



***HIGH-PERFORMANCE POWER-EFFICIENT X86-64 SERVER AND DESKTOP PROCESSORS
Using the core codenamed
“Bulldozer”***


**Sean White
19 August 2011**



THE DIE
Overview, Floorplan, and
Technology



THE DIE | Photograph

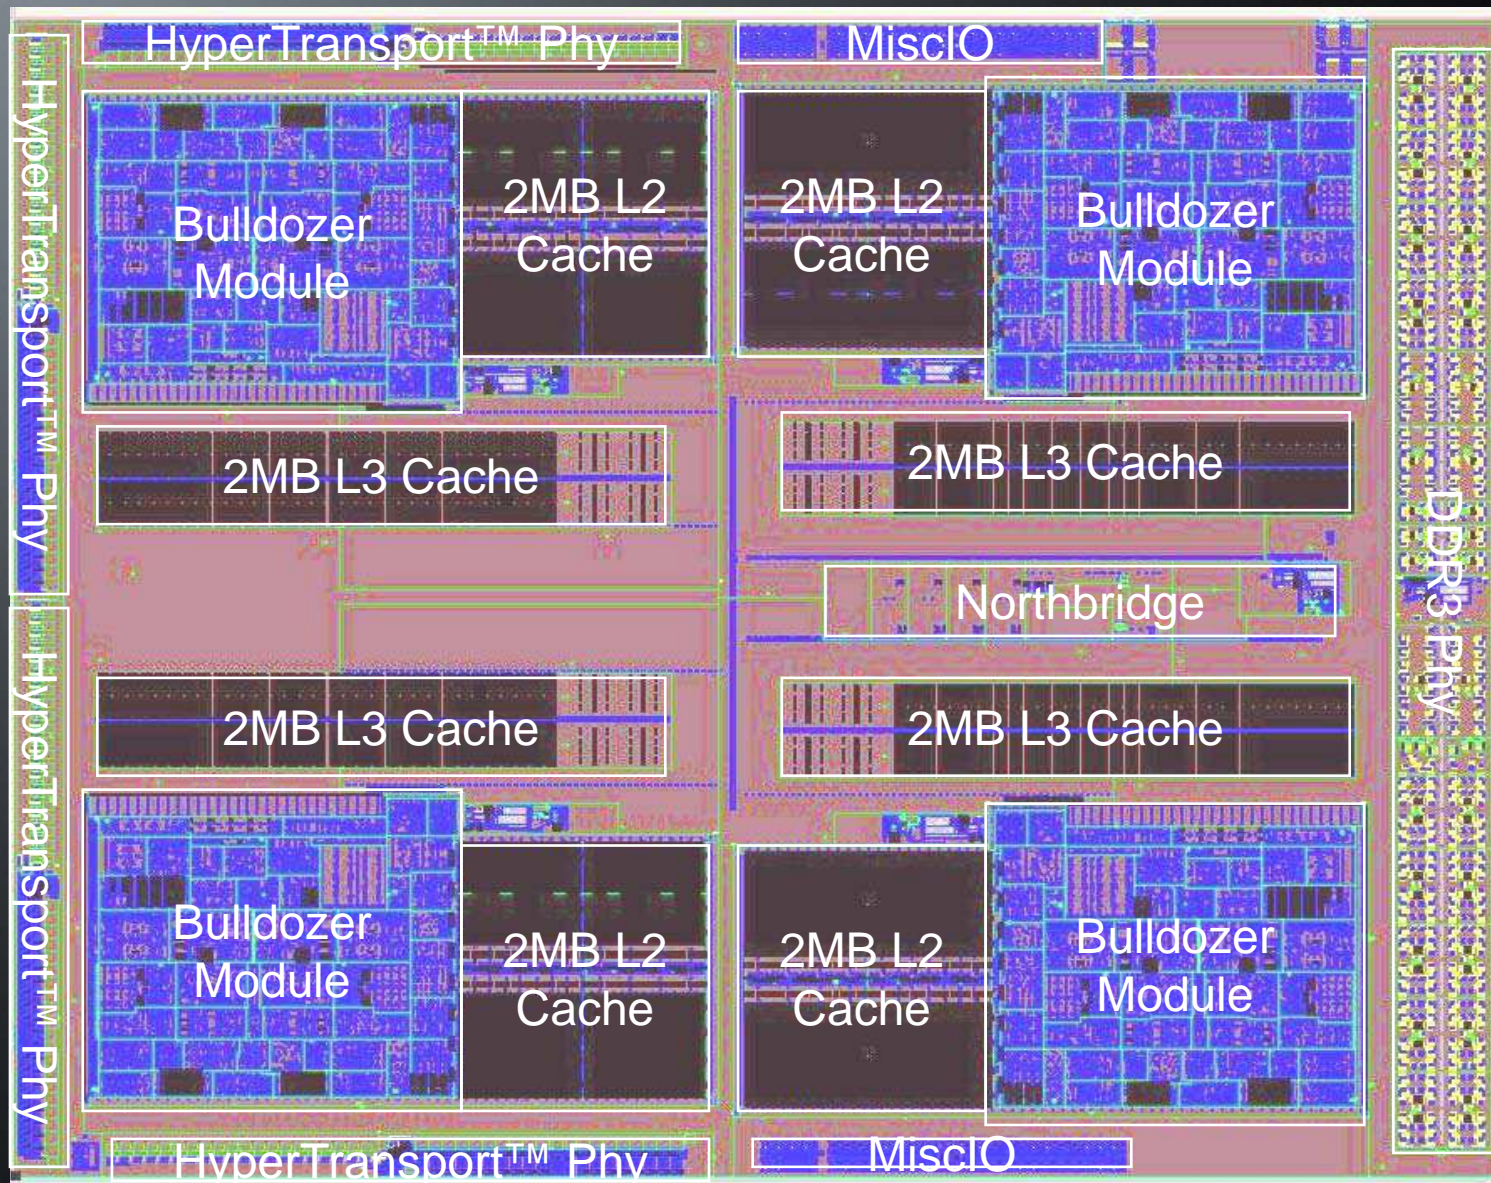


THE DIE / What's on it?

- Eight “Bulldozer” cores
 - High-performance, power-efficient AMD64 cores
 - First generation of a new execution core family from AMD (Family 15h)
 - Two cores in each Bulldozer module
 - 128 KB of Level1 Data Cache, 16 KB/core, 64-byte cacheline, 4-way associative, write-through
 - 256 KB of Level1 Instruction Cache, 64 KB/Bulldozer module, 64-byte cacheline, 2-way associative
 - 8 MB of Level2 Cache, 2 MB/Bulldozer module, 64-byte cacheline, 16-way associative
- Integrated Northbridge which controls:
 - 8 MB of Level3 Cache, 64-byte cacheline, 16-way associative, MOESI
 - Two 72-bit wide DDR3 memory channels
 - Four 16-bit receive/16-bit transmit HyperTransport™ links



THE DIE | Floorplan (315 mm²)



THE DIE / Process Technology


- 32-nm Silicon-On-Insulator (SOI) Hi-K Metal Gate (HKMG) process from GlobalFoundries
- 11-metal-layer-stack
- Low-k dielectric
- Dual strain liners and eSiGe to improve performance.
- Multiple VT (HVT, RVT, LVT) and long-channel transistors.

