


Cache-Aware Roofline Model: Performance, Power and Energy-Efficiency

Aleksandar Ilić,


Frederico Pratas and Leonel Sousa



INESC-ID/IST,
Universidade de Lisboa, Portugal



Introduction



THIS IS ALL ABOUT BUILDING THE ROOFS



HITTING THE ROOF GETTING THE MAXIMUM ARCHITECTURES VS. APPLICATIONS


- Multi-core CPUs, accelerators, diverse compute capabilities, complex memory hierarchy...
- Complex applications, different compute and memory requirements
- Optimize and characterize applications for specific architectures

MAXIMUM ATTAINABLE PERFORMANCE VS. POWER/ENERGY
CONSUMPTION


⇒ HOW FAR CAN WE GO?

⇒ WHAT ARE THOSE MAXIMUMS FOR PERFORMANCE, POWER, ENERGY, EFFICIENCY...?

STATIC DAG-BASED SCHEDULING

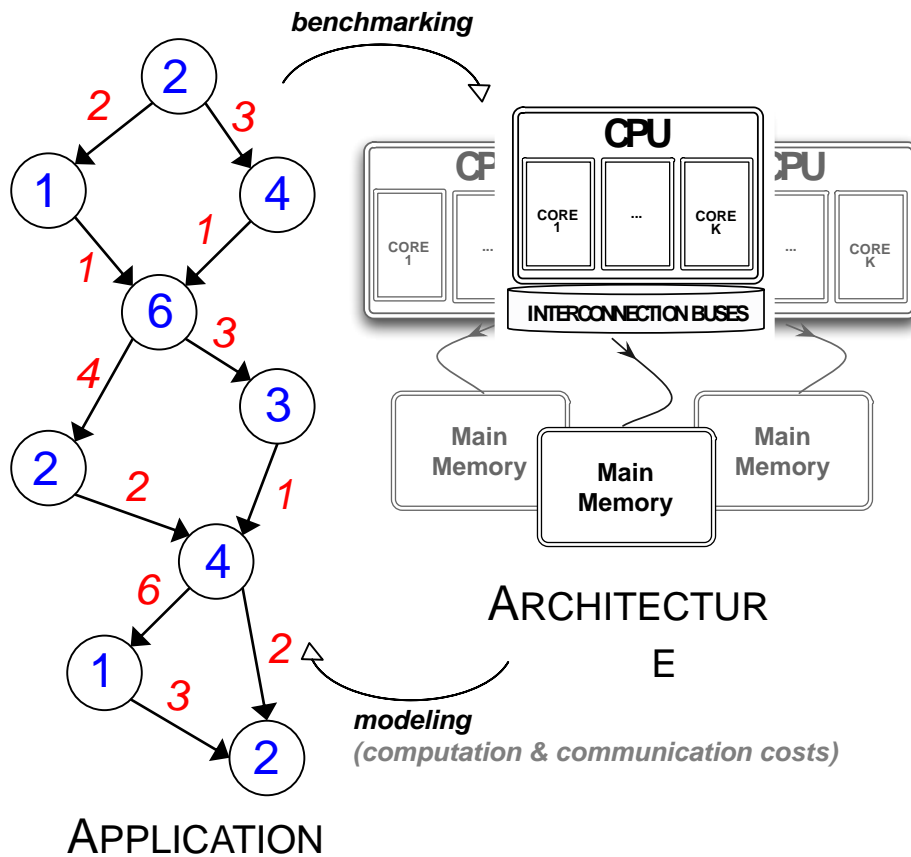


APPLICATION

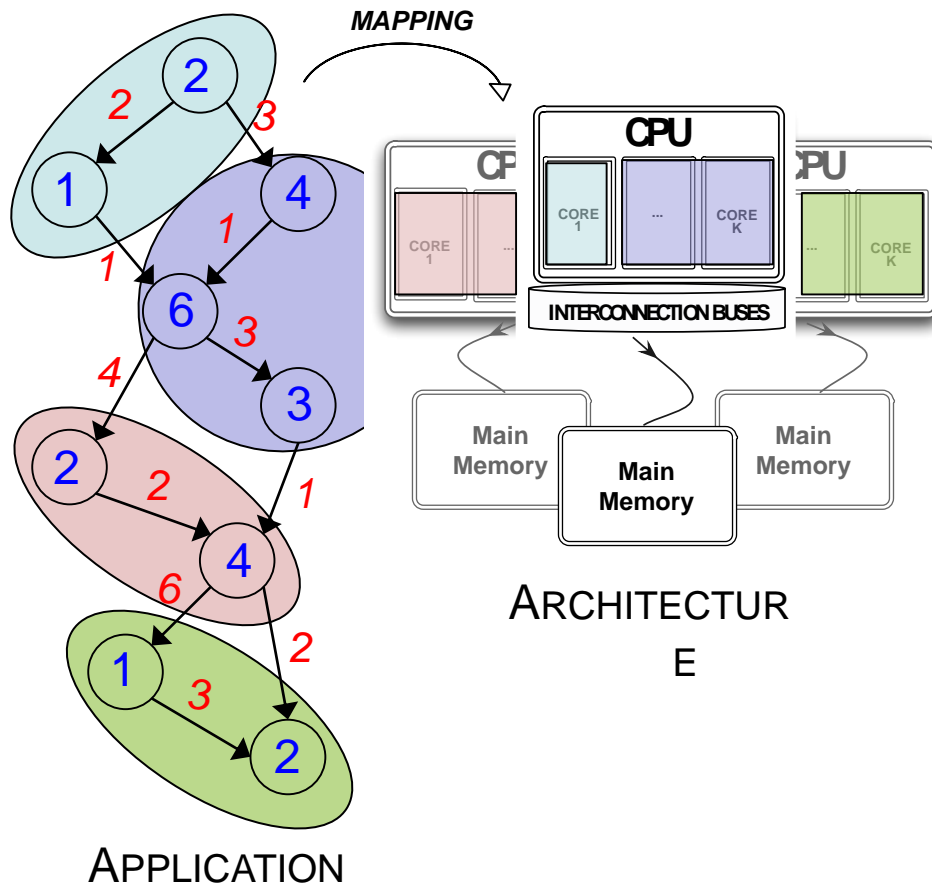


ARCHITECTURE

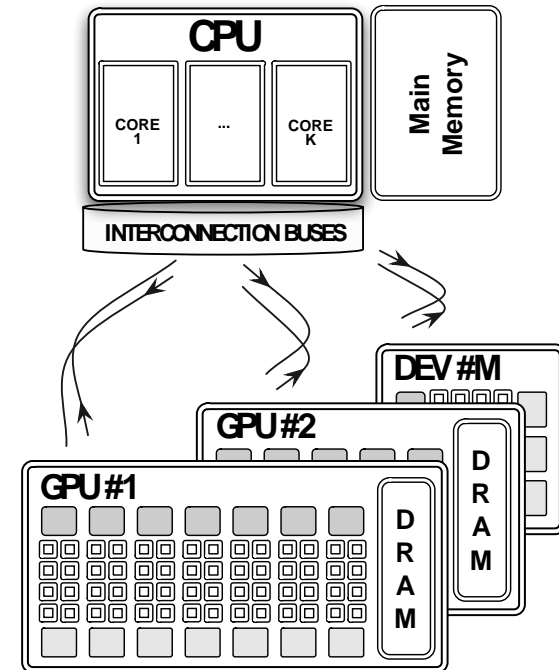



STATIC DAG-BASED SCHEDULING



STATIC DAG-BASED SCHEDULING







STATIC DAG-BASED SCHEDULING DYNAMIC DLT-BASED SCHEDULING







ARCHITECTURE

STATIC DAG-BASED SCHEDULING DYNAMIC DLT-BASED SCHEDULING



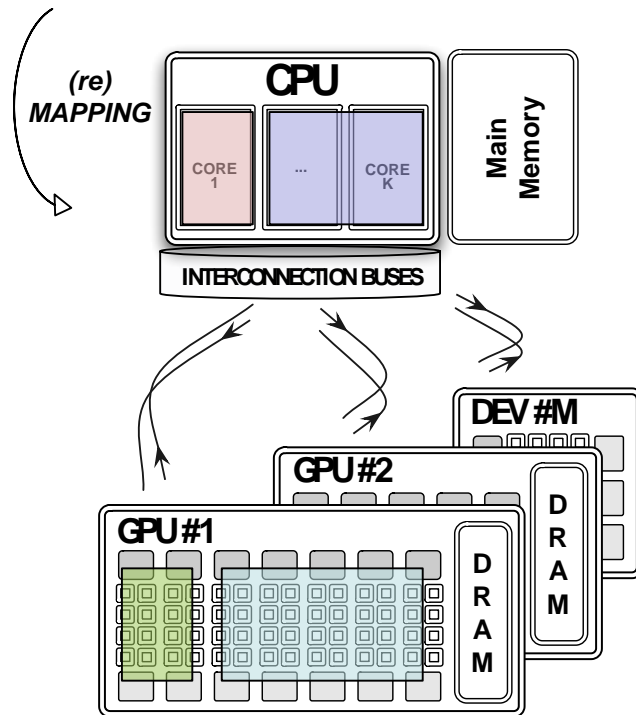



ARCHITECTURE

STATIC DAG-BASED SCHEDULING DYNAMIC DLT-BASED SCHEDULING



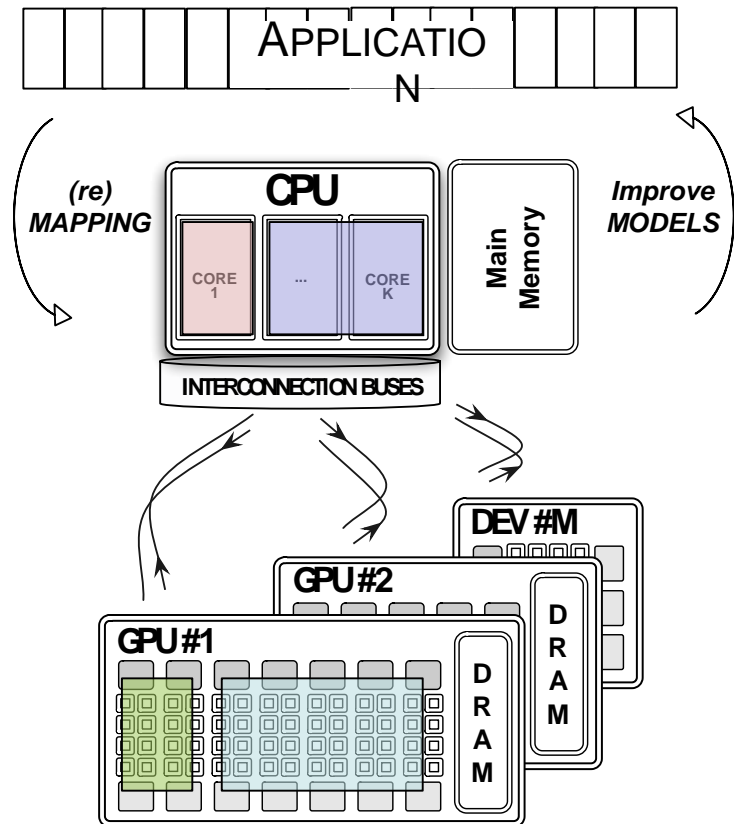
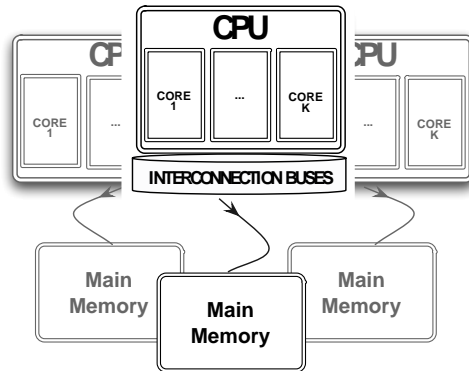

ARCHITECTUR

STATIC DAG-BASED SCHEDULING DYNAMIC DLT-BASED SCHEDULING



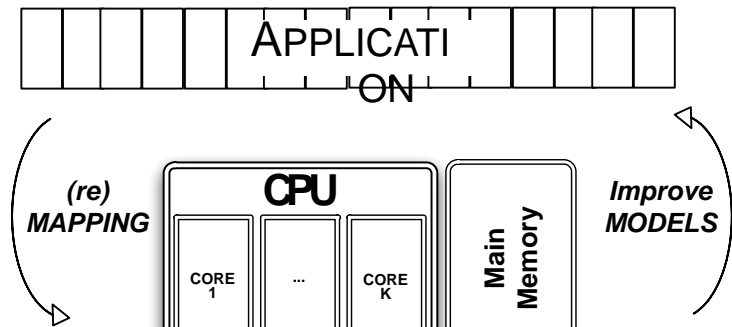
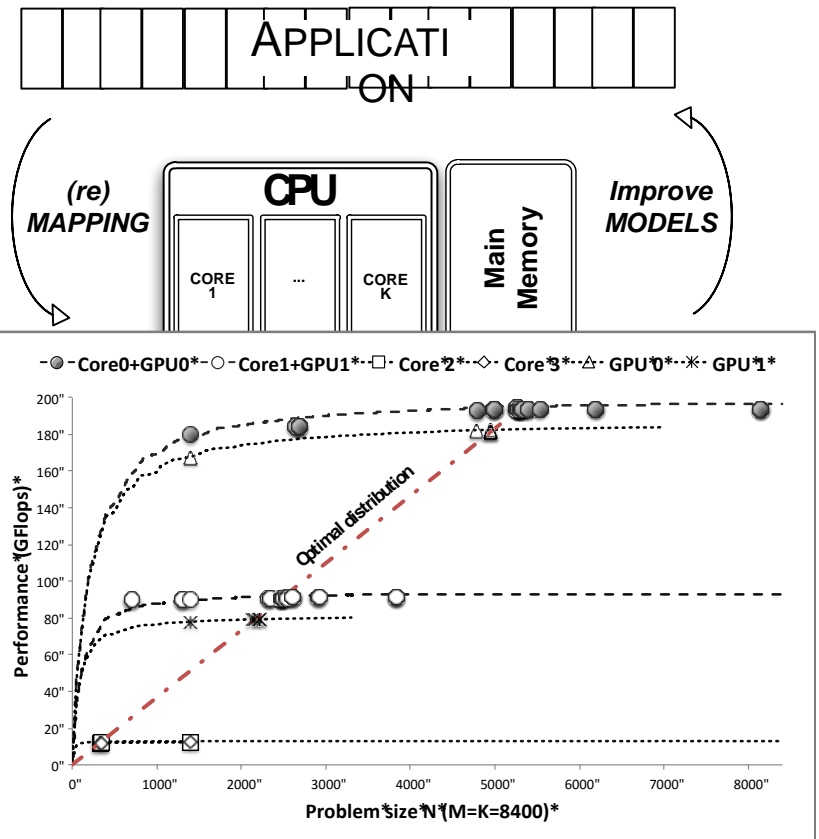


ARCHITECTURE

STATIC DAG-BASED SCHEDULING DYNAMIC DLT-BASED SCHEDULING




ARCHITECTUR

STATIC DAG-BASED SCHEDULING DYNAMIC DLT-BASED SCHEDULING



ARCHITECTUR

STATIC DAG-BASED SCHEDULING DYNAMIC DLT-BASED SCHEDULING



REACHING THE MAXIMUM IS LIMITED BY

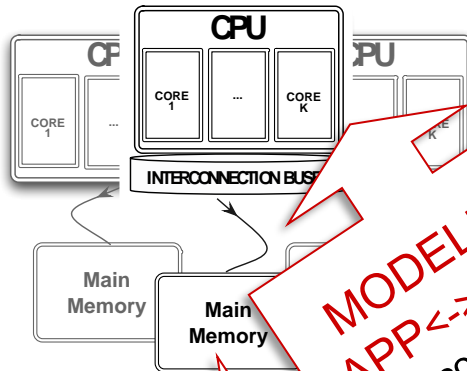

- Application characteristics and demands
- Architecture capabilities to satisfy them
- Trade-offs, assumptions ...

