

GENERAL COMMENTS ON CFEB-SCA OPERATION

The 25 ns clock used for CFEB-SCA operation comes from DMB and is driven by the LHC 40 MHz clock¹.

Buckeye output is continuously sampled by SCA every 50 ns. CFEB can be configured to digitize 8 or 16 samples. Below, 8-sample digitization will be assumed. Extrapolation to 16 cells is trivial.

Sync Reset releases all SCA blocks/cells and forces the SCA controller to start over from the first cell in the SCA block 5.

SCA 8-cell blocks are extracted from the pool of available ones according to Gray code. Typically, cycling goes over 4 blocks, e.g., 1, 3, 2, 0². In case of very high rates, Gray code sequence may not always be possible (due to shortage of available blocks)—the best possible choice is then made. One cannot mask/skip bad SCA cells or blocks, even if they are known to be bad.

An arriving LCT locks 2 or 3 SCA blocks, depending on whether this LCT falls in the last 7 cells (Case 1, Fig. 1) or the 1st cell of the block (Case 2, Fig. 2)³. These blocks will not be available for sampling the Buckeye output signal until they are released. LCT phase with respect to SCA 50 ns clock is also locked at that time (LCT can arrive every 25 ns).

If no L1As come to CFEB in $X \pm 1$ BXs after LCT, the LCT-locked blocks are released and returned back to the pool of available blocks. The number X is currently hard coded in CFEB firmware to be 116, or $T_X = 2900$ ns. This is OK for beam tests, but too long for the 3.2 μ s CMS L1 latency—therefore, in future it will be changed accordingly. The ± 1 BX window is fixed and will not change.




Fig. 1. Locking SCA blocks and cells by LCT and L1A*LCT coincidence: Case 1.

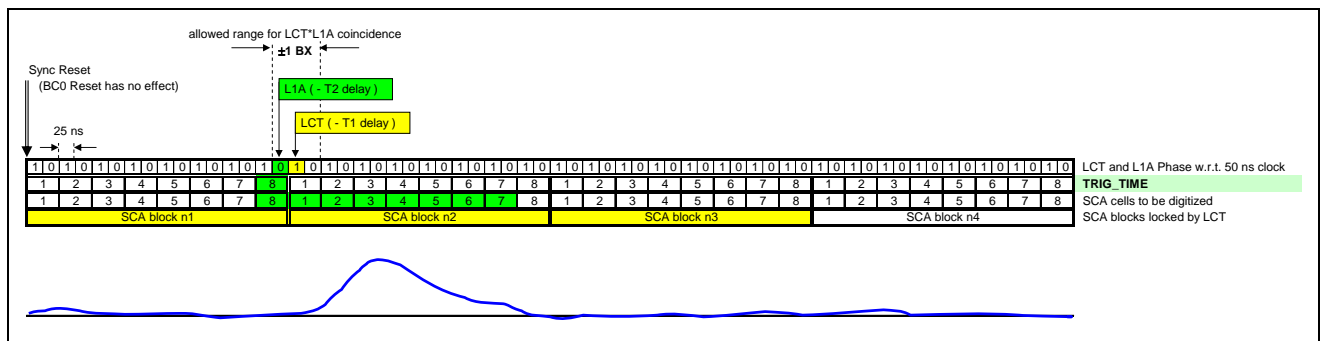


Fig. 2. Locking SCA blocks and cells by LCT and L1A*LCT coincidence: Case 2.

¹ Note: CFEB-Comparator 25 ns clock is also driven by the same 40 MHz LHC clock, but it comes to CFEB via TMB and, therefore, will have an arbitrary but fixed phase with respect to the CFEB-SCA clock.

² If the LCT decision latency were < 20 BXs (< 500 ns), one could get away by cycling over only 2 SCA blocks (800 ns). But the current CLCT latency from the moment track goes through CSC to the CLCT signal coming back to CFEB is ~ 34 BXs (850 ns). The plan is to cut it back by ~ 7 BXs.

