

# Roofline Performance Model

**JaeHyuk Kwack**

Argonne National Laboratory

# Outline



**Introduction to the roofline model**



**Hands-on example on Aurora**

# Introduction to the Roofline Model


A short version of SC Tutorial Slides generated by Samuel Williams, LBNL (swwilliams@lbl.gov)  
Please join SC25 tutorial for Roofline model for more details and practices  
<https://sc25.conference-program.com/presentation/?id=tut115&sess=sess253>

**We spend millions of dollars porting  
applications to CPUs and GPUs...**

***How do we know if we are getting our  
money's worth?***


# Getting our money's worth?

- Really a question of getting good performance on application benchmarks
- Imagine profiling a mix of GPU-accelerated benchmarks ...
- Performance (GFLOP/s) alone may not be particularly insightful



# Are we getting good performance?

- We could compare performance to a CPU...
  - Speedup may seem random
  - Aren't GPUs always 10x faster than a CPU?
  - If not, what does that tell us about architecture, algorithm or implementation?
- **'Speedup' provides no insights into architecture, algorithm, or implementation.**
- **'Speedup' provides no guidance to CS, AM, applications, procurement, or vendors.**



# Are we getting good performance?

- Instead of speedup, we could take a CS approach and look at performance counters...

- Record microarchitectural events on CPUs/GPUs
- Use architecture-specific terminology
- May be broken
- We may be able to show correlation between events, but...

- ...providing actionable guidance to CS, AM, applications, or procurement can prove elusive.

