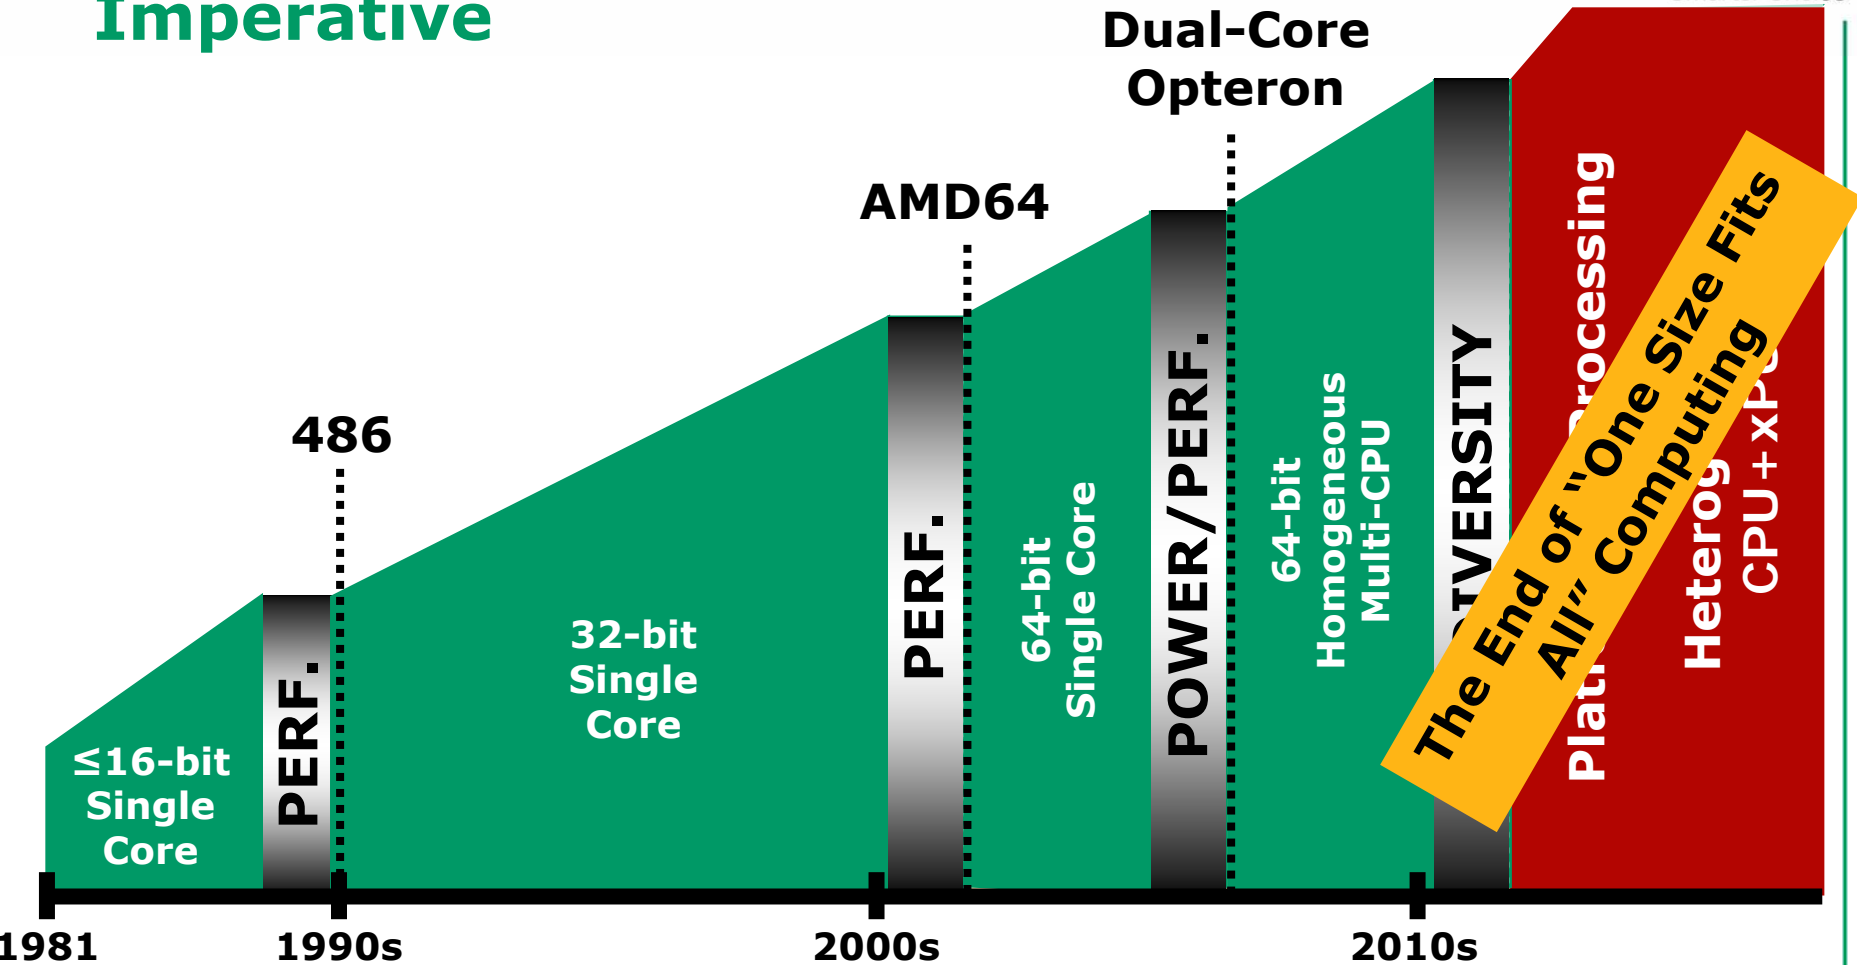




CoEHT Symposium

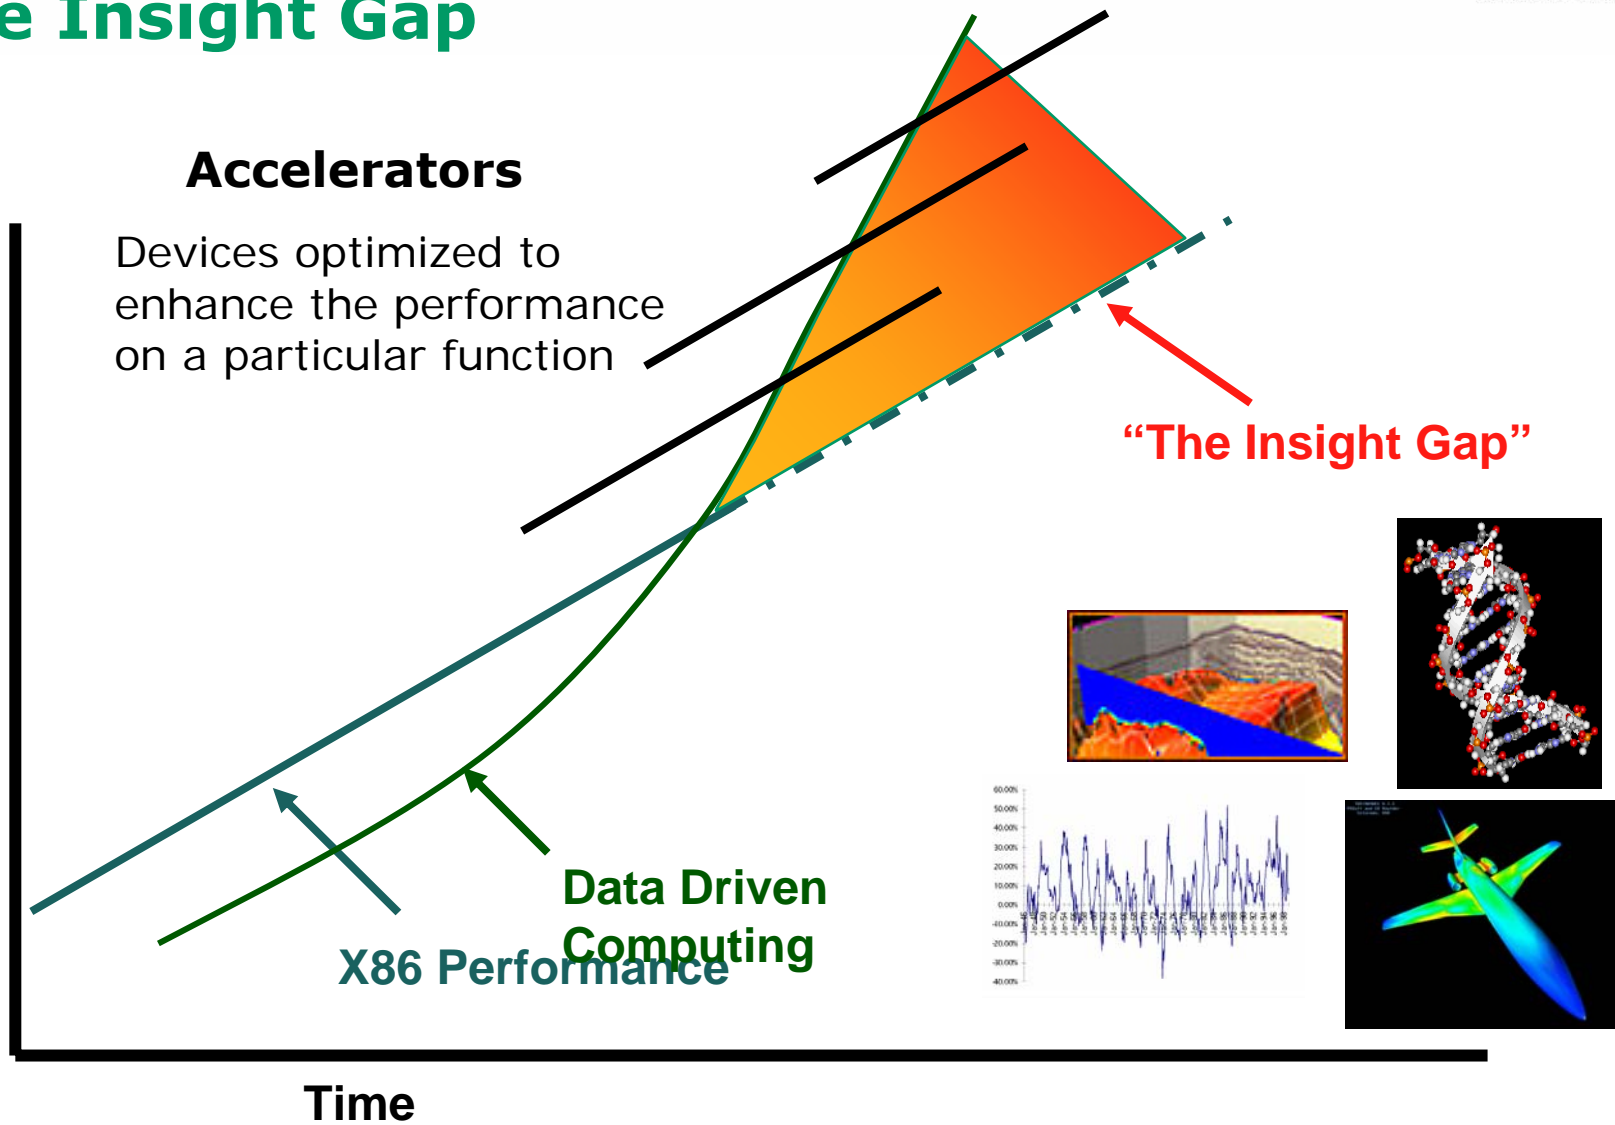
February 16, 2007
Douglas O'Flaherty

The Heterogeneous Processing Imperative



By the end of the decade, homogenous multi-core becomes increasingly inadequate

Industry Landscape: The Insight Gap