³ Note that locking 3 SCA blocks for an LCT falling in the first cell of a block is driven by technical conveniences rather than by a real necessity. Indeed, if LCT_phase is 1, one needs to lock blocks n1 and n2; otherwise—n2 and n3.

If a L1A signal does come within the ± 1 BX window (by the CMS L1 trigger design, only one L1A signal can occur within 3 BXs), three things happen.

First, DAV-signal (data available) is sent to DMB. DAV signal is in sync with the L1A signal and is transmitted via a dedicated wire (not used for CFEB data transmission). CFEB DAV bits appear in the DMB headers/trailers.

Second, the block cell corresponding to L1A set a corresponding bit in TRIG_TIME⁴ high and the phase of L1A with respect to 50 ns clock (L1A_PHASE) is recorded. Taking, Fig.1 as an illustration: TRIG_TIME=0000.0100 and L1_PHASE=0. Effectively, the combination of TRIG_TIME and L1A_PHASE is a 4-bit counter of bunch crossings from the time of the last Sync Reset to the L1A signal. Note that DMB also has a 4-bit counter that gets reset only by Sync Reset and, therefore, its content must correlate with that of CFEBs, albeit some constant shift, and, consequently, allows for assigning L1A_number and BXN as recorded by DMB to CFEB data. LCT_PHASE is also recorded (if L1A and LCT phases are different, LCT arrival time has 50 ns uncertainty).

Third, the LCT-locked blocks get frozen for the time needed to digitize and transmit time samples: $T_0(n) + 150$ ns (Flash ADC) \times 16 (strips/ADC) \times $N_{\text{samples}} = T_0(n) + 19.2$ μ s for 8-sample readout. The cells to be digitized are counted starting from the cell corresponding to L1A (TRIG_TIME). Digitized samples are not buffered on CFEB, but rather sent (pushed) to DMB right away one by one as they become available. SCA blocks are released back in the pool block-by-block when all needed cells per block are digitized⁵. The overhead $T_0(n)$ depends on the position n of first cell to be digitized in the 1st block: $T_0(n)=400$ ns + (150 ns) * (TRIG_TIME - 1).

If a second LCT \times L1A arrives and requires the same block to be frozen, the overlapping “*to-digitize*” cells are tagged as OVERLAPPED (bit 14 in CFEB data format). They are not digitized for the second time, but rather copied in the overlap FIFO on DMB at the time of transmitting data corresponding to the first LCT \times L1A coincidence and are reused when the second event data package is prepared by DMB. Also, TRIG_TIME may have one more bit (corresponding to the second L1A \times LCT coincidence) set high, as long as the second coincidence falls in the same SCA block as the first one. E.g. (using Fig.1), if the second L1A arrives in 200 ns after the first one and has LCT \times L1A coincidence, the TRIG_TIME for the first event will be 0100.0100.

TRIG_TIME bits are the same for all samples belonging to the same SCA block. Samples from consecutive blocks will have TRIG_TIME bits set by occurrences of L1A \times LCT coincidence corresponding to those blocks. At low rates, one expects only one bit set high in TRIG_TIME corresponding to cells of the first SCA block, while all consecutive cells will typically have TRIG_TIME=0.

If a third (fourth/etc.) LCT \times L1A comes within 400 ns after the first LCT \times L1A, this would require same SCA cells used in three events, which cannot be done in the present design. Instead, such occasions⁶ would result in:

- a) no DAV signal being generated
- b) MOVLP bit (multiple overlap) being transmitted to DMB (via a dedicated wire)
- c) no CFEB data being sent to DMB
- d) TRIG_TIME might be acquiring another bit set high.

After transmitting the full package of 96 digitized charges corresponding to one time sample, CFEB adds 4 trailer words and goes to the next time sample. CFEB part that performs digitization and data transmission does not care of event boundaries—it just digitizes time samples marked for digitization and sends (pushes) the data to DMB.

