jFEX

These slides are meant to be a basis for discussions on Feb. 08/09

Constraints & Numerology

Assumption: one crate, several modules.

- Each module covers full phi, limited eta range
- Data sharing with immediate neighbour only, for practical reasons
  - sensible baseline is eta coverage of 0.8 per module
  - for +/- 3.2 eta coverage we need 8 modules
  - maximum environment size of 0.9 in eta
  - maximum environment requires 100% duplication !
- Each module will carry several FPGAs
  - assume 8 FPGAs, covering 0.8×0.8 (η×φ) each
    - total of 64 FPGAs for jFEX processor
    - compares to total of 32 jet FPGAs for current JEP
    - Seems reasonable due to intended improvement on granularity
- go for high density optoelectrical components
- design for short electrical traces of high speed links
- Latency aware

Components for 2018

Conservative approach: plan for devices only that are in the pipeline already now

- Large(ish) Vitex-7
  - SSI devices have higher latency. Avoid them unless there is considerable benefit in terms of cost
  - Plan for XC7VX550T
    - Mid range (smaller than the device envisaged for L1Topo)
    - Foot print compatible, larger device exists
    - Support up to 13 Gb/s (dependent on speed grade/price tag)
    - Nothing currently known about possible gaps in line rate range
- 12-channel opto devices at ~10Gb/s
  - can reasonably expect that microPOD will be a viable option by then
  - Hope for slight increase in data rate
- Expect to use ~48-fibre bundles

Basically design jFEX along the lines of L1Topo, with higher FPGA count and board-level data paths optimised for sliding windows algorithm

Data replication

<< If local design density permits duplication of a small fraction of signals, it will just as well allow for large scale duplication. Just a matter of money >>

- No replication of any source into more than two sinks
- Forward duplication only
- Avoid any data retransmission
- Fan-out in eta handled at source only (DPS)
  - Try to persuade DPS to do duplication at the parallel end (on-FPGA), using additional MGTs
  - Possible fall-back: optical fibre splitting, if optical power budget permits (probably more expensive and bulky than parallel duplication)
- Fan-out in phi handled at destination only
  - Consider passive electrical splitting of 10Gb/s signals
  - Active signal fan-out would compromise design density
  - Fall-back: retransmission on parallel links → latency penalty of more than 1 tick (SelectIO SerDes).

Floor plan of jFEX module

- ATCA
- 8 processors
- 4 microPODs each
- Passive fan-out
- Small amount of control logic
- Electrical backplane basically unused