HOW FAR CAN WE GO?

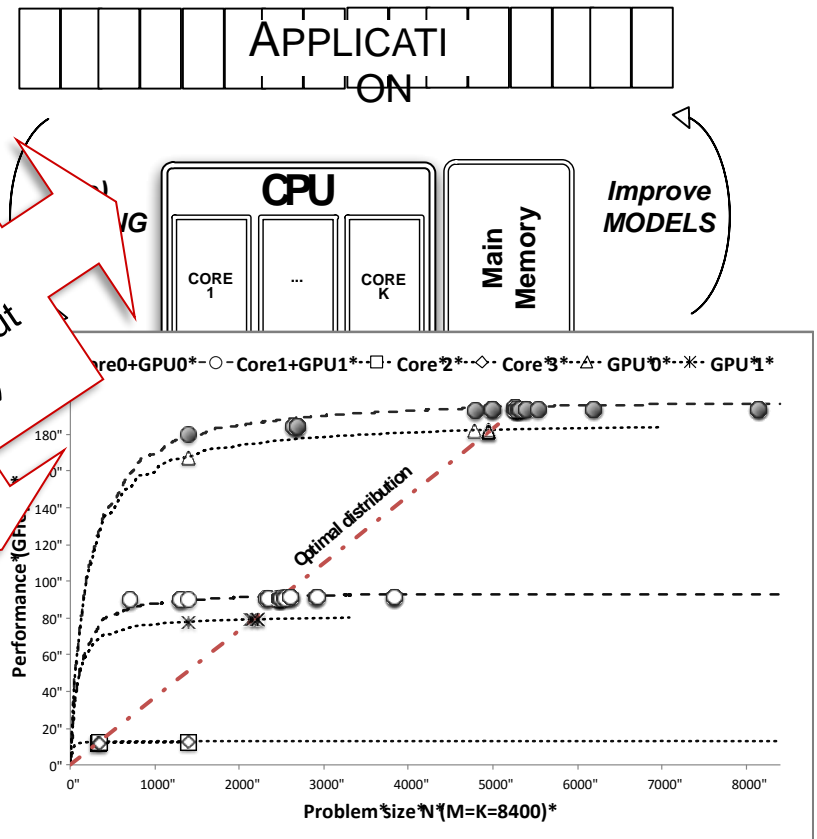


ARCHITECTUR

STATIC DAG-BASED SCHEDULING DYNAMIC DLT-BASED SCHEDULING



**MODELING
APP ↔ ARCH**
(not complex, but
insightful)



ARCHITECTUR

THE CACHE-AWARE ROOFLINE MODEL:

- PERFORMANCE*
- POWER
- EFFICIENCY

*A. Ilic, F. Pratas and L. Sousa "Cache-ware Roofline Model: Upgrading the Loft", *IEEE Computer Architecture Letters* (2013)

Multi-core CPU



SEVERAL (IDENTICAL) PROCESSING CORES


- PARALLELISM ACROSS THE PROCESSING CORES

Multi-core CPU



SEVERAL (IDENTICAL) PROCESSING CORES

- Parallelism across the processing cores
- INSTRUCTION LEVEL PARALLELISM (PIPELINING)




Multi-core CPU



SEVERAL (IDENTICAL) PROCESSING CORES

- Parallelism across the processing cores
- INSTRUCTION LEVEL PARALLELISM (PIPELINING)




Multi-core CPU



SEVERAL (IDENTICAL) PROCESSING CORES

- Parallelism across the processing cores
- INSTRUCTION LEVEL PARALLELISM (PIPELINING)




Multi-core CPU



SEVERAL (IDENTICAL) PROCESSING CORES

- Parallelism across the processing cores
- INSTRUCTION LEVEL PARALLELISM (PIPELINING)



Multi-core CPU




SEVERAL (IDENTICAL) PROCESSING CORES

- Parallelism across the processing cores
- INSTRUCTION LEVEL PARALLELISM (PIPELINING)

HANDS-ON: IVY BRIDGE AT 3.5 GHz (I7 3770K)

- Double-precision Floating-point operation (FLOP)
 - Multiplication(MUL) or addition (ADD)

- **Throughput:** 1 instruction/clock



Instruction Type	FLOPS per Instr.	Performance (GFLOPS/s)	
		1 Core	4 Cores
64 bits	1	3.5 (1flop x 3.5GHz)	14
128 bits (SSE)	2	7	28
256 bits (AVX)	4	14	56

Multi-core CPU




SEVERAL (IDENTICAL) PROCESSING CORES

- Parallelism across the processing cores
- INSTRUCTION LEVEL PARALLELISM (PIPELINING)

HANDS-ON: IVY BRIDGE AT 3.5 GHZ (I7 3770K)

- Double-precision Floating-point operation (FLOP)
 - Multiplication(MUL) or addition (ADD)

- **Throughput:** 1 instruction/clock




Instruction Type	FLOPS per Instr.	Performance (GFLOPS/s)	
		1 Core	4 Cores
64 bits	1	3.5 (1flop x 3.5GHz)	14
128 bits (SSE)	2	7	28
256 bits (AVX)	4	14	56

Multi-core CPU



SEVERAL (IDENTICAL) PROCESSING CORES

- Parallelism across the processing cores
- Instruction level parallelism (pipelining)
- **IN-CORE PARALLELISM (several ports for different ops)**




← Independent instructions can run in parallel

Multi-core CPU



SEVERAL (IDENTICAL) PROCESSING CORES

- Parallelism across the processing cores
- Instruction level parallelism (pipelining)
- **IN-CORE PARALLELISM (several ports for different ops)**



HANDS-ON: IVY BRIDGE AT 3.5 GHz (I7 3770K)

- Double-precision FLOPs
- **Throughput:** 1 instruction/clock


Instruction Type	FLOPS per Instr.	Performance (GFLOPS/s)	
		1 Core	4 Cores
64 bits	1	3.5 (2flops x 3.5GHz)	14
128 bits (SSE)	2	7	28
256 bits (AVX)	4	14	56

Multi-core CPU



SEVERAL (IDENTICAL) PROCESSING CORES

- Parallelism across the processing cores
- Instruction level parallelism (pipelining)
- **IN-CORE PARALLELISM (several ports for different ops)**



HANDS-ON: IVY BRIDGE AT 3.5 GHz (I7 3770K)

- Double-precision FLOPs

- **Throughput: 2 FP instructions/clock**


Instruction Type	FLOPS per Instr.	Performance (GFLOPS/s)	
		1 Core	4 Cores
64 bits	1	3.5 → 7 <small>(2flops x 3.5GHz)</small>	14 → 28
128 bits (SSE)	2	7 → 14	28 → 56
256 bits (AVX)	4	14 → 28	56 → 112

Multi-core CPU



SEVERAL (IDENTICAL) PROCESSING CORES


- Parallelism across the processing cores
- Instruction level parallelism (pipelining)
- **IN-CORE PARALLELISM (several ports for different ops)**



i7 3770K Ivy Bridge	Performance (GFlops/s)*	Bandwidth L1→C (GB/s)*
1 Core	28	?
4 Cores	112	?

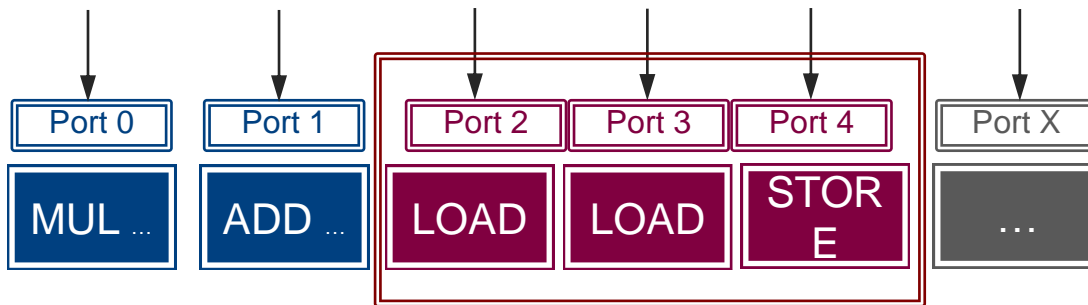
*256-bit AVX double-precision floating-point instructions

Multi-core CPU



SEVERAL (IDENTICAL) PROCESSING CORES

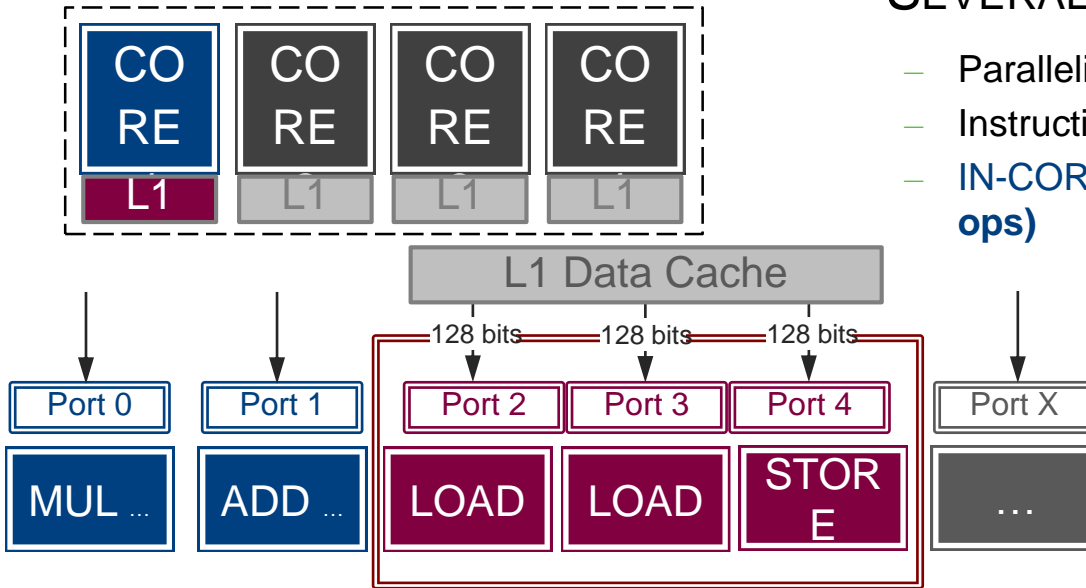
- Parallelism across the processing cores
- Instruction level parallelism (pipelining)
- **IN-CORE PARALLELISM (several ports for different ops)**



i7 3770K Ivy Bridge	Performance (GFlops/s)*	Bandwidth L1→C (GB/s)*
1 Core	28	?
4 Cores	112	?

*256-bit AVX double-precision floating-point instructions

Multi-core CPU



SEVERAL (IDENTICAL) PROCESSING CORES


- Parallelism across the processing cores
- Instruction level parallelism (pipelining)
- **IN-CORE PARALLELISM (several ports for different ops)**

i7 3770K Ivy Bridge	Performance (GFlops/s)*	Bandwidth L1→C (GB/s)*
1 Core	28	?
4 Cores	112	?

*256-bit AVX double-precision floating-point instructions

WHAT IS THE PEAK (THEORETICAL) BANDWIDTH OF THE L1→CORE COMMUNICATION BUS?

Multi-core CPU



SEVERAL (IDENTICAL) PROCESSING CORES

- Parallelism across the processing cores
- Instruction level parallelism (pipelining)
- **IN-CORE PARALLELISM (several ports for different ops)**

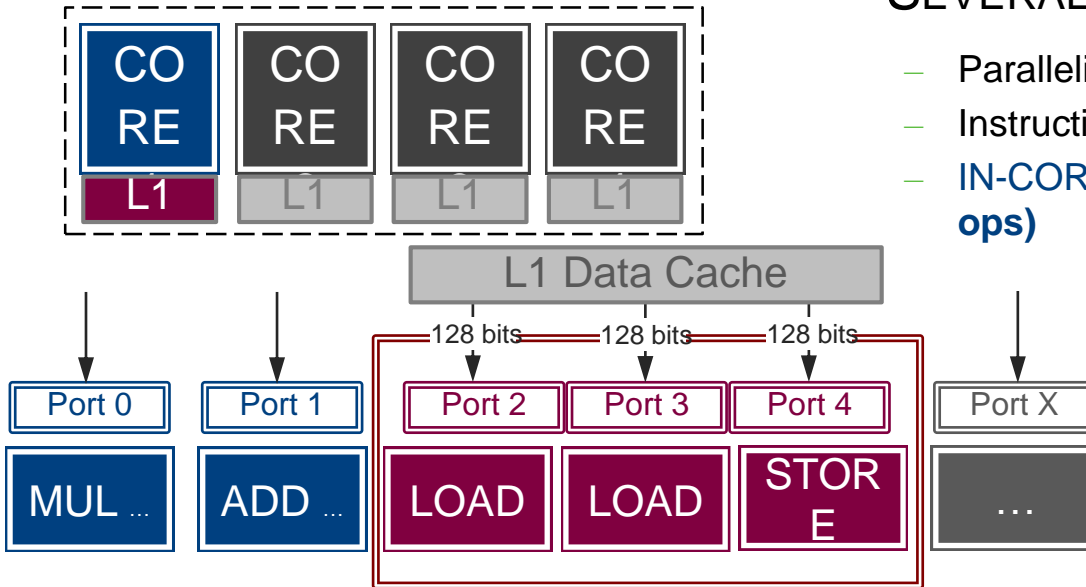
i7 3770K Ivy Bridge	Performance (GFlops/s)*	Bandwidth L1→C (GB/s)*
1 Core	28	168
4 Cores	112	?