| Layer | Type | Pitch |
|------------------|------|--------|
| CPP ¹ | - | 130 nm |
| M01 | 1x | 104 nm |
| M02 | 1x | 104 nm |
| M03 | 1x | 104 nm |
| M04 | 1.3x | 130 nm |
| M05 | 1.3x | 130 nm |
| M06 | 2x | 208 nm |
| M07 | 2x | 208 nm |
| M08 | 2x | 208 nm |
| M09 | 4x | 416 nm |
| M10 | 16x | 1.6 μm |
| M11 | 16x | 1.6 μm |

¹ Contacted poly pitch



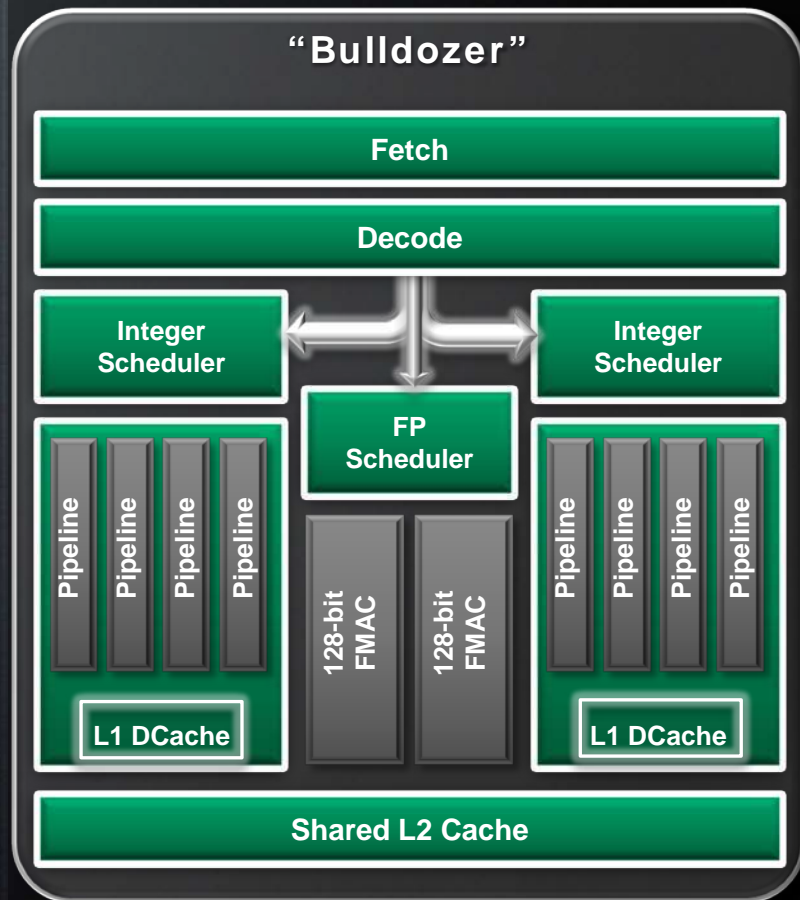
THE BULLDOZER CORE
A few highlights



THE BULLDOZER CORE | What's different about this core?

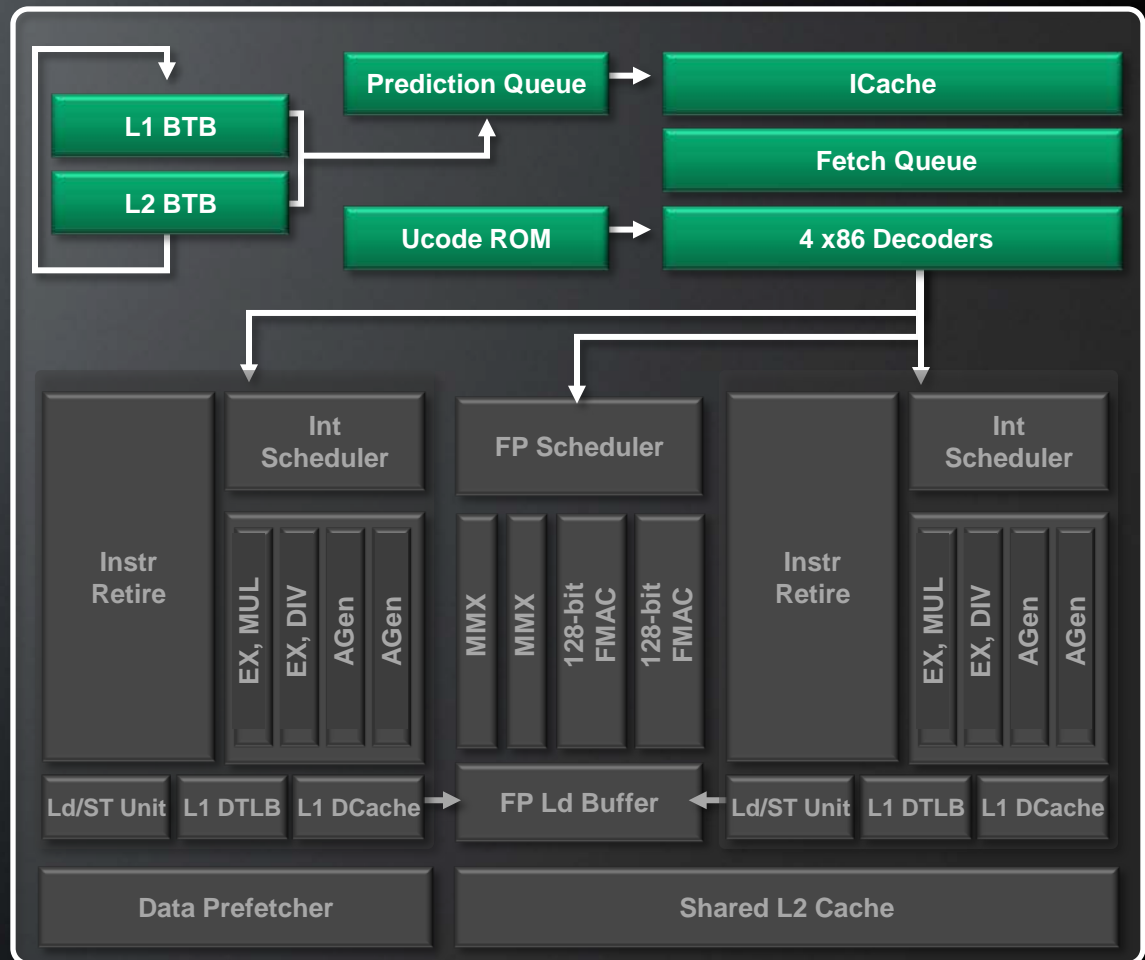
A Bulldozer module contains 2 cores

- Functions with high utilization that can't be shared without significant compromises exist for each core
- Ex: Integer pipelines, Level1 data caches
- Other functions are shared between the cores
- Ex: Floating point pipelines, Level2 cache
- This allows two cores to each use a larger, higher-performance function (ex: floating point unit) as they need it for less total die area than having separate, smaller functions for each core



