



T4240 DMA/PCIE Performance Analysis Using Cadence Palladium


Wai Chee Wong
Sr.Member of Technical Staff
Freescale Semiconductor

Raghu Binnamangalam
Sr.Technical Marketing Engineer
Cadence Design Systems




Freescale, the Freescale logo, AlliVec, C-5, CodeTEST, CodeWarrior, ColdFire, C-Ware, the Energy Efficient Solutions logo, mobileGT, PowerQUICC, QorIQ, StarCore and Symphony are trademarks of Freescale Semiconductor, Inc., Reg. U.S. Pat. & Tm. Off. BeeKit, BeeStack, ColdFire+, CoreNet, Flexis, Kinetics, MXC, Platform in a Package, Processor Expert, QorIQ Converge, Qorivva, QUICC Engine, SMARTMOS, TurboLink, VortiQa and Xtrinsic are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2011 Freescale Semiconductor, Inc.

DAC 2013, Austin, TX




Agenda

- Company Overview
- Use of emulation at Freescale
- Palladium usage for emulation model under test
- Performance case studies
- Experiences using Palladium system
- Summary



Company Overview

- Global leader in embedded processing solutions
- Key player in automotive, consumer, industrial and networking markets
- Manufacturer of microcontroller, controllers, sensors
- Support multiple active projects at any time
 - Design centers from US, Canada, Israel, India and China
 - Applications includes networking, automotive, connectivity and industrial
 - Collaboration effort across multiple functional teams



