ATLAS L1Calo Pre-processor compressed S-Link data formats

D.P.C. Sankey

Rutherford Appleton Laboratory, Didcot, Oxon., OX11 0QX, UK

Version 1.5, March 14, 2008

Introduction

The overall structure of the data formats for the ATLAS Level-1 Calorimeter Trigger is described in the project specification document for the Read-out Driver[1]. This note solely expands on the details for the Pre-processor compressed and super-compressed algorithms, starting with the preliminary compressed coding algorithm (version 1.00) dating from July 2006, up to version 1.03.
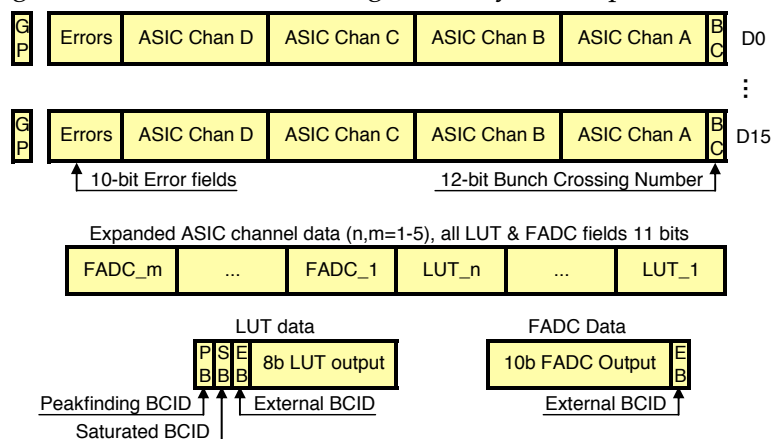

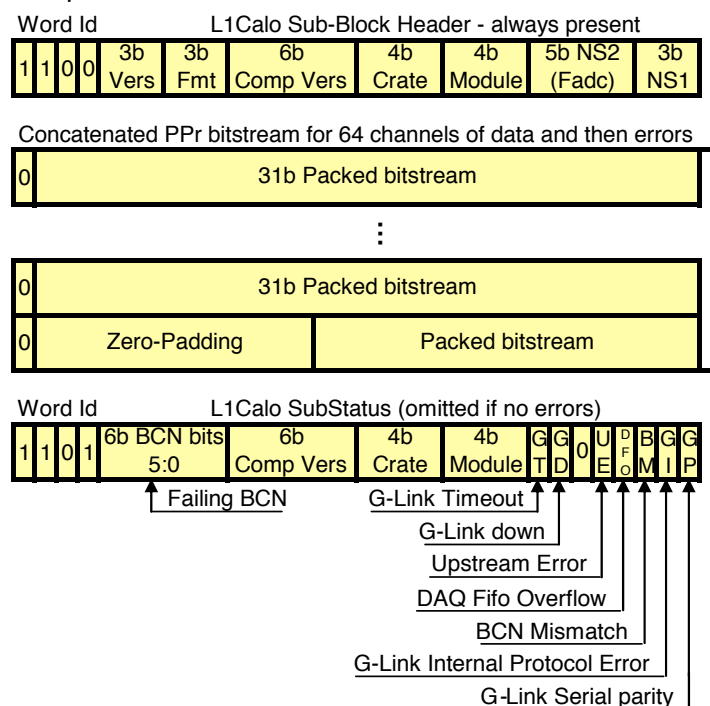Figure 1   Overview of input G-Link format

Figure 2   Overview of output S-Link format

The general structure of the input Pre-processor G-Link format is shown in figure 1 and of the compressed S-Link output format in figure 2; for both there is a single block of data for each Pre-processor, corresponding to 64 channels of data and associated error information from the 16 MCM's on each Pre-processor. For the output format the incoming time ordered data arriving in parallel are transformed into a packed serial bit stream in channel order. The bit stream consists of the channel

data (ASIC channels A from pins D0 to D15 in order, then similarly ASIC channels B through D) immediately followed by the errors.

