

# Inside Intel® Core™ Microarchitecture (Nehalem)

Ronak Singhal  
Senior Principal Engineer  
Intel Corporation  
Hot Chips 20  
August 26, 2008

# Legal Disclaimer

- INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL® PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER, AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF INTEL® PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT. INTEL PRODUCTS ARE NOT INTENDED FOR USE IN MEDICAL, LIFE SAVING, OR LIFE SUSTAINING APPLICATIONS.
- Intel may make changes to specifications and product descriptions at any time, without notice.
- All products, dates, and figures specified are preliminary based on current expectations, and are subject to change without notice.
- Intel, processors, chipsets, and desktop boards may contain design defects or errors known as errata, which may cause the product to deviate from published specifications. Current characterized errata are available on request.
- Nehalem, Merom, Wolfdale, Harpertown, Tylersburg, Penryn, Westmere, Sandy Bridge and other code names featured are used internally within Intel to identify products that are in development and not yet publicly announced for release. Customers, licensees and other third parties are not authorized by Intel to use code names in advertising, promotion or marketing of any product or services and any such use of Intel's internal code names is at the sole risk of the user
- Performance tests and ratings are measured using specific computer systems and/or components and reflect the approximate performance of Intel products as measured by those tests. Any difference in system hardware or software design or configuration may affect actual performance.
- Intel, Intel Inside, Intel Core, Intel Xeon, Intel Core2 and the Intel logo are trademarks of Intel Corporation in the United States and other countries.
- \*Other names and brands may be claimed as the property of others.
- Copyright © 2008 Intel Corporation.

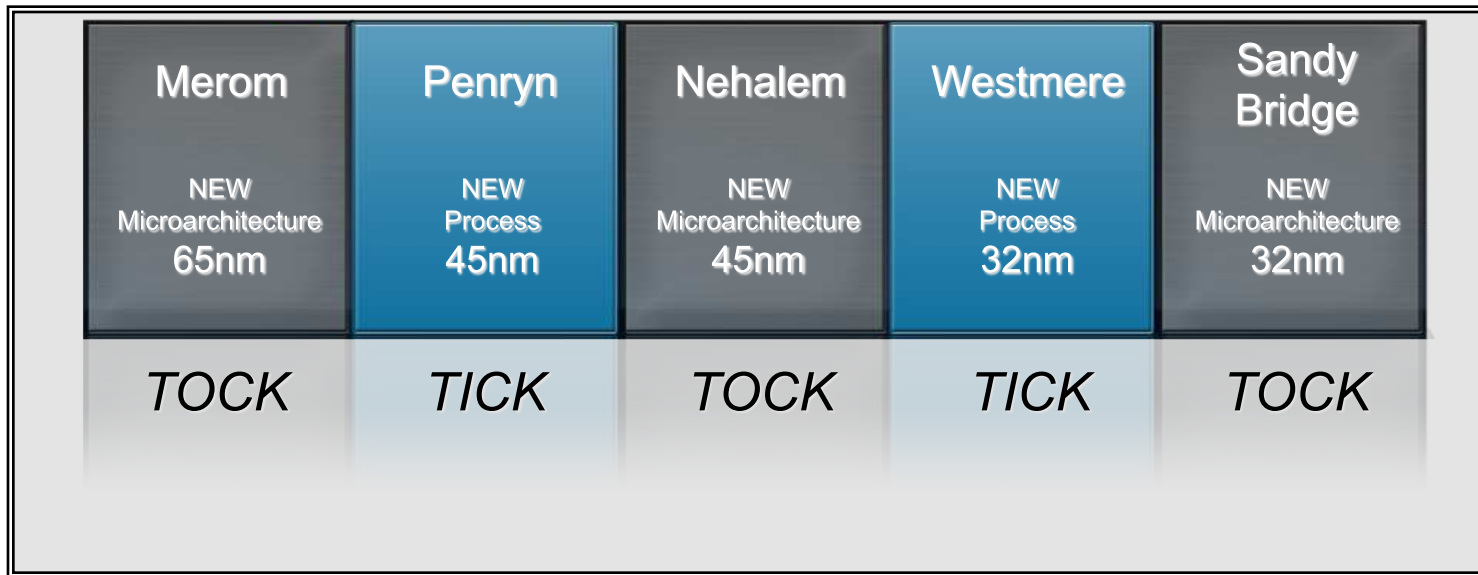
# Risk Factors

This presentation contains forward-looking statements that involve a number of risks and uncertainties. These statements do not reflect the potential impact of any mergers, acquisitions, divestitures, investments or other similar transactions that may be completed in the future. The information presented is accurate only as of today's date and will not be updated. In addition to any factors discussed in the presentation, the important factors that could cause actual results to differ materially include the following: Demand could be different from Intel's expectations due to factors including changes in business and economic conditions, including conditions in the credit market that could affect consumer confidence; customer acceptance of Intel's and competitors' products; changes in customer order patterns, including order cancellations; and changes in the level of inventory at customers. Intel's results could be affected by the timing of closing of acquisitions and divestitures. Intel operates in intensely competitive industries that are characterized by a high percentage of costs that are fixed or difficult to reduce in the short term and product demand that is highly variable and difficult to forecast. Revenue and the gross margin percentage are affected by the timing of new Intel product introductions and the demand for and market acceptance of Intel's products; actions taken by Intel's competitors, including product offerings and introductions, marketing programs and pricing pressures and Intel's response to such actions; Intel's ability to respond quickly to technological developments and to incorporate new features into its products; and the availability of sufficient supply of components from suppliers to meet demand. The gross margin percentage could vary significantly from expectations based on changes in revenue levels; product mix and pricing; capacity utilization; variations in inventory valuation, including variations related to the timing of qualifying products for sale; excess or obsolete inventory; manufacturing yields; changes in unit costs; impairments of long-lived assets, including manufacturing, assembly/test and intangible assets; and the timing and execution of the manufacturing ramp and associated costs, including start-up costs. Expenses, particularly certain marketing and compensation expenses, vary depending on the level of demand for Intel's products, the level of revenue and profits, and impairments of long-lived assets. Intel is in the midst of a structure and efficiency program that is resulting in several actions that could have an impact on expected expense levels and gross margin. Intel's results could be impacted by adverse economic, social, political and physical/infrastructure conditions in the countries in which Intel, its customers or its suppliers operate, including military conflict and other security risks, natural disasters, infrastructure disruptions, health concerns and fluctuations in currency exchange rates. Intel's results could be affected by adverse effects associated with product defects and errata (deviations from published specifications), and by litigation or regulatory matters involving intellectual property, stockholder, consumer, antitrust and other issues, such as the litigation and regulatory matters described in Intel's SEC reports. A detailed discussion of these and other factors that could affect Intel's results is included in Intel's SEC filings, including the report on Form 10-Q for the quarter ended June 28, 2008.