*256-bit AVX double-precision floating-point instructions

WHAT IS THE PEAK (THEORETICAL) BANDWIDTH OF THE L1→CORE COMMUNICATION BUS?

- $3 \times 128 \text{ bits} = 48 \text{ bytes}$ can be transferred at the same time (per core)
- bus operates at the frequency of the processor $\rightarrow 3.5 \text{ GHz}$
- $\Rightarrow 48 \text{ bytes} \times 3.5 \text{ GHz} = \mathbf{168 \text{ GB/s (1 Core)}}$

Multi-core CPU



SEVERAL (IDENTICAL) PROCESSING CORES

- Parallelism across the processing cores
- Instruction level parallelism (pipelining)
- **IN-CORE PARALLELISM (several ports for different ops)**

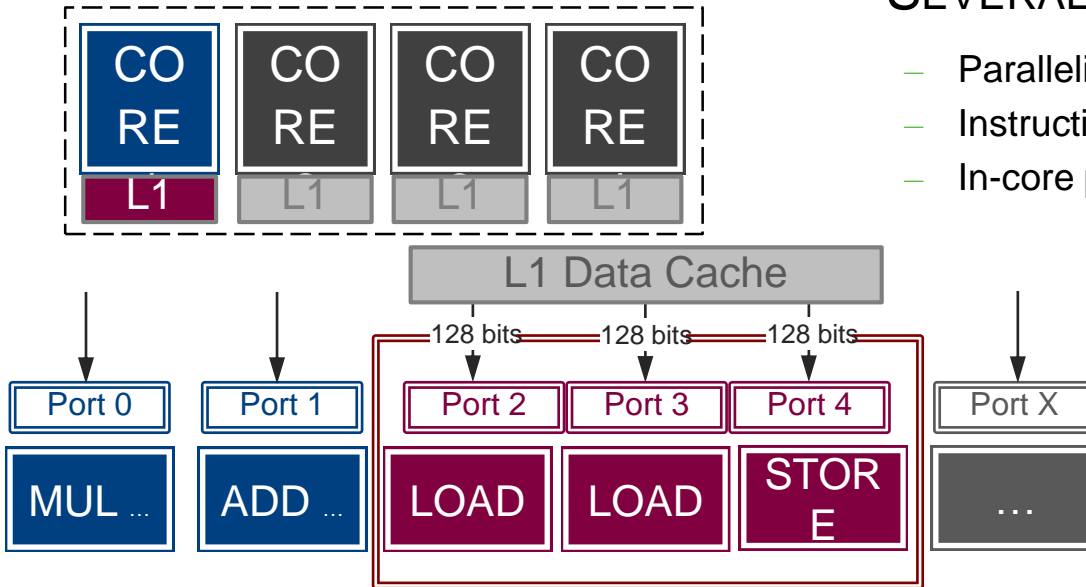
i7 3770K Ivy Bridge	Performance (GFlops/s)*	Bandwidth L1→C (GB/s)*
1 Core	28	168
4 Cores	112	672

*256-bit AVX double-precision floating-point instructions

WHAT IS THE PEAK (THEORETICAL) BANDWIDTH OF THE L1→CORE COMMUNICATION BUS?

- $3 \times 128 \text{ bits} = 48 \text{ bytes}$ can be transferred at the same time (per core)
- bus operates at the frequency of the processor $\rightarrow 3.5 \text{ GHz}$
- $\Rightarrow 48 \text{ bytes} \times 3.5 \text{ GHz} = \mathbf{168 \text{ GB/s (1 Core)}}$
- $\Rightarrow 4 \times 168 \text{ GB/s} = \mathbf{672 \text{ GB/s (4 cores)}}$

Multi-core CPU



SEVERAL (IDENTICAL) PROCESSING CORES

- Parallelism across the processing cores
- Instruction level parallelism (pipelining)
- In-core parallelism (several ports for different ops)


i7 3770K Ivy Bridge	Performance (GFlops/s)*	Bandwidth L1→C (GB/s)*
1 Core	28	168
4 Cores	112	672

*256-bit AVX double-precision floating-point instructions

We can't get higher than this!

What does it tell about the **attainable performance?**

Multi-core CPU



SEVERAL (IDENTICAL) PROCESSING CORES

- Parallelism across the processing cores
- Instruction level parallelism (pipelining)
- In-core parallelism (several ports for different ops)

In general, real applications mix different number of **flops** and **bytes**:

```
// matrix multiplication example
for i=1 to M
  for j=1 to N
    for k=1 to K
      C[i,j] += A[i,k]*B[k,j]
```


i7 3770K Ivy Bridge	Performance (GFlops/s)*	Bandwidth L1→C (GB/s)*
1 Core	28	168
4 Cores	112	672

*256-bit AVX double-precision floating-point instructions

We can't get higher than this!

What does it tell about the **attainable performance?**

Multi-core CPU



SEVERAL (IDENTICAL) PROCESSING CORES

- Parallelism across the processing cores
- Instruction level parallelism (pipelining)
- In-core parallelism (several ports for different ops)


In general, real applications mix different number of flops and bytes:

```
// matrix multiplication example
for i=1 to M
  for j=1 to N
    for k=1 to K
      C[i,j] += A[i,k]*B[k,j]
```

IF memory operations and computations are serially performed:




BUT they are actually executed in parallel (interleaved):



i7 3770K Ivy Bridge	Performance (GFlops/s)*	Bandwidth L1→C (GB/s)*
1 Core	28	168
4 Cores	112	672

*256-bit AVX double-precision floating-point instructions

Multi-core CPU



SEVERAL (IDENTICAL) PROCESSING CORES

- Parallelism across the processing cores
- Instruction level parallelism (pipelining)
- In-core parallelism (several ports for different ops)

In general, real applications mix different number of flops and bytes:

```
// matrix multiplication example
for i=1 to M
  for j=1 to N
    for k=1 to K
      C[i,j] += A[i,k]*B[k,j]
```


i7 3770K Ivy Bridge	Performance (GFlops/s)*	Bandwidth L1→C (GB/s)*
1 Core	28	168
4 Cores	112	672

*256-bit AVX double-precision floating-point instructions


IF memory operations and computations are serially performed:



BUT they are actually executed in parallel (interleaved):



Multi-core CPU



SEVERAL (IDENTICAL) PROCESSING CORES

- Parallelism across the processing cores
- Instruction level parallelism (pipelining)
- In-core parallelism (several ports for different ops)

In general, real applications mix different number of flops and bytes:

```
// matrix multiplication example
for i=1 to M
  for j=1 to N
    for k=1 to K
      C[i,j] += A[i,k]*B[k,j]
```

i7 3770K Ivy Bridge	Perf. $[F_P]$ (GFlops/s)*	Bwidth L1→C $[B_P]$ (GB/s)*
1 Core	28	168
4 Cores	112	672


*256-bit AVX double-precision floating-point instructions

IF memory operations and computations are serially performed:




$$T_C = \text{\#flops}/F_P$$

BUT they are actually executed in parallel (interleaved):



$$\Rightarrow T_T = \max\{T_C, T_M\} = \max\{\text{\#flops}/F_P, T_M\}$$

Multi-core CPU



SEVERAL (IDENTICAL) PROCESSING CORES

- Parallelism across the processing cores
- Instruction level parallelism (pipelining)
- In-core parallelism (several ports for different ops)

In general, real applications mix different number of flops and bytes:

```
// matrix multiplication example
for i=1 to M
  for j=1 to N
    for k=1 to K
      C[i,j] += A[i,k]*B[k,j]
```


i7 3770K Ivy Bridge	Perf. $[F_P]$ (GFlops/s)*	Bwidth L1→C $[B_P]$ (GB/s)*
1 Core	28	168
4 Cores	112	672

*256-bit AVX double-precision floating-point instructions

IF memory operations and computations are serially performed:



BUT they are actually executed in parallel (interleaved):




$$\Rightarrow T_T = \max\{T_C, T_M\} = \max\{\#flops/F_P, \#bytes/B_P\}$$

$$T_C = \#flops/F_P$$

$$T_M = \#bytes/B_P$$

Multi-core CPU



SEVERAL (IDENTICAL) PROCESSING CORES

- Parallelism across the processing cores
- Instruction level parallelism (pipelining)
- In-core parallelism (several ports for different ops)


In general, real applications mix different number of **flops** and **bytes**:

```
// matrix multiplication example
for i=1 to M
  for j=1 to N
    for k=1 to K
      C[i,j] += A[i,k]*B[k,j]
```

i7 3770K Ivy Bridge	Perf. $[F_p]$ (GFlops/s)*	Bwidth L1→C $[B_p]$ (GB/s)*
1 Core	28	168
4 Cores	112	672


*256-bit AVX double-precision floating-point instructions

Memory operations and computations are executed in parallel (interleaved):



$$\Rightarrow T_T = \max\{T_C, T_M\} = \max\{\text{\#flops}/F_p, \text{\#bytes}/B_p\}$$

Multi-core CPU



SEVERAL (IDENTICAL) PROCESSING CORES

- Parallelism across the processing cores
- Instruction level parallelism (pipelining)
- In-core parallelism (several ports for different ops)


In general, real applications mix different number of **flops** and **bytes**:

```
// matrix multiplication example
for i=1 to M
  for j=1 to N
    for k=1 to K
      C[i,j] += A[i,k]*B[k,j]
```

i7 3770K Ivy Bridge	Perf. $[F_P]$ (GFlops/s)*	Bwidth L1→C $[B_P]$ (GB/s)*
1 Core	28	168
4 Cores	112	672

*256-bit AVX double-precision floating-point instructions

Memory operations and computations are executed in parallel (interleaved):




$$\Rightarrow T_T = \max\{T_C, T_M\} = \max\{\text{\#flops}/F_P, \text{\#bytes}/B_P\}$$

Attainable Performance (F_A) of the architecture:

$$F_A = \text{\#flops}/T_T$$

Multi-core CPU



SEVERAL (IDENTICAL) PROCESSING CORES

- Parallelism across the processing cores
- Instruction level parallelism (pipelining)
- In-core parallelism (several ports for different ops)


In general, real applications mix different number of **flops** and **bytes**:

```
// matrix multiplication example
for i=1 to M
  for j=1 to N
    for k=1 to K
      C[i,j] += A[i,k]*B[k,j]
```

i7 3770K Ivy Bridge	Perf. $[F_P]$ (GFlops/s)*	Bwidth L1→C $[B_P]$ (GB/s)*
1 Core	28	168
4 Cores	112	672

*256-bit AVX double-precision floating-point instructions

Memory operations and computations are executed in parallel (interleaved):



$$\Rightarrow T_T = \max\{T_C, T_M\} = \max\{\text{\#flops}/F_P, \text{\#bytes}/B_P\}$$

Attainable Performance (F_A) of the architecture

$$F_A = \text{\#flops}/T_T$$

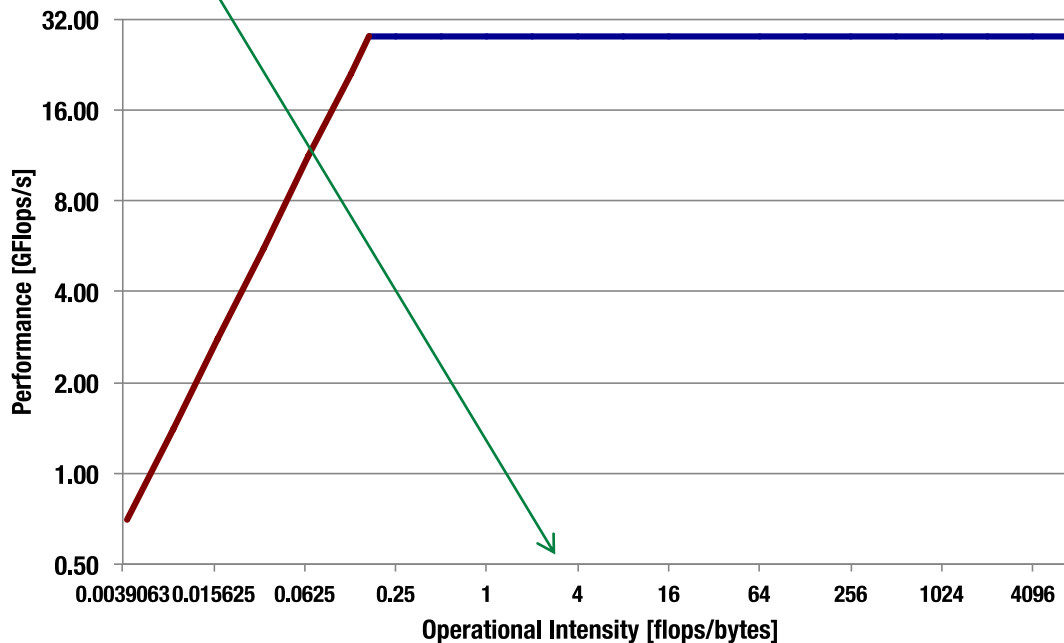
Roofline

Cache-aware Roofline Model



CACHE-AWARE ROOFLINE MODEL - insightful performance model of multi-core architectures relates:

- 1) Maximum Attainable Performance ($F_A = \#flops / T_T$)
- 2) Operational Intensity ($I = \#flops / \#bytes$).



