





2nd Generation Intel® Core™ Processor Family: Intel® Core™ i7, i5 and i3



Oded Lempel, Sandy Bridge Design Management

Agenda

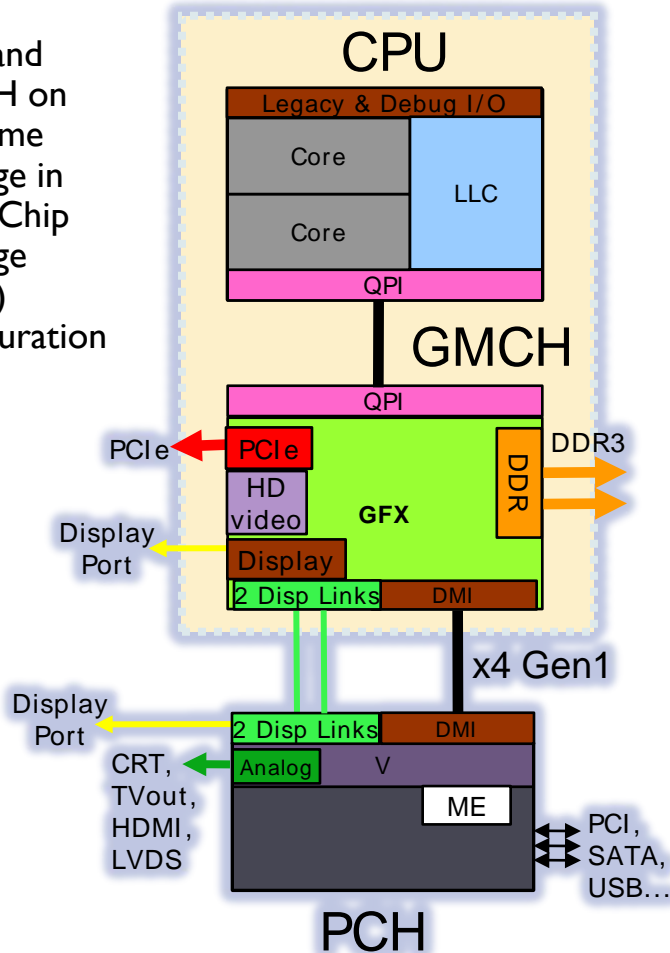
- ▶ CPU Overview
- ▶ System Agent, Ring Architecture and System Integration
- ▶ Core Enhancements and IA Extensions
- ▶ Integration Challenges



System On Chip Integration

2010 Platform

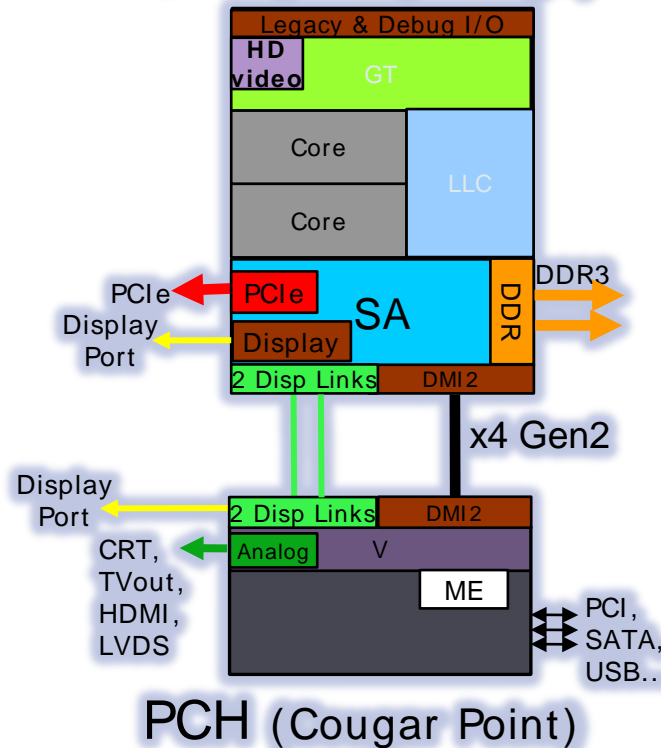
- CPU and GMCH on the same package in Multi-Chip Package (MCP) configuration



2011 Platform

CPU (Sandy Bridge)

- CPU and GMCH integrated into a single chip
- PCH remains as separate die outside of CPU package



Intel® Core™ Microarchitecture

Greater Performance/Lower Power Consumption

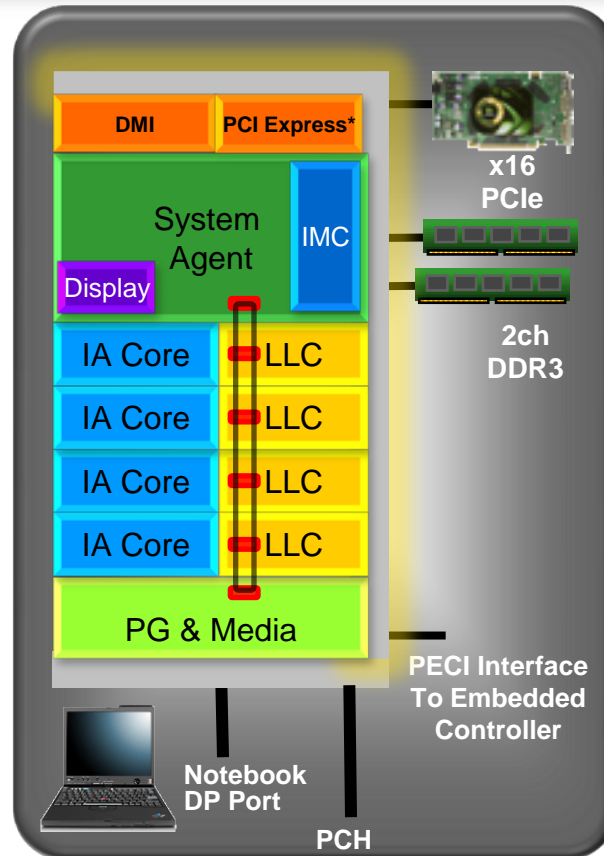
Intel® Hyper-Threading Technology
4 Cores / 8 Threads
2 Cores / 4 Threads

Intel® Advanced Vector Extension (Intel® AVX)

High Bandwidth Last Level Cache (LLC)

High Bandwidth / Low Latency Modular Interconnect

Next Generation Processor Graphics (PG) and Media



Embedded Display Port (eDP)

Integrated Memory Controller (iMC) 2ch DDR3

Discrete Graphics Support
1x16 or 2x8


Next Generation Intel® Turbo Boost Technology

Substantial Performance Improvement

Integrates CPU, Graphics Core, Memory Controller and PCI Express on 32nm Process



2nd Gen Intel® Core™ Microarchitecture




System Agent (SA), Ring Architecture and System Integration




Challenges of Multi-Core Integration

- ▶ **Scalability**
 - ▶ Performance should scale with number of Cores
- ▶ **Robustness**
 - ▶ Bandwidth should not limit large scale integration
- ▶ **Modularity**
 - ▶ Core is used in Server products
 - ▶ Interface to Ring needs to be robust
 - ▶ Be able to quickly generate derivatives
 - ▶ Variable number of cores
 - ▶ Varying Processor Graphics sizes
- ▶ **Power Performance Controllability**
 - ▶ Tight power management control of all components, providing better granularity and deeper idle/sleep states



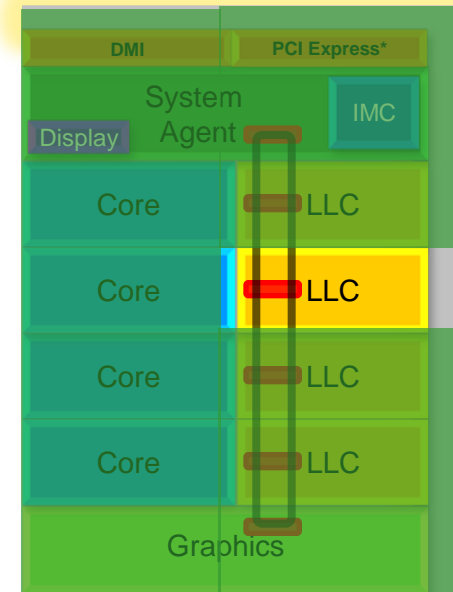
Cache Architecture

- ▶ High BW Last Level Cache, **shared** among Cores and Graphics
 - ▶ Inclusive Multi-Bank LLC, 64B Cache Line, 16-way associative
 - ▶ IA cores and LLC run at the same variable frequency
 - ▶ Cache Size and BW scale with the number of IA cores
- ▶ Significant performance boost, saves memory bandwidth and power
- ▶ Each Cache “slice” holds 1/N of LLC (1.5 or 2 MB)
 - ▶ Addresses are hashed among slices



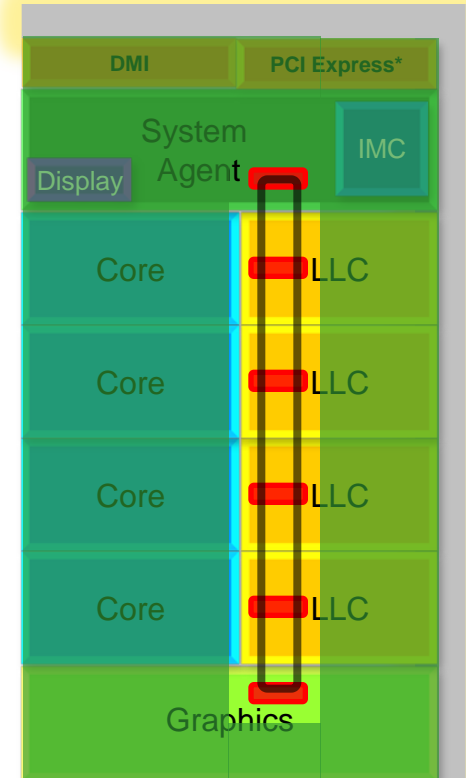
Cache Box

- ▶ **Interface block within LLC slice**
 - ▶ Between **Core/Graphics/Media** and the **Ring**
 - ▶ Implements the **ring logic, arbitration, cache controller**
 - ▶ Communicates with **System Agent** for LLC misses, external snoops, non-cacheable accesses
- ▶ **Full cache pipeline** in each cache box
 - ▶ Physical Addresses are **hashed at the source** to **prevent hot spots** and increase bandwidth
 - ▶ Maintains **coherency and ordering** for the addresses **that are mapped to it**
 - ▶ LLC is fully inclusive with “**Core Valid Bits (CVB)**” – eliminates unnecessary snoops to cores
 - ▶ Per-core CVB indicates whether the core needs to be snooped for a given cacheline
- ▶ Runs at **core voltage/frequency**, scales with cores



Scalable Ring On-Die Interconnect

- ▶ **Ring-based** interconnect between Cores, Graphics, Last Level Cache (LLC) and System Agent domain
- ▶ Composed of **4 physical rings**
 - ▶ 32 Byte **Data** ring, **Request** ring, **Acknowledge** ring and **Snoop** ring
 - ▶ Fully pipelined at **core frequency/voltage**: bandwidth, latency and power scale with cores
- ▶ Massive ring **wire routing** runs over the LLC with no area impact
- ▶ Access on ring always picks the **shortest path** – minimize latency
- ▶ **Distributed arbitration**, sophisticated ring protocol to handle coherency, ordering, and core interface
 - ▶ Arbitration happens in each of the Ring Stops which span the multiple slices
 - ▶ Example in backup slides
- ▶ **Scalable to servers** with large number of processors



Coherency and Ordering


- ▶ Ring does not maintain transaction ordering
 - ▶ Transactions flow in parallel and out-of-order
- ▶ Ordering requirement maintained by agents sending the requests
 - ▶ IA Core, Processor Graphics, System Agent
- ▶ Cache Coherency protocol is based on QPI
 - ▶ MESI based Source Snooping

Robustness – Distributed coherency & ordering
Scalability – Bandwidth, Latency & Performance
Modularity – to number of Cores



Lean and Mean System Agent (SA)

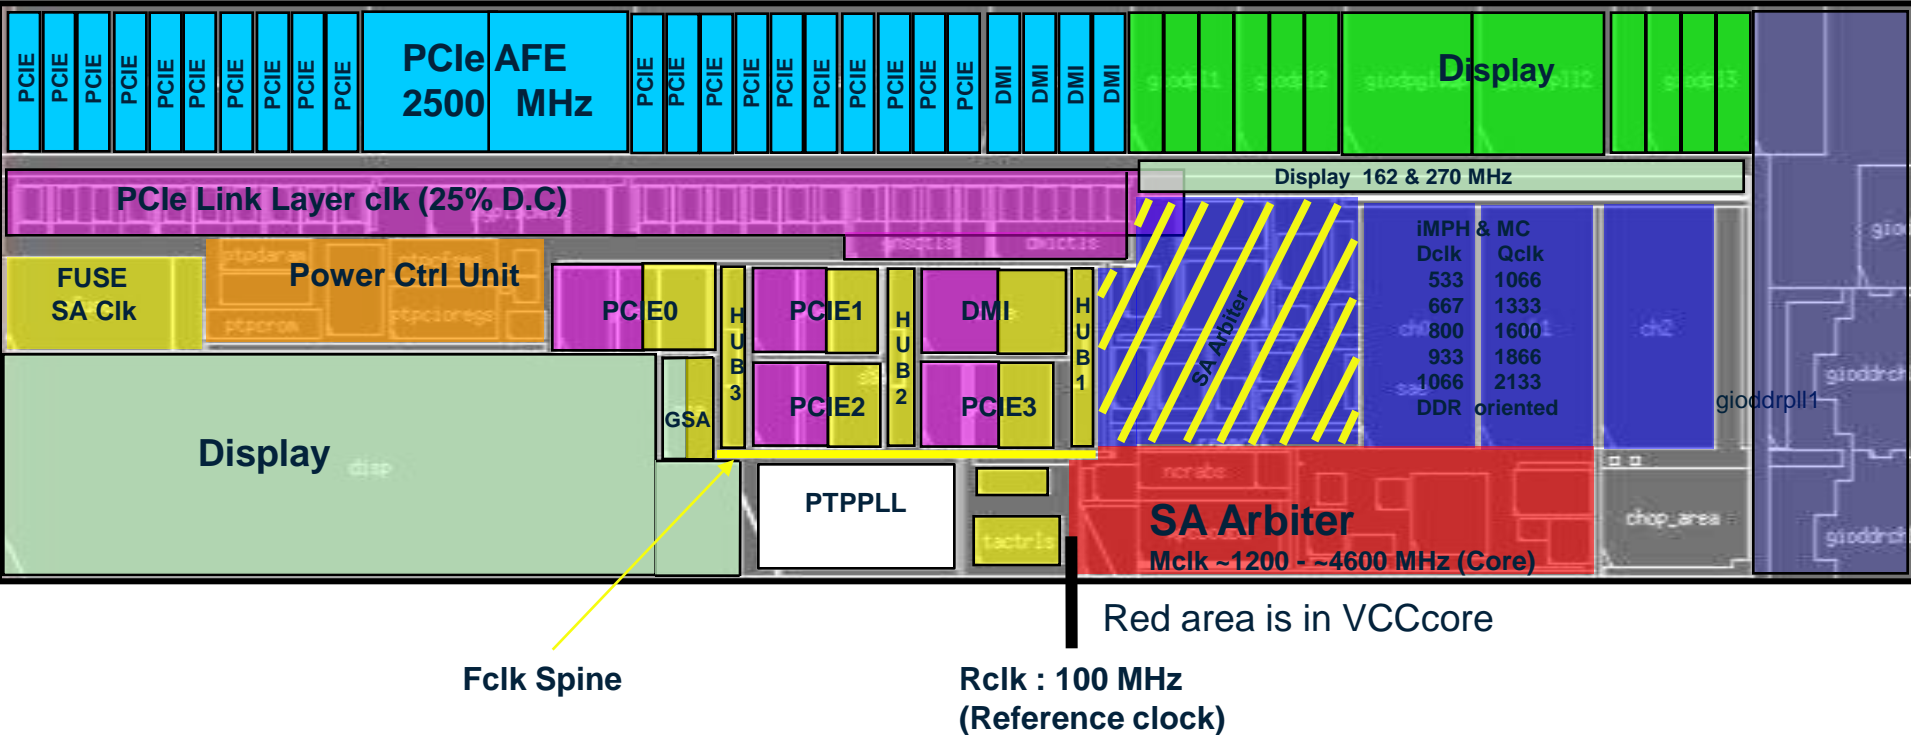
- ▶ PCI Express, DMI, Memory Controller, and Display Engine integrated on-die
 - ▶ Display Engine is part of the Integrated Graphics PCI Device
- ▶ Smart integration with the ring
 - ▶ Provides cores/Graphics /Media with high BW, low latency to DRAM/IO for best performance
 - ▶ Handles IO-to-cache coherency
 - ▶ Direct Memory Access (DMA) from devices to memory will snoop the CPU cache hierarchy
 - ▶ Address conflicts (multiple requests associated with the same cacheline in a given window of time) handled in SA
- ▶ Extensive power and thermal management for PCI Express and DDR
 - ▶ Links can go down to low power states
 - ▶ Thermal throttling available, invoked based on temperature calculations