# Agenda

- Intel® Core™ Microarchitecture (Nehalem) Philosophy
- CPU Core Features
- New Platform Architecture
- Power Management

# Tick-Tock Development Model



<sup>1</sup>Next generation Intel® Xeon® processor (Wolfdale)  
45nm next generation Intel® Core™ microarchitecture (Penryn)  
Next generation Intel® Xeon® processor (Harpertown)  
Intel® Core™ Microarchitecture (Nehalem)  
Intel® Microarchitecture (Westmere)  
Intel® Microarchitecture (Sandy Bridge)

All products, dates, and figures are preliminary and are subject to change without notice.



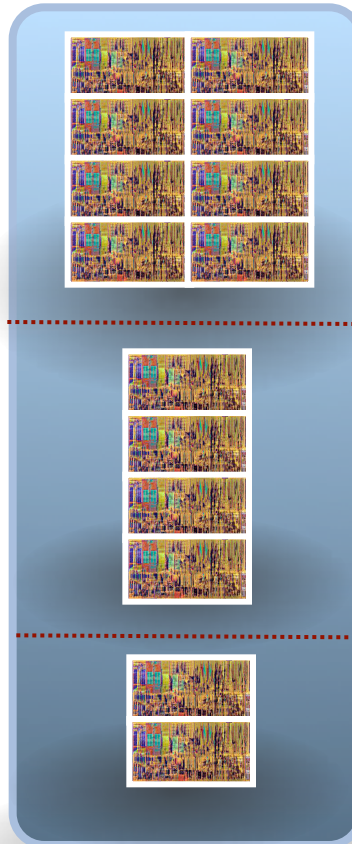

# Scalable Cores

Same core for all segments

Common software optimization

Common feature set

Intel® Core™  
Microarchitecture  
(Nehalem)  
45nm



Servers/Workstations

Energy Efficiency, Performance, Virtualization, Reliability, Capacity, Scalability

Desktop

Performance, Graphics, Energy Efficiency, Idle Power, Security

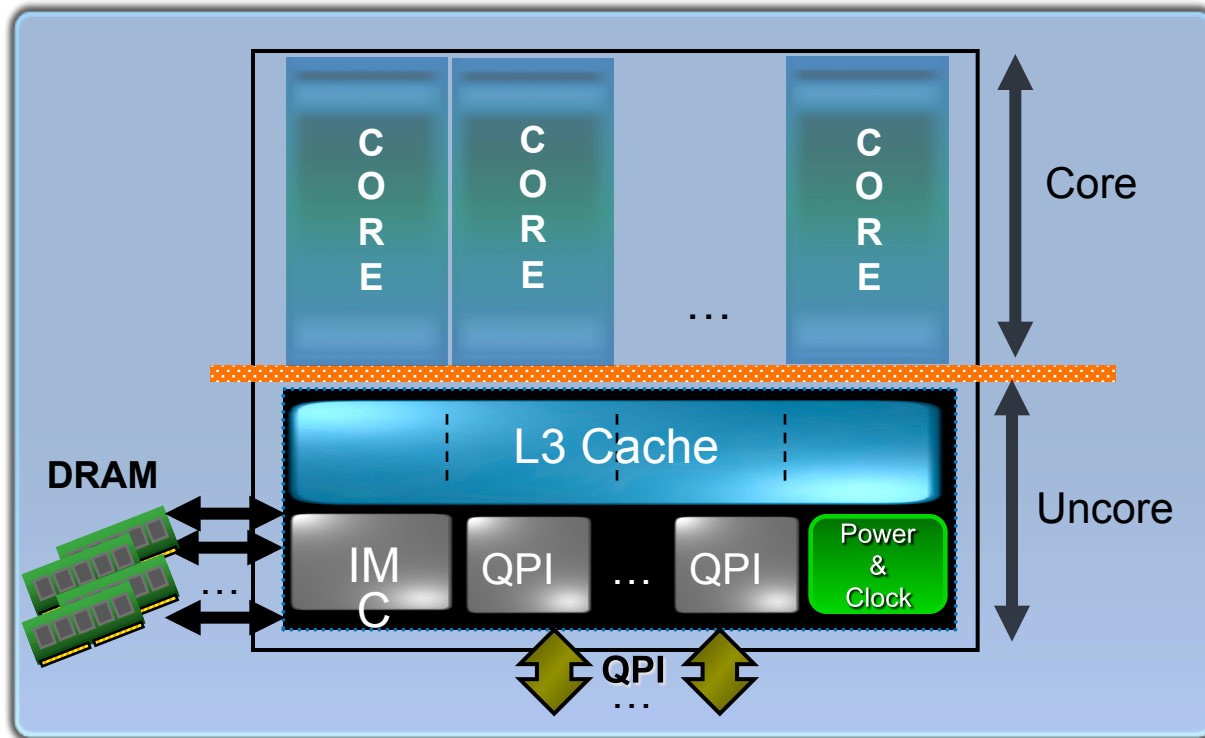
Mobile

Battery Life, Performance, Energy Efficiency, Graphics, Security


Optimized cores to meet all market segments



# Designed For Modularity



Differentiation in the “Uncore”:




Optimal price / performance / energy efficiency  
for server, desktop and mobile products





## First Intel® Core™ Microarchitecture (Nehalem)-based products: Intel® Core™ i7 processor

- Quad-core
- 731 million transistors
- 8MB 3<sup>rd</sup> Level Cache
- Simultaneous Multi-Threading
- New SSE4.2 Instructions
- Integrated DDR3 Memory Controller