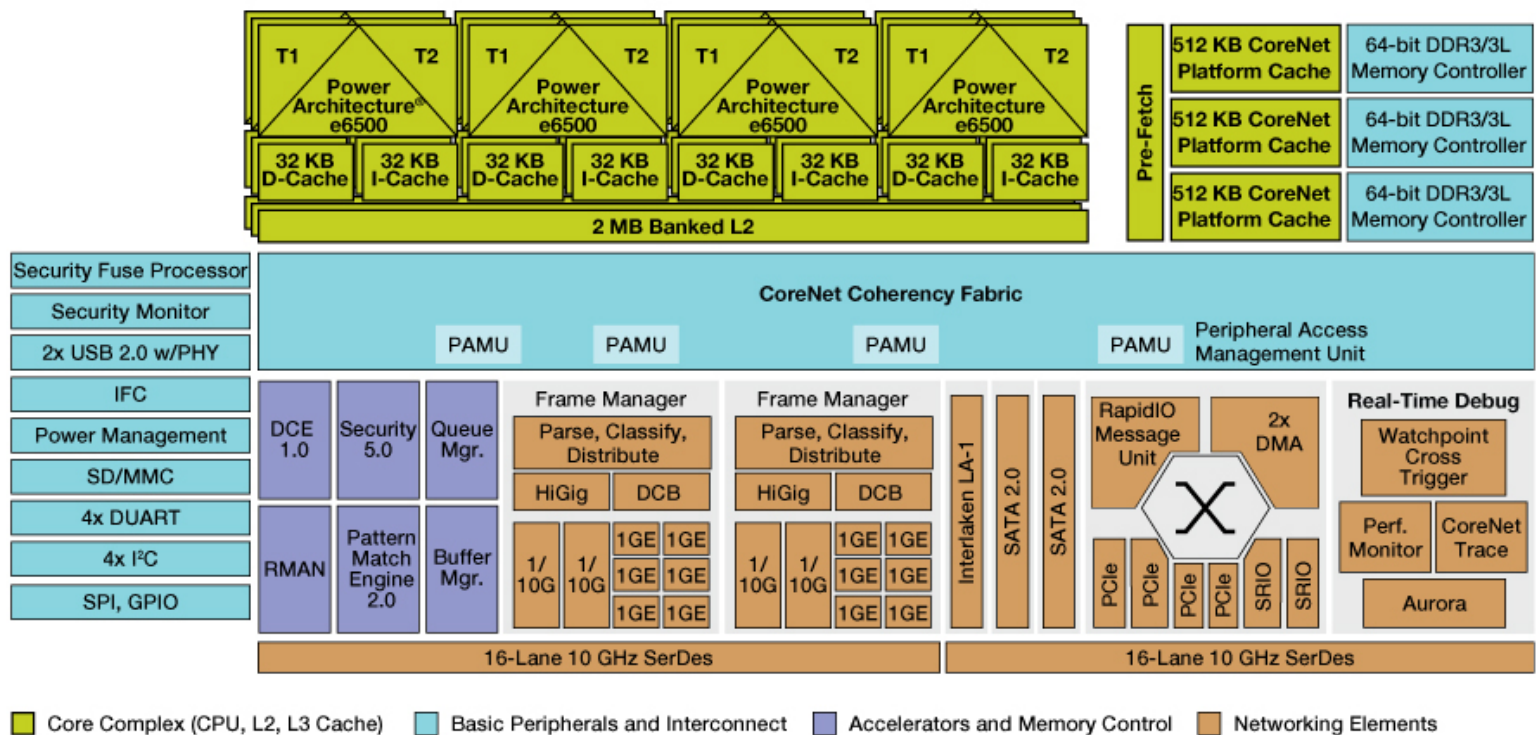
Use of Emulation Model at Freescale

- Functional verification
- Power analysis
- Failure reproduction and analysis
- Software validation (uboot, linux, dink, codewarrior firmware, etc..)
- System performance validation
 - Proof Point Application (ppa)
 - Defacto
 - § Life 4, coremark, NASA

T4240 Virtual Core Communications Processor

- Public Information

QorIQ T4240 Communications Processor




http://www.freescale.com/webapp/sps/site/prod_summary.jsp?code=T4240



Need for Emulation of T4240

- T4240 is an advanced communications processor benefitting a wide variety of applications due to its advanced processing, I/O integration and power management features
- New twelve, dual threaded e6500 cores, running at 1.8GHz
- Expanded platform cache and additional core capacity
- System is highly software centric, hence waiting for Si to do actual testing can be time detrimental
 - Emulation for early complete HW-SW system co-verification



T4240 Emulation

- T4240 design mapped to Cadence Palladium series
- Full Chip capacity: 140 million gates
- Palladium III operating range: ~500kHz
- Palladium XP operating range: 800KHz-1MHz
- Compile time: ~ 3-4 hours






