FeTCP: Hardware-Based Light Weight TCP/IP for 10 Gigabit Ethernet

Jie Zhang, Xiaoshan Jiang
The State Key Laboratory of Particle Detection and Electronics,
Institute of High Energy Physics, CAS, China
E-mail: zhj@ihep.ac.cn

1·Motivation

- Benefit from technology development
  - 10/40 Gbit/s Ethernet network
  - New PCs with multicore CPUs
- New uTCA/ATCA architecture
  - Simple interface to custom readout electronics
  - Reliable and in-order data delivery
  - Available IP Core for Altera and Xilinx FPGA
- Data bandwidth requirement
  - up to 10 Gbit/s
- Commercial Products
  - ~~Price~~
  - ~~Maintainability~~

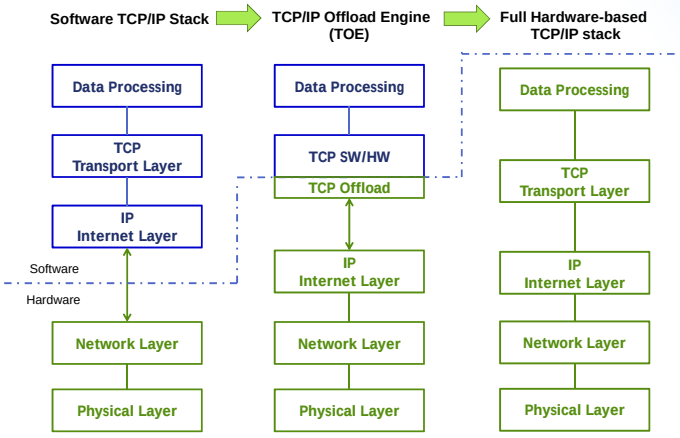

Comparison of software and hardware-based TCP/IP stacks

Benefits of Full Hardware-based TCP/IP stack

- Real-time transmission
  - As fast as hardware limit
- Resources consumption
- Clock synchronization
- BOM reduction
- Power optimization

2· IMPLEMENTATION

Simplification

- The main data flow is only in one direction from the front-end to the receiving PC.
- The DAQ network topology is fixed and the network bandwidth is enough to avoid packet congestion.

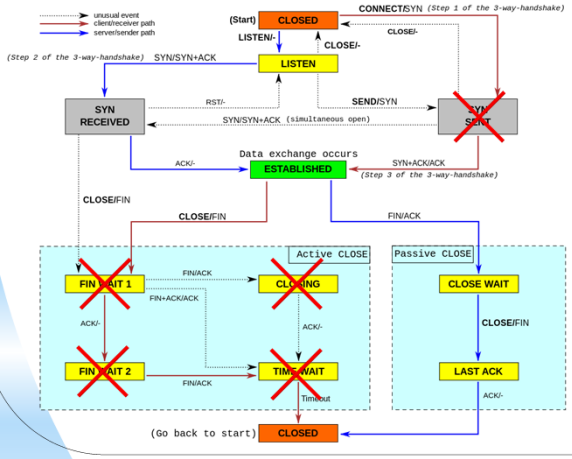

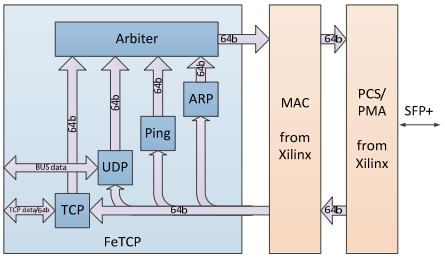

- Packets arriving from SFP+ transceiver enter the parameterizable IP solution PCS/PMA (10GBASE-R) and MAC (10GEMAC) block which is provided by Xilinx.
- The APR, Ping, UDP and TCP block decode the packet and check for validity (CRC, MAC, IP address and protocol match) independently.
- The ARP and Ping block only process ARP request and ICMP echo-request packets.
- The UDP block processes IPbus and controls the external bus.
- The TCP block is the main block and processes TCP. The FeTCP, acts as server, waits client to connect. Only one socket is accepted when the connection is established.
- The Arbiter block arbitrates and selects among sources and sends the packet to the MAC block.

- Looking from FPGA Side, FeTCP looks like a FIFO
- Looking from Client(PC) side, it looks like a Server

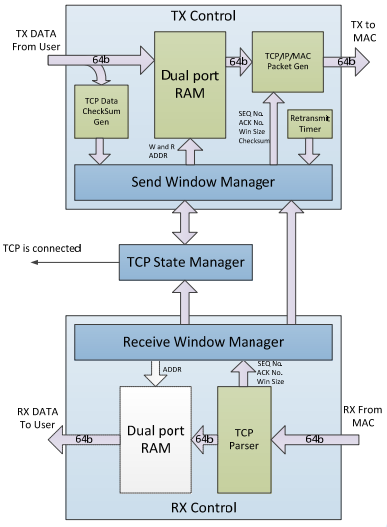

The block with dotted line is needed if we implement bidirectional transmission, unrealized by now

3· PRELIMINARY RESULT

- Full hardware-based TCP processor
- Only one external device—an SFP+. No other devices are required.
- Resource utilization

xc7k325t-2ffg900			
Slice Logic Utilization	**Used**	**Available**	**Utilization**
Number of Slice Registers	6,504	407,600	1%
Number of Slice LUTs	7,615	203,800	3%
Number of occupied Slices	3,056	50,950	5%
Number of RAMB36E1/FIFO36E1s	9	455	2%
Number of RAMB18E1/FIFO18E1s	0	890	0%
Number of BUFG/BUFGCTRLs	4	32	12%

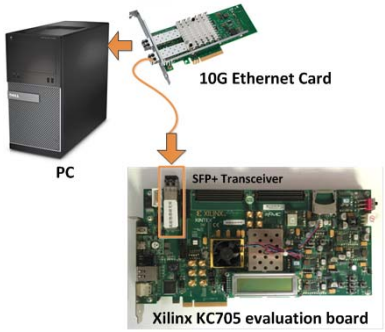

The preliminary test shows the throughput of FeTCP is about 400MBps, which has certain distance to the maximum possible TCP throughput.

TX Latency: 217.6ns

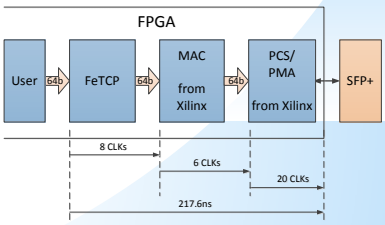

Running TCP at the maximum link speeds requires some software and hardware performance tuning. Optimization is in progress.