Maintaining Cache Coherency with AMD Opteron™ Processors using FPGA's

Parag Beeraka | February 11, 2009

Outline

- Introduction
- FPGA Internals
- Platforms
- Results
- Future Enhancements
- Conclusion

Introduction

Two ways to maintain Cache Coherency

- Memory is coherent with **AMD Opteron™**

- Ability to Cache **AMD Opteron™ 's** Memory

Our Approach

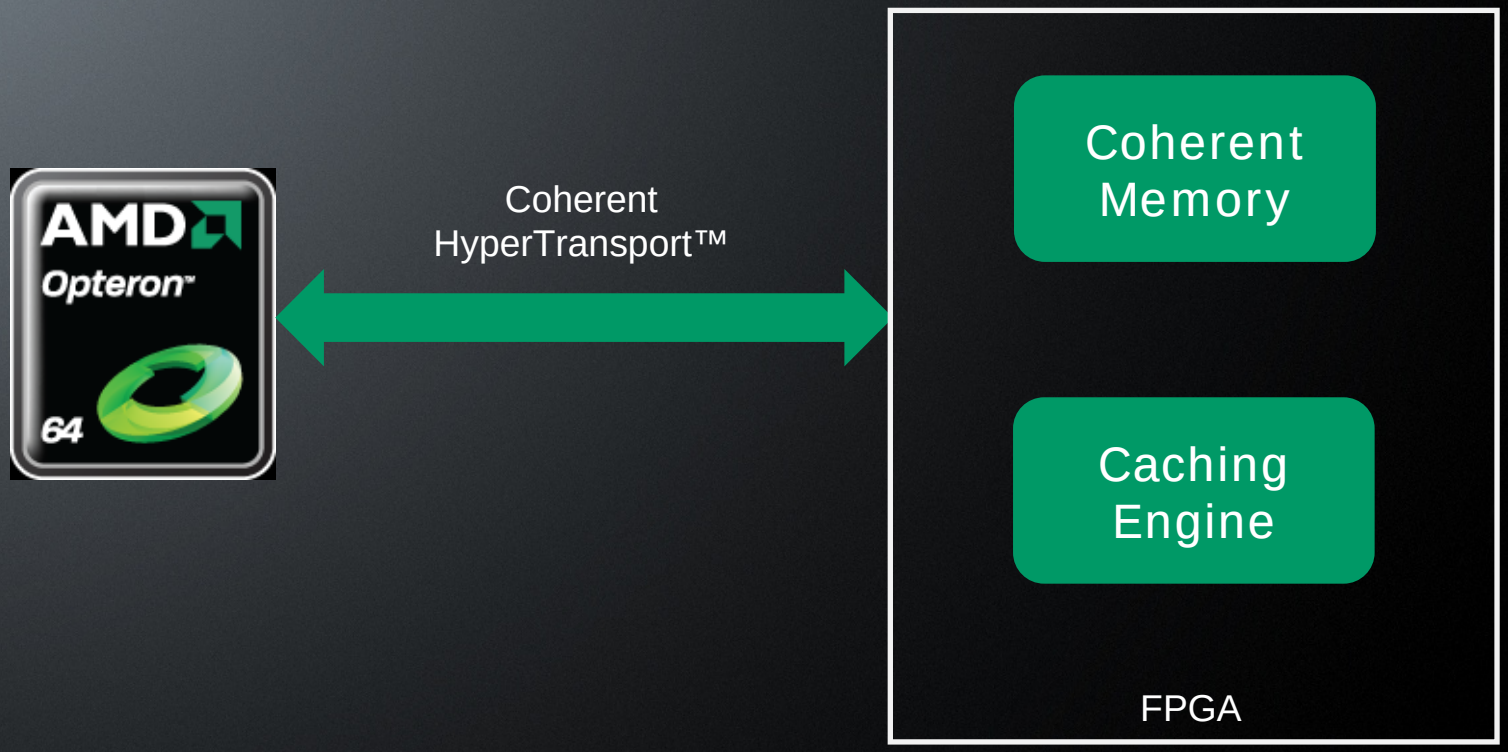

FPGA Internals

