Trigger Hardware Development

Modular Trigger Processing Architecture

Matt Stettler, Magnus Hansen
CERN
Costas Foudas, Greg Iles, John Jones
Imperial College

3/2007

Traditional Trigger Processors

- Generally pushed the state of the art
  - In speed and density
  - Large hardware designs
  - This will not change in the future
- Fully custom one-off systems
  - Complex and difficult to modify
    - Difficult to maintain over the life of the experiment
    - Designer's knowledge difficult to pass on
  - Lack clean internal interfaces
    - Modular to a degree, but based on ad-hoc bus schemes
  - Generally don't scale well
    - Dataflow defined by hardware architecture
    - Inflexible by design

Current Systems

- FPGA technology adopted rapidly
  - Current FPGAs have sufficient performance
    - Speed and density
  - Allow considerable flexibility in logic design
    - Very large devices available
    - Problem becomes physical concentration of data
- Dataflow is still hardware based
  - Some designs have used FPGAs as data routers
    - Inefficient at this task
    - Large number required increases design risk considerably
  - Many simply define the dataflow in hardware
    - Large parallel buses
- Large circuit boards required
  - Expensive to develop in both time and money

Desired Attributes for new Systems

- Modular
  - Reasonably fine grained
    - Smaller circuit boards
    - Easier and less expensive to develop
  - Should scale well
- Well defined internal interfaces
  - Allows modules to be developed independently
    - Perhaps shared across projects
  - At least electrically compatible
- Flexible in both logic and interconnect
  - Retain flexibility of FPGAs
  - Add complementary data routing flexibility
    - Modify data flow without altering hardware
    - Perhaps give the ability of dynamic reconfiguration

Relevant Current Technology

- Large FPGAs with built in SERDES links
  - Very high density logic
  - SERDES I/O allows physical concentration of data
    - 16-24 3+Gbps links available on large devices
    - Normal I/O pins support 1+Gbps links
- LVDS/CML crosspoint switches
  - Up to 144x144 non blocking matrices
  - Asynchronous, Protocol agnostic
    - Multi rate switching supported intrinsically
  - Currently available at 4+Gbps, 10Gbps announced
- Serial Memory becoming available
  - FBDIMM modules designed for server market
    - Support very large arrays (>100GB)
    - Electrically compatible with LVDS/CML