Torrenza and the Pareto Distribution

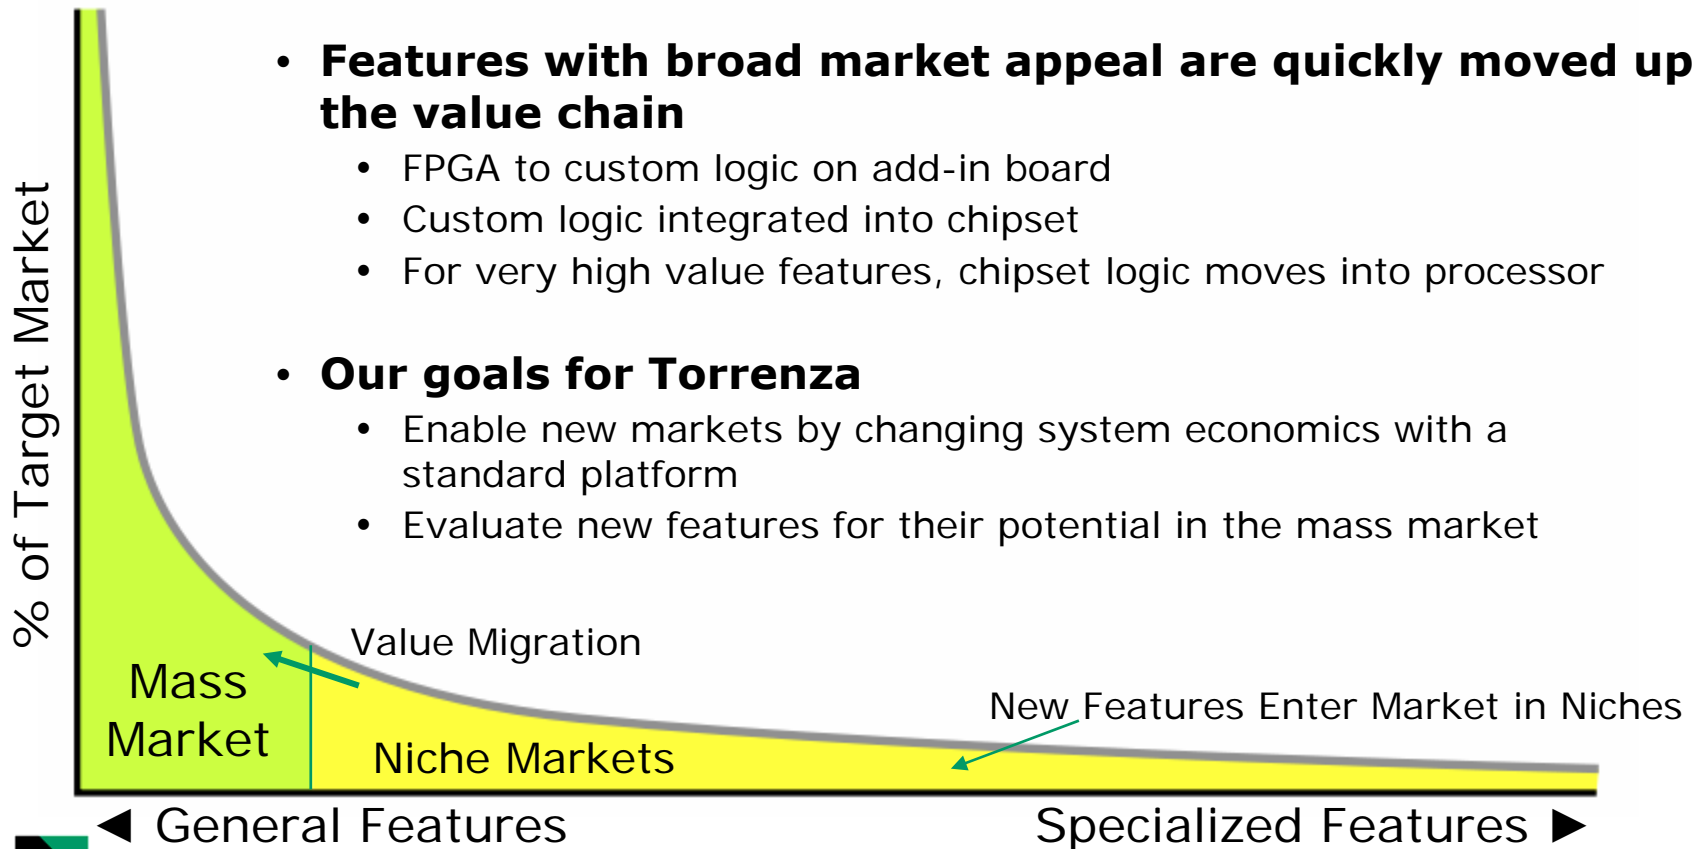
- **New features are introduced in niche markets**
 - Some features will never reach broad market appeal, but will have stable niche markets over time
 - The sum of those niche market opportunities is itself a considerable market opportunity

