

Comparison of data formats

Jan C. Bernauer for the organizers

Streaming Readout VII, 2020




RBRC

RIKEN BNL Research Center



**Stony Brook
University**



Thank you all for sending me the slides!

- ▶ Full slides are attached at the end. I'll show here some selections of things I noticed.

Microslice format:

STS/MUCH																																																																
Header	8b	0xDD																Header Format ID = 0xDD																Has to be accessed with the FLES IPC through the Microsilice and MicrosiliceDescriptor classes																														
	8b	0x01																Header Format vers. = 0x01																																														
	16b	DPB MAC LSB																Equipment ID (16b) = DPB MAC LSB																																														
	16b	XX																Status and error flags (16b)																																														
	8b	0x10 [STS] or 0x50 (MUCH)																SubSystem ID (8b)																																														
	8b	0x20																SubSystem Format vers. (8b)																																														
	64b	multiple of TS_MSB length (1.6 us), HAS TO MATCH other subsystems in experiment																Microsilice index/start time in ns (64b)																																														
	32b	calculated by FLJM FW core																CRC-32C (Castagnoli polynomial) of data content (32b), calculated by FLJM FW core																																														
32b	Nb messages * 4																Data content size in bytes (32b) = NbB																																															
64b	XX																Offset in buffer in bytes (64b)																																															
Bits format																																																																
Message	63	62	61	60	59	58	57	56	55	54	53	52	51	50	49	48	47	46	45	44	43	42	41	40	39	38	37	36	35	34	33	32	31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1 - 0	1st TS_MSB																TS_MSB cycle message containing current TS_MSB counter overflow cycles (minimum of 19b for 1 month)																																															
...	TS_MSB + empty for 64b padding + Hits, TS_MSB suppressed if no hit																																																															
n	Nothing OR end of microsilice message																End of microsilice message OR last Hit																																															

- ▶ Byte aligned, padded to 8 bytes
- ▶ A microslice is part of a buffer? (Offset?)
- ▶ CRC

Timeslice Header/Trailer:

Table 1: Detector Specific MUX

31-0	NU	frame type 4b	spill number (11b)	SrcID (16b)
63-32	SliceID (18b)			Slice Length(14b)
95-64	Time of Slice			
127-96	NmbOfImages (24b)			NU (8b)
...	payload			
end	1	Slice CRC (31)		

31 30 29 28 27 26 25 24 23 22 21 20 19 18 17 16 15 14 13 12 11 10 9 8 7 6 5 4 3 2 1 0

payload: Image Header/Trailer:

31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
ImageID 14bit														StationID 3bit			ViewID 3bit			DataType 6bit				NU 4bit		S					
Payload 32bit																															
Payload 32bit																															
Payload 32bit																															
CRC 32bit																															

- ▶ bit packed
- ▶ 2xCRC

MBS subevent header (big endian representation):

```
typedef struct{

    int      l_dlen;          /* Data length +2 in words */

    short    i_subtype;      /* Subtype */

    short    i_type;         /* Type number */

    char     h_control;      /* Processor type code */

    char     h_subcrate;     /* Subcrate number */

    short    i_procid;       /* Processor ID [from setup] */

} s_veshe;
```

White Rabbit full timestamp header (WRTS):

```
int      sub-system id      /* (32 bits, multiples of 0x100) */

short    0x03E1             /* (16 bit fixed code)*/

short    WRTS_L16           /* WRTS bits 00-15 */

short    0x04E1             /* (16 bit fixed code)*/

short    WRTS_M16           /* WRTS bits 16-31 */

short    0x05E1             /* (16 bit fixed code)*/

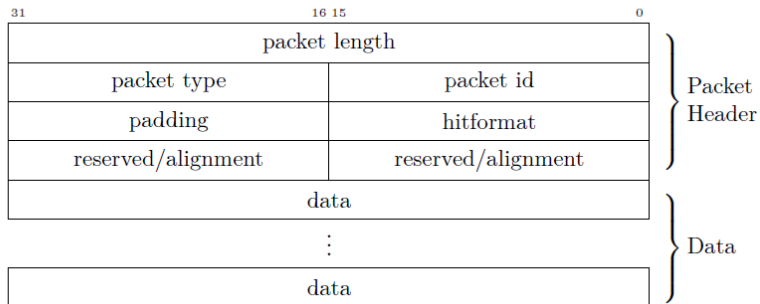
short    WRTS_H16           /* WRTS bits 32-47 */

short    0x06E1             /* (16 bit fixed code)*/

short    WRTS_X16           /* WRTS bits 48-61 */
```

Custom hit messages with "local" timestamps (TS, relative to WRTS header)

Table 1: The Packet Header.



