

COMBOv2 - Hardware Accelerators for High-Speed Networking

Jiří Novotný, Martin Žádník

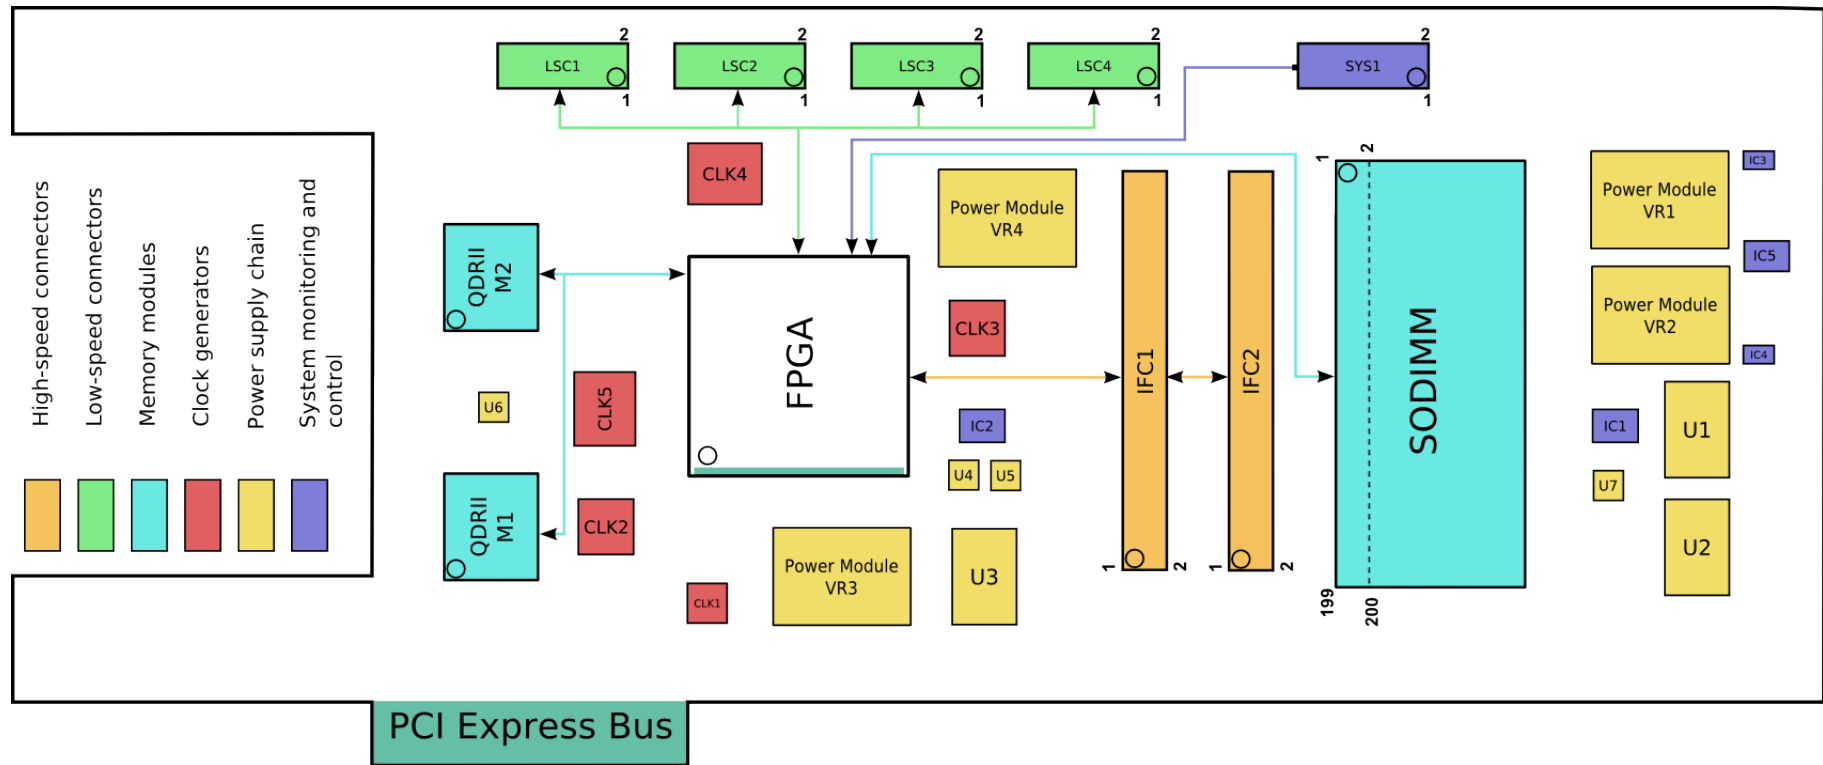
novotny@ics.muni.cz, izadnik@fit.vutbr.cz