Efficient Peripheral Device Integration



System Agent Clocking

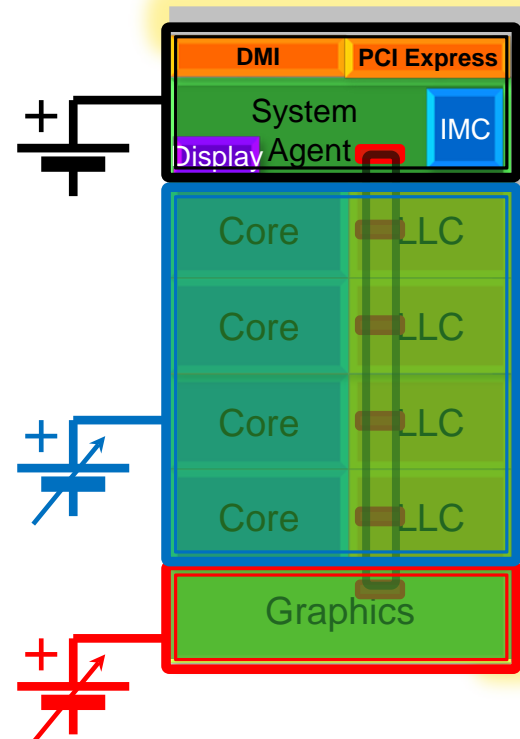


Chipset Complexities on High End CPU Product



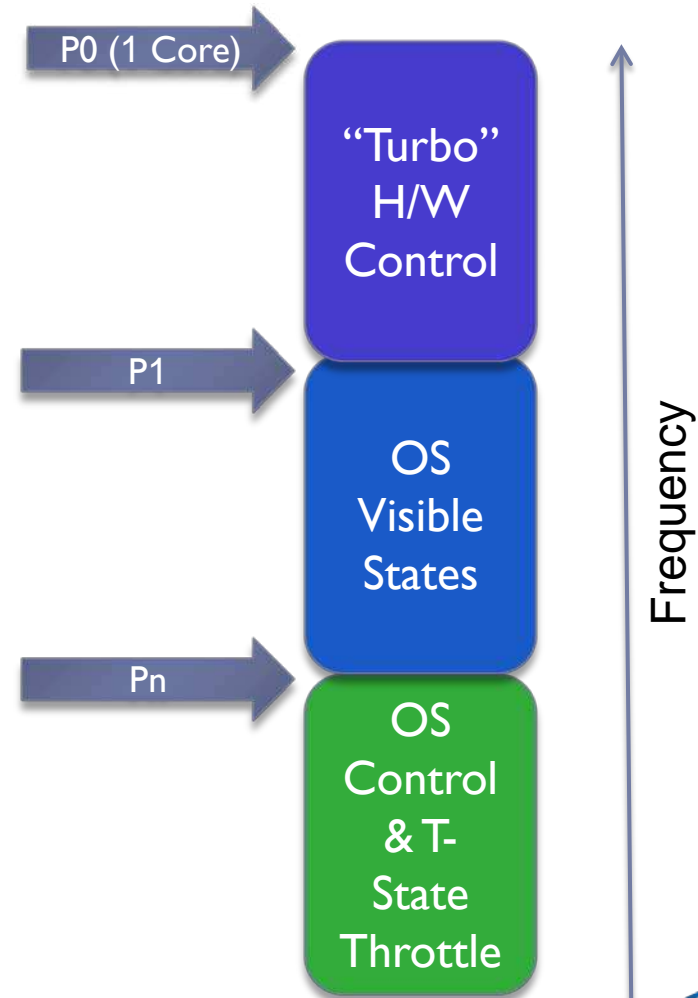
Power Management

- ▶ Separate voltage and frequency from ring/cores, Display integration for better battery life
 - ▶ Three separate power/frequency domains: SA (Fixed), Cores/Ring, Graphics (Variable)
- ▶ GPU is on a dedicated power plane
 - ▶ Can be powered down independently of the IA cores
 - ▶ Graphics turbo allows for higher frequency on both desktop and mobile parts
- ▶ Power Budget Management
 - ▶ Power Control Unit (PCU) – Programmable uController, handles all power management and reset functions in the chip
 - ▶ **Dynamically redistribute** power between IA & PG
 - Power budget allocated based on workload requirements

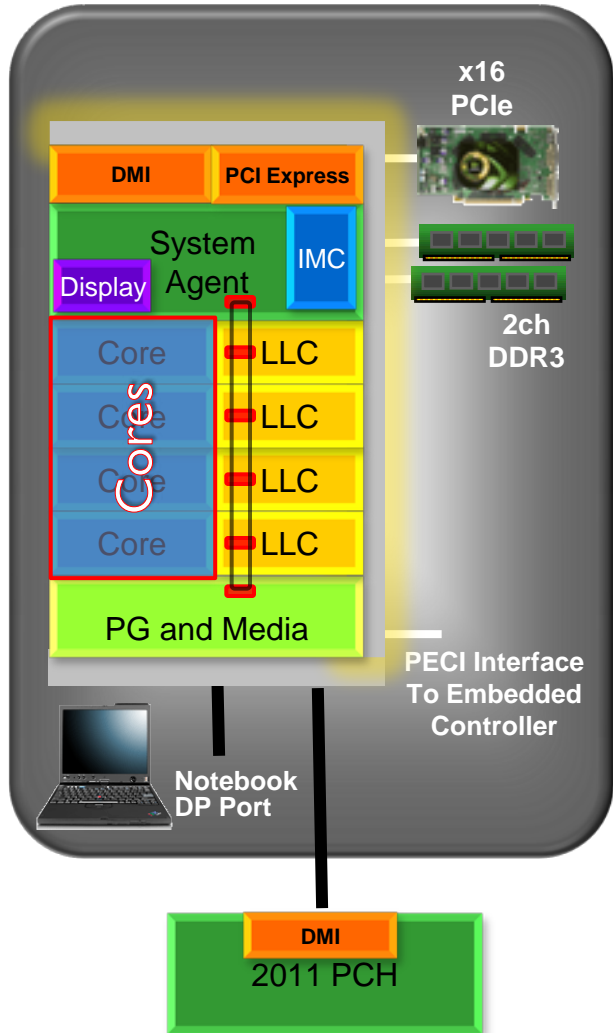


Intel® Turbo Boost Technology 2.0

- ▶ Performance on Demand
 - ▶ CPU can run faster than base operating frequency
 - ▶ Working within power, current and temperature constraints
 - ▶ Dependent on number of active cores
 - Highest frequency is with 1 core (gets all the headroom)
- ▶ Processor operates at power level higher than rated upper power limit (TDP) for short durations
 - ▶ P0 is Maximum Possible Frequency
 - ▶ P1 is Guaranteed Frequency
 - ▶ Pn is Energy Efficient Frequency



2nd Gen Intel[®] Core[™] Microarchitecture



Intel[®] Core[™] Enhancements and IA Extensions

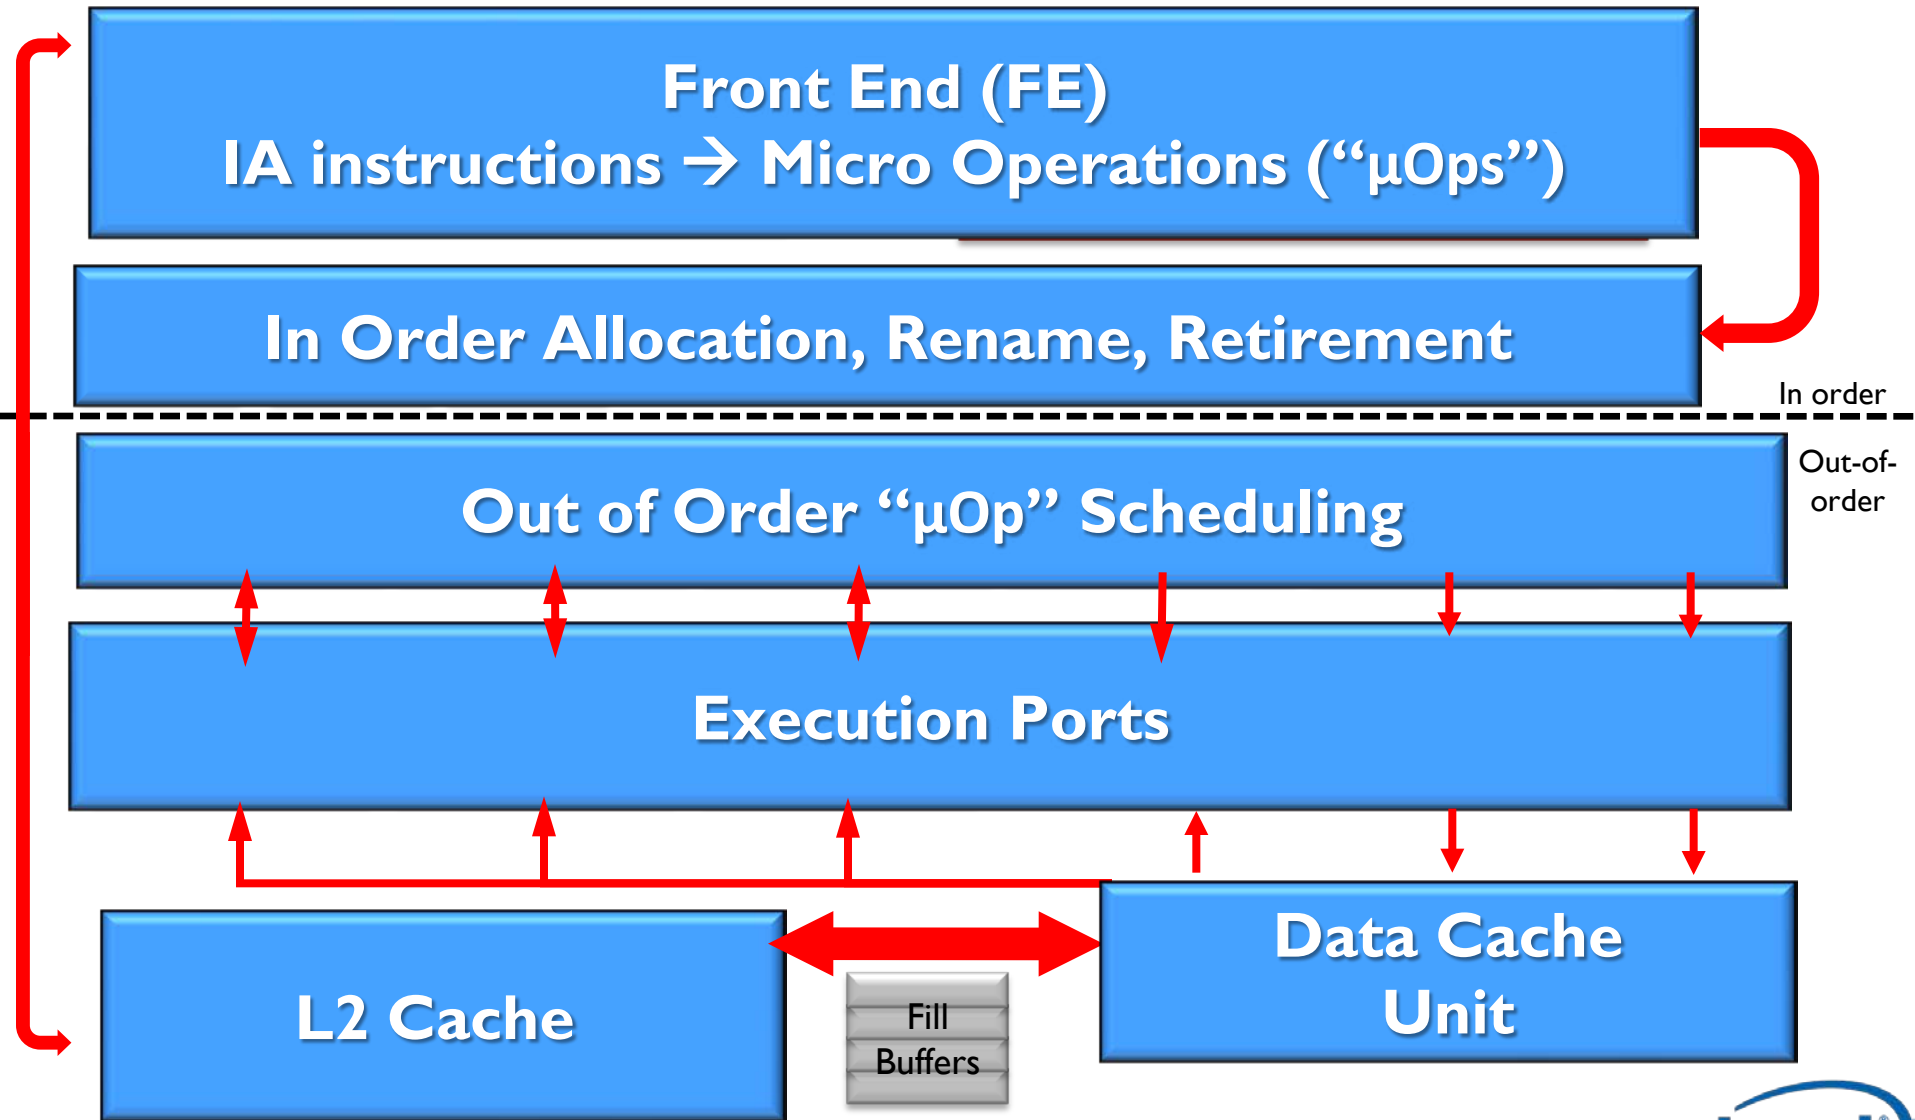


Sandy Bridge Processor Core

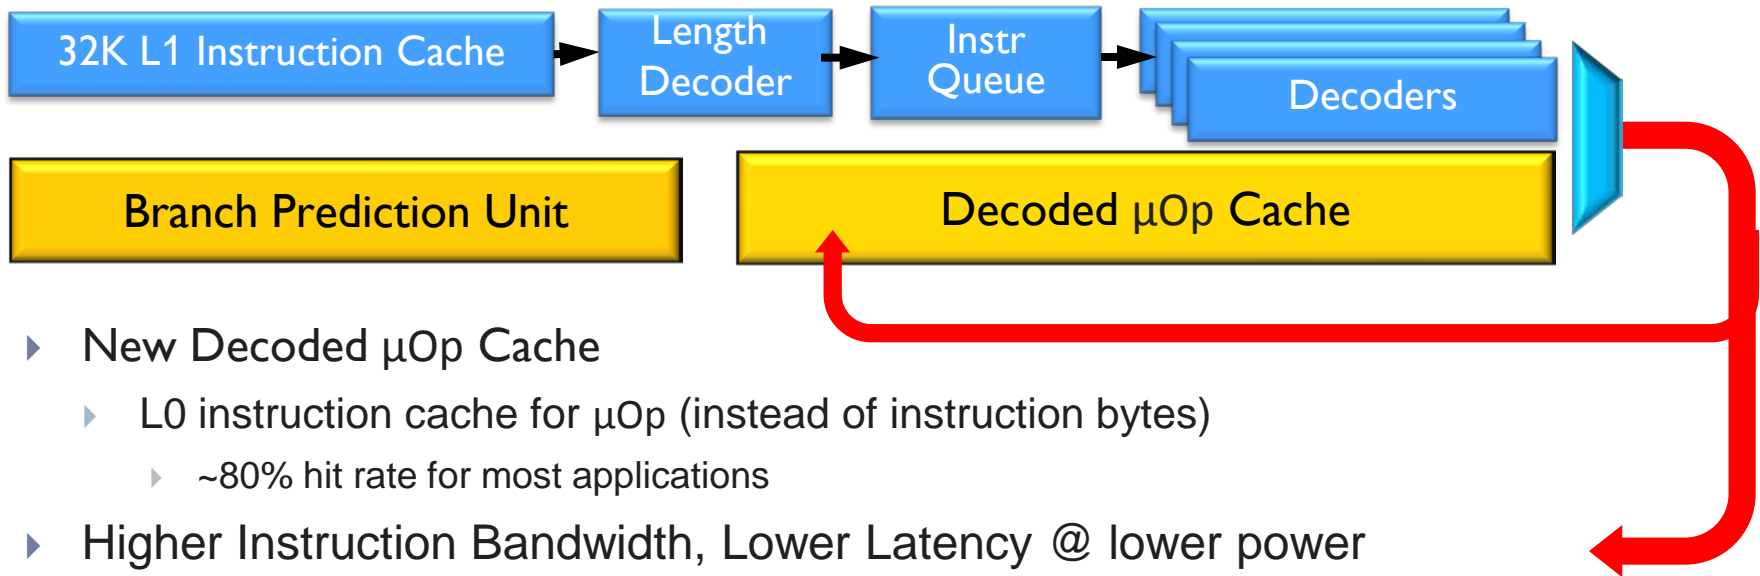
- ▶ **Converged** building block for Mobile, Desktop, and Server
- ▶ **Power-Performance** micro-architecture features
 - ▶ Significant redesign relative to the previous generation Core
 - ▶ “**Cool**” features – better than linear performance/power
 - ▶ “**Really Cool**” features – gain performance, reduce power
- ▶ **ISA Extensions** for important new usages
 - ▶ Floating Point and Throughput
 - ▶ Intel® Advanced Vector Extensions (Intel® AVX) - Significant boost for selected compute intensive applications
 - ▶ Security
 - ▶ AES (Advanced Encryption Standard) throughput enhancements
 - ▶ Large Integer RSA speedups
 - ▶ OS/VMM and Server related features
 - ▶ State save/restore optimizations