Proposal for CMS HCAL Muon and Quiet bit distribution

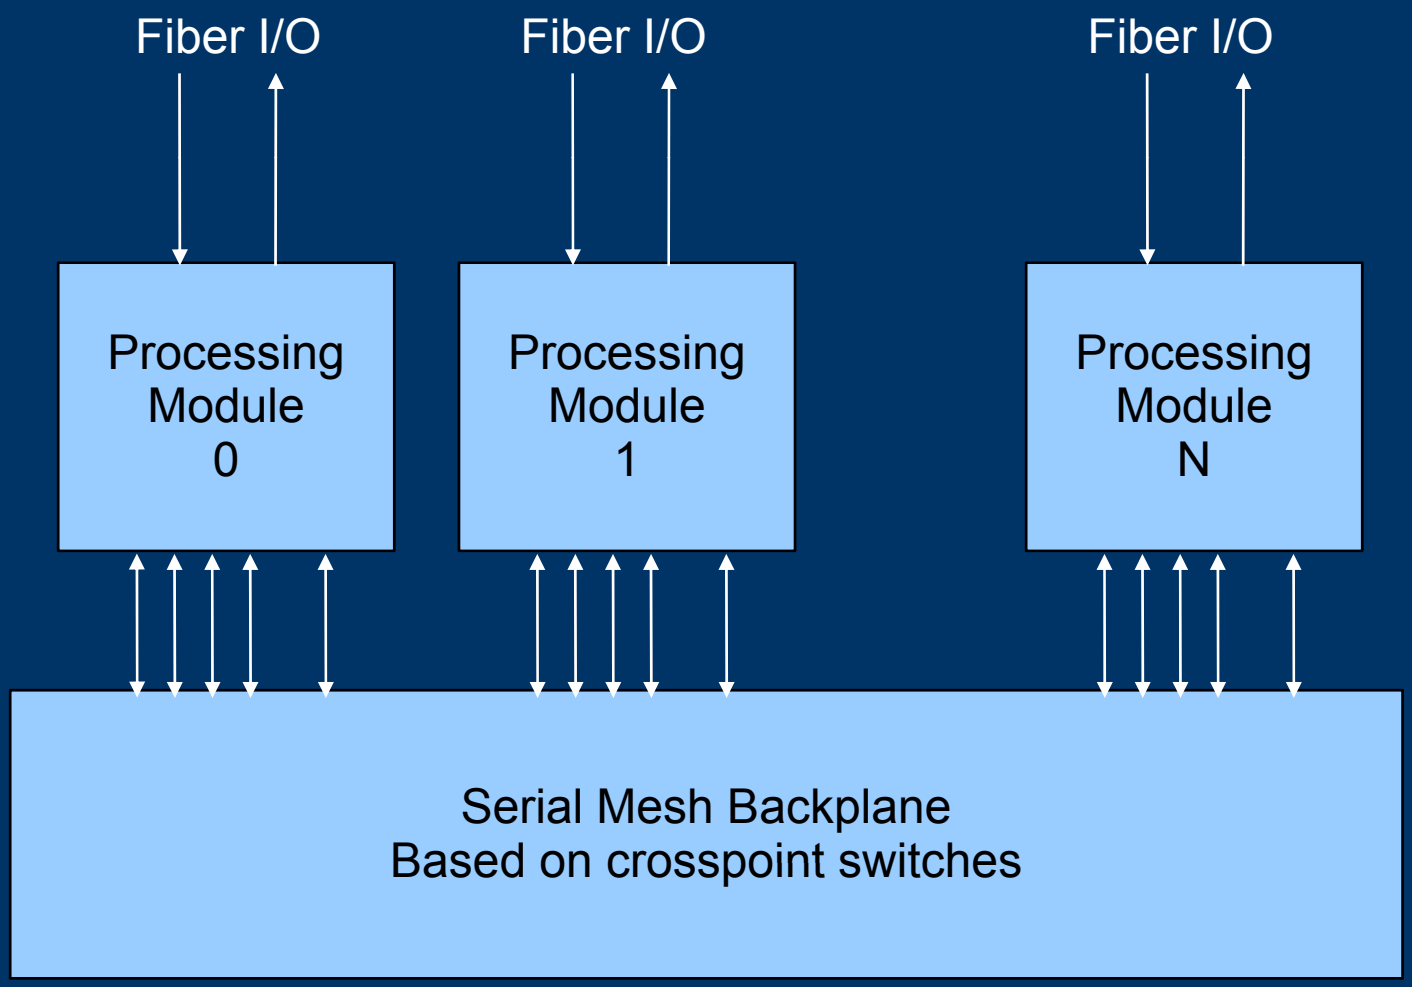

Implementation

- Proposed for CMS Global Calorimeter Trigger HCAL muon & quiet bit distribution
  - This system is ready for layout now
- Micro TCA based, 12 slot chassis
  - 20 high speed links per slot
- Processing module
  - Based on Xilinx V5LX110T
  - Mindspeed M21141 72x72 crosspoint
  - MFP Fiber interface (SNAP-12)
- Backplane
  - Commercial unit possible, but limiting
  - Custom unit based on Mindspeed M21161 144x144
    - Programmable topology
    - 500 Gbps maximum bandwidth