- ▶ short aligned, 32 bit padded
- ▶ no crc?

Take away / Questions

- ▶ Rarely footer (which might be nice for CRC in a streaming setup?)
- ▶ Mostly byte aligned
- ▶ Typical fields: Length (first?), event nr/timing, channel
- ▶ Length of all, or payload?
- ▶ Since we won't have "events", how do we specify the time?
 - ▶ Absolute wall clock time?
 - ▶ Bunch Nr (+offset for finer timing?)

Streaming readout for multi-purpose DAQ system MBS

data packets from free running systems:

- free running data acquisition systems sends formatted **MBS sub-events** of hits
- each sub-event is **headed by White Rabbit Time Stamp (WRTS) - 1 ns units, starting from 1970**
- each sub-event contains **data from “many” hits** (MBS container)
- **each hit has TS** of variable size, but significantly smaller than WRTS
- each hit TS must have a **sufficient correlation to full WRTS** header
- **hit data format** inside MBS container has no dependency for time sorting and **can be chosen freely by each detector/sub-system**

MBS subevent header (big endian representation):

```
typedef struct{  
  
    int        l_dlen;           /* Data length +2 in words */  
  
    short      i_subtype;        /* Subtype */  
  
    short      i_type;           /* Type number */  
  
    char       h_control;        /* Processor type code */  
  
    char       h_subcrate;       /* Subcrate number */  
  
    short      i_procid;         /* Processor ID [from setup] */  
  
} s_veshe;
```

White Rabbit full timestamp header (WRTS):

```
int        sub-system id      /* (32 bits, multiples of 0x100) */  
  
short      0x03E1             /* (16 bit fixed code)*/  
  
short      WRTS_L16           /* WRTS bits 00-15 */  
  
short      0x04E1             /* (16 bit fixed code)*/  
  
short      WRTS_M16           /* WRTS bits 16-31 */  
  
short      0x05E1             /* (16 bit fixed code)*/  
  
short      WRTS_H16           /* WRTS bits 32-47 */  
  
short      0x06E1             /* (16 bit fixed code)*/  
  
short      WRTS_X16           /* WRTS bits 48-61 */
```

Custom hit messages with “local” timestamps (TS, relative to WRTS header)

.....

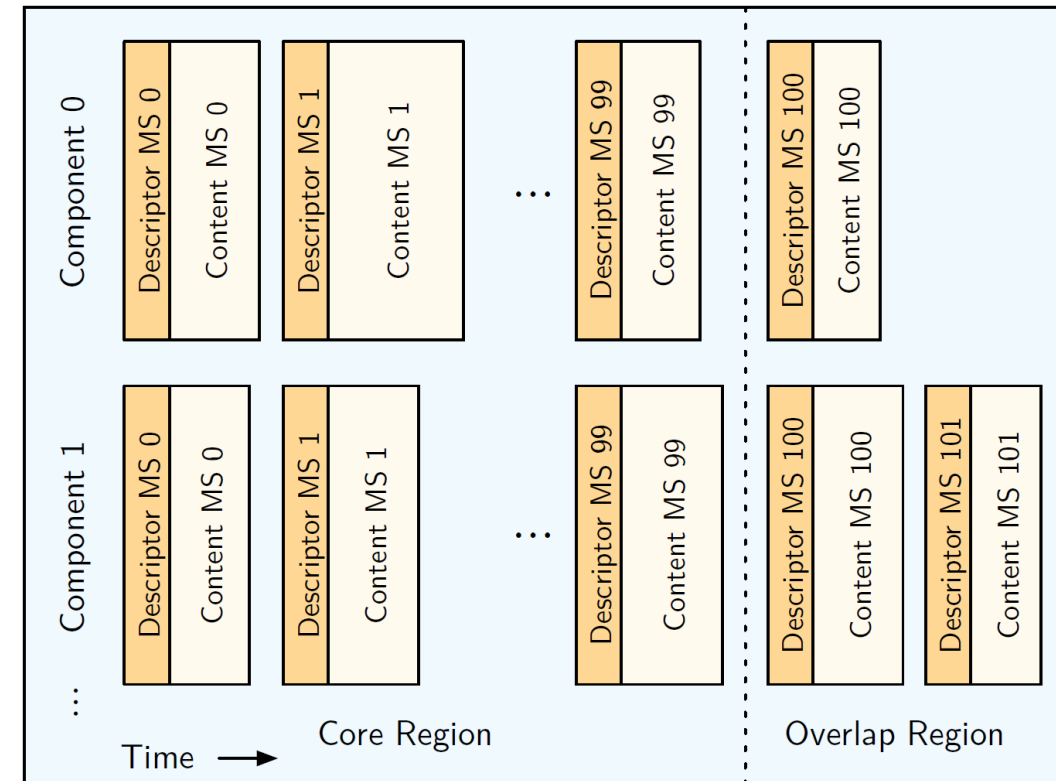
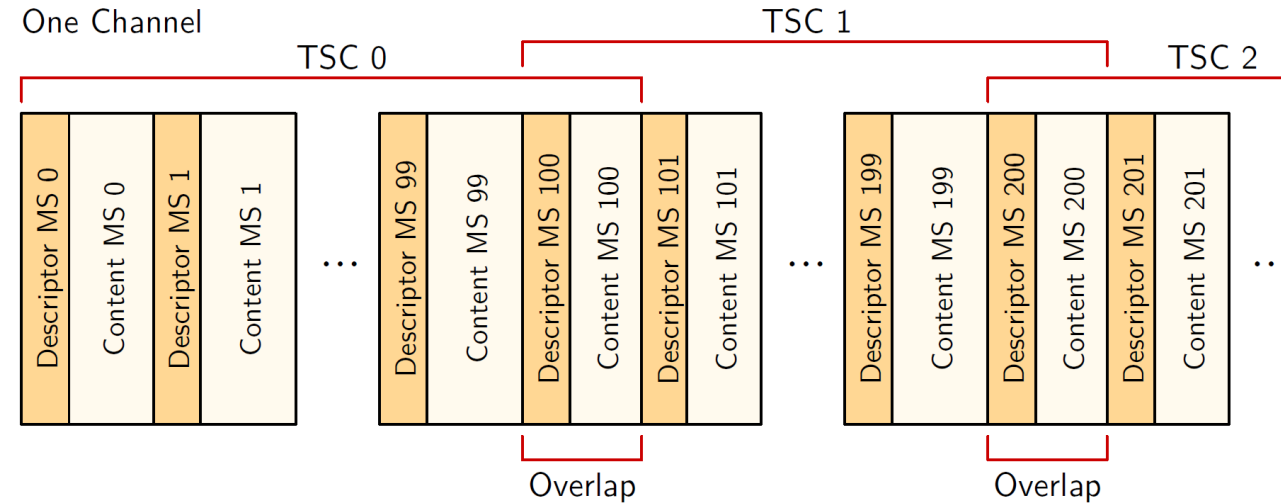
CBM: Data containers for transport and analysis