Proof Point Application



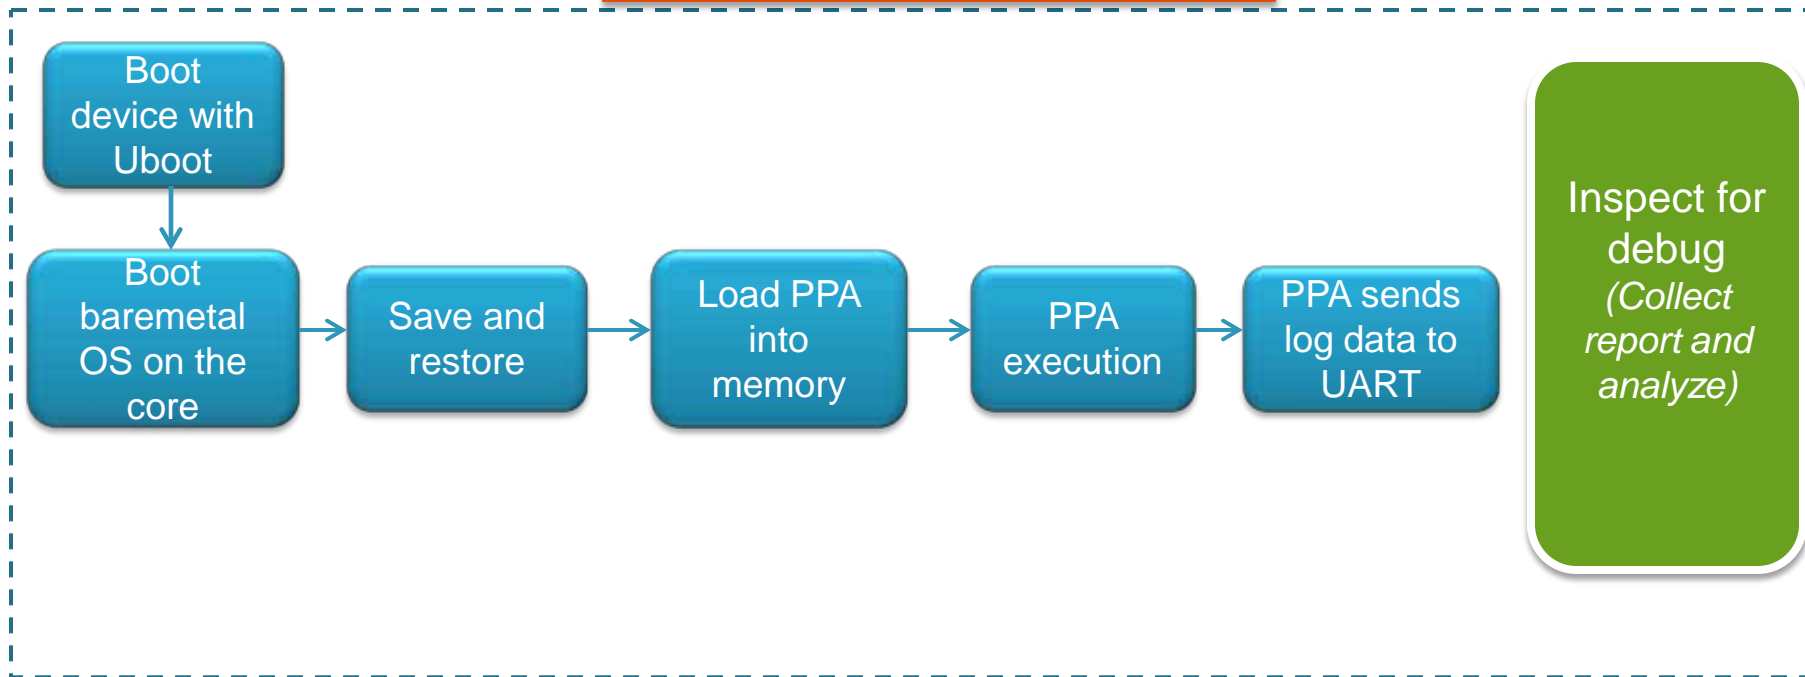
Freescale, the Freescale logo, AlliVec, C-5, CodeTEST, CodeWarrior, ColdFire, C-Ware, the Energy Efficient Solutions logo, mobileGT, PowerQUICC, QorIQ, StarCore and Symphony are trademarks of Freescale Semiconductor, Inc., Reg. U.S. Pat. & Tm. Off. BeeKit, BeeStack, ColdFire+, CoreNet, Flexis, Kinetics, MXC, Platform in a Package, Processor Expert, QorIQ Converge, Qorivva, QUICC Engine, SMARTMOS, TurboLink, VortiQa and Xtrinsic are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2011 Freescale Semiconductor, Inc.