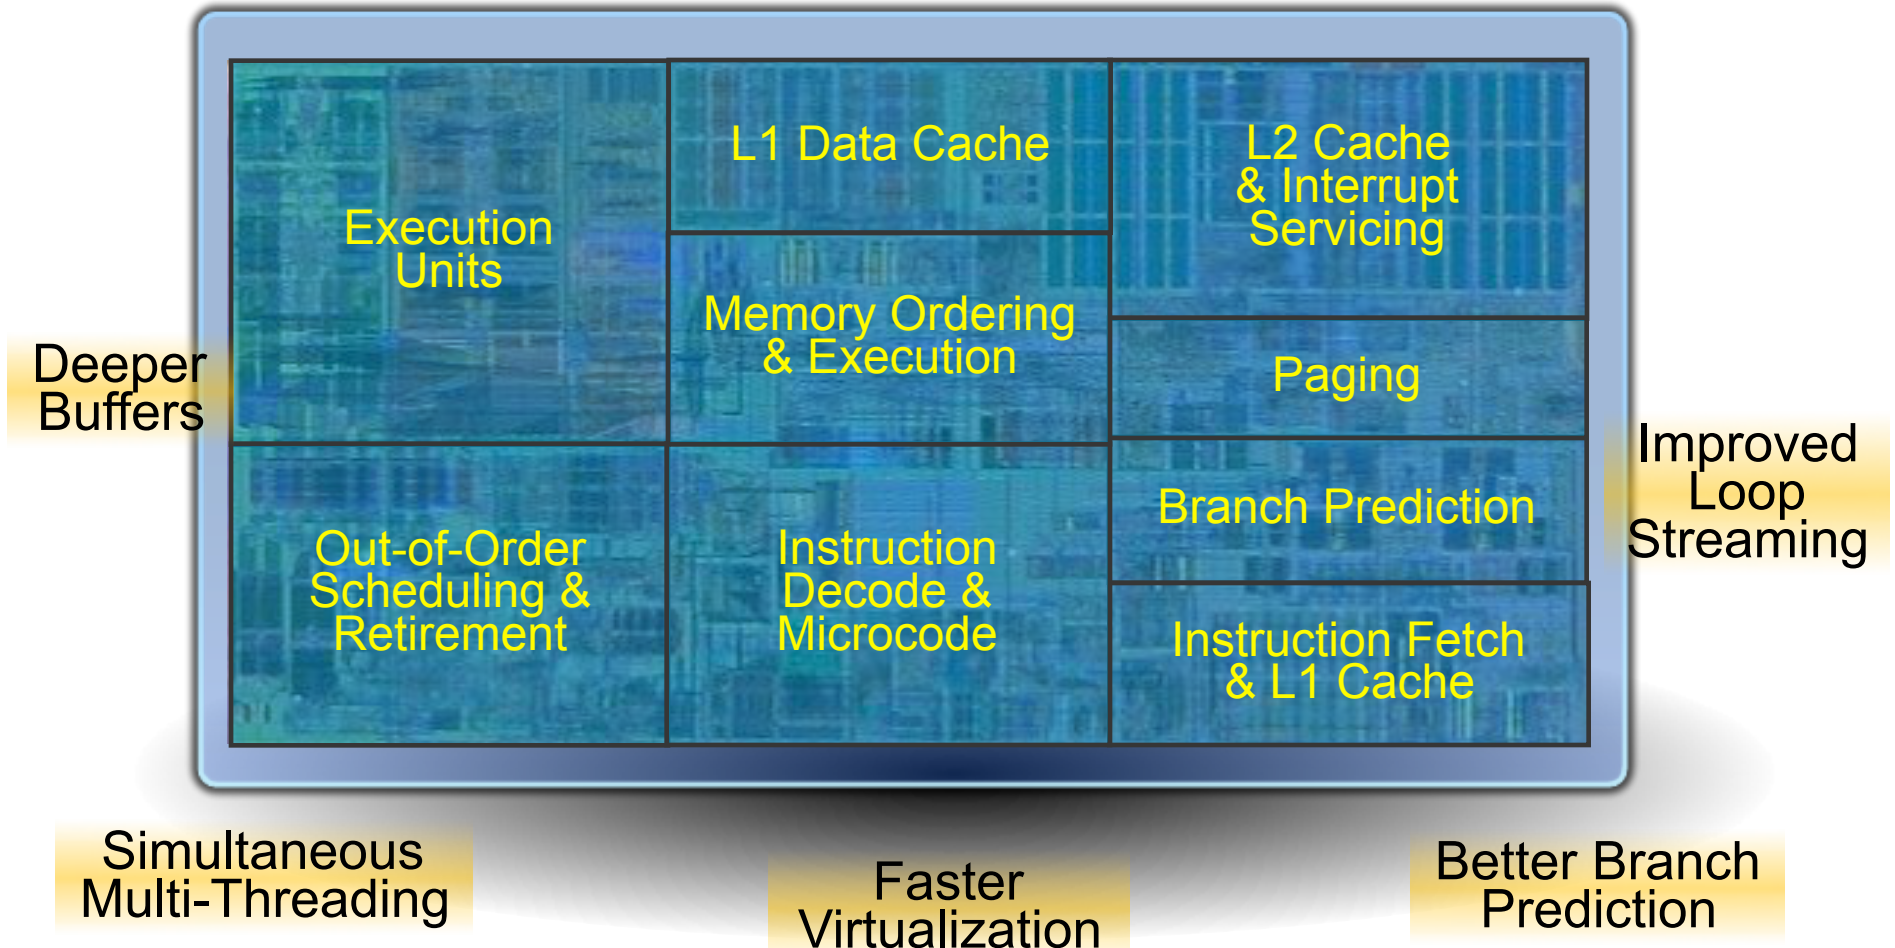


# Designed for Performance

New SSE4.2  
Instructions

Improved Lock  
Support

Additional Caching  
Hierarchy



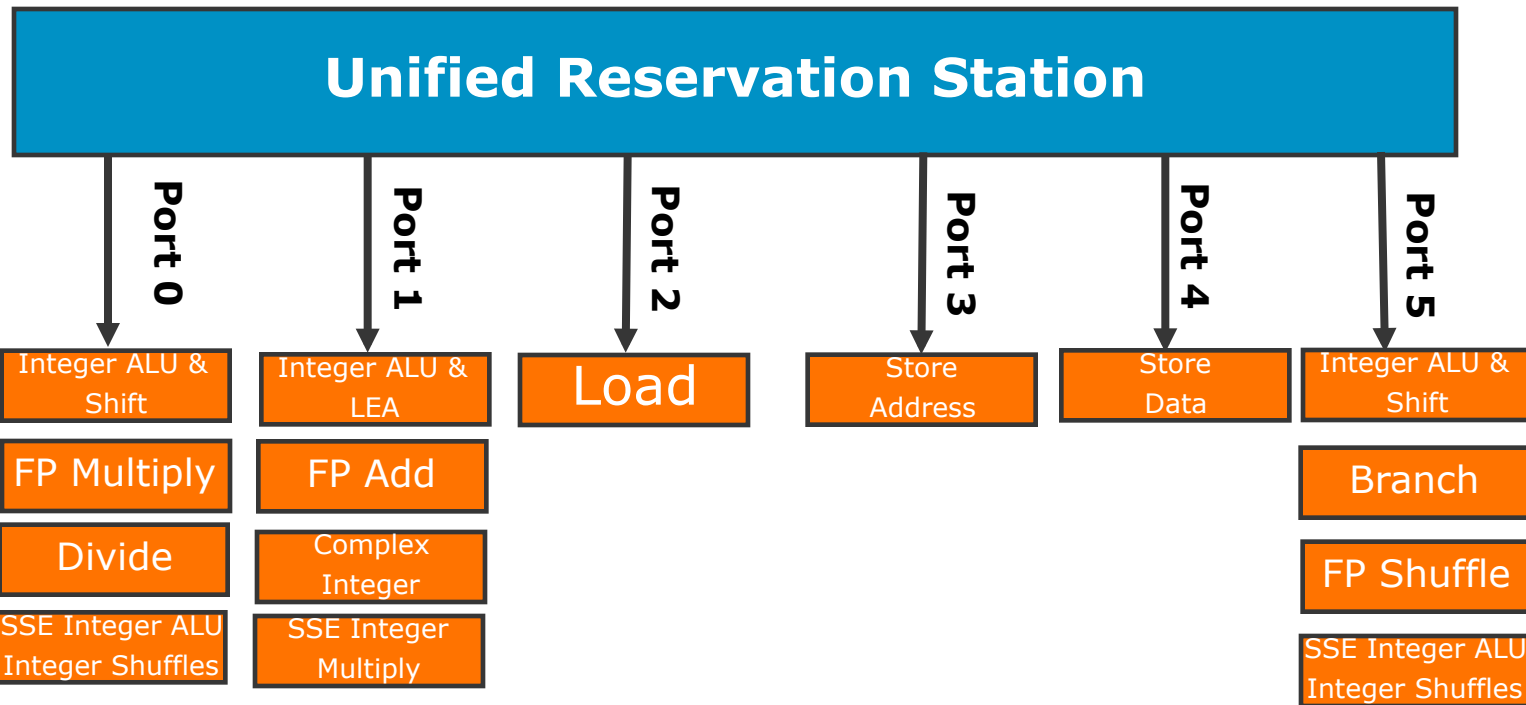
# Execution Unit Overview

## Unified Reservation Station

- Schedules operations to Execution units
- Single Scheduler for all Execution Units
- Can be used by all integer, all FP, etc.



## Execute 6 operations/cycle

- 3 Memory Operations
  - 1 Load
  - 1 Store Address
  - 1 Store Data
- 3 "Computational" Operations



# Intel® Core™ Microarchitecture (Nehalem) Processors

- Foundation is existing Intel® Core™ microarchitecture
- Focus on improving performance and power efficiency
- Key Performance Features
  - Improved Branch Prediction
    - L2 Branch Predictor
    - Advanced Renamed Return Stack Buffer
  - Increased Parallelism
    - 33% larger instruction window
  - Improved Memory Transaction Handling
    - Fast 16-byte unaligned accesses
    - New 2<sup>nd</sup> level TLB
    - Improved Lock Handling




<sup>1</sup>Intel® Pentium® M processor (formerly Dothan)  
Intel® Core™ microarchitecture (formerly Merom)  
Intel® Core™ microarchitecture (Nehalem)

Intel® Pentium® 4 processor  
Intel® Core™ Duo processor  
Intel® Core™ microarchitecture (Nehalem)-based processor

# New 3-level Cache Hierarchy

- 1<sup>st</sup> level caches
  - 32kB Instruction cache
  - 32kB Data Cache
  - Support more cache misses in parallel
- 2<sup>nd</sup> level Unified Cache
  - 256 kB per core
  - Designed for very low latency
- 3<sup>rd</sup> level Unified Cache
