

Data Formats for Phase-1 Upgrade

Draft version 0.4 (2 March, 2012)

Table of Contents

Table of Contents.....	1
Introduction.....	1
CP/JEP processor crates.....	1
Backplane transfer to CMX	1
Jet merger.....	1
Energy merger.....	3
EM, τ /hadron mergers.....	3
Fiber links from CMX to L1Topo	5
Fiber links from Muon CTP interface (MUCTPI) to L1Topo.....	5
Topological processor (L1Topo)	5
CTP interface	5

Introduction

This is a working document with the purpose of specifying data transfer between components in the Phase-1 upgrade of the L1Calo trigger and the topological processor (L1Topo). Working assumptions in this document include:

- Common merger modules (CMMs) in the CP/JEP processor subsystems are replaced by a CMX with all legacy interfaces, plus optical outputs to L1Topo.
- 160 Mbit/s backplane data transfer between CPMs/JEMs and CMX
- 6.4 Gbit/s serial data links from CMX to L1Topo inputs over 12-fiber optical cable bundles.

CP/JEP processor crates

Backplane transfer to CMX

Fourteen (14) CPMs or sixteen (16) JEMs in each processor crate send real-time results to two CMX modules at each end of the backplane. Each CPM/JEM to CMX link consists of 25 parallel point-to-point lines, for a total of 50 output lines per module. For Day-1 running, data are transmitted at 40 Mbit/s on 24 of the lines, with an odd parity bit on the 25th. In the strawman TP design, 24 lines carry data at 160 Mbit/s, and the 25th line carries a data clock for synchronization at the receiving end.

The total data payload for each of the two processor module outputs is therefore a 24-bit parallel word with four words per LHC bunch crossing, for a data payload of 96 bits per bunch crossing per CPM/JEM-to-CMX link..

Jet merger

Each JEM covers a core region of 4×8 (32) Jet elements (each element corresponding to 0.2×0.2 in eta and phi, respectively). This core region is surrounded by a larger “environment” of 7×11 jet elements duplicated from the core regions of

bordering JEMs. For each Jet element the JEM receives separate 9-bit electromagnetic and hadronic Et sums of four trigger towers from the PreProcessor. These sums are separately noise-suppressed by applying thresholds, and then summed together to produce a single 10-bit Et value with a least count of 1 GeV for use in the jet and energy-sum algorithms. The energy-sum and jet-finding algorithms are implemented in two separate, large FPGAs.

To identify a Jet ROI, a jet “core” is first found using a sliding window algorithm with a 2×2 window moving in steps of 1 jet element. A core 2×2 window has an Et sum that is a local maximum compared with its eight neighbors. Once a 2×2 local maximum has been identified, Et sums of jet windows comprising 2×2 , 3×3 and 4×4 jet elements (.4, .6 and .8 in eta-phi) that surround that core¹ are compared against defined thresholds. These Et sums are effectively truncated to 10 bits, with a least count of 1 GeV. The Day-1 jet algorithm accommodates up to eight independent jet definitions, each comprising a jet window size and minimum Et sum.

The output of the current jet algorithm is therefore 8 sets (one for each ROI “subregion”) of 8 threshold bits, each indicating whether a certain jet definition has been met. Two additional bits give the position of the ROI within that subregion, for a total of 10 bits per ROI.

A summary of the information available from Day-1 algorithms is as follows:

# of ROIs	Bits / ROI	Total bits / JEM
8 jet subregions	8 threshold bits + 2 position bits	80
Available bits at 160 Mbit/s:		96

The Day-1 output to the jet CMM is at 40 Mbit/s, so the data volume is 24 bits per BC, plus one bit of odd parity. For JEMs reporting only jets in the central portion of the calorimeter, these results are eight 3-bit multiplicities corresponding to the numbers of ROIs found that passed each of the eight jet threshold definitions:

b0-2	b3-5	b6-8	B9-11	B12-14	B15-17	b18-20	b21-23	bit 24
Thr 0	Thr 1	Thr 2	Thr 3	Thr 4	Thr 5	Thr 6	Thr 7	Parity

For JEMs that report both central and forward (FCAL) jets, this format is modified to 12 2-bit multiplicities (8 central jet definitions “CJ”, 4 FCAL jet definitions “FC”):

0-1	2-3	4-5	6-7	8-9	10-11	12-13	14-15	16-17	18-19	20-21	22-23	24
CJ0	CJ1	CJ2	CJ3	CJ4	CJ5	CJ6	CJ7	FC0	FC1	FC2	FC3	Pty

¹ In the case of the 2×2 jet window, it is simply the same as the core window. The 4×4 jet window selected is the one with the core window at its center. For 3×3 jet windows, there are four windows containing the 2×2 core. The Et sums of all four compared against the jet threshold, and a jet ROI is identified if at least one of these sums passes.

For the upgraded backplane speed of 160 Mbit/s, the available backplane bandwidth becomes 96 data bits per BC, transmitted over 24 single-ended lines. The twenty-fifth line is used to transmit a 40 MHz data clock encoded with a parity bit.

