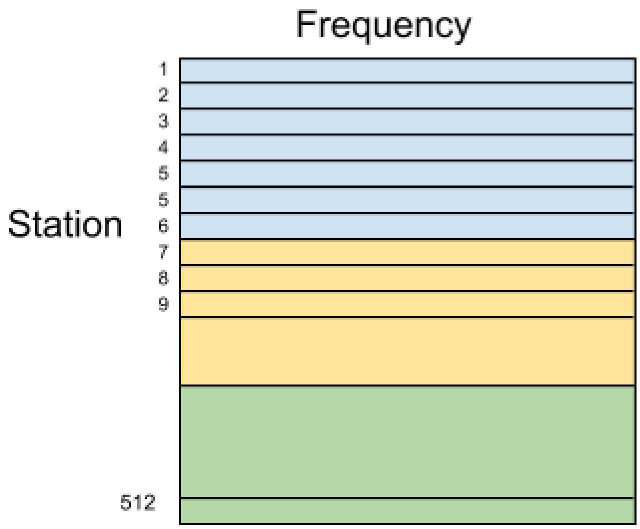

Figure 4: *A three subarray configuration*

Each subarray can be processed independently in the FPGA. Data routing for each station is also controlled independently of any other. Thus the routing for the stations in one subarray can be configured for standard 5.4kHz processing while that for another subarray can be configured for zoom modes. Subarrays of stations are independent entities that are separable from each other. Routing solutions for a zoom mode or “standard” mode for the full array define the routing required for each station. This same routing is used for stations in a subarray to implement the zoom mode or standard mode in the subarray.

Subarrays in the correlator have one limitation and that is due to the granularity of the correlator. The correlator has a minimum block size N. For 512 stations there are 512/N blocks and $(512/N)^2/2$ iterations of the correlator are needed to process a single frequency channel. This sets the compute power of the correlator. For any number of stations in a subarray from 512 to 512-(N-1) the same compute power is required. Any further subarrays using the remaining stations requires extra compute power in the correlator. If the extra compute power is not available then these extra (small) subarrays cannot be processed.

Zoom Mode Routing

A difference between the Separate Filterbank approach and Atomic COTS occurs for zoom modes. In the Separate Filterbank approach the basic unit of bandwidth is 43.4kHz and in Atomic COTS it is 781kHz. With 43.4kHz bands the data for a zoom band (4MHz or greater) can be distributed over many FPGAs. The output data rates do not saturate the output bandwidth of the FPGA, even for zoom bands with frequency resolutions of 0.226Hz. For Atomic COTS the bandwidth of channels is 18 times greater and the output from the FPGA could saturate for 0.266Hz zooms. The data rate is reduced by a second data distribution stage which spreads the data for high resolution zoom channels across a number of FPGAs.

Substations

Substations do not have a single data plane representation. Consider just the case of a single subarray of substations. It can have anywhere between 513 to 10,032 substations and each represents a different correlator design. A 16 antenna ATCA design is different to 27 antenna VLA design, which is different to a 128 tile MWA design. For the SKA the correlator is asked to adapt to a 20:1 variation in the number of substations (compensated by a reduction in bandwidth to keep compute requirement constant). The resulting data planes are very different, as shown in Figure 5.