  - Size depends on # of cores
  - **Inclusive** cache
  - Core valid bits for minimizing snoop traffic




## ***Why Inclusive?***

- Cache acts as a snoop filter
- Only snoops cores when necessary
- Provides **Scalability**
- Minimizes **Latency**

# Other Key Features

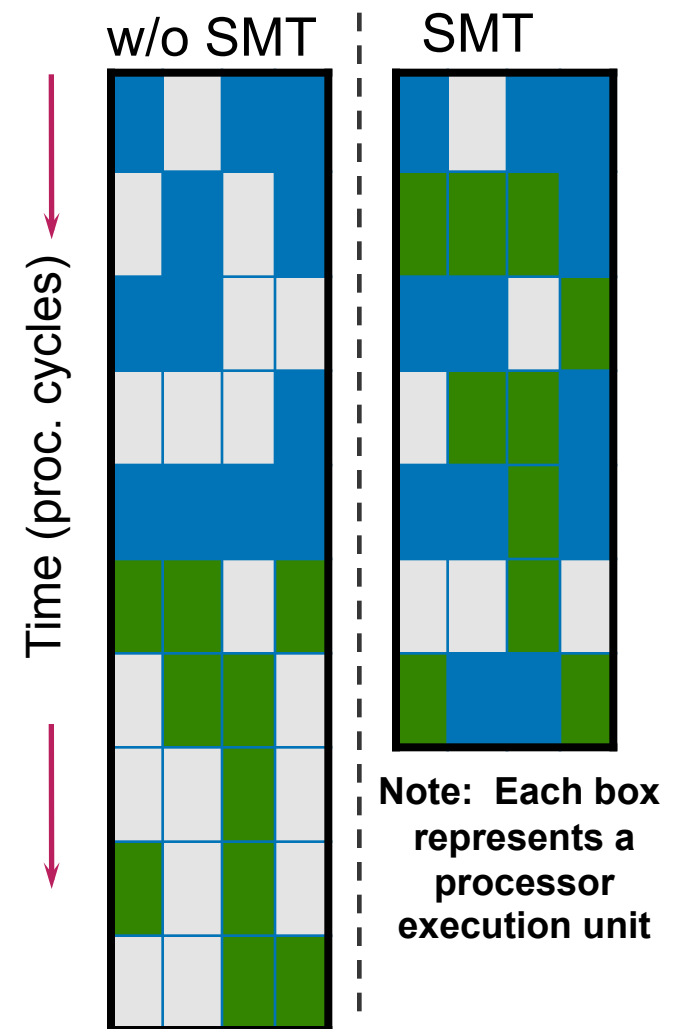
- New instructions (SSE4.2)
  - XML/String/text processing
  - CRC32
  - POPCNT
- Virtualization
  - Best virtualized performance starts w /best native performance
  - Goals:
    - Reduce # of transitions between host/guest
    - Reduce latency of transitions between host/guest
  - Features:
    - Microarchitecture: 40% reduction in “round-trip latency” vs. prior products
    - Architecture
      - Extended Page Table (EPT): Eliminate exists from guest due to page table management
      - VPID: Eliminate TLB flushes on host/guest transitions



<sup>1</sup>Intel® Core™ microarchitecture (formerly Merom)  
45nm next generation Intel® Core™ microarchitecture (Penryn)  
Intel® Core™ microarchitecture (Nehalem)

# Intel® Hyper-Threading Technology

- Also known as Simultaneous Multi-Threading (SMT)
  - Run 2 threads at the same time per core
- Take advantage of 4-wide execution engine
  - Keep it fed with multiple threads
  - Hide latency of a single thread
- Most **power efficient** performance feature
  - Very low die area cost
  - Can provide significant performance benefit depending on application
  - Much more efficient than adding an entire core
- Intel® Core™ microarchitecture (Nehalem) advantages
  - Larger caches
  - Massive memory BW