Core Microarchitecture



New FE – Decoded μ Op Cache




- ▶ New Decoded μ Op Cache
 - ▶ L0 instruction cache for μ Op (instead of instruction bytes)
 - ▶ ~80% hit rate for most applications
- ▶ Higher Instruction Bandwidth, Lower Latency @ lower power
 - ▶ Decoders can be shut down when not needed
 - ▶ Decoded μ Op cache throughput 32B/cycle
 - ▶ More Cycles sustaining 4 instruction/cycle
- ▶ Coupled with improved Branch Prediction Unit
 - ▶ Able to “stitch” across taken branches in the control flow

**“Really Cool” Feature –
Save Power while Increasing Performance**

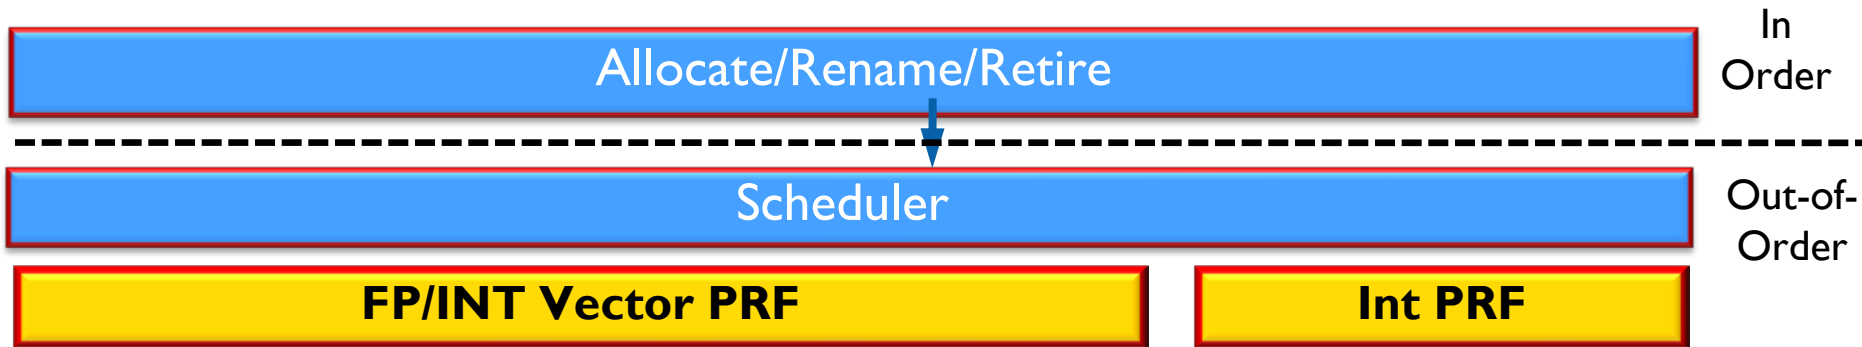
Intel® AVX – Doubling the FLOPs

- ▶ Extend SSE FP instruction set to 256 bits operand size
 - ▶ Intel AVX extends all 16 XMM registers to 256bits
- ▶ New, non-destructive source syntax
 - ▶ VADDPS ymm1, ymm2, ymm3
- ▶ 256-bit Multiply + 256-bit ADD + 256-bit Load per clock...
- ▶ New Operations to enhance vectorization
 - ▶ Broadcasts
 - ▶ Masked load & store



“Cool Feature” –
Wide vectors and non-destructive source: more work, fewer instructions
Extending the existing state is area and power efficient

New OOO – Physical Register File

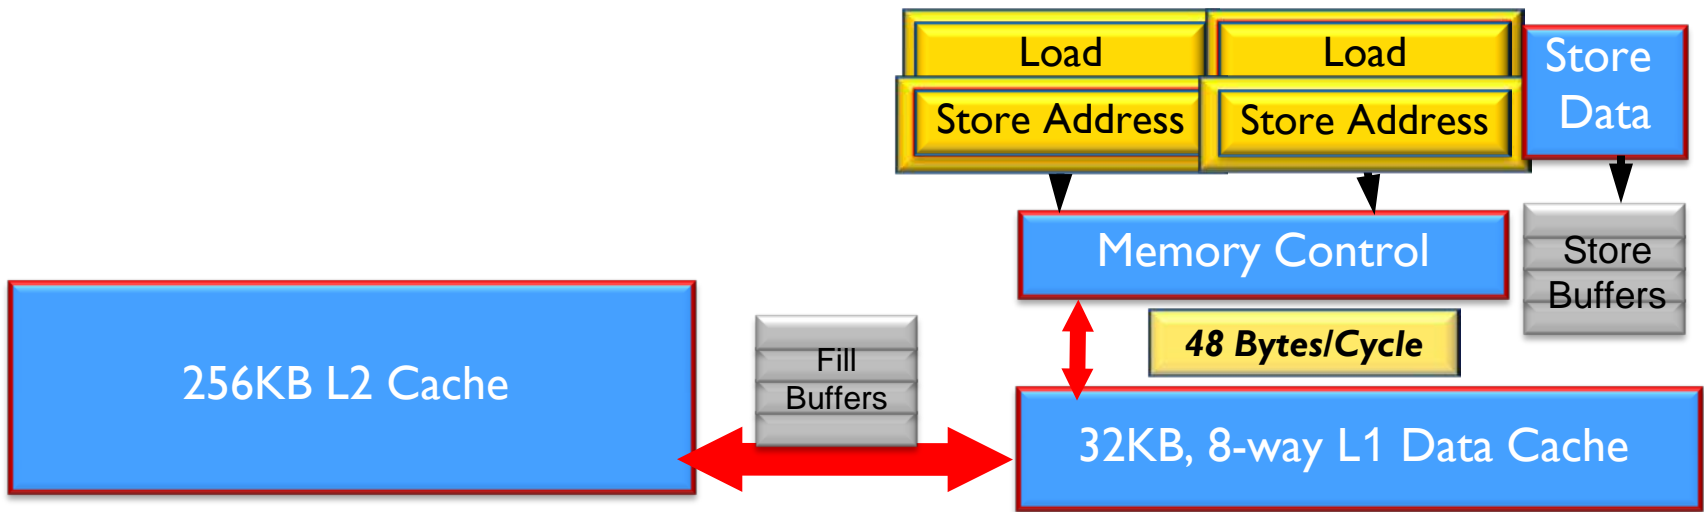


- ▶ Move to a Physical Register File (PRF)
 - ▶ A physical register file stores μOp operands
 - ▶ As μOp flows in the OOO engine it only carries pointers to the operands
 - ▶ No movement after calculation
 - ▶ In previous architecture every μOp had a copy of every operand it needed
- ▶ Allow significant increase in buffer sizes
 - ▶ Parallelism window ~33% larger

“Cool” Feature – Better than linear performance/power
Key enabler for Intel® Advanced Vector Extensions (Intel® AVX)

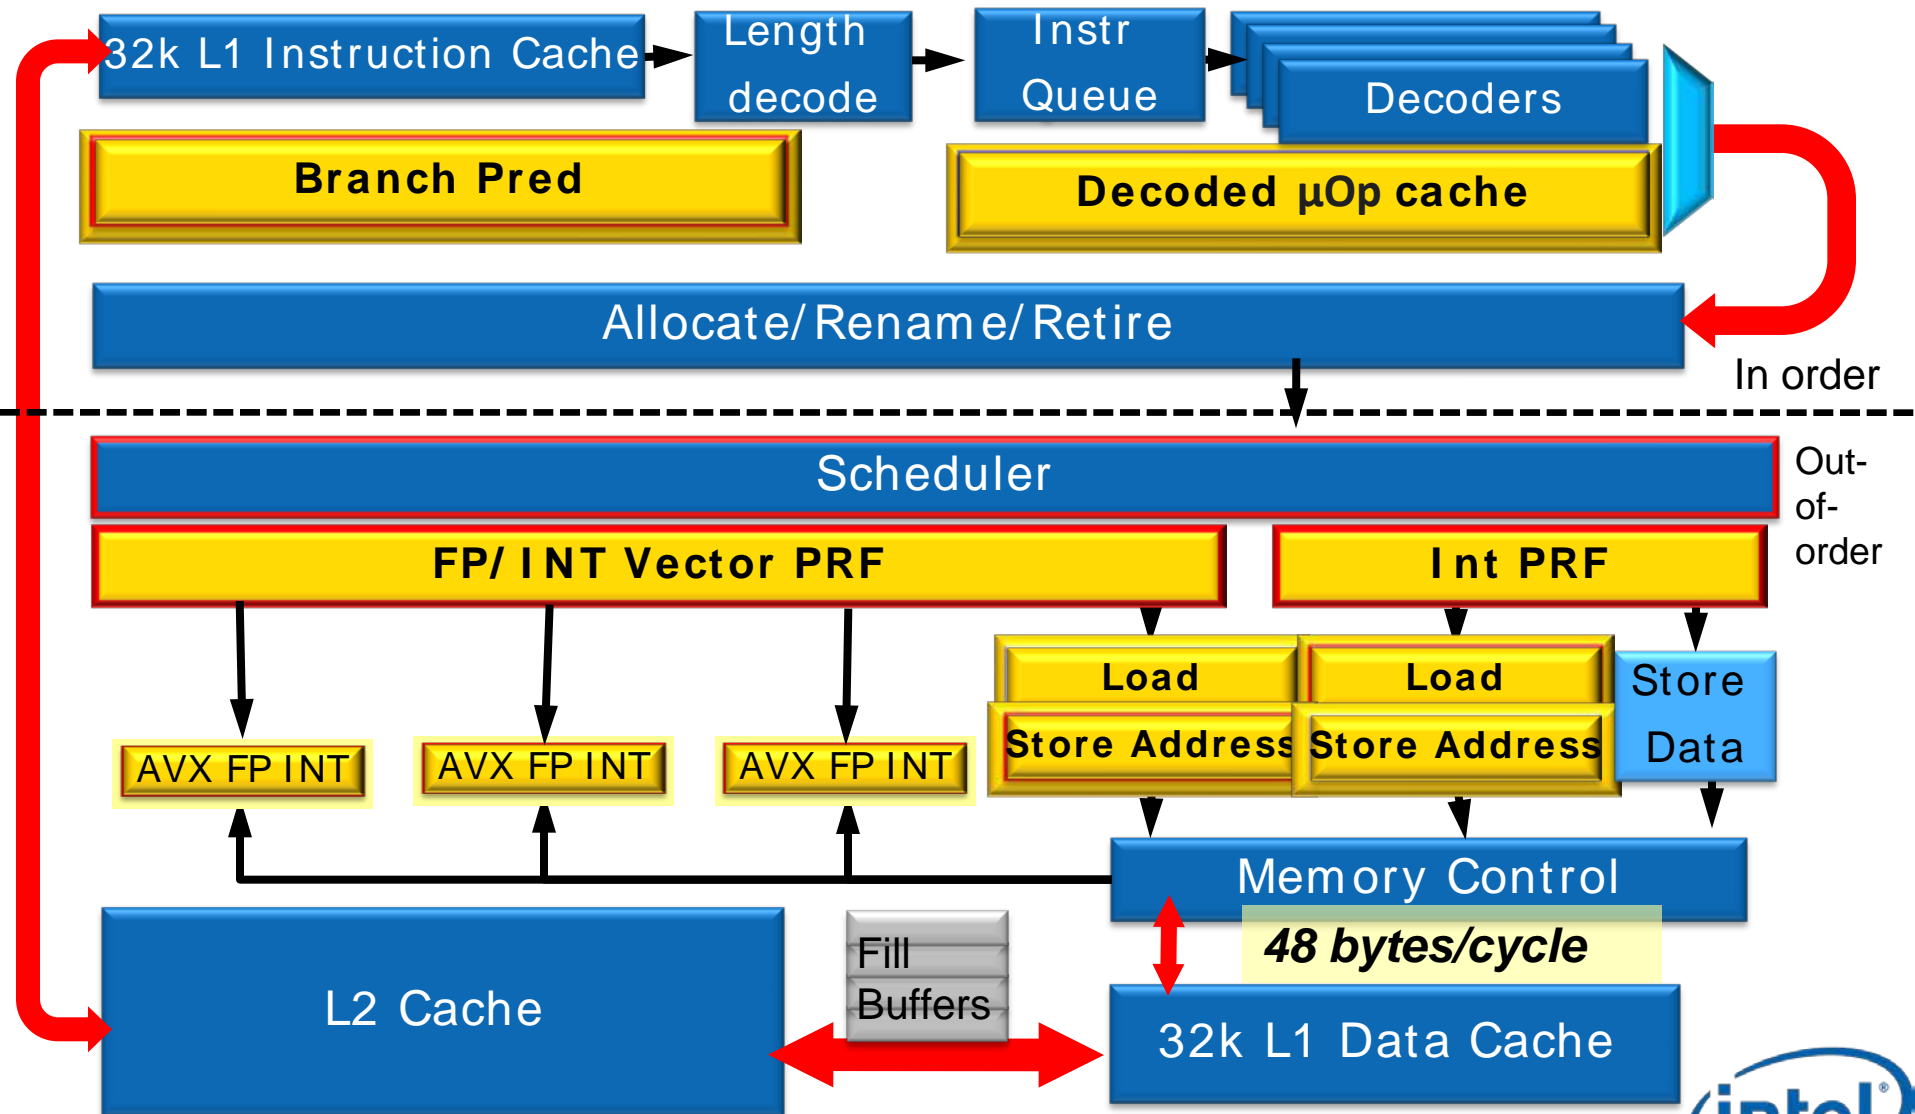


New Memory Cluster



- ▶ Memory Unit serves three data accesses per cycle
 - ▶ Upgrade separate load and store pipes to be multi-purpose
 - ▶ Two read requests of up to 16B and 1 store request of up to 16B
 - ▶ 2nd Load Port essential to keep up with Intel® AVX memory requirements

Putting it Together – Core Microarchitecture

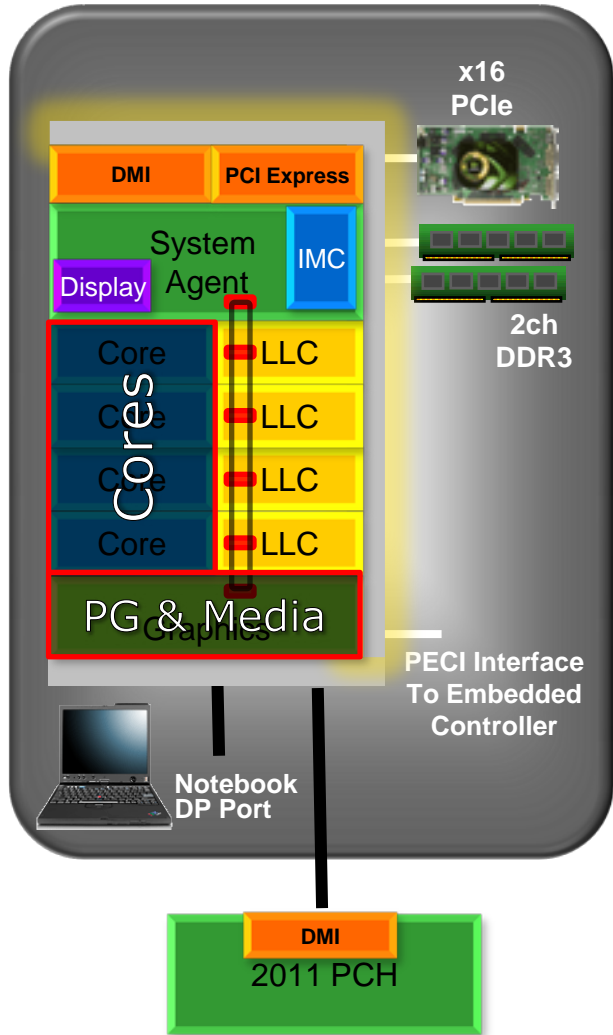


Other Architectural Extensions

- ▶ **Cryptography Instruction Throughput Enhancements**
 - ▶ Throughput for AES instructions introduced in Intel® Core™ microarchitecture (formerly codenamed Westmere)
- ▶ **Large Number Arithmetic Throughput Enhancements**
 - ▶ ADC (Add with Carry) throughput doubled
 - ▶ Multiply (64-bit multiplicands with 128-bit product)
 - ▶ ~25% speedup on existing RSA binaries!
- ▶ **State Save/Restore Enhancements**
 - ▶ New state added in Intel® Advanced Vector Extensions (Intel® AVX)
 - ▶ HW monitors features used by applications
 - ▶ Only saves/restores state that is used



2nd Gen Intel[®] Core[™] Microarchitecture



Integration Challenges



Process Technology Challenges

- ▶ IA Core requires **Fast** process for performance, **Low Leakage** for TDP Power
 - ▶ C-State residency for Idle power
 - ▶ CPU low per mode state
- ▶ Processor Graphics requires **Fast** process for performance, **Ultra Low Leakage** for Power
 - ▶ C-State residency for Idle power
- ▶ System Agent Requires **Ultra Low leakage** for Idle Power
 - ▶ “Mostly On”

Advanced manufacturing technology enables multi Transistor flavors on single die (low leakage & ultra low leakage)

Advanced dynamic Power Management enables flexibility – lower P1 (guaranteed) frequency for higher performance!