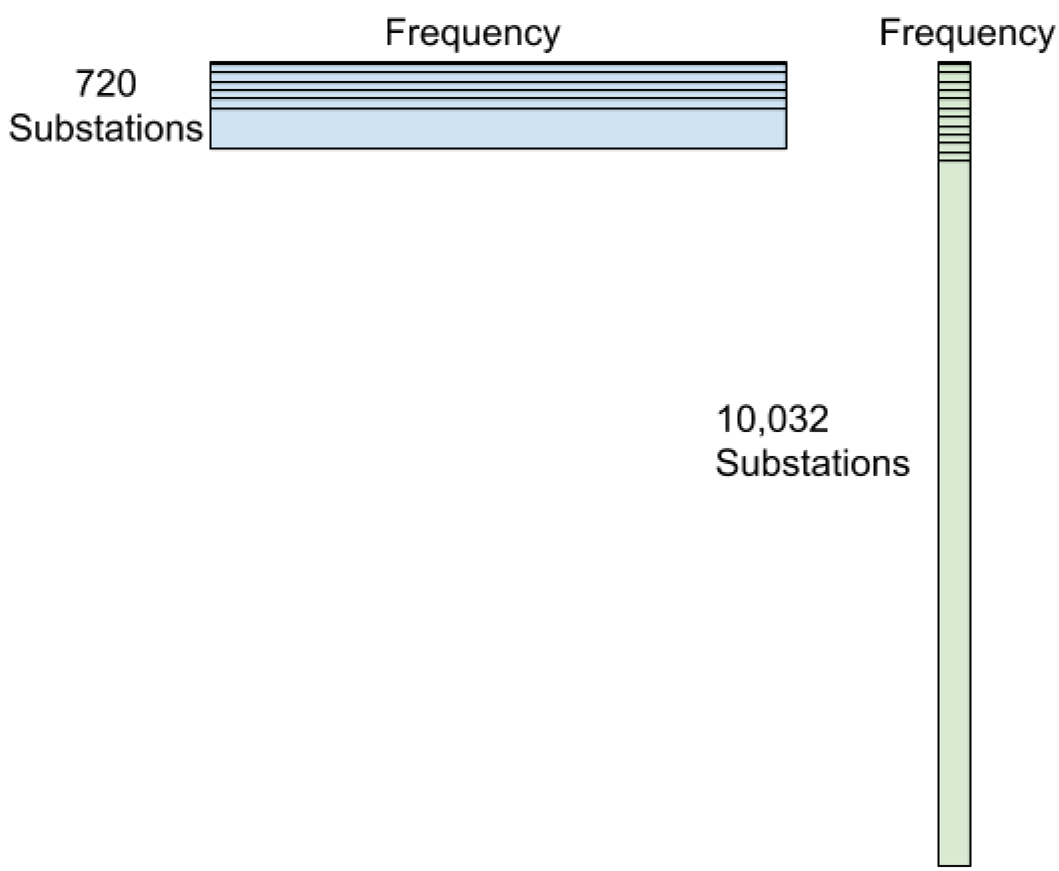

Figure 5: *Some different Substation-Frequency Data Planes*

As the number of substations grows the number of FPGAs needed to process a given bandwidth grows. Both approaches solve this problem by reducing the bandwidth to an FPGA as the number of substations increases, by up to a factor 384 for 10,032 substations. For the Atomic COTS each FPGA implements a filterbank on a single coarse channel for a number of substations. This is similar to the operation of VFC in the Separate Filterbank approach. The fine channels are then redistributed across FPGAs via the P4 switch. This is similar to the VFC to VSP data redistribution. Analysis for 720 to 10032 substations is present in the documentation of Atomic COTS.

For the Separate Filterbank approach only the 1024 and 2048 substation case had been analysed. However, it uses the same technique of spreading frequency channels across multiple FPGAs as Atomic COTS. For substations the total data rate decreases as the number of substations increases. For example, in the 10032 substation case one coarse channel is processed and the data rate to a VSP is 5% of that compared to a 512 station correlator. Each VSP has more than enough internal capacity to distribute the data evenly across FPGAs for processing. However a full analysis has not been done and full implementation is not guaranteed. In particular, the design was for a fixed routing structure and the design would need to change to add the high degree of routing flexibility needed. This will add to the firmware development cost.

For both designs some inefficiencies arise because of this flexibility. The size of the group of FPGAs is generally not a submultiple of the number of FPGAs. There will remain a set of FPGAs that do not form a group sufficient to process a full coarse channel. Possibly a part channel can be processed. For an arbitrary number of substations the resulting bandwidth, to match compute resources, is not usually a multiple of 5.4kHz (the standard frequency resolution). Further, it is desirable to have the total number of fine channels be a multiple of the number of FPGAs imposing further constraints on “optimal” configurations. Building a correlator with sufficient spare capacity would solve this problem.

Note, as with stations, subarrays within a set of identical substations are separable. That once a set of substations, all with the same bandwidth is instantiated the set can be distributed across any number of subarrays.

Substations - Zoom Modes

There is no explicit requirement for zoom bands with substations. Also zoom bands are limited by the maximum data rate to SDP. Even for standard observation it is not possible to transmit full bandwidth to SDP. As the number of substations increases in a subarray the number of visibilities per frequency channel increases quadratically. Thus as the total number of frequency channels possible drops rapidly. Each zoom band generates 17,280 frequency channels and for substation this is three times the rate of station data. This means not even a single zoom band could be supported with 1024 substations by the bandwidth to SDP.

