CMX:
OPTIONS FOR CMX INPUT AND L1TOPO OUTPUT MODULES

Wojtek Fedorko

CMX (base) Firmware functional blocks

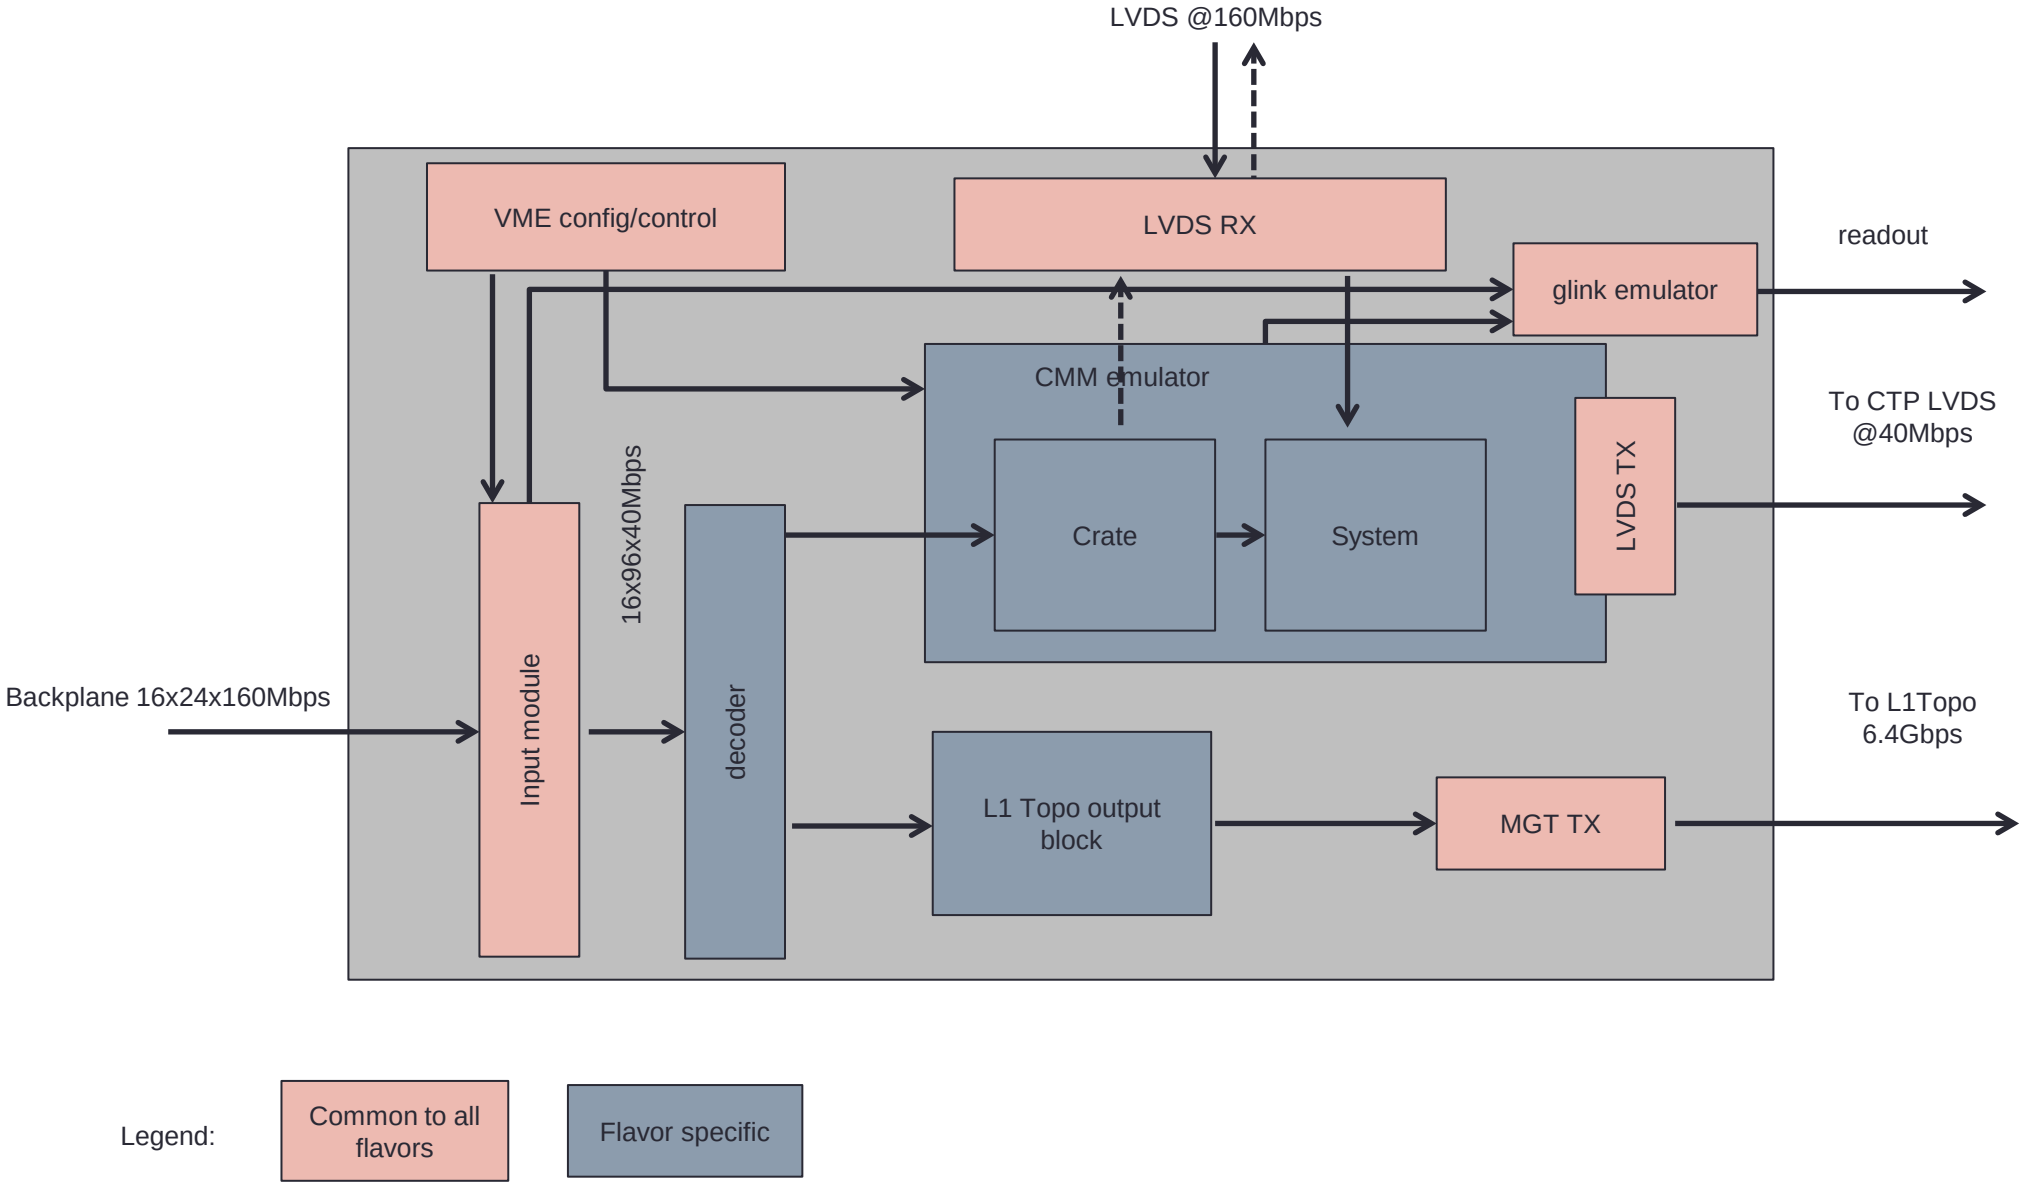

Input module options

- 16 clock domains of the processor inputs + 1 ‘system’ clock domain
- Scheme 1 ‘Demultiplex first’ (then synchronise)
  - Data captured using the forwarded 80MHz clock using IDDR buffer
  - IDDR demultiplexes to 80Mbps
  - Further demultiplexing to 40Mbps using same forwarded clock
  - Synchronisation to system domain at 40 MHz
- Scheme 2 ‘Synchronise first’ (then demultiplex)
  - Data captured using the forwarded 80MHz clock using IDDR buffer
  - IDDR demultiplexes to 80Mbps
  - Synchronisation to system domain at 80MHz (e.g. using first word fall through FIFOs)
  - Further demultiplexing to 40Mbps in the
  - Interfaces
  - Coordination and Implementation

Input module schemes: Pros and Cons

- ‘Demultiplex first’:
  - Latency – need to wait for all data to arrive before further processing
  - Framing pattern – need special pattern once to tell demux words 0,1 from 2,3
  - Wide window (25 ns) to latch data into system domain
  - Simplifies design of decoder
  - Implemented and tested on ML605 (minor mods may be needed)
- ‘Synchronise first’:
  - Potentially less robust (narrow window to cross clock domains)
  - May complicate decoder design
  - Data available for further processing 0.5 BC earlier