- **Features with broad market appeal are quickly moved up the value chain**

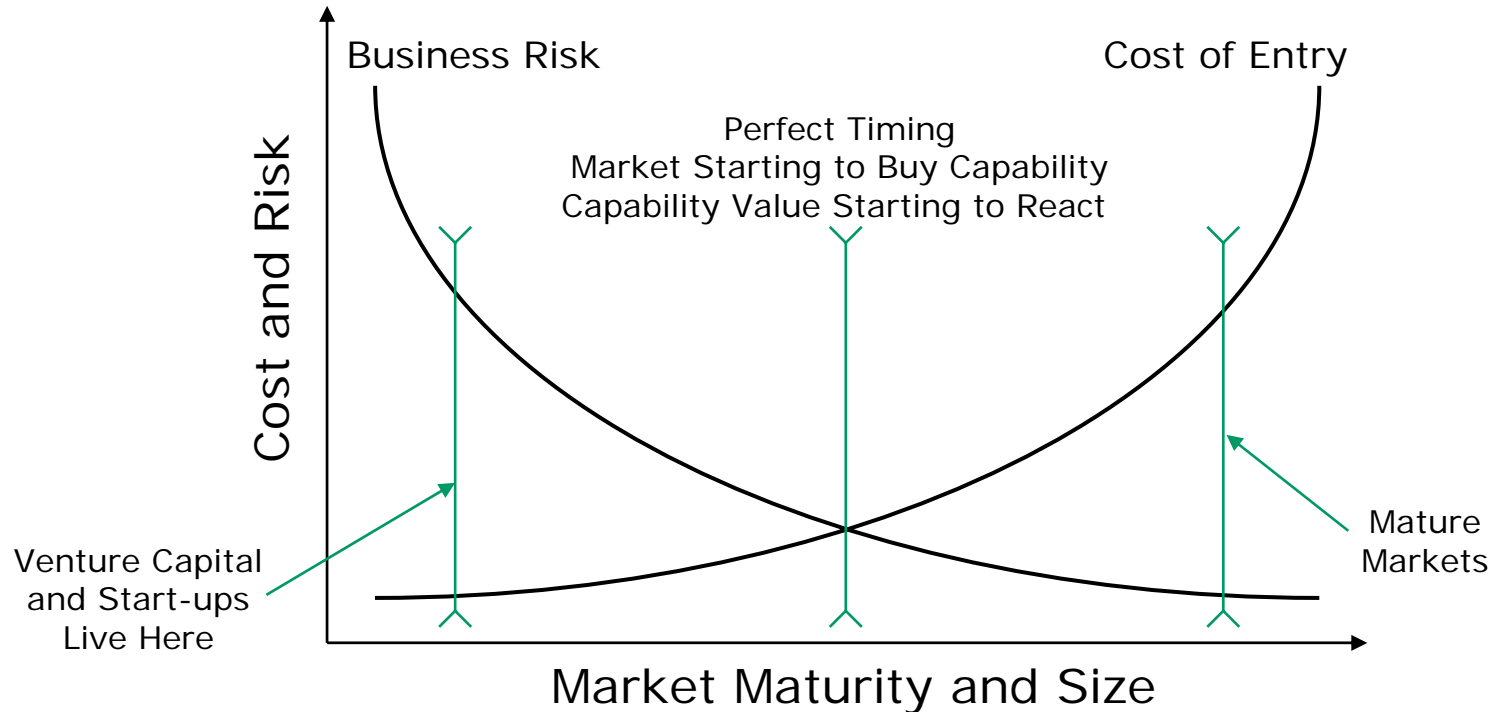
- FPGA to custom logic on add-in board
- Custom logic integrated into chipset
- For very high value features, chipset logic moves into processor