⁴ Prior to the 2003 beam tests, TRIG_TIME cell was set by LCT.

⁵ All decisions on locking/releasing SCA blocks are taken every 400 ns in synch with the Buckeye signal sampling cycle (= one SCA block). This implies that it may take anywhere between 0 and 400 ns to release an SCA-block after the last needed sample in it was digitized.

⁶ IMPORTANT: CMS L1A rate rules (see footnote on page 5) imply that such compact trains of L1A signals (3 L1As within 400 ns) can never happen. However, with 16-cell readout the window becomes 800 ns and triple-L1A bursts within 800 ns may be possible. Note that in beam tests conditions such overlaps are not excluded for either of the two, 8- or 16-cell, readout scenarios.

B-words:

If SCA runs out of free SCA blocks, corresponding time windows when no free SCA blocks were available for signal sampling are tracked and in the case when LCT×L1A coincidence requires to digitize some samples that were never recorded, special B-words are generated in place of ADC data corresponding to the lost samples. More specifically, in place of 100 words for each unavailable sample, transmitted are 4 B-words with structure of 1011.001x.TRIG.TIME, where 8 TRIG.TIME bits correspond to the usual TRIG_TIME and x=SCA_FULL at the time of transmitting the corresponding word. Note that the bit 15 is non-zero unlike as in the case of normal CFEB data.

If SCA does run out of free SCA blocks, it does not mean that *all* samples will be necessarily lost (the pool of available blocks can change every 400 ns). Therefore, the length of CFEB data block will be: (4 B-words)*(number of lost samples) + (100 words)*(number of digitized samples) = N 16-bit words. For 8-sample readout mode, a set of B-words may show up in the beginning or at the end of the CFEB data block. For the 16-sample readout mode, B-words may show up in the middle of the CFEB data block—they would replace 8 consecutive time samples in this case.

The DAV bit is always sent in these cases, even if one ends up with not a single digitized time sample transmitted.

State of the NF_SCA counter (number of free SCA blocks):

After Sync Reset:	=12 (short time)
As we start filling the first block	=11 (400 ns)
As we start filling second block	=10 (next 400 ns)
After filling second block (in 400 ns)	= 9 (some time)

If no LCT arrived within 800 ns after we finished with filling out the first block, release the first block	=10
If LCTs do not arrive ever	= 9 → 10 → 9 → 10 → ...

If LCT arrived, lock 2 blocks	= 9
As we keep filling/releasing sampling blocks	= 8 → 7 → 8 → 7 → 8 → ...

If L1A does not arrive in T _x , release 2 LCT-locked blocks	= 9 → 10 → 9 → 10 →
If L1A does arrive, keep LCT-locked blocks and start digitization	= 8 → 7 → 8 → 7 → 8 → ...
As digitization proceeds, "finished-with" blocks are released	
- after releasing the first block	= 9 → 8 → 9 → 8 → ...
- after finishing with the second	=10 → 9 → 10 → 9 → ...

Highlighted are the typical expected NF_SCA states locked during CFEB data transmission at low hit rate.

CFEB-2005 DATA FORMAT

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0	x1	y1	12 ADC bits + under/overflow bit (#12)												
0	x1	y1	...												
0	x1	y1	...												
0	x1	y1	...												
0	x1	y1	...												
0	x1	y2	...												
0	x1	y2	...												
0	x1	y2	...												
0	x1	y2	...												
0	x1	y2	...												
0	... strips go in the Gray code order: {0,1,3,2,6,7,5,4,12,13,15,14,10,11,9,8} ...														
0	x1	y16	...												
0	x1	y16	...												
0	x1	y16	...												
0	x1	y16	...												
0	x1	y16	...												
0	15-bit CRC for 96 words above (only lowest 13 bits of 96 words checked)														
0	1	1	1	a	b	c	d	LCT_PIPE_CNT(4)				NF_SCA(4)			
0	1	1	1	CFEB_L1A(6)					L1A_PIPE_CNT(5)					e	
0	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0

