Education as a New Direction for Standards in Reconfigurable Computing

Dr. Eric A. Stahlberg
Computational Science Director, Wittenberg University
President, OpenFPGA
estahlberg@wittenberg.edu

Today's Presentation

- Context for reconfigurable computing today
- Lessons, challenges and opportunities
- Economics sustains progress
- Education effects change
- Changing times
- Aims and opportunities

Reconfigurable Computing Goal

- *"Pervasive use of reliable reconfigurable technologies across computing platforms and environments"*
- Why?
  - FPGAs are power efficient
  - FPGAs perform well in several areas of growing need
  - FPGAs can be greatly optimized
  - All options for affordable performance improvement will eventually be tapped for HPC and Supercomputing

Has FPGA Interest Faded?

- It's all relative
  - Attention diverted to GPUs by researchers
  - Attention diverted to GPUs by industry
    - Nvidia, AMD, Microsoft, several startups
  - Startup costs for GPUs lower
  - OpenCL has moved forward
  - Intellectual property opportunities greater
  - Software development faster than hardware design
    - Hardware Engineering – limit functionality to minimize sources of error while maximizing performance
    - Software Development – expand functionality to maximize possible areas of use
    - Not yet pulling fully together

Has Interest in FPGA Standards Ended?

- GPU interest shows greater immediate potential
  - Near term return on investment (ROI) appears better
  - Changing industry priorities
- Consolidation and migration
  - DRC moved into security vertical
  - XtremeData shifts into data vertical
  - Yet several vendors remain
- Significant barriers still to be addressed
- Possible apprehension due to GPU popularity