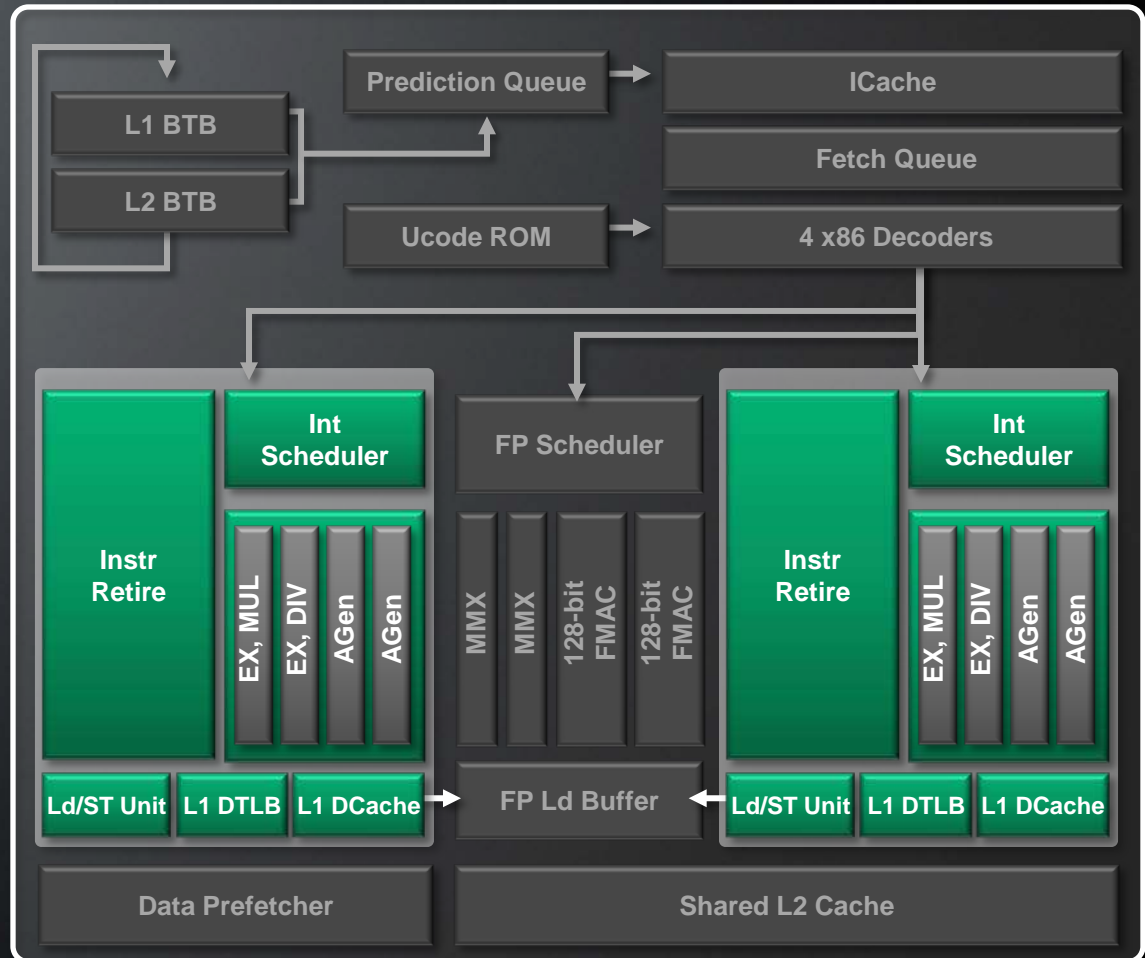
THE BULLDOZER CORE / Microarch: Shared Frontend

- Decoupled predict and fetch pipelines
- Prediction-directed instruction prefetch
- Instruction cache: 64K Byte, 2-way
- 32-Byte fetch
- Instruction TLBs:
 - Level1: 72 entries, mixed page sizes
 - Level2: 512 entries, 4-way, 4K pages
- Branch fusion



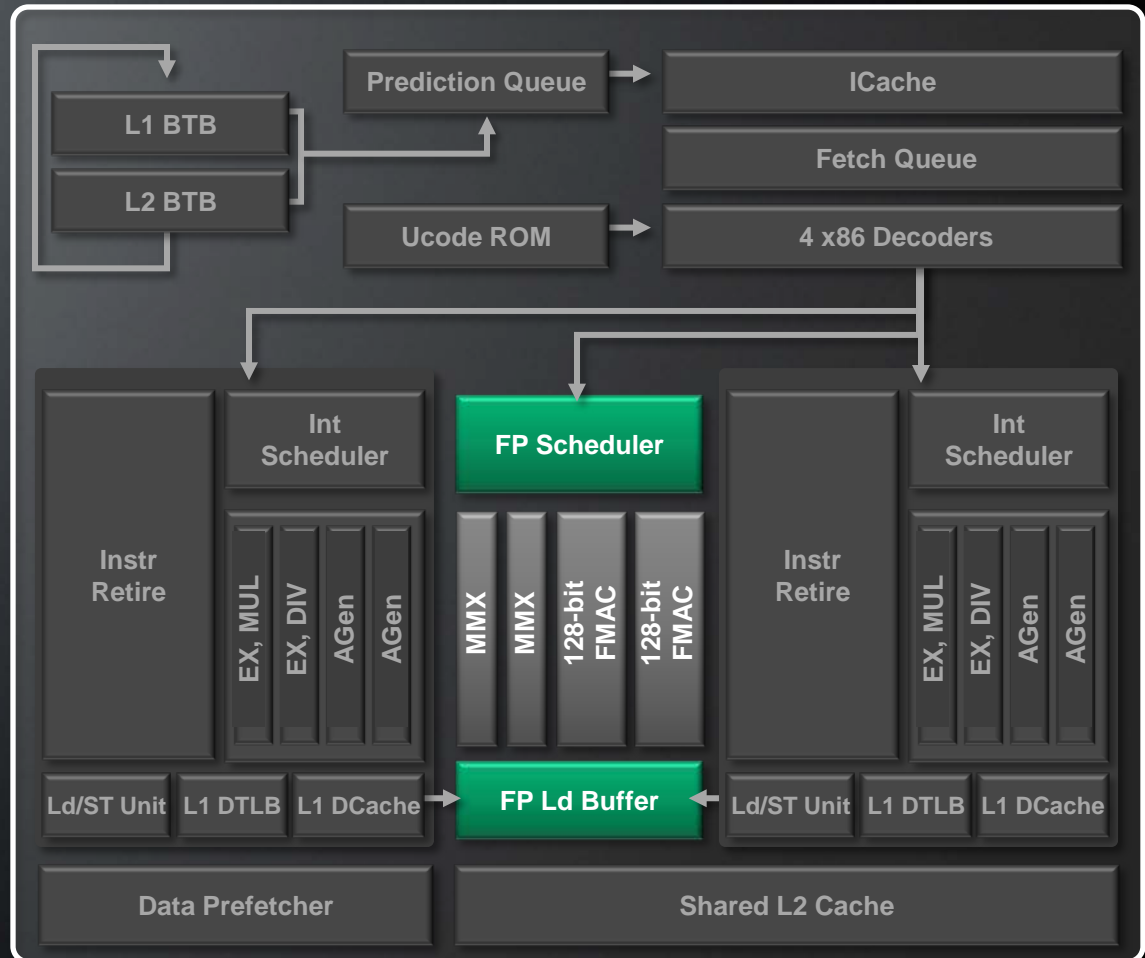
THE BULLDOZER CORE / Microarch: Separate Integer Units

