FPGA-Acceleration on COTS x86 Platforms

University of Mannheim,
16 Feb 2007

Today's Agenda

- XtremeData Corporate & Team background
- Why FPGAs in COTS x86?
  - Issues and XDI Solution
- FPGA acceleration markets
  - FPGAs in HPC
- Summary

XtremeData: Corporate History…

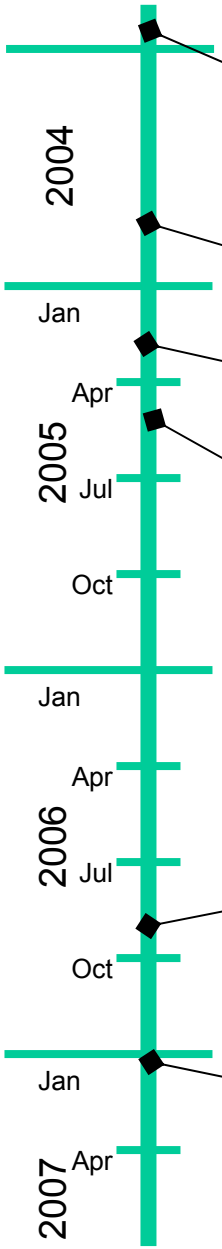

- Incorporated 2003, Seed funds raised.
- Market research & POC completed: target markets identified.
- SeriesA raised and development started with two teams: hardware in Chicago and software in Bangalore, India.
- System architecture defined: commodity hardware platform, accelerator and database engine.
- FPGA Module offered as a stand-alone product, press releases; strategic partnerships made, shipments started…
- SeriesB fund raise closing Jan 2007 for Go-To-Market financing

Team Background

Ravi Chandran, CEO

- BE Electronics, India, MS EE, University of Texas, Arlington, MBA, Kellogg School, Northwestern University, IL
- President, Binary Machines, Inc., Schaumburg, IL
- COO, VP of Engineering., Bio-Imaging Research, Inc., Lincolnshire, IL (www.bio-imaging.com)

  20+ years of product development & design services in medical & industrial (NDT) imaging markets. 20+ years experience with Toshiba Medical Systems – 20% of worldwide CT scanner installed base.

1 in 5 CT scanners worldwide have an imaging system designed by our team

Images courtesy of: Toshiba Medical Systems, Philips Medical Systems & BIR Inc.

Vision

XtremeData's vision is to build "Accelerated Computing Appliances"

"Appliance" → **implies**

- Easy installation – "plug and use"
- No disruption to existing process

"Accelerated Computing" → **implies** → x86 CPU ("PC") + FPGA

Strategy

Our strategy is to enable “Accelerated Computing Appliances” by:

1. **coupling off-the-shelf x86 hardware (“PC’s”)**

2. **with FPGA accelerators**

3. **via a software middleware layer that enables ease-of-use.**

› **We believe that the combination of these 3 key concepts gives us a compelling and sustainable price/performance advantage over the long term.**

FPGAs in Computing: Market Environment & Challenges

Market Environment....~2002

High-Performance Embedded Systems

- Specialized CPUs & DSPs + FPGA
- Specialized interconnect (Myrinet, Race++, RapidIO..)
- Custom boards, backplanes

Low Volume / High Cost

High Performance

Commodity "PC" systems

- x86 CPUs
- Standard interconnect (PCI-X, GigE)

Low Performance

High Volume / Low Cost

Market Environment….today

High-Performance Embedded Systems

Outperformed by ~3Ghz x86 CPUs

Outperformed by FPGAs at high-end

Outperformed by IB

- Specialized CPUs & DSPs + FPGA
- Specialized interconnect (Myrinet, Race++, RapidIO..)
- Custom boards, backplanes

Low Volume / High Cost

High Performance

Commodity "PC" systems

- x86 CPUs – multi-core, 3Ghz
- Standard interconnect (PCIe, IB, 10GigE)

Best choice: x86 CPU+FPGA.

How to do this?