Adoption Takes Time

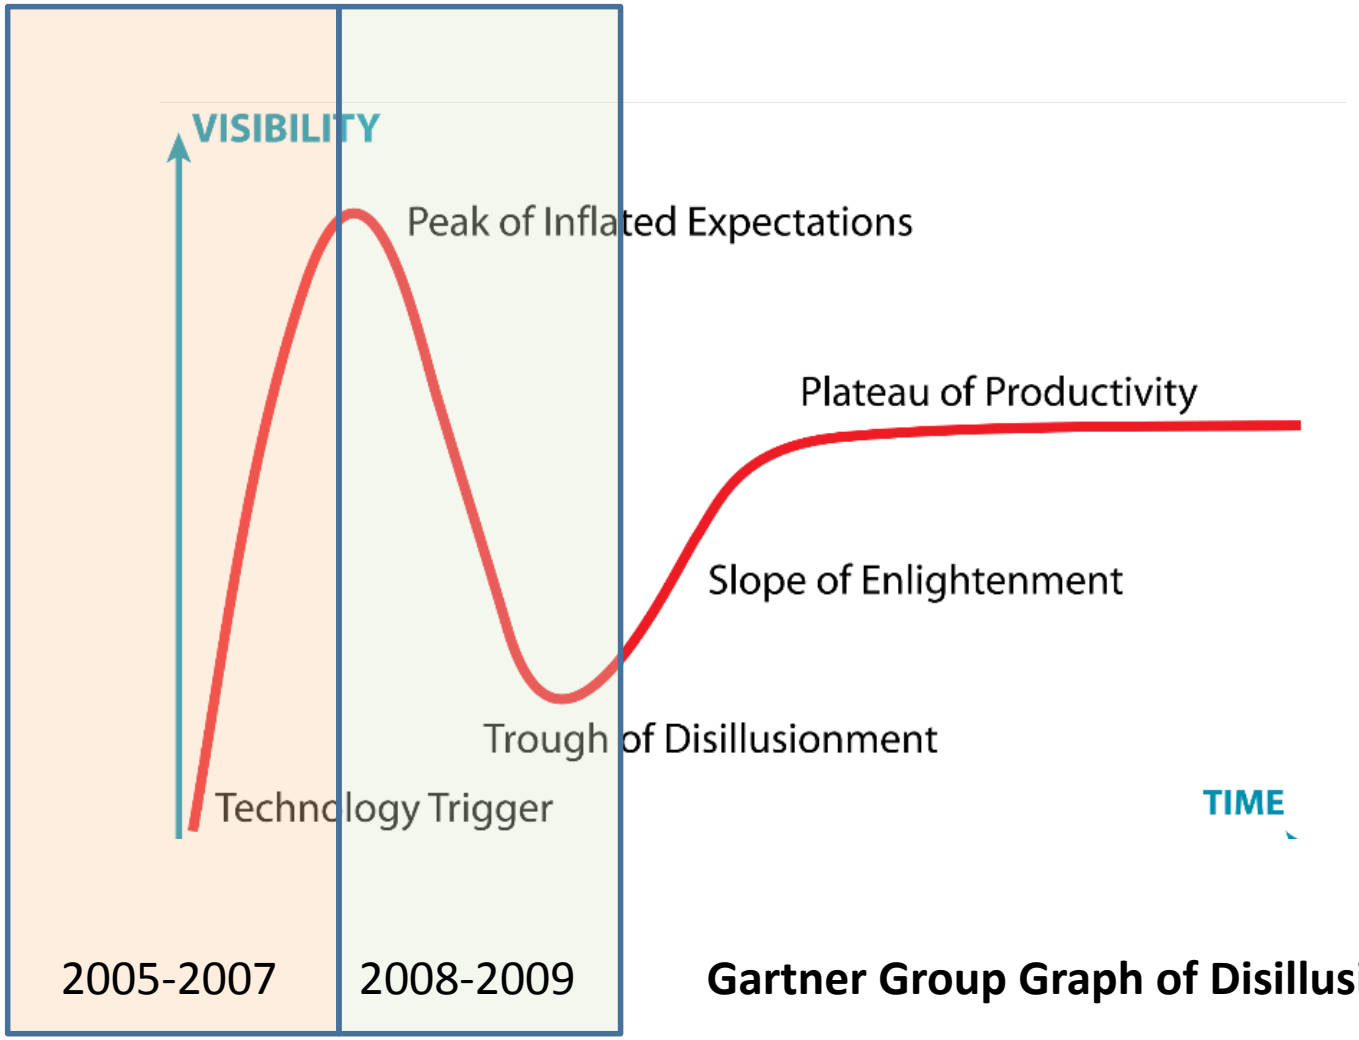

Gartner Group Graph of Disillusionment

Selected HPC Trends

- HPC will resume growth in mid 2010
- Global leadership pursuits will spur efforts forward
- Evolution before revolution
- GPUs gain traction on x86 growth
- Power/cooling issues continue to grow, but not rapidly
- Cloud computing is coming
- Commoditization will continue to level playing field