Process Technology Challenges

- ▶ New IOs (PCI Express, DMI, DDR) added to mainstream CPUs in top end process
 - ▶ Voltage supply issues
 - ▶ IO requires higher voltage (translates to higher power)
 - ▶ Non-Scalable design with process generations
 - ▶ Gain in area with newer process is much less for the IO
- ▶ Multiple Analog devices
 - ▶ PLLs due to multiple clocking
 - ▶ IOs
- ▶ Isolations and Crossovers due to multiple asynchronous power planes and clock domains

Many IOs still remain in the PCH – Strict Architectural criteria for IO and Analog integration

Power-Performance vs. Complexity tradeoffs



Summary



32nm Next Generation Core[®]
Microarchitecture

Processor Graphics

System Agent, Ring Architecture
and Other Innovations

Intel[®] AVX

Performance and Power Efficiency



Q&A



Glossary

Acronym	Term
QPI	QuickPath Interconnect
LLC	Last Level Cache
GMCH	Graphics and Memory Controller Hub
LVDS	Low Voltage Differential Signalling
DMI	Direct Media Interface
TDP	Thermal Design Point
DP	Display Port
eDP	Embedded Display Port
IA	Intel Architecture
PG	Processor Graphics
OOO	Out-Of-Order
PCH	Platform Controller Hub

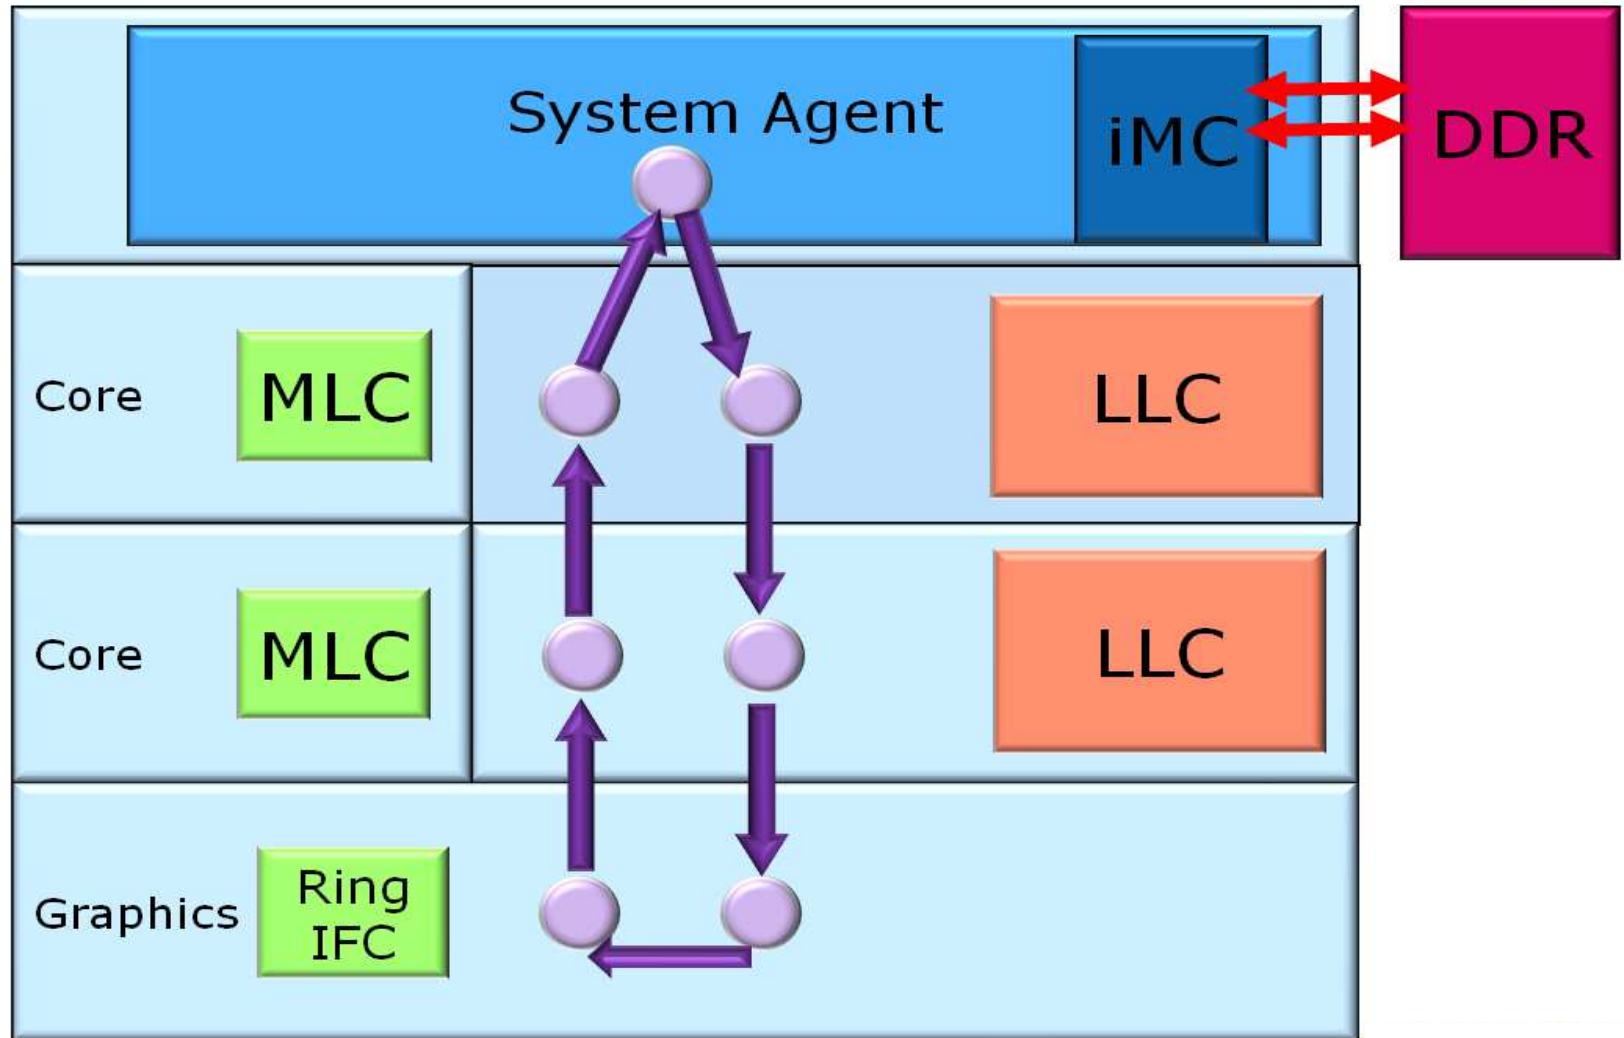


Glossary (cont)

Acronym	Term
AVX	Advanced Vector Extension
iMC	Integrated Memory Controller
EV	Electrical Validation
DFD	Design For Debug