FPGA Internals

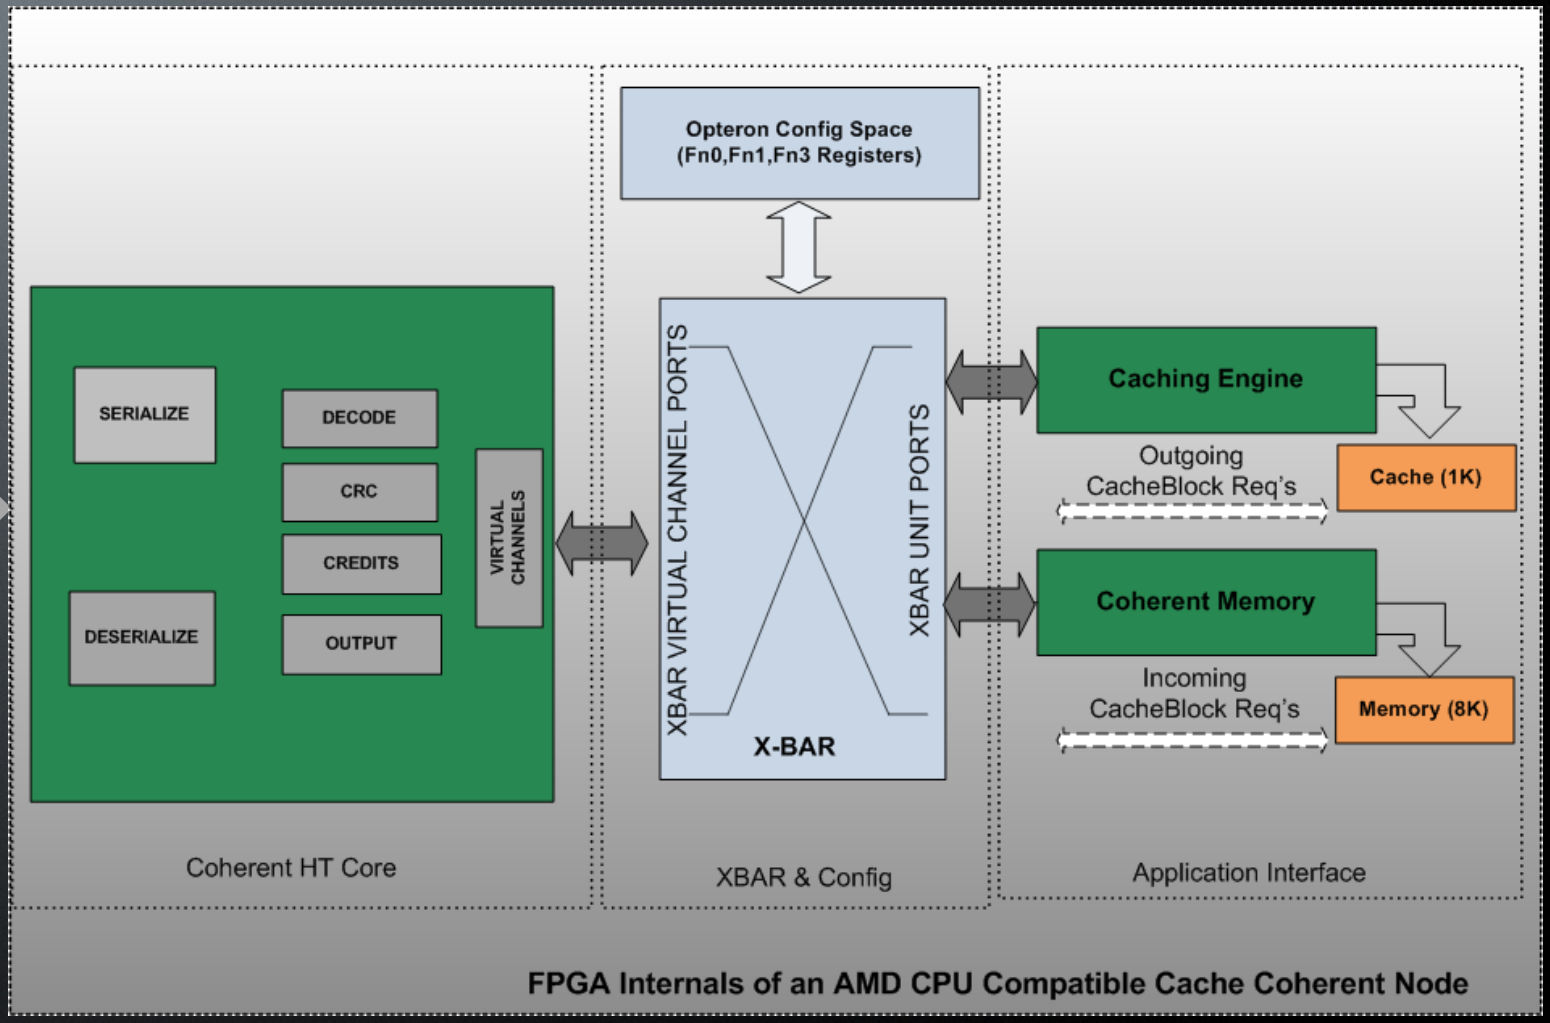

FPGA Internals

- Coherent HyperTransport™ core from University of Heidelberg
  - Follows *HyperTransport*™ Protocol
- X-Bar and Configuration Space
  - Configuration Space Register Strapping's
- FPGA Caching Engine (CPU) Core
  - Ability to request Cache Lines from **AMD Opteron™**
- FPGA Coherent Memory (MCT) Core
  - Memory is coherent with **AMD Opteron™** CPU's