- Thread retire logic
- Physical Register File (PRF)-based register renaming
- Unified scheduler per core
- Way-predicted 16K Byte Level1 Data cache
- Data TLB: 32-entry fully associative
- Fully out-of-order load/store
 - 2 128-bit loads/cycle
 - 1 128-bit store/cycle
 - 40-entry Load queue
 - 24-entry Store queue



THE BULLDOZER CORE | Microarch: Shared FPU

- Co-processor organization
- Reports completion back to parent core
- Dual 128-bit Floating Point Multiply/Accumulate (FMAC) pipelines
- Dual 128-bit Packed Integer pipelines
- PRF-based register renaming
- A single floating point scheduler for both cores



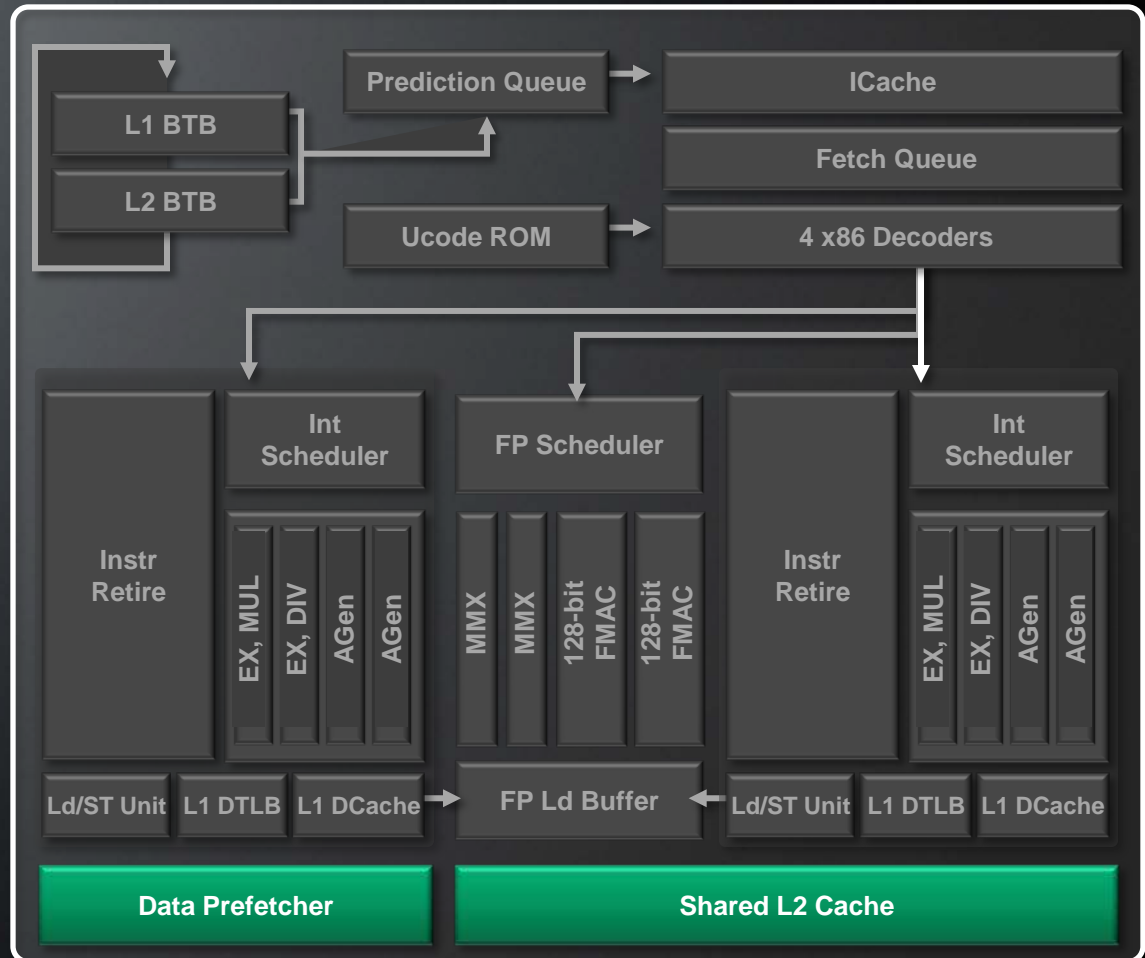
THE BULLDOZER CORE / New Floating Point Instructions

| New Instructions | Applications/Use Cases |
|---|---|
| SSSE3, SSE4.1, SSE4.2 (AMD and Intel) | <ul style="list-style-type: none">• Video encoding and transcoding• Biometrics algorithms• Text-intensive applications |
| AESNI PCLMULQDQ (AMD and Intel) | <ul style="list-style-type: none">• Application using AES encryption• Secure network transactions• Disk encryption (MSFT BitLocker)• Database encryption (Oracle)• Cloud security |
| AVX (AMD and Intel) | Floating point intensive applications: <ul style="list-style-type: none">• Signal processing / Seismic• Multimedia• Scientific simulations• Financial analytics• 3D modeling |
| FMA4 (AMD Unique) | HPC applications |
| XOP (AMD Unique) | <ul style="list-style-type: none">• Numeric applications• Multimedia applications• Algorithms used for audio/radio |