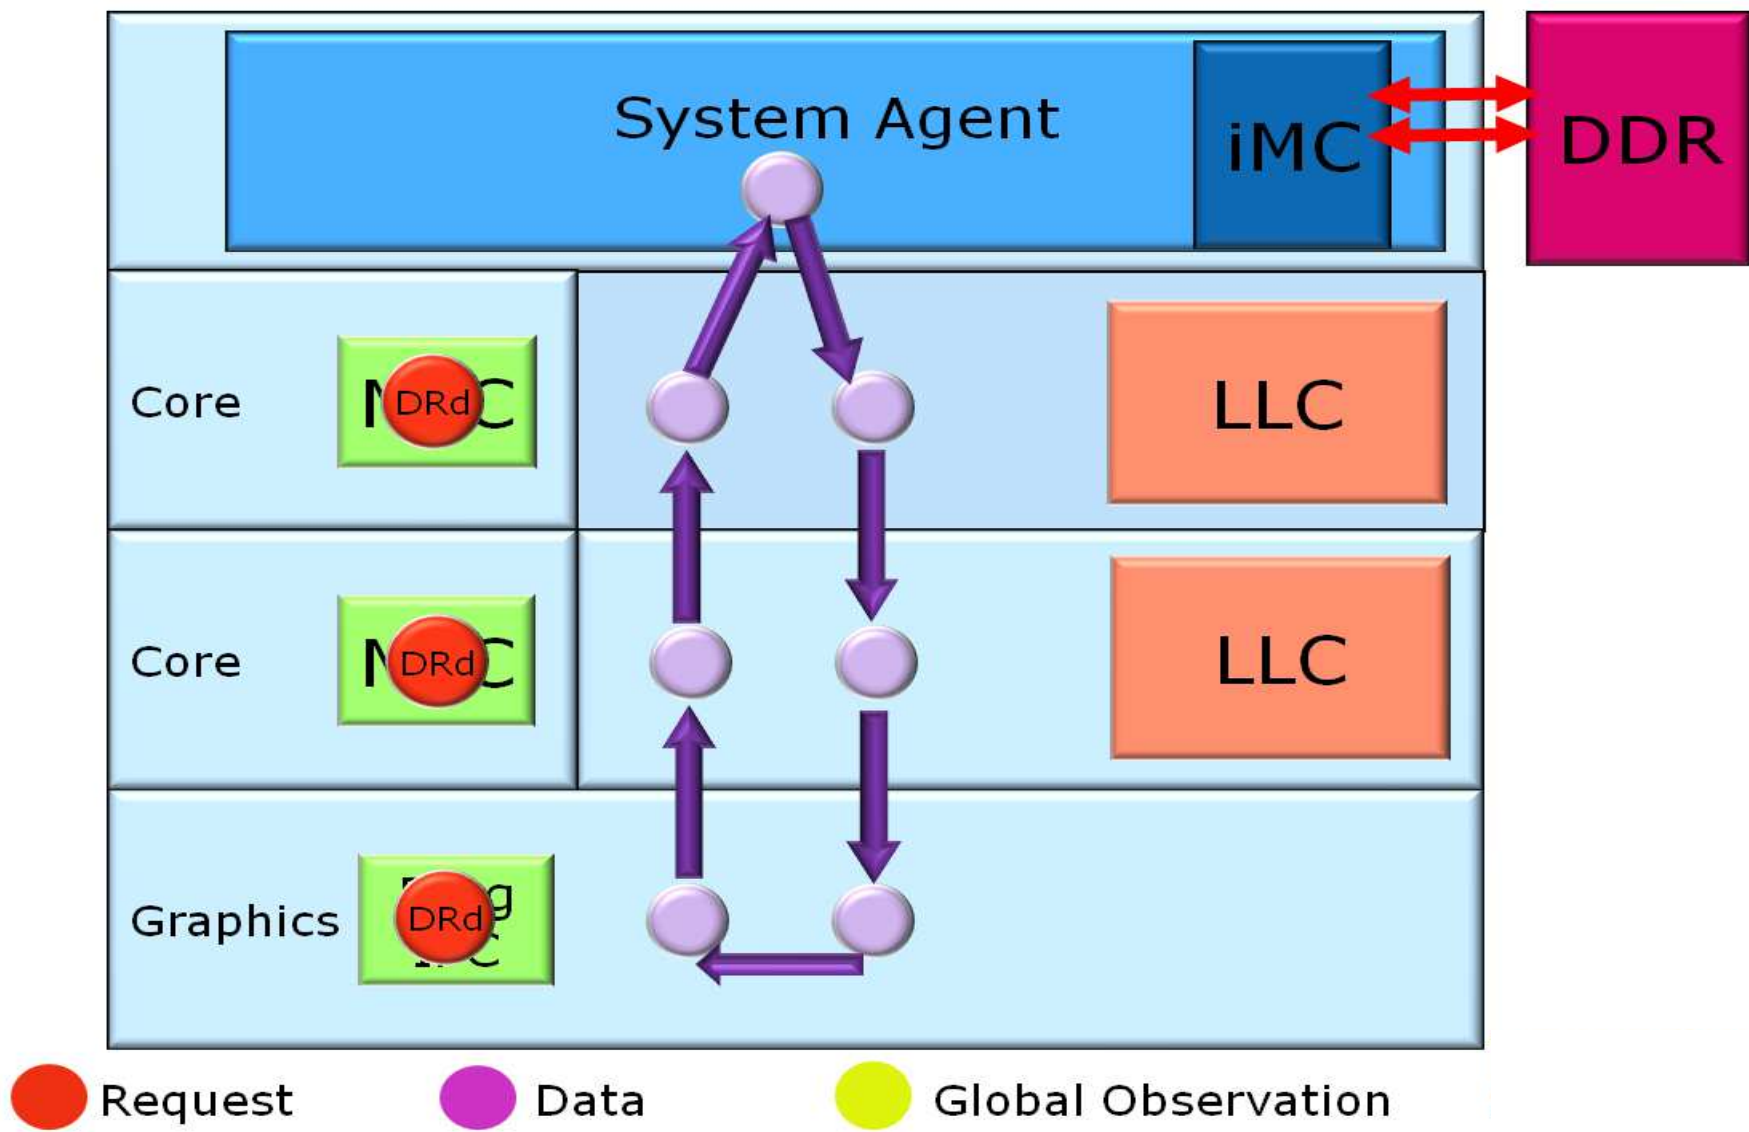


Ring Illustration: Clean LLC Hit

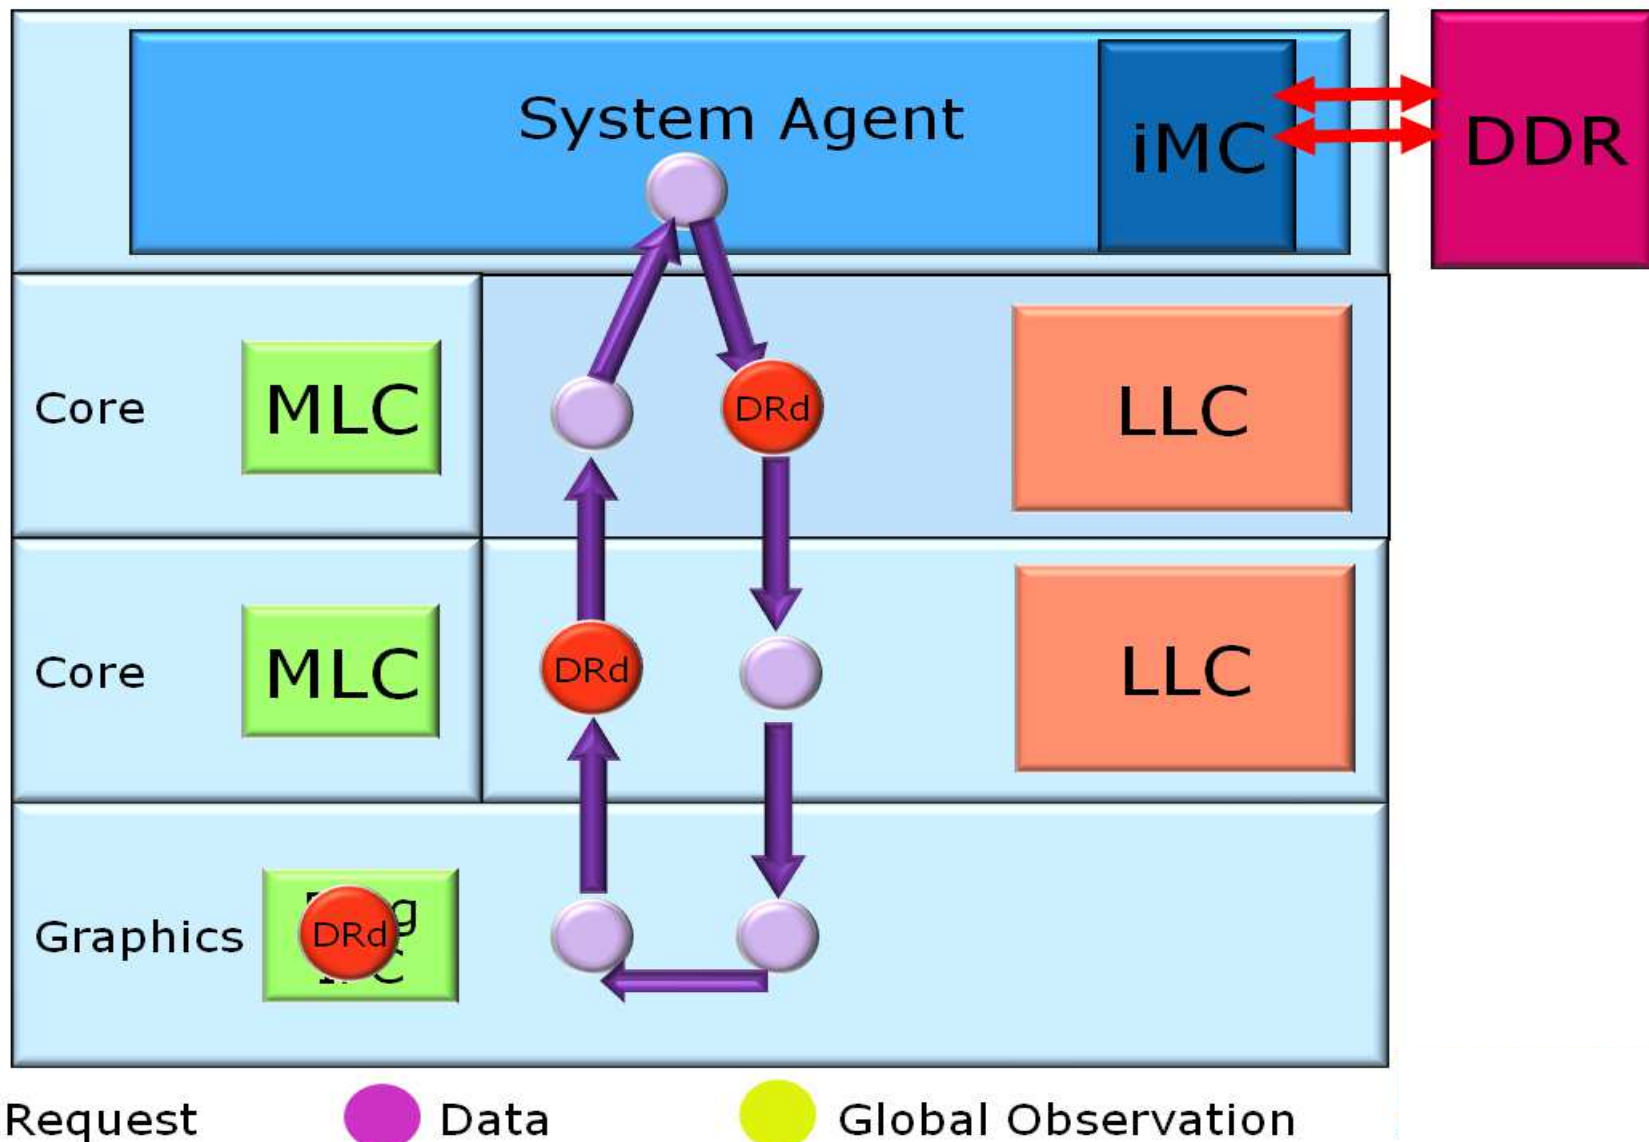


● Request ● Data ● Global Observation

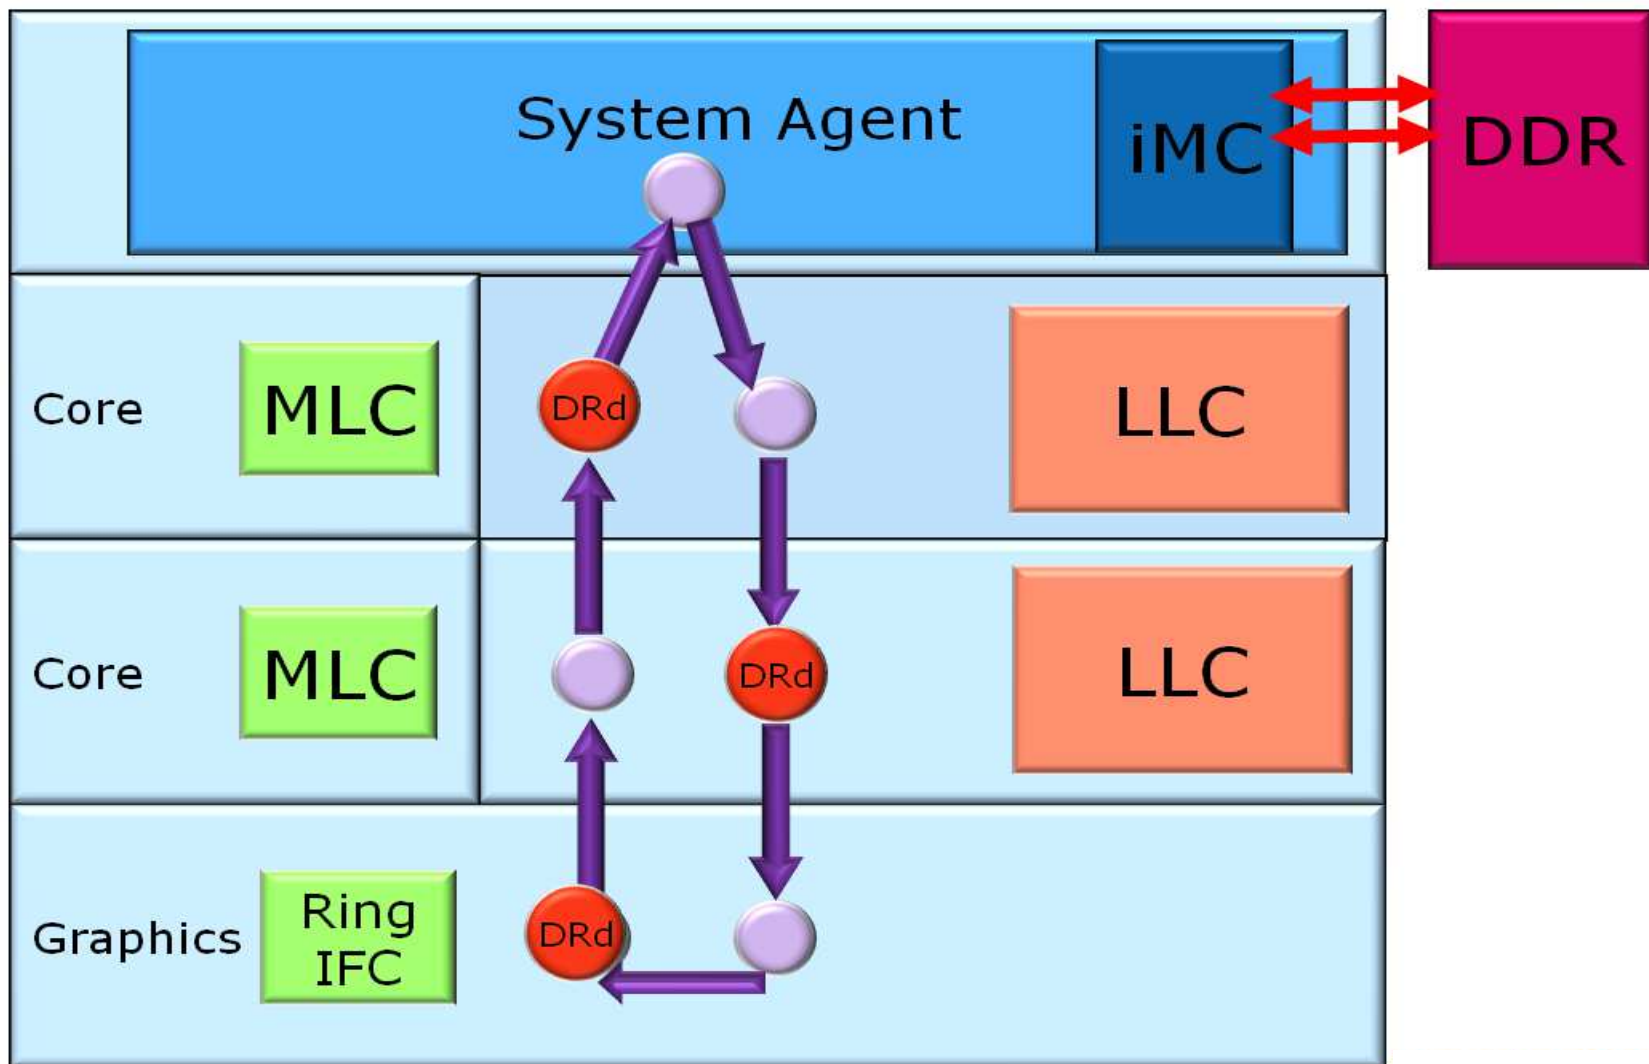
Ring Illustration: Clean LLC Hit



Ring Illustration: Clean LLC Hit



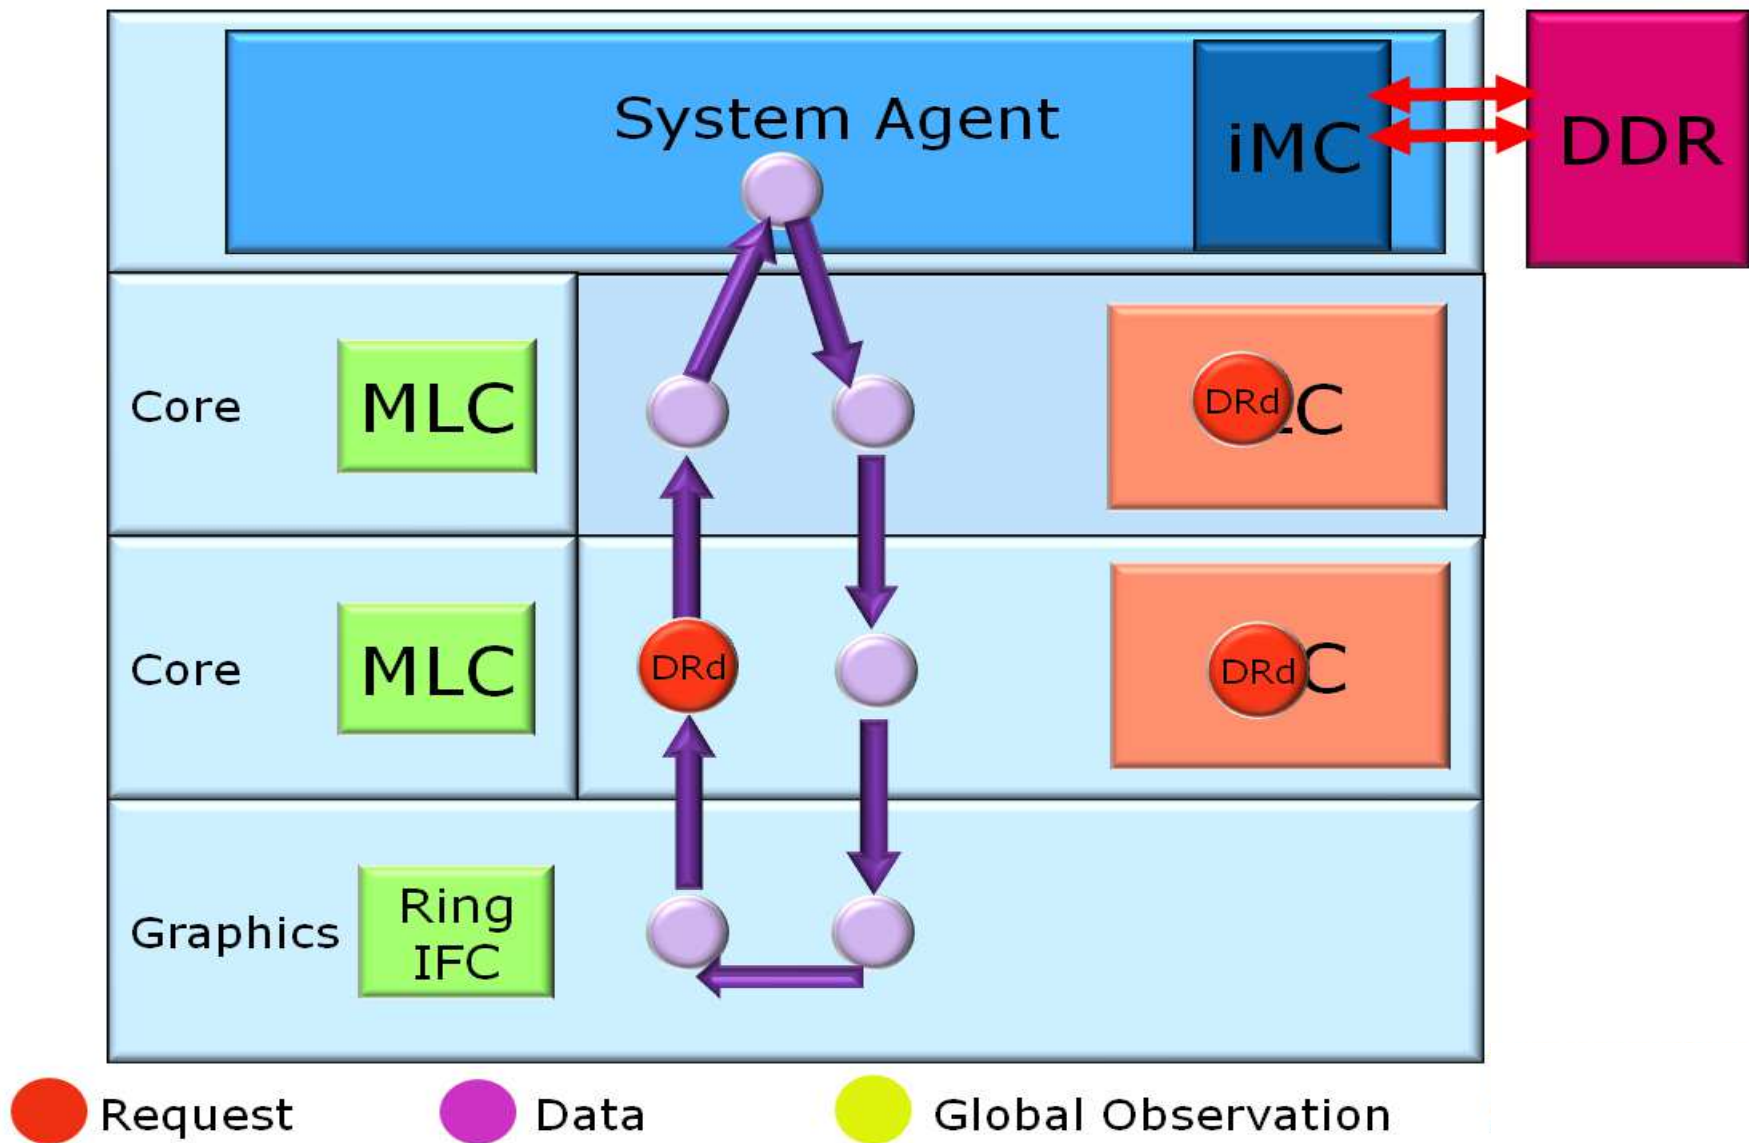
Ring Illustration: Clean LLC Hit



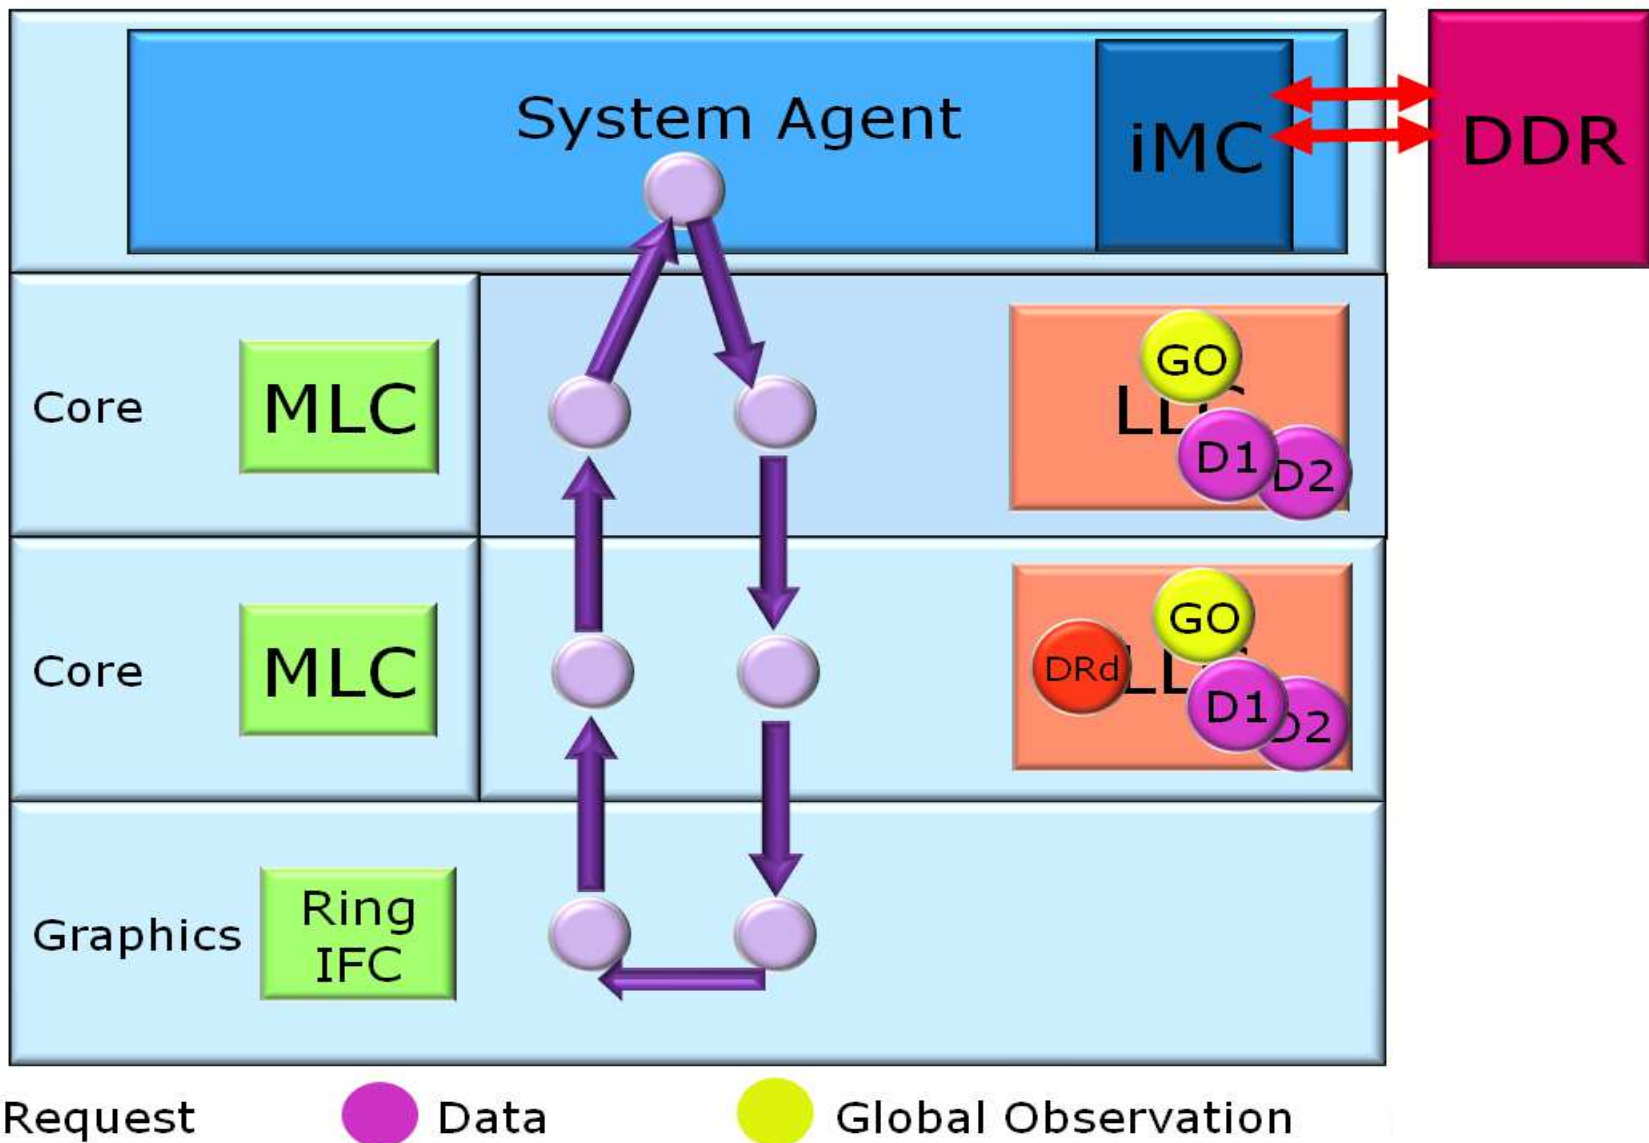
● Request ● Data ● Global Observation



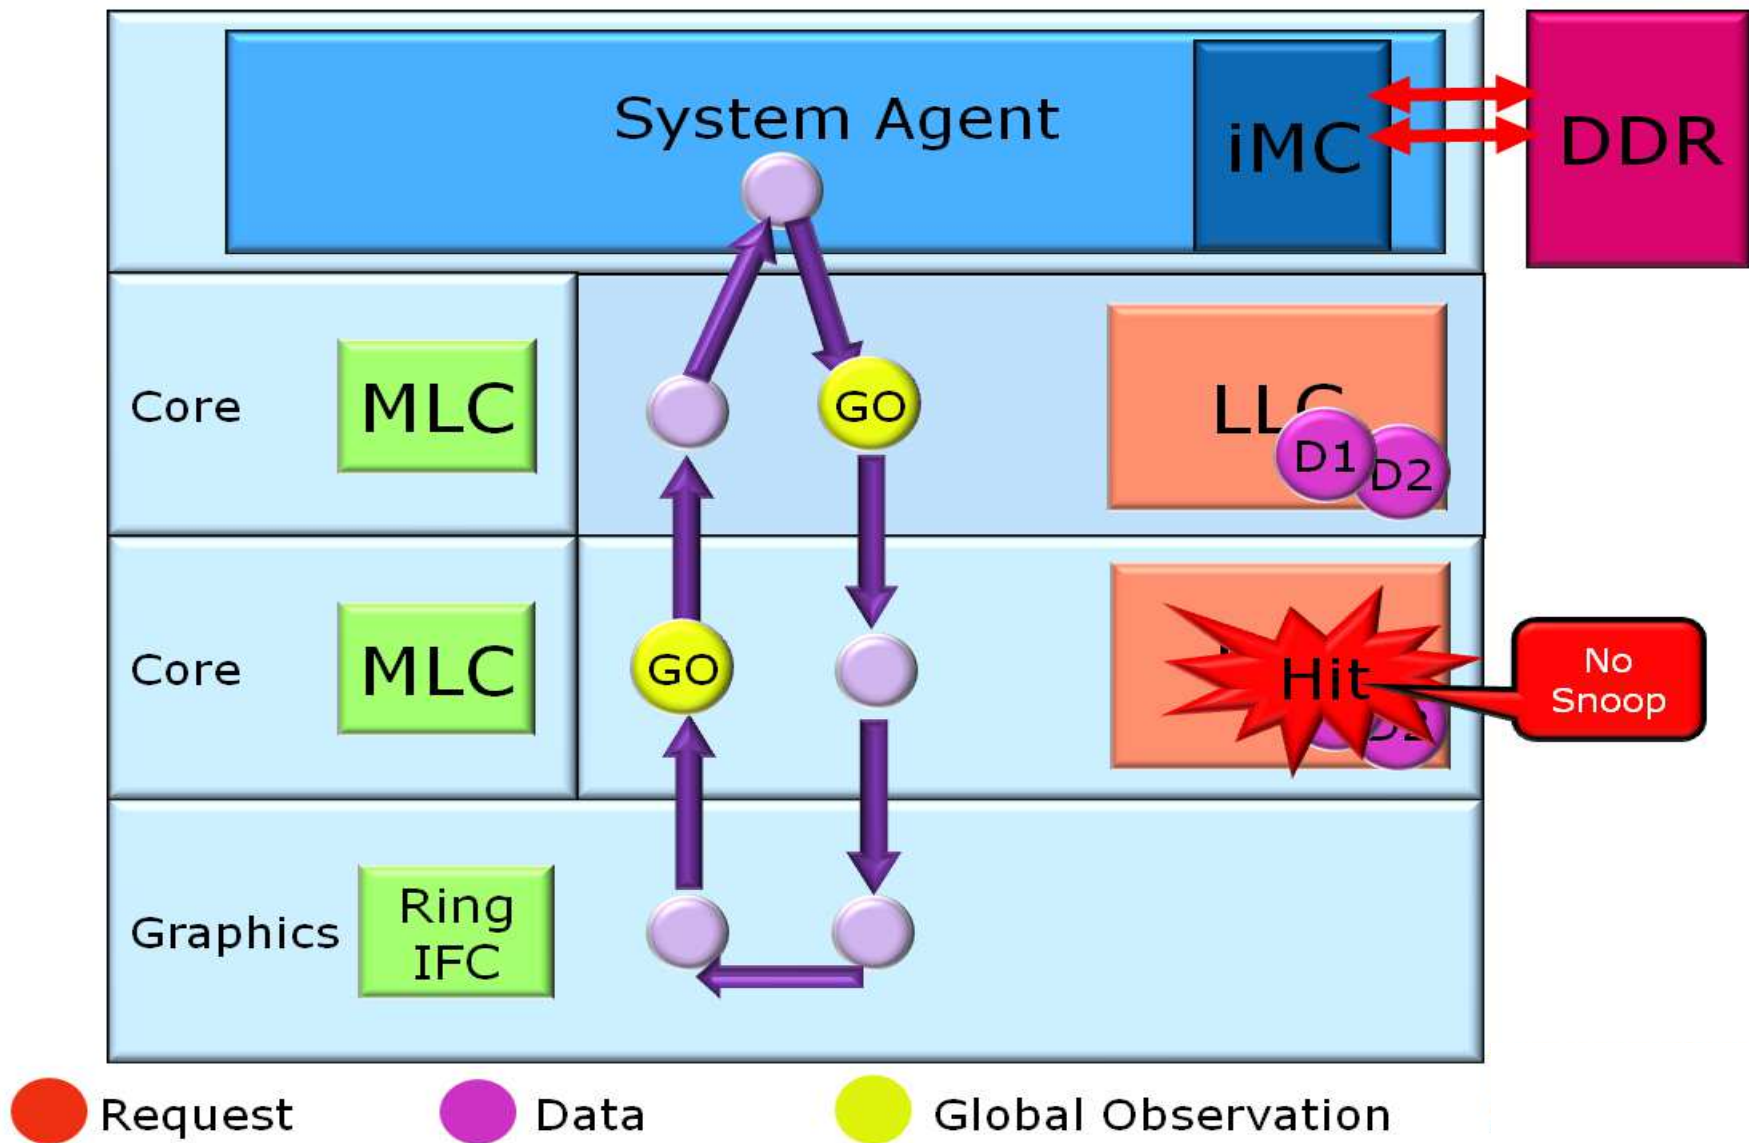
Ring Illustration: Clean LLC Hit



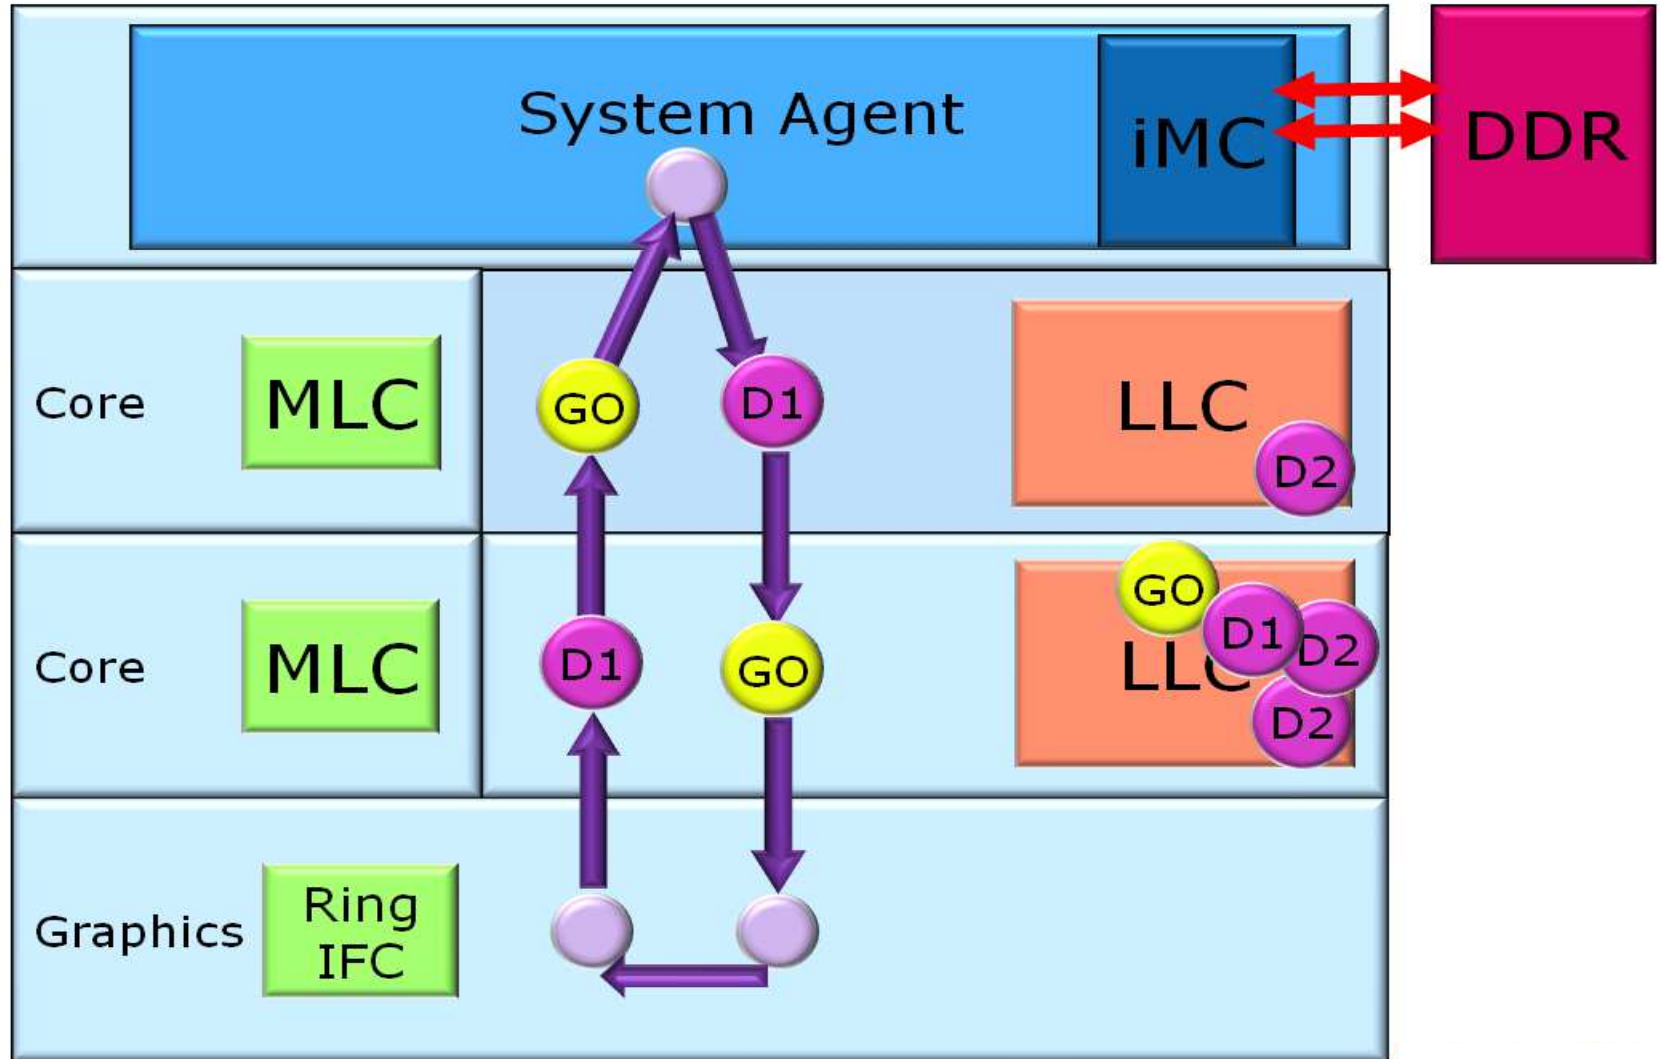
Ring Illustration: Clean LLC Hit



Ring Illustration: Clean LLC Hit

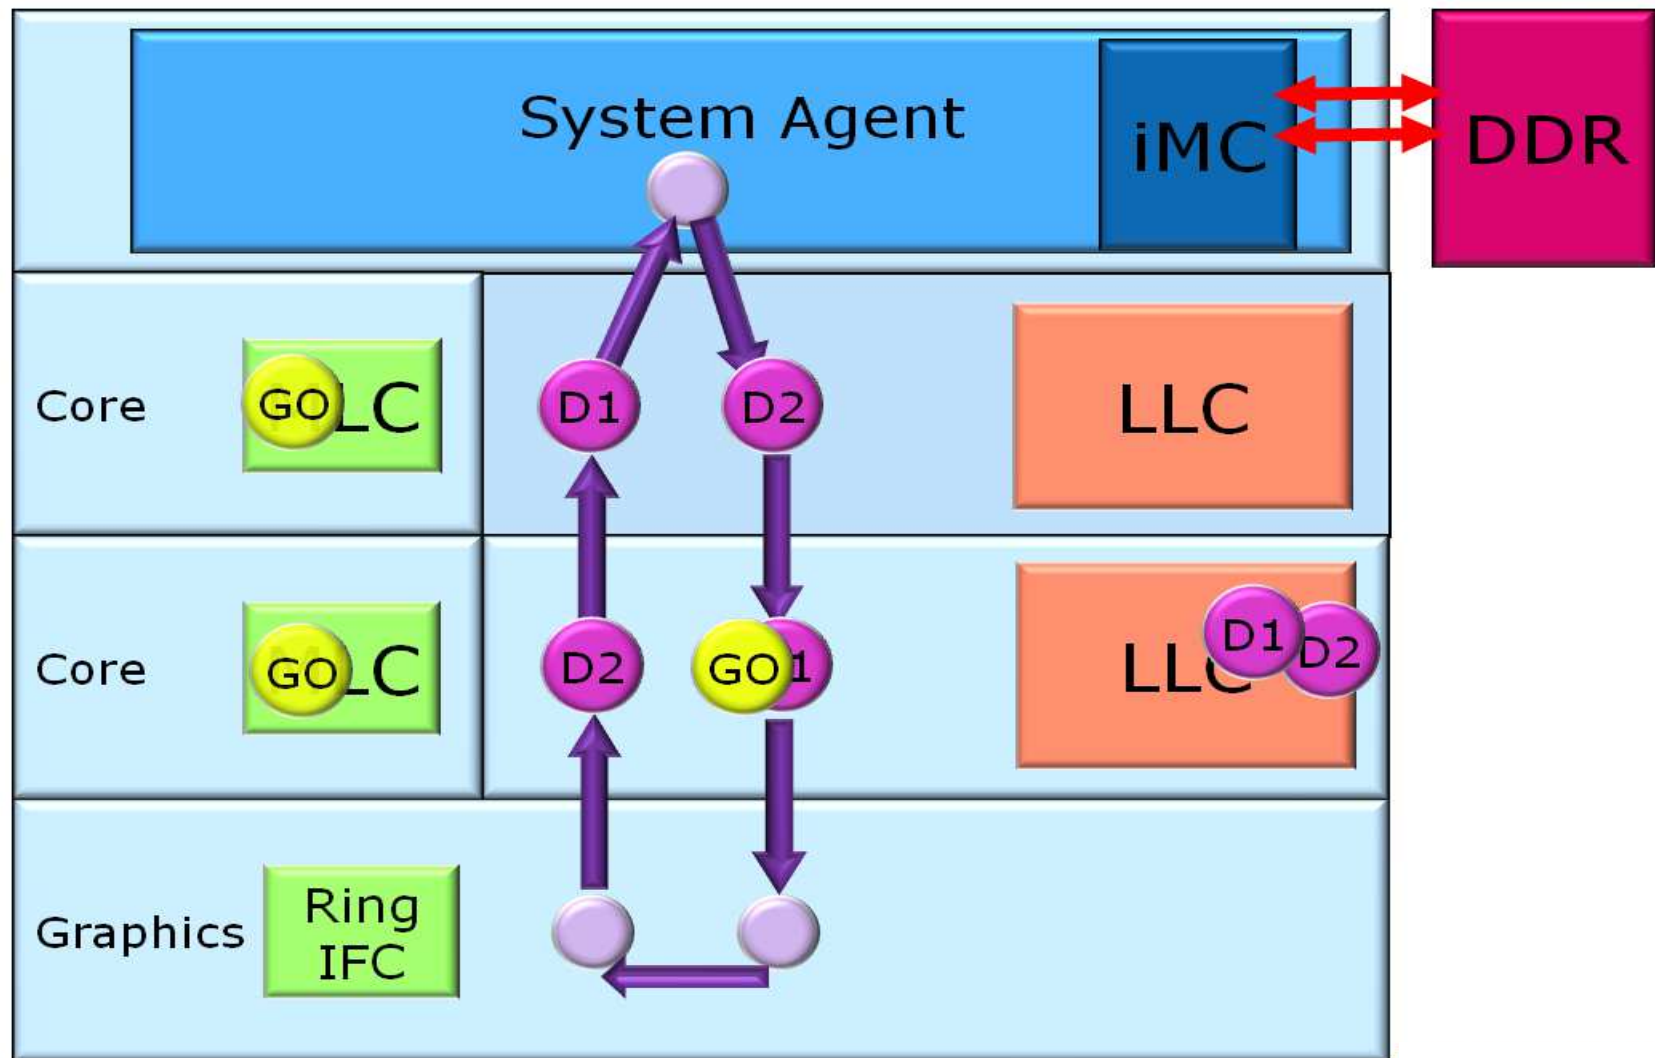


