

Timing in the Streaming Interface

| | Organisatie / Organization | Datum / Date |
|---|----------------------------|--------------|
| Auteur(s) / Author(s): Kooistra | ASTRON | |
| Controle / Checked: Schoonderbeek | ASTRON | |
| Goedkeuring / Approval: Schoonderbeek | ASTRON | |
| Autorisatie / Authorisation: Handtekening / Signature Gunst | ASTRON | |

© ASTRON 2011

All rights are reserved. Reproduction in whole or in part is prohibited without written consent of the copyright owner.

Distribution list:

| Group: | Others: |
|---|--|
| Gijs Schoonderbeek Andre Gunst Dion Kant Daniel van der Schuur Raj Rajan Thilak | Harro Verkouter (JIVE) Jonathan Hargreaves (JIVE) Salvatore Pirruccio (JIVE) |

Document history:

| Revision | Date | Author | Modification / Change |
|----------|------------|-------------|-----------------------|
| 0.1 | 2011-03-15 | E. Kooistra | Creation |
| | | | |
| | | | |
| | | | |

Table of contents:

| | | |
|-------|----------------------------------|---|
| 1 | Introduction..... | 4 |
| 1.1 | Purpose | 4 |
| 1.2 | Overview..... | 4 |
| 1.2.1 | Data path | 4 |
| 1.2.2 | Sync pulse | 5 |
| 1.2.3 | Block sequence number | 5 |
| 2 | Timing definitions | 6 |
| 2.1 | Streaming interface record | 6 |
| 2.2 | Sync pulse | 6 |
| 2.3 | BSN | 7 |
| 3 | Related timing aspects..... | 8 |
| 3.1 | Wideband processing..... | 8 |
| 3.2 | Data path control | 8 |
| 3.2.1 | Restarting a signal path | 8 |

Terminology:

| | |
|-------------|---|
| ACC | Accumulator |
| ADC | Analogue to Digital Converter |
| BF | Beam Former |
| DP | Data Path, the whole DSP processing for all signal paths |
| FFT | Fast Fourier Transform |
| LCU | Local Control Unit (control computer) |
| MM | Memory Mapped |
| PFB | Poly-phase Filter Bank |
| PHY | Physical layer |
| PPS | Pulse Per Second |
| Signal path | One ADC output signal |
| SOSI | Source Out Sink In |
| ST | Streaming |
| UNB | https://svn.astron.nl/UniBoard_FP7/UniBoard/ |
| UTC | Coordinated Universal Time |
| X | Correlator |

References:

1. "Specification for module interfaces using VHDL records", ASTRON-RP-380, E. Kooistra
2. "Practical information on the LOFAR station timing", LOFAR-ASTRON-RPT-312, E. Kooistra
3. "RSP Firmware Design Description", LOFAR-ASTRON-SDD-018, W. Poiesz, E. Kooistra, W. Lubberhuizen
4. "Sync Pulse Usage in CASPER DSP Blocks", 2008, H. Chen, P. McMahon, A. Parsons
5. "Avalon Interface Specifications", mnl_avalon_spec.pdf, www.altera.com

1 Introduction

1.1 Purpose

The document defines the sync pulse field and the block sequence number field that are added to the streaming interface [1] to provide sample accurate timing throughout the entire DSP chain from ADC to final data path (DP) output. The definitions in this document are based on the ideas from LOFAR [2] and from CASPER [4].

1.2 Overview

1.2.1 Data path

Figure 1 shows a typical data path processing system from multiple input signal paths to output auto powers and cross powers per frequency band and per beam direction. Figure 1 has the focus on the down-sampling functions (indicated by the downward arrows), because these are important for the data path timing.