Combination of all blocks → Cache Coherent Node

Coherent HyperTransport™ Core

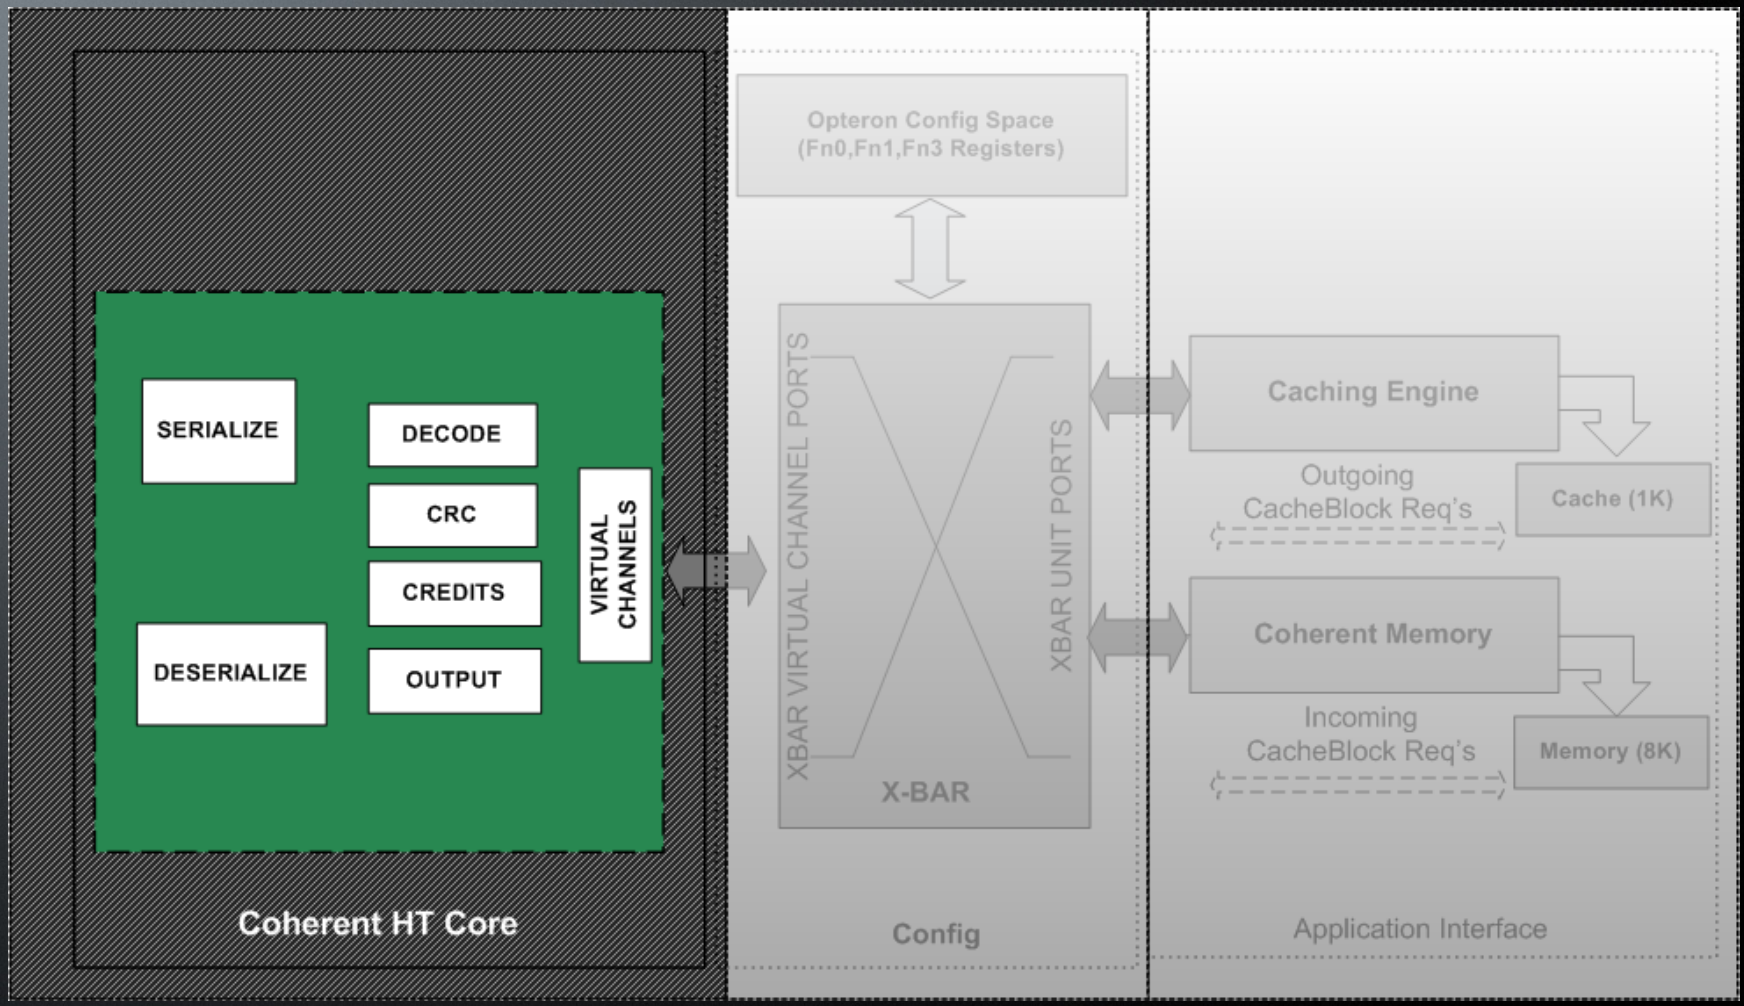

Coherent HyperTransport™ Core

- Follows full *HyperTransport*™ 1.0 Protocol
  - Takes care of HT Init, CRC's, NOP's, VC's etc
- Developed by University of Heidelberg
- Supports 400 / 800 Mbps Link Speed per CAD/ CTL Line
- Supports 8bit / 16bit Link Width
- Ability to decode and send all commands in coherent HyperTransport™ specification