Selected and paraphrased from IDC presentation 2010 Predictions

The Changing HPC Landscape

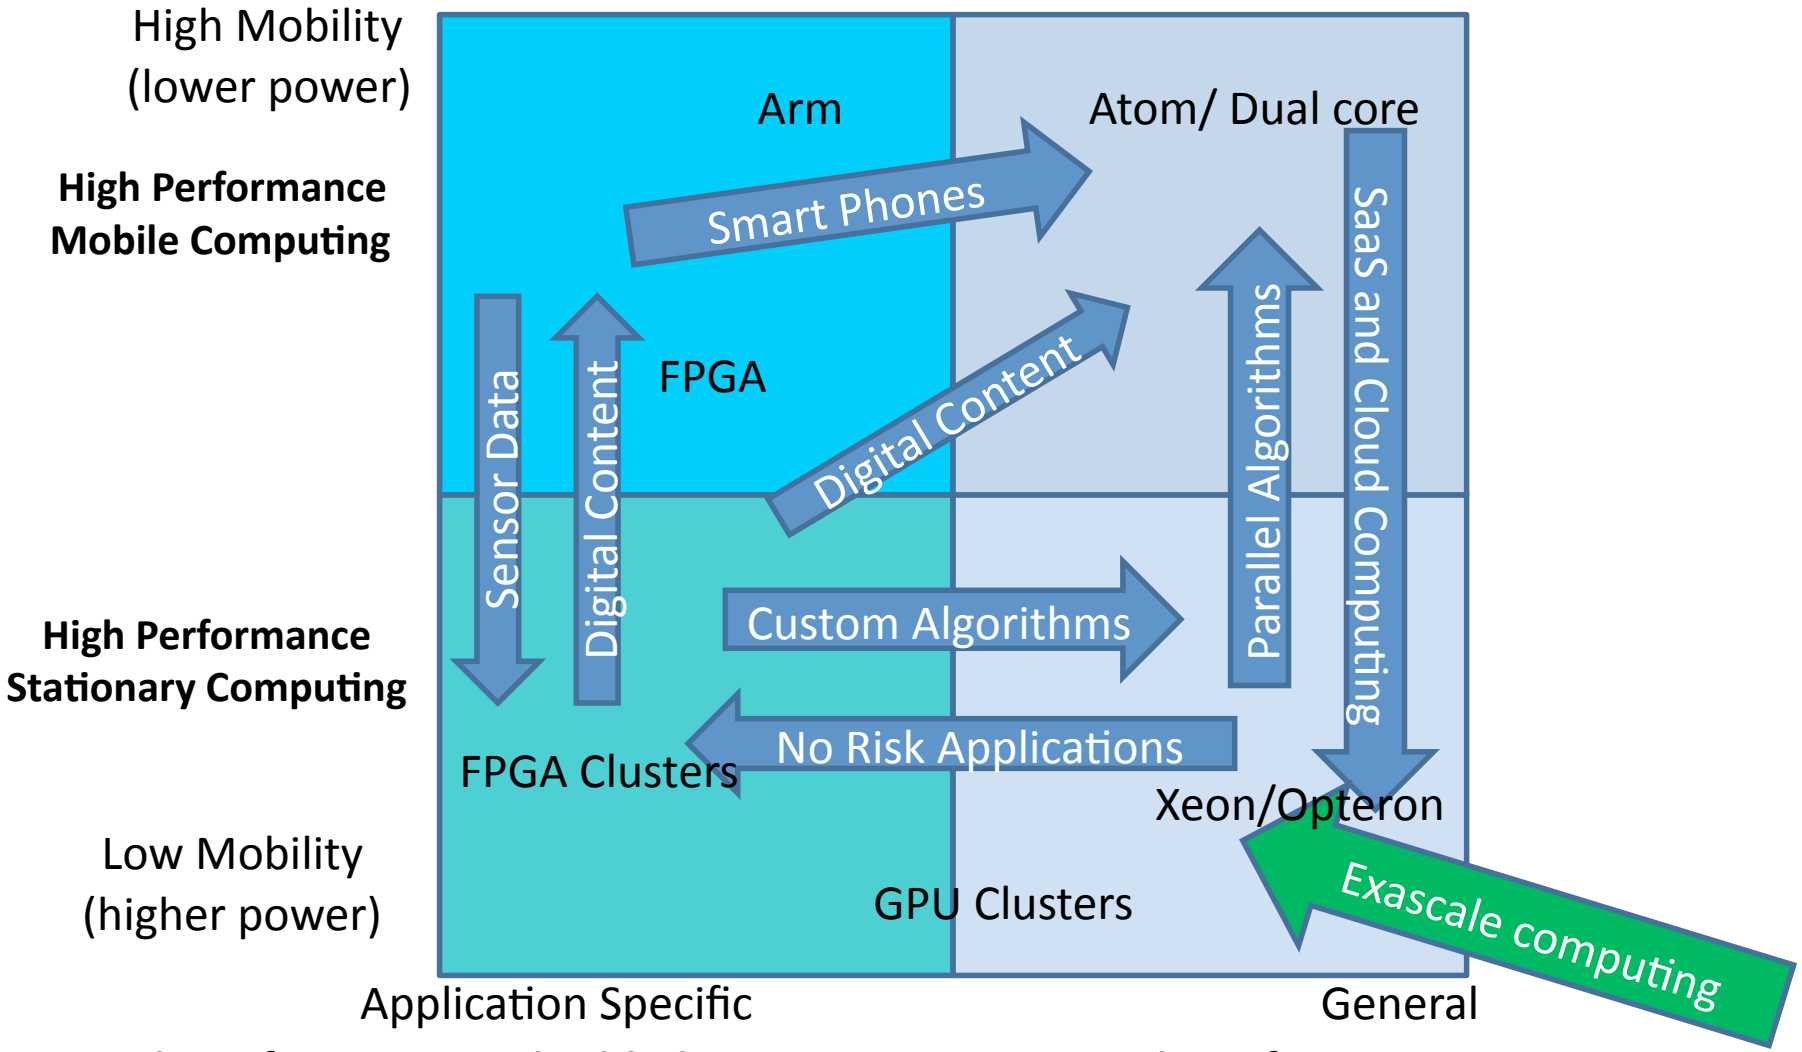

Reconfigurable Supercomputing Today

- Are FPGAs dead?
  - Will CPUs solve all problems?
  - Will GPUs fill in the gaps?
  - Will performance continue to matter?
- FPGAs are not dead.
  - Parallel processing wall will force code and system redesign
  - GPUs are providing incentive to redevelop code bases
  - Products boosting productivity of existing resources are expected to be attractive
  - Data volumes are growing and FPGAs excel at data processing
  - Power efficiency is increasingly important
  - FPGAs are growing in capability per unit cost
- The door remains open for reconfigurable supercomputing

Drivers of Change

- Innovation sparks progress
- Business value sustains opportunities
- Return on Investment (ROI) is key
  - A greater ROI wins the day

Return on Investment Sustains Progress

Scenario 1 :
(niche market)

$$\frac{\text{Reliable products in small volume and larger margin}}{\text{Tool costs + device costs + development costs}}$$

Scenario 2:
(mass market)

$$\frac{\text{Reliable products in large volume and smaller margin}}{\text{Tool costs + device costs + development costs}}$$

How to improve the FPGA Application ROI?

Reconfigurable Computing and Return on Investment

- Investment Costs
  - Cost of tools (not open, and competition limited)
  - Cost of devices (getting cheaper, but still expensive)
  - Costs to develop (time to create applications still high)
- Return Potential
  - Scenario 1: Reliable products of large margin and small volume
  - Scenario 2: Reliable products of small margin and large volume

GPU Lessons for a Successful ROI

- Reduce costs of tools to create applications
  - CUDA was available for free
- Have inexpensive devices for development
  - GPUs were/are cheap to develop with
- Shorten time to develop applications
  - Short cycle times
  - Solutions could be readily used on other compatible devices
- Grow the number of reliable applications
  - Efforts abound with increasing numbers of enhanced applications and libraries

$$\textbf{Good ROI} = \frac{\textbf{Reliable products in large volume and small margin}}{\textbf{Low tool costs + low device costs + short development time}}$$

Education Hits Multiple ROI Aspects

- Education increases confidence in solutions
  - Reliable benchmark information to grow the market
  - Accountability for reconfigurable
- Education reduces development time