The expanded contents of the real-time jet output will include:

- Eight ‘presence’ bits, indicating in which subregions jet RoIs were identified
- Four 2-bit ‘fine position’ fields, one for each identified RoI
- Four 10-bit transverse energy sums for jet size 1
- Four 10-bit transverse energy sums for jet size 2

Possible high-speed JEM-CMX readout formats are:

Bits:	0-23								24
Word									
0	P 0-3		fp1	fp2	P 4-7		fp3	fp4	clk/ pty
1	Et jet1 size1		Et jet1 size2		Et jet2 size1		Et jet2 size2		
2									
3	Et jet3 size1		Et jet3 size2		Et jet4 size1		Et jet4 size2		

Or:

Bits:	0-23								24	
Word										
0	P 0-3					P 4-7		Et jet2 size2		clk/ pty
1	fp1	Et jet1 size1		Et jet1 size2		Et jet2 size1				
2	fp2	Et jet3 size1		Et jet3 size2		Et jet4 size1		Et jet4 size2		
3	fp3							fp4		

The latter format, while less regular than the first, has the advantage that all information on the first RoI has been completely received in the CMX by the second backplane word, and the second RoI is received by the third word. This may allow CMX →L1Topo latency to be reduced slightly.

Energy merger

EM, τ /hadron mergers

Each CPM covers a core region of 4×16 trigger towers ($.1 \times .1$ segmentation in eta and phi, with a least count of 1 GeV). The core region is internally partitioned into a 2×8 array of sixteen 2×2 cluster subregions, each of which can potentially contain at most one ROI. Physically the cluster processing algorithms are implemented in eight large FPGAs (known historically as “CP chips”). Two additional “merger” FPGAs each collect data from the CP chips for half of the threshold bits and send final results to the CMM.

The EM and tau algorithms are similar, and indeed 8 of the 16 cluster definitions can be internally selected to perform either (the other 8 are EM-only). A sliding window

algorithm with “core” window of 2×2 trigger towers tests for a sum of two towers within that window that are adjacent in eta or phi, and that exceed a minimum Et threshold. Trigger towers in the surrounding 4×4 ring (and the hadronic towers behind the EM “core”) are also compared with maximum isolation thresholds.

For each cluster subregion, the current cluster algorithm produces 16 threshold bits, each corresponding to an independent EM or τ /hadron ROI definition. Two additional bits give the fine position of the cluster within this subregion to a resolution of 0.1.

A summary of the information available from Day-1 algorithms is as follows:

# of ROIs	Bits per ROI	Total bits
16 subregions	16 thresholds (plus 2 fine-position bits)	256 (plus position)

Each CP crate has two merger modules, so half of the real-time output data output is divided between them.

Each CP chip reports 8 of the 16 thresholds for its two subregions to each of two “merger” FPGAs on the CPM. This is done for every bunch crossing via two 16-bit parallel output ports running at 40 Mbit/s. For Day-1 running, each merger chip produces eight 3-bit multiplicities corresponding to its set of 8 thresholds, and transmits them to the CMM at 40 Mbit/s:

b0-2	b3-5	b6-8	b9-11	b12-14	b15-17	b18-20	b21-23	bit 24
Thr 0	Thr 1	Thr 2	Thr 3	Thr 4	Thr 5	Thr 6	Thr 7	Parity

For upgrade running at 160 Mbit/s, simulation studies suggest encoding up to five EM and hadronic clusters per CPM/CMX as follows:

- 16 presence bits (1L, 1R, 2L, ..., 8R)
- Five 8-bit cluster ET values (Cluster ET)
- Five 2-bit fine positions (fp)
- Five 2-bit EM isolation bits (ei)
- Five 2-bit hadronic isolation bits (hi)
- Five 2-bit hadronic veto bits (hv)

ROIs would be identified as EM or hadronic clusters satisfying the lowest cluster ET with the weakest isolation cuts.

A possible 160 Mbit/s format might be:

Bits	0-23										24
0	Presence bits 1L – 8R					ei1	hi1	hv1	fp1	clk/ pty	
1	Cluster ET 1		Cluster ET 2			ei2	hi2	hv2	fp2		
2	Cluster ET 3		Cluster ET 4			ei3	hi3	hv3	fp3		
3	Cluster ET 5		ei4	hi4	hv4	fp4	ei5	hi5	hv5		

Note that for hadronic clusters, the hadronic veto is not used, and would be left blank.

Fiber links from CMX to L1Topo

Each CMX in a processor crate receives feature data from the fourteen (14) CPMs or sixteen (16) JEMs, and transmits them to the TP over 12-fiber optical transmission modules at a serial rate of 6.4 Gbit/s. The serial data streams are produced by multi-Gbit transceivers in the FPGA, using 8b/10b encoding to maximize the data payload per fiber to 128 bits/BC/fiber.

Fiber links from Muon CTP interface (MUCTPI) to L1Topo

The current topological trigger proposal includes provisions for including muon ROIs, which would be provided by new muon interface octant modules (MIOCT) in the muon CTP interface. While such boards have yet to be specified, it may be a useful exercise to consider the data that might be expected (or requested).

Sixteen MIOCT modules each receive and process muon candidates from one octant (0.8 in phi) of either the A or C side. Within the barrel detector the muons are reported with $.1 \times .1$ eta-phi granularity, but only one muon ROI can be found per $.2 \times .2$ area. The coverage of the TGC part of the muon system extends to an eta of 2.4, so the number of equivalent $.2 \times .2$ muon “subregions” in a single MIOCT might be 12×4 in eta and phi, respectively, for a total of 48.

Topological processor (L1Topo)

CTP interface