THE BULLDOZER CORE / Microarch: Shared Level 2 Cache

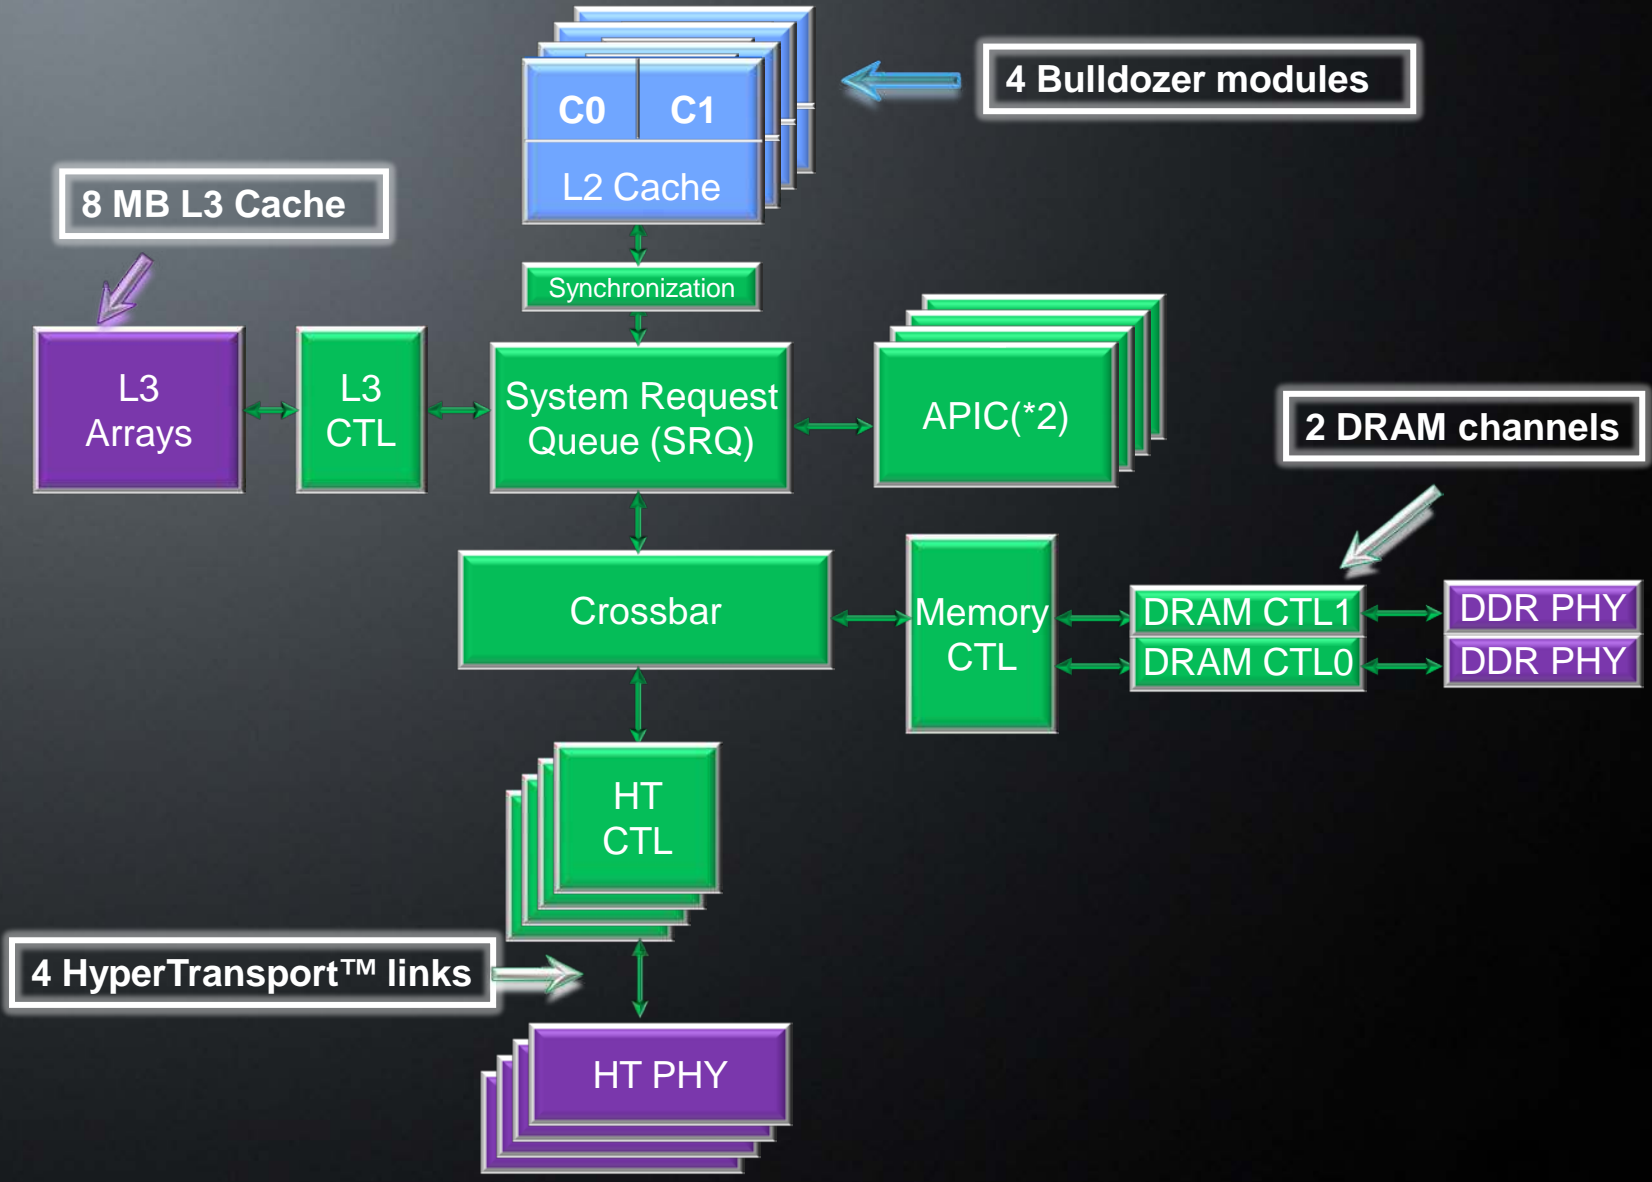
- 16-way unified L2 cache
- L2 TLB and page walker
 - 1024-entry, 8-way
 - Services both Instruction and Data requests
- Multiple data prefetchers
- 23 outstanding L2 cache misses for memory system concurrency



THE NORTHBRIDGE
A few of its highlights



THE NORTHBRIDGE | Simplified Block Diagram



THE NORTHBRIDGE | SRQ, Crossbar, Memory Controller

● System Request Queue and Crossbar

- The northbridge SRQ and crossbar form the traffic hub of the design, routing from any requestor to the appropriate destination.
- Examples: A core request to the L3 cache; a request from another die or chip from a HyperTransport™ link to the memory controller, etc.

● Memory Controller

- Each memory controller enforces cache coherency and proper order of operations for the memory space it “owns”, distributing this responsibility across multi-die or multi-chip systems.
- This allows multi-die or multi-chip systems to scale rather than bottleneck at a single coherency/ordering choke point in the system.
- Obviously, the memory controller also sends the read/write requests to the DRAM controllers.

THE NORTHBRIDGE | DRAM Controllers, L3 Cache, Probe Filter

● DDR3 DRAM Controllers

- Two per die
- Supports Unbuffered (UDIMM), Registered (RDIMM) and Load-Reduced (LRDIMM) memory, 1.50V, 1.35V and 1.25V depending on the processor package, with frequencies up to DDR3-1866
- New power-saving features: Low-voltage DDR3, fast self-refresh entry with internal clocks gated, aggressive precharge power down, tristate addr/cmd/bank w/chip-select deassertion, throttle activity for thermal or power reduction, put RX circuits in standby when no reads active, etc.


● L3 Cache and Probe Filter

- Up to 8 MB of Level3 cache, normally shared, can be partitioned
- L3 data ECC protected (single-bit correct, double-bit detect)
- Probe Filter: the northbridge can filter coherency traffic to improve system bandwidth
- Space in the L3 array is used to hold Probe Filter data when it's enabled

THE NORTHBRIDGE | Four 16-bit HyperTransport™ Links

- High-speed unidirectional signaling, 16-bits receive/16-bits transmit
 - A 16-bit link can be ungangged into two 8-bit links
- Up to 6.4 Gigatransfers/sec/bit (16-bit link = 12.8 GByte/s RX, 12.8 TX)
 - Maximum link frequency depends on package, chipset capability and printed circuit board signal integrity
- Number of links vary with package
 - AM3+ has a single link to the chipset
 - C32 has three links for chipset and/or dual-processor coherent interconnect
 - G34 has four external links for chipset and/or multi-processor coherent interconnect (plus internal links between the dice in the multi-chip module)
- In retry mode, bit errors on the link are detected on CRC error and packets are reissued for highly reliable operation.
- Power consumption is minimized based on configuration (ex: link width/frequency) and activity.