➤ MicroSlice(μ S) = self-contained data container from a single CRI and for a fixed period of time

- Output of the CRI in CBM => generated in FW
- Constant length in experiment time
- Typical period of time: 10's of μ s to ms
- Length adapted to data format, data rate (beam condition), container efficiency and network performance

➤ TimeSlice(TS) = container collecting the μ S of all CRI cards in the setup and for a given number of μ S


- Assembled by FLENET (CBM DAQ prototype)
 - ⇒ generated in SW
- Typical number of μ S per source: 10-1000
 - ⇒ time range: ms to s
- Includes overlap μ S to avoid analysis losses close to edge
- Length adapted to match memory resource in processing nodes
- One TS for the full setup for each time interval
- Input unit for Online analysis: TS are distributed to the processing nodes for reconstruction, event building, physics analysis and selection



CBM: Data sources, Data Format and constraints on containers

- Data sources = ASICs and/or FPGAs, self-triggered and free-streaming
- Messages = representation of the signals generated in data sources, not context free as optimized for best compromise between resolution and bandwidth usage
- Typical stream organization: Periodic context messages with MSB of timestamp + Hit messages with LSB & ADC, TOT, chan ...
- MicroSlice(μ S) = self-contained data container from a single CRI and for a fixed period of time, granularity of length choice depends on context messages interval

Example for Silicon Tracker prototype:



Microslice format:

STS/MUCH																																																																
Header	8b	0xDD																																Header Format ID = 0xDD												Has to be accessed with the FLES IPC through the Microslice and MicrosliceDescriptor classes																		
	8b	0x01																																Header Format vers. = 0x01																														
	16b	DPB MAC LSB																																Equipment ID (16b) = DPB MAC LSB																														
	16b	XX																																Status and error flags (16b)																														
	8b	0x10 (STS) or 0x50 (MUCH)																																SubSystem ID (8b)																														
	8b	0x20																																SubSytem Format ver. (8b)																														
	64b	multiple of TS_MSB length (1.6 us), HAS TO MATCH other subsystems in experiment																																Microslice index/start time in ns (64b)																														
	32b	calculated fy FLIM FW core																																CRC-32C (Castagnoli polynomial) of data content (32b), calculated by FLIM FW core																														
	32b	Nb messages * 4																																Data content size in bytes (32b) = NbB																														
	64b	XX																																Offset in buffer in bytes (64b)																														
Message	Bits format																																																															
	63	62	61	60	59	58	57	56	55	54	53	52	51	50	49	48	47	46	45	44	43	42	41	40	39	38	37	36	35	34	33	32	31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1 - 0	1st TS_MSB																																TS_MSB cycle message containing current TS_MSB counter overflow cycles (minimum of 19b for 1 month)																															
...	TS_MSB + empty for 64b padding + Hits, TS_MSB suppressed if no hit																																																															
n	Nothing OR End of microslice message																																End of microslice message OR last Hit																															