- **Our goals for Torrenza**

- Enable new markets by changing system economics with a standard platform
- Evaluate new features for their potential in the mass market



Early vs. Late Market Entry Managing Risk and Cost



- **New capabilities carry inherent business risk**
- **Managing risk is a business philosophy**
 - Venture Capital accepts high risk for potentially high returns
 - Most mature enterprises only accept low risk, but end up paying top-dollar for what are essentially new divisions
- **Torrenza gives the industry an early assessment of emerging capabilities**

Torrenza Platform Challenges

End-Customers

Barriers to Application Performance

- Heterogeneous computing & application software
- Synchronization & Communication overhead



OEM & Channel

Barriers to Support

- Proliferation of SKUs
- Economic challenges to differentiate



Innovators

Barriers of Entry

- Inflexibility of legacy platforms
- Need innovation roadmap from concept to deployment



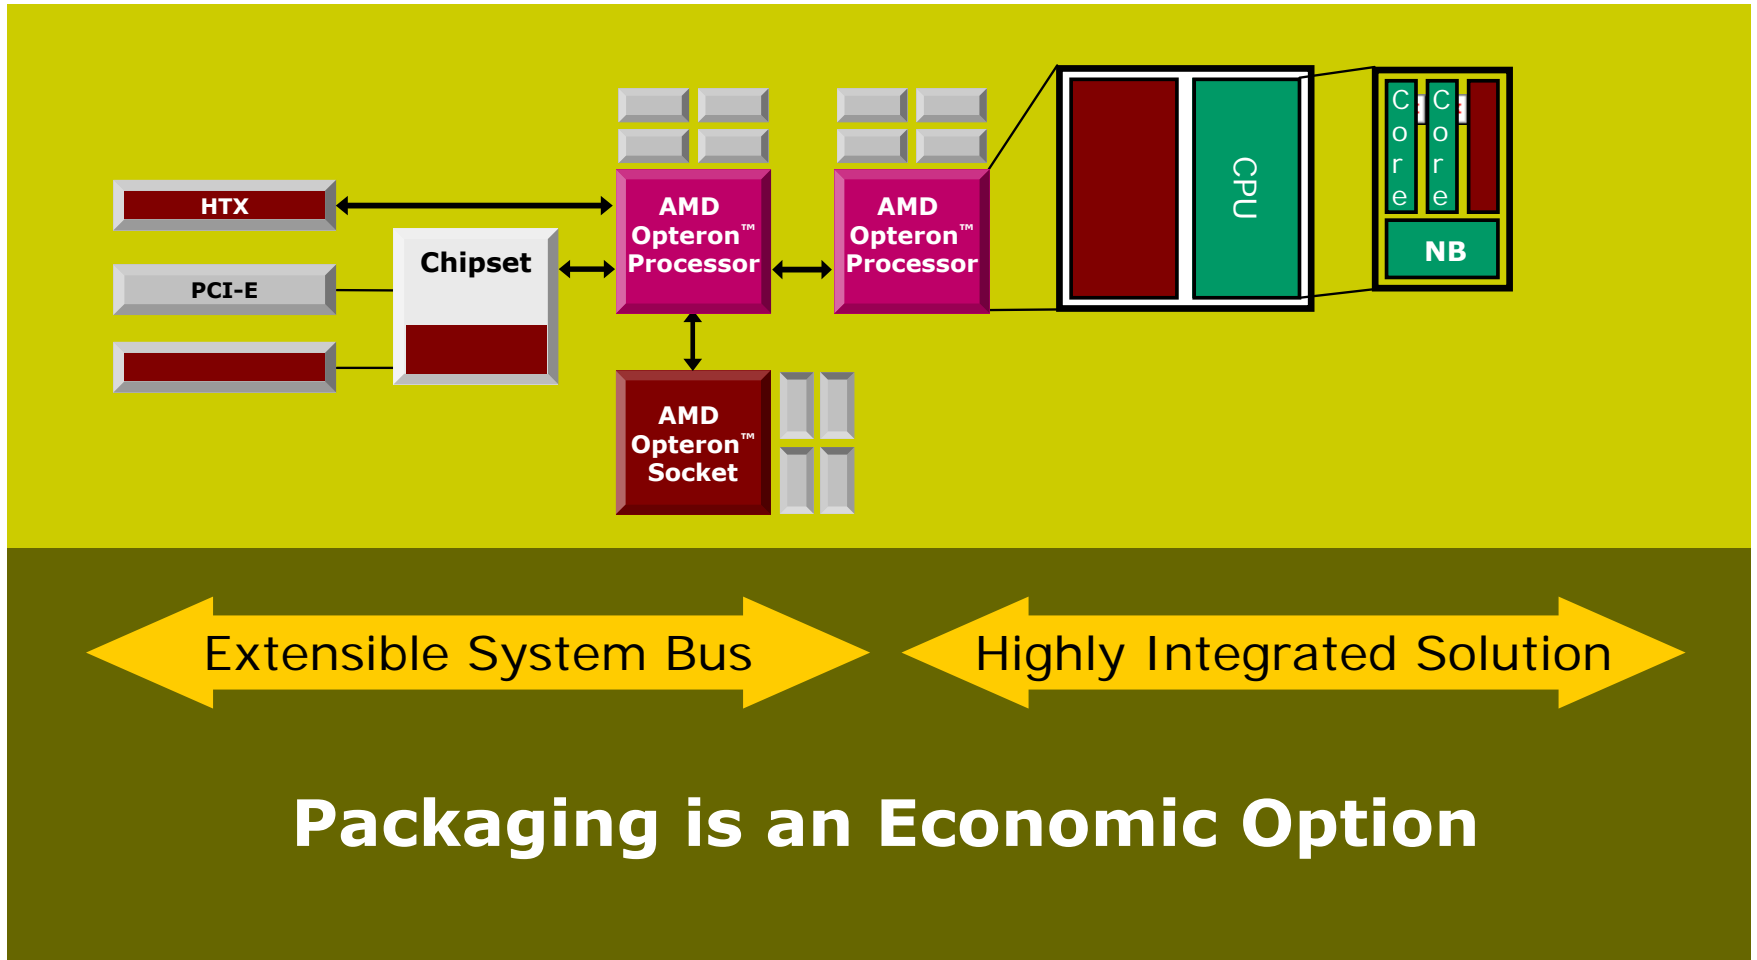
Torrenza Platform Challenges

Enrich the IT industry with AMD64 technology as the **OPEN, STANDARDS BASED Innovation Platform** for our customers and end users.