POWER MANAGEMENT
A few key capabilities

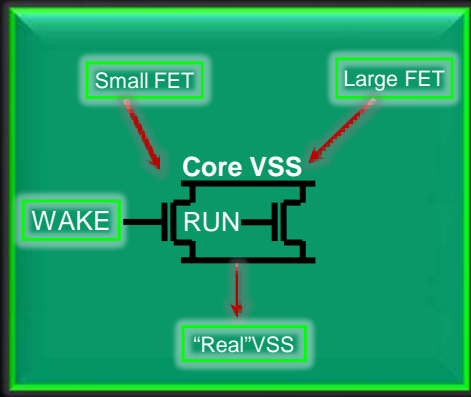
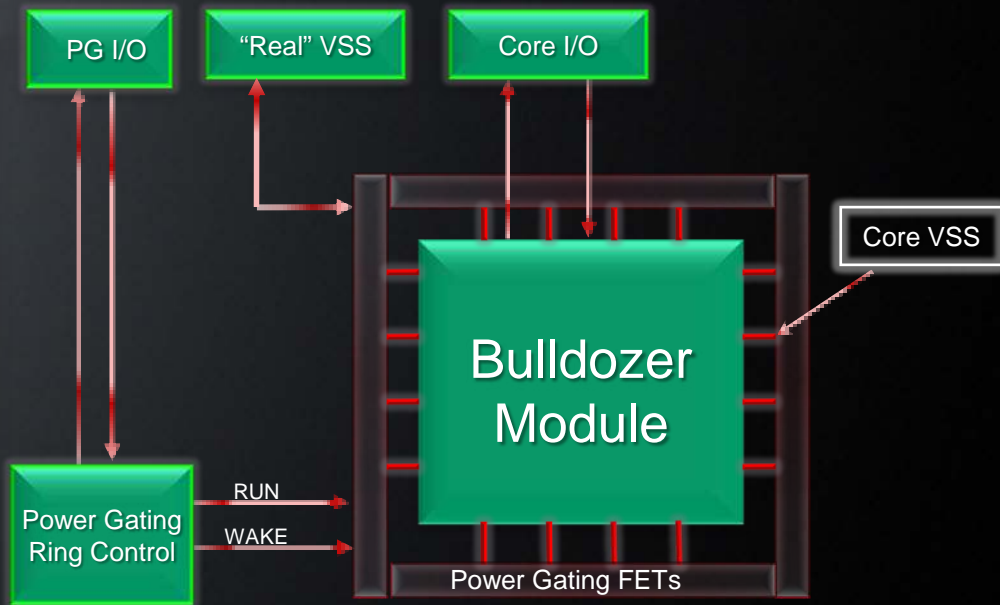


POWER MANAGEMENT | Key Features and Functions

- **The Bulldozer core is architected to be power-efficient**
 - Minimize silicon area by sharing functionality between two cores
- **All blocks and circuits have been designed to minimize power (not just in the Bulldozer core)**
 - Extensive flip-flop clock-gating throughout design
 - Circuits power-gated dynamically
- **Numerous power-saving features under firmware/software control**
 - Core C6 State (CC6)
 - Core P-states/AMD Turbo CORE
 - Application Power Management (APM)
 - DRAM power management
 - Message Triggered C1E

POWER MANAGEMENT / Core C6 State (CC6)


- Core C6: if a core isn't active, remove power
- Implemented in this physical design by a power gating ring that isolates the Core VSS for each Bulldozer module from the "Real" VSS
- CC6 entry: when both Bulldozer cores in the module are idle, flush caches and dump register state to CC6 save space, then gate Core VSS
- CC6 exit: ungate Core VSS, reload CC6 saved state, resume execution (ex: service interrupts, etc.)



POWER MANAGEMENT / P-states, AMD Turbo CORE

Core P-states specify multiple frequency and voltage points of operation

- Higher frequency P-states deliver greater performance but require higher voltage and thus more power
- The hardware and operating system vary which P-state a core is in to deliver performance as needed, but use lower frequency P-states to save power whenever possible




AMD Turbo CORE: when the processor is below its power/thermal limits the frequency and voltage can be boosted above the normal maximum and stay there until it gets back to the power/thermal limits




TURBO CORE | Base, All Core Boost and Max Boost


Base frequency with TDP headroom



All core boost activated



Max turbo activated



**All
Core
Boost**


When there is TDP headroom in a given workload, AMD Turbo CORE technology is automatically activated and can increase clock speeds across **all cores**.

**Max
Turbo
Boost**

When a lightly threaded workload sends half the Bulldozer modules into C6 sleep state but also requests max performance, AMD Turbo CORE technology can increase clock speeds for the active Bulldozer modules.



THE 1-SOCKET DESKTOP AM3+ PROCESSOR CODENAMED “ZAMBEZI”



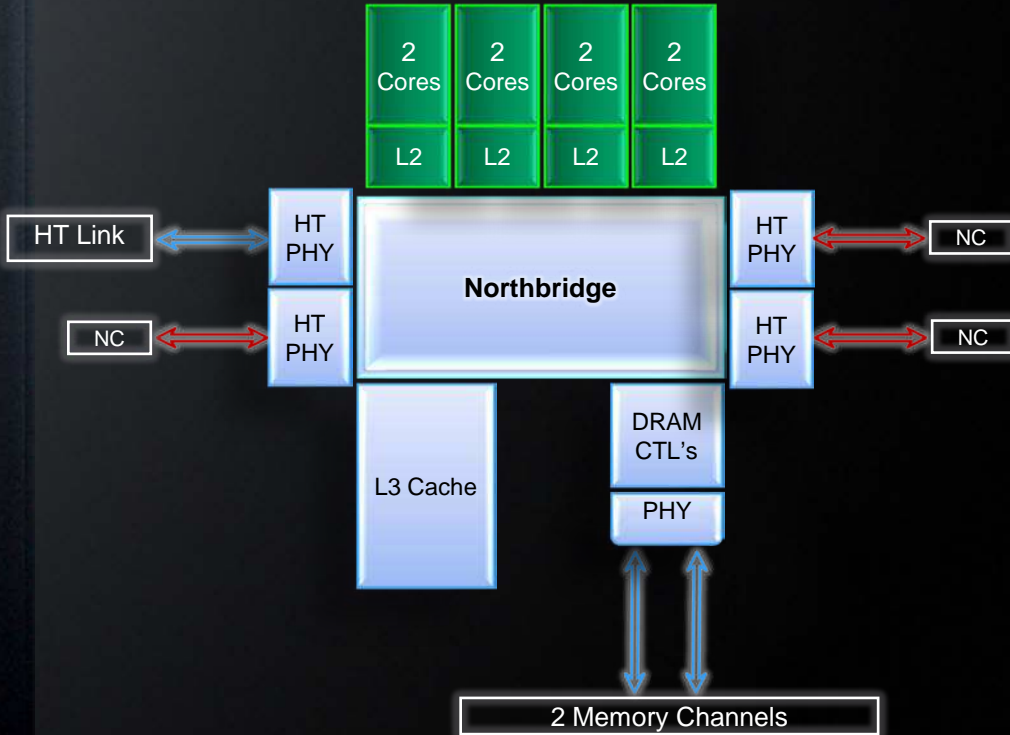
- 940-pin lidded micro PGA package
- 1.27mm pin pitch
- 31 row x 31 column pin array
- C4 die attach