  - Increases number of knowledgeable application engineers
  - Improves cohesion of combined software development and hardware design projects
- Successful educational programs inspire
  - Inexpensive technology and tools
  - Portability of educational examples across time and technology
  - Contributions to a broader body of knowledge

Accelerators to Applications:

Supercharging the Undergraduate Computer Science Curriculum

- NSF funded (CISE) project started August 2009
- Pervasive parallel/accelerated computing for undergraduates
- Goal is to provide broader background of future 'Application Engineers' to better develop tomorrow's applications
  - Emphasizing reliability and performance
  - Integrating generality
  - Blending innovative solutions to real-world problems
- Partners Wittenberg University and Clemson University
  - PI's: Eric Stahlberg and Melissa Smith
  - Wittenberg faculty: Steven Bogaerts, Brian Shelburne, Kyle Burke
- Emphasize integration across domains
  - undergraduate and graduate experience
  - disciplines (computer science and engineering with computational science applications)

Open fpga

wittenberg UNIVERSITY

Key Project Elements

- Enhancements for six computer science courses
  - **Introduction to Programming** (using GPU libraries)
  - **Computer Organization/ Computer Architecture** (multi-core designs and processor communication topics)
  - **Algorithms** (multi-level parallelism and performance evaluation)
  - **Programming Languages** (include examination of emerging parallel languages)
  - **Data Structures** (multi-threaded instances for common data structures)
  - **Software Engineering** (components, APIs and robustness)
- Impacts three computational science courses
  - **Computational models and methods**
  - **Bioinformatics**
  - **Computational chemistry**
- Summer research experiences
  - Undergraduates working with faculty and graduate students in summers
  - Industry participation to validate experiences
  - Validation with internships

Early Progress

- First regional HMC workshop planned April 8-9, 2010
- Smith group has analyzed optimization techniques across technologies (paper due out next month)
- Profiling complete for several applications
  - LAMMPS, GROMACS, CPMD, Neural Networks
- Information distribution systems nearly ready (via OpenFPGA)
- First major push for undergraduates will be summer 2010
  - New educational content for undergraduates
  - Projects and common APIs