Proof Point Application (PPA)

- Background
 - PPA is a software framework which is built on top of a baremetal OS
 - Provides the infrastructure for performance analysis


System flow during performance analysis






PPA performance analysis

- System performance depends on ability to balance both HW and SW features effectively
- Regular SW cannot run on emulator, hence PPA
 - PPA is an application requiring a light weight OS
- Early performance analysis using PPA in emulator is the gating factor to tape-out
 - Design should meet requirement spec



DMA Overview

- 3 Modes of operation
 - Basic direct
 - Basic chaining
 - Extended chaining
- “X” dedicated channel engine per DMA controller
- Basic direct mode
 - DMA transfer data based on register value
 - SW does all the initialization, and set the CS bit to start channel operation. HW set CB (Channel Busy) at the beginning of the transfer, clear it at the end of the transfer
 - One time data transfers



DMA operation


- Basic Chaining Mode:
 - software first builds a table of link descriptors in memory
 - link descriptor contains parameters such as source address, destination address and transfer size
 - multiple entries of link descriptors are allowed
- Software then initialize the Current Link Descriptor Address Register (CLNDAR)
 - DMA controller loads descriptors from memory, decodes, and then performs a transfer using the link descriptor information
 - If the current descriptor is the last one, the DMA controller stop after executing the current link descriptor. Otherwise the DMA controller reads the next link descriptor from memory and begins another DMA transfer