- **Essential design objectives**
 - ◆ develop brand-new family of HW accelerators based on experiences gained during the development of previous generation COMBO cards
 - ◆ prepare architecture concept with straightforward extension path and superior performance characteristics
 - ◆ ensure flexibility and mutual compatibility throughout COMBOV2 family
 - ◆ COMBOV2 family lifespan expectancy is **5 to 7 years**
- **Acceleration platform optimized for networking domain**
 - ◆ built around main processing board and collection of add-on cards
 - ◆ high-performance at **10Gbps** minimum, with the imminent perspective of **100Gbps** data processing speed
 - ◆ quick and easy implementation of new algorithms
- **Modular system consists of**
 - ◆ mother cards
 - ◆ interface cards
 - ◆ low speed cards

- **Powerful system with rich connectivity features**
 - ◆ challenging computations with FPGA Virtex-5, QDRII and DDR2 memory
 - ◆ attachment of high-speed network interfaces or general-purpose data links
- **Technology highlights**
 - ◆ adoption of Virtex-5 FPGA devices
 - ◆ 8-lane PCI Express communication bus (up to 16Gbps in each direction)
- **Unique design flexibility**
 - ◆ single PCB layout supports various members from Virtex-5 LXT family
 - ◆ LX50T, LX85T, LX110T or LX155T FPGA can be selected to match the actual design needs or target production cost
- **Sophisticated boot and management solution**
 - ◆ diverse initialization, configuration and monitoring tasks in one place
 - ◆ system based on small FPGA Spartan-3E, commodity FLASH and PSRAM