Application Interface

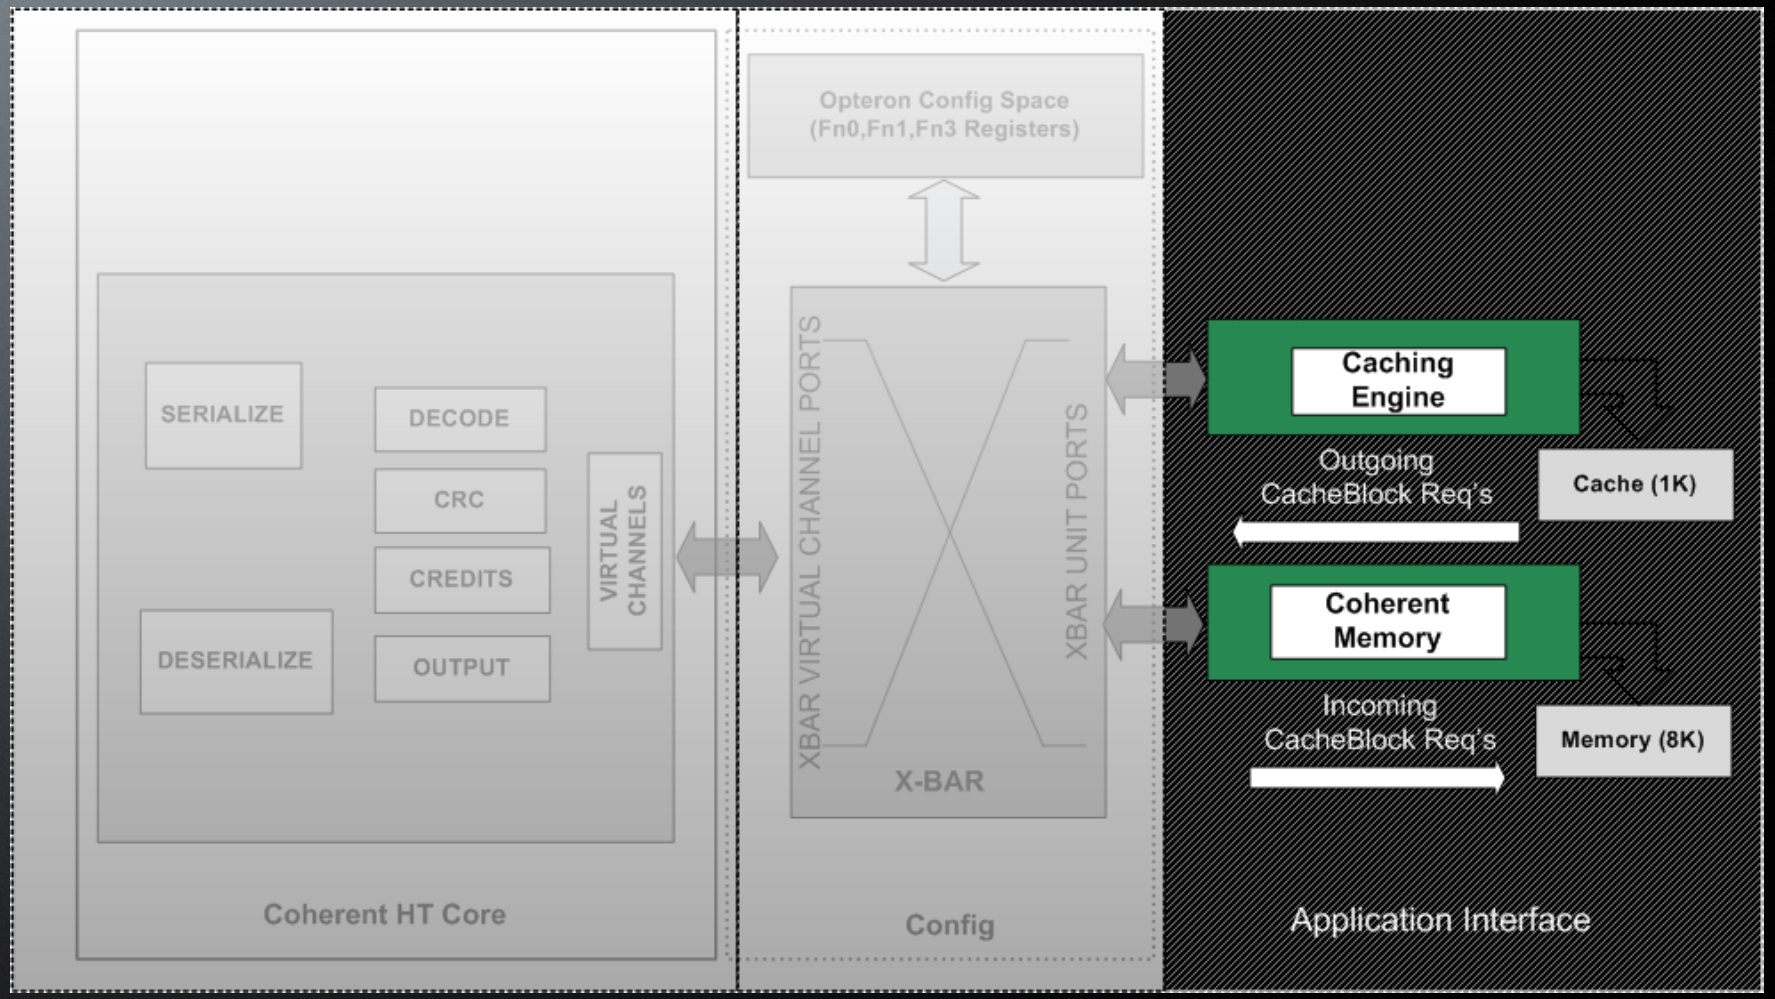

Caching Engine (CPU) Core

- Core capable of accessing Opteron™ Cache Lines maintaining Cache Coherency
- Direct Mapped Cache (1K)
  - 16 Cache Lines (Each Cache Line - 64 bytes)
- Implements “MOESI” State Machine for Cache Coherency
- Support Write Back, Write Thru and Non Cacheable Memory Type
- Handles incoming System Interrupts, Snoop Requests and Lock Transactions
- Handles Probe Filter related Requests and Responses (New AMD Opteron™ feature)