A Bright Future for Reconfigurable Supercomputing

- Current cycle of FPGA adoption is approaching productivity phase
- FPGAs are increasingly more capable
- Power and cooling issues are not solved
- Parallel programming is still a challenge
- GPUs have opened the door for software and system redesign
  - ORNL convened Heterogeneous Multi-core Computer Consortium and met in January spanning technologies
- Exa-scale computing is demanding new approaches
- Formation of HiPEAC and other organizations
- Priorities in EU funding for RC
- NovoG consortium (just yesterday)
- This conference continues

What's Ahead for Reconfigurable Computing

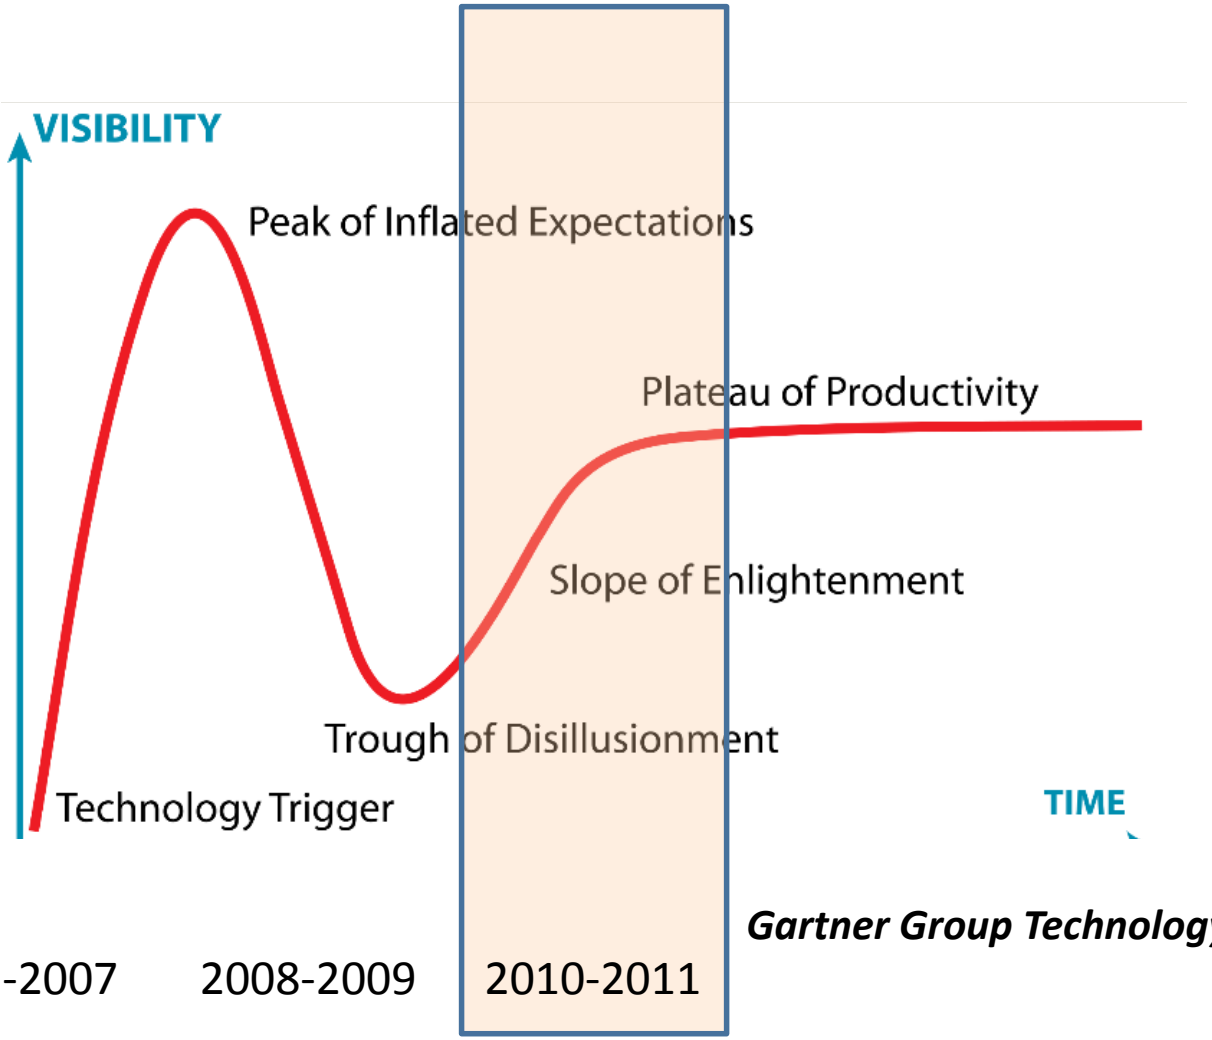

Gartner Group Technology Adoption Cycle

EU Priorities in RC for 2010-2014

- Focusing on innovation and prototypes
- Emphasizing multi-core and parallelism
- Driving industry needs
  - Network infrastructure, robotics, content and media
  - Components and systems engineering
- Transition to multi-core across the spectrum
- Remain focused on single host systems
- 45 million Euro budget with call open summer 2010
- Areas of emphasis
  - Parallelism
  - **International collaborations**
  - Customization and tool chains
  - **Virtualization**
  - **New architecture and alternatives to von Neuman solutions**

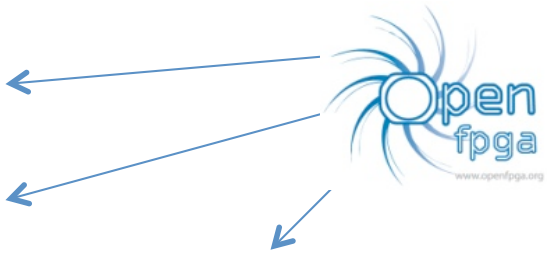

Exascale – Reconfigurable Supercomputing?