High Performance

High Volume / Low Cost

Take x86 CPU back to embedded world

Bring FPGA forward to x86-COTS world

FPGA in COTS x86: Challenges

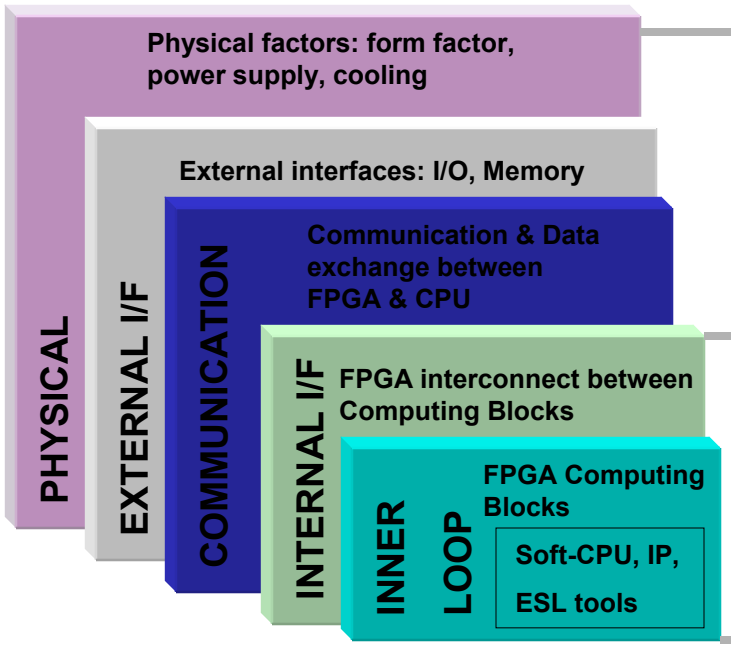

System Architect issues

- ***How to integrate FPGA into x86-COTS?***

Design engineer (s/w & h/w) issues

- ***How to transition from DSP world to FPGAs?***

Our Solution

Dual-socket AMD Opteron Motherboard:

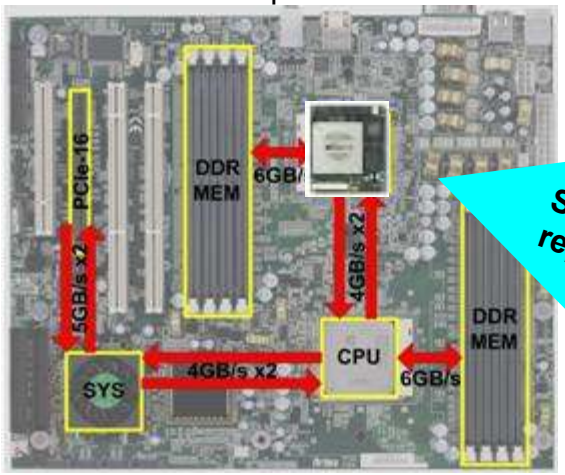

Idea: build a simple, minimalist board with interfaces to HyperTransport and memory: (Patent Pending)

drop-in replacement for an AMD Opteron with *no* changes to motherboard!

Simply remove Opteron & replace with FPGA Module !

- **FPGA uses all motherboard resources meant for CPU:**
  - **HyperTransport Links, Memory interface, power supply, heat-sink**
- **Usable with any AMD Opteron (or future Intel CPU) server**
- **Mix & match FPGAs, CPUs on quad-socket systems**

› Usable in rack-mount or high-density "blade" server systems (including ATCA), where plug-in boards are not feasible

Today's 940-pin solution…

Mechanical
- Plugs directly into socket-940
- Fits within AMD-specified retention frame
- 68 x 60 mm form factor
- Can use off-the-shelf Opteron™ heat sink

HyperTransport Interfaces (HT)
- Multiple HT interfaces
- 16 bits wide @ 800 M Transfers/s
- Bridging to additional XD1000™modules