**Simultaneous multi-threading enhances performance and energy efficiency**





# SMT Performance Chart



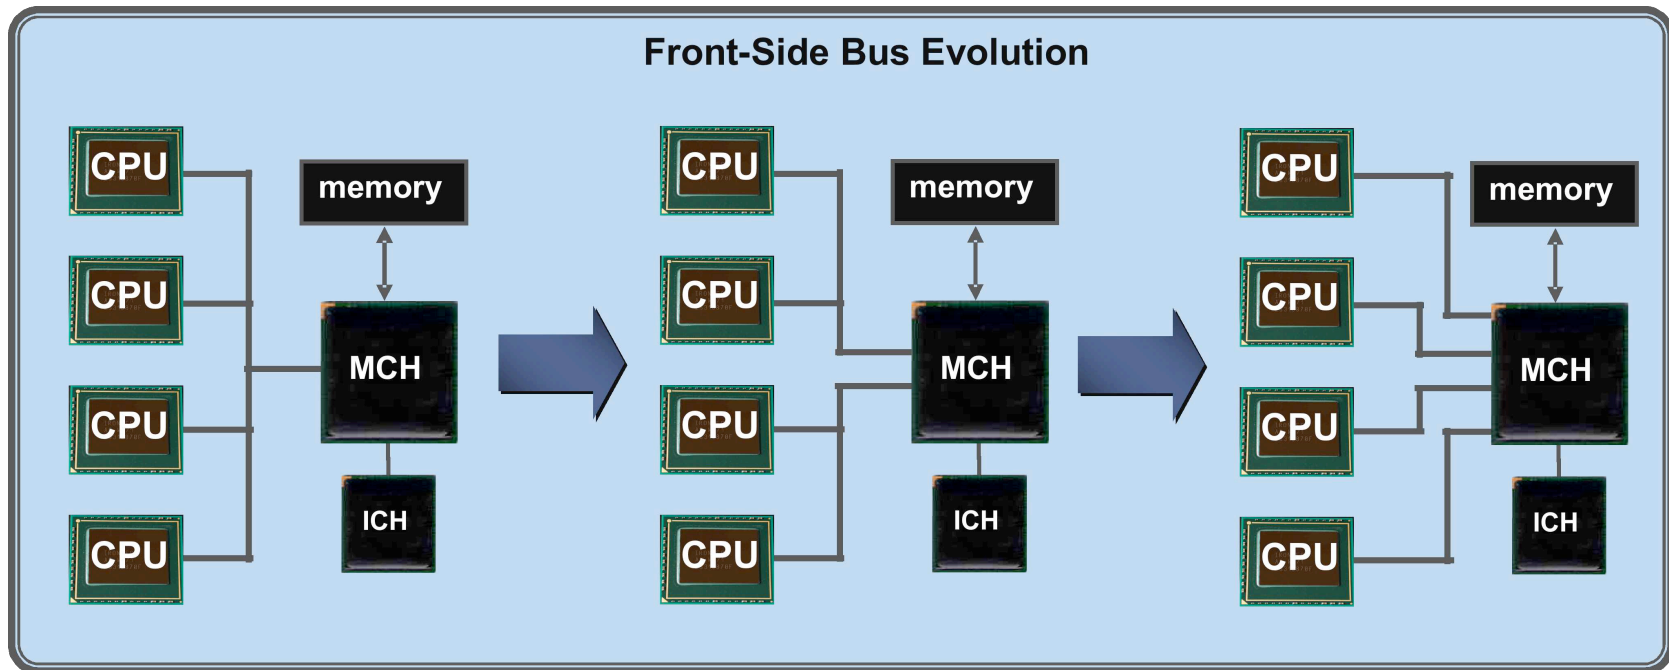
Floating Point is based on SPECfp\_rate\_base2006\* estimate  
 Integer is based on SPECint\_rate\_base2006\* estimate

SPEC, SPECint, SPECfp, and SPECrate are trademarks of the Standard Performance Evaluation Corporation.  
 For more information on SPEC benchmarks, see: <http://www.spec.org>

Source: Intel. Configuration: pre-production Intel® Core™ i7 processor with 3 channel DDR3 memory. Performance tests and ratings are measured using specific computer systems and / or components and reflect the approximate performance of Intel products as measured by those tests. Any difference in system hardware or software design or configuration may affect actual performance. Buyers should consult other sources of information to evaluate the performance of systems or components they are considering purchasing. For more information on performance tests and on the performance of Intel products, visit <http://www.intel.com/performance/>




# Today's Platform Architecture



# New Platform Architecture

- Integrated Memory Controller
  - Native DDR3
  - Massive memory **bandwidth**
  - Very **low memory latency**
- Intel® QuickPath Interconnect (Intel® QPI)
  - New point-to-point interconnect
    - Socket to socket
    - Socket to chipset
  - Build **scalable** solutions
  - High Bandwidth, low latency
  - Speeds up to 6.4 GT/sec initially
    - **25.6 GB/sec** per link

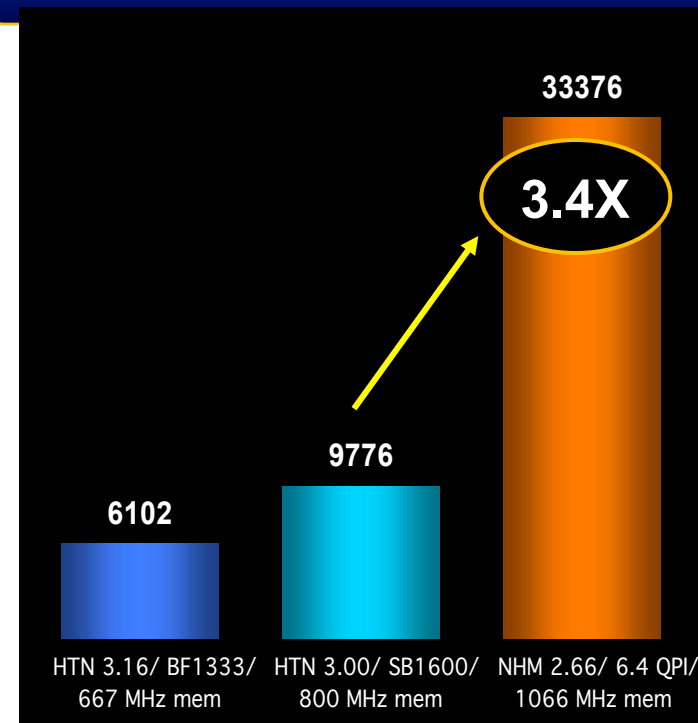


**Significant performance leap from new platform**

# Memory Bandwidth – Initial Intel® Core™ Microarchitecture (Nehalem) Products

- 3 memory channels per socket
- $\geq$  DDR3-1066 at launch
- Massive **memory BW**
- **Scalability**
  - Design IMC and core to take advantage of BW
  - Allow performance to scale with cores
    - Core enhancements
      - Support more cache misses per core
      - Aggressive hardware prefetching w/ throttling enhancements
    - Example IMC Features
      - Independent memory channels
      - Aggressive Request Reordering