Ring Illustration: Clean LLC Hit



● Request ● Data ● Global Observation

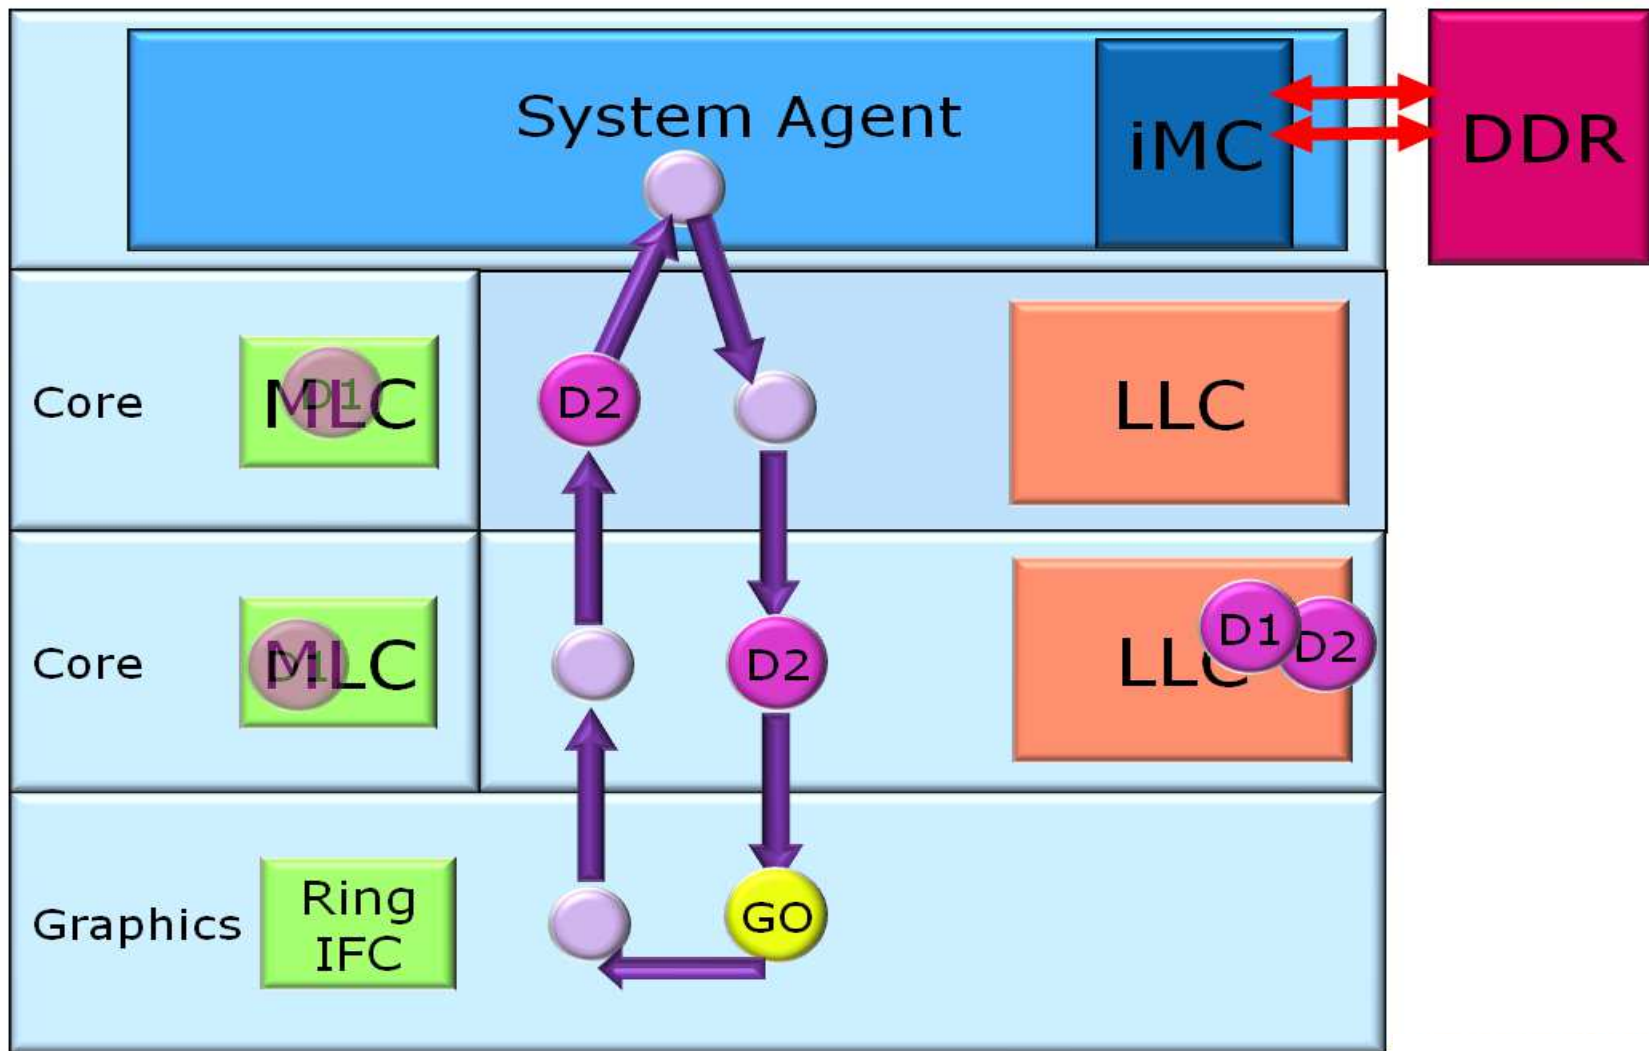
Ring Illustration: Clean LLC Hit



● Request ● Data ● Global Observation

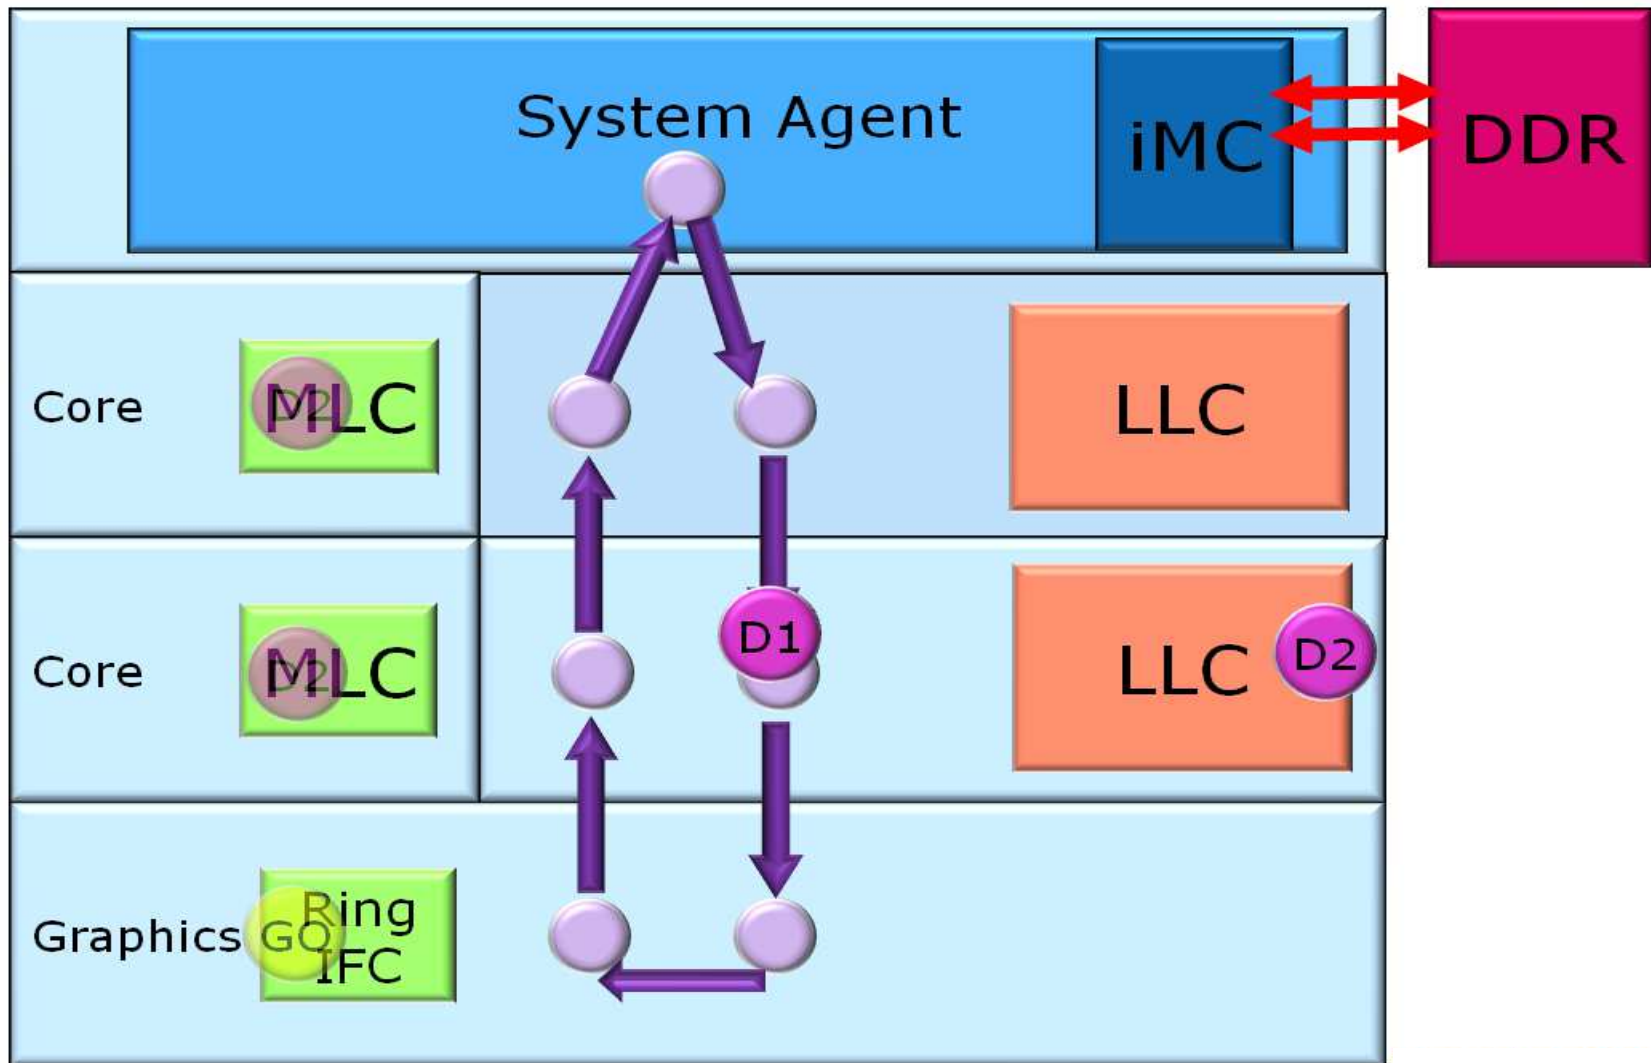


Ring Illustration: Clean LLC Hit



● Request ● Data ● Global Observation

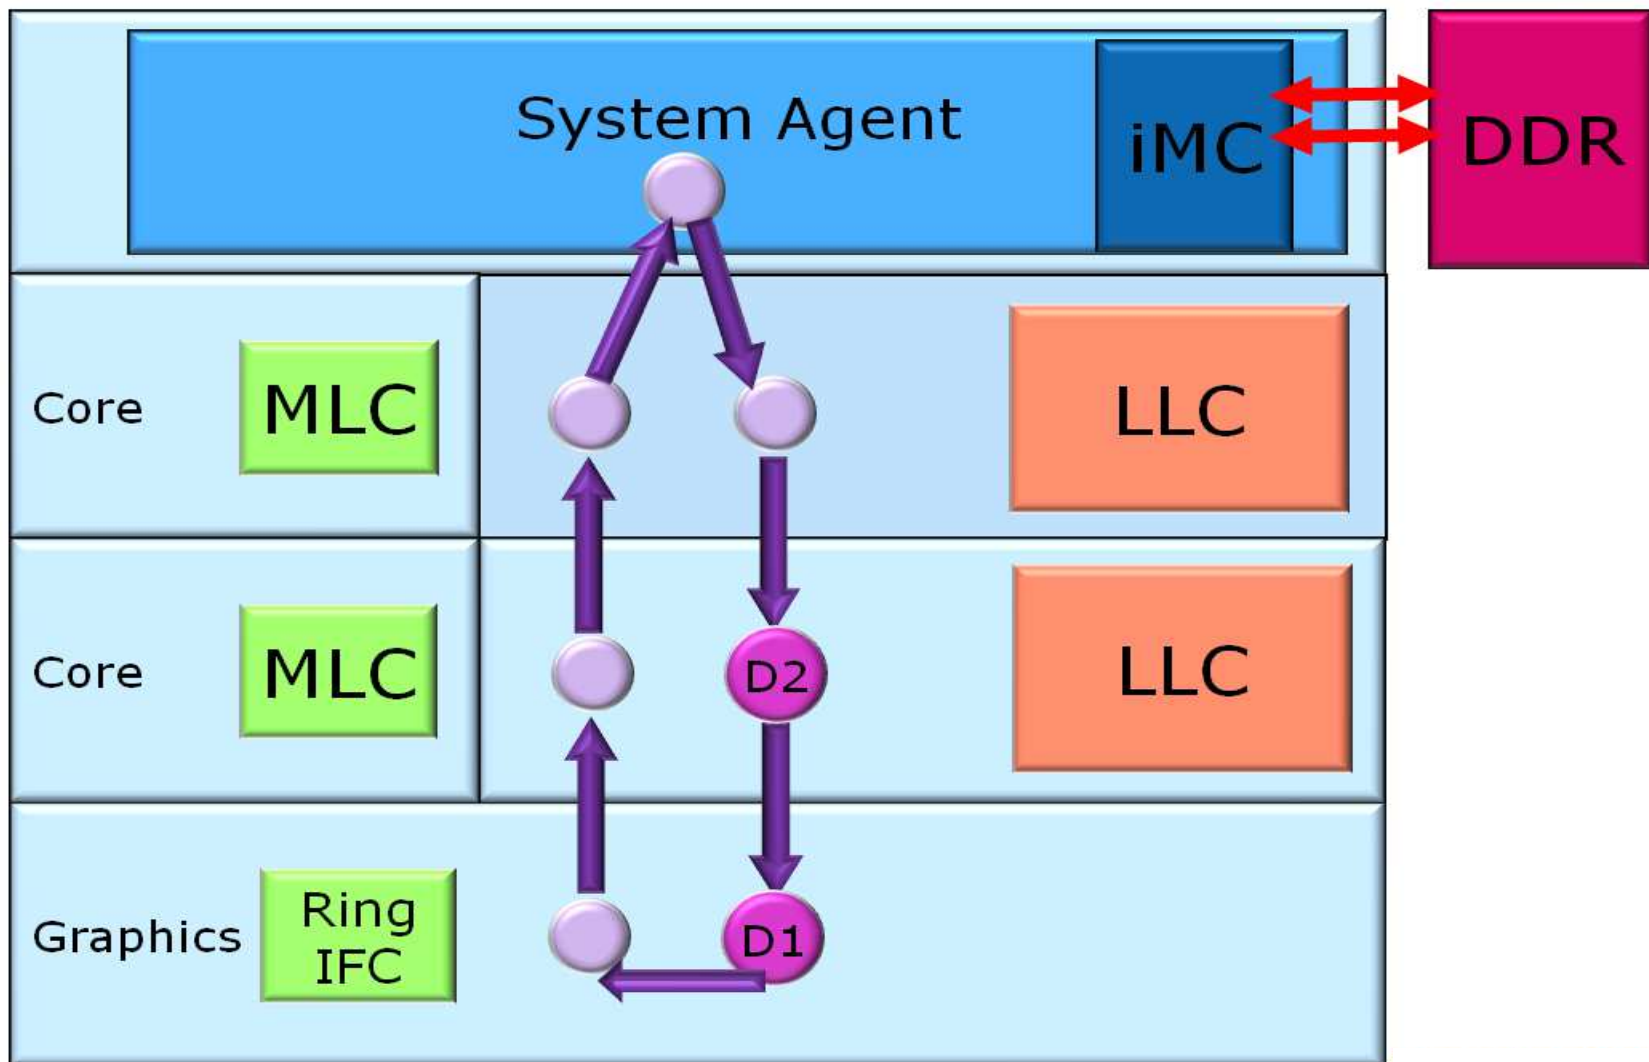
Ring Illustration: Clean LLC Hit



● Request ● Data ● Global Observation



Ring Illustration: Clean LLC Hit



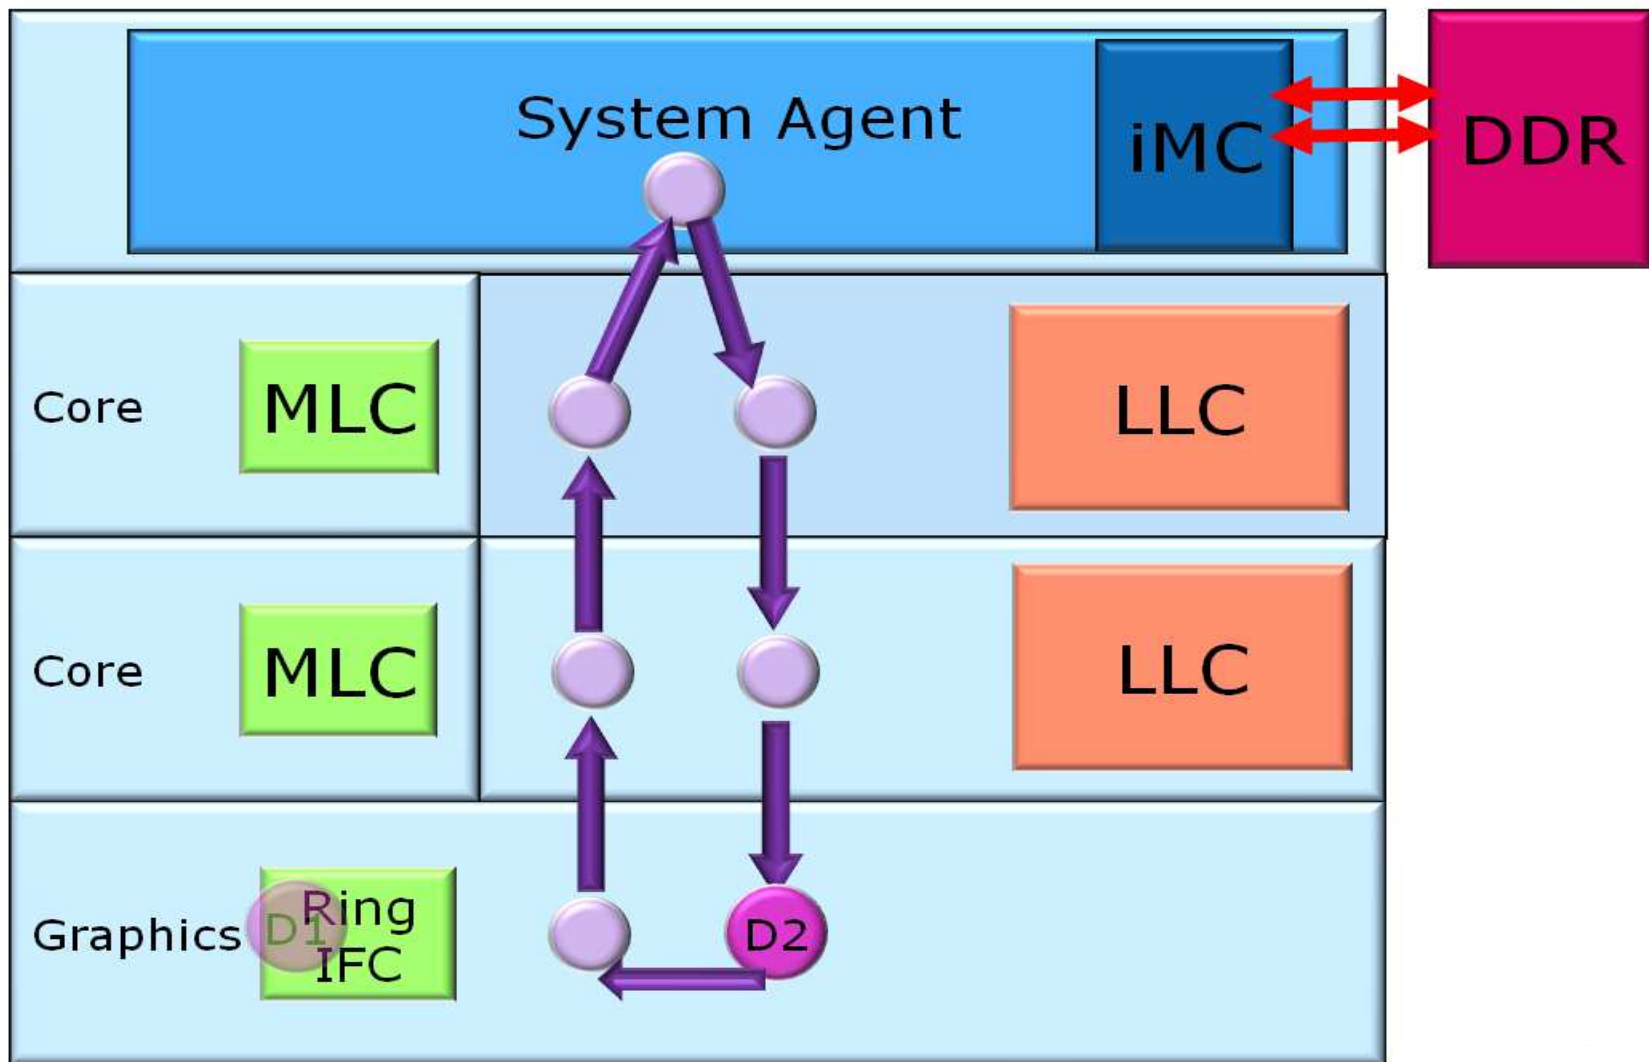
● Request

● Data

● Global Observation



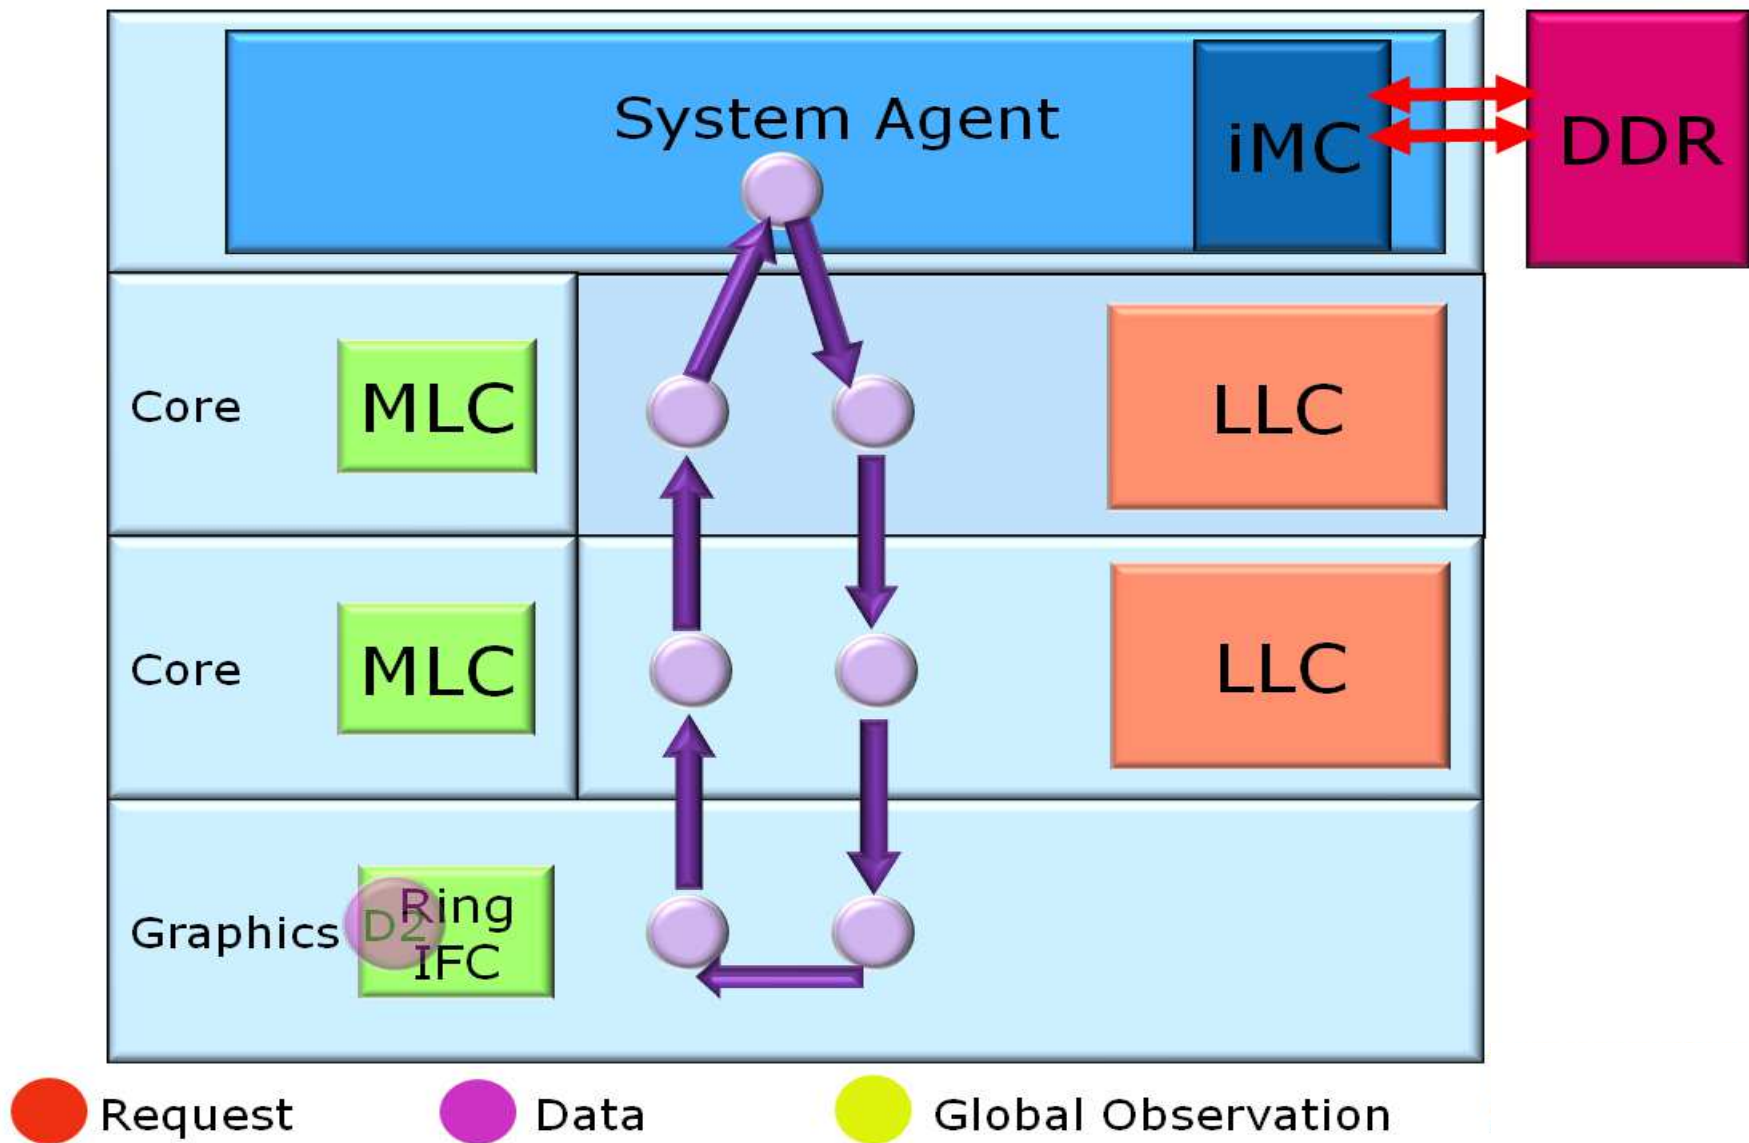
Ring Illustration: Clean LLC Hit