Performance Case Study 1

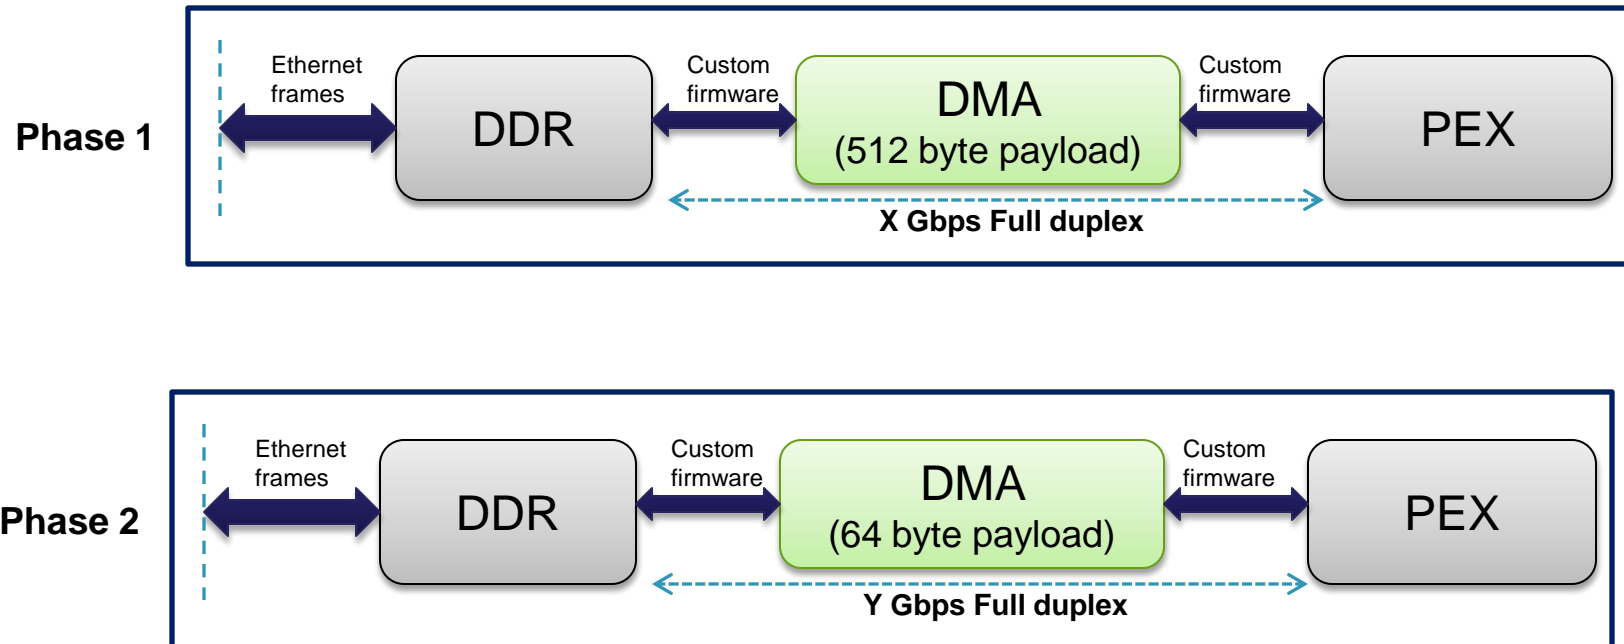



Freescale, the Freescale logo, AlllVec, C-5, CodeTEST, CodeWarrior, ColdFire, C-Ware, the Energy Efficient Solutions logo, mobileGT, PowerQUICC, QorIQ, StarCore and Symphony are trademarks of Freescale Semiconductor, Inc., Reg. U.S. Pat. & Tm. Off. BeeKit, BeeStack, ColdFire+, CoreNet, Flexis, Kinetics, MXC, Platform in a Package, Processor Expert, QorIQ Converge, Qorivva, QUICC Engine, SMARTMOS, TurboLink, VortiQa and Xttrinsic are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2011 Freescale Semiconductor, Inc.



Performance case study 1 - Use case


Network Interface Card (NIC)





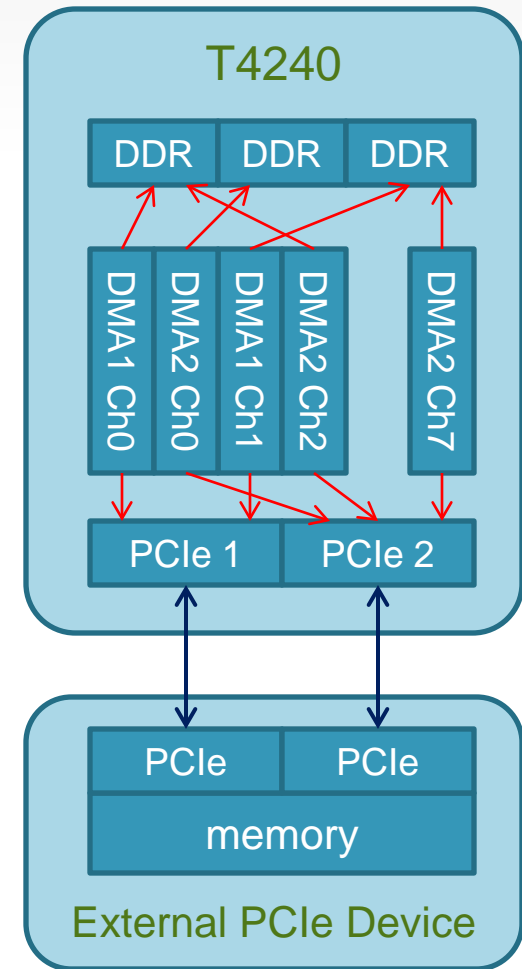
Setup for performance case study 1

- DMA transfers data in Chaining mode
 - DDR to DDR, or
 - DDR to PEX, or
 - PEX to DDR
- System Baseline
 - DMA descriptors locked to L3
 - Core/platform/DDR @ 1800/800/2133(MHz)
 - 3 way DDR interleaving with 4KB granule
- Payload size range of 64, 128, 256, 512 and 1024 byte
- Number of channel ranges from 1, 2 4, 8 or 16



Evaluation Phases

- Part I – DDR/DDR
 - Simple to setup
 - Sufficient to validate performance related to the known CHB issues
- Part II – DDR/PCle
 - Require an external PCle device to sink outbound PCle traffic






Part I : DDR/DDR



Freescale, the Freescale logo, AlliVec, C-5, CodeTEST, CodeWarrior, ColdFire, C-Ware, the Energy Efficient Solutions logo, mobileGT, PowerQUICC, QorIQ, StarCore and Symphony are trademarks of Freescale Semiconductor, Inc., Reg. U.S. Pat. & Tm. Off. BeeKit, BeeStack, ColdFire+, CoreNet, Flexis, Kinetics, MXC, Platform in a Package, Processor Expert, QorIQ Converge, Qorivva, QUICC Engine, SMARTMOS, TurboLink, VortiQa and Xtrinsic are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2011 Freescale Semiconductor, Inc.