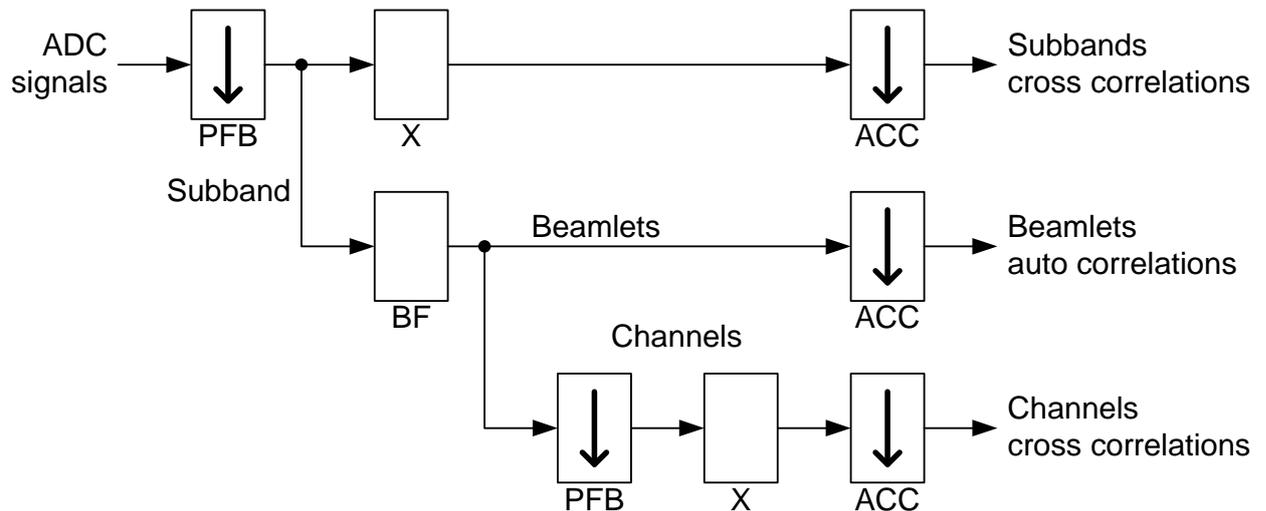


Figure 1: Typical data path processing

1.2.2 Sync pulse

The sync pulse is a periodic signal that aligns the streaming data to a starting time t_0 . By indicating t_0 with a central timing pulse multiple signal paths have the same t_0 . By using a central pulse that is aligned to the UTC top of second all signal paths are aligned to UTC. The central pulse is typically a PPS signal. Therefore a control computer must inform the data path processing which single PPS pulse in time to use for t_0 . Typically the data path processing will use the first PPS pulse after it got enabled.

The period of the sync pulse is completely independent of the PPS period. Instead the sync pulse period depends on the largest time-variant interval in the data path processing. The size of the time-variant interval depends on the time variant processing steps. Time variant processing occurs where data samples get processed in blocks of samples, because then it depends on what sample the block starts. This is similar to the time variant aspects of clock division. Functional causes for block processing are due to down-sampling and occur at:

- ADC time series data filter bank that outputs down-sampled frequency subbands
- Subband filter bank that outputs down-sampled frequency channels for the correlator
- Number of accumulated samples for the integration time of the statistics

Time-variant blocks can also occur for implementation reasons and these may increase the sync period:

- Transpose (corner turner, reorder) functions that operate on blocks of samples
- Framing blocks of data to be able to transport them via PHY links or store them in memory

In [4] it is stated that the sync pulse interval also depends on the number of taps of the FIR filter structure in the Poly-phase Filter Bank. This is probably not true, because this filtering is a time invariant operation, so the number of taps does not affect the sync pulse interval.

1.2.3 Block sequence number

It is not necessary to count the input ADC samples at the ADC sample rate. It is sufficient to count blocks of data using a Block Sequence Number (BSN), because the down-sample factors are known. At t_0 the BSN defaults to 0 unless it is initialized by the control computer and then the BSN increments at every sync pulse. The BSN field has sufficient width to count blocks for years.

The resolution of the BSN is equal to the total down-sample factor. Hence the sync pulse period or resolution of the BSN field depends on the kind of data that it refers to in the data path, e.g. subband data, channel data or integrated data. However at all stages $BSN = 0$ aligns to t_0 with an absolute accuracy equal to the original sample period of the ADC. This original ADC sample period value is fixed and known, so it does not need to be passed on with the data.

Together the UTC time at t_0 , the known ADC sample period, the sync pulse period and the corresponding BSN uniquely define the sample time at each stage of the DP processing.

2 Timing definitions

2.1 Streaming interface record

The streaming interface is defined by a VHDL record that contains all the signals that can be needed for handling data transfer between a source and a sink. This record is called SOSI for Source-Out-Sink-In and is similar to the streaming interface definitions for the Altera Avalon interface [5]. For data path processing the SOSI record is now extended with a 'sync' field and a 'bsn' field as shown below. It is not mandatory to use all fields.

```

TYPE t_dp_sosi IS RECORD -- Source Out or Sink In
  sync      : STD_LOGIC;
  bsn       : STD_LOGIC_VECTOR(c_dp_stream_bsn_w-1 DOWNTO 0); -- 48 bit
  data      : STD_LOGIC_VECTOR(c_dp_stream_data_w-1 DOWNTO 0);
  valid     : STD_LOGIC;
  sop       : STD_LOGIC;
  eop       : STD_LOGIC;
  empty     : STD_LOGIC_VECTOR(c_dp_stream_empty_w-1 DOWNTO 0);
  channel   : STD_LOGIC_VECTOR(c_dp_stream_channel_w-1 DOWNTO 0);
  err       : STD_LOGIC_VECTOR(c_dp_stream_error_w-1 DOWNTO 0);
END RECORD;

```

2.2 Sync pulse

The sync pulse starts or restarts the processing of a DP function. If the sync pulse does not occur or if it occurs at the expected cycle, then the DP functions just continues. If the sync pulse occurs at an unexpected cycle then the DP function restarts. This scheme allows the whole DP processing to continue with only one initial sync pulse and to be realigned at any time. However, typically the sync pulse is applied periodically at the minimum time-variant interval period. Each DP function must propagate the sync pulse with the same latency as that it imposes on the output data, this in order to maintain synchronization throughout the entire DP processing chain.

Figure 2 show the timing of the sync pulse with respect to the data valid. The valid signal indicates whether a clock cycle carries valid data or not. The signals a, b, c and d are all correct sync pulse signals.