sample 1
strip 1st {#0}
all 6 layers {#3,1,5,6,4,2}

sample 1
strip 2nd {#1}
all 6 layers {#3,1,5,6,4,2}

sample 1
strip 16th {#8}
all 6 layers {#3,1,5,6,4,2}

<-- word 97

<-- word 98

<-- word 99

<-- dummy word to make
total number of words (100),
divisible by 4

Repeat this 100-word structure for each next time sample, i.e. 8 or 16 times

4 B-words are sent in place of 100 data word for each missing SCA sample...

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	1	1	0	0	1	f	TRIG_TIME(8)							
1	0	1	1	0	0	1	f	TRIG_TIME(8)							
1	0	1	1	0	0	1	f	TRIG_TIME(8)							
1	0	1	1	0	0	1	f	TRIG_TIME(8)							
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0

Event Size:

If no B-words present, = NCFEBs × Nsamples × 100 × 2 Bytes
= 1.6 kB/CFEB for 8-sample readout (double for 16-sample readout)

If only B-words present, = NCFEBs × Nsamples × 4 × 2 Bytes
= 64 B/CFEB for 8-sample readout (double for 16-sample readout)

x(i)		If =0, this time sample overlaps with the next event and must be reused by DMB in the next event packet. This bit must be same in all 96 appearances per time sample. Expect 1 at low rates.
y(1:8)	TRIG_TIME(8)	The least significant bit in this 8-bit section unpacked from the first SCA sample tags the cell in the SCA block that corresponds to arrival of L1A, e.g.: 0000.0100 corresponds to the cell 3. Effectively, TRIG_TIME and L1A_PHASE make a counter of BXs since the last Sync Reset to L1A. See text for interpretation of more significant bits in the first time sample. Also, TRIG_TIME bits associated with time samples belonging to different SCA blocks are set independently and are likely to be different!
y(9:12)	SCA_BLK(4)	SCA block this sample belongs to.
y(13)	L1A_PHASE	25-ns phase of 50-ns clock at L1A arrival (1 for the first half of 50 ns).
y(14)	LCT_PHASE	25-ns phase of 50-ns clock at LCT arrival (1 for the first half of 50 ns). <i>Note: If LCT phase \neq L1A phase, then position of LCT relative to L1A is ambiguous.</i>
y(15)	SCA_FULL	If =1, SCA is currently full (no cells for signal sampling available). SCA data may be lost, but not necessarily in this event. <i>Note: SCA_FULL is evaluated at the time of transmitting the corresponding 16-bit word, i.e. it can be different for different time samples and even can flip within 6 sequential words corresponding to strip 9 (15th being transmitted). SCA full condition will go away by itself if LCT and/or LCT\timesL1A rates drop (no resets needed).</i>
y(16)	TS_FLAG	Number of samples to digitize (0 corresponds to 8 samples, 1—to 16 samples)
	CRC(15)	15-bit CRC word (generated by CFEB)
	NF_SCA(4)	Number of free SCA blocks (not used for sampling and not locked by LCT or L1A \times LCT). At low rates and at the time of transmitting data from CFEB to DMB, one must expect NF_SCA be mostly 7-8-9.
	LCT_PIPE_CNT(4)	LCT pipeline buffer length (at the time the word is transmitted), or number of <u>all</u> blocks locked by LCT and not yet locked by LCT \times L1A. Since LCT locks blocks for no longer than T_x , the max count is $T_x/(400\text{ ns}) = 8$. At low rates, typically one must expect 0 at the time of transmitting CFEB data
	L1A_PIPE_CNT(5) = bits(4:1) \times 8 ^{bit(5)}	L1A pipeline buffer length (at the time word is transmitted, or number of <u>all</u> blocks locked by LCT \times L1A. Since digitization takes no longer than 40 μ s (16 samples) and taking into account CMS L1A rate constraints ⁷ , the max count is $4 \times 6 + 3 = 27$. At low rates, typically, one must expect this bit be 2/1 during transmitting the first $N_{\text{sample}} - 1$ samples and 0/1 in the last CFEB trailer.
b	LCT_PIPE_EMPTY	LCT pipeline buffer empty (=1)
d	LCT_PIPE_FULL	LCT pipeline buffer full (=1); can never be full ($8 < 2^4 = 16$).
a	L1A_PIPE_EMPTY	L1A pipeline buffer empty (=1)
c	L1A_PIPE_FULL	L1A pipeline buffer full (=1); can never be full ($27 < 2^8 = 256$)
e	L1A_PIPE_WARNING	More than 32 L1A's in the L1A pipeline (=1)
	CFEB_L1A(6)	L1A number as recorded by CFEB