Caching Engine (CPU) Core

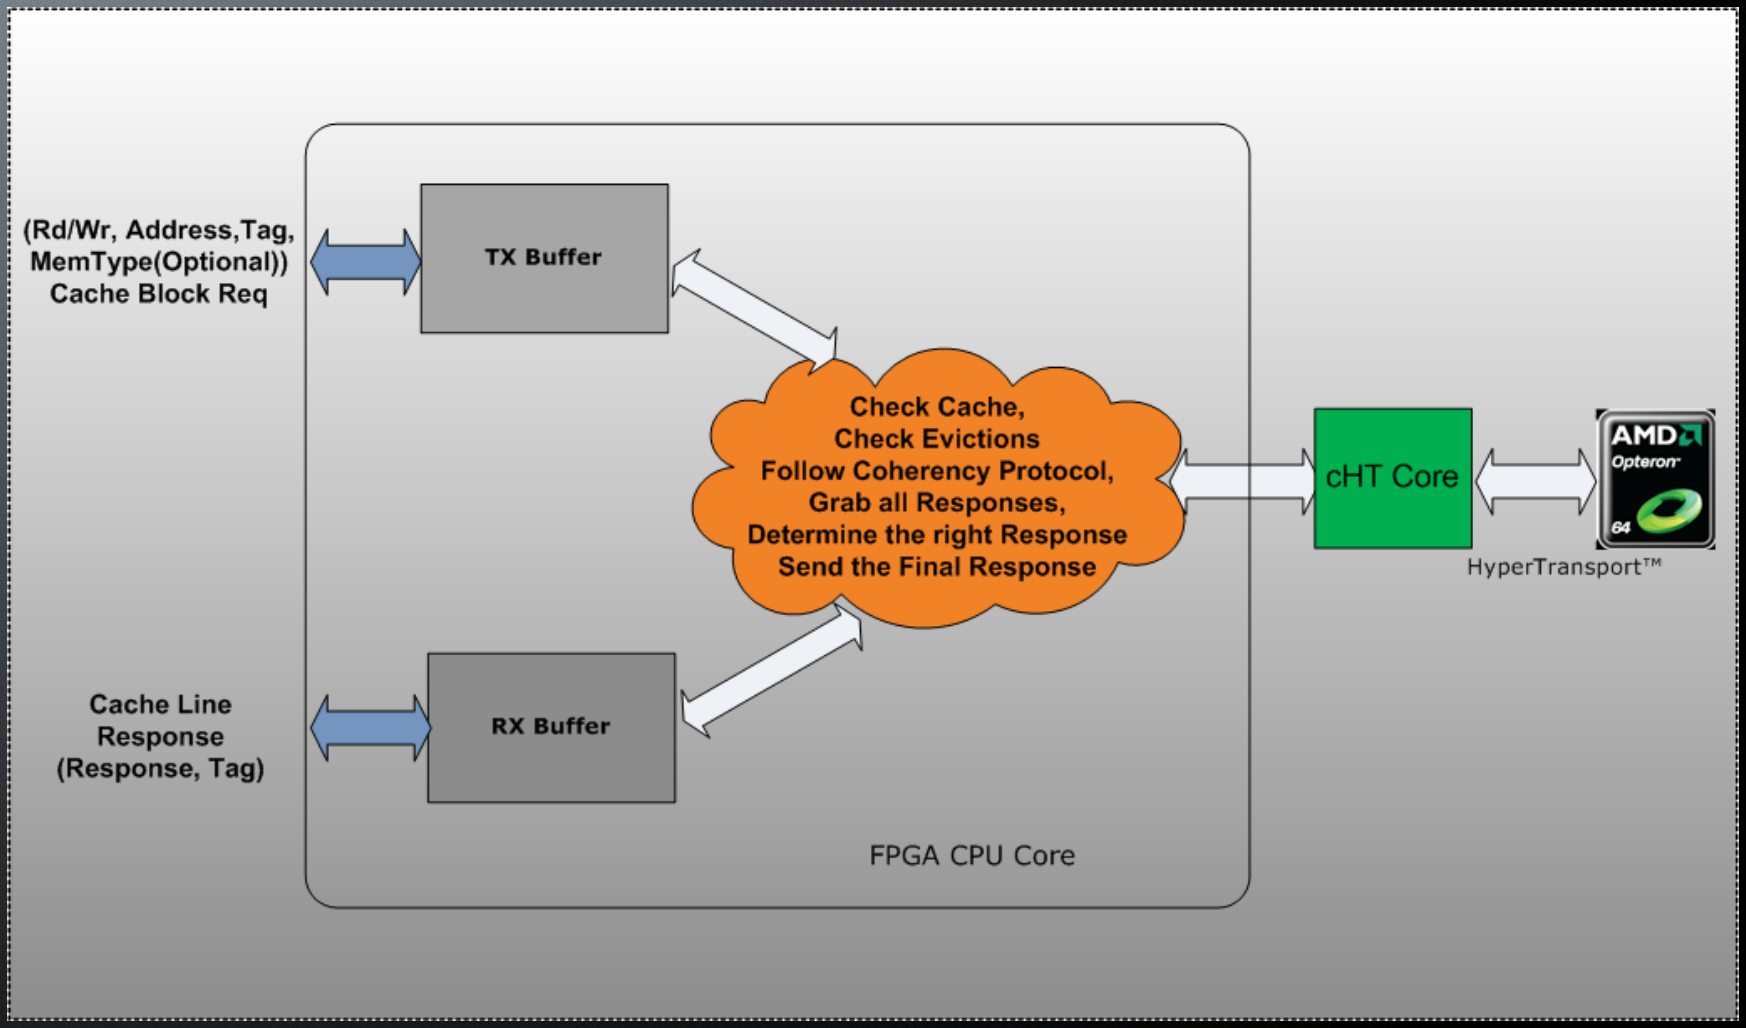

Caching Engine (CPU) Core

- Automatic SrcTag, Node Information Insertion into Requests
- Handles 32 Requests in-flight (Max for HyperTransport™)
- Handles Out – of – Order Responses for Requests
- Optional 1K Internal Buffers on Tx and Rx Interfaces

Coherent Memory (MCT) Core

- Core capable of having Cache Coherent Memory
- Developed by AMD Dresden Design Center and University of Heidelberg
- 8K Memory - Supports 128 Cache Lines
- All incoming Cache Block Requests are Routed to Memory Controller
- FPGA Coherent Memory (MCT) can also be cached by Caching Engine (CPU) Core