- Characteristics of future Exa-scale solutions
  - Ultra-efficient
  - Use of big market consumer market technologies
    - Embedded and low-power ideas from mobile devices
  - Specific targeted solutions
  - **Green Flash** Feasibility Study at LBNL
    - Goal to influence HPC design
    - Design for low power and greater concurrency
    - Lower clock rates
    - Simpler cores
    - Tailor design to application

From Horst Simon presentation April 2009

Green Flash Approach

- Choose application targets first
- Design systems for applications
- Leverage power-efficient embedded technology
- Design hardware, software, and algorithms together using auto-tuning and emulation
- Achieve Exa-scale sooner and more efficiently
- Expected to augment existing HPC systems

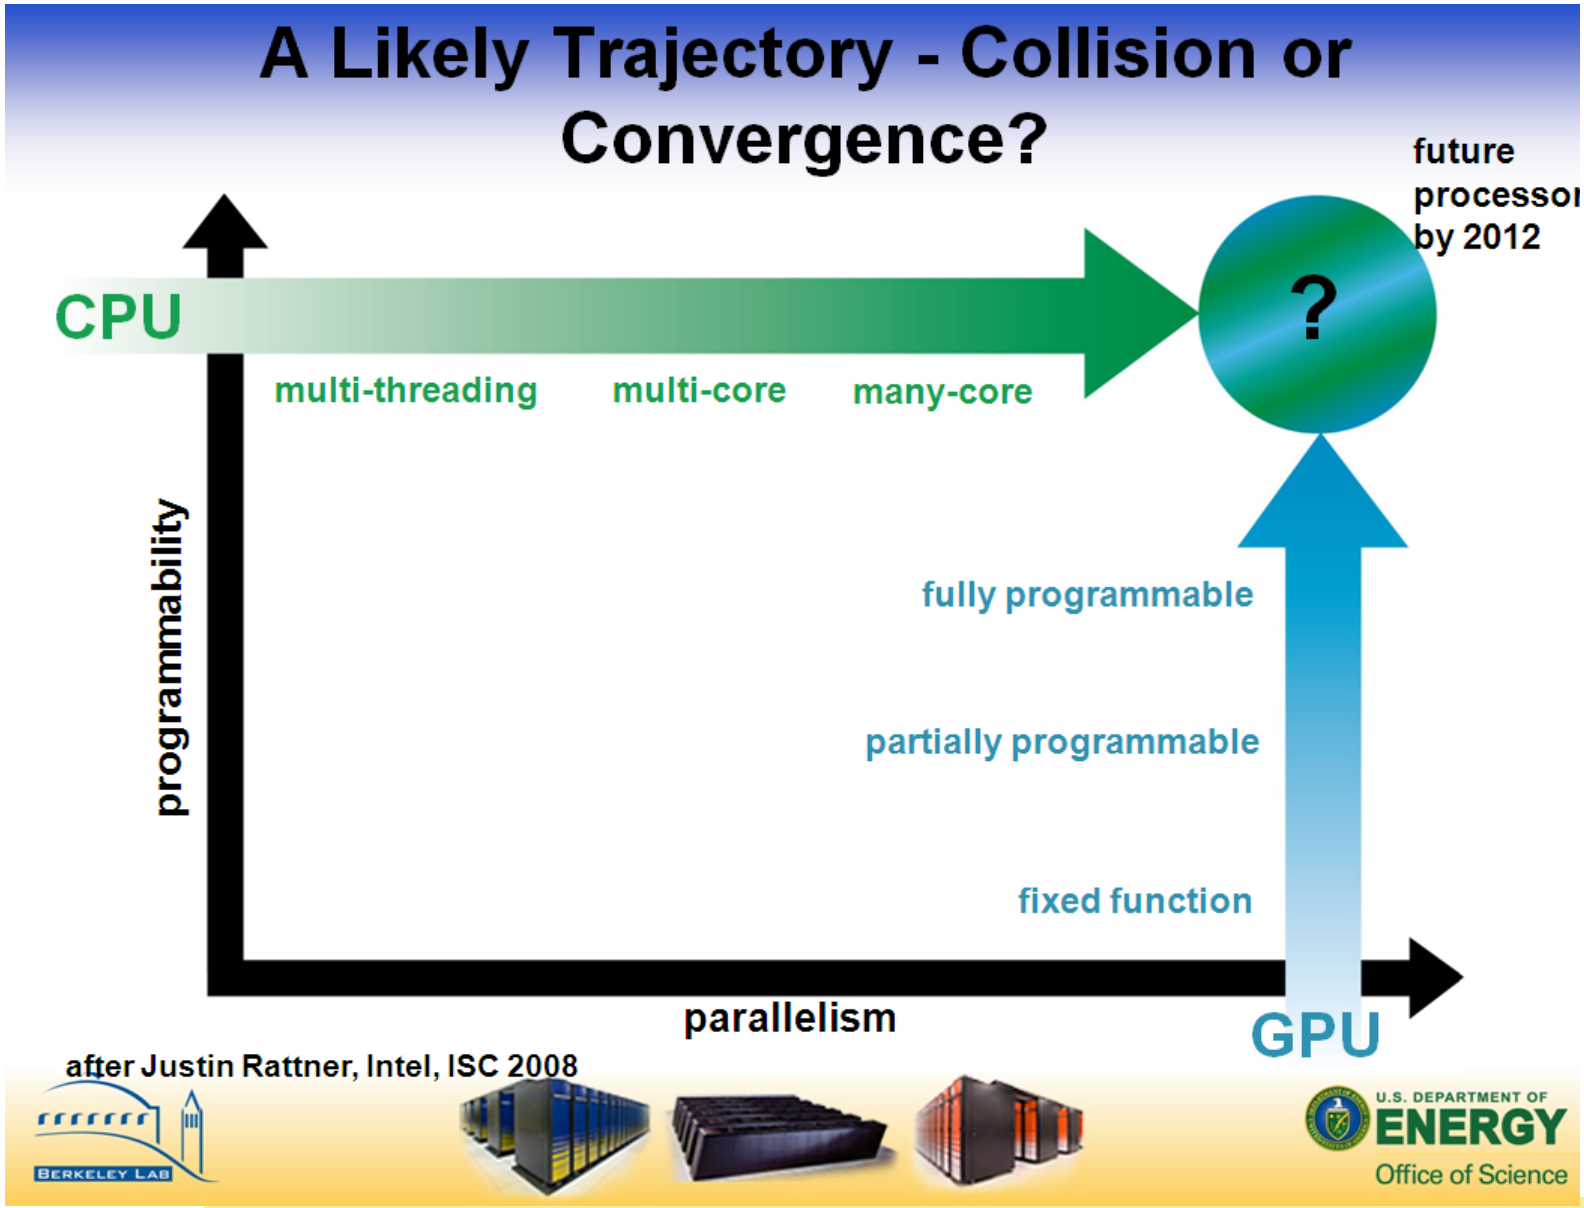

Source: Horst Simon Presentation from April 2009

Time to review…

What technology embodies?

- Massive core counts
- Full programmability
- Power efficiency
- Adaptability
- On-the-fly tuning

Reconfigurable Technology

Open fpga
www.openfpga.org
wittenberg UNIVERSITY

Top Barriers to RC Supercomputing

1. Programming (languages and models, ease of programming)
2. **Standards (communications, heterogeneous interconnection, interfaces, portability)**
3. Tools
4. **Education**
5. Costs
6. Development time

(source OpenFPGA survey, November 2009)

Programmability Status

	OpenMP	MPI	OpenCL	CUDA	Brook+	libspe	VHDL/Verilog	Mitrion-C
Target platform	CPU CBEA	CPU	CPU GPU CBEA	GPU (NV)	GPU (AMD)	CBEA	FPGA	FPGA
Availability	Win Linux Mac	Win Linux Mac	Win Linux Mac	Win Linux Mac	Win Linux	Linux	Win Linux Mac	Win Linux Mac
Abstraction	Pragmas	API	API	API, compiler	API, compiler	API, compiler	API	Compiler
Host language	C, C++, Fortran	C, C++, Fortran	C	C, C++	C++	C, C++, Fortran	"C-like"	C
Kernel language	—	—	C99-based	C99-based, some C++	C99-based	C, Fortran, almost C++	—	C
Memory model	Shared	Distributed	~ PGAS	~ PGAS	Data streams	~ PGAS	all	all
Data-parallelism	Global view	—	SPMD, SIMD	SPMD	SPMD	MPMD	all	all
Task-parallelism	—	Explicit	Full	Streams	Streams	Explicit	all	all
Ease of use	★★★	★	★★	★★	★★	★	★	★
Maturity	★★★	★★★	★	★★★	★★	★★★	★★★	★★

Source: Brodtkorb, Dyken, Hagen, Hjelmervik and Storaasli, Scientific Computing

*note – not all programming environments for FPGAs were evaluated.

OpenFPGA Overview

- Non-profit international organization promoting standards in reconfigurable computing
- Combined industry, government, academic, international at grass roots level and above
- Member and sponsor supported US 501c3
- Board includes ORNL, GE, SRC, ZIB, OSC, Optngn, NCSA, Wittenberg, and independents
- Over 500 participant mailing list, with over 40 countries
- Recognizing and promoting progress in standards for RC computing internationally
- Working groups and online services
- Membership drive forming for 2010