Memory Interface
- 128 bits wide DDR-333 memory
- 5.4 GBytes/s bandwidth
- Up to four 4GB DIMMs of ECC memory

SRAM
- 8 MB of Zero Bus Turn-around (ZBT) SRAM
- 800 Mbytes/s bandwidth
- 32 bits wide with parity
- 5 clock cycle latency for reads @ 200MHz

FPGA Configuration
- Auto FPGA configuration on power-up
- Host triggered FPGA reconfiguration

Monitoring
- FPGA mastered I2C bus
- Voltage monitoring
- Temperature monitoring

Test Support
- JTAG test port
- 4 programmable LEDs
- 8 programmable test pads

Flash ROM
- 32 MB of CFI FLASH
- Use for FPGA configuration files, or application data

Development Package
- HyperTransport core
- Memory controller core
- Linux device driver
- FPGA messaging infrastructure

› **Newer Opteron socket solutions on the roadmap…**

XtremeData, Inc.
COMPUTING REDEFINED

FPGA-Acceleration Markets

FPGA acceleration markets

"Embedded" Computing

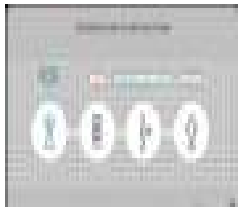

Medical Imaging

- Toshiba
- GE Medical
- Siemens
- Phillips

Telecom

- Motorola
- Nokia
- Ericcson

High Performance Computing

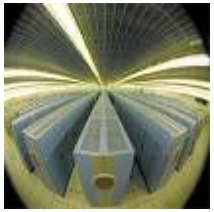

Financial, Database, Geoscience

- Cray
- SGI
- Mercury
- Linux NetworX

Scientific BioInformatics

- TimeLogic
- Tarari
- Progeniq

Emerging market: Video

Broadcast

- Set Top Box
- Video on Demand
- IPTV

Consumer

- DVD creation
- HD Video

› **Some examples of companies / applications that are using FPGA acceleration today**

FPGAs in High-Performance Computing (HPC)

HPC:
"Burning" Issues

" … more than 80% of data centers are already constrained by electrical power, physical space, or cooling capacity. Simply adding **more of the same kinds of systems** is clearly no solution, …"

[Sun Whitepaper on Throughput Computing, Nov 2005]

" … IBM Fellow Bernard Meyerson told the crowd at the Hot Chips conference yesterday that he expects a power crisis of sorts to occur in the server market come 2007. That's when the overall cost of powering and cooling all the servers in the US will outpace the amount of money spent on new server sales.…"

[The Register, 23 Aug 2006]

" … Google is rumored to have a million servers around the world and, according to a knowledgeable source, is already the top electricity user in at least one large U.S. state. …."

[Fortune, 1 May 2006]

"Just think about where there are windmills, dams, and other natural power sources around the world, and that's where you're going to see server farms," Ray Ozzie, Chief Software Architect, Microsoft

[Fortune, 1 May 2006]

› **Performance/Watt is key…..FPGAs are a viable alternative.**

FPGAs in HPC: Challenges

- HDL-based design flow NOT acceptable!
- A software-oriented programming model is NECESSARY

- High-level FPGA design tool (ESL) is a must: C-based, MatLab, etc.,?
- We have ongoing efforts to enable high-level design flow…

ESL: Our first step: PSP for ImpulseC

C Source Code

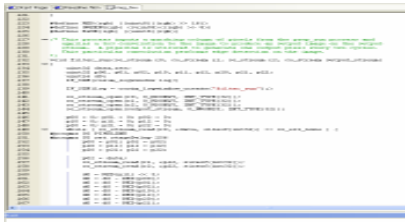

Single and double-precision IEEE 754 floating-point arithmetic supported in FPGA and inferred from code