The data for each ASIC channel (consisting of 5 samples of FADC data and 1 sample of LUT data, corresponding to 66 bits of input data per channel) are encoded in a formatted block and the output bit stream for the data from all the channels consists of a concatenated string of these blocks.

In order to identify the version of the algorithm used and the compression mode, the major part of the version number of the coding algorithm is included in the sub-block header in the data stream as the field 'Vers' with the minor part as 'Comp Vers' and the format mode, compressed (2) or super-compressed (3), is included as the field 'Fmt'.

There are an additional three parameters for the algorithm, namely the threshold below which data are suppressed in super-compressed mode, the time offset of the LUT data with respect to the start of the FADC data and the nominal FADC baseline lower bound. In the ROD's these are all read directly from the Input Channel Compression Control Register, along with the compression mode, and are required to be constant during a run.

Starting with L1Calo ROD fragment event format minor version[2] number 0x1002 these are encoded on output in the User Header as illustrated in figure 3. The FADC baseline lower bound is the field 'Lower Bound' and the time offset of the LUT data with respect to the start of the FADC data is the difference between 'Triggered FADC slice number' and 'Triggered LUT slice number'.
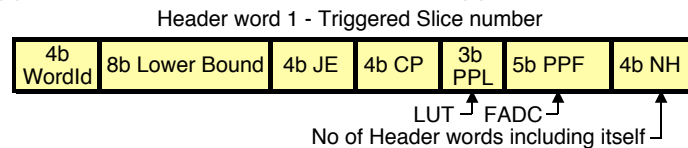

Header word 1 - Triggered Slice number

Figure 3   L1Calo User Header for event format minor version 0x1002

Version 1.03 of the Pre-processor compressed and super-compressed S-Link coding algorithms

Version 1.03 defines compressed and super-compressed formats for the channel data along with a format for the error block. In compressed mode all channels are present; in super-compressed mode channels with non-extant LUT data and all FADC data below the threshold are suppressed.

Channel data

In super-compressed mode the block for each ASIC channel begins with a single bit indicating the presence ('1') or suppression ('0') of the following data block. In compressed mode all channels are present so this bit is unnecessary. There are then seven different output format types to encode the block of data for the channel. The data for each channel are encoded in the smallest of these format types capable of fully representing the data.

The block if present consists of a header, between 4 and 8 bits long, then LUT data and external BCID bits if non-zero and finally the FADC samples. The FADC samples are in chronological order except that the first sample and the smallest are interchanged and are encoded as the smallest sample, then the differences with respect to this of the remainder.

The shortest two format types have a header 4 bits long. These two formats are targeted at noise and require that the LUT data and external BCID bits are all zero. The format of the header is the sample number of the minimum FADC sample (in the range 0-4), added to either 0 for the shortest format (format 0) or 5 for the second shortest (format 1).

There then follow the FADC samples as described above. The minimum value must be in the range FADC baseline lower bound to lower bound plus 15 and is encoded as the offset from the lower bound in a 4 bit field. The remaining are encoded as each sample less the above minimum value, in 4 fixed

width fields. The FADC field width for format 0 is 2 bits and for format 1 3 bits, illustrated in figure 4, giving a total length of 16 bits for format 0 and 20 bits for format 1.

Format 0

2b	2b	2b	2b	4b	4b
FADC samples				Min	Header

Format 1

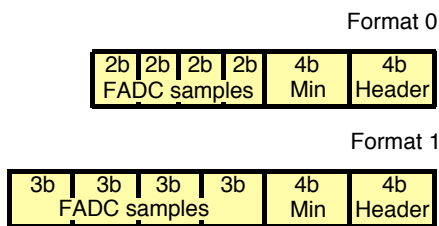| 3b	3b	3b	3b	4b	4b
FADC samples				Min	Header

Figure 4   Compressed formats 0 and 1

For the remaining 4 normal formats there are an additional 2 bits to the header. The first 4 bits encode the sample number of the minimum FADC sample as before, now added to 10, giving a value in the range 10 to 14. The next 2 bits encode as '00' for format 2 through to '11' for format 5.