● Request ● Data ● Global Observation



Ring Illustration: Clean LLC Hit



Validation Challenges

- ▶ Full-Chip
 - ▶ Model size
 - ▶ Turnaround time
 - ▶ Stimuli controllability
- ▶ Knowledge and Expertise
 - ▶ The product is much more multidisciplinary
 - ▶ Stronger Architectural and Logic orientation
- ▶ Agility to late changes



Post Silicon Validation and Debug


- ▶ IOs
 - ▶ On lead process → Complex Electrical Validation (EV)
 - ▶ Probing becomes more and more complex and expensive
- ▶ Controllability and Visibility due to integration
 - ▶ Legacy FSB based debug and validation – irrelevant since FSB used to connect CPU and GMCH (now integrated within the same die)
 - ▶ Monolithic Concurrency
- ▶ Platform level power management
- ▶ Strong reliance of Hardware and Software to Time To Market (TTM)
- ▶ Customer enabling

New Design for Debug (DFD)
Strong reliance on Firmware
Enhanced customer support




Post Silicon Validation and Debug

Sandy Bridge introduces the **Generic Debug eXternal Connection (GDXC)**, a debug bus that allows monitoring the traffic between the IA cores, PG, caches and SA on the processor internal ring.



Pwr Management



GDXC allows chip, system or software debuggers to sample Ring and SA data traffic as well as protocol control signals

Time stamp

PCIe2 signals