OpenFPGA: Taking Research to Practice

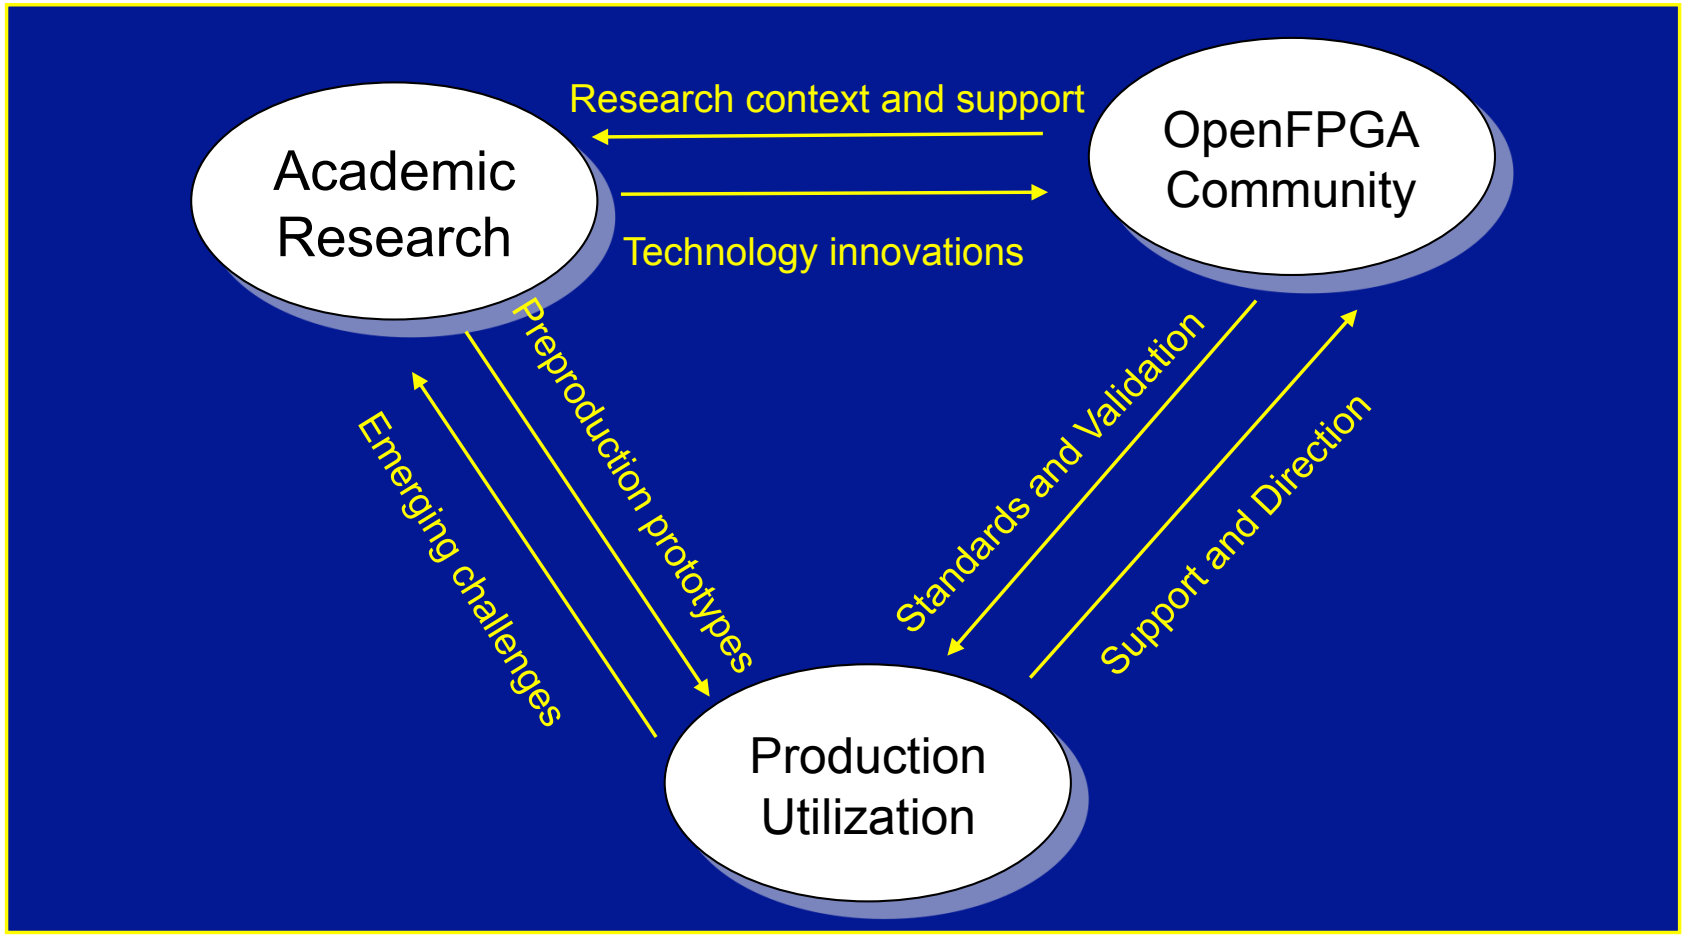

OpenFPGA Working Groups

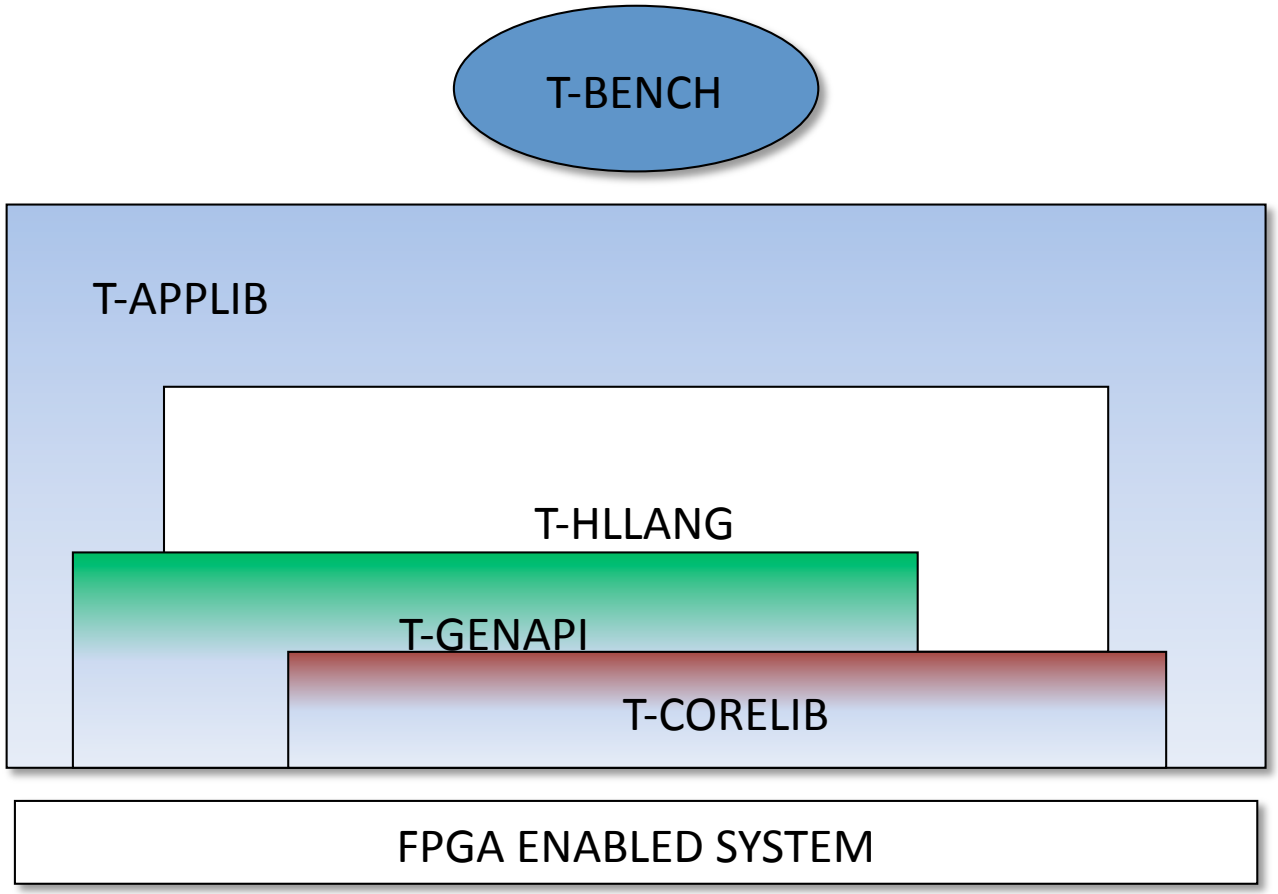

OpenFPGA Efforts to Improve Sustainability

- Early emphasis to reduce development time
  - APIs for portability (Gen API)
  - Core standards to enhance reusability (CoreLib)
  - Application libraries for key areas (Applib)
  - Validation Lab for members
- Efforts still pending
  - Increasing number of reliable applications
  - Reducing costs of devices
  - Reducing costs of tools

Aims for 2010

- Build a unified global body of knowledge around reconfigurable supercomputing
  - Application information, benchmarking, examples, etc.
- Education as a means to motivate standards, sharing, portability and low cost
- Education as a means to grow numbers of deployed solutions
- Lower barriers to broader adoption to improve ROI

Visit www.openfpga.org and sign up on the participant list

OpenFPGA Priorities for 2010

- Next stages for early initiatives
  - Reduce net development time with common API
  - Improve reuse with core standards definitions
  - Continue with award at HPRCTA workshop at SC10
- Build participation, membership and sponsorship
  - Provide new services of value to the community
  - Reenergize working groups
- Build on current interests and directions
  - Increase number of reliable applications
  - Improve quality of available performance information
  - Bring focus to broader community education
  - Bring focus to validation and verification

Specific OpenFPGA Efforts

- APIs – GenAPI working group
  - Bring forward a first standard that all agree to use
  - standards-based clusters will be appealing (*source IDC)
- Cores – CoreLib working group
  - Bring forward a base set of knowledge on each core/component
- Reliable applications – Applib working group
  - Create reliable FPGA libraries for use
- Component registration – Software and online service
  - Register components for tracking and visiblity
  - Arriving online in April – join the OpenFPGA participant list
- Refocused Benchmarking working group forming
- Education working group forming

RAC System Efforts Continue

- Register components with trusted third party
  - Upload metadata
  - Register specific file (VHDL, bitstreams, libraries, etc.)