Stream Bandwidth – Mbytes/Sec (Triad)

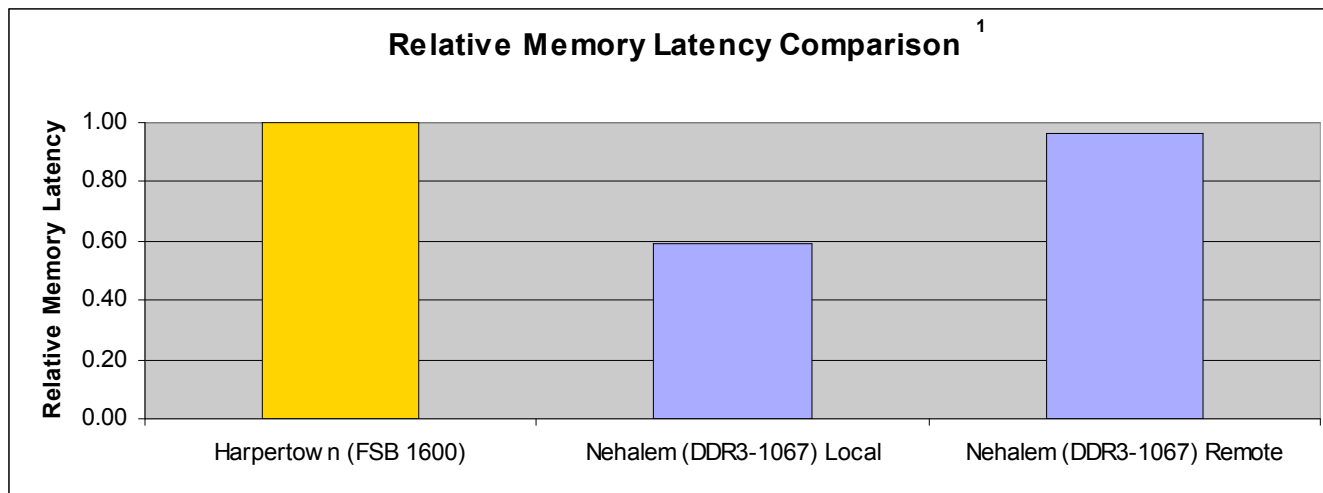


Source: Intel Internal measurements – August 2008

**Massive memory BW provides performance and scalability**

# 2-socket Memory Latency Comparison

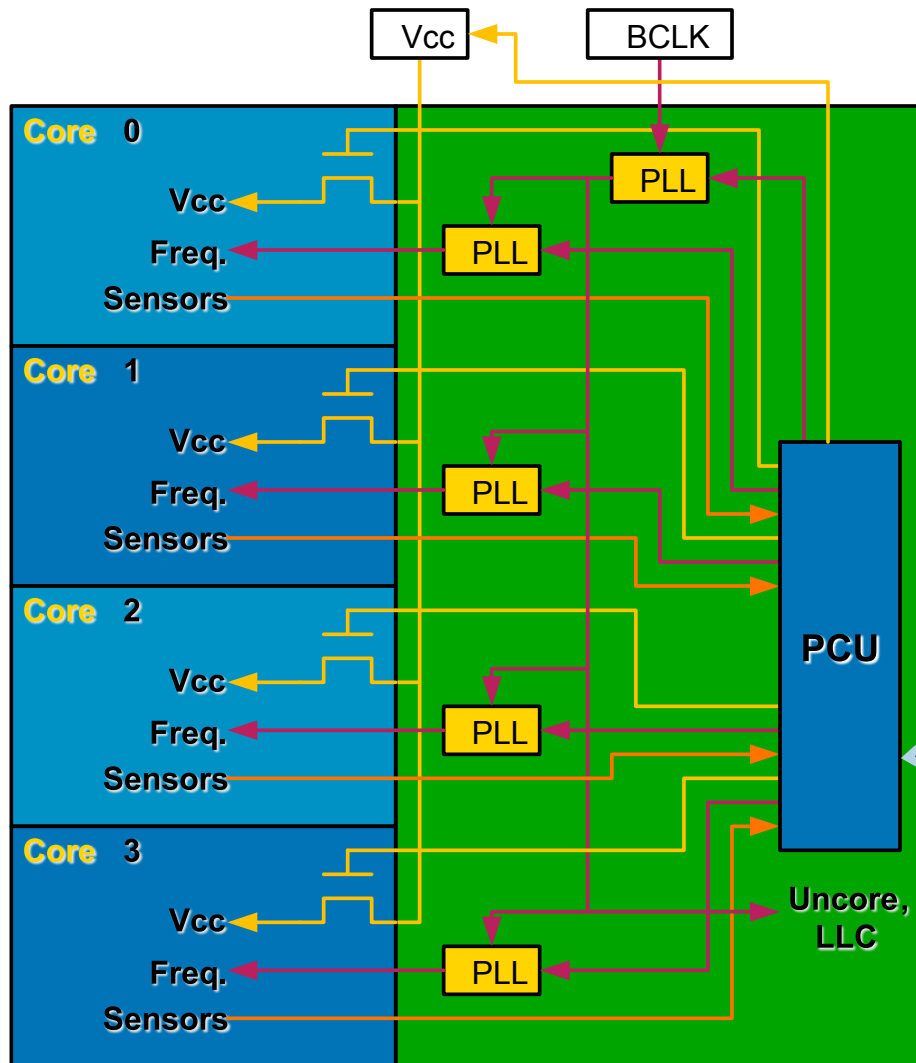
- **Low memory latency** critical to high performance
- Design integrated memory controller for low latency
- Design cache hierarchy for quick snoop response time
- NUMA: Need to optimize both local and remote memory latency
- Intel® Core™ microarchitecture (Nehalem) delivers:
  - Huge reduction in local memory latency
  - Even remote memory latency is fast
- Effective memory latency depends per application/OS
  - NHM has lower latency regardless of mix of local/remote traffic



<sup>1</sup>Next generation Intel® Xeon® processor (Harpertown)  
Intel® Core™ microarchitecture (Nehalem)