  - No framing pattern needed

Output module

- 3072 bits sent/BC on 24 TXs = 128 bits/TX
- 5.14 Gbps data rate, 6.4 Gbps line rate with 8b/10b encoding
- Each GTX accepts 16bit word at 320 MHz
- 2 clock domains comprising 3 neighboring ‘quads’ each
- 2 MMCMs seem to be necessary at this line rate
- Input is a vector of 3072 bits at 40 Mbps (or multiplicity e.g. 1536 @ 80 Mbps) translated from TO records.
- Vector is ‘condensed’ non-0 TO in the top rest ‘0’.
- 1-1 mapping between position in the vector and TX (e.g. first 128 bits go on TX 0 , next on TX 1…)
- If byte is identically 0 K28.5 is sent instead.
- TOs sent one after another without respecting byte or word boundary
- ‘Unused’ bits at the end set to 0.
- Q: Should we distribute the load among TXs? (reduce occupancy on TX 0 to reduce latency? What is preferable for L1Topo?)

Updated power estimate

- SAIF activity specification:

Source	Voltage (V)	Total Current (A)	Dynamic (A)	Quiescent (A)
Vccint	1	5.1	1.7	3.3
Vccaux	2.5	0.7	0.1	0.6
Vcco25	2.5	0.0	0.0	0.0
MGTAVcc	1	1.8	1.2	0.5
MGTAVtt	1.2	1.4	0.0	1.4

Power (W)
Total:
10.2
Dynamic:
3.1
Quiescent:
7.1

+0.4 W dynamic power consumption from the CMM emulation logic (thanks to Pawel and Sam)

- Very small dynamic power estimated
- GTX implemented – steady pattern on all TX (D0.0-K28.5)
- No synchronization to system clock domain yet