COMPASS++/AMBER FriDAQ Protocol

Timeslice Header/Trailer:

Table 1: Detector Specific MUX

31-0	NU	frame type 4b	spill number (11b)											SrcID (16b)																		
63-32	SliceID (18b)													Slice Length(14b)																		
95-64	Time of Slice																															
127-96	NmbOfImages (24b)																		NU (8b)													
...	payload																															
end	1	Slice CRC (31)																														
	31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0

payload: Image Header/Trailer:

31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
ImageID 14bit														StationID 3bit				ViewID 3bit			DataType 6bit						NU 4bit				S
Payload 32bit																															
Payload 32bit																															
Payload 32bit																															
CRC 32bit																															

1 Data Formats

1.1 The Packet Header

The composition of the full 16-byte header is shown in table 1.

Table 1: The Packet Header.

31	16 15	0	
packet length			} Packet Header
packet type	packet id		
padding	hitformat		
reserved/alignment	reserved/alignment		
data			} Data
⋮			
data			

The length is measured in 32-bit units (DWords), which allows us to have packet lengths of up to 16 GBytes, although most packets are substantially smaller. Each packet is a multiple of 32bit units (so each data structure starts at least on a 32 bit boundary). We usually choose higher data alignment boundaries (64bit or even 128bit). Aligning the data blocks to the prevailing CPU data bus widths (64bit at present) speeds up the processing of data.

The fields in the header are

packet length the overall length of the packet structure in Dwords

packet id a unique identifier for the packet that says which unit generated the packet

packet type indicates the fundamental storage size in the packet, expressed in bytes (1 (character data), 2 (16bit), 4 (32 bit),...). This field is also known as the “swap unit” in case the data payload has to be byte-swapped for a different CPU architecture. It also gives the unit for the padding value.

hitformat This value identifies an algorithm to decode the data payload so the decoded data can be accessed by a set of standard APIs.

padding The amount of additional data added to bring the packet size to the desired alignment boundary.

2 reserved/alignment fields Those fields can hold 2 16bit values as needed to verify the proper alignment of various data blocks. They are set to 0 if unused.

31	16	15	0
packet length = 6			
type = 2		id = 1001	
padding = 1		hitformat = 3002	
0x3A		0x79CE	
40		20	
0		55	

} Packet Header
 } Data

Table 2: An example of a (fictitious) 64-bit aligned packet that holds the three 16bit values 20, 40, and 55, and a combined alignment value of 0x3A79CE.

31	16	15	8	7	0
Event length					
reserved				event type	
Event Sequence					
Run Number					
Time/Alignment 0					
Time/Alignment 1					
Luminosity Block					
reserved					
data					
⋮					
data					

Event Header

Data

Table 3: The structure of an Event Header.

Table 2 shows the composition of a (fictitious) packet with id 1001 that holds three 16-bit values (so it has the packet type 2) 20, 40, and 55. In order to maintain the 64bit alignment of the data, an additional 16bit word is added, which gives a padding value of 1.

1.2 The Event Header

While in transit between components, for example between a SEB and a Buffer Box, a number of packets is preceded by a *Event Header*, sometimes also called the *Frame Header*.

Table 3 shows the structure of the event header. The event length is again given in

	31	16 15	8 7	0	
0x0000890c	event length = 0x890c = 35084				} Event Header
0x00000002	0x000000			type = 2	
0x00000002	Event Sequence = 2				
0x00001051	Run Number= 0x1051 = 4177				
0x00000000	Time field 1 = 0				
0x5be1e129	Time field 2 = 0x5be1e129 = 1541529897				
0x00000000	Luminosity Block = 0				
0x00000000	reserved = 0				
	packet data				} Data
⋮					
	packet data				

Table 4: A hex-dump of an actual Event Header and its structure. The event type 2 denotes streaming data. Because the first time field is 0, the 2nd word is interpreted as a Unix time (1541529897). This corresponds to a date of Nov 6, 2018, 13:44:57, when the data were taken.

DWords (32bit length).

The Event header structure has two general-purpose time and alignment fields. If the first alignment field is 0, the second field is interpreted as a Unix time (32bits). Else the two fields are interpreted as system-specific alignment data.