Building the Cache-aware Roofline Model




i7 3770K Ivy Bridge	Perf. [F_P] (GFlops/s)*	Bwidth L1→C [B_P] (GB/s)*
1 Core	28	168
4 Cores	112	672

*256-bit AVX double-precision floating-point instructions



Building the Cache-aware Roofline Model





i7 3770K Ivy Bridge	Perf. [F_P] (GFlops/s)*	Bwidth L1→C [B_P] (GB/s)*
1 Core	28	168
4 Cores	112	672

*256-bit AVX double-precision floating-point instructions





Building the Cache-aware Roofline Model



i7 3770K Ivy Bridge	Perf. [F_P] (GFlops/s)*	Bwidth L1→C [B_P] (GB/s)*
1 Core	28	168
4 Cores	112	672

*256-bit AVX double-precision floating-point instructions

Building the Cache-aware Roofline Model

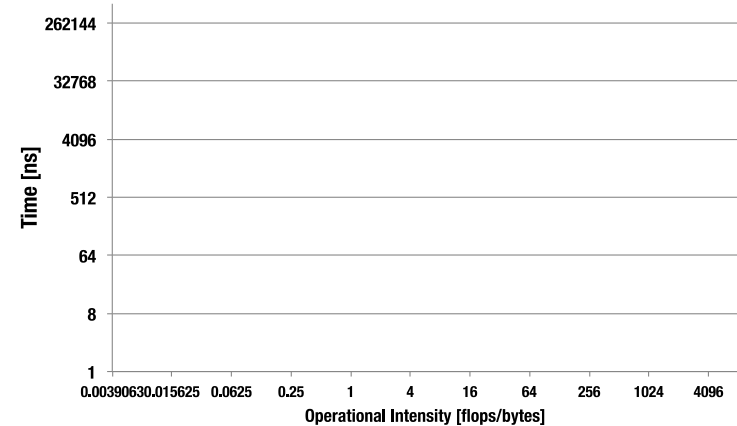



$$I = \text{flops/bytes} = 0.0083 = 8 \text{ flops} / 960 \text{ bytes} [1\text{MAD}/(10 \times (2\text{LD} + \text{ST}))^*]$$

i7 3770K Ivy Bridge	Perf. $[F_P]$ (GFlops/s)*	Bwidth L1→C $[B_P]$ (GB/s)*
1 Core	28	168
4 Cores	112	672

*256-bit AVX double-precision floating-point instructions

Building the Cache-aware Roofline Model



$$I = \text{flops/bytes} = 0.0083 = 8 \text{ flops} / 960 \text{ bytes} \quad [1\text{MAD}/(10 \times (2\text{LD} + \text{ST}))^*]$$



i7 3770K Ivy Bridge	Perf. $[F_P]$ (GFlops/s)*	Bwidth L1→C $[B_P]$ (GB/s)*
1 Core	28	168
4 Cores	112	672

*256-bit AVX double-precision floating-point instructions

$$T_C = \# \text{flops} / F_P = 8 \text{ flops} / 28 = 0.29 \text{ ns}$$



Building the Cache-aware Roofline Model



$$I = \text{flops/bytes} = 0.0083 = 8 \text{ flops} / 960 \text{ bytes} \quad [1\text{MAD}/(10 \times (2\text{LD} + \text{ST}))^*]$$

i7 3770K Ivy Bridge	Perf. $[F_P]$ (GFlops/s)*	Bwidth L1→C $[B_P]$ (GB/s)*
1 Core	28	168
4 Cores	112	672



*256-bit AVX double-precision floating-point instructions

$$T_C = \# \text{flops} / F_P = 8 \text{ flops} / 28 = 0.29 \text{ ns}$$

$$T_M = \# \text{bytes} / B_P = 960 \text{ bytes} / 168 = 5.71 \text{ ns}$$



Building the Cache-aware Roofline Model



$$I = \text{flops/bytes} = 0.0083 = 8\text{flops}/960\text{bytes} [1\text{MAD}/(10 \times (2\text{LD} + \text{ST}))^*]$$

i7 3770K Ivy Bridge	Perf. $[F_P]$ (GFlops/s)*	Bwidth L1→C $[B_P]$ (GB/s)*
1 Core	28	168
4 Cores	112	672

*256-bit AVX double-precision floating-point instructions



$$T_C = \#flops/F_P = 8\text{flops}/28 = 0.29 \text{ ns}$$

$$T_M = \#bytes/B_P = 960\text{bytes}/168 = 5.71 \text{ ns}$$

$$T_T = \max\{T_C, T_M\} = 5.71 \text{ ns}$$



Building the Cache-aware Roofline Model



$$I = \text{flops/bytes} = 0.0083 = 8 \text{flops}/960 \text{bytes} [1\text{MAD}/(10 \times (2\text{LD} + \text{ST}))^*]$$

i7 3770K Ivy Bridge	Perf. [F_P] (GFlops/s)*	Bwidth L1→C [B_P] (GB/s)*
1 Core	28	168
4 Cores	112	672

*256-bit AVX double-precision floating-point instructions



$$T_C = \# \text{flops} / F_P = 8 \text{flops} / 28 = 0.29 \text{ ns}$$

$$T_M = \# \text{bytes} / B_P = 960 \text{bytes} / 168 = 5.71 \text{ ns}$$

$$T_T = \max\{T_C, T_M\} = 5.71 \text{ ns}$$



Building the Cache-aware Roofline Model



$$I = \text{flops/bytes} = 0.0083 = 8\text{flops}/960\text{bytes} [1\text{MAD}/(10 \times (2\text{LD} + \text{ST}))^*]$$

i7 3770K Ivy Bridge	Perf. [F_P] (GFlops/s)*	Bwidth L1→C [B_P] (GB/s)*
1 Core	28	168
4 Cores	112	672

*256-bit AVX double-precision floating-point instructions

$$T_C = \#flops/F_P = 8\text{flops}/28 = 0.29 \text{ ns}$$



$$T_M = \#bytes/B_P = 960\text{bytes}/168 = 5.71 \text{ ns}$$

$$T_T = \max\{T_C, T_M\} = 5.71 \text{ ns}$$

$$F_A = \#flops/T_T = 8\text{flops}/5.71 \text{ ns} = 1.4 \text{ Gflops/s}$$



Building the Cache-aware Roofline Model



$$I = \text{flops/bytes} = 0.0083 = 8 \text{ flops} / 960 \text{ bytes} \quad [1\text{MAD}/(10 \times (2\text{LD} + \text{ST}))^*]$$

i7 3770K Ivy Bridge	Perf. [F_P] (GFlops/s)*	Bwidth L1→C [B_P] (GB/s)*
1 Core	28	168
4 Cores	112	672

*256-bit AVX double-precision floating-point instructions

$$T_C = \# \text{flops} / F_P = 8 \text{ flops} / 28 = 0.29 \text{ ns}$$



$$T_M = \# \text{bytes} / B_P = 960 \text{ bytes} / 168 = 5.71 \text{ ns}$$

$$T_T = \max\{T_C, T_M\} = 5.71 \text{ ns}$$

$$F_A = \# \text{flops} / T_T = 8 \text{ flops} / 5.71 \text{ ns} = 1.4 \text{ Gflops/s}$$



Building the Cache-aware Roofline Model



$I = \text{flops/bytes} = 0.067 = 64\text{flops}/960\text{bytes} [8\text{MAD}/(10 \times (2\text{LD} + \text{ST}))^*]$

i7 3770K Ivy Bridge	Perf. $[F_P]$ (GFlops/s)*	Bwidth L1→C $[B_P]$ (GB/s)*
1 Core	28	168
4 Cores	112	672

*256-bit AVX double-precision floating-point instructions

$T_C = \text{\#flops}/F_P = ?$



$T_M = \text{\#bytes}/B_P = ?$

$T_T = \max\{T_C, T_M\} = ?$

$F_A = \text{\#flops}/T_T = ?$



Building the Cache-aware Roofline Model



$I = \text{flops/bytes} = 0.067 = 64\text{flops}/960\text{bytes} [8\text{MAD}/(10 \times (2\text{LD} + \text{ST}))^*]$

i7 3770K Ivy Bridge	Perf. [F_P] (GFlops/s)*	Bwidth L1→C [B_P] (GB/s)*
1 Core	28	168
4 Cores	112	672

*256-bit AVX double-precision floating-point instructions

$T_C = \text{\#flops}/F_P = 64\text{flops}/28 = 2.29 \text{ ns}$



$T_M = \text{\#bytes}/B_P = 960\text{bytes}/168 = 5.71 \text{ ns}$

$T_T = \max\{T_C, T_M\} = 5.71 \text{ ns}$

$F_A = \text{\#flops}/T_T = ?$



Building the Cache-aware Roofline Model



$I = \text{flops/bytes} = 0.067 = 64\text{flops}/960\text{bytes} [8\text{MAD}/(10 \times (2\text{LD} + \text{ST}))^*]$

i7 3770K Ivy Bridge	Perf. [F_P] (GFlops/s)*	Bwidth L1→C [B_P] (GB/s)*
1 Core	28	168
4 Cores	112	672

*256-bit AVX double-precision floating-point instructions

$T_C = \text{\#flops}/F_P = 64\text{flops}/28 = 2.29 \text{ ns}$



$T_M = \text{\#bytes}/B_P = 960\text{bytes}/168 = 5.71 \text{ ns}$

$T_T = \max\{T_C, T_M\} = 5.71 \text{ ns}$

$F_A = \text{\#flops}/T_T = ?$



Building the Cache-aware Roofline Model



$I = \text{flops/bytes} = 0.067 = 64\text{flops}/960\text{bytes} [8\text{MAD}/(10 \times (2\text{LD} + \text{ST}))^*]$

i7 3770K Ivy Bridge	Perf. [F_P] (GFlops/s)*	Bwidth L1→C [B_P] (GB/s)*
1 Core	28	168
4 Cores	112	672

*256-bit AVX double-precision floating-point instructions

$T_C = \text{\#flops}/F_P = 64\text{flops}/28 = 2.29 \text{ ns}$



$T_M = \text{\#bytes}/B_P = 960\text{bytes}/168 = 5.71 \text{ ns}$

$T_T = \max\{T_C, T_M\} = 5.71 \text{ ns}$

$F_A = \text{\#flops}/T_T = 64\text{flops}/5.71 \text{ ns} = 11.2 \text{ Gflops/s}$



Building the Cache-aware Roofline Model



$I = \text{flops/bytes} = 0.067 = 64\text{flops}/960\text{bytes} [8\text{MAD}/(10 \times (2\text{LD} + \text{ST}))^*]$

i7 3770K Ivy Bridge	Perf. $[F_P]$ (GFlops/s)*	Bwidth L1→C $[B_P]$ (GB/s)*
1 Core	28	168
4 Cores	112	672

*256-bit AVX double-precision floating-point instructions

$T_C = \text{\#flops}/F_P = 64\text{flops}/28 = 2.29 \text{ ns}$



$T_M = \text{\#bytes}/B_P = 960\text{bytes}/168 = 5.71 \text{ ns}$

$T_T = \max\{T_C, T_M\} = 5.71 \text{ ns}$

$F_A = \text{\#flops}/T_T = 64\text{flops}/5.71 \text{ ns} = 11.2 \text{ Gflops/s}$



Building the Cache-aware Roofline Model



$I = \text{flops/bytes} = 0.016 = 160\text{flops}/960\text{bytes} [20\text{MAD}/(10 \times (2\text{LD} + \text{ST}))^*]$

i7 3770K Ivy Bridge	Perf. $[F_P]$ (GFlops/s)*	Bwidth L1→C $[B_P]$ (GB/s)*
1 Core	28	168
4 Cores	112	672

*256-bit AVX double-precision floating-point instructions

$T_C = \text{\#flops}/F_P = ?$



$T_M = \text{\#bytes}/B_P = ?$

$T_T = \max\{T_C, T_M\} = ?$

$F_A = \text{\#flops}/T_T = ?$



Building the Cache-aware Roofline Model



$I = \text{flops/bytes} = 0.016 = 160\text{flops}/960\text{bytes}$ [20MAD/((10x(2LD+ST))*)]

i7 3770K Ivy Bridge	Perf. [F_P] (GFlops/s)*	Bwidth L1→C [B_P] (GB/s)*
1 Core	28	168
4 Cores	112	672

*256-bit AVX double-precision floating-point instructions

$T_C = \text{\#flops}/F_P = 160\text{flops}/28 = 5.71 \text{ ns}$



$T_M = \text{\#bytes}/B_P = 960\text{bytes}/168 = 5.71 \text{ ns}$

$T_T = \max\{T_C, T_M\} = 5.71 \text{ ns}$ ($T_C = T_M$)

$F_A = \text{\#flops}/T_T = ?$



Building the Cache-aware Roofline Model



$I = \text{flops/bytes} = 0.016 = 160\text{flops}/960\text{bytes}$ [20MAD/((10x(2LD+ST))*)]

i7 3770K Ivy Bridge	Perf. [F_P] (GFlops/s)*	Bwidth L1→C [B_P] (GB/s)*
1 Core	28	168
4 Cores	112	672

*256-bit AVX double-precision floating-point instructions

$T_C = \text{\#flops}/F_P = 160\text{flops}/28 = 5.71 \text{ ns}$



$T_M = \text{\#bytes}/B_P = 960\text{bytes}/168 = 5.71 \text{ ns}$

$T_T = \max\{T_C, T_M\} = 5.71 \text{ ns}$ ($T_C = T_M$)

$F_A = \text{\#flops}/T_T = ?$



Building the Cache-aware Roofline Model



$I = \text{flops/bytes} = 0.016 = 160\text{flops}/960\text{bytes}$ [20MAD/((10x(2LD+ST))*)]

i7 3770K Ivy Bridge	Perf. [F_P] (GFlops/s)*	Bwidth L1→C [B_P] (GB/s)*
1 Core	28	168
4 Cores	112	672

*256-bit AVX double-precision floating-point instructions

$T_C = \text{\#flops}/F_P = 160\text{flops}/28 = 5.71 \text{ ns}$



$T_M = \text{\#bytes}/B_P = 960\text{bytes}/168 = 5.71 \text{ ns}$

$T_T = \max\{T_C, T_M\} = 5.71 \text{ ns}$ ($T_C = T_M$)

$F_A = \text{\#flops}/T_T = 160\text{flops}/5.71 \text{ ns} = 28 \text{ Gflops/s}$



Building the Cache-aware Roofline Model



$$I = \text{flops/bytes} = 0.016 = 160\text{flops}/960\text{bytes} \quad [20\text{MAD}/(10 \times (2\text{LD} + \text{ST}))^*]$$

i7 3770K Ivy Bridge	Perf. [F_P] (GFlops/s)*	Bwidth L1→C [B_P] (GB/s)*
1 Core	28	168
4 Cores	112	672

*256-bit AVX double-precision floating-point instructions

$$T_C = \#flops/F_P = 160\text{flops}/28 = 5.71 \text{ ns}$$



$$T_M = \#bytes/B_P = 960\text{bytes}/168 = 5.71 \text{ ns}$$

$$T_T = \max\{T_C, T_M\} = 5.71 \text{ ns} \quad (T_C = T_M)$$

$$F_A = \#flops/T_T = 160\text{flops}/5.71 \text{ ns} = 28 \text{ Gflops/s}$$



Building the Cache-aware Roofline Model



$I = \text{flops/bytes} = 0.016 = 160\text{flops}/960\text{bytes}$ [20MAD/((10x(2LD+ST))*)]

i7 3770K Ivy Bridge	Perf. [F_P] (GFlops/s)*	Bwidth L1→C [B_P] (GB/s)*
1 Core	28	168
4 Cores	112	672