DMA Throughput (DDR-to-DDR)

8 outstanding vs. 16 outstanding

Performance Improvement from 8 to 16 outstanding read requests


Throughput (Gbps)	Number of DMA Channels				
Payload Size (Byte)	1	2	4	8	16
64	99.00%	100.77%	98.30%	122.57%	180.65%
128	99.72%	102.98%	93.67%	124.75%	157.62%
256	99.43%	98.86%	96.70%	118.30%	149.14%
512	100.23%	101.20%	105.90%	138.53%	151.95%
1024	99.26%	101.16%	123.13%	140.16%	145.93%



Part II: DDR/PCle




Freescale, the Freescale logo, AlliVec, C-5, CodeTEST, CodeWarrior, ColdFire, C-Ware, the Energy Efficient Solutions logo, mobileGT, PowerQUICC, QorIQ, StarCore and Symphony are trademarks of Freescale Semiconductor, Inc., Reg. U.S. Pat. & Tm. Off. BeeKit, BeeStack, ColdFire+, CoreNet, Flexis, Kinetics, MXC, Platform in a Package, Processor Expert, QorIQ Converge, Qorivva, QUICC Engine, SMARTMOS, TurboLink, VortiQa and Xtrinsic are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2011 Freescale Semiconductor, Inc.



DMA Throughput (DDR/PCIe) - Setup

- | Model Used and Experiment Setup
 - T4 rev 1.0 design by large
 - Optimistic (and unrealistic) PCIe read latency

- | DMA Configuration
 - Dedicated mode – each DMA controller is dedicated to transfer data in a single direction
 - DMA1 – DDR2PEX
 - DMA2 – PEX2DDR
 - Mixed mode – each DMA channel is dedicated to transfer data in a single direction; each DMA controller may transfer data in opposite direction
 - Dedicated mode vs. Mixed mode
 - | Dedicated mode has better throughput
 - | Mixed mode has balanced transfer rate for traffic of opposite direction



Observation

- | Throughput of DDR-to-PEX is much better than PEX-to-DDR
 - Each DMA can only support “X” outstanding reads
 - Overall throughput may increase if there are more read buffers in DMA


- | Implication
 - Just adding additional DMA channels won't help
 - Having a DMA descriptor cache should help
 - | Less read buffers are used for descriptor fetching, leaving more read buffers for payload
 - | Improvement is expected to be small because DMA is mainly waiting for data from PEX but not for descriptors



