TFC Partitioning Support and Status

Beat Jost
Cern EP

Trigger and DAQ Architecure

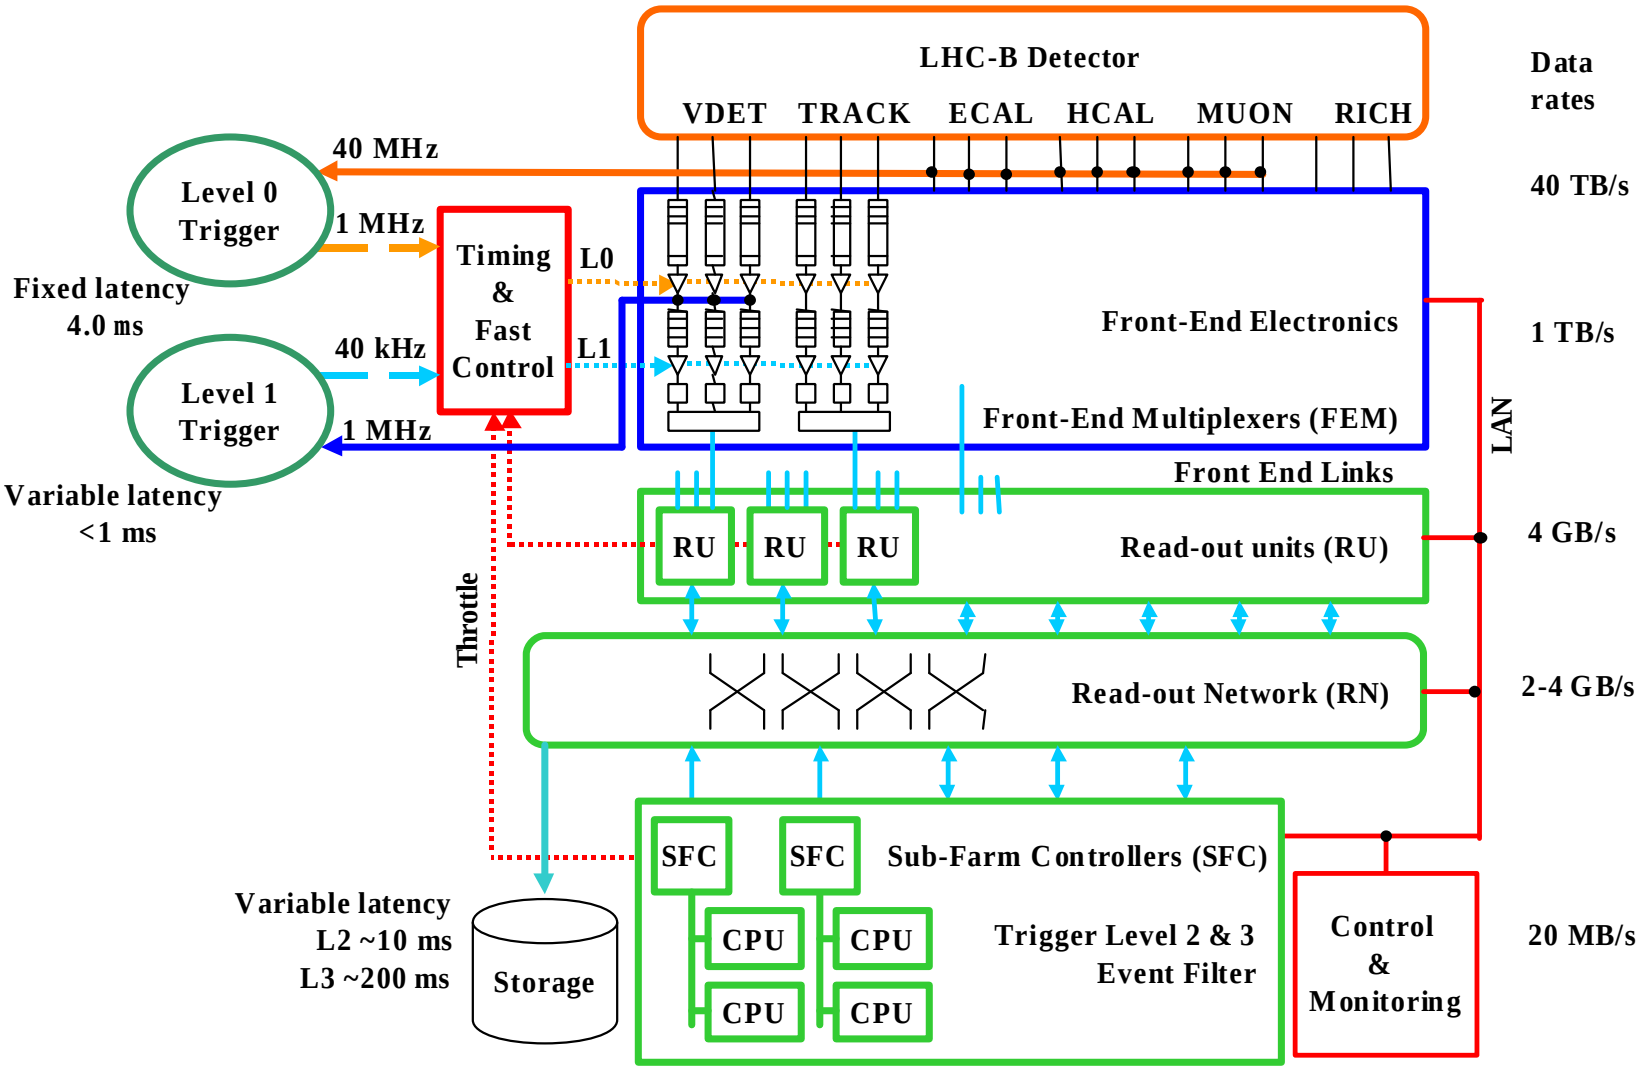

TFC Architecture

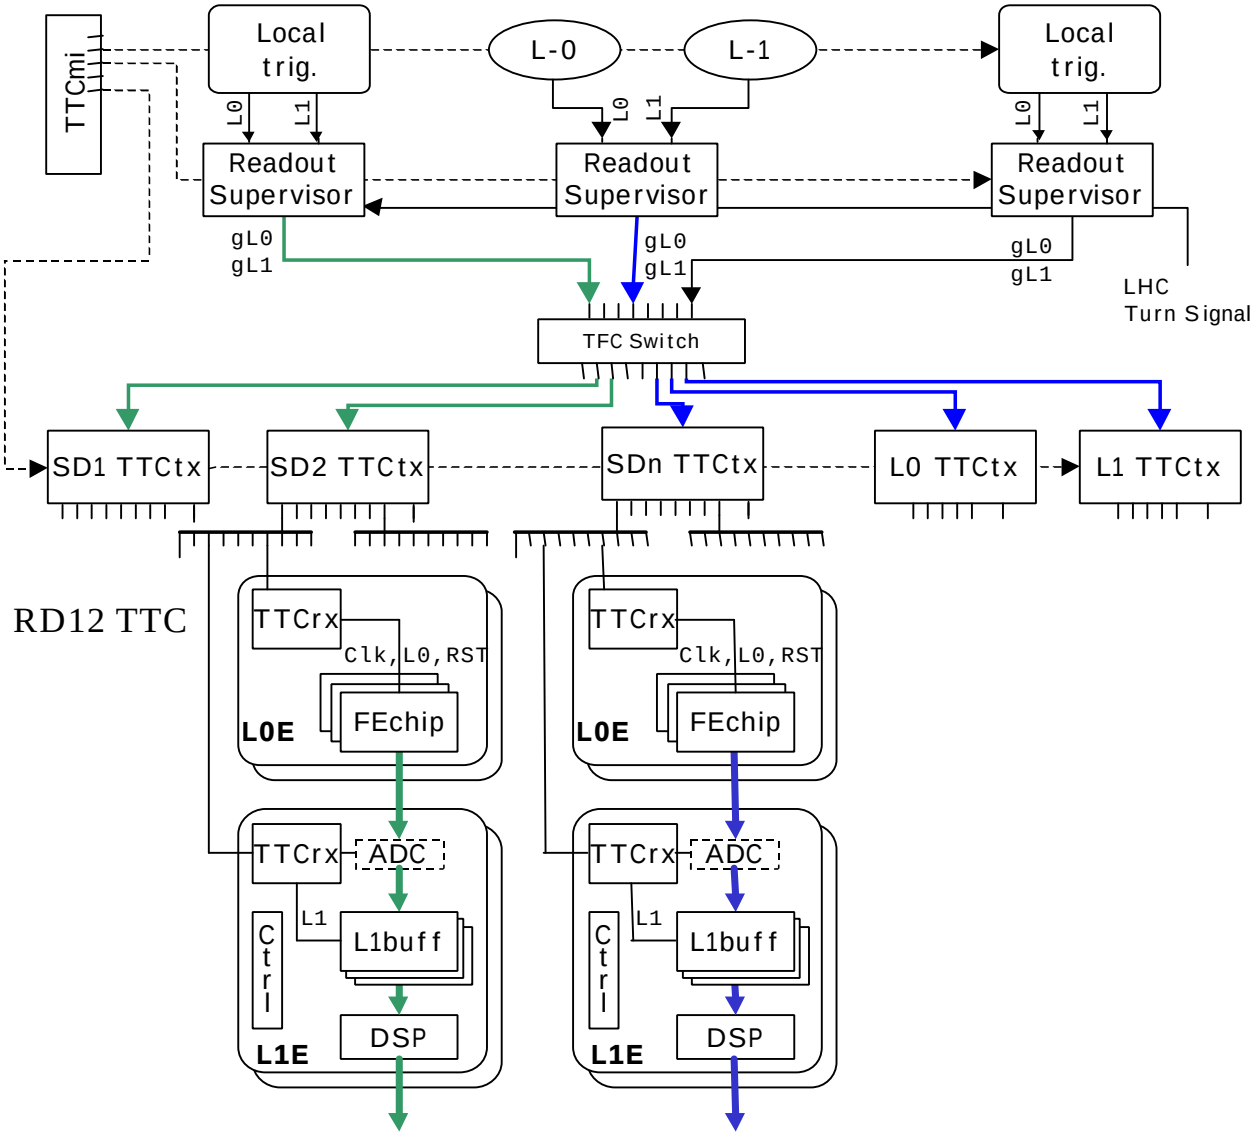

Signal Generation and Handling

Signal Distribution

Signal Reception

Partitioning

Definitions

- ***Partitioning*** is the sharing of the online system (ECS, TFC, DAQ) by independently and concurrently operated *Partitions*. Within a partition operation is coherent
- A ***domain*** is a functional component of the online system (ECS, TFC, DAQ, L2/L3 Farm)
- A ***Partition*** is defined as a collection or aggregation of ***partitioning grains***, where a partitioning grain is the smallest component in the domain that can be controlled independently plus a fraction of shared resources (e.g. Readout Network, controls network,…)
- In the TFC domain the partitioning grain is a L0 Electronics chip.