For format 2 there is then one final bit to the header indicating the presence or otherwise of LUT data, set to '0' indicating no following LUT data or '1' to indicate that the subsequent 3 bits encode the bottom 3 bits of the LUT sample. The assumption is that no external BCID bits are set and the presence of non-zero LUT data implies that the peak-finding BCID bit is set.

Format 2

4b	4b	4b	4b	4b	3b	1 + 2 + 4b
FADC samples				Min	LUT*	Header
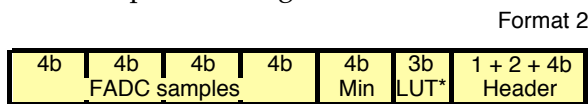
*) 3b LUT field suppressed if LUT data zero

Figure 5   Compressed format 2

The FADC samples are then encoded as for formats 0 and 1 as the minimum value and the differences with respect to this of the remainder ordered as above in 5 fields 4 bits wide. This gives a total length of 30 bits when there are extant LUT data, otherwise 27 bits, illustrated in figure 5.

In comparison to the fixed length formats above, the remaining three normal formats are variable length. For the FADC data they consist of either short fields where the value to be encoded is less than 16, or long fields, where the length of the field is given by the format number, this latter being determined by the largest value to be encoded. In the case of format 3 the values to be encoded all fit into 6 bits, in the case of format 4 8 bits and finally format 5 the full 10 bits.

Each of these formats starts with an 8 bit header consisting of 4 bits encoding the number of the minimum FADC sample added to 10, then 2 bits encoding the format number as given above, finally 2 bits, the lower of which indicates the presence or otherwise of LUT data, set to '0' indicating no following LUT data, the higher similarly the presence or otherwise of the external BCID bits for the FADC samples if these cannot be correctly derived from the external BCID bit on the LUT sample and those for the other samples set to zero.

If present, the LUT data are encoded as 11 bits exactly as in the input data as illustrated in figure 1. For the external BCID bits, these if not derivable as above are encoded as 5 bits in chronological order.

Finally the FADC samples are encoded as the smallest sample first, then the differences with respect to this of the remainder ordered as described above. In each case if the resultant value is less than 16 (in the case of the minimum the offset from FADC baseline lower bound), it is encoded as a short field, otherwise a long field (in the case of the minimum, without any offset).

The short field consists of 1 bit set to '0' followed by the least significant 4 bits of the resultant value. The long fields consist of 1 bit set to '1', then either 6 (format 3), 8 (format 4) or 10 bits (format 5).

These are illustrated in figure 6. They correspond to a minimum channel block length of 33 bits (albeit these data would encode as one of the preceding formats) and an absolute maximum length of 79 bits.
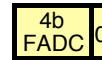
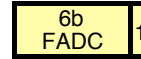| FADC samples…	5b BCID*	11b LUT*	2 + 2 + 4b Header

*) 11b LUT field suppressed if LUT data zero

*) 5b BCID field suppressed if external BCID bits derivable

FADC fields: Short field (LSB = 0)

4b FADC	0

Format 3 long field (LSB = 1)

6b FADC	1

Format 4 long field (LSB = 1)

8b FADC	1
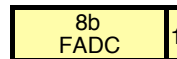
Format 5 long field (LSB = 1)
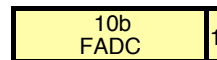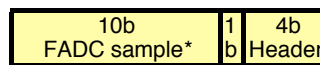
10b FADC	1

Figure 6   Compressed formats 3 to 5

The final format, format 6, illustrated in figure 7, encodes channels where the FADC data are all equal and all LUT and external BCID information are zero, as is the case, for example, for disabled channels.

In this case the 4 bit header is set to 15. If the FADC contents are zero there then follows a single bit set to '0', otherwise a single bit set to '1' and then the FADC value encoded in 10 bits.