Micro TCA Crate/Backplane

Processing module Block diagram

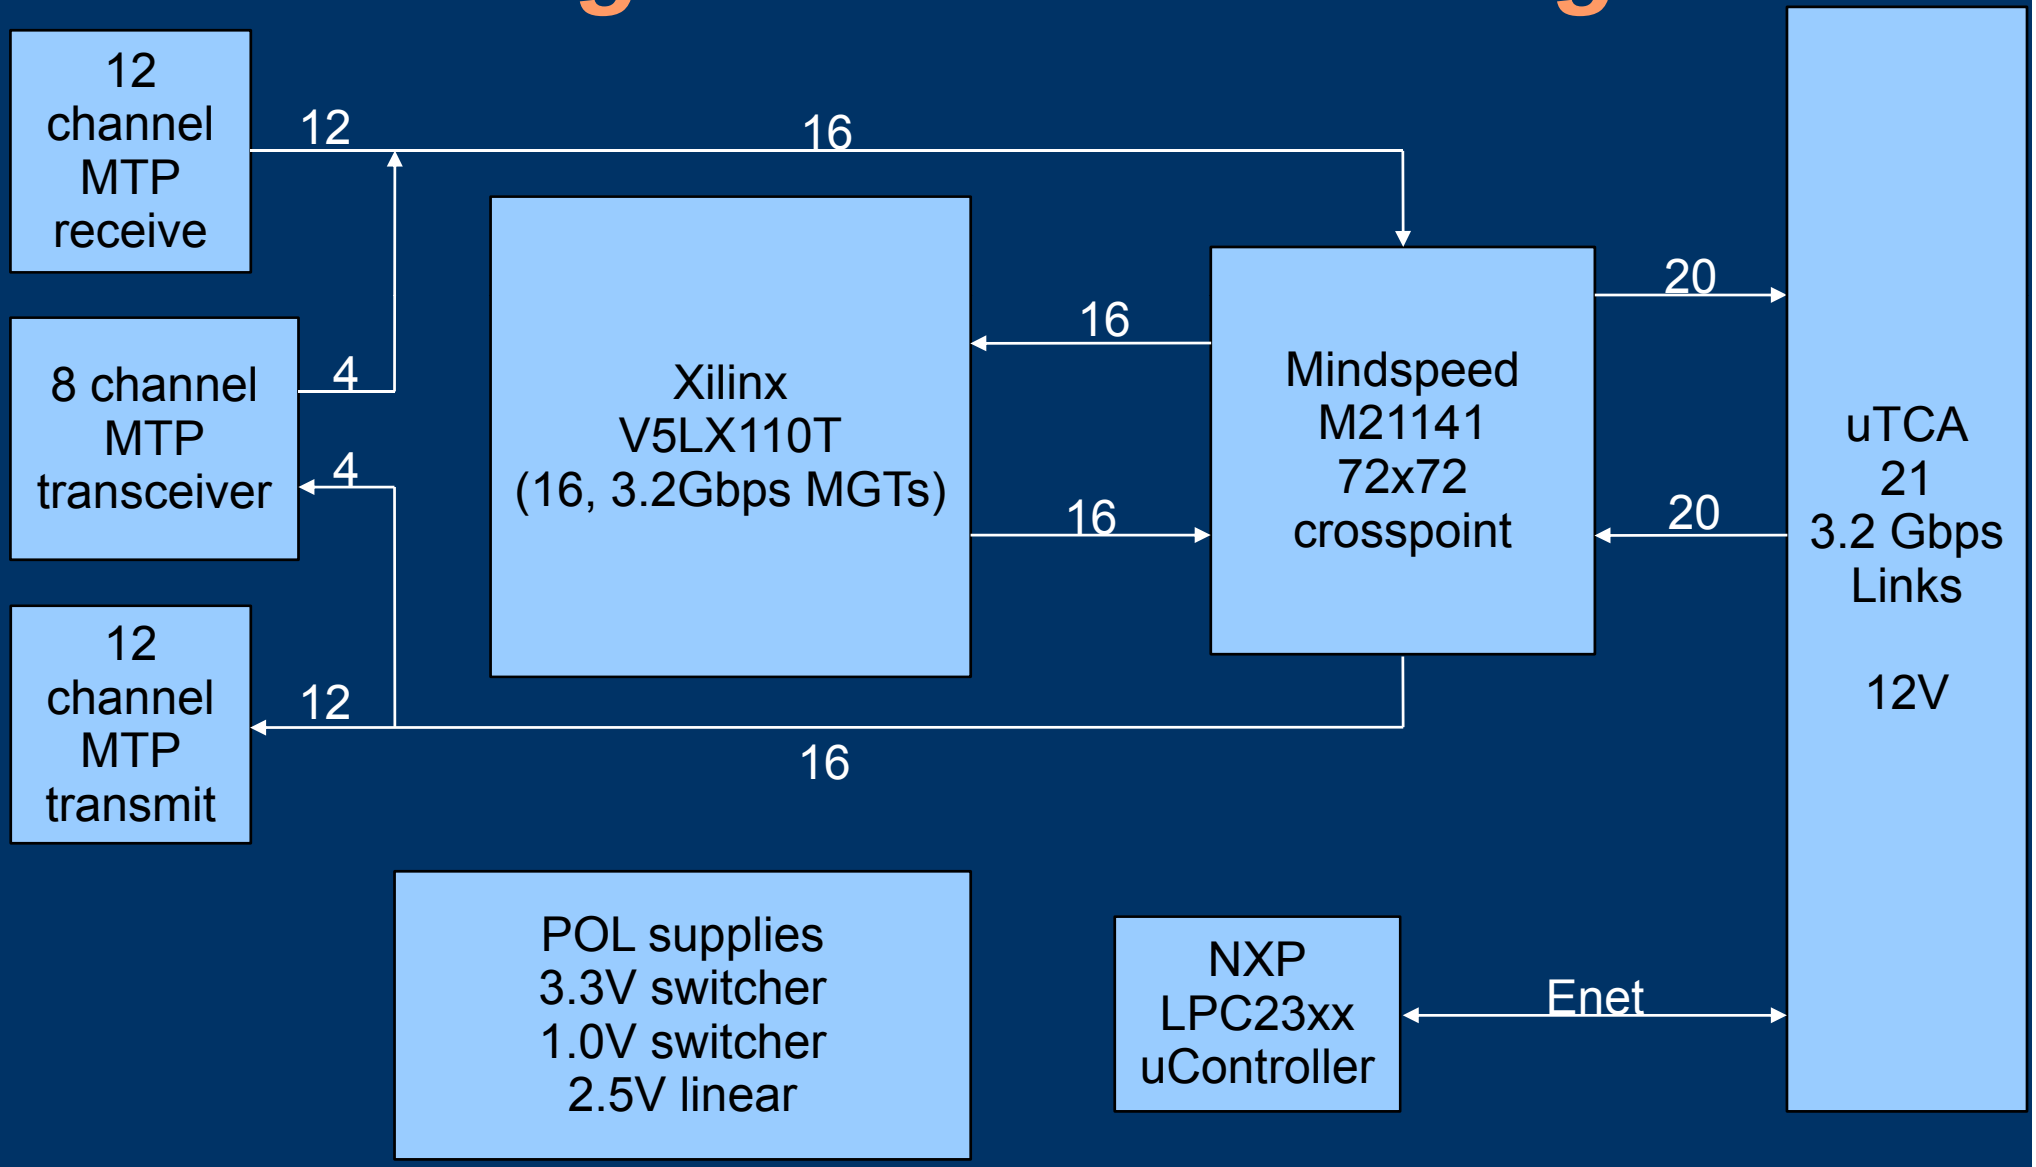

Key Features of Processing Module

- Xilinx Virtex 5LX110T FPGA
  - 16 low power MGTs, Superior logic density/speed
  - Standard I/O 1Gbps capable
- Data I/O direct from fiber
  - 16 channels full duplex @ 3.2 Gbps
  - Fiber utilized for inter-crate data links
  - Same fiber modules used on previous system
- Crosspoint routes to FPGA, backplane, and fiber
  - 1:N data replication supported at wire speed
  - FPGA output data sent to backplane or fiber
  - Data sharing between systems intrinsically supported
- Slow control via standard Ethernet
  - Based on NXP 23xx uController
  - Standard Ethernet protocols