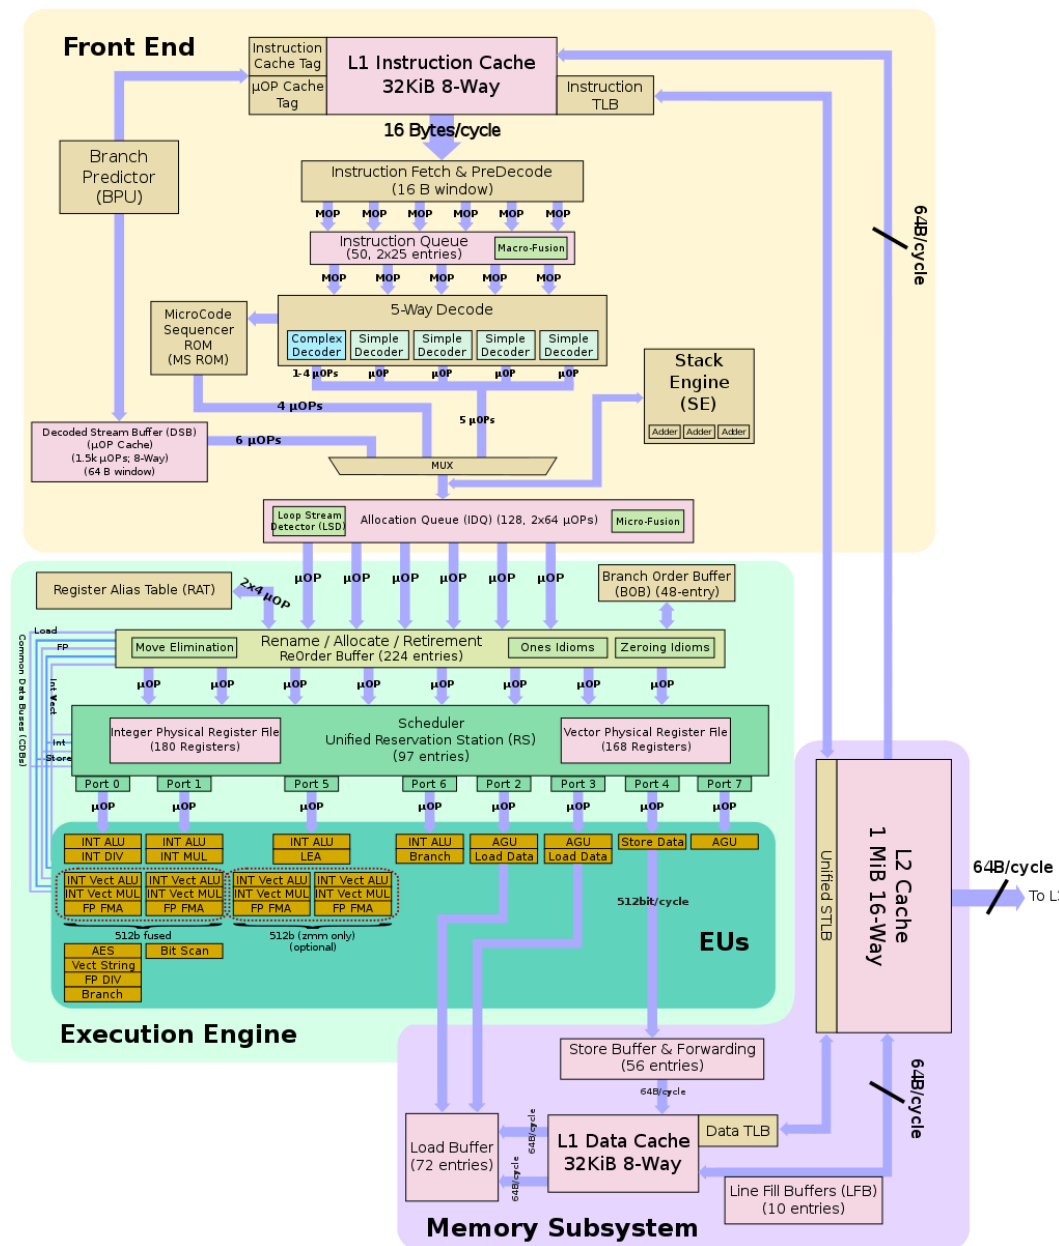
```
.  
. .  
. .  
FRONTEND_RETIRED.LATENCY_GE_8_PS  
FRONTEND_RETIRED.LATENCY_GE_16_PS  
FRONTEND_RETIRED.LATENCY_GE_32_PS  
RS_EVENTS.EMPTY_END  
FRONTEND_RETIRED.L2_MISS_PS  
FRONTEND_RETIRED.L1I_MISS_PS  
FRONTEND_RETIRED.STLB_MISS_PS  
FRONTEND_RETIRED.ITLB_MISS_PS  
ITLB_MISSES.WALK_COMPLETED  
BR_MISP_RETIRED.ALL_BRANCHES_PS  
IDQ.MS_SWITCHES  
FRONTEND_RETIRED.LATENCY_GE_2_BUBBLES_GE_1_PS  
BR_MISP_RETIRED.ALL_BRANCHES_PS  
MACHINE_CLEARS.COUNT  
MEM_LOAD_RETIRED.L1_HIT_PS  
MEM_LOAD_RETIRED.FB_HIT_PS  
MEM_LOAD_UOPS_RETIRED.L1_HIT_PS  
MEM_LOAD_UOPS_RETIRED.HIT_LFB_PS  
MEM_INST_RETIRED.STLB_MISS_LOADS_PS  
MEM_UOPS_RETIRED.STLB_MISS_LOADS_PS  
MEM_LOAD_RETIRED.L2_HIT_PSMEM_LOAD_UOPS_RETIRED.L2_HIT_PS  
MEM_LOAD_RETIRED.L3_HIT_PS  
MEM_LOAD_UOPS_RETIRED.LLC_HIT_PS  
MEM_LOAD_UOPS_RETIRED.L3_HIT_PS  
MEM_LOAD_RETIRED.L3_MISS_PS  
MEM_LOAD_UOPS_RETIRED.LLC_MISS_PS  
MEM_LOAD_UOPS_MISC_RETIRED.LLC_MISS_PS  
MEM_LOAD_UOPS_RETIRED.L3_MISS_PS  
MEM_INST_RETIRED.ALL_STORES_PS  
MEM_UOPS_RETIRED.ALL_STORES_PS  
ARITH.DIVIDER_ACTIVE  
ARITH.DIVIDER_UOPS  
ARITH.FPU_DIV_ACTIVE  
INST_RETIRED.PREC_DIST  
IDQ.MS_UOPS  
INST_RETIRED.PREC_DIST  
. . .
```

# Are we getting good performance?

- We could take the computer architect's approach and build a simulator to understand performance nuances...