Mother Card - overview



Based on Virtex-5 LXT FPGA family

Memories

- 2x QDRII memories
- SODIMM connector – DDR2 DRAM, RLDRAM

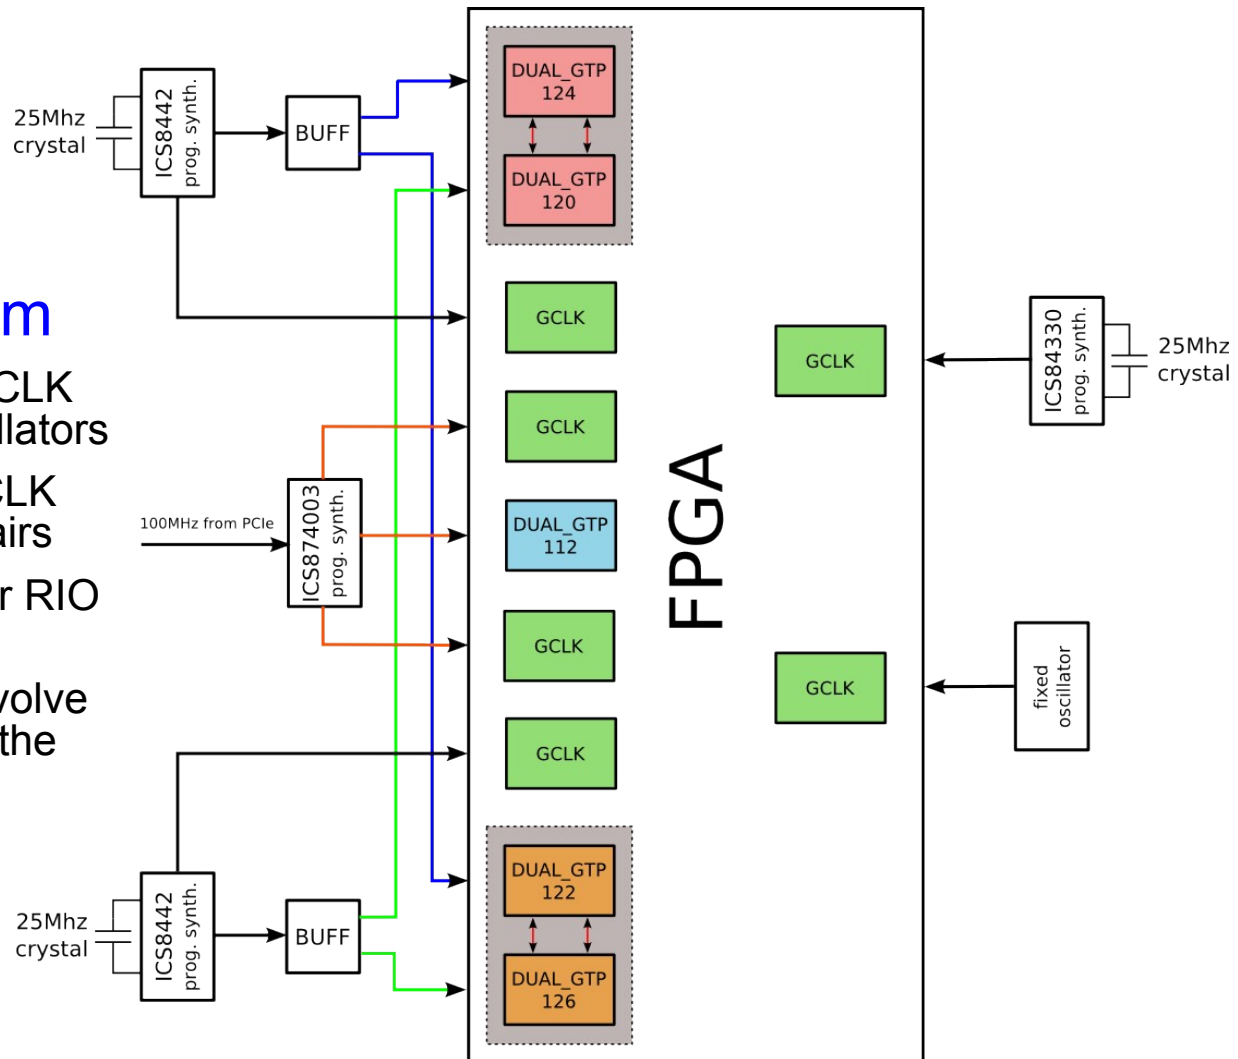
Interfaces

- high speed interconnection (IFC)
- low speed connector (LSC)
- system monitor (SYS)

Mother Card - clocking system

Flexible clock system

- Low-jitter programmable CLK generators and fixed oscillators
- On board distribution of CLK signals with differential pairs
- Dedicated CLK signals for RIO blocks and FPGA fabrics
- System operation may involve several frequencies from the interval **100-500MHz**



QDRII SRAM memory



- Available capacity **144Mbit** in two memory chips, where the PCB layout actually supports **576Mbit** with future chips
- Working frequencies up to **250MHz** for DDR data transfers
- Theoretical bandwidth may reach **18Gbps** at 250MHz with low latency
- Typical utilization:
 - ◆ **fast data buffers with low latency**
 - ◆ **routing tables**
 - ◆ **state memory**

SODIMM memory slot



- **DDR2 DRAM** modules with **2GB** of maximum capacity
- Working frequencies up to **267MHz** for DDR data transfers
- Theoretical bandwidth may reach **32Gbps** at 250MHz
- Custom **RLDRAM** modules with **864Mbit** of storage space can be plugged into SODIMM slot
- Typical utilization:
 - ◆ **temporary packet storage**
 - ◆ **flow measurement informations**
 - ◆ **fast packet FIFO with RLDRAM modules**

Specification

- 4 RocketIO channels, 36 LVDS pairs, 4 CLK pairs, 14-bit wide management bus
- 8 PWR extension pins for add-on cards power distribution
- LVDS pairs are symmetrically placed in two I/O banks of FPGA, RIO channel bonding
- Separation of differential pairs with ground for better signal integrity
- Theoretical data bandwidth of 28Gbps in each direction with a single IFC connector
- Mother card includes two IFC connectors

Typical applications

- Four 1Gbps interfaces connected through RIO lines
- Up to 16 interfaces at 1Gbps connected through LVDS pairs
- One 10Gbps XAUI interface using dedicated RIO lines
- One 10Gbps interface connected through LVDS pairs
- Connection of external boards for improved system scalability