Custom Backplane block diagram

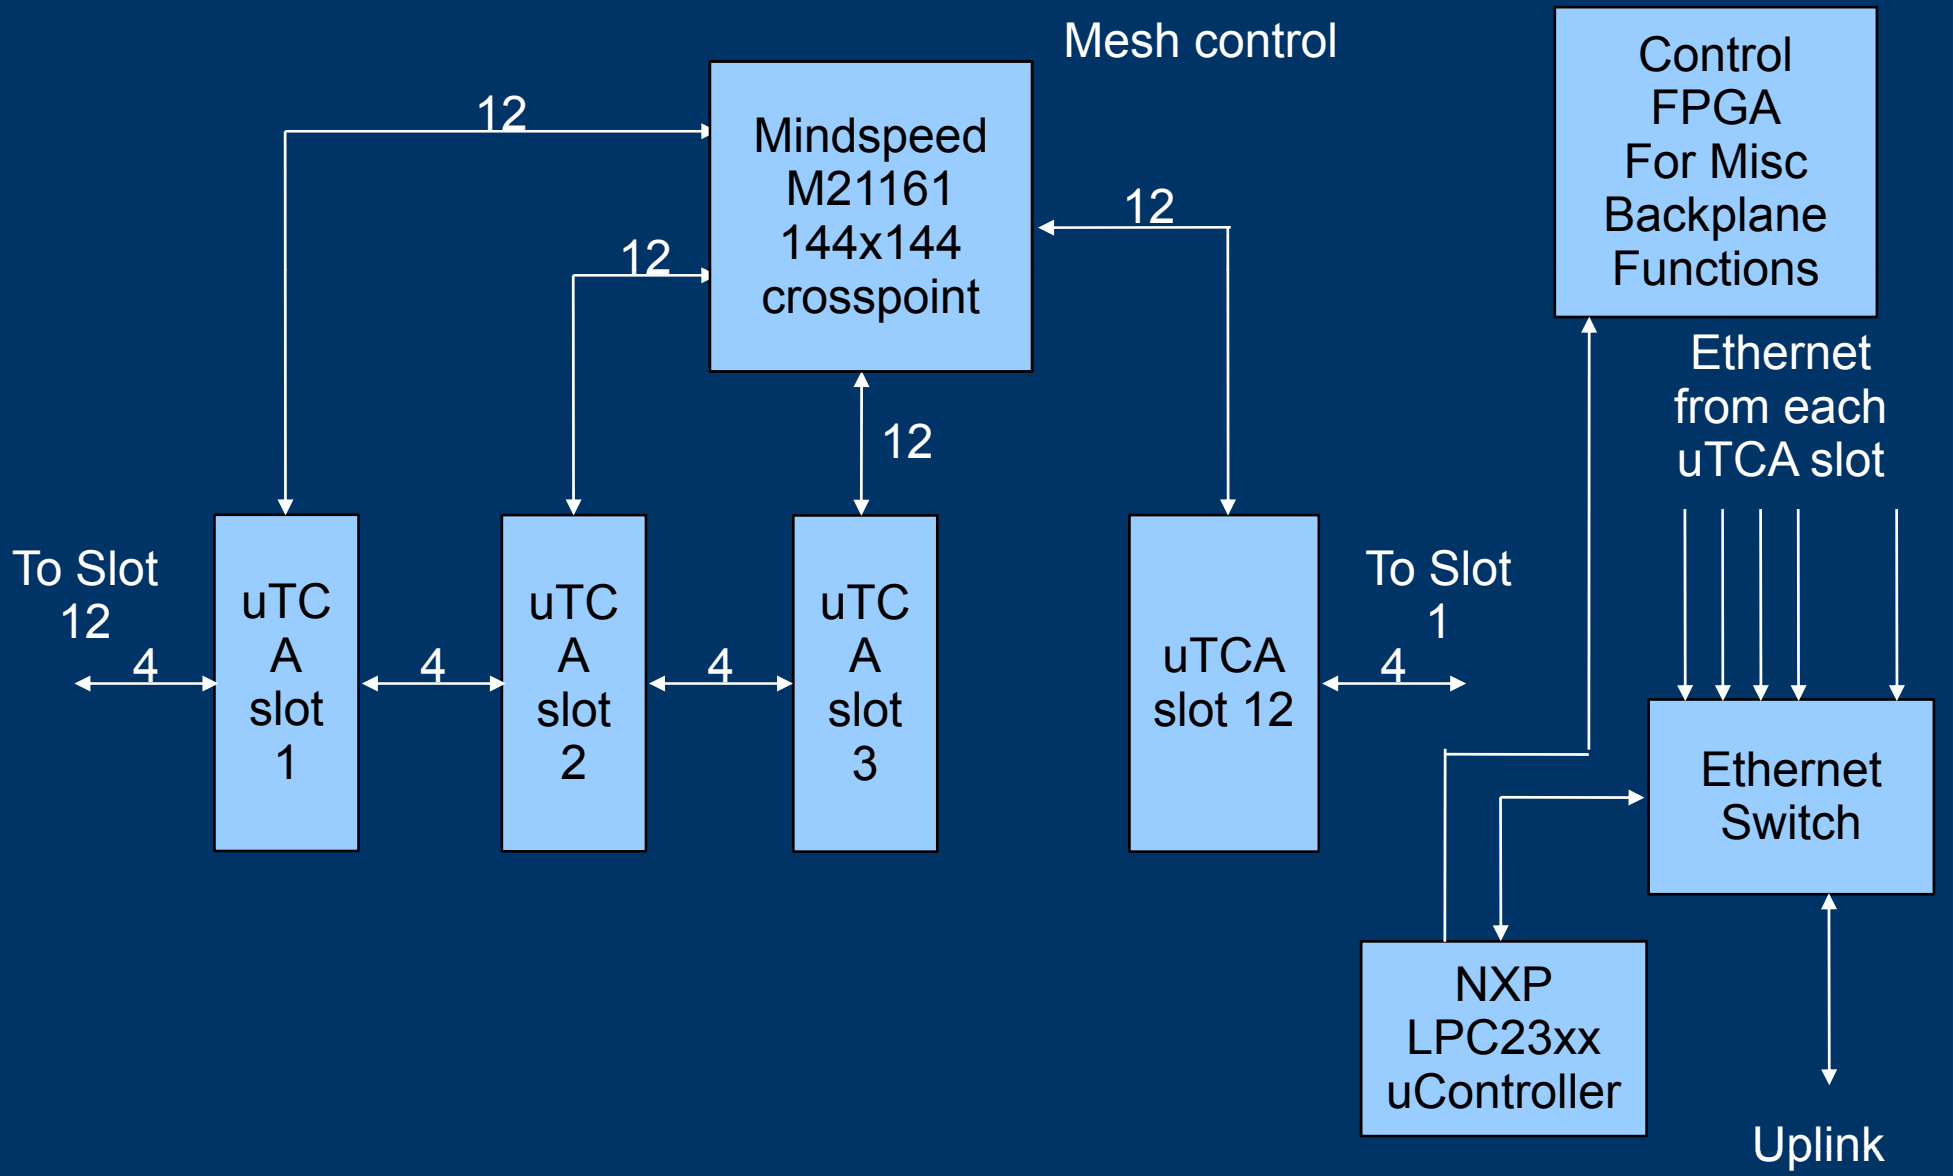

Key Features of Custom Backplane

- Active backplane based on crosspoint switches
  - Mindspeed M21161
  - Active backplane eliminates congestion
    - uTCA Hub slots cannot support full mesh connections
- Programmable topology
  - Mesh, star, etc based on crosspoint configuration
- 20 links/slot for high speed data communication
  - M21161 supports 4.1Gbps data rate
    - System will be limited to 3.2 Gbps by FPGA/MFP Optics
- 1 link/slot for 100Mb Ethernet control interface
  - Star topology interface, switch IC on backplane
  - Backplane configuration via Ethernet also
    - Same uController as processing module

FBDIMM Architecture
(future Interest)