Format 6

10b FADC sample*	1b	4b Header

*) 10b FADC field suppressed if sample zero

Figure 7   Compressed format 6

Error information

As illustrated in figure 1 the incoming G-Link data ends with a 10 bit error block on each pin corresponding to each MCM. The detailed format of this block is shown in figure 8.

G-Link serial parity
RemFPGA FIFO corrupt
Bunch Number mismatch
Event Number mismatch
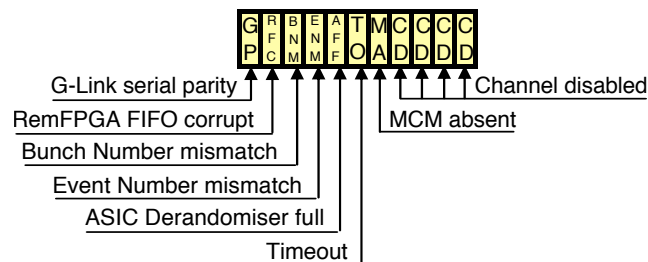ASIC Derandomiser full
Timeout
MCM absent
Channel disabled

Figure 8   Expanded error block, including parity bit

The first five of these bits constitute status information per MCM rather than actual errors. There is additionally the parity bit terminating the data block; the results of the parity test are treated as an additional error bit, giving six error bits per MCM.

The output error format begins with a two bit header; the first (status bit) is set to '1' if any status bits were non-zero, the second (error bit) similarly set to '1' if any error bits were non-zero.

If any of these bits are set there then follows a 16 bit map to indicate which MCM(s) contained the non-zero bits. Then for each MCM so indicated there is the 5 bit status block (if the overall status bit is set) and/or the 5+1 error block (if the overall error bit is set).

Substatus word

If the overall error bit above is set or any other error detected the substatus word as defined in figure 2 is written; any upstream errors are mapped onto the 'Upstream Error' bit, parity errors onto the 'G-Link Serial Parity' bit.

In the case of no errors the substatus word is suppressed.

Version 1.02 of the Pre-processor compressed S-Link coding algorithm

This version did not support the super-compressed format.

Version 1.01 of the Pre-processor compressed S-Link coding algorithm

Prior to resolution of the method for transmitting the time offset of the LUT data with respect to the FADC data and the nominal pedestal level set in the Pre-processors to the ROD, the assumption was that the LUT sample corresponds to the middle sample of FADC data and that the nominal pedestal is set to 20 FADC counts.

In the sub-block header in the data stream the number of FADC slices (NS2) and of LUT slices (NS1) were encoded as 4 bit fields.

More importantly, the treatment of external BCID bits was different. Rather than treating the external BCID bit in the LUT sample and that on the corresponding FADC sample as potentially different, the assumption was that these were the same. For the other external BCID bits, these if non-zero are encoded as 4 bits in chronological order, stepping over the bit from the FADC sample corresponding to the LUT sample.

Version 1.00 of the Pre-processor compressed S-Link coding algorithm

Version 1.00 was the first version to run in hardware.

It defined six different output formats to encode an input stream of 5 samples of FADC data and 1 sample of LUT data per channel. As for version 1, the data for each channel are encoded in the smallest of these formats capable of fully representing the data.

These six formats corresponded to formats 0 to 5 for version 1 above, with one important difference. This is that, rather than the FADC slices being in chronological order except that the first sample and the smallest are interchanged, the slices are encoded as the smallest sample first, then the differences with respect to this of the remainder in chronological order.

There was no format 6 to encode data where all the FADC samples are identical; rather data are encoded as format 0 or 3 depending on the sample value.

Finally there was no support for outputting the incoming error block from the Pre-processor or for producing the S-Link substatus word as this functionality was not immediately critical for the first hardware tests.

---

[1]   ATLAS Level-1 Calorimeter Trigger-Read-out Driver Project Specification version 1.1
       <http://hepwww.rl.ac.uk/Atlas-L1/Modules/ROD/ROD-spec-version1_1.pdf>

[2]   Section 5.6 in "The raw event format in the ATLAS Trigger & DAQ"
       <https://edms.cern.ch/document/445840>