- Freedom of Choice
- Longevity and Stability
- Differentiated and Upgradeable



Acceleration Adoption Model

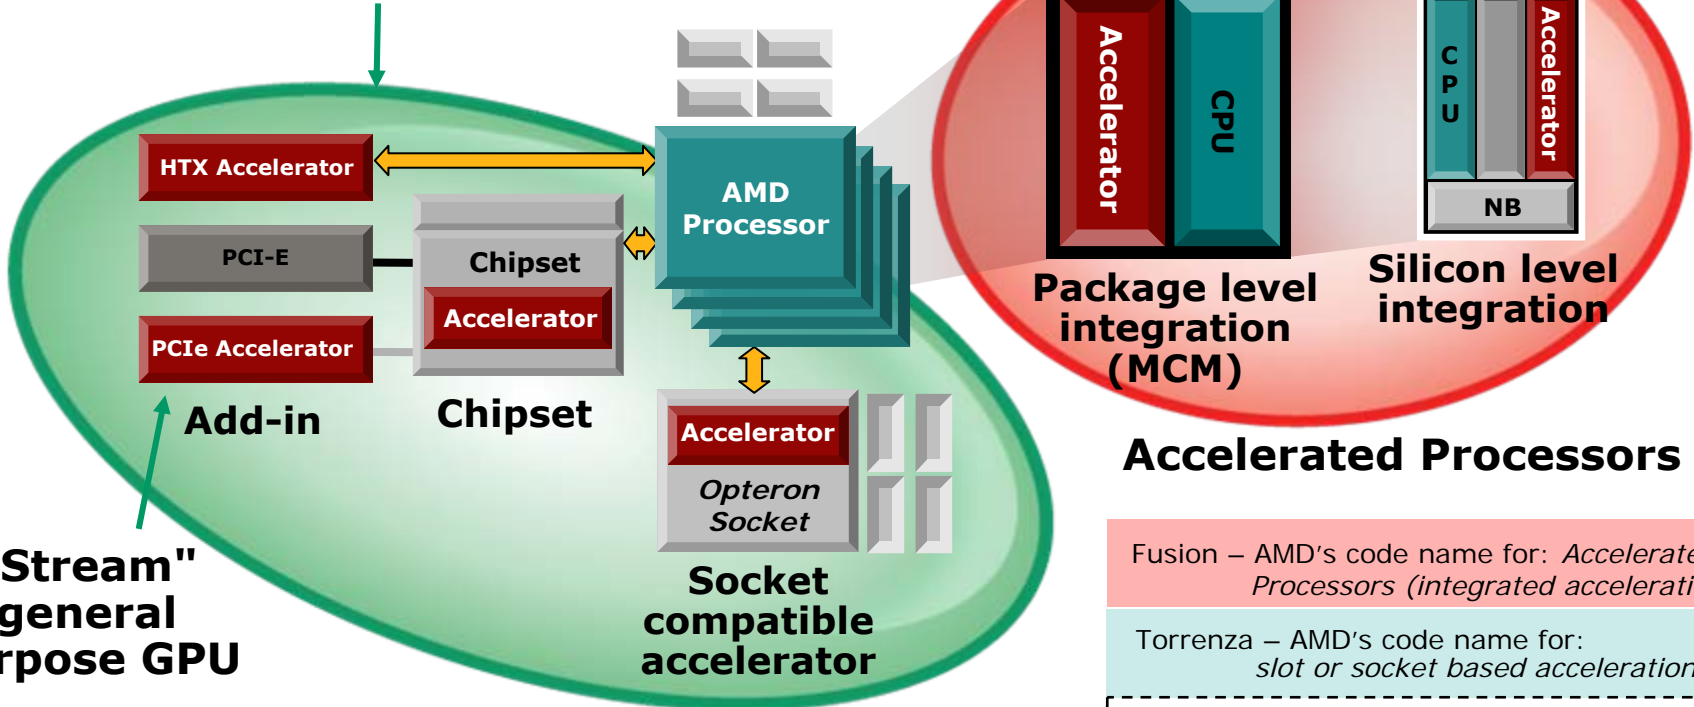


Continuum of Solutions

Accelerated Computing

"Fusion"

"Torrenza"