⁷ No more than: 1 L1A per 75 ns, 2 L1A's per 625 ns, 3 L1A's per 2.5 μ s, 4 L1A's per 6 μ s,

Rev 3 Changes and References:

- Word99 is completely redefined (J. Gu's email from Aug 25, 2005)

Rev 2 Changes and References:

- J. Gu's comments on the previous draft are incorporated (comments received May 18, 2004)

Rev 1 References:

- More recent updates (private communications with J. Gu., S. Durkin, J. Gilmore, April 2004)
- <http://www.physics.ohio-state.edu/~cms/ddu/ddu.html> (last updated 1 May 2003)
- http://www.physics.ohio-state.edu/~cms/dmb/dataformat_ddu.html (last updated 16 Jan 2003)

APPENDIX: CFEB DATA FORMAT (Prior to September 2005)

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0	x1	y1	12 ADC bits + under/overflow bit (#12)												
0	x1	y1	...												
0	x1	y1	...												
0	x1	y1	...												
0	x1	y1	...												
0	x1	y2	...												
0	x1	y2	...												
0	x1	y2	...												
0	x1	y2	...												
0	x1	y2	...												
0	... strips go in the Gray code order: {0,1,3,2,6,7,5,4,12,13,15,14,10,11,9,8} ...														
0	x1	y16	...												
0	x1	y16	...												
0	x1	y16	...												
0	x1	y16	...												
0	x1	y16	...												
0	15-bit CRC for 96 words above (only lowest 13 bits of 96 words checked)														
0	1	1	1	a	b	c	d	LCT_PIPE_CNT(4:7)				NF_SCA(0:3)			
0	1	1	1	e	f	g	h	L1_PIPE_CNT(0:7)							
0	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1

sample 1
strip 1st {#0}
all 6 layers {#3,1,5,6,4,2}

sample 1
strip 2nd {#1}
all 6 layers {#3,1,5,6,4,2}

sample 1
strip 16th {#8}
all 6 layers {#3,1,5,6,4,2}

<-- word 97
<-- word 98
<-- word 99
<-- dummy word to make total number of words (100), divisible by 4

Repeat this 100-word structure for each next time sample, i.e. 8 or 16 times

4 B-words are sent in place of 100 data word for each missing SCA sample...

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	1	1	0	0	1	f	TRIG_TIME(8)							
1	0	1	1	0	0	1	f	TRIG_TIME(8)							
1	0	1	1	0	0	1	f	TRIG_TIME(8)							
1	0	1	1	0	0	1	f	TRIG_TIME(8)							

Event Size range:

If no B-words present, = NCFEBs × Nsamples × 100 × 2 Bytes
= 1.6 kB/CFEB for 8-sample readout (double for 16-sample readout)