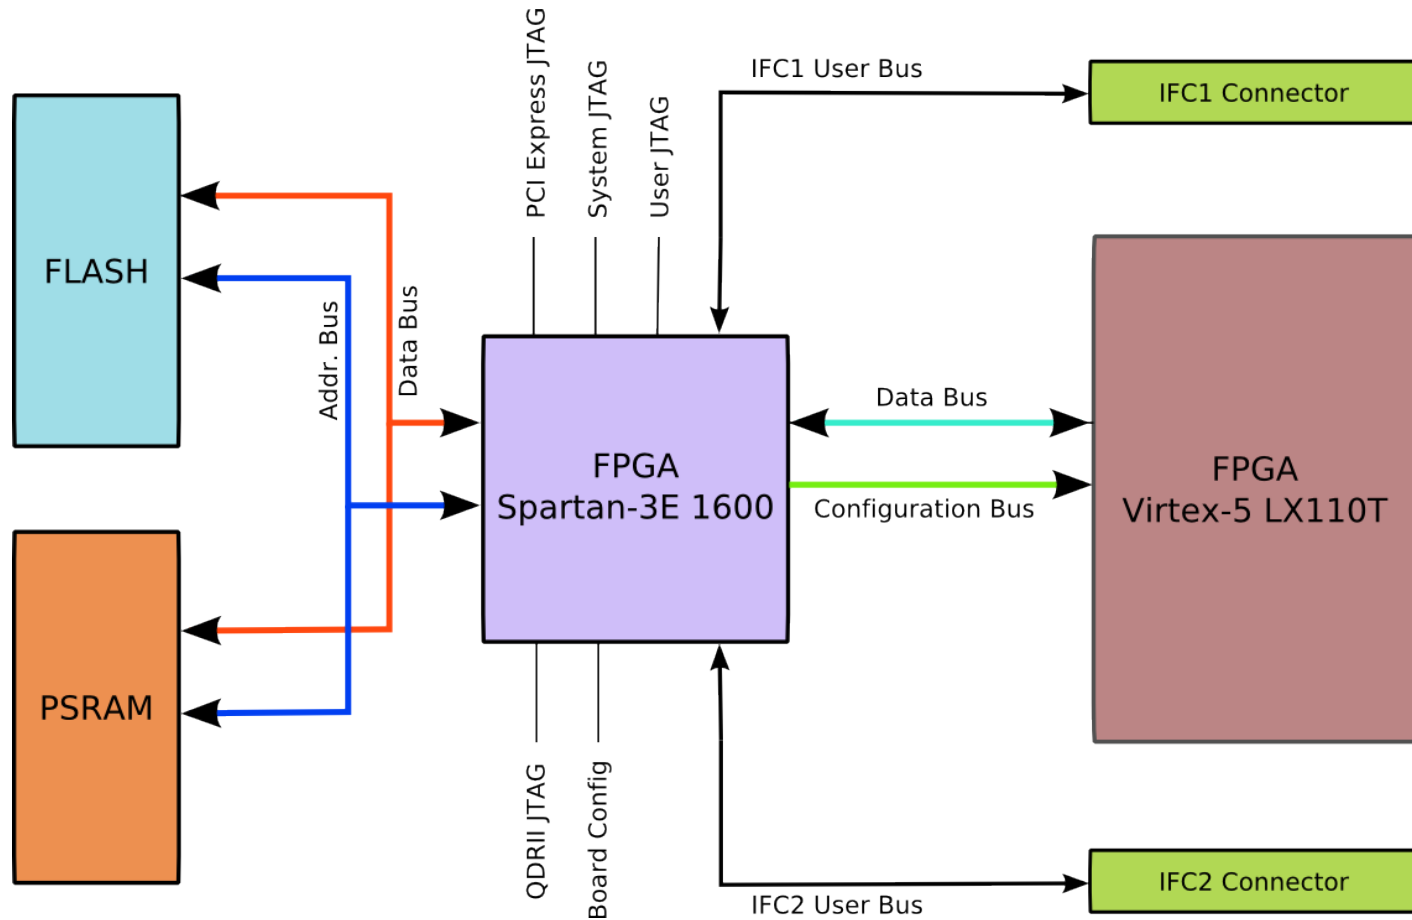
Specification

- 10 LVDS pairs for custom usage, 1 auxiliary pair
- 8 PWR extension pins for add-on modules power distribution
- Separation of differential pairs with ground for better signal integrity
- Theoretical data bandwidth of 4Gbps in each direction with a single LSC connector
- Mother card includes four LSC connectors