# Power Control Unit

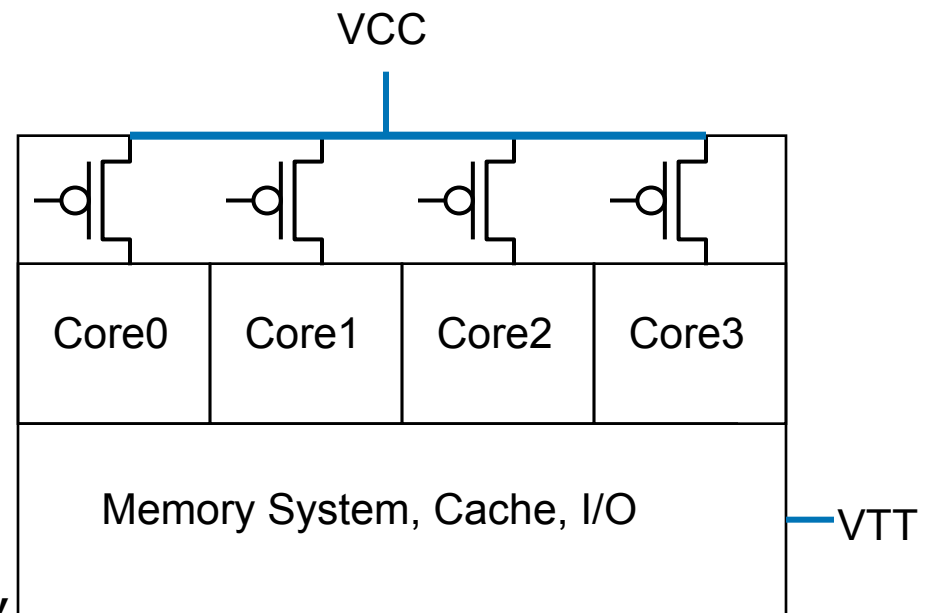


Integrated proprietary microcontroller  
Shifts control from hardware to embedded firmware  
Real time sensors for temperature, current, power  
Flexibility enables sophisticated algorithms, tuned for current operating conditions



# Integrated Power Gate

- Integrated power switch between VR output and core voltage supply
  - Very low on-resistance
  - Very high off-resistance
  - Much faster voltage ramp than external VR
- Enables per core C6 state
  - Individual cores transition to  $\sim 0$  power state
  - Transparent to other cores, platform, software



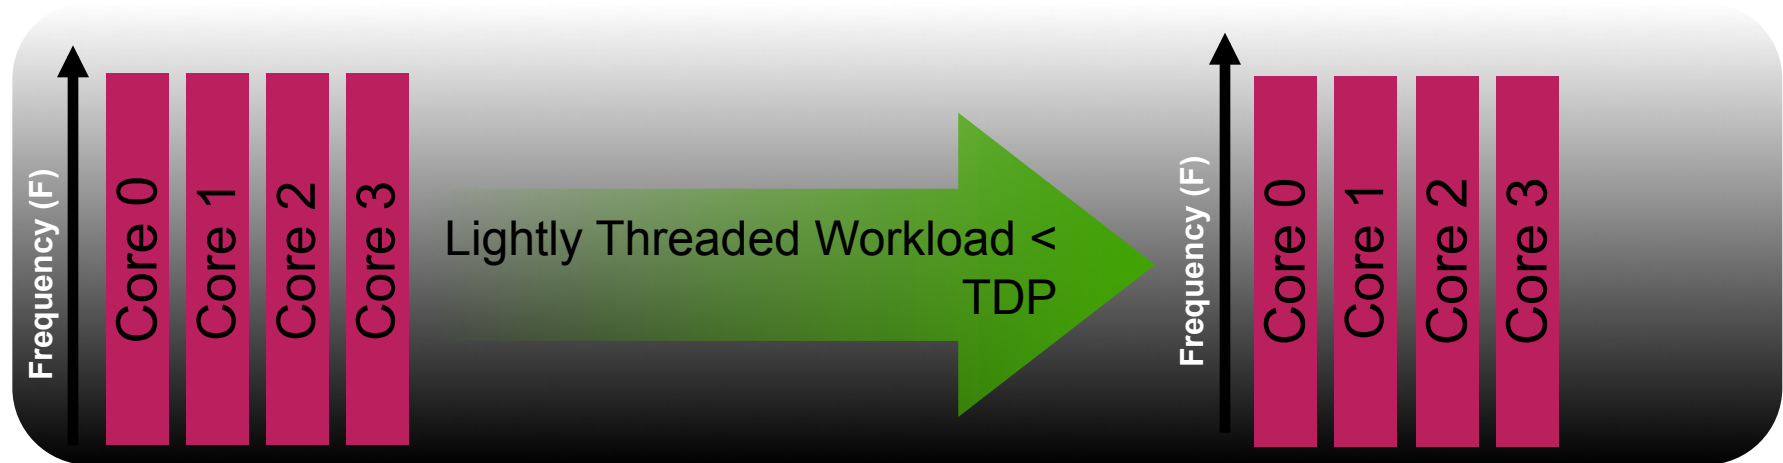
***Close collaboration with process technology to optimize device characteristics***