If only B-words present, = NCFEBs × Nsamples × 4 × 2 Bytes
= 64 B/CFEB for 8-sample readout (double for 16-sample readout)

x(i)		If =0, this time sample overlaps with the next event and must be reused by DMB in the next event packet. This bit must be same in all 96 appearances per time sample. Expect 1 at low rates.
y(1:8)	TRIG_TIME(8)	The least significant bit in this 8-bit section unpacked from the first SCA sample tags the cell in the SCA block that corresponds to arrival of L1A, e.g.: 0000.0100 corresponds to the cell 3. <i>Effectively, TRIG_TIME and L1A_PHASE make a counter of BXs since the last Sync Reset to L1A.</i> See text for interpretation of more significant bits in the first time sample. Also, TRIG_TIME bits associated with time samples belonging to different SCA blocks are set independently and are likely to be different!
y(9:12)	SCA_BLK(4)	SCA block this sample belongs to.
y(13)	L1A_PHASE	25-ns phase of 50-ns clock at L1A arrival (1 for the first half of 50 ns).
y(14)	LCT_PHASE	25-ns phase of 50-ns clock at LCT arrival (1 for the first half of 50 ns). <i>Note: If LCT phase \neq L1A phase, then position of LCT relative to L1A is ambiguous.</i>
y(15)	SCA_FULL	If =1, SCA is currently full (no cells for signal sampling available). SCA data may be lost, but not necessarily in this event. <i>Note: SCA_FULL is evaluated at the time of transmitting the corresponding 16-bit word, i.e. it can be different for different time samples and even can flip within 6 sequential words corresponding to strip 9 (15th being transmitted).</i> <i>SCA full condition will go away by itself if LCT and/or LCT\timesL1A rates drop (no resets needed).</i>
y(16)	TS_FLAG	Number of samples to digitize (0 corresponds to 8 samples, 1—to 16 samples)
	CRC(15)	15-bit CRC word (generated by CFEB)
	NF_SCA(4)	Number of free SCA blocks (not used for sampling and not locked by LCT or L1A \times LCT). At low rates and at the time of transmitting data from CFEB to DMB, one must expect NF_SCA be mostly 7-8-9.
	LCT_PIPE_CNT(4)	LCT pipeline buffer length (at the time the word is transmitted), or <i>number of all blocks locked by LCT and not yet locked by LCT\timesL1A</i> . Since LCT locks blocks for no longer than T_x , the max count is $T_x/(400\text{ ns})=8$. At low rates, typically one must expect 0 at the time of transmitting CFEB data
b	LCT_PIPE_MT	LCT pipeline buffer empty (=1)
d	LCT_PIPE_FULL	LCT pipeline buffer full (=1); can never be full ($8 < 2^4=16$).
	L1_PIPE_CNT(8)	L1A pipeline buffer length (at the time word is transmitted, or <i>number of all blocks locked by LCT\timesL1A</i> Since digitization takes no longer than 40 μ s (16 samples) and taking into account CMS L1A rate constraints ⁸ , the max count is $4 \times 6 + 3 = 27$. At low rates, typically, one must expect this bit b 2/1 during transmitting the first $N_{\text{sample}}-1$ samples and 0/1 in the last CFEB trailer.
a	L1_PIPE_MT	L1A pipeline buffer empty (=1)
c	L1_PIPE_FULL	L1A pipeline buffer full (=1); can never be full ($27 < 2^8=256$)
f	SCA_FULL	SCA full condition (=1) <i>Same as the serialized bit y(15). However, the check is done at time of transmitting the word 99.</i>
h	LCT_POP	This bit indicates that a used CSA block has become available. <i>This bit latches LCT_POP state that pops up only for 25 ns just before LCT_PIPE_CNT decreases by 1. Chances of seeing it equal 1 in data stream are very slim...</i>
e	BUSY	This bit indicates that the system is digitizing data (normally must be 1, and 0 at the very end of an event at low rates).
g	CFEB_PUSH	This bit indicates that data is being sent to DMB. In data stream, must be 1 (except, maybe, for the very beginning or end of the event due to intrinsic CFEB delay offsets).

⁸ No more than: 1 L1A per 75 ns, 2 L1A's per 625 ns, 3 L1A's per 2.5 μ s, 4 L1A's per 6 μ s,