- For a DAQ activity partitions are defined in terms of L0 Electronics chips and the ***partitioning rules*** determine which additional components are needed to make up a fully functional partition.

Partitioning Grains (examples)

Domain	Partitioning Grain	Comments
TFC	L0E Chip	
Readout	FEMs, Rus, RS',...	Added automatically by partitioning rules
ECS	Controls Device (PS, RUs, FEMs, SFCs,...)	
L2/L3 Farm	SFC	

Example of Partitioning Rules

- In the TFC, domain choosing L0E-chips automatically adds the appropriate TTCrx'and TTCtx'to the partition
- Subsequently all the L1-Electronics connected to the added L0E has to be added to the partition.
- All the FEM components connected to the L1E will be added so will be the RUs that are connected to the FEMs
- Finally a certain number of SFCs will be needed together with the connected CPUs
- upstream of the TTCtx a Readout Supervisor will be added (which one depends on the sub-detector(s) to which the TTCtx'belong)
- In order to allow coherent running of the system, the RS has to be connected to the TTCtx'in the partition (path throughTFC switch).
- In addition all the components above need to be controlled and monitored, hence the corresponding control infrastructure will be added
- If necessary the appropriate DCS components (HV, LV) are added

Boundary Conditions induced by Partitioning

- No mixing at any level above the Readout Network of data originating from equipment connected to different TTCtx'
  - Level-0 Electronics (chips can be removed, but not run independently)
  - Level-1 Electronics
  - FEMs
  - RUs
  - Throttle ORs

  All these belong to only one active partition at any one time and receive data originating from only one TTCtx.