X-BAR and Configuration Space

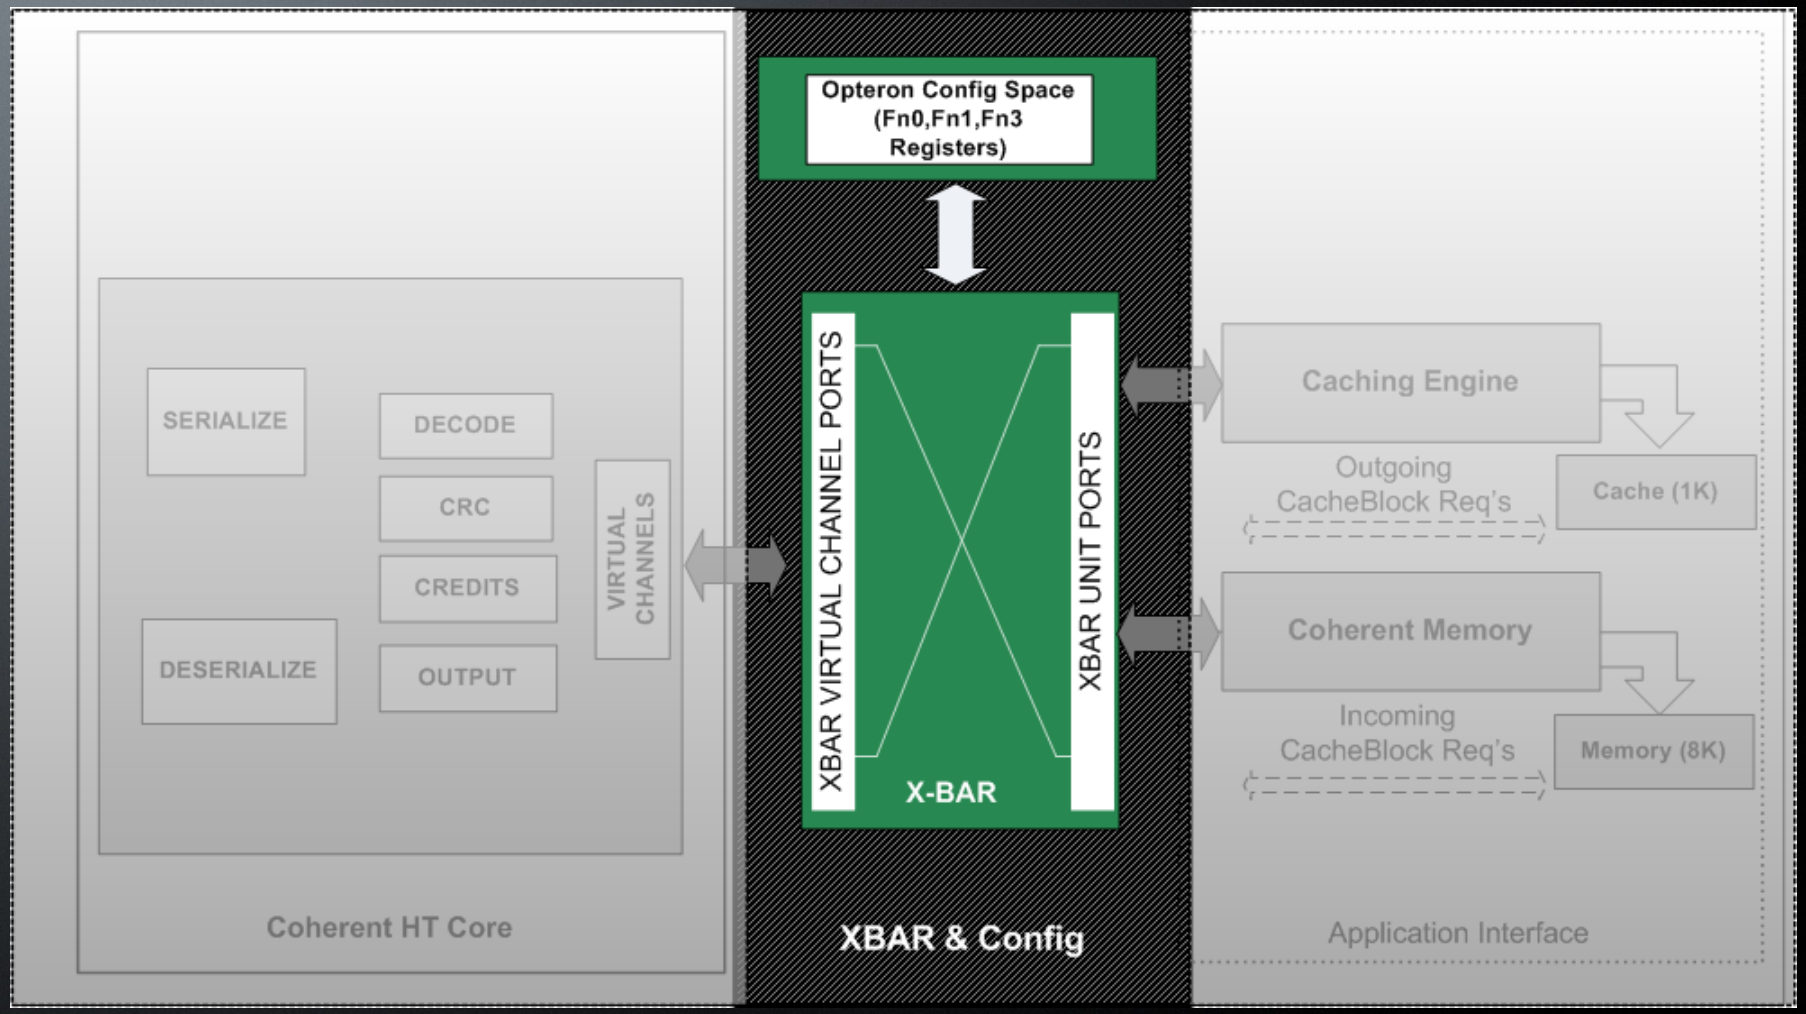

Configuration Space and X-BAR

Configuration Space

- Similar to AMD Family 10 Architecture Configuration Space
- Supports Fn0, Fn1, Fn3 and Fn4 Registers
- Added support for HyperTransport™ 3.0

X-BAR

- Developed by University Of Heidelberg
- Routes packets between Units (Caching Engine (CPU), Coherent Memory (MCT) Core) and Coherent HT Core
- Uses Routing Table Registers from Configuration to Route Packets to Units

Platforms

FPGA Cache Coherent Node Platform

- Two Xilinx Virtex 5 FPGA's
- Supports 4 16-bit HyperTransport™ Links with HMZD Connectors
- USB Interface for Back End Control

