A low-latency, high-performance versatile SerDes Interface IP

Dr. Mondrian Nüssle IP-SOC2017, December 2017, Grenoble mondrian.nuessle@extoll.de



EXTOLL Background

- German high-performance computing hardware (networking) and IP design company
- Spin-off from U. of Heidelberg
- Based in Mannheim







EXTOLL started with HPC network designs
+ Added IP & Design Services





Low-latency interconnects

- Low-latency networking and interconnects important
 - Parallel computing (MPI,...)
 - Other applications as well!
- EXTOLL Tourmalet is the lowest-latency HPC interconnect (network) solution today
 - 0.810µs application-toapplication MPI latency
 - (Ethernet, typical: 5-10µs...)







EXTOLL ASIC – Tourmalet







- ASIC specifically designed for HPC
- x16 PCIe gen3 connectivity
- High sustained message rate (> 80M messages/s)
- Low latency message exchange (~0.8us)
- 7 Links, 12 lanes of up to 100Gb/s per direction and link
- 8.9GB/s MPI bandwidth
- 640 MHz Clock frequency
- 270M transistors
- ~60ns Hop latency
- >100 SerDes instances!



SerDes

- SerDes IP is one of the fundamental building blocks of todays high-speed networking
- EXTOLL found available IP offerings not always optimal:
 - High data rates often only on latest nodes
 - Not cost efficient!
 - Not optimized for low latency applications
 - ...
- Design of in-house SerDes technology

















Digital-centric design

- High speed SerDes architecture based mainly on digital logic
- Complemented by advanced verification and modelling methodology
- Minimum number of analogue and full custom components
- Keep complexity in synthesizable RTL code
- Enforce consistency between model and implementation
- Various digital control and tuning loops for robust performance
- Flexible, robust architecture; easy to migrate to other technologies; adaptation to customer applications







Important Features

- SerDes PHY is combined from transceiver lanes and clock generation block
- Flexible number of transceivers in one block
- Common high speed PLL for line rates @ 2.5 to 28 Gbps
- Programmable transmitter with equalizer (4 tap FIR)
- Programmable linear RX equalizer (CTLE)
- Programmable discrete RX equalizer (5 tap DFE)
- Digital Clock Data Recovery (CDR)
- Comprehensive suite of calibration circuits/loops
- Support for PCIe specific features (Far end RX detect, Electrical Idle,..)
- Diagnostic features:
 - Pattern generators
 - Concurrent Eye Monitor for equalization and channel analysis
 - Far End and Near Loopbacks
 - Analog testbus
- Available for 28nm/22nm processes!







Block Diagram





Low-Latency!

Examples of low-latency optimizations:

- TX Side
 - No sync FIFOS needed for multi-lane implementations
 - No additional stage for FIR equalizer
- RX Side
 - no time lost in 5-tap DFE due to quarter rate architecture
 - O-stage bitslip logic for word alignment instead of 1-2 stage barrel shifter logic





Design state & Silicon

- 8-lane test chip
- Tape-out in TSMC 28nm HPC+ in summer 2017
- Package and test board design

Done!











Summary

- Unique low-latency features
- Support for latest PCIe, Ethernet, et al. Rates (16Gbps, 25.x, 28.x,...)
- High data rates (more than 12.5Gbps!) available for planar 28/22nm processes (cost efficient nodes)
- Silicon proven (TSMC 28nm HPC+)
- Cost efficient licensing









Thank You!

Contact us: info@extoll.de

Come and see us at the exhibition!