*256-bit AVX double-precision floating-point instructions

$T_C = \text{\#flops}/F_P = 160\text{flops}/28 = 5.71 \text{ ns}$



$T_M = \text{\#bytes}/B_P = 960\text{bytes}/168 = 5.71 \text{ ns}$

$T_T = \max\{T_C, T_M\} = 5.71 \text{ ns}$ ($T_C = T_M$)

$F_A = \text{\#flops}/T_T = 160\text{flops}/5.71 \text{ ns} = 28 \text{ Gflops/s}$



Building the Cache-aware Roofline Model



$I = \text{flops/bytes} = 1 = 960\text{flops}/960\text{bytes} [120\text{MAD}/(10 \times (2\text{LD} + \text{ST}))^*]$

i7 3770K Ivy Bridge	Perf. $[F_P]$ (GFlops/s)*	Bwidth L1→C $[B_P]$ (GB/s)*
1 Core	28	168
4 Cores	112	672

*256-bit AVX double-precision floating-point instructions

$T_C = \text{\#flops}/F_P = ?$



$T_M = \text{\#bytes}/B_P = ?$

$T_T = \max\{T_C, T_M\} = ?$

$F_A = \text{\#flops}/T_T = ?$



Building the Cache-aware Roofline Model



$I = \text{flops/bytes} = 1 = 960\text{flops}/960\text{bytes} [120\text{MAD}/(10 \times (2\text{LD} + \text{ST}))^*]$

i7 3770K Ivy Bridge	Perf. [F_P] (GFlops/s)*	Bwidth L1→C [B_P] (GB/s)*
1 Core	28	168
4 Cores	112	672

*256-bit AVX double-precision floating-point instructions

$T_C = \#flops/F_P = 960\text{flops}/28 = 34.29 \text{ ns}$



$T_M = \#bytes/B_P = 960\text{bytes}/168 = 5.71 \text{ ns}$

$T_T = \max\{T_C, T_M\} = 34.29 \text{ ns}$

$F_A = \#flops/T_T = 960\text{flops}/34.29\text{ns} = 28 \text{ Gflops/s}$



Building the Cache-aware Roofline Model



$I = \text{flops/bytes} = 1 = 960\text{flops}/960\text{bytes} [120\text{MAD}/(10 \times (2\text{LD} + \text{ST}))^*]$

i7 3770K Ivy Bridge	Perf. [F_P] (GFlops/s)*	Bwidth L1→C [B_P] (GB/s)*
1 Core	28	168
4 Cores	112	672

*256-bit AVX double-precision floating-point instructions

$T_C = \#flops/F_P = 960\text{flops}/28 = 34.29 \text{ ns}$



$T_M = \#bytes/B_P = 960\text{bytes}/168 = 5.71 \text{ ns}$

$T_T = \max\{T_C, T_M\} = 34.29 \text{ ns}$

$F_A = \#flops/T_T = 960\text{flops}/34.29\text{ns} = 28 \text{ Gflops/s}$



Building the Cache-aware Roofline Model



$I = \text{flops/bytes} = 64 = 61440 \text{flops} / 960 \text{bytes} [7680 \text{MAD} / ((10 \times (2 \text{LD} + \text{ST}))^*)]$

i7 3770K Ivy Bridge	Perf. [F_P] (GFlops/s)*	Bwidth L1→C [B_P] (GB/s)*
1 Core	28	168
4 Cores	112	672

*256-bit AVX double-precision floating-point instructions

$T_C = \# \text{flops} / F_P = 7680 \text{flops} / 28 = 2194.29 \text{ ns}$



$T_M = \# \text{bytes} / B_P = 960 \text{bytes} / 168 = 5.71 \text{ ns}$

$T_T = \max\{T_C, T_M\} = 2194.29 \text{ ns}$

$F_A = \# \text{flops} / T_T = 28 \text{ Gflops/s}$



Building the Cache-aware Roofline Model



$I = \text{flops/bytes} = 64 = 61440 \text{flops} / 960 \text{bytes} \text{ [7680MAD} / ((10 \times (2\text{LD} + \text{ST}))^*)]$

i7 3770K Ivy Bridge	Perf. [F_P] (GFlops/s)*	Bwidth L1→C [B_P] (GB/s)*
1 Core	28	168
4 Cores	112	672

*256-bit AVX double-precision floating-point instructions

$T_C = \text{\#flops} / F_P = 7680 \text{flops} / 28 = 2194.29 \text{ ns}$



$T_M = \text{\#bytes} / B_P = 960 \text{bytes} / 168 = 5.71 \text{ ns}$

$T_T = \max\{T_C, T_M\} = 2194.29 \text{ ns}$

$F_A = \text{\#flops} / T_T = 28 \text{ Gflops/s}$



Building the Cache-aware Roofline Model



$I = \text{flops/bytes} = 8192 = 7864320 \text{ flops} / 960 \text{ bytes} \text{ [983040MAD}/(10 \times (2\text{LD} + \text{ST}))^*]$

i7 3770K Ivy Bridge	Perf. [F_P] (GFlops/s)*	Bwidth L1→C [B_P] (GB/s)*
1 Core	28	168
4 Cores	112	672

*256-bit AVX double-precision floating-point instructions

$T_C = \text{\#flops}/F_P = 280868.58 \text{ ns}$


$T_M = \text{\#bytes}/B_P = 5.71 \text{ ns}$

$T_T = \max\{T_C, T_M\} = 280868.58 \text{ ns}$

$F_A = \text{\#flops}/T_T = 28 \text{ Gflops/s}$



Building the Cache-aware Roofline Model



APPLICATION


CHARACTERIZATION

- Memory-bound applications
- Compute-bound applications

i7 3770K Ivy Bridge	Perf. [F_P] (GFlops/s)*	Bwidth L1→C [B_P] (GB/s)*
1 Core	28	168
4 Cores	112	672

*256-bit AVX double-precision floating-point instructions

Building the Cache-aware Roofline Model



APPLICATION

CHARACTERIZATION


- Memory-bound applications
- Compute-bound applications

Application is a SINGLE POINT in the Cache-Aware Roofline Model!

i7 3770K Ivy Bridge	Perf. [F_P] (GFlops/s)*	Bwidth L1→C [B_P] (GB/s)*
1 Core	28	168
4 Cores	112	672

*256-bit AVX double-precision floating-point instructions


Building the Cache-aware Roofline Model



MODEL FOR 1 CORE

MODEL FOR 4 CORES?

Multi-core CPU




i7 3770K Ivy Bridge	Perf. [F_P] (GFlops/s)*	Bwidth L1→C [B_P] (GB/s)*
1 Core	28	168
4 Cores	112	672

*256-bit AVX double-precision floating-point instructions




Building the Cache-aware Roofline Model



MODEL FOR 4 CORES

MODEL FOR 1 CORE


Multi-core CPU



i7 3770K Ivy Bridge	Perf. [F_P] (GFlops/s)*	Bwidth L1→C [B_P] (GB/s)*
1 Core	28	168
4 Cores	112	672

*256-bit AVX double-precision floating-point instructions


Building the Cache-aware Roofline Model



MODEL FOR 4 CORES

MODEL FOR 1 CORE


Multi-core CPU



i7 3770K Ivy Bridge	Perf. [F_P] (GFlops/s)*	Bwidth L1→C [B_P] (GB/s)*
1 Core	28	168
4 Cores	112	672

*256-bit AVX double-precision floating-point instructions

Building the Cache-aware Roofline Model



MODEL FOR 4 CORES


MODEL FOR 1 CORE

As for now, we just considered L1 bandwidth!

i7 3770K Ivy Bridge	Perf. [F_P] (GFlops/s)*	Bwidth L1→C [B_P] (GB/s)*
1 Core	28	168
4 Cores	112	672


*256-bit AVX double-precision floating-point instructions

Multi-core CPU



Multi-core Architectures - Memory Hierarchy -

Multi-core CPU



MEMORY HIERARCHY

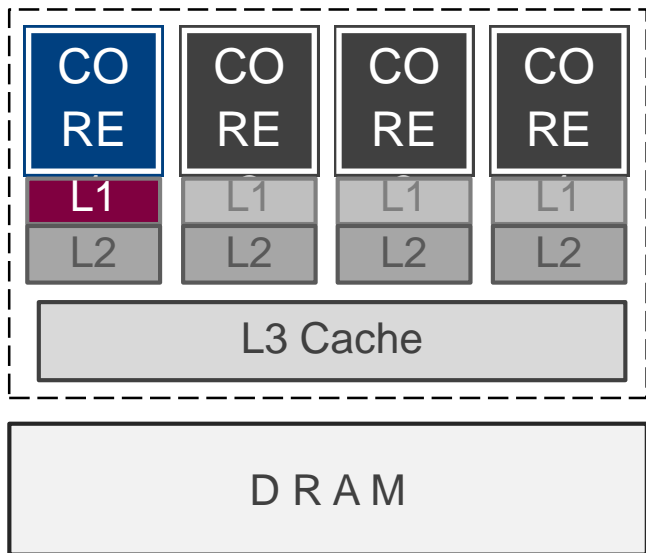
- Set of on-chip caches: private (L1, L2) or shared (L3)
- Global memory (DRAM)
- Caches hide the latency when accessing DRAM (also between successive cache levels)

i7 3770K Ivy Bridge	Perf. $[F_P]$ (GFlops/s)*	Bwidth L1→C $[B_P]$ (GB/s)*
1 Core	28	168
4 Cores	112	672

*256-bit AVX double-precision floating-point instructions

Multi-core Architectures - Memory Hierarchy -

Multi-core CPU



MEMORY HIERARCHY

- Set of on-chip caches: private (L1, L2) or shared (L3)
- Global memory (DRAM)
- Caches hide the latency when accessing DRAM (also between successive cache levels)

CACHE-AWARE ROOFLINE MODEL


- Peak FP performance and L1 bandwidth obtained from processor's specifications (bottom table)
- **We need bandwidth from all other memory levels to the Core?**

i7 3770K Ivy Bridge	Perf. $[F_P]$ (GFlops/s)*	Bwidth L1→C $[B_P]$ (GB/s)*
1 Core	28	168
4 Cores	112	672

*256-bit AVX double-precision floating-point instructions

Multi-core Architectures - Memory Hierarchy -

Multi-core CPU



MEMORY HIERARCHY

- Set of on-chip caches: private (L1, L2) or shared (L3)
- Global memory (DRAM)
- Caches hide the latency when accessing DRAM (also between successive cache levels)

CACHE-AWARE ROOFLINE MODEL

- Peak FP performance and L1 bandwidth obtained from processor's specifications (bottom table)
- **We need bandwidth from all other memory levels to the Core?**

MICRO-BENCHMARKS FOR ARCHITECTURE CHARACTERIZATION

i7 3770K Ivy Bridge	Perf. [F _P] (GFlops/s)*	Bwidth L1→C [B _P] (GB/s)*
1 Core	28	168
4 Cores	112	672

*256-bit AVX double-precision floating-point instructions


```
// AVX Assembly code: 2 Loads + 1 Store
```

```
vmovapd 0(%rax), %ymm0
vmovapd 32(%rax), %ymm1
vmovapd %ymm2, 64(%rax)
vmovapd 96(%rax), %ymm3
vmovapd 128(%rax), %ymm4
vmovapd %ymm5, 160(%rax)
```


...

Multi-core Architectures - Memory Hierarchy -

Multi-core CPU



Memory bandwidth variation for AVX, SSE, and DP scalars



// AVX Assembly code: 2 Loads + 1 Store

```
vmovapd 0(%rax), %ymm0
vmovapd 32(%rax), %ymm1
vmovapd %ymm2, 64(%rax)
vmovapd 96(%rax), %ymm3
vmovapd 128(%rax), %ymm4
vmovapd %ymm5, 160(%rax)
```


...

i7 3770K Ivy Bridge	Perf. [F _P] (GFlops/s)*	Bwidth L1→C [B _P] (GB/s)*
1 Core	28	168
4 Cores	112	672


*256-bit AVX double-precision floating-point instructions

Multi-core Architectures - Memory Hierarchy -

Multi-core CPU



Memory bandwidth variation for AVX, SSE, and DP scalars



// AVX Assembly code: 2 Loads + 1 Store


```
vmovapd 0(%rax), %ymm0
vmovapd 32(%rax), %ymm1
vmovapd %ymm2, 64(%rax)
vmovapd 96(%rax), %ymm3
vmovapd 128(%rax), %ymm4
vmovapd %ymm5, 160(%rax)
```

i7 3770K Ivy Bridge	Perf. [F _P] (GFlops/s)*	Bwidth L1→C [B _P] (GB/s)*
1 Core	28	168
4 Cores	112	672


*256-bit AVX double-precision floating-point instructions

Multi-core Architectures - Memory Hierarchy -

Multi-core CPU



Memory bandwidth variation for AVX, SSE, and DP scalars



How to measure?

```
// Configured Performance Counters
```

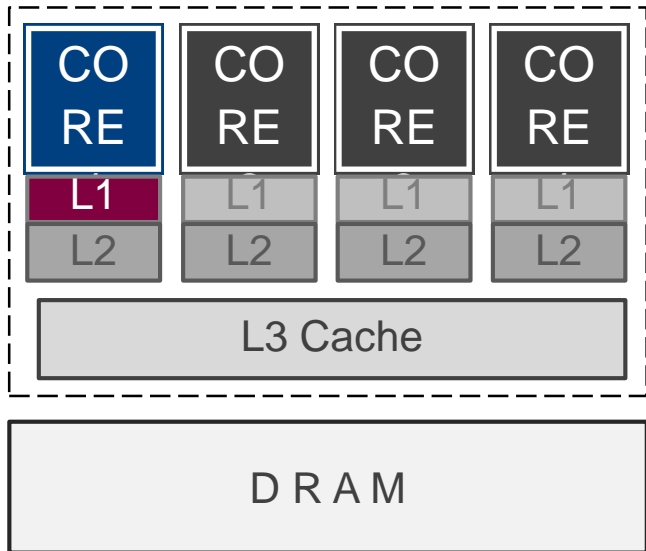
```
CPU_CLK_UNHALTED.CORE/REF
MEM_UOP_RETIRED.ALL_LOADS
MEM_UOP_RETIRED.ALL_STORES
...
```

```
// AVX Assembly code: 2 Loads + 1 Store
```