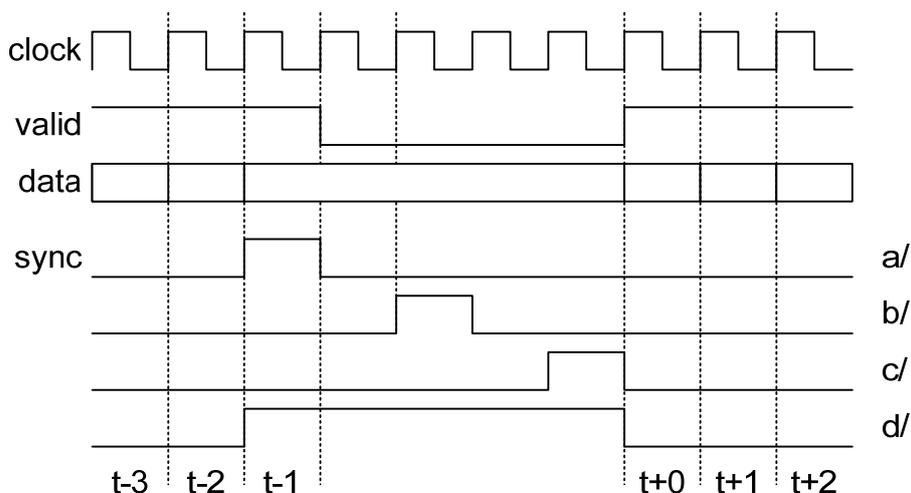


Figure 2: Timing of data valid and data sync [3]

2.3 BSN

The BSN depends on the stage of the DP processing. Based on section 1.2.2 there typically are the following kinds of block sequence numbers:

- Subband sequence number
- Channel sequence number
- Statistics sequence number

The BSN increments per sync pulse period, even if the sync pulse does not occur. Figure 3 shows the BSN timing with respect to the sync pulse.

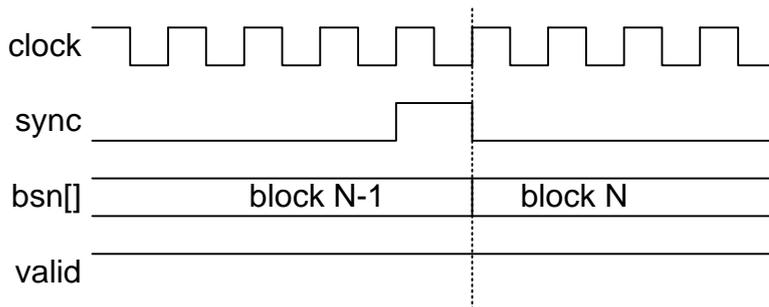


Figure 3: BSN timing

For packetized data the BSN gets transported in a header field along with the data. The receiver knows the type of data, so it can also relate the BSN to t_0 .

3 Related timing aspects

3.1 Wideband processing

For wideband signal processing the ADC sample frequency can be an integer factor P larger than the data path (DP) processing clock frequency. This factor P is a division factor per signal path, so it can cause a timing uncertainty of up to $P-1$ sample periods between these signals paths, if the division is not synchronized. The assumption is that the division is synchronized somehow (by using a central reference). The sync pulse period in the DP processing clock domain will be a factor P less than the down-sample factor of the ADC samples. For example, if the ADC sample frequency is 800 Msps and the DP processing clock frequency is 200 MHz then $P = 4$, and for a real input PFB with a down-sample factor of 1024 the sync pulse period in the 200 MHz clock domain is then 256 cycles.

3.2 Data path control

The communication between the control computer and the data path processing logic has some timing uncertainty, in the order of < 1 ms. In LOFAR this timing uncertainty was brought back again to ADC sample period accuracy by letting the local control computer (LCU) send the control data in one second and then let the data path apply it in the next second when the internal PPS occurred. This internal PPS accounts for the data path latency and with some margin around the external PPS the LCU can control the data path processing with maximum timing accuracy. However this dependency on the PPS is not necessary, as argued in section 1.2.2, and can even be awkward.

Therefore instead the LCU – DP control can be managed using the BSN. The LCU roughly knows the current BSN number in the data path, e.g. because it can read it out from the DP. For each control message to the DP the LCU adds the BSN number that states when the corresponding task in the DP should apply the control data. The task in the DP has a double buffer (similar as in LOFAR). The LCU control data is kept in one part of the buffer and when the sync for the corresponding BSN occurs, then the buffer swaps and the control data gets applied. If necessary the buffer could be made deeper than a double buffer to support more sophisticated scheduling at the expense of more logic. For control data that can be applied immediately a single buffer is enough. This scheme of scheduling the LCU – DP control based on the BSN allows any scheduling interval per DP function with a resolution of the sync pulse period and independent of the PPS.

3.2.1 Restarting a signal path

A signal path in the entire DP processing can be restarted when the sync pulse coincides again with the external PPS. The LCU can then initialize the BSN for that signal path to the current BSN value of the other signal paths in the DP processing.