Connectivity (L0E, L1E,FEMs, RUs,Throttle ORs) has to follow the TTC system

- SFC's
  - Partition grain of the CPU farm
  - Associated to only one active partition at the time, but can be re-assigned to another partitions at another time.

Partitioning Support in the TFC System

- The main ingredient to support partitioning in the TFC system is the TFC Switch
- It allows to associate any set of outputs (TTCtx') to any one of the Readout Supervisors connected to the inputs

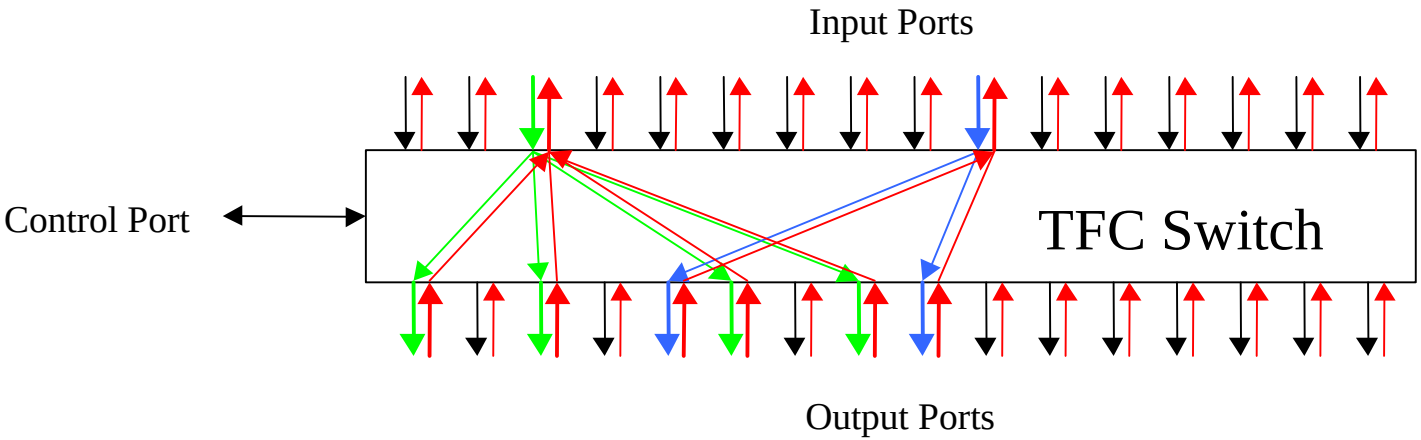

- The association is done through the controls interface
- There are also two reverse paths where the output towards the RS is the OR of the inputs. This is to propagate the L0 and L1 hardware throttle signals

Specifications of the TFC Switch (I)

TFC Path

Item	Specification	Comments
Number of input ports	16	Or more
Number of output ports	16	Or more
Input/output port characteristics	Compatible with RD12 TTC	
Jitter introduced by switch	<100 ps	To be checked ~5 ns might be viable
Phase difference between output ports	<100 ps	To be checked ~5 ns might be viable
Delay from input to output	<10 ns	

Specifications of the TFC Switch (II)

Throttle Path

Item	Specification	Comments
Number of input ports	2x16	L0/L1 Throttle
Number of output ports	2x16	L0/L1 Throttle
Input/output port characteristics	LVDS	
Logical Operation	Output is Logical OR of inputs	L0/L1 separately
Counters Increment Width Clear	One per input and output port Assertion of throttle on corresponding port 32-64 bits Through controls port	
Timers Start Stop Latching Resolution	One per input and output port Assertion of throttle on corresponding port De-assertion of Throttle Until restarted <25 ns	 History
Delay from input to output	<50 ns	

Monitoring

History of throttle assertions

- Basic idea is to store transitions with time-stamp in a FIFO (similar to modern Logic State Analyzers)
  - Example for L0 Throttle (L1 Throttle analogous)