Mixed Subarray

The correlator requirements do not preclude a mix of subarrays, these are subarrays of different bandwidths and/or number of stations/substations. Beamforming still occurs on subarrays of stations but not on a subarray of substations. With stations the input bandwidth is always 384 channels and the frequency/station diagram is always like Figure 3 for a maximal array. A maximal "array" uses all available compute resources leaving none for any other subarray.

Examples are a 300MHz @ 512 stations, a 150MHz @ 720 substation subarray, a 50MHz @ 1248 substation subarray, and 1 coarse channel @ 10032 substation subarray. A maximal array for a given number of substations (or stations) has a bandwidth that can be processed. Here the bandwidth is used to define the array. So a 50MHz maximal array processes data for 1248 substations. If bandwidths that included a fraction of a coarse channel are allowed there are 9520 possible maximal arrays (512 to 10032 substations). Each possible maximal array for a given bandwidth has its own, and different, description in terms of a frequency/station diagram (Figure 5).

But although a maximal array uses all the correlator compute resources a subarray derived from a maximal array does not. Thus it is possible to mix a subarray(s) from one maximal array with a subarray(s) from a different maximal array. An example of mixed subarrays is a 256 station@300MHz subarray and a 700 substations@75MHz subarray.

A given set of subarrays may need the routing for a number of different maximal arrays. The routing for each maximal array is unique to that maximal array. The data from a single LFAA station may belong to a number of subarrays and different routing strategies may be needed for different input channels. Mixed subarrays will demand considerable flexibility in the data routing system. This flexibility is inherent in the Atomic COTS but will need to be added to the design of routing firmware for the Separate Filterbank Approach. The limiting case is the largest maximal array where data from a single coarse channel for 10032 substations must be distributed across all FPGAs. This is very different to standard observing with 512 stations.

To determine if all Use Cases can be implemented all that is needed is for each FPGA to be able to process the subarrays for each maximal array. Without any testing of the correlator, the approach Low.CBF has taken is very flexible and this is possible. The conclusion is that both approaches Separate Filterbank and Atomic COTS, can handle subarrays from multiple maximal arrays (multiple full maximal arrays is not possible). The bounds to the size of subarrays is defined by requirement SKA1-SYS_REQ-3368, which has $Ns^2Ncc \le Kcorr$ and puts a limit on the compute required of the correlator (Ns number of substations, Ncc number of channels, and Kcorr a constant).

Note, each maximal array requires its own data routing implementations and correlator processing. Each maximal array that is allowed will require extra control software, firmware development, and the associated verification.

Verifying Use Cases

The following conclusions apply to both the Separate Filterbank approach (ECP-200028 Enhanced Low Correlator/Beamformer architecture for engineering subarrays) and Atomic COTS (ECP-200039).

The separability of correlator and beamformer function means once a Use Case for beamformers is established it is valid for all modes of the correlator. Similarly correlator modes do not affect the beamformer.

Beamformer

For the beamformer the total input data is 512 stations and 384 LFAA channels. The effect of subarrays is to reduce the number of possible inputs to a beam and station beams may reduce the available bandwidth. At most the effect is to reduce the workload. So the beamformer works just as well with any number of subarrays or stationbeams.

Correlator - Stations

This also applies to the correlator processing 512 stations: any combination of station beams, standard or zoom bands, and subarrays is allowed. Both approaches share the same correlator design.

Correlator - Substations

The greatest difficulty for the systems is with substation as each number of substation (across all subarrays of a given substation type) requires a different data routing and memory storage solution. Considering the case of a single subarray, each time the number of substations in the subarray changes a different routing solution is needed. Internally the correlator must be reconfigured to generate the larger visibility set. Also, in general, it is not always possible to fully match the substation and bandwidth that meet requirements accurately to the FPGA resources.

Mixed Subarray

Mixing of subarrays from differing maximal arrays is possible. Initial analysis of the resulting complex Use Cases indicates both the Separate Filterbank and Atomic COTS approach can implement mixed subarray. However, the added work for each such Use Case is not quantifiable at present.