ZAMBEZI | Performance Desktop Processor

“Zambezi”, for the AM3+ Platform

- AM3+ socket infrastructure adds:
 - Support for low-voltage DRAM
 - Increased ILDT current for higher frequency HyperTransport™ link (2.0A maximum per Gen3 link)
 - Increase in IDDR current (4.0A maximum)
- Older AM3 processors plug-in compatible with AM3+ motherboards
- 2 memory channels, unbuffered DIMMs, up to DDR3-1866
- 1 HyperTransport™ link, up to 5.2 GT/s
- For use with AMD 9-Series chipsets
 - AMD 990FX
 - AMD 990X
 - AMD 970
 - AMD SB950



***THE 1 OR 2-SOCKET
SERVER C32 PROCESSOR
CODENAMED “VALENCIA”***



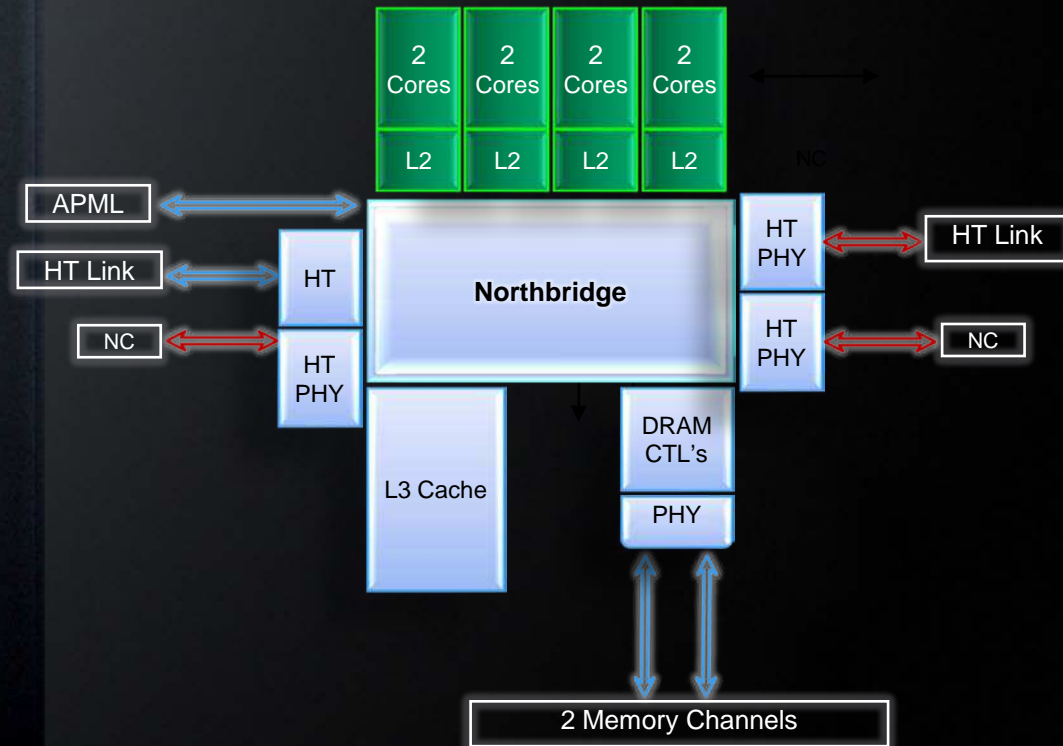
- 1207-land lidded micro LGA package
- 1.10mm land pitch
- 35 row x 35 column land array
- C4 die attach



VALENCIA / 1-2 Socket Server Processor

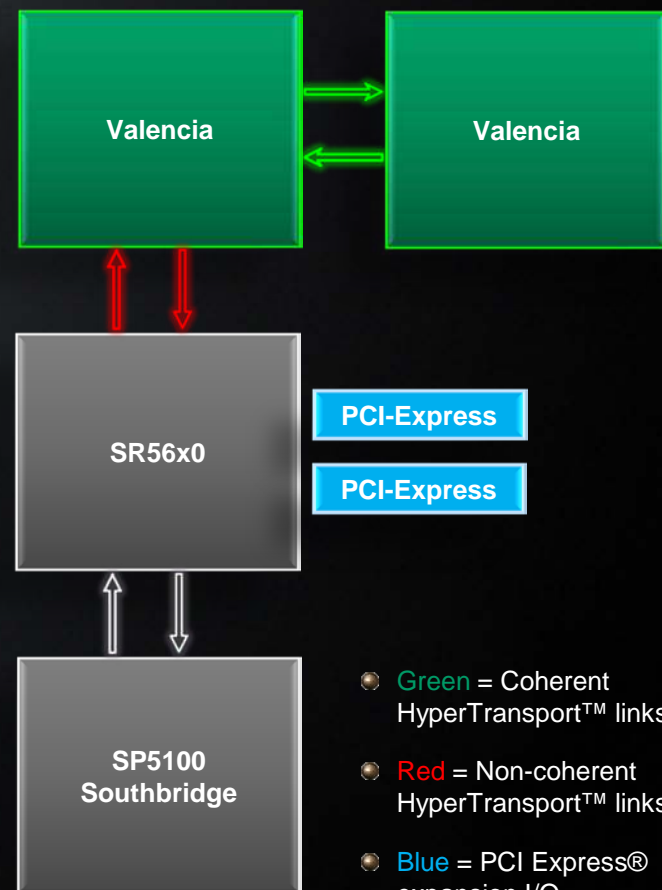
“Valencia”, for the C32 Platform

- Valencia is compatible with existing C32 motherboards (AMD Opteron™ 4000 series processor-based platform) with appropriate BIOS update
- 2 memory channels, UDIMM, RDIMM or LRDIMM, up to DDR3-1600
- 3 HyperTransport™ links, up to 6.4 GT/s
 - Most designs use only 2 links to achieve lower Thermal Design Power (TDP)
- Advanced Platform Management Link (APML)
- 1 or 2 socket systems
- For use with AMD server chipsets
 - AMD SR5690
 - AMD SR5670
 - AMD SR5650
 - AMD SP5100



VALENCIA / 2-socket system example

- System design optimized to provide maximum performance for minimum cost and power for 1-2 socket servers
- Up to 16 cores in a two socket system
- Two DDR3 Memory channels per socket
- Northbridge expansion I/O
 - AMD SR5690: 42 PCI Express® lanes
 - AMD SR5670: 30 PCI Express® lanes
 - AMD SR5650: 22 PCI Express® lanes
- SP5100 Southbridge: SATA, PCI, USB