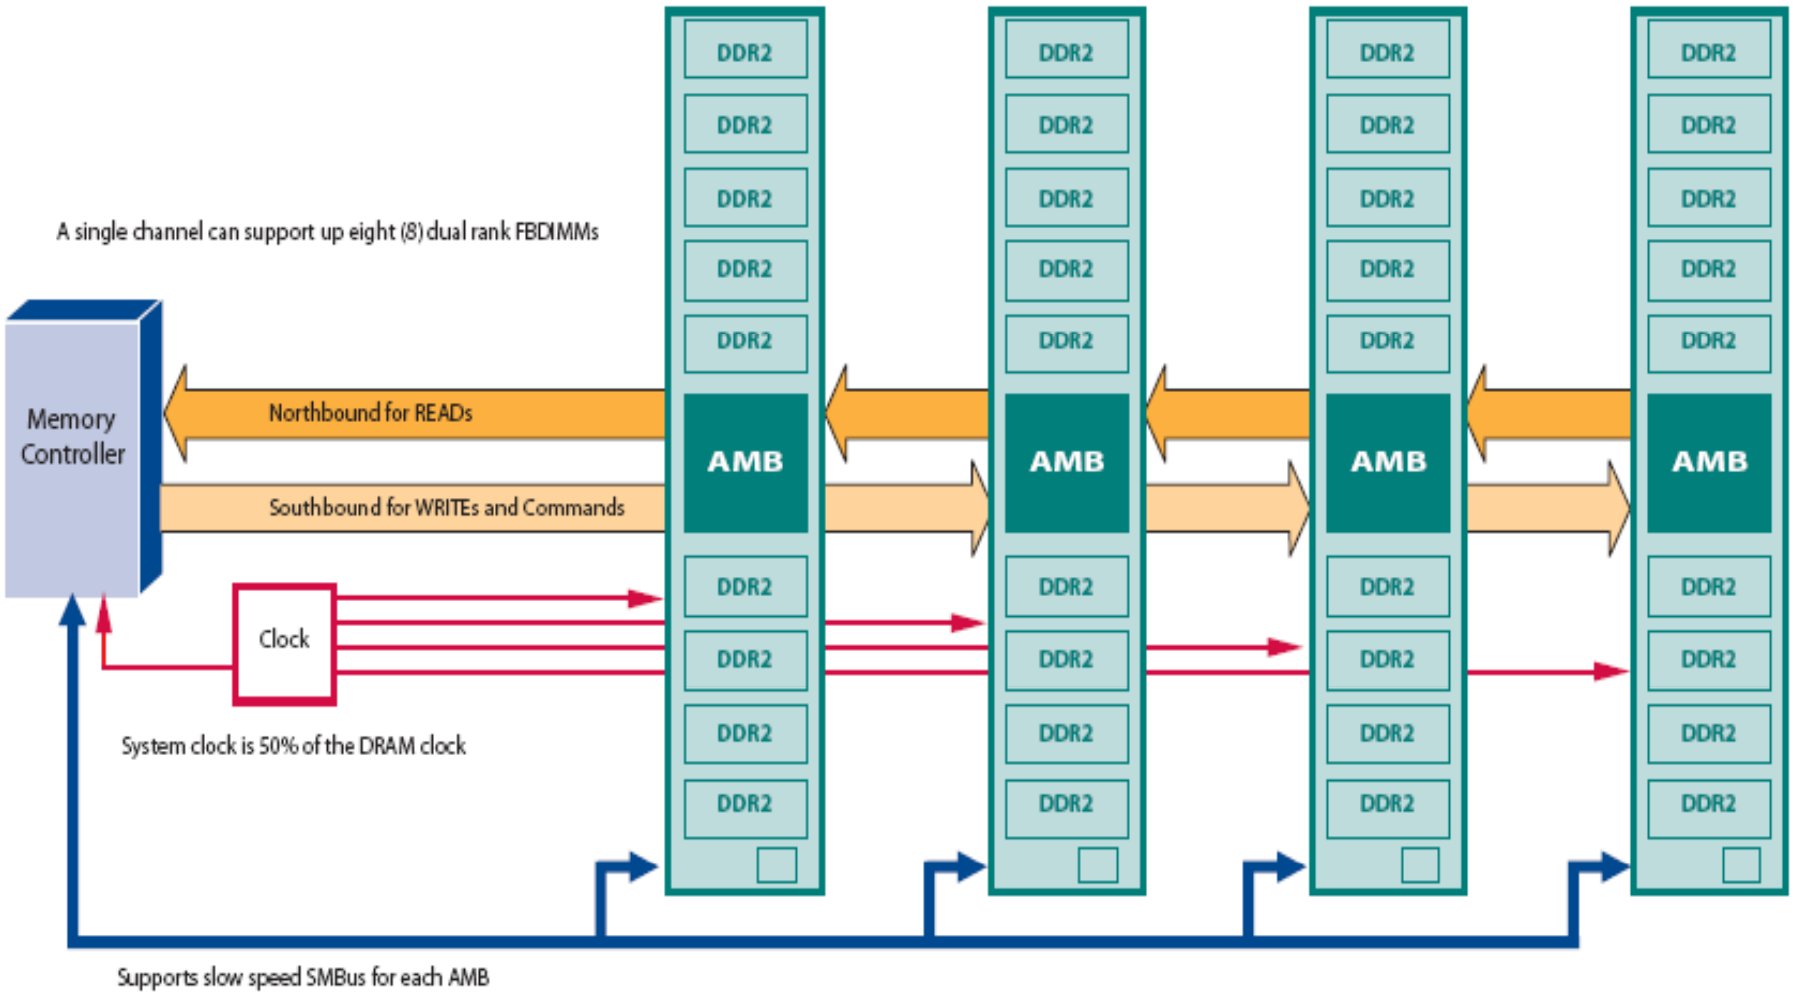

Key Features of FBDIMM

- Serial memory is designed for large blocks
  - Packet based logical model
  - Optimized for block transfers
- Utilize Advanced Memory Buffers (AMB)
  - Control local banks of DDR2 devices
  - Create daisy chain of banks
- Interface to host is via CML links
  - Currently working at 3.2 Gbps, eventually up to 9.6
  - Utilize 12-14 full duplex links
- Up to 8 AMBs can be daisy chained
  - Achieves 192GB with current technology
- Can be combined with serial mesh
  - Very large distributed memory arrays possible

Serial Link Latency

- Historically the limiting factor
  - Deserialization has dominated latency
  - This has improved with new FPGA families
- Current performance (Virtex5)
  - 10-12 byte clocks for receive
    - This is 2.5-3 40MHz clocks in our implementation
  - 5-7 byte clocks for transmit
    - 1-2 40MHz clocks
- Still significant
  - Biases FPGA choice to fewer large devices
  - Number of “hops” needs careful attention
- Crosspoints do not contribute to latency
  - Nearly wire speed (ps delay) - insignificant
  - No serialization/deserialization