Accelerated Processors

Fusion – AMD's code name for: *Accelerated Processors (integrated acceleration)*

Torrenza – AMD's code name for: *slot or socket based acceleration*

Stream – *Specific example of a GPGPU accelerator under Torrenza*

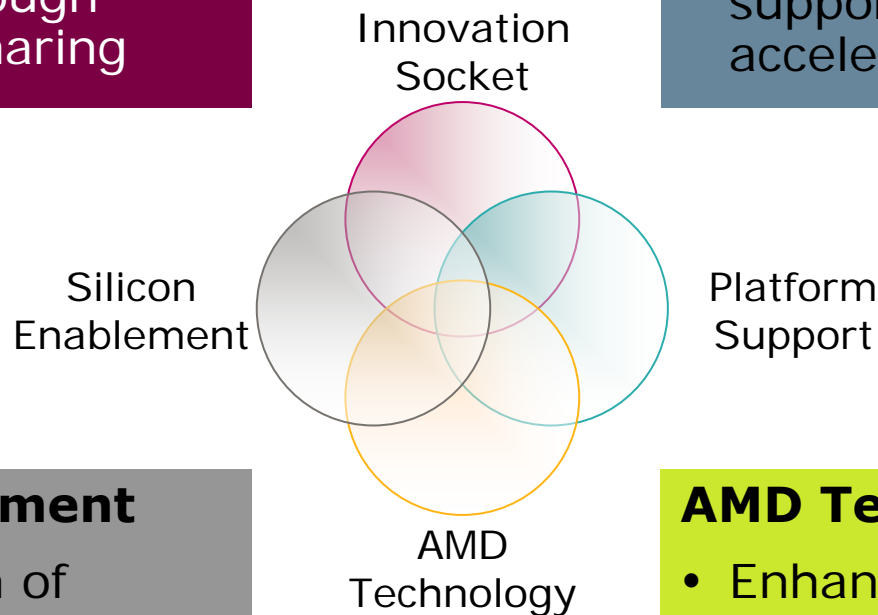
Torrenza Initiative Elements

Innovation Socket

- Propagating *Freedom of Choice* through technology sharing

Platform Support

- Standards platform support for accelerators



Silicon Enablement

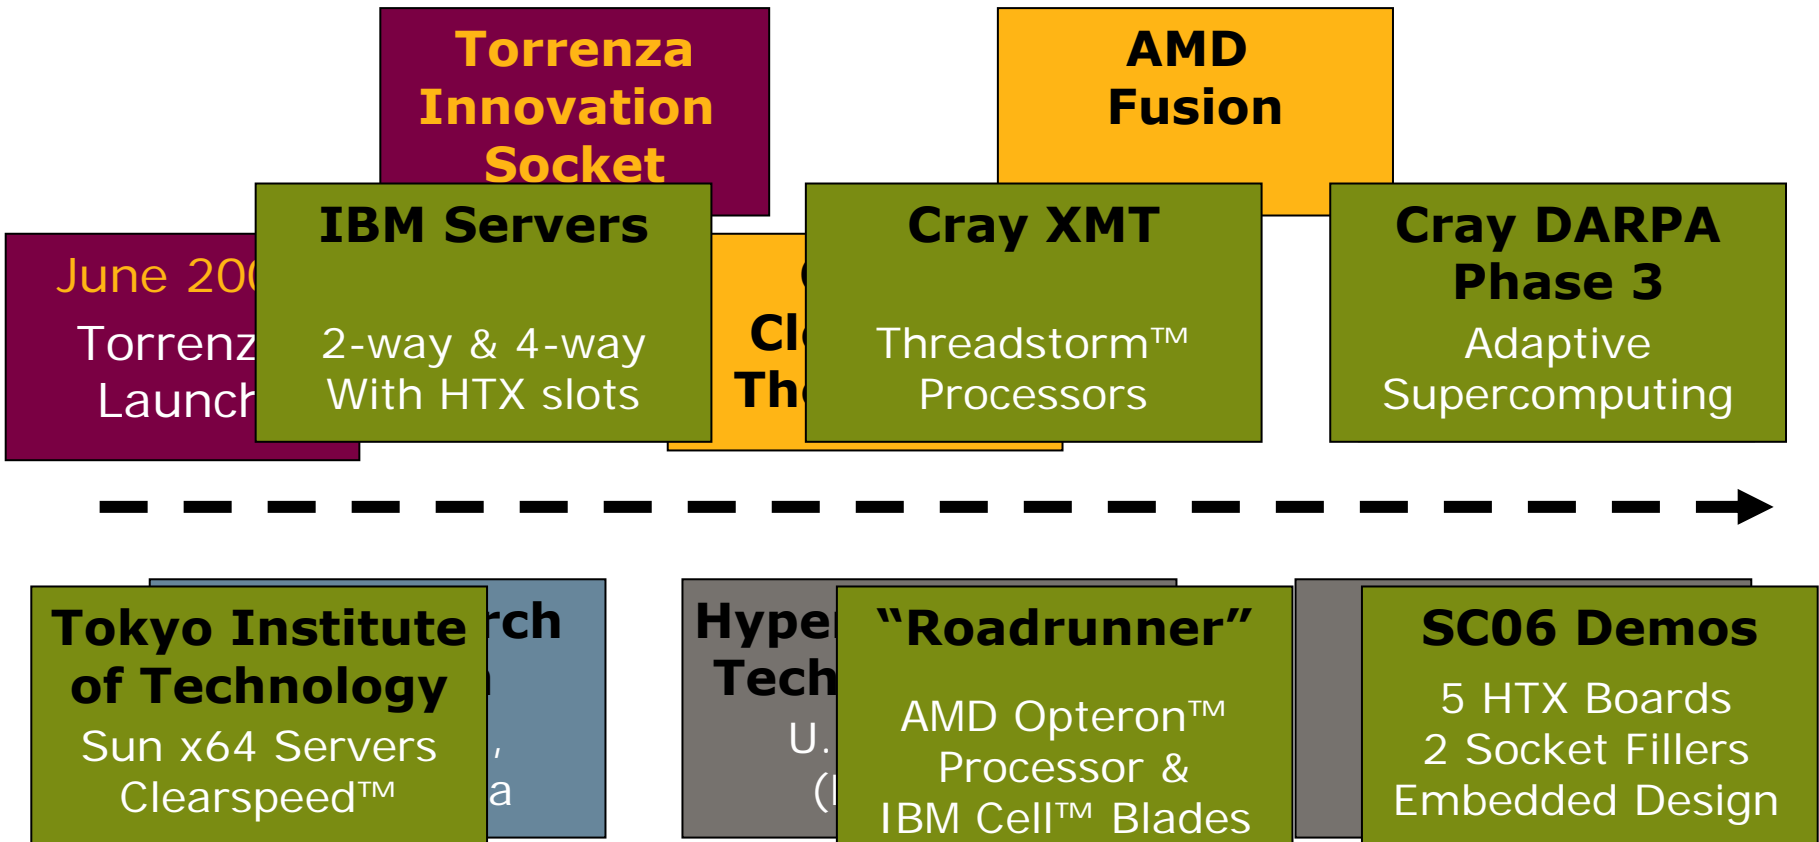

- An ecosystem of silicon development and support

AMD Technology

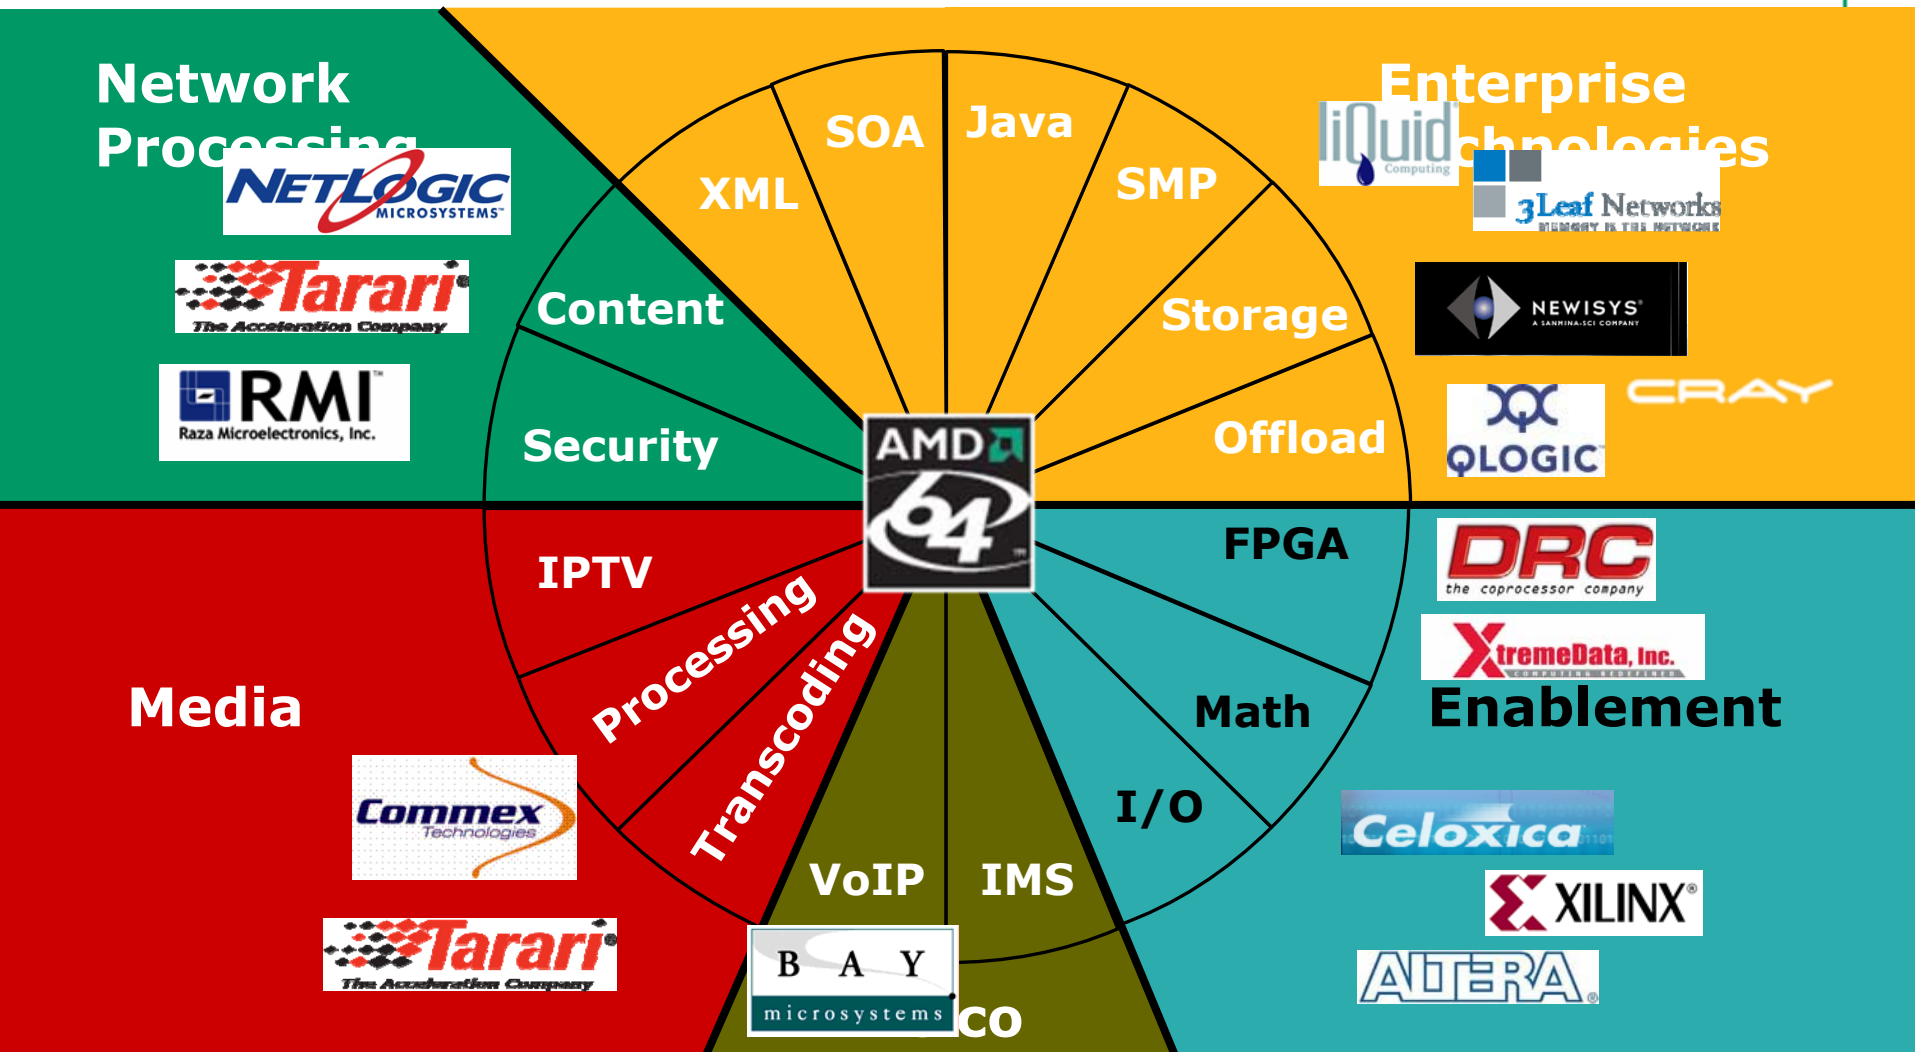
- Enhancing the platform & protocols for accelerators

Torrenza 2006 Timeline

Announcements



Advancing AMD's Torrenza Initiative



Reconfigurable Computing

Commercial Implementations



Center of Excellence

UNIVERSITÄT
MANNHEIM

HT, cHT & HTX
Reference Designs

Research Program



Platform & Programming Support



**This presentation is an exploration,
not a roadmap of a committed plan.**

On Acceleration Protocols

Offload Model: Asynchronous

(Survey of current accelerators)

- Majority of current base
 - E.g. GPU, Ethernet NIC, Appliances
- Have developed worked out latency hiding mechanisms
 - Large on-card memory enables buffering for large block transfers
 - Highly managed buffer & context techniques (e.g., HW multithreading)
- DMA is key function
 - Includes one or more system memory copies
 - Driver techniques to address issue in market e.g. Chimney
- Explicitly coded Communication Channels
 - Using available IO functions

Co-Processing Model: Synchronous

(Survey of accelerators)

Significant number of new entrants & evolving solutions

- Computational Offload
 - HPC, image processing, analytics, algorithm in chip
- Reliable delivery at wire speed
 - Messaging, Classification, Security, Virtualization
- **Applicability of accelerator is dependant upon problem size & intensity**
 - Acceleration must overcome start-up time (transfer latency & synchronization)
 - More synchronous communication between system thread and device processing
- **Compute efficiency requires moving data more often**
 - Accelerators are taking less time to process and need to increase concurrency
 - Larger, more complex context increases memory footprint
 - Overlapping transfers hides latency and increases accelerator throughput
- **More random access to system memory**
 - Increasing number of cores, threads & virtual machines within system
 - Many problems are not regularly ordered
 - Latency hiding is most effective on large, or highly predictable, work units

Goal for auto-generated offload requires support of smaller work-unit granularity