Typical applications

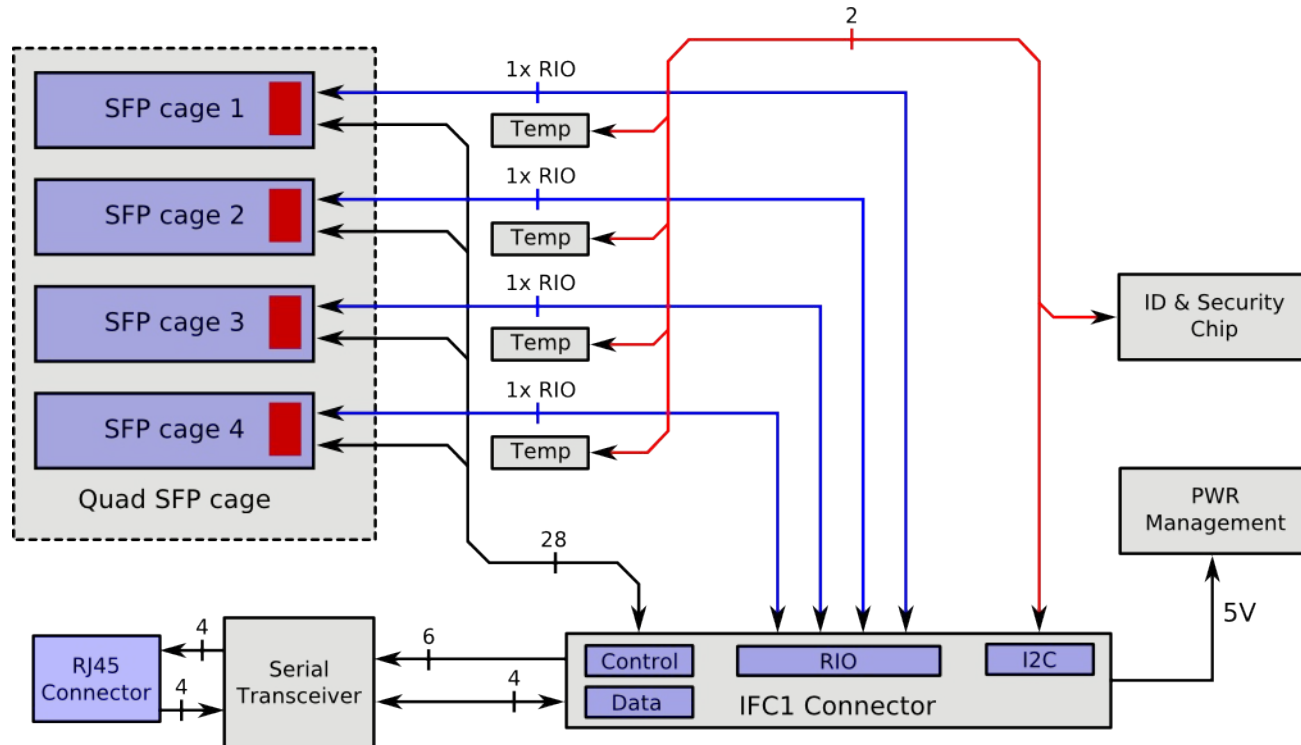
- Up to 4 interfaces at 1Gbps in each direction
- Up to 8 interfaces at 1Gbps for network traffic interception
- Time-stamp module connection
- SATA module connection
- Other peripheral devices

Booting and Configuration



- Spartan-3E takes care of the boot phase, reconfiguration and system management
- FLASH memory keeps bitstreams for Spartan-3E, Virtex-5 or even add-on cards
- Spartan-3E is booting in fallback BPI mode from FLASH and then configures Virtex-5
- Rapid design update through PSRAM+UserJTAG and background transfer to FLASH

1Gbps Interface Card



Attached to mother card by IFC connector

Interfaces

- Quad SFP cage – 4x 1Gbps interface
- RJ45 connector for GPS connection

Others

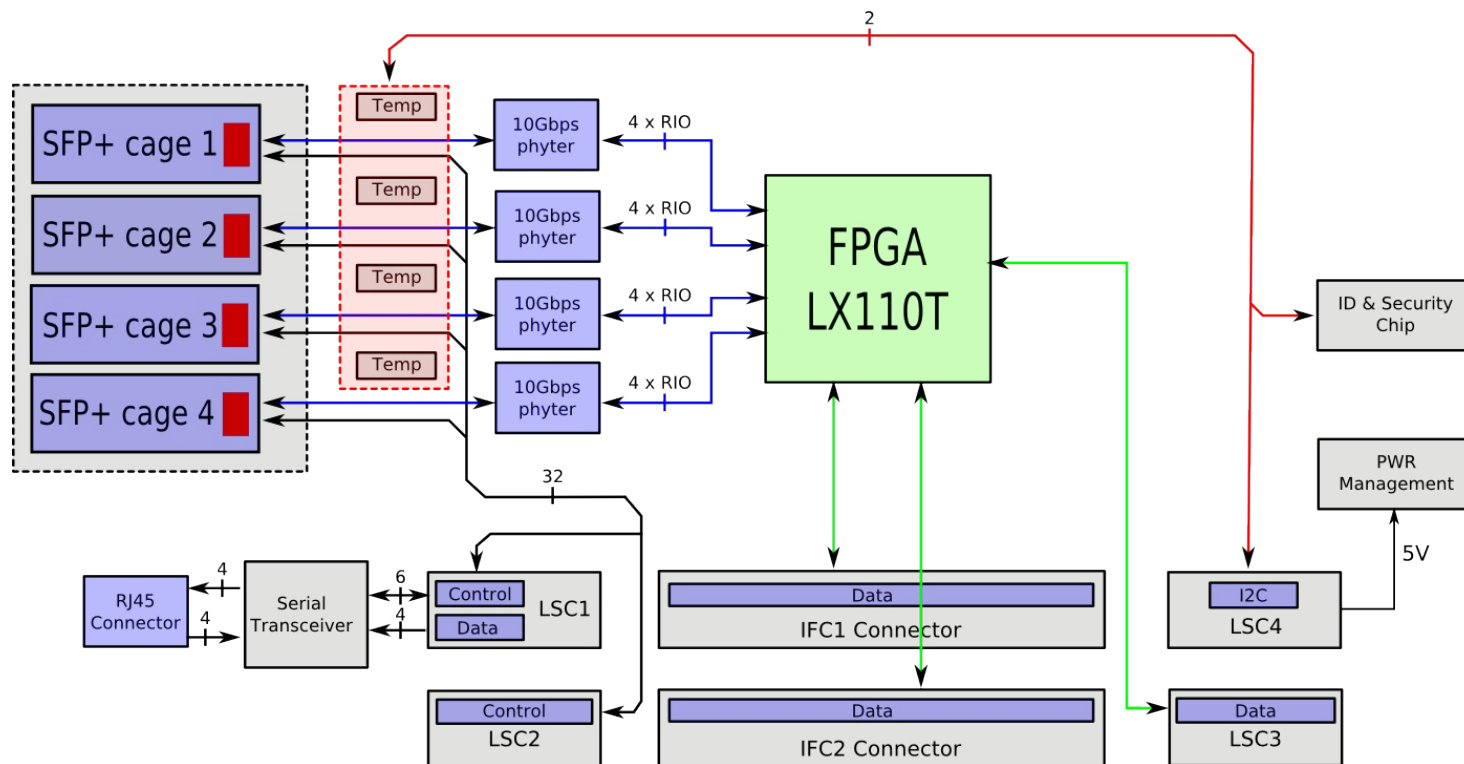
- Temperature monitoring
- Card identification and security