FPGA Cache Coherent Node Test Platform

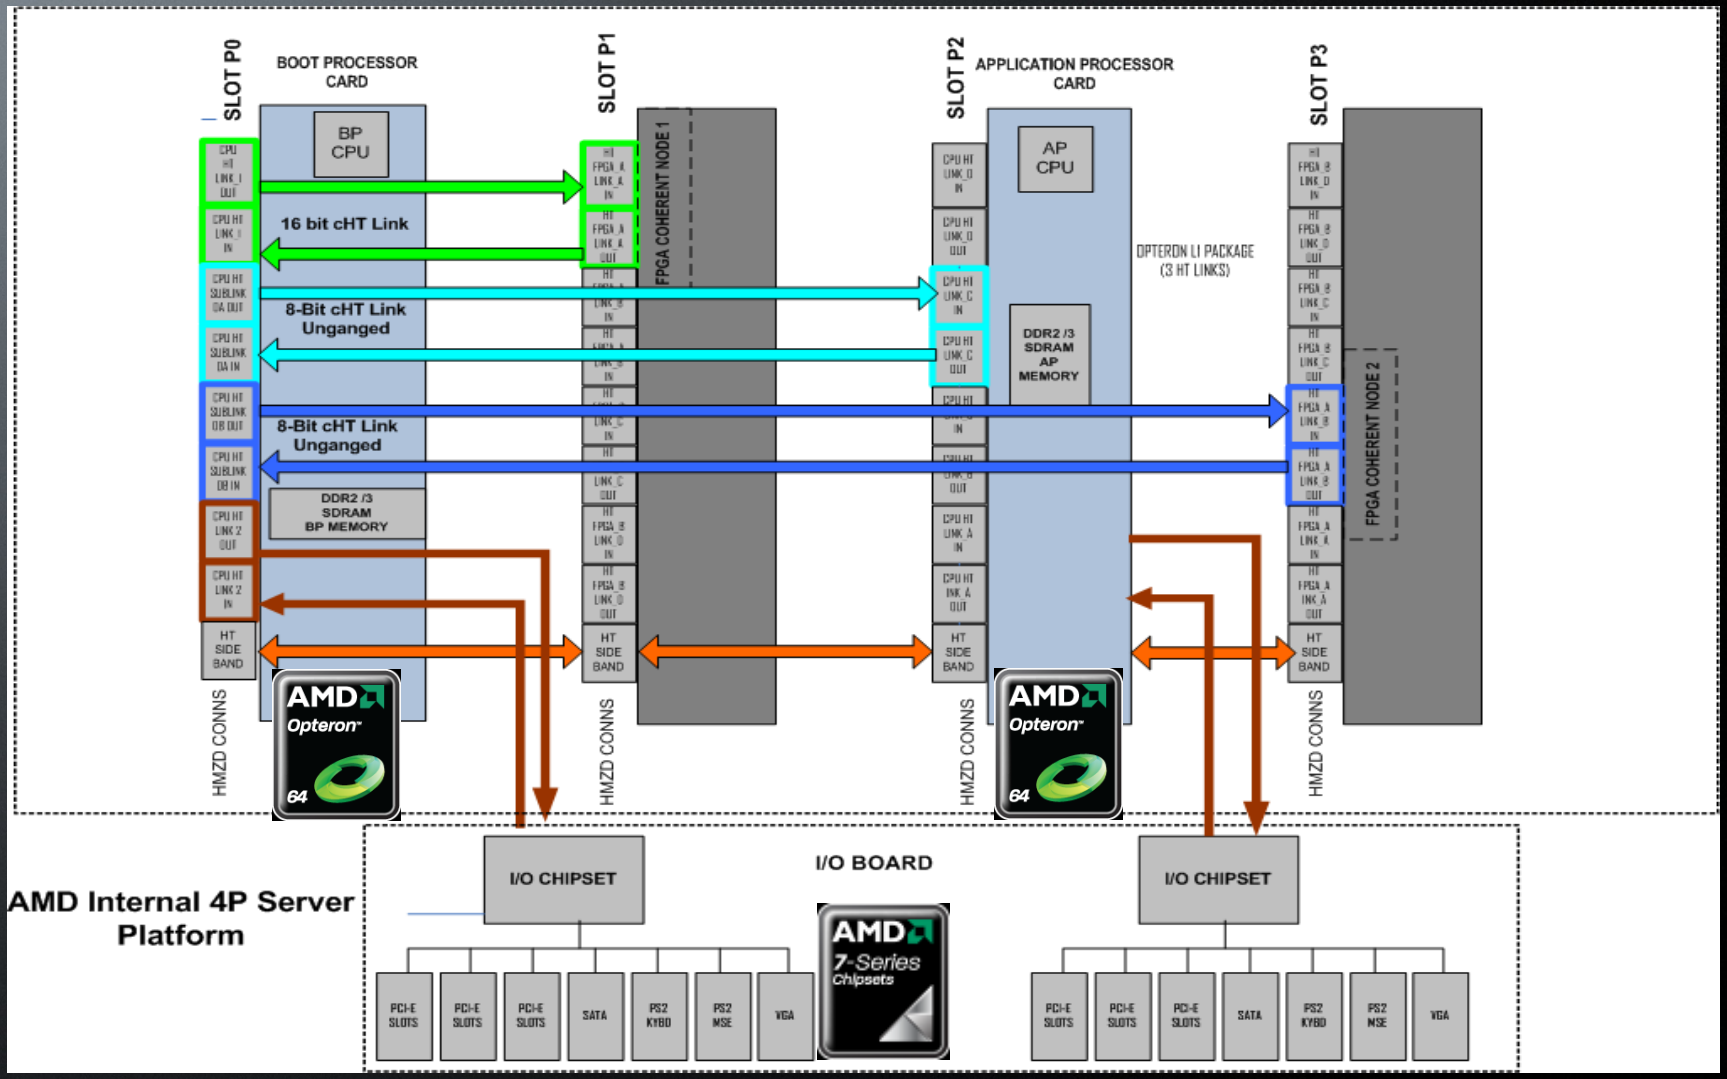

Platform BIOS

- Made changes to AGESA (AMD) and IBV (Phoenix, AMI) Bios Code for Coherent Node Configuration
  - Used DeviceID for configuring HT Fabric Topology
  - Bypass APIC Initialization
  - Bypass SMM Processing Routines
  - Bypass Power Management Routines
- Gained access to FPGA Coherent Memory Core (MCT) by defining a Memory Hole and changing DRAM Routing Tables