- Modern architectures are incredibly complex
- Simulators may perfectly reproduce performance
- Lots of information interpretable only by computer architects
- worse, might incur  $10^6\times$  slowdowns

- Provide no insights into quality or limits of algorithm or implementation.
- Provide no guidance to CS, AM, application developers.

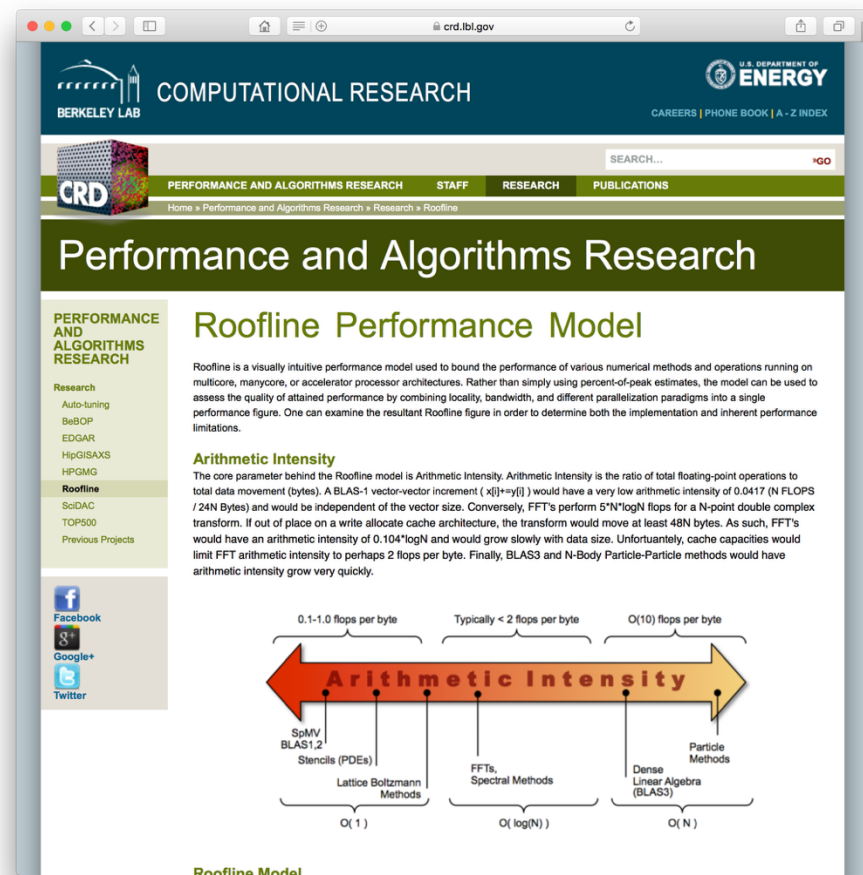




# What's missing...


- Each community speaks their own language and develops specialized tools/methodologies
- Need common mental model of application execution on target system
- Sacrifice accuracy to gain...
  - Architecture independence / extensibility
  - Readily understandable by broad community
  - Intuition, insights, and guidance to CS, AM, apps, procurement, and vendors

➤ **Roofline is just such a model**




<https://crd.lbl.gov/roofline>

- Assume HW/SW can perfectly overlap communication and computation
- Which takes longer?
  - Data Movement
  - Computation




$$\text{Time} = \max \left\{ \begin{array}{l} \# \text{FP ops} / \text{Peak GFLOP/s} \\ \# \text{Bytes} / \text{Peak GB/s} \end{array} \right.$$

- Assume HW/SW can perfectly overlap communication and computation
- Which takes longer?
  - Data Movement
  - Computation
- Is performance limited by compute or data movement?



$$\frac{\text{Time}}{\text{\#FP ops}} = \max \left\{ \begin{array}{l} 1 / \text{Peak GFLOP/s} \\ \text{\#Bytes} / \text{\#FP ops} / \text{Peak GB/s} \end{array} \right.$$