# Intel® Core™ Microarchitecture (Nehalem) Turbo Mode

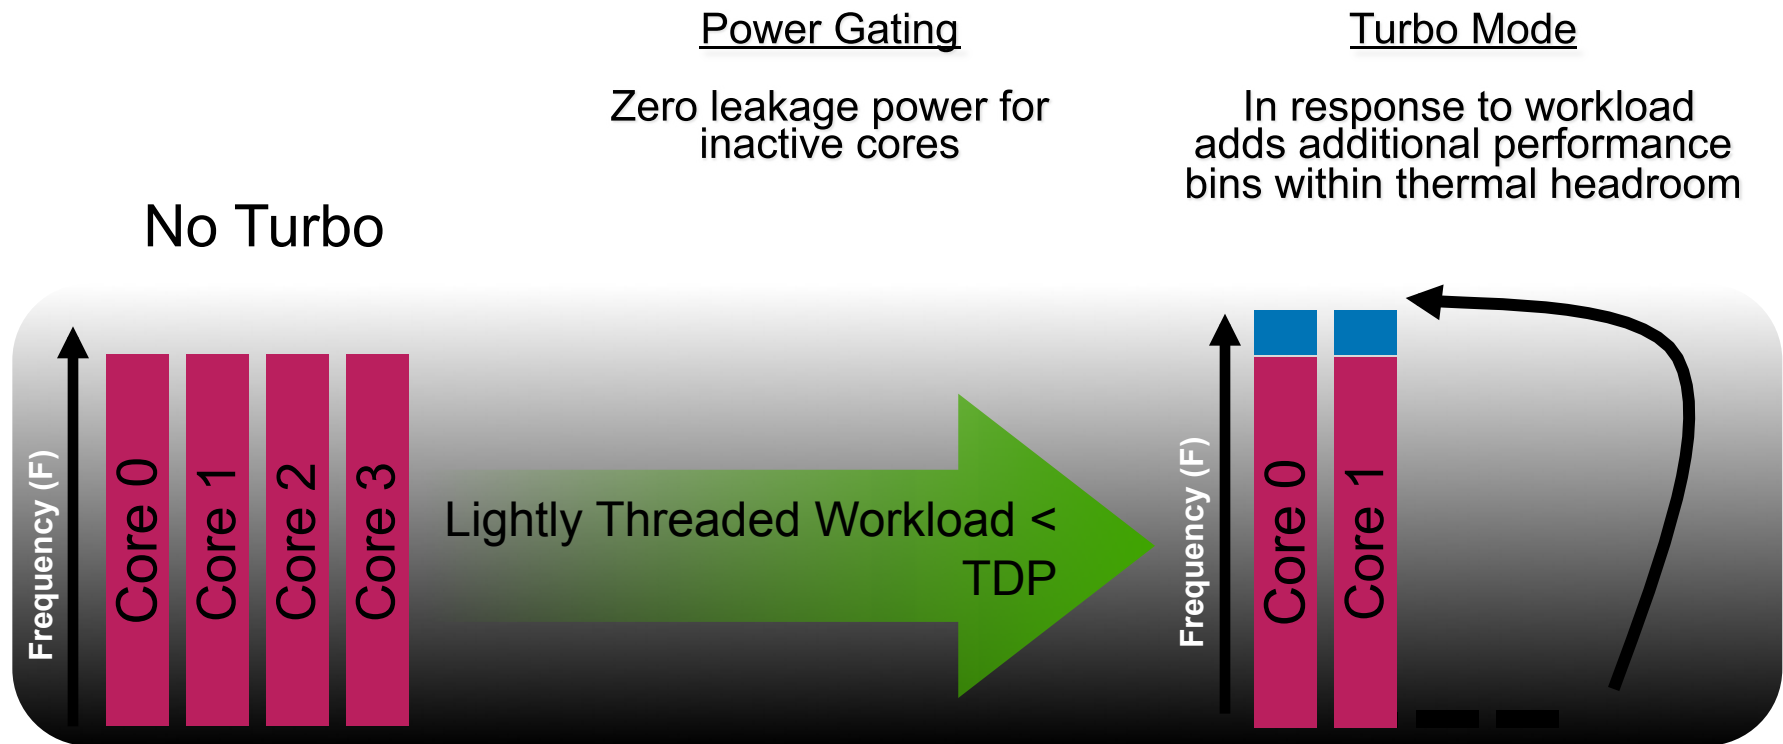
## Power Gating

Zero leakage power for  
inactive cores

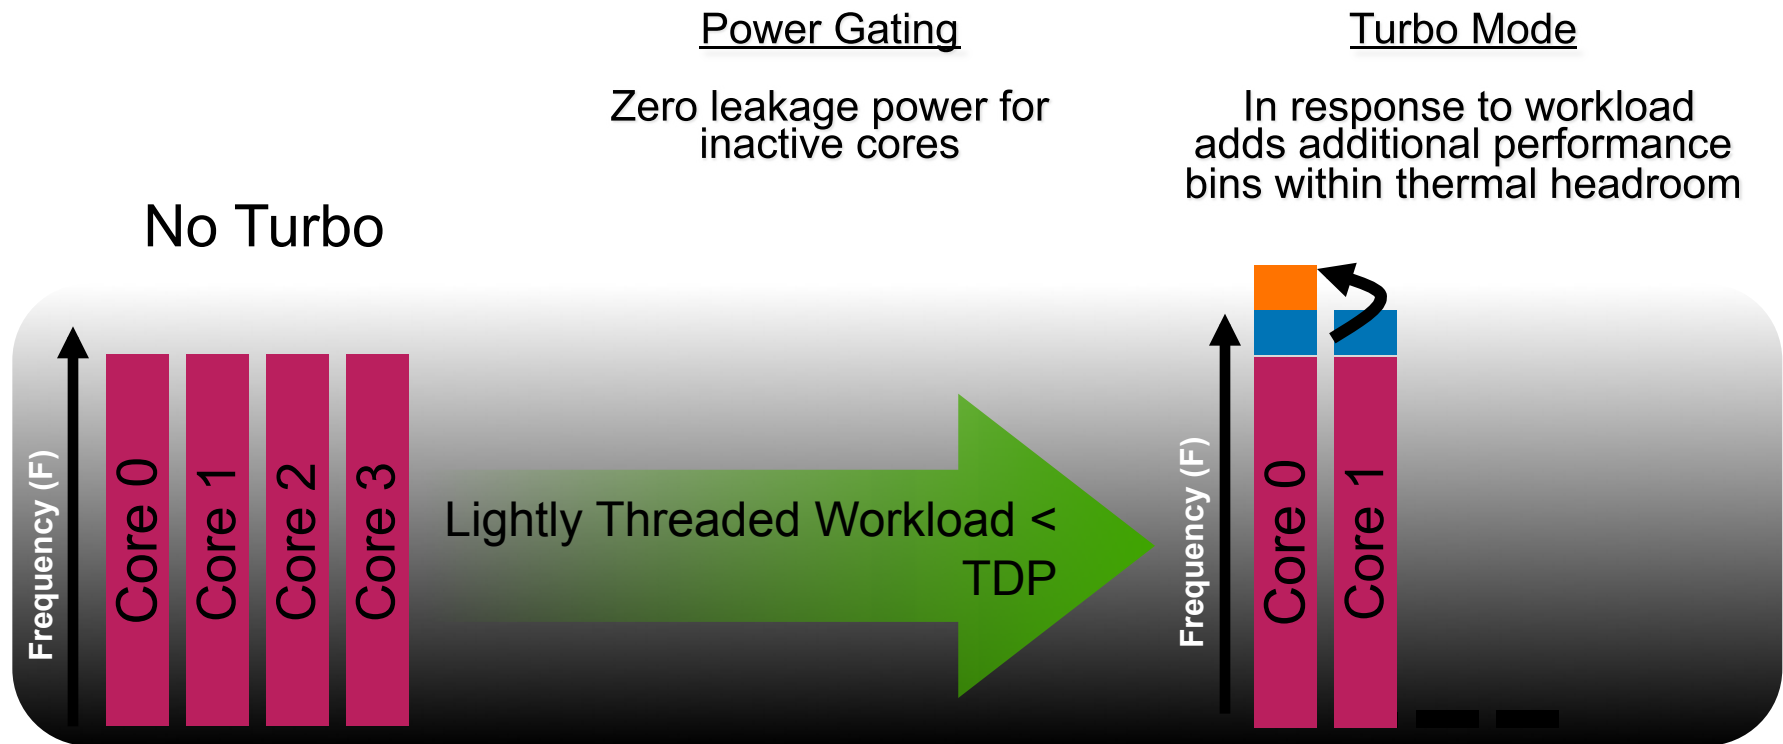
No Turbo



# Intel® Core™ Microarchitecture (Nehalem) Turbo Mode



# Intel® Core™ Microarchitecture (Nehalem) Turbo Mode



Dynamically Delivering Optimal Performance and Energy Efficiency



# Summary

- Intel® Core™ microarchitecture (Nehalem)
  - The 45nm Tock
- Designed for
  - ***Power Efficiency***
  - ***Scalability***
  - ***Performance***
- Key Innovations
  - Enhanced Processor Core
  - Brand New Platform Architecture
  - Sophisticated Power Management