Platform Experiments

- Verified on Barcelona and Shanghai range of Processors on AMD Internal Server Platforms
- BIOS support for various 4P Configuration's
  - 1 AMD Opteron™, 1 FPGA Coherent Node
  - 2 AMD Opteron™ CPU's , 2 FPGA Coherent Nodes in a Star Topology
- Developed Linux Device Driver to read and write to FPGA's coherent memory (Verify FPGA Coherent Memory Core)
- Developed scripts to access Opteron™ Cache Lines using USB Back End Interface (Verify FPGA Caching Engine)

Results

Performance Metrics – Caching Engine (CPU)

Data Transfer Metrics for 8 bit HT Link with 400Mbps Line Rate

HT 8 bit 400Mbps
Max Bandwidth
- 800MB/s

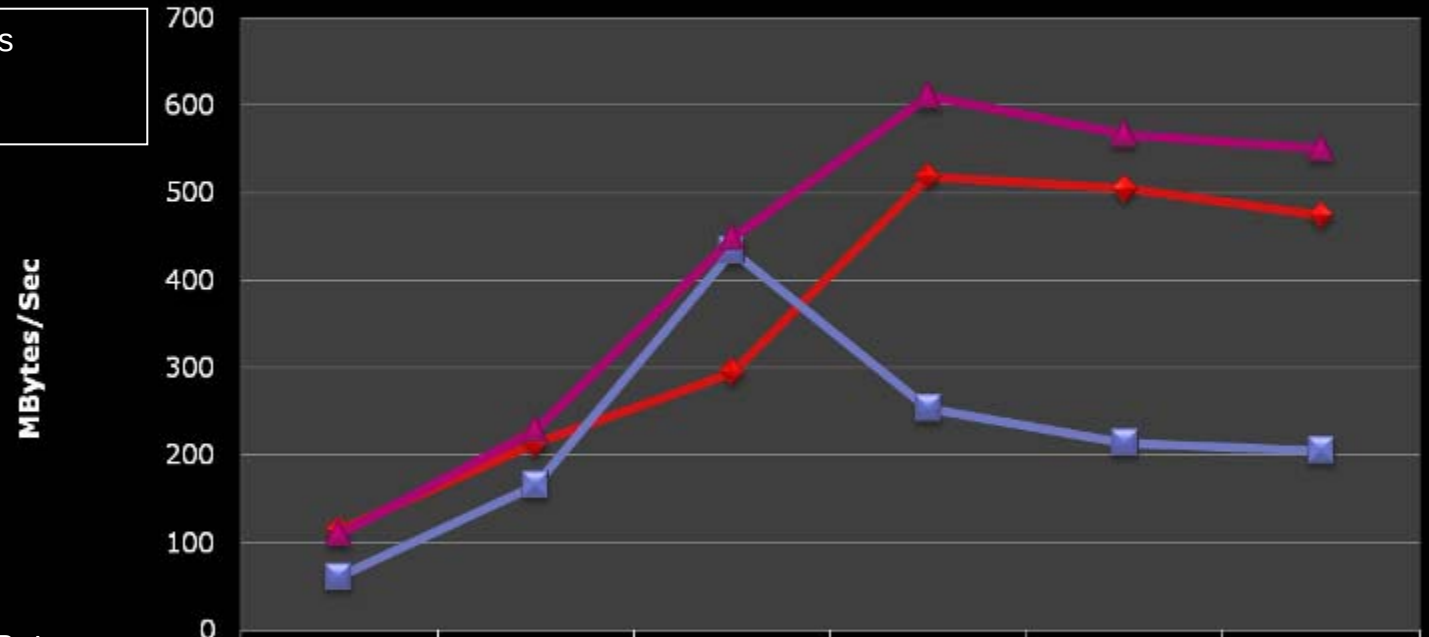

Bytes	64	256	1024	4096	16384	65536
Cache Read (Mbytes / sec)	114.25	213.25	295	518.375	504.5	474
Own Cache Block(Mbytes/sec)	61.5	166.125	433.875	253.875	214.375	205.625
Cache Read with Probe Filter (Mbytes/sec) (*Expected)	110.25	228.5	449	611.25	567.25	550.625

Future Enhancements & Conclusion

Future Enhancements

- Try Coherent HyperTransport™ Core with 16 bit Link Width and HT400 on same platform
- Perform Experiments in full 4P system with 2 FPGA Cache Coherent Nodes
- Verify on real system with next generation HT Core from University Of Heidelberg
- Plan to deliver whole package with next generation HT Core for Coherent Licensee's
- Port to other FPGA Vendor Devices
- Run Benchmarks with real applications

Conclusion

- Developed a Cache Coherent Node that can be used over Coherent HyperTransport™ with AMD Opteron™ Processors using FPGA's in various environments
  - Silicon Validation
  - Simulation BFM's Validation
  - Accelerated Computing (Co – Processing using FPGA's)

Danke!!

Questions ?

Trademark Attribution

AMD, the AMD Arrow logo and combinations thereof are trademarks of Advanced Micro Devices, Inc. in the United States and/or other jurisdictions. Other names used in this presentation are for identification purposes only and may be trademarks of their respective owners.

AMD
The future is fusion