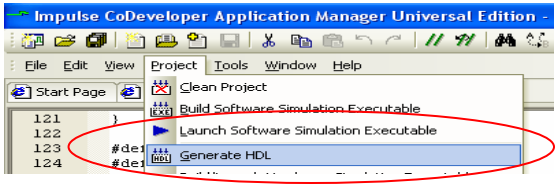

Compile C to HDL

Impulse C Module

XDI Memory Controllers

XDI Hypertransport Interface

Calls SOPC Builder drives it with scripts

SOPC Builder

All Components Connected via Avalon Fabric

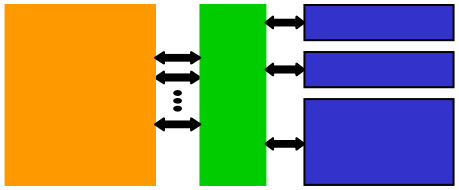

1) All RTL for FPGA on XD1000
2) Complete Quartus II Project
3) Complete Quartus II Compile Script

› **No user-supplied HDLsource code required**

XtremeData, Inc.
COMPUTING REDEFINED

Some FPGA Demo applications…

HPC: Financial Analytics

- Key Applications
  - Derivatives Trading
  - Black Scholes Model
  - BGM/LIBOR Market Model
  - Monte Carlo Simulations
- System Requirements
  - Reduce total power consumption
    - Cooling capacity is a major concern!
  - Run in Tier 1 blade/chassis servers
  - Accelerate cycle for trading & risk management decisions

Black-Scholes models the behavior of the price of an asset (share price of a traded stock, for instance) as a stochastic process which is described linear parabolic partial differential equation:

$$\frac{\partial V}{\partial t} + \frac{1}{2}\sigma^2 S^2 \frac{\partial^2 V}{\partial S^2} + rS\frac{\partial V}{\partial S} - rV = 0$$

› Accelerate mathematical modeling starting from HLL source code

HPC: Financial Analytics…

Monte Carlo Black-Scholes Simulation

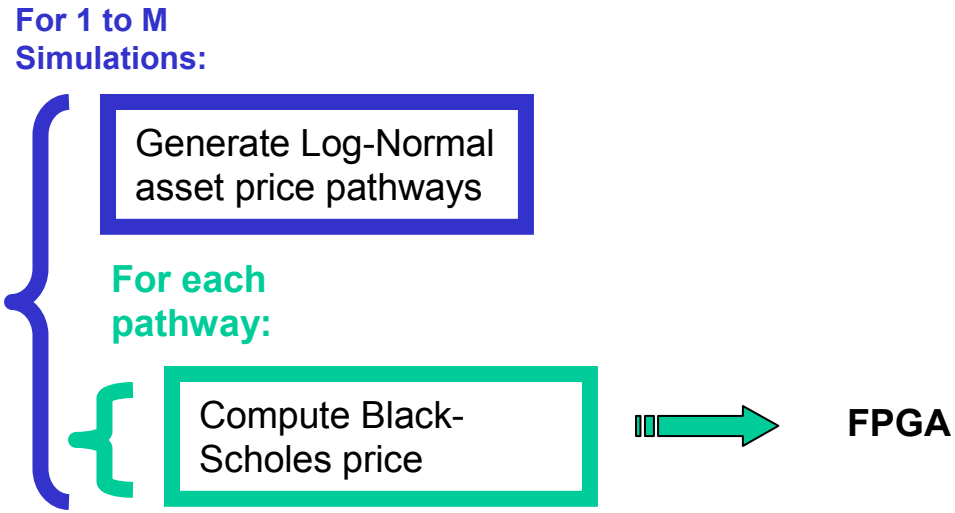

Whitepaper on this case study available on XtremeData, ImpulseC and HT Consortium websites.