Others

- Temperature monitoring
- Card identification and security

10Gbps Interface Card – part 2.



Attached to mother card by two IFC and four LSC connectors

Interfaces	Others
------------	--------


- Four XFP cages – 4x 10Gbps interface
- RJ45 connector for GPS connection

- FPGA LX110T to transform Rocket IO to LVDS pairs
- Temperature monitoring
- Card identification and security



- FPGA LX110T to transform Rocket IO to LVDS pairs
- Temperature monitoring
- Card identification and security

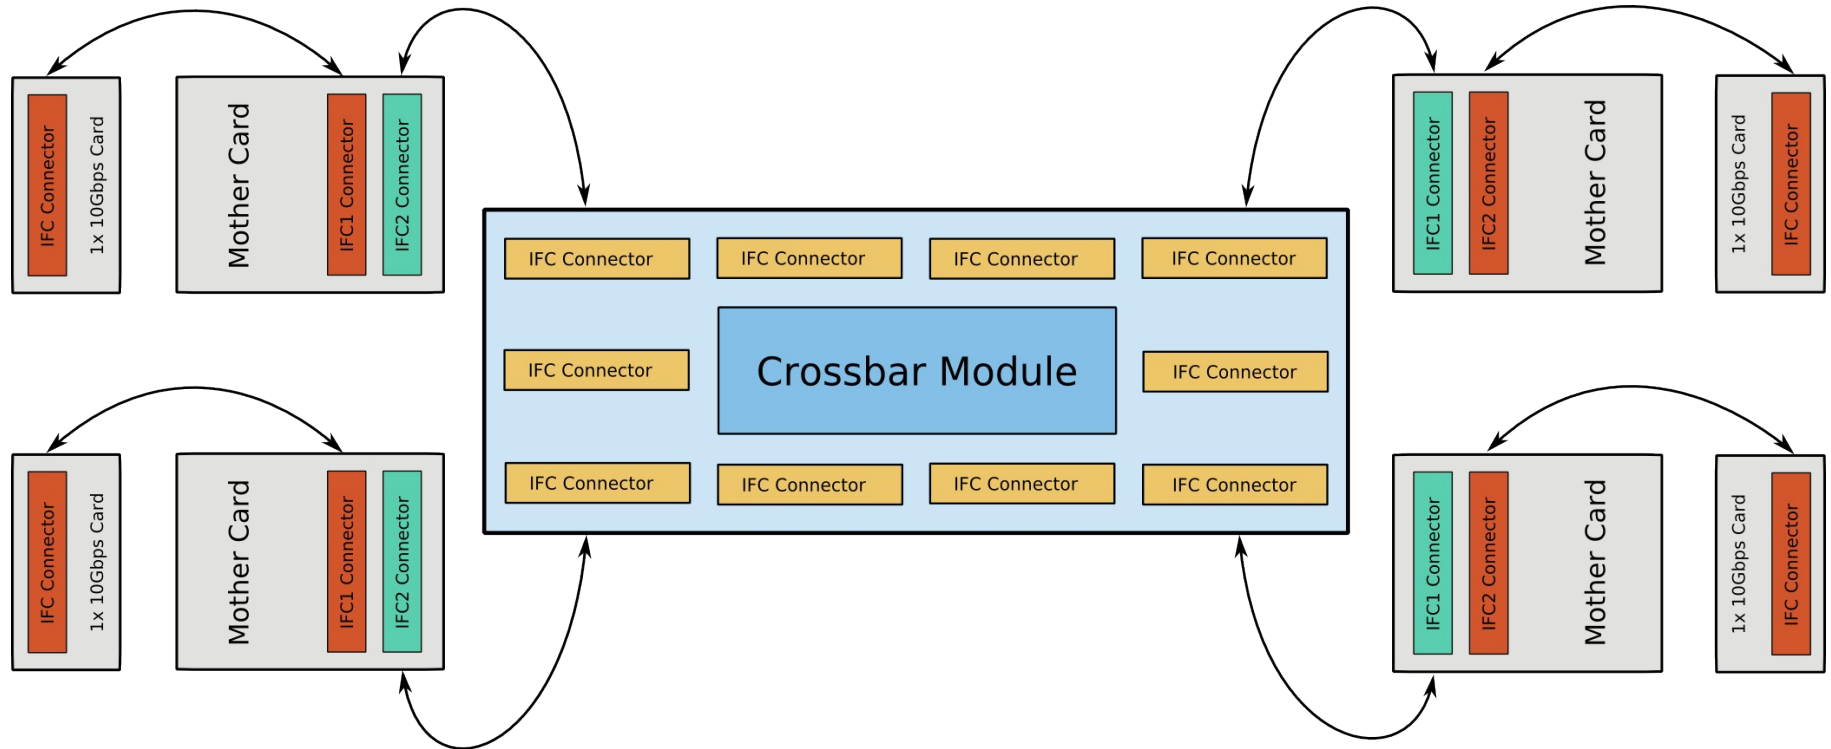
Chain of cards for stream processing



- Required number of mother cards can be linked together to get higher performance for specific applications