- Assume HW/SW can perfectly overlap communication and computation
- Which takes longer?
  - Data Movement
  - Computation
- Is performance limited by compute or data movement?




$$\frac{\text{\#FP ops}}{\text{Time}} = \min \left\{ \begin{array}{l} \text{Peak GFLOP/s} \\ (\text{\#FP ops} / \text{\#Bytes}) * \text{Peak GB/s} \end{array} \right.$$

- Assume HW/SW can perfectly overlap communication and computation
- Which takes longer?
  - Data Movement
  - Computation
- Is performance limited by compute or data movement?

$$\text{GFLOP/s} = \min \left\{ \begin{array}{l} \text{Peak GFLOP/s} \\ \text{AI} * \text{Peak GB/s} \end{array} \right.$$

*Arithmetic Intensity (AI) = measure of data locality*



# Arithmetic Intensity


- Measure of data locality (data reuse)
- Ratio of Total Flops performed to Total Bytes moved
- For the DRAM Roofline...
  - Total Bytes to/from DRAM
  - Includes all cache and prefetcher effects
  - Can be very different from total loads/stores (bytes requested)
  - Equal to ratio of sustained GFLOP/s to sustained GB/s (time cancels)

# (DRAM) Roofline Model

$$\text{GFLOP/s} = \min \begin{cases} \text{Peak GFLOP/s} \\ \text{AI} * \text{Peak GB/s} \end{cases}$$

AI (Arithmetic Intensity) = FLOPs / Bytes (moved to/from DRAM)

- Plot bound on **Log-log scale** as a function of AI (data locality)




# (DRAM) Roofline Model

$$\text{GFLOP/s} = \min \begin{cases} \text{Peak GFLOP/s} \\ \text{AI} * \text{Peak GB/s} \end{cases}$$

AI (Arithmetic Intensity) = FLOPs / Bytes (moved to/from DRAM)

- Plot bound on **Log-log scale** as a function of AI (data locality)
- Roofline tessellates the locality-performance plane into five regions...



