NRC·CNRC

SKA1 Mid Telescope

Scan Configuration and Interface Mid.CBF to SDP

Sonja Vrcic

17. March, 2020

Architecture CoP

National Research Council Canada
Conseil national de recherches Canada

Canada

Overview

- Interface Mid.CBF to SDP
- Mid.CBF
  - Correlation, visibilities
  - Scan configuration
  - Some aspects of Mid.CBF design
- Who selects output links where to transmit visibilities
- Current Status: ICD and Implementation
- Roll-out Plan (for output links)

Interface CSP to SDP

- 300-000000-002 SKA1 MID SDP-CSP Interface Control Document
- Latest version: Revision 06, 2019-02-15, available in eB, not signed.

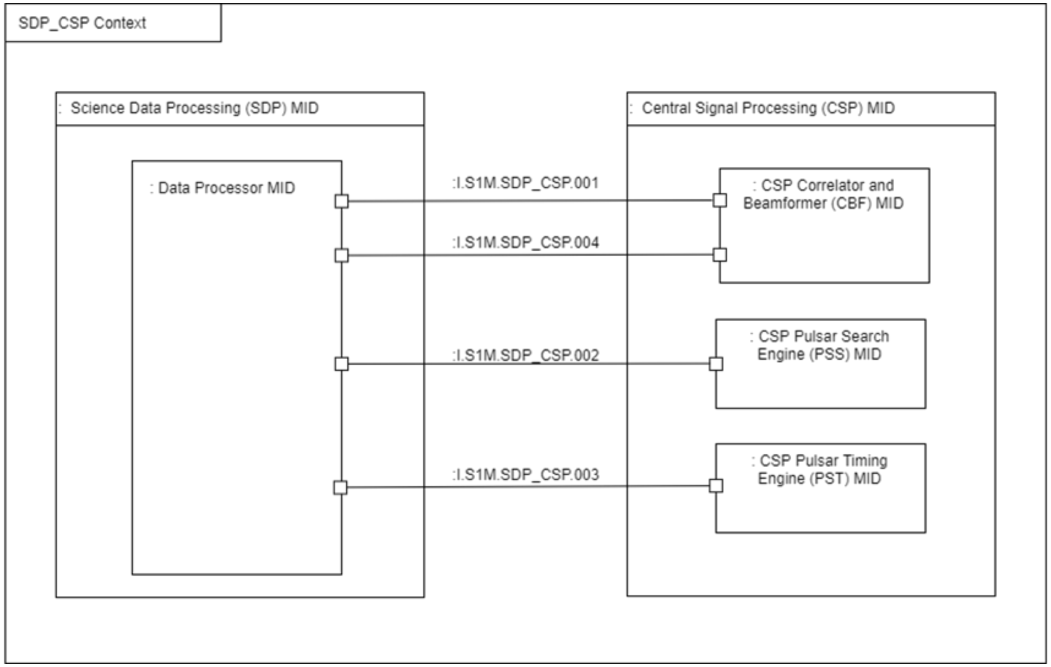

Interface ID Number	Interface Description
I.S1M.SDP_CSP.001	SKA1-Mid SDP_CSP Visibility Data Interface
I.S1M.SDP_CSP.002	SKA1-Mid SDP_CSP Pulsar Search Data Interface
I.S1M.SDP_CSP.003	SKA1-Mid SDP_CSP Pulsar Timing Data Interface
I.S1M.SDP_CSP.004	SKA1-Mid SDP_CSP Transient Buffer Interface

The transfer of the data routing information supplied by SDP to CSP via TM, is not described in this ICD.
Refer to:
ICD Mid SDP to TM and
ICD Mid CSP to TM.

ICD SDP to CSP – Physical Context

- The physical transport between the CSP and the SDP is provided by the SaDT (Signal and Data Transport).
- SADT provides the H/W infrastructure as well as end-to-end transport, carrier protocol, and data integrity for transport from CSP-Mid to SDP-Mid.
- Physical interfaces I.S1M.SADT_CSP.005 and I.S1M.SADT_SDP.003 are used to transport Visibility, Transient Buffer and VLBI data.
- I.S1M.SADT_CSP.005 consist of 80 x 100 Gpbs links.

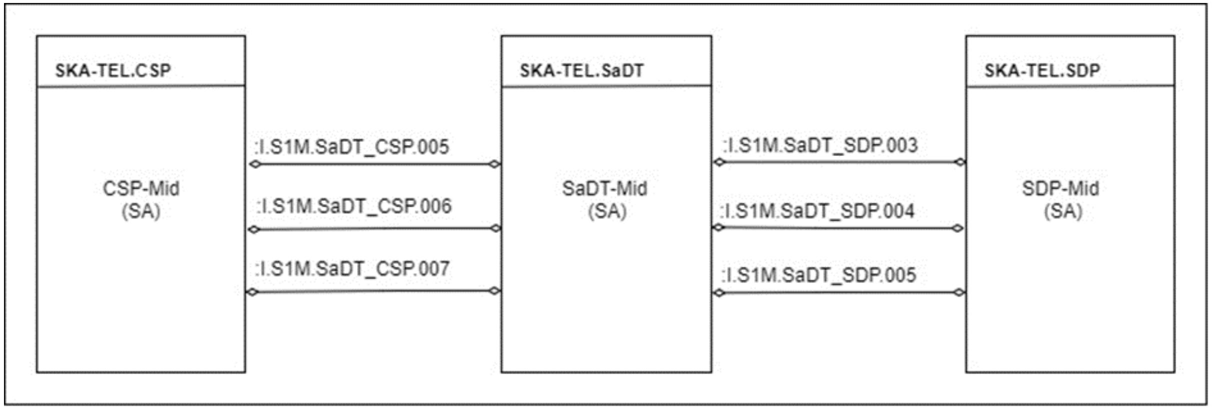

SDP – CSP Visibility Data Interfaces I.S1M.SDP_CSP.001

- This interface will carry real-time streaming visibility data unidirectionally from the CSP to the SDP instances.
- A visibility is defined as the complex cross-correlation produced between two individual receiving elements (real or beamformed) for a single spectral channel.
- For the purpose of this interface the definition of a visibility includes autocorrelations.

Visibility Data Frame

- The visibility data frame is a multi-dimensional array of visibilities with the following axes:
  - Frequency Channels
  - Phase bins (in SKA1 CSP_Mid must be 1)
  - Baselines
  - Polarisation products
- The dimensions of the visibility data frame are determined on a per scan basis and remain fixed for the duration of a scan.
- A visibility data frame for each channel is transmitted on regular cadence as determined by the **correlator integration time**.
  - Configurable on per scan basis.
  - Range: 0.14s - 1.4s. Allowed values are integer multiples of the minimum integration time.

Scan – informal definition

- Scan is a period during which the Correlator-Beamformer configuration does not change.

Protocol Stack

- CDR version of the ICD defines SPEAD as the application layer protocol.
- NZApp and CIPA team are investigating use of RDMA for this interface, current proposal is to use a custom-defined protocol and UDP to transmit visibilities over 100Gbps link.
  - Work in progress !

Mid.CBF Implementation - Correlation

- Input received from receptors is channelized and correlated.
- Coarse channelization produces Frequency Slices, 200MHz of bandwidth.
- Each Frequency Slice (FS) can be forwarded to one or more Frequency Slice Processors (FSPs) for further processing.
- Each FSP can be configured to perform correlation or beamforming (for PSS, PST, or VLBI).
- Each FSP configured for correlation can produce the full visibility set for 200MHz of bandwidth.
- Each FSP configured for correlation can correlate 200MHz of bandwidth for all SKA1 MID dishes (i.e. baselines) and all sub-arrays.

Frequency Slice Processor (FSP)

Correlation, Number of Channels, Channel Averaging

- Each FSP configured for correlation produces **14880 fine channels** across the correlated bandwidth (Frequency Slice or Zoom Window).

  Frequency Slice bandwidth FSBW = 200 MHz

  Zoom Window bandwidth ZWBW = FSBW / 2^n; n=0, 1, 2, …, 6

- Channels are evenly spaced in frequency.
- Mid.CBF provides 27 FSPs (26 + 1 for redundancy).
- Channel averaging can be used to reduce the amount of output data.
- As a part of a scan configuration, for each FSP, Mid.CBF receives the table that for each group of 744 channels (20 groups per FSP), indicates the channel averaging factor.
  - ❖Range 0, 1, 2, 3, 4, 6, 8

    0 means do not send channels to SDP

    1 means no averaging

    2 means average two adjacent channels

    3 means average three adjacent channels, and so on.

Frequency Slice Processor (FSP)
Correlation and access to output links

- A Frequency Slice Processor (FSP) consists of 20 TALON-DX boards, each board is equipped with STRATIX-10 FPGA.
- In an FSP configured for correlation, each **FPGA**:
  - Performs correlation for 14880/20 = 744 fine channels.
  - Performs averaging in time and frequency.
  - For each channel (after averaging) generates a Visibility Data Frame with the full set of visibilities for all the baselines in the sub-array.
  - Transmits Visibility Data Frames.
    - Each Visibility Data Frame can be transmitted over one of the 4 output links it is connected to.
    - Each FPGA can send ~25Gbps on each output link.

Note:

Overall, each FSP can transmit Visibility Data over all 80 links, but each Visibility Data Frame (visibilities for a single channel) can be transmitted over one of 4 output links.

If Output Links are selected by OET or SDP

- In the case that selection of the output links is performed by Observation Execution Tool (or SDP), the output link for each channel should be specified as a part of scan configuration.

The following should be provided for each FSP :

- Channel averaging factor per group of 744 channels (array of 20 integers – this is already in place).
- Output link per channel, for each of 14880 fine channels (less when averaging in frequency is performed).
- When channel averaging factor is >1, the result of averaging is assigned the channel ID of the 1st channel in the averaged group.

If output links are selected by Mid.CBF

- If preferred, Mid.CBF could implement fixed distribution of visibilities for each subarray evenly across the output links.
  - For example: ¼ of channels generated by one FPGA transmitted over one of the 4 output links FPGA has access to.
  - For each averaging factor, distribution over output links can be known in advance.
  - Calculation of the link utilization is simple and straight-forward, even in the presence of multiple independently controlled sub-arrays.
  - When averaging in frequency is used, the load may not be evenly spread across all the links (e.g. if fine channels are required for some groups and averaging is performed for others).

- If fixed distribution is not required, Mid.CBF might be able to re-distribute channels if one of the links failed (although this should be rare event).

- General observations:
  - Observation Preparation and Execution tools shall keep track of the total Mid.CBF output and shall not send to Mid.CBF requests that exceed capacity of the output links.
  - Mid.CBF shall be able to detect and reject invalid requests.

FPGA	channels	LINK	channels	LINK	channels	LINK	channels	LINK
1	0 - 185	1	186 - 371	2	372 - 557	3	558 - 743	4
2	744 - 929	5	923 - 1115	6	1116 - 1301	7	1302 - 1487	8
3	1488 - 1673	9	1674 - 1859	10	1860 - 2045	11	2046 - 2231	12
4	2231 - 2417	13	2418 - 2603	14	2604 - 2789	15	2790 - 2975	16
5		17		18		19		20
6		21		22		23		24
7		25		26		27		27
8		29		30		[illegible]		32
9		33		34		35		36
10		37		[illegible]		39		40
11		41		42		42		44
12		45		[illegible]		47		48
13		49		50		51		52
14		53		[illegible]		55		56
15		57		58		59		60
16		61		62		63		64
17		65		66		67		68
18		69		70		71		72
19		73		74		75		76
20		77		78		79		80

An incomplete table for distribution of the channels generated on a single FSP over the output links

Conclusion

- In general, Mid.CBF Team is willing to implement any solution that does not require unreasonable effort to implement.
- But before we update the ICDs and implementation it would be good to carefully consider all implications.
- The following 3 slides describe the current status of the ICDs and MVP implementation.

NRC·CNRC

Current Status: ICD and MVP Implementation

Mid.CBF distributes the output of each subarray across the output links, and for each channel, publishes the link used, as follows:

- Scan ID
  - For each FSP:
    - FSP ID
    - FS ID
    - For each channel:
      - Channel ID
      - Channel BW
      - Channel Central Frequency
      - CBF Output Link

- Total number of channels (for 28 FSPs) can be up to 14,880 x 28 = 416,640.

Format:

https://confluence.skatelescope.org/display/SE/CSP_Mid+Scan+Configuration+for+Correlation

Suggestion:

JSON is verbose, we should consider more compact format, for example TANGO pipes.

Current Status: ICD and MVP Implementation

SDP responds to Mid.CBF by publishing destination addresses for each channel:

- Scan ID
  - For Each FSP:
    - FSP ID
    - For each channel:
      - Channel ID
      - Phase Bin ID (currently must be 1)
      - SDP MAC Address
      - SDP IP Address
      - SDP Port

NRC·CNRC

Current Status: ICD and MVP Implementation Message flow (simplified)

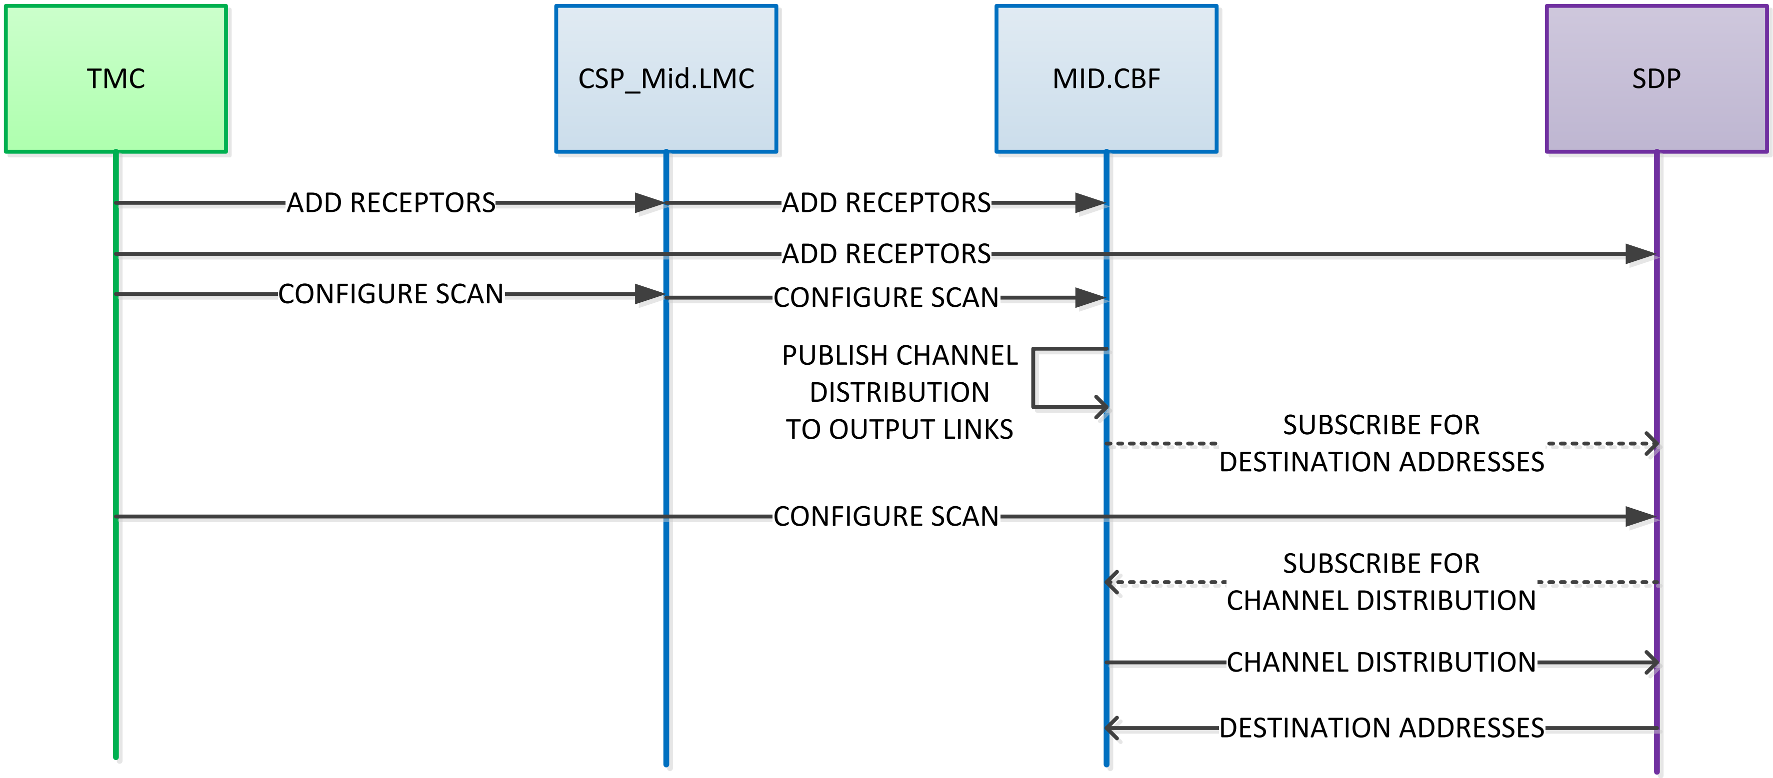

Roll-out Plan

❖AA 00.5 and AA1 (TDC)

- 4 antennas / 8 antennas
- Correlated bandwidth 800 MHz
- One Mid.CBF output link is sufficient to handle all visibilities.

❖AA2, AA3 and AA4 (Mid.CBF)

- Output links can be rolled-out in increments of 20.
  For example: 20, 40, 80

Thank you

Sonja Vrcic
Tel: 250-497-2309
Sonja.Vrcic@nrc-cnrc.gc.ca
www.nrc-cnrc.gc.ca