  - Get a reference ID for each component
  - Preserves IP investment
- Raise overall visibility of efforts and expertise
- New features being planned for future releases
- MRSC 2009 – command line version announced
  - Use internally to define and keep metadata
  - Transfer metadata to public search site
- MRSC 2010 – web-version entering friendly user phase
  - Annotate and upload object directly to website

OpenFPGA Support

- Membership and Sponsorship
  - OpenFPGA is all volunteer
  - Memberships provide sustaining support
  - Sponsorships support major and strategic initiatives
- Grant Support
  - OpenFPGA is partnering in emerging grant opportunities
  - Supports building Open FPGA body of knowledge
- Community Involvement
  - Participate at no cost
  - Contribute to projects
  - Exploring new connections (e.g. HiPEAC in Europe)
- Interest continues to grow
- Visit **www.openfpga.org and** sign-up as free participant or paid member

HPC Key Challenges Remain in 2010

- Weak application performance improvements
- Effective highly parallel programming
- System imbalance
- Power and space usage
- Software licensing
- Ease-of-use

Source: IDC HPC reports – Feb 2010

A Challenge to the Community

- Opportunity is Now
- Take on the barriers of programming, tools, and reducing overall cost
- Get involved locally and globally
- Contribute to the common body of knowledge
- Support broader efforts
- Reduce the costs of tools to create applications
- Look increasingly at volume markets

Acknowledgements

- Melissa Smith and her research group at Clemson
- OpenFPGA board, sponsors and members
- Ohio Future Jobs
- National Science Foundation CISE division
- Wittenberg University and colleagues
- US Department of Energy
- MRSC program committee for the opportunity today

Thank you.

estahlberg@wittenberg.edu
estahlberg@openaccelerator.org
http://www.openfpga.org

Additional Barriers to Sustainable RC

- Other barriers mentioned in survey
  - Technical
    - Memory size and speed
    - Bandwidth between FPGA and general-purpose front end
    - High level languages to HDL
    - Adaptability
  - Non-technical
    - Software
    - Availability
    - Robust codes for end-users
    - Industry buy-in
    - Simplistic usage, invisible to users
    - Sales
  - Personnel
    - Productivity