- Each section executes a partial task and the intermediate results are exchanged via high-speed IFC connectors

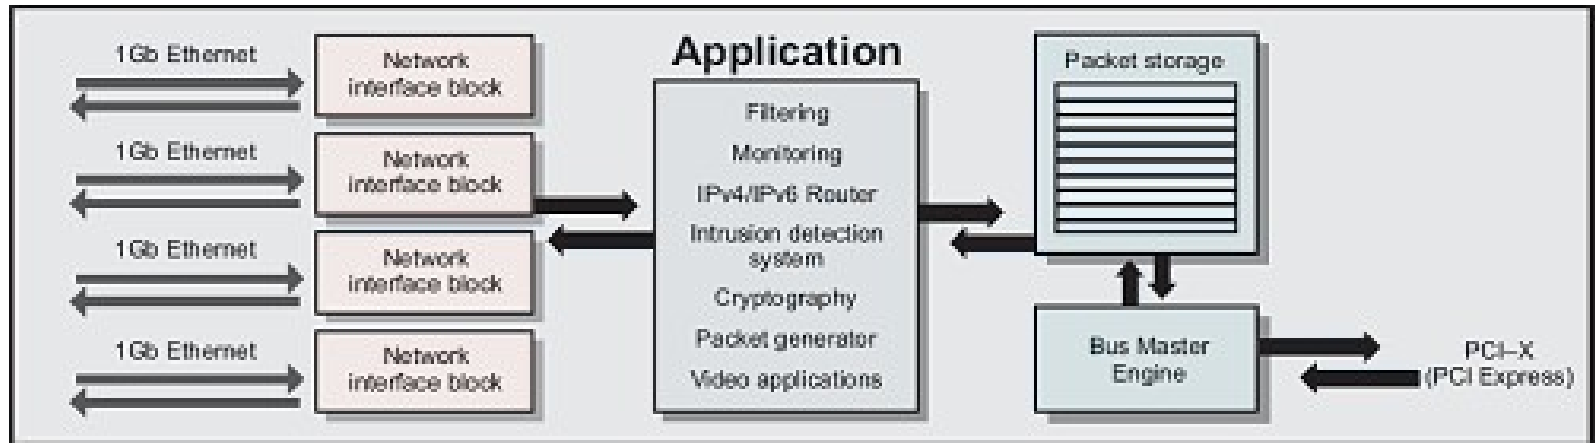
Platform scalability – part 2.

Direct connection using crossbar module



- The possibility to establish 10Gbps direct communication channel between any pair of mother cards
- Switching core based on dedicated ASIC or FPGA-like technology
- Rapid data exchange and cooperation for large-scale applications

- High performance scalable platform for rapid development of FPGA applications



- Basic Features

- 1 G and 10 G Ethernet support
- PCI, PCI-X and PCI Express x1, x4, x8 support
- Vendor abstraction layer
- Customizable device drivers