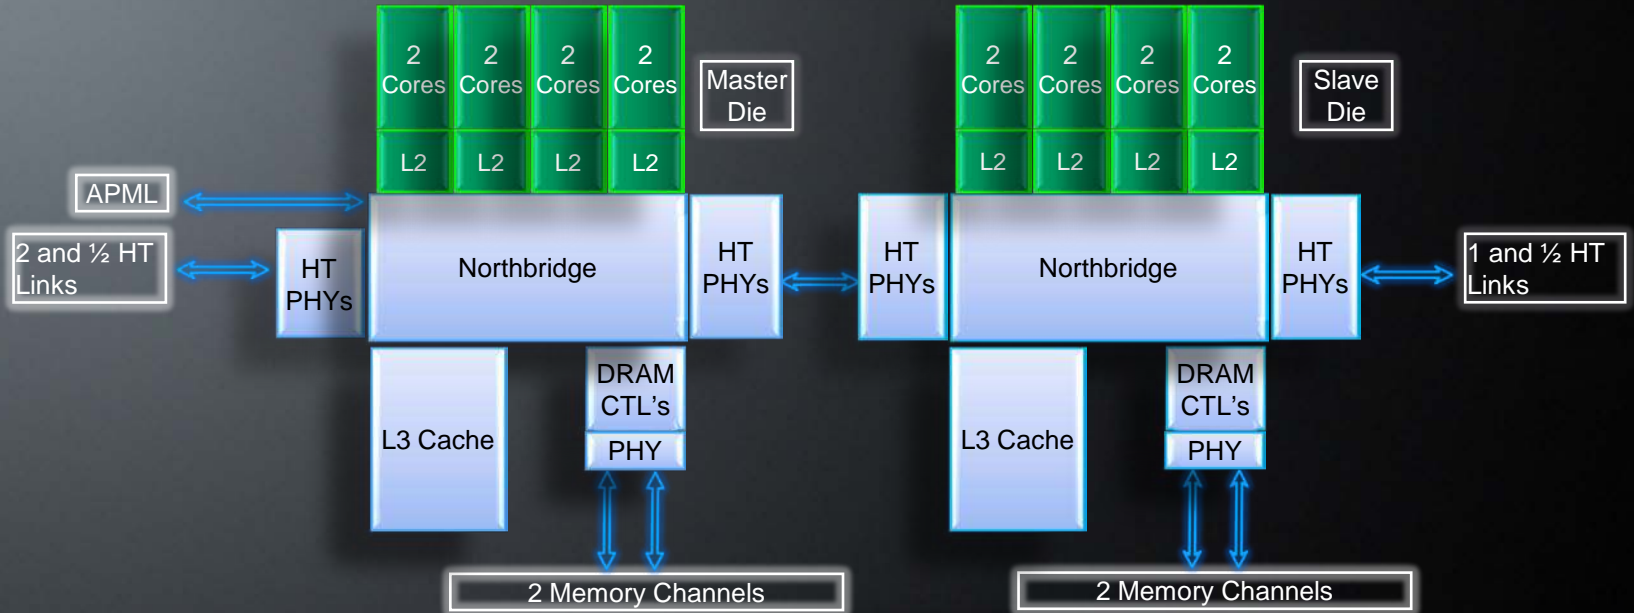
**THE 1 TO 4-SOCKET
SERVER G34 PROCESSOR
CODENAMED
“INTERLAGOS”**



- 1944-land lidded LGA package
- 1.00mm land pitch
- 57 row x 40 column land array
- C4 die attach



INTERLAGOS / 1-4 Socket Server Processor



“Interlagos”, for the G34 Platform

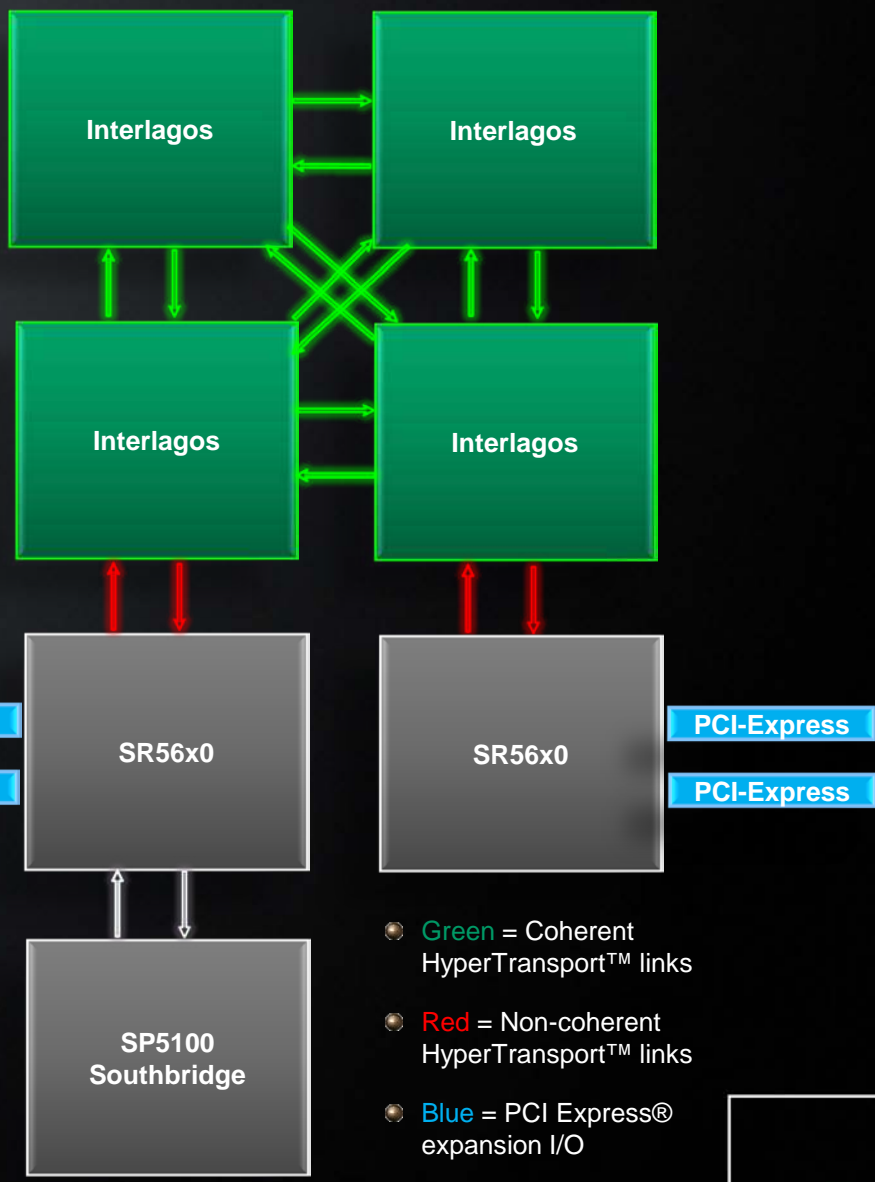
- Two die in a multichip module for 1 to 4 socket systems
- Compatible with existing G34 motherboards (AMD Opteron™ 6000 series processor-based platform) with appropriate BIOS update
- OS views it as one multi-core processor with up to 16 cores

- Up to 16 MB of combined Level3 cache
- 4 memory channels, UDIMM, RDIMM or LRDIMM, up to DDR3-1600
- 4 external HyperTransport™ links, up to 6.4 GT/s (plus internal die-to-die links)
- Advanced Platform Management Link (APML)
- For use with AMD server chipsets
 - AMD SR56x0, AMD SP5100



INTERLAGOS / 4-socket system example

- System design optimized to provide maximum performance for 1-4 socket servers
- Up to 64 cores in a four socket system to support highly multi-threaded workloads
 - Database, Web, Virtualization, Cloud Computing, High-Performance Computing, etc.
- 16 DDR3 channels w/4 sockets
 - To support demanding memory-bandwidth-intensive workloads
- High-performance PCI Express® links via SR56x0 chipset to support demanding I/O-intensive workloads (high-speed network, storage, etc.)



Trademark Attribution

AMD, the AMD Arrow logo, AMD Opteron 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.

PCI Express is a registered trademark of PCI-SIG.

©2011 Advanced Micro Devices, Inc. All rights reserved.