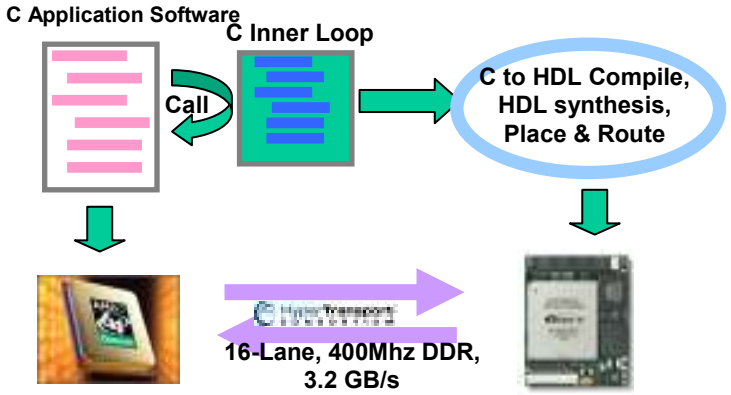

› **Demonstrated 20x speedup today, 50x possible very soon**

Video: H.264 Encoding

- Hi-Def Video Encoding/Decoding
  - Broadcast
  - Content Creation
  - DVD Authoring
- SD video: (PAL) 704 x 576 at 25 fps
- CPUs
  - CPU (~2 GHz) maxed out at 10 fps
  - Will not linearly scale up for HD
- FPGAs
  - 32 fps
  - 14K LEs out of 140K = ~10% utilization of biggest 2006 FPGA.
  - Can scale up to 9x with 6 cores: ~300 fps for SD.
  - **CPU @ 10 fps will encode 2-hour movie in 5 hours**
  - **FPGA @ 300 fps will encode in 10 mins !**
- Can easily handle HD video: 1920 x1080 at 30 fps , with < 50% utilized
- 2007 FPGA 2x larger 2006

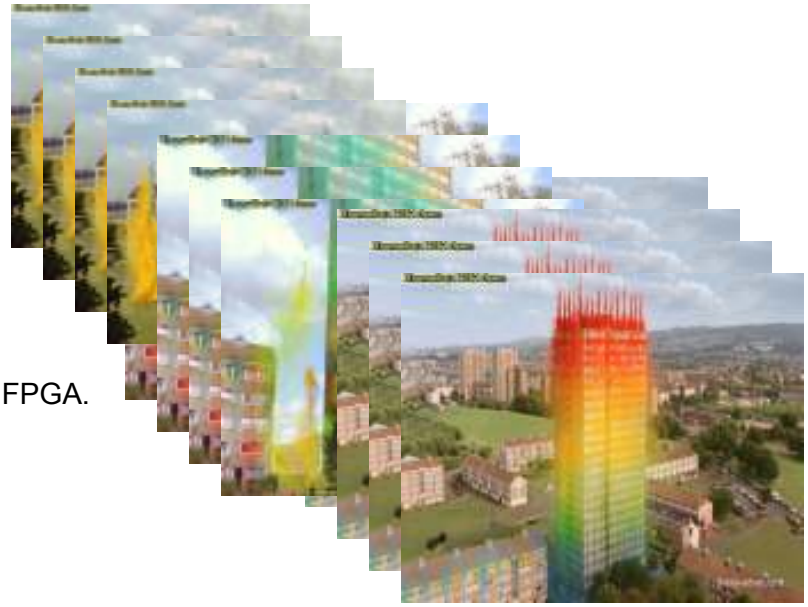

FPGA Module: In the news…

› XtremeData featured in AMD Annual Analyst Day presentation on "Torrenza" initiative, 1 June 2006

› "Altera Launches High-Performance Computing University Program with AMD, Sun and XtremeData", press release, 21 Aug 2006

› Plus much more exposure in print and internet media….

University Program

- **Objectives:**
  - Research and drive adoption of:
    - FPGA co-processing
    - Medical imaging, data analytics, text searches, network security, bioinformatics, & energy apps.
    - Programming tool chain
- **Partners:**
  - Altera, AMD, Sun, XtremeData
- **20 Kits Planned (Maybe more in 2007)**
  - U. of Illinois, Stanford, U. of Florida, Purdue, CMU, UCLA, Boston U, U of Mannheim, etc.

XtremeData XD1000
FPGA co-processor module

AMD Opteron

Sun Ultra 40 Workstation

ALTERA. AMD Sun microsystems XtremeData, Inc.

Thank You!