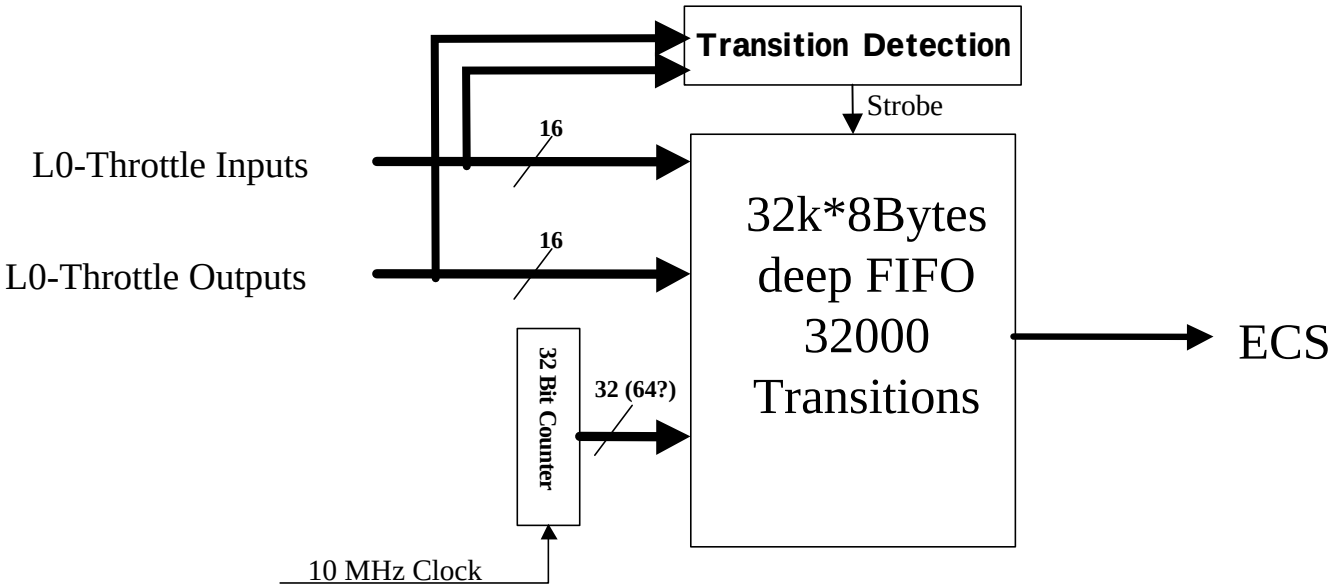

- Pros:
  - significantly less memory needed
- Cons:
  - more software effort to get throttle history

Implementation

- The TFC Switch' functionality will probably be implemented in two modules
  - one for the TFC path
    - possibly based on LUTs or ECL components (depending on jitter requirements)
  - one for the throttle path
    - probably based on FPGA
    - two separate but identical modules, one for the L0 throttle and one for the L1 throttle
- The form-factor of each module could very well be a 1-3 U rack-mount device
  - Front-Panel space

Other Components

- To support hardware throttling we need a unit that simply ORs a number of inputs (e.g. 16) to one output
  - Specifications
    - very similar to throttle path of the TFC switch
    - only one output
    - only needs to disable inputs
    - same characteristics concerning monitoring information
    - See also throttling session…

Status of TFC

- RD12
  - designing/building “final” version of TTCtx (max 448 Receivers)
  - designing/building TTCmi (machine interface)
    clock fanout/BC Reset. Max. 90 primary outputs. Very low jitter (10 ps)
  - “final” version of TTCrx in the pipeline (see Joergen’s presentation)
- LHCb TFC
  - TFC switch
    - specifications “ready”. First thoughts of implementation started
      (LHCb note in preparation)
    - First prototype in <6 months
  - Readout Supervisor
    - Specifications (LHCb note) in preparation (see Richard’s presentation)
    - first prototype in 12-15 months
  - Throttle OR
    - thoughts on specifications (very similar to switch) started
    - first prototypes in ~6 months

Issues to be discussed

- Architecture fixed?
- Partitioning scheme agreed?
- Monitoring capabilities of TFC switch and Throttle OR