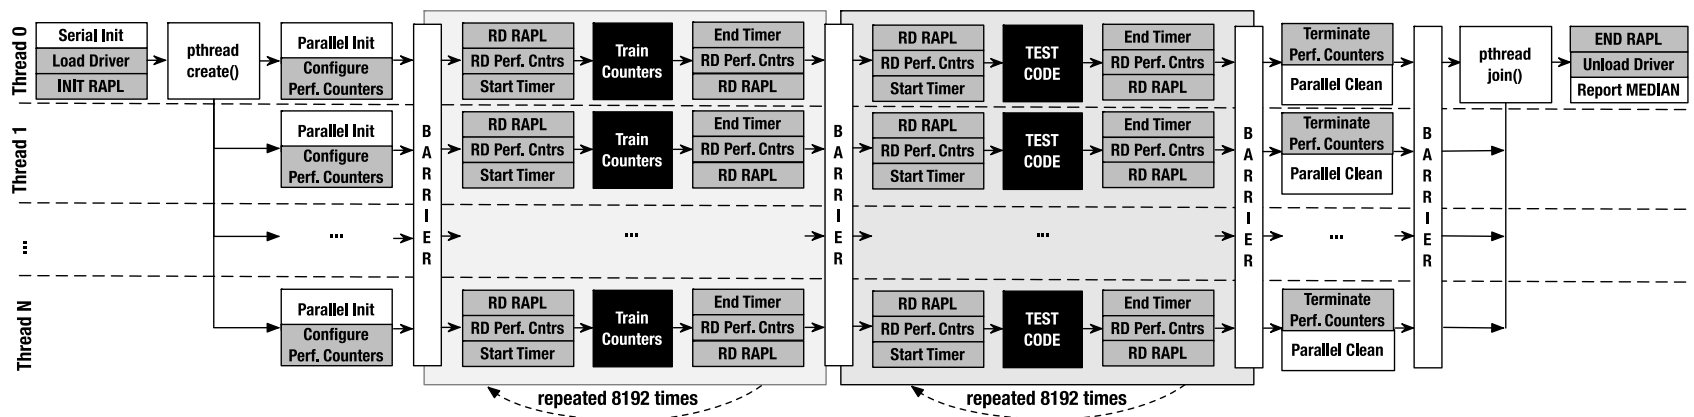

```
vmovapd 0(%rax), %ymm0
vmovapd 32(%rax), %ymm1
vmovapd %ymm2, 64(%rax)
vmovapd 96(%rax), %ymm3
vmovapd 128(%rax), %ymm4
vmovapd %ymm5, 160(%rax)
...
```

Multi-core Architectures - Memory Hierarchy -

Multi-core CPU



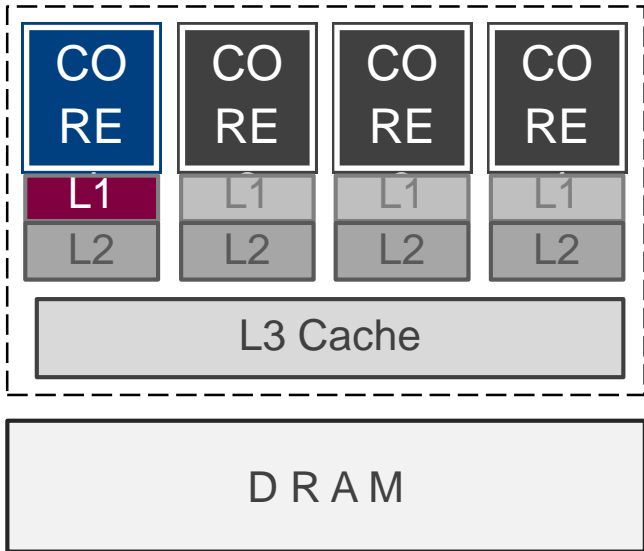
Memory bandwidth variation for AVX, SSE, and DP scalars




Multi-core Architectures - Memory Hierarchy -



Multi-core CPU



Performance variation for AVX



How to measure?

// Configured Performance Counters

```
CPU_CLK_UNHALTED.CORE/REF
FP_OPS_EXE_SSE_SCALAR_DBL
FP_OPS_EXE_SSE_FP_PACKED_DBL
SIMD_FP_256_PACKED_DBL
...
```

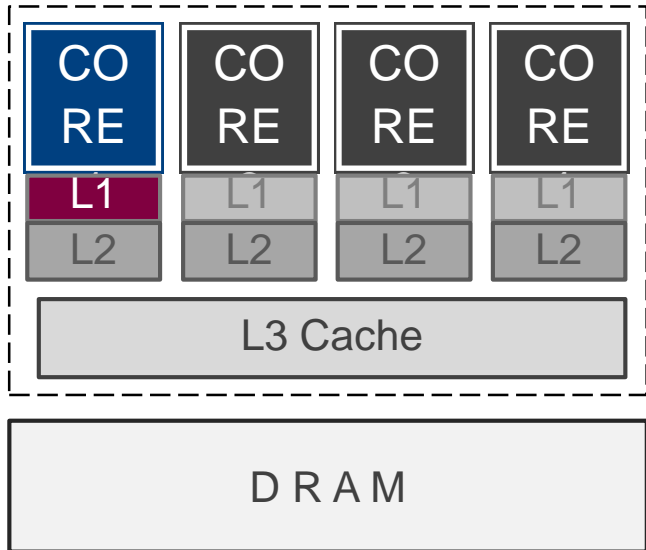
// AVX Assembly code: 2 Loads + 1 Store

```
vmulpd    %ymm0, %ymm0, %ymm0
vaddpd    %ymm1, %ymm1, %ymm1
vmulpd    %ymm2, %ymm2, %ymm2
vaddpd    %ymm3, %ymm3, %ymm3
vmulpd    %ymm4, %ymm4, %ymm4
vaddpd    %ymm5, %ymm5, %ymm5
...
```




Multi-core Architectures - Memory Hierarchy -

Multi-core CPU



Performance variation for AVX



// AVX Assembly code: 2 Loads + 1 Store


```
vmulpd  %ymm0, %ymm0, %ymm0
vaddpd  %ymm1, %ymm1, %ymm1
vmulpd  %ymm2, %ymm2, %ymm2
vaddpd  %ymm3, %ymm3, %ymm3
vmulpd  %ymm4, %ymm4, %ymm4
vaddpd  %ymm5, %ymm5, %ymm5
```

i7 3770K Ivy Bridge	Perf. [F_p] (GFlops/s)*	Bwidth L1→C [B_p] (GB/s)*
1 Core	28	168
4 Cores	112	672

*256-bit AVX double-precision floating-point instructions


Cache-Aware Roofline Model

- Putting it all together -




i7 3770K Ivy Bridge	Perf. [F_P] (GFlops/s)*	Bwidth L1 → C [B_P] (GB/s)*
1 Core	28	168
4 Cores	112	672

*256-bit AVX double-precision floating-point instructions




Cache-aware Roofline Model: Hands On



- Insightful **single plot model**
 - Shows performance limits of multicores
 - Redefined OI: flops and bytes as seen by core
 - Constructed once per architecture
- Considers **complete memory hierarchy**
 - Influence of caches and DRAM to performance
- Applicable to **other types of operations**
 - not only floating-point
- **Useful for:**
 - **Application** characterization and optimization
 - **Architecture** development and understanding

Cache-aware Roofline Model: Hands On



• Total Cache-aware Roofline Model

- Includes **all transitional states** (traversing the memory hierarchy and filling the pipeline)
- Single-plot modeling for **different** types of compute and memory **operations**

• Insightful **single plot model**

- Shows performance limits of multicores
- Redefined OI: flops and bytes as seen by core
- Constructed once per architecture

• Considers **complete memory hierarchy**

- Influence of caches and DRAM to performance


• Applicable to **other types of operations**

- not only floating-point

• Useful for:

- **Application** characterization and optimization
- **Architecture** development and understanding

Cache-aware Roofline Model: Hands On



• Total Cache-aware Roofline Model

- Includes **all transitional states** (traversing the memory hierarchy and filling the pipeline)
- Single-plot modeling for **different** types of compute and memory **operations**

• Insightful **single plot model**

- Shows performance limits of multicores
- Redefined OI: flops and bytes as seen by core
- Constructed once per architecture

• Considers **complete memory hierarchy**

- Influence of caches and DRAM to performance


• Applicable to **other types of operations**

- not only floating-point

• Useful for:

- **Application** characterization and optimization
- **Architecture** development and understanding

Cache-aware Roofline Model: Hands On



• Total Cache-aware Roofline Model

- Includes **all transitional states** (traversing the memory hierarchy and filling the pipeline)
- Single-plot modeling for **different** types of compute and memory **operations**

• Insightful **single plot model**

- Shows performance limits of multicores
- Redefined OI: flops and bytes as seen by core
- Constructed once per architecture

• Considers **complete memory hierarchy**

- Influence of caches and DRAM to performance


• Applicable to **other types of operations**

- not only floating-point

• Useful for:

- **Application** characterization and optimization
- **Architecture** development and understanding

Cache-aware Roofline Model: Hands On



• Total Cache-aware Roofline Model

- Includes **all transitional states** (traversing the memory hierarchy and filling the pipeline)
- Single-plot modeling for **different** types of compute and memory **operations**

• Insightful **single plot model**

- Shows performance limits of multicores
- Redefined OI: flops and bytes as seen by core
- Constructed once per architecture

• Considers **complete memory hierarchy**

- Influence of caches and DRAM to performance

• Applicable to **other types of operations**

- not only floating-point

• Useful for:

- **Application** characterization and optimization
- **Architecture** development and understanding


Cache-Aware Roofline Model vs. State-of-the-Art



WHAT IS HOT (WHAT IS NOT)?
- APPLICATION CHARACTERIZATION -


The Original Roofline Model*

- **Multi-cores:** Powerful cores and memory hierarchy (caches and DRAM)



- **Performance:** Computations (*flops*) and communication (*bytes*) overlap in

time


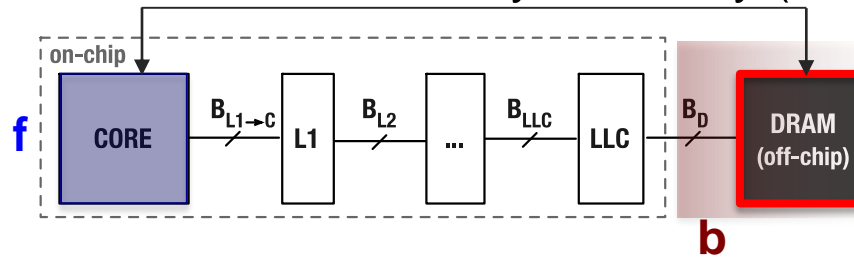


* Williams, S., Waterman, A. and Patterson, D., "Roofline: An insightful visual performance model for multicore architectures", Communications of the ACM (2009)


The Original Roofline Model: Hands On



- **Multi-cores:** Powerful cores and memory hierarchy (caches and DRAM)



APP-D (data traffic from DRAM)



$$I = (\sum f_i) / (\sum b_i)$$


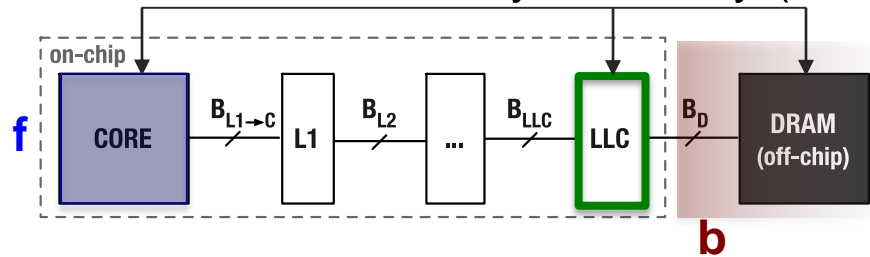
I is constant

* Williams, S., Waterman, A. and Patterson, D., "Roofline: An insightful visual performance model for multicore architectures", Communications of the ACM (2009)


The Original Roofline Model: Hands On



- **Multi-cores:** Powerful cores and memory hierarchy (caches and DRAM)



APP-L3 (data fits in L3)




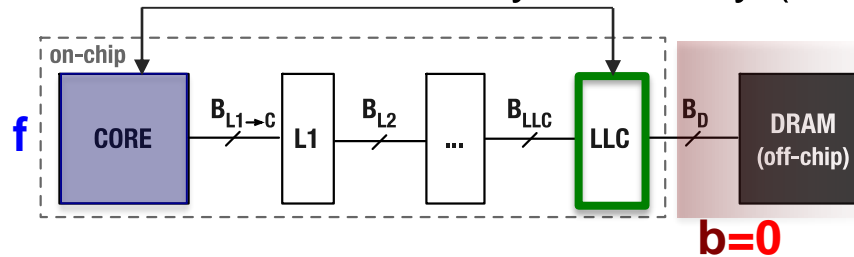
$$I_1 = f_1 / b_1$$

* Williams, S., Waterman, A. and Patterson, D., "Roofline: An insightful visual performance model for multicore architectures", Communications of the ACM (2009)


The Original Roofline Model: Hands On



- **Multi-cores:** Powerful cores and memory hierarchy (caches and DRAM)



APP-L3 (data fits in L3)



$$I_1 = f_1 / b_1$$



$$I_2 = (f_1 + f_2) / b_1$$

* Williams, S., Waterman, A. and Patterson, D., "Roofline: An insightful visual performance model for multicore architectures", Communications of the ACM (2009)


The Original Roofline Model: Hands On



- **Multi-cores:** Powerful cores and memory hierarchy (caches and DRAM)



APP-L3 (data fits in L3)



$$I_1 = f_1 / b_1$$

$$I_2 = (f_1 + f_2) / b_1$$

$$I_i = (\sum f_i) / b_1$$


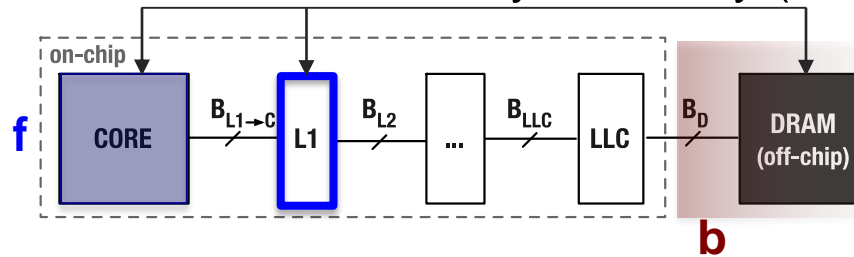
I is variable

* Williams, S., Waterman, A. and Patterson, D., "Roofline: An insightful visual performance model for multicore architectures", Communications of the ACM (2009)


The Original Roofline Model: Hands On



- **Multi-cores:** Powerful cores and memory hierarchy (caches and DRAM)



APP-L1 (data fits in L1)



$$I_1 = f_1 / b_1$$

$$I_2 = (f_1 + f_2) / b_1$$

$$I_i = (\sum f_i) / b_1$$

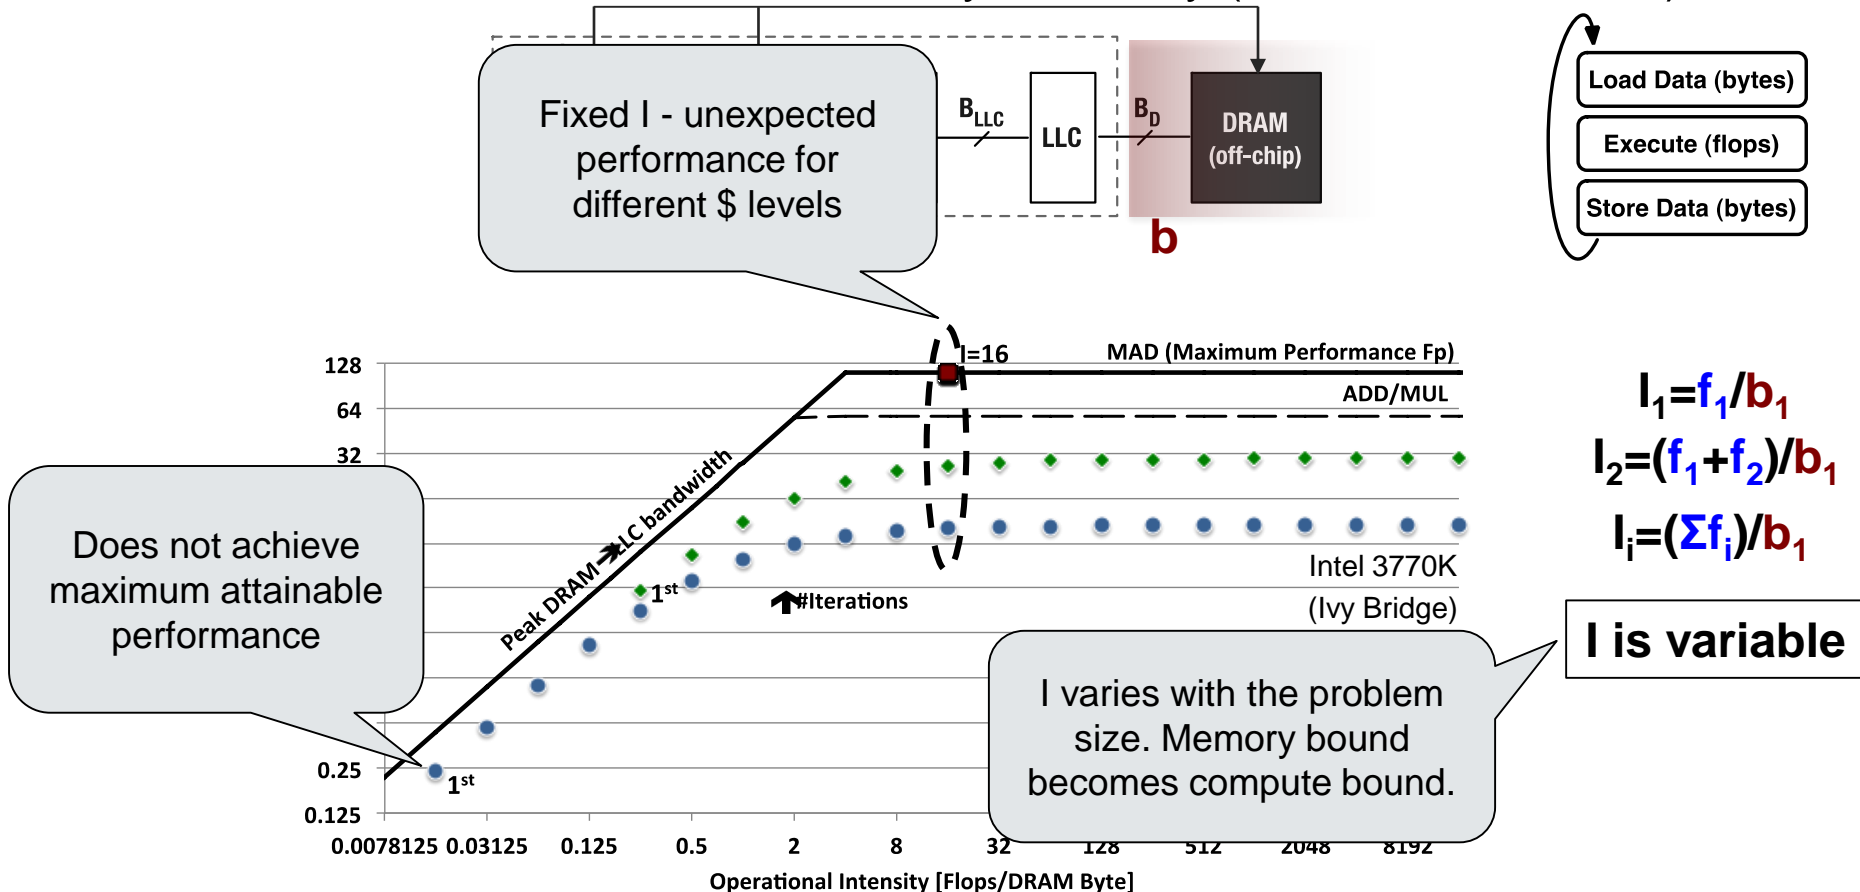
I is variable

* Williams, S., Waterman, A. and Patterson, D., "Roofline: An insightful visual performance model for multicore architectures", Communications of the ACM (2009)

The Original Roofline Model: Hands On




- **Multi-cores:** Powerful cores and memory hierarchy (caches and DRAM)





* Williams, S., Waterman, A. and Patterson, D., "Roofline: An insightful visual performance model for multicore architectures", Communications of the ACM (2009)

Cache-aware Roofline Model



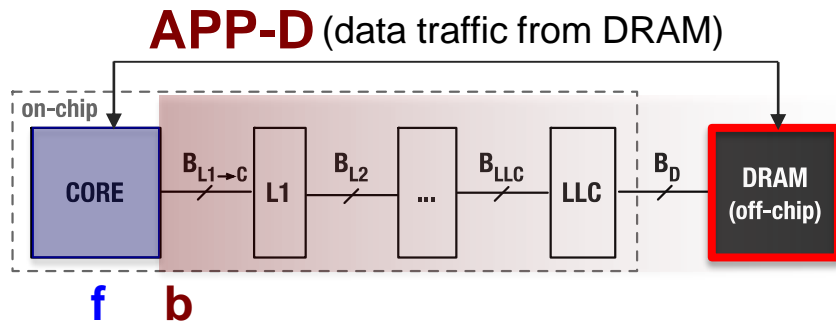

- **Multi-cores:** Powerful cores and memory hierarchy (caches and DRAM)



- **Performance:** Computations (*flops*) and communication (*bytes*) overlap in time





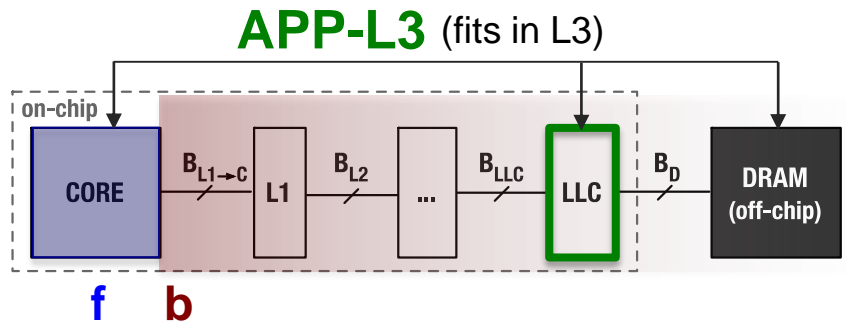

Cache-aware Roofline Model: Hands On



$$I = (\sum f_i) / (\sum b_i)$$

I is constant

Cache-aware Roofline Model: Hands On



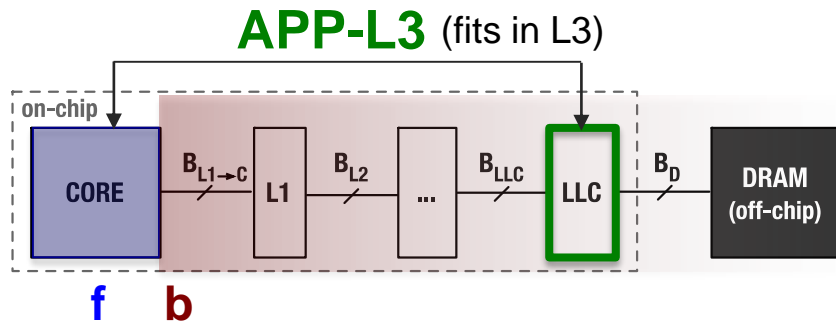



$$I = (\sum f_i) / (\sum b_i)$$

I is constant

* Ilic, A., Pratas, F. and Sousa, L., "Cache-aware Roofline Model: Upgrading the Loft", IEEE Computer Architecture Letters, 2013

Cache-aware Roofline Model: Hands On







$$I = (\sum f_i) / (\sum b_i)$$

I is constant

* Ilic, A., Pratas, F. and Sousa, L., "Cache-aware Roofline Model: Upgrading the Loft", IEEE Computer Architecture Letters, 2013

Cache-aware Roofline Model: Hands On



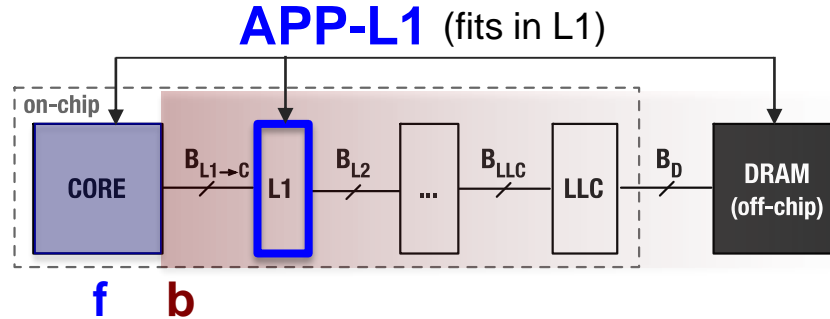



$$I = (\sum f_i) / (\sum b_i)$$

I is constant

* Ilic, A., Pratas, F. and Sousa, L., "Cache-aware Roofline Model: Upgrading the Loft", IEEE Computer Architecture Letters, 2013


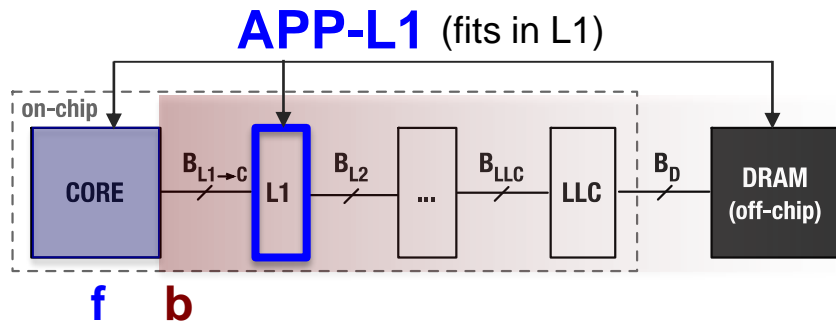

Cache-aware Roofline Model: Hands On




$$I = (\sum f_i) / (\sum b_i)$$

I is constant

Cache-aware Roofline Model: Hands On



Achieves maximum attainable performance is always memory bound.



'I' does not vary. The performance tends to the cache level ceiling.

$$I = (\sum f_i) / (\sum b_i)$$

I is constant

* Ilic, A., Pratas, F. and Sousa, L., "Cache-aware Roofline Model: Upgrading the Loft", IEEE Computer Architecture Letters, 2013

Cache-Aware Roofline Model vs. State-of-the-Art




WHAT IS HOT (WHAT IS NOT)
APPLICATION CHARACTERISTICS

Hands-on:
Application
Optimization


Practical Example: Dense Matrix Multiplication




Cache-aware Roofline Model



Original Roofline Model




```
// matrix multiplication example
for i=1 to M
  for j=1 to N
    for k=1 to K
      C[i,j] += A[i,k]*B[k,j]
```




1) **Basic implementation:** All matrices stored in row-major order.

Practical Example: Dense Matrix Multiplication


Cache-aware Roofline Model



Original Roofline Model




```
// matrix multiplication example
for i=1 to M
  for j=1 to N
    for k=1 to K
      C[i,j] += A[i,k]*B[k,j]
```




1) **Basic implementation:** All matrices stored in row-major order.

Practical Example: Dense Matrix Multiplication


Cache-aware Roofline Model



Original Roofline Model




```
// matrix multiplication example
for i=1 to M
  for j=1 to N
    for k=1 to K
      C[i,j] += A[i,k]*B[k,j]
```



1) **Basic implementation:** All matrices stored in row-major order.


Practical Example: Dense Matrix Multiplication

Cache-aware Roofline Model




- app in the **compute bound** region
- mainly limited by DRAM
- can be optimized to hit higher cache levels

Original Roofline Model



- app in the **memory bound** region
- mainly limited by DRAM
- can be optimized up to the slanted part


```
// matrix multiplication example
// VER 1: Row major matrices
for i=1 to M
  for j=1 to N
    for k=1 to K
      C[i,j] += A[i,k]*B[k,j]
```



1) **Basic implementation:** All matrices stored in row-major order.


Practical Example: Dense Matrix Multiplication

Cache-aware Roofline Model




- app in the **compute bound** region
- almost hits **L3**
- **can be further optimized** to hit higher cache levels

Original Roofline Model



- app in the **memory bound** region
- performance *hits the roof* of the model
- suggests that the **optimization is finished**


```
// matrix multiplication example
// VER 1: Row major matrices
// OPT 2: Transpose B matrix
for i=1 to M
  for j=1 to N
    for k=1 to K
      C[i,j] += A[i,k]*B[k,j]
```



2) **Transposition:** One matrix is transposed into column-major


Practical Example: Dense Matrix Multiplication

Cache-aware Roofline Model




- performance is further improved
- breaking the cache level ceilings towards the roof

Original Roofline Model



- optimization process finished


```
// matrix multiplication example
// OPT 3: Blocking for L3
for i=1 to M
  for j=1 to N
    for k=1 to K
      C[i,j] += A[i,k]*B[k,j]
```



3) Blocking for L3: All matrices are blocked to efficiently exploit L3


Practical Example: Dense Matrix Multiplication

Cache-aware Roofline Model




- performance is further improved
- breaking the cache level ceilings towards the roof

Original Roofline Model




- optimization process finished

```
// matrix multiplication example
// VER 1: Row major matrices
// OPT 2: Transpose B matrix
// OPT 3: Blocking for L3
// OPT 4: Blocking for L2
// OPT 5: Blocking for L1
```




Practical Example: Dense Matrix Multiplication

Cache-aware Roofline Model




- OPT 6 achieves **near theoretical performance**

Original Roofline Model




- moves to the compute bound region
(shift in operational intensity)

```
// matrix multiplication example
// VER 1: Row major matrices
// OPT 2: Transpose B matrix
// OPT 3: Blocking for L3
// OPT 4: Blocking for L2
// OPT 5: Blocking for L1
// OPT 6: Highly optimized (MKL)
```




Practical Example: Dense Matrix Multiplication

Cache-aware Roofline Model




- OPT 6 achieves **near theoretical performance**

Original Roofline Model



- moves to the compute bound region
(shift in operational intensity)



```
// matrix multiplication example
// VER 1: Row major matrices
// OPT 2: Transpose B matrix
// OPT 3: Blocking for L3
// OPT 4: Blocking for L2
// OPT 5: Blocking for L1
// OPT 6: Highly optimized (MKL)
```






Cache-aware Roofline Model: Use Cases



Application Characterization



Online Monitoring




* Antão, D., Taniça, L., Ilić, A., Pratas, F., Tomás, P., and Sousa, L., "Monitoring Performance and Power for Application Characterization with Cache-aware Roofline Model", PPAM'13

THE CACHE-AWARE ROOFLINE MODEL:

- PERFORMANCE
- POWER*
- EFFICIENCY

- Ilić, A., Pratas, F. and Sousa, L., “Beyond the Roofline: Power, Energy and Efficiency Modeling for Multicores” (submitted)



POWER MODELING FOR 3 DIFFERENT DOMAINS (RAPL-BASED):

1. POWER OF CORES (P_C)

- consumed by components within the cores

2. UNCORE POWER (P_U)

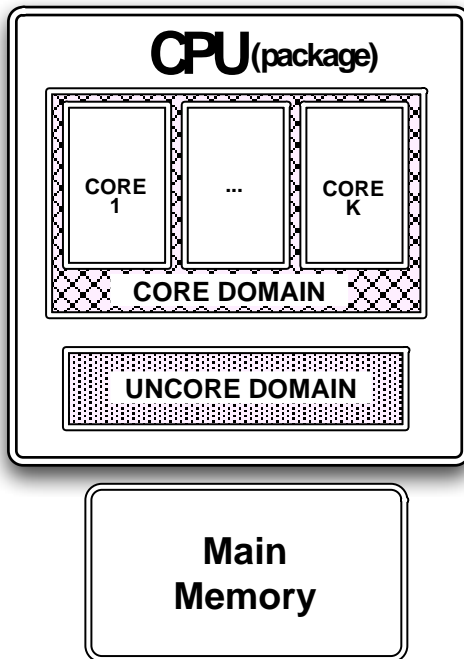
- consumed by all other (non-processing) parts of the chip, e.g., off-chip memory controller

3. PACKAGE POWER (P_P)

- the power of the complete processor chip

Power Roofline Model

1. POWER OF CORES (P_C) 2. UNCORE POWER (P_U) 3. PACKAGE POWER (P_P)



POWER ROOFLINE MODEL RELATES
POWER CONSUMPTION WITH OPERATIONAL
INTENSITY ($I=f/b$)

⇒ Average Power Consumption must be considered
– during the *time interval*, $T(I)$, in which the (Roofline) performance is obtained!

⇒ Power Contributions of both **memory operations** and **FP operations** vary with two factors:

1. The number of executed operations
2. The contribution of each during the time interval $T(I)$


Power Roofline Model




1. POWER OF CORES (P_C)

2. UNCORE POWER (P_U)

3. PACKAGE POWER (P_P)




$$I = f/b = 1/128$$




Power Roofline Model



- 1. POWER OF CORES (P_C)
- 2. UNCORE POWER (P_U)
- 3. PACKAGE POWER (P_P)



$l = f/b = 1/128$





Power Roofline Model




1. POWER OF CORES (P_C)

2. UNCORE POWER (P_U)

3. PACKAGE POWER (P_P)




$l = f/b = 1/128$




Power Roofline Model



- 1. POWER OF CORES (P_C)
- 2. UNCORE POWER (P_U)
- 3. PACKAGE POWER (P_P)



$I = f/b = 4/128$





Power Roofline Model




1. POWER OF CORES (P_C)

2. UNCORE POWER (P_U)

3. PACKAGE POWER (P_P)



$$I = f/b = 4/128$$




Power Roofline Model





1. POWER OF CORES (P_C)

2. UNCORE POWER (P_U)

3. PACKAGE POWER (P_P)



$$I = f/b = 4/128$$





Power Roofline Model




1. POWER OF CORES (P_C)

2. UNCORE POWER (P_U)

3. PACKAGE POWER (P_P)



$$I = f/b = 16/128$$





Power Roofline Model




1. POWER OF CORES (P_C)

2. UNCORE POWER (P_U)

3. PACKAGE POWER (P_P)



$$I = f/b = 16/128$$





Power Roofline Model




1. POWER OF CORES (P_C)

2. UNCORE POWER (P_U)

3. PACKAGE POWER (P_P)



$I = f/b = 64/128$






Power Roofline Model



1. POWER OF CORES (P_C)

2. UNCORE POWER (P_U)

3. PACKAGE POWER (P_P)




Power Roofline Model




1. POWER OF CORES (P_C)

2. UNCORE POWER (P_U)


3. PACKAGE POWER (P_P)



MAXIMUM OVERLAP = MAXIMUM POWER



$$I = f/b \approx 512/128$$





Power Roofline Model




1. POWER OF CORES (P_C)

2. UNCORE POWER (P_U)

3. PACKAGE POWER (P_P)



$$I = f/b = 1024/128$$





Power Roofline Model




1. POWER OF CORES (P_C)

2. UNCORE POWER (P_U)

3. PACKAGE POWER (P_P)



$I = f/b = 4096/128$





Power Roofline Model




1. POWER OF CORES (P_C)

2. UNCORE POWER (P_U)

3. PACKAGE POWER (P_P)



$$I = f/b = 65536/128$$






Power Roofline Model



1. POWER OF CORES (P_C)

2. UNCORE POWER (P_U)

3. PACKAGE POWER (P_P)





Power Roofline Model



1. POWER OF CORES (P_C)

2. UNCORE POWER (P_U)

3. PACKAGE POWER (P_P)



THE MOST “DESIRABLE” PERFORMANCE POINT IS THE WORST IN THE POWER DOMAIN


Power Roofline Model




1. POWER OF CORES (P_C)

2. UNCORE POWER (P_U)

3. PACKAGE POWER (P_P)




THE IMPACT OF THE OFF-CHIP MEMORY CONTROLLER POWER REDUCES WITH THE NUMBER OF MEMORY OPERATIONS AND THEIR CONTRIBUTION WHEN COMPARED TO THE FP OPERATIONS



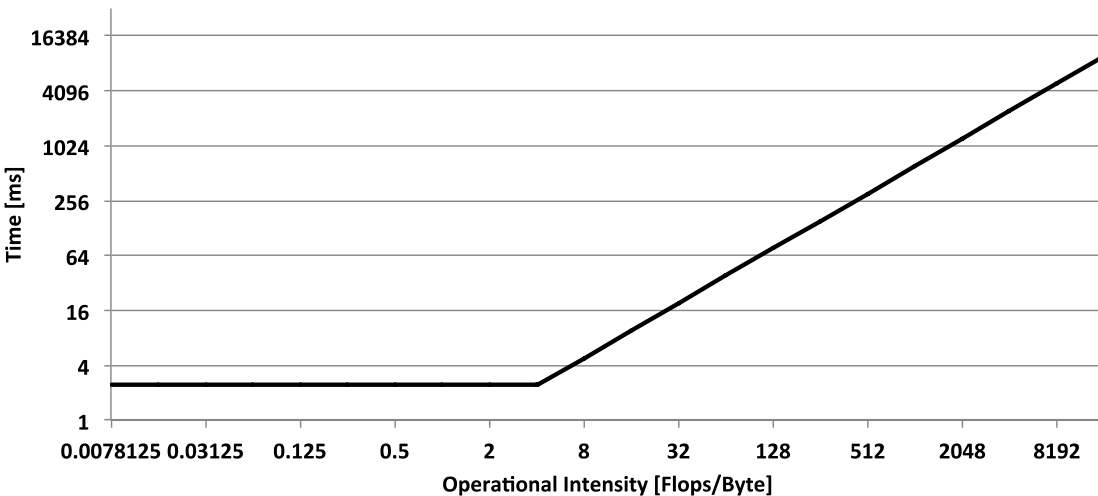
Power Roofline Model



- 1. POWER OF CORES (P_C)
- 2. UNCORE POWER (P_U)
- 3. PACKAGE POWER (P_P)





THE PACKAGE POWER DEPENDS ON SUPERPOSITION WITH THE POWER OF CORES, THE UNCORE POWER, AND THE OTHER COMPONENTS ON THE CHIP



Power Cache-aware Roofline Model





- Different power domains: $Cores + Uncore = Package$




Power Cache-aware Roofline Model




- **Performance:** Computations (*flops*) and communication (*bytes*) overlap in time
- **Power consumption:** *Superposed* contributions of *flops* and *bytes*




- **Performance:** Computations (*flops*) and communication (*bytes*) overlap in time
- **Power consumption:** Superposed contributions of *flops* and *bytes*
- **Total Power Cache-aware Roofline model**




- **Performance:** Computations (*flops*) and communication (*bytes*) overlap in time
- **Power consumption:** Superposed contributions of *flops* and *bytes*
- **Total Power Cache-aware Roofline model**



- **Performance:** Computations (*flops*) and communication (*bytes*) overlap in time
- **Power consumption:** Superposed contributions of *flops* and *bytes*
- **Total Power Cache-aware Roofline model**



- **Performance:** Computations (*flops*) and communication (*bytes*) overlap in time
- **Power consumption:** Superposed contributions of *flops* and *bytes*
- **Total Power Cache-aware Roofline model**







THE CACHE-AWARE ROOFLINE MODEL:

- PERFORMANCE
- POWER
- EFFICIENCY*

** Ilić, A., Pratas, F. and Sousa, L., "Beyond the Roofline: Power, Energy and Efficiency Modeling for Multicores"*


Efficiency Cache-Aware Roofline Model



Efficiency Cache-Aware Roofline Model




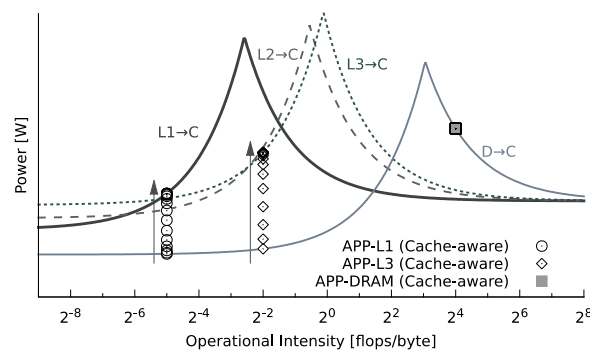

Energy-efficiency



Application Behavior






Cache-aware Roofline Models*






* Ilić, A., Pratas, F. and Sousa, L., "Cache-aware Roofline model: Upgrading the loft", IEEE Computer Architecture Letters (2013)

* Ilić, A., Pratas, F. and Sousa, L., "Beyond the Roofline: Power, Energy and Efficiency Modeling for Multicores" (submitted)

Cache-aware Roofline Models*



Original Roofline Models**



* A. Ilić, F. Pratas, and L. Sousa, "Cache-aware Roofline model: Upgrading the loft", *IEEE Computer Architecture Letters* (2013)




** S. Williams, et.al. "Roofline: An insightful visual performance model for multicore architectures", *Comm. of the ACM* (2009)

** J. Choi, D. Bedard, R. Fowler, and R. Vuduc. "A roofline model of energy", *IPDPS* (2013/2014)




Cache-aware Roofline Model: Use Cases



Application Characterization



Online Monitoring



* Ilić, A., Pratas, F. and Sousa, L., "Beyond the Roofline: Power, Energy and Efficiency Modeling for Multicores" (submitted)
 * Antão, D., Taniça, L., Ilić, A., Pratas, F., Tomás, P., and Sousa, L., "Monitoring Performance and Power for Application Characterization with Cache-aware Roofline Model", PPAM'13

BUNCH OF CACHE-AWARE ROOFLINE MODELS (EXPERIMENTALLY VERIFIED)

- **(Total) Performance**
- **(Total) Power Roofline Models:** for several domains, i.e., power of cores, uncore power and complete package power
- **Energy Roofline Model:** Time + Power Domains
- **Energy-Efficiency Roofline Model:** Performance + Power Domains
- **EDP-based Roofline Model:** Performance + Energy Domains

ALL MODELS OBTAINED WITHIN A SINGLE TEST PROCEDURE

- THE SAME TIME NEEDED AS FOR CONSTRUCTING THE PERFORMANCE ROOFLINE MODEL

FUTURE WORK

- INTEGRATION OF THE PERFORMANCE CARM IN INTEL TOOLS
- GPUS, ARMS, COMPLETE SYSTEM ...

Questions?

Thank you!

Further readings:

A. Ilic, F. Pratas, and L. Sousa, ***“Cache-aware Roofline model: Upgrading the loft”***, IEEE Computer Architecture Letters, CAL (2013)

A. Ilic, F. Pratas, and L. Sousa, ***“CARM: Cache-Aware Performance, Power and Energy-Efficiency Roofline Modeling”***, Intel CATC (2015)

L. Taniça, A. Ilic, P. Tomás, and L. Sousa, ***“SchedMon: A Performance and Energy Monitoring Tool for Modern Multi-cores”***, MuCoCoS/Euro-Par (2014)

D. Antão, L. Taniça, A. Ilic, F. Pratas, P. Tomás, and L. Sousa, ***“Monitoring Performance and Power for Application Characterization with Cache-aware Roofline Model”***, PPAM (2013)

A. Ilic, F. Pratas, and L. Sousa, ***“Beyond the Roofline: Power, Energy and Efficiency Modeling for Multicores”*** (#\$%&)