Need improved communications channels

Accelerator Communication/Synchronization Models

- Traditional approaches for I/O acceleration have typically been based on **Asynchronous** implementations
 - I/O device performance DMA for both inbound and outbound data transfers
 - Notification from I/O to host is by interrupt
 - Notification from host to I/O is typically via kernel-mode driver
- These approaches are appropriate for many I/O-related functions, but are too restrictive for computational acceleration
 - An efficient **Synchronous** mode of operation makes fine-grain offloading much easier and more effective
 - Host processor “baby-sits” accelerator, performing command/control and some computation
 - Polling is appropriate – the accelerator is doing most of the work, so overlap is a bonus, not a requirement
 - **Latency**, **Overhead**, and **Concurrency** are critical attributes of the interconnect

Sync Latency: Orders of Magnitude



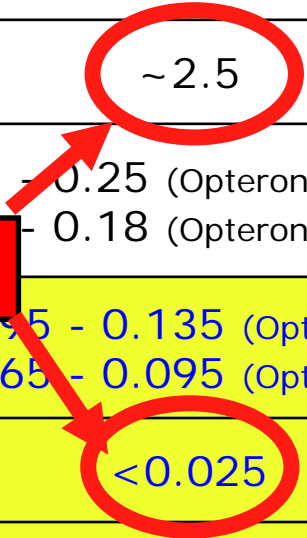
Type of Connection	Short Message Latency	
	(microseconds)	(CPU Cycles)
IP over WAN	~100,000	240,000,000
IP over LAN	~1,000	2,400,000
Unix Pipes in SMP (SUSE 9.1 on 2-socket Opteron)	>10	24,000
Unix Sockets in SMP (SUSE 9.1 on 2-socket Opteron)	~20 (TCP) ~15 (UDP)	48,000 36,000
User-mode Sockets in SMP (Scali)	~2.5	6000
Simple lock-free shared memory producer/consumer handoff	0.21 - 0.25 (Opteron 4 socket) 0.15 - 0.18 (Opteron 2 socket)	500 - 600 350 - 430
Cache-to-cache intervention	0.095 - 0.135 (Opteron 4s) 0.065 - 0.095 (Opteron 2s)	225 - 325 150 - 225
One-way HT probe latency (per hop)	<0.025	<60
On-Chip core-to-core latency	~0.0025	~6

Sync Latency: Orders of Magnitude



Type of Connection	Short Message Latency	
	(microseconds)	(CPU Cycles)
IP over WAN	~100,000	240,000,000
IP over LAN	~1,000	2,400,000
Unix Pipes in SMP (SUSE 9.1 on 2-socket Opteron)	>10	24,000
Unix Sockets in SMP (SUSE 9.1 on 2-socket Opteron)	~20 (TCP) ~15 (UDP)	48,000 36,000
User-mode Sockets in SMP (Scali)	~2.5	6000
Simple lock-free shared memory	0.21 - 0.25 (Opteron 4 socket) 0.15 - 0.18 (Opteron 2 socket)	500 - 600 350 - 430
Cache-to-cache intervention	0.095 - 0.135 (Opteron 4s) 0.065 - 0.095 (Opteron 2s)	225 - 325 150 - 225
One-way HT probe latency (per hop)	<0.025	<60
On-Chip core-to-core latency	~0.0025	~6

100x slower than hardware!



Why are Current Architectures so SLOW?

- Communication is not a feature of the architecture – it is a **side effect** of memory references
 - Hardware cannot optimize what it cannot identify
 - Caches reduce **capacity** misses, but no benefit for **communication**
- No “meta-information” is available to combine **communication** and **synchronization**
 - Multiple operations required: “data” plus “flag” with ordered updates
- Invalidate-based coherence protocols provide no way to “push” results to consumer
 - Data is always dirty in the producer’s cache
- Transparent Caching provides no way to exploit hardware dataflow mechanisms to “hand off” data from thread to thread
 - Polling/spinlocks rather than delaying responses until data is ready

Torrenza Technology Future

Attributes

- Optimize communication between system elements
 - CPU to CPU
 - Device to CPU
- Efficient use of platform bandwidth
 - Including system bandwidth, not just device link
- Lowest latency for complete operation
 - Work on cache line sized messages
 - Separate mechanisms for synchronization, small & large data movement
- Work within existing device enumerations
 - IOMMU & similar security functions
- Exposing future memory commands to device
 - Typically only available via cache-coherent HyperTransport™ technology

Two Routes for Data Movement

The increasing need for synchronization & small messages as a complement to bulk data transfer, not a replacement

- Small messages based on cache lines
- DMA for bulk data movement

The Open Philosophy

AMD's open architecture philosophy encourages innovation

The Technology Benefits

*Industry-standard platforms
Manufacturing efficiencies
Roadmap stability*



The Business Benefits

*Greater customer satisfaction
Better total cost of ownership
Greater flexibility*

Public Torrenza Participants

Altera AMI Bay Networks Cadence

Celoxica Commex Technologies Cray DRC Exegy Flextronics

GDA Technology IBM Linux Networx Liquid Computing NetLogic Newisys Phoenix

QLogic RMI Sun U. Mannheim Tarari 3 Leaf Networks Tyan Xilinx XtremeData

An open architecture is the key to enterprise success

Disclaimer and Trademark Attribution

The information contained herein is subject to change and may be rendered inaccurate for many reasons, including, but not limited to product and roadmap changes, component and motherboard version changes, new model and/or product releases, product differences between differing manufacturers, software changes, BIOS flashes, firmware upgrades, or the like. AMD assumes no obligation to update or otherwise correct or revise this information. However, AMD reserves the right to revise this information and to make changes from time to time to the content hereof without obligation of AMD to notify any person of such revisions or changes.

AMD MAKES NO REPRESENTATIONS OR WARRANTIES WITH RESPECT TO THE CONTENTS HEREOF AND ASSUMES NO RESPONSIBILITY FOR ANY INACCURACIES, ERRORS OR OMISSIONS THAT MAY APPEAR IN THIS INFORMATION.

AMD SPECIFICALLY DISCLAIMS ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR ANY PARTICULAR PURPOSE. IN NO EVENT WILL AMD BE LIABLE TO ANY PERSON FOR ANY DIRECT, INDIRECT, SPECIAL OR OTHER CONSEQUENTIAL DAMAGES ARISING FROM THE USE OF ANY INFORMATION CONTAINED HEREIN OR FOR THE PERFORMANCE OR OPERATION OF ANY PERSON, INCLUDING, WITHOUT LIMITATION, ANY LOST PROFITS, BUSINESS INTERRUPTION, DAMAGE TO OR DESTRUCTION OF PROPERTY, OR LOSS OF PROGRAMS OR OTHER DATA, EVEN IF AMD IS EXPRESSLY ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

© 2006 Advanced Micro Devices, Inc. All rights reserved. AMD, the AMD Arrow logo, AMD Opteron, and combinations thereof are trademarks of Advanced Micro Devices, Inc. HyperTransport is a licensed trademark of the HyperTransport Technology Consortium. Other names are for informational purposes only and may be trademarks of their respective owners.