Performance Case Study 2

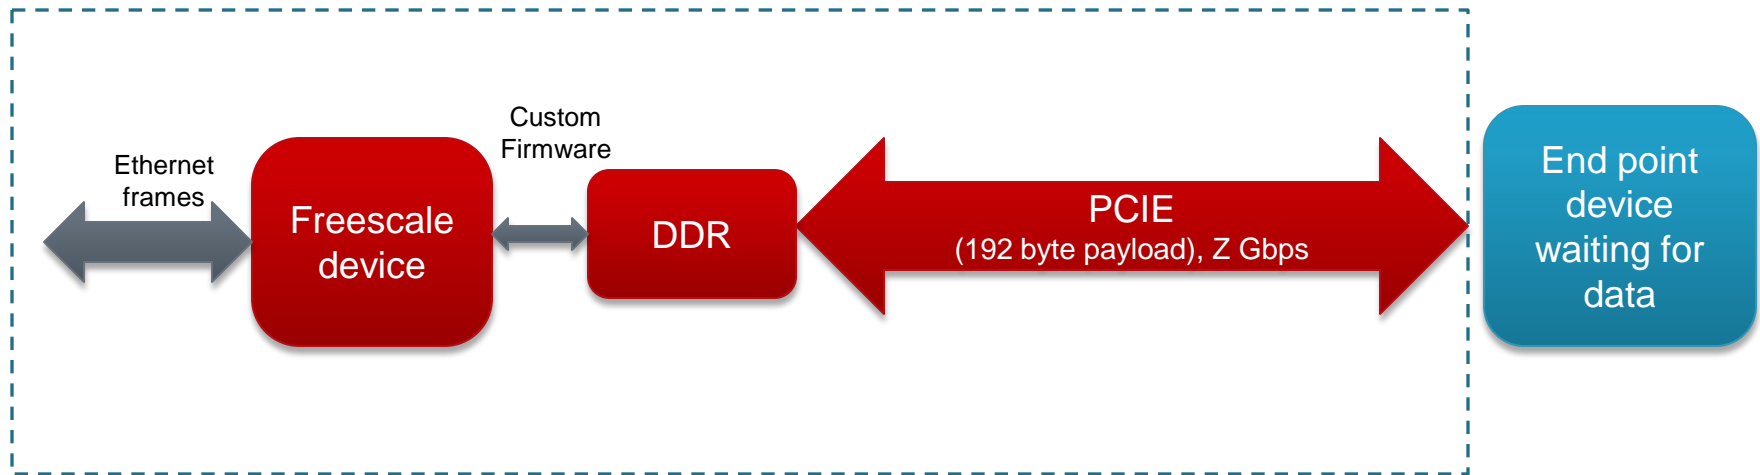


Freescale, the Freescale logo, AlllVec, C-5, CodeTEST, CodeWarrior, ColdFire, C-Ware, the Energy Efficient Solutions logo, mobileGT, PowerQUICC, QorIQ, StarCore and Symphony are trademarks of Freescale Semiconductor, Inc., Reg. U.S. Pat. & Tm. Off. BeeKit, BeeStack, ColdFire+, CoreNet, Flexis, Kinetics, MXC, Platform in a Package, Processor Expert, QorIQ Converge, Qorivva, QUICC Engine, SMARTMOS, TurboLink, VortiQa and Xttrinsic are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2011 Freescale Semiconductor, Inc.



Performance case study 2 - Use case

Video streaming application






Experiences




Freescale, the Freescale logo, AlllVec, C-5, CodeTEST, CodeWarrior, ColdFire, C-Ware, the Energy Efficient Solutions logo, mobileGT, PowerQUICC, QorIQ, StarCore and Symphony are trademarks of Freescale Semiconductor, Inc., Reg. U.S. Pat. & Tm. Off. BeeKit, BeeStack, ColdFire+, CoreNet, Flexis, Kinetics, MXC, Platform in a Package, Processor Expert, QorIQ Converge, Qorivva, QUICC Engine, SMARTMOS, TurboLink, VortiQa and Xtrinsic are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2011 Freescale Semiconductor, Inc.




What worked well?

- Sub-Block level analysis along with complete system level verification using emulation
- Accumulated emulation data from previous designs helped greatly in designing and identifying bugs
- Waveform analysis using emulation data helped save weeks worth of efforts in analyzing performance bottlenecks
 - Not possible to test everything
 - Know where to measure and how to measure!



Experiences and Suggested improvements

- Explore hybrid verification with virtual platforms for faster emulation runs of large systems
- Understanding traffic patterns very important for max. performance
- Software co-verification integral and has to be completed early



Summary

- Performance validation and analysis is becoming more important because of the advancements in SOC and reuse methodology
- Emulation is a good platform because of fast execution time as compared to simulation, full visibility of internal signals, and the fast turnaround time of model build
- Mapping system level performance goals to SoC level is achievable. Once the mapping is done, the value of validating performance goal before tape out is extremely valuable.