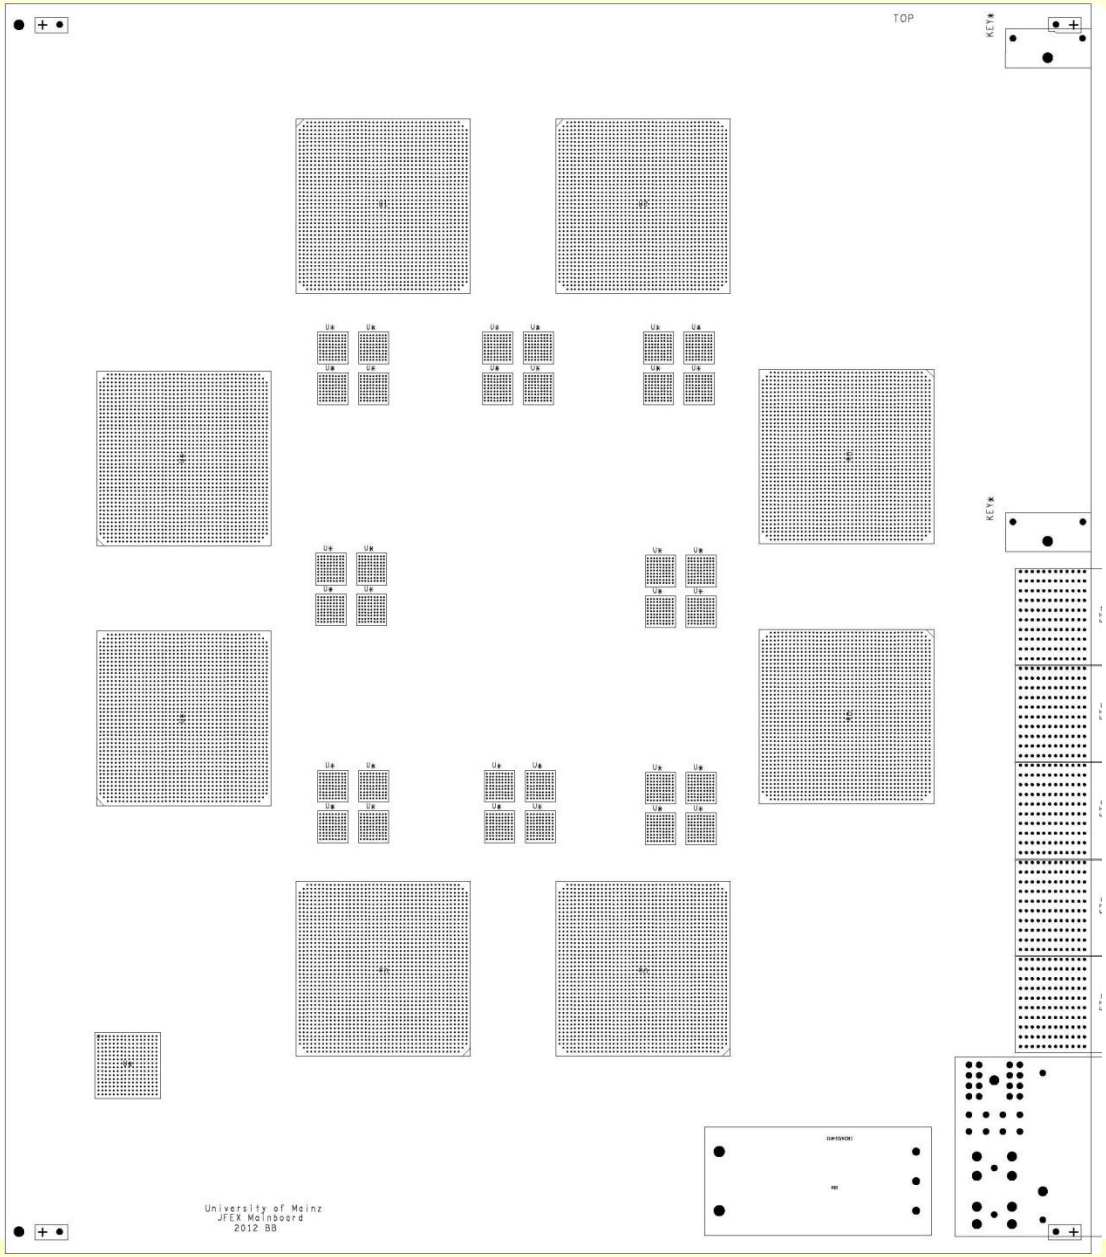

Bandwidth vs. granularity

Let's do the math for the suggested scheme. Go for the extreme. Do a rough estimate:

- Environment 0.9×0.9
- Each FPGA receives fully duplicated data in eta and phi: 1.6×1.6 worth of data required for a core of 0.8×0.8
- 256 bins @ 0.1 η×φ
- Maximum aggregate bandwidth (payload @ 10Gb/s line rate) of chosen FPGA is 80*8Gb/s=640Gb/s
- 640/256 Gb/s per bin = 2.5Gb/s, ie. 62 bit
- Assuming BX multiplexing, that's basically doubled to 125bit, that's probably more than 8 energies per bin

→ We will be able to route highly granular data into the jFEX processor chips, at large environment

Fibre count / conclusion

- Due to full duplication in phi direction, exactly half the signals are routed into the modules on fibres
  - 320 fibres
  - 27 × 12-channel opto receivers
  - 7 × 48-way fibre bundles / MTP connectors

- The 8-module jFEX seems possible with ~2013's technology
- Allows for both excellent granularity and large environment
- Can DPS deliver the required data ?
- How to arrive at sufficiently dense fibre contents for phase 2 (eta orientation of tile RODs ?) ?
- Is passive electrical splitting feasible ?

→ Start to explore technologies and feasibility soon