- Supported Boards



Xilinx ML555



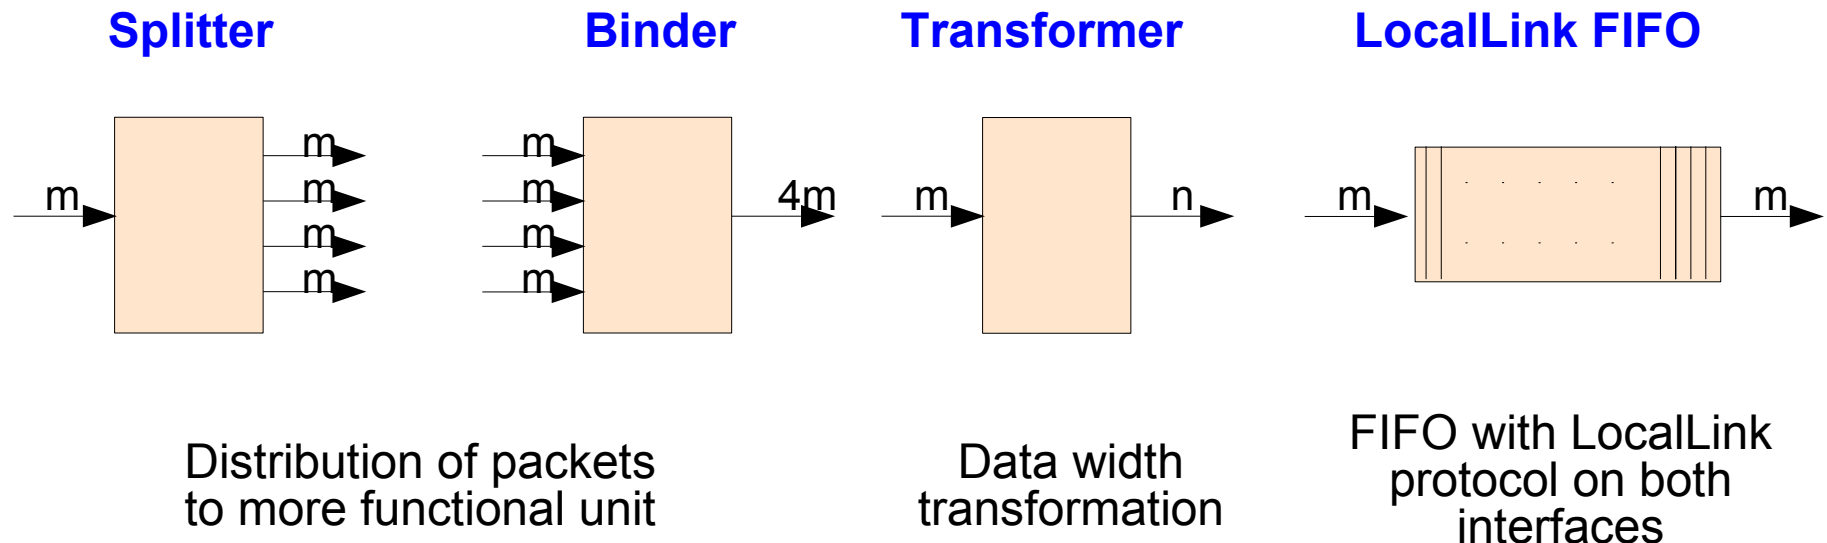
COMBO6X

COMBOv2

- Connection to PCI, PCI-X and PCI-Express
Support of **Bus Master DMA transfer** initialized from adapter
 - ◆ Controlled by **PowerPC processor** in Virtex-II technology
 - ◆ **Implemented in logic** for Virtex 5 technology
- Prepared packet reception and transmit buffers with LocalLink ifc.
 - ◆ NetCOPE takes care of the transfers between card and host memory
- All data transfers between FPGA components and the host RAM provided by prepared interconnection system
 - ◆ **High throughput Internal bus** with full PCI Express throughput
 - ◆ **Resource-saving Local bus** allows to connect large number of components

- Based on LocalLink protocol
 - ◆ Header, Payload, Footer – user data with control information
 - ◆ Configurable data width: 8, 16, 32, ...

Set of prepared components



- UNIX based operating systems (Linux, *BSD)
 - ◆ Standard network interface
TCP/IP stack or PCAP library support
Operating system and software application accesses the FPGA board like a common NIC card
 - ◆ Application specific interface
Transfers of arbitrary data frames between card and host memory
Examples of usage:
 - Packets with additional control data (for example Time-Stamps)
 - Aggregated frames for effective DMA transfers
 - Application data – statistical information about network traffic
- Hardware abstraction – independence of target platform
- Drivers are under GPL licence and can be changed to support end user application

- ★ Advanced router for up to 40Gbps networks
- ★ L2 switch with 2x 10Gbps ports and 16x 1Gbps ports
- ★ Load balancer from 2x 10Gbps ports to 16 different hosts
- ★ QoS measurements for up to 40Gbps networks
- ★ Network monitoring and analysis system
- ★ Intrusion detection system
- ★ And many others...



Thank you for attention!