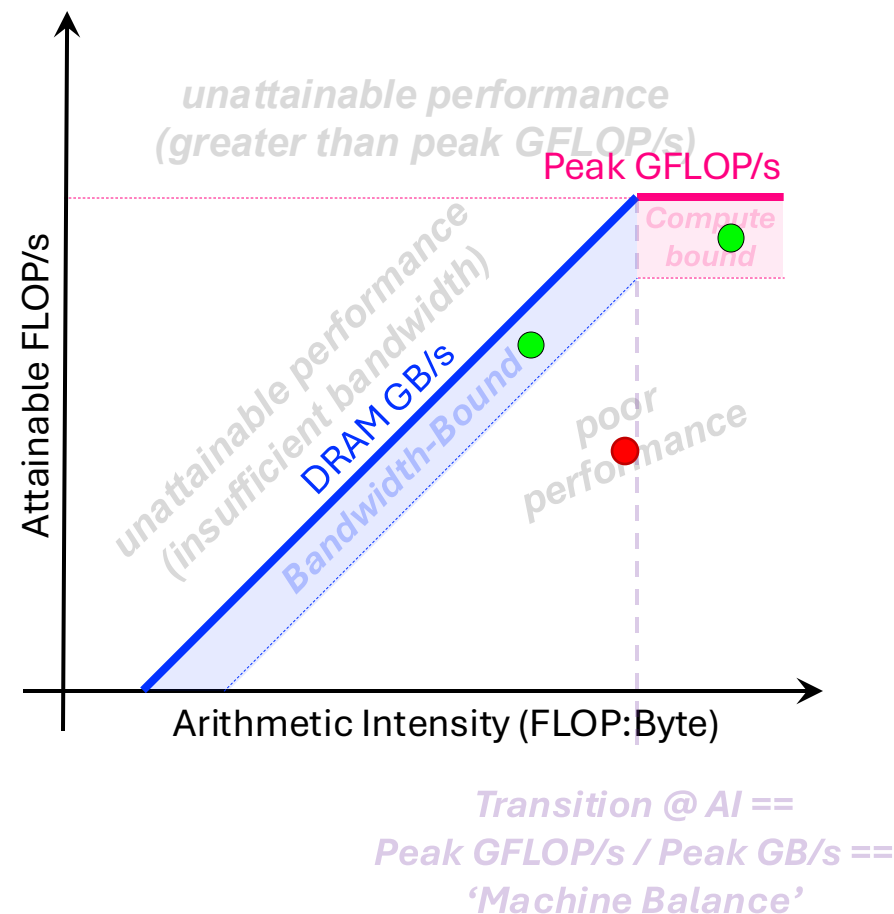


# (DRAM) Roofline Model

$$\text{GFLOP/s} = \min \begin{cases} \text{Peak GFLOP/s} \\ \text{AI} * \text{Peak GB/s} \end{cases}$$

AI (Arithmetic Intensity) = FLOPs / Bytes (moved to/from DRAM)

- Plot bound on **Log-log scale** as a function of AI (data locality)
- Roofline tessellates the locality-performance plane into five regions...
- Measure application (AI, GF/s) and plot in the 2D locality-performance plane.




# Roofline Examples

# Roofline Example #1

- Typical machine balance is 5-10 FLOPs per byte...
  - 40-80 FLOPs per double to exploit compute capability
  - Artifact of technology and money
  - **Unlikely to improve**
- Consider STREAM Triad...

```
#pragma omp parallel for
for(i=0;i<N;i++){
  Z[i] = X[i] + alpha*Y[i];
}
```

  - 2 FLOPs per iteration
  - Transfer 24 bytes per iteration (read X[i], Y[i], write Z[i])
  - **AI = 0.083 FLOPs per byte == Memory bound**

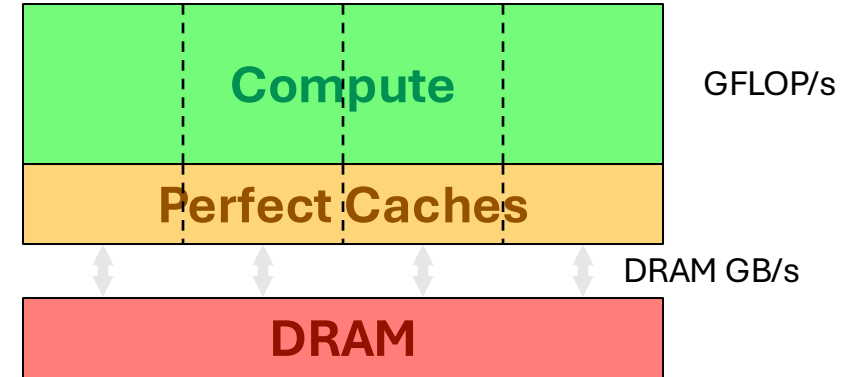


# Roofline Example #2

- Conversely, 7-point constant coefficient stencil...

```
#pragma omp parallel for
for(k=1;k<dim+1;k++){
for(j=1;j<dim+1;j++){
for(i=1;i<dim+1;i++){
    new[k][j][i] = -6.0*old[k ][j ][i ]
                  + old[k ][j ][i-1]
                  + old[k ][j ][i+1]
                  + old[k ][j-1][i ]
                  + old[k ][j+1][i ]
                  + old[k-1][j ][i ]
                  + old[k+1][j ][i ];
}}}

```




# Roofline Example #2

## ■ Conversely, 7-point constant coefficient stencil...

- 7 FLOPs
- 8 memory references (7 reads, 1 store) per point
- **$AI = 7 / (8 \cdot 8) = 0.11$  FLOPs per byte**  
**(measured at the L1)**


```
#pragma omp parallel for
for(k=1;k<dim+1;k++){
  for(j=1;j<dim+1;j++){
    for(i=1;i<dim+1;i++){
      new[k][j][i] = -6*old[k][j][i]
        + old[k][j][i-1]
        + old[k][j][i+1]
        + old[k][j-1][i]
        + old[k][j+1][i]
        + old[k-1][j][i]
        + old[k+1][j][i]
    }
  }
}
```



# Roofline Example #2

- Conversely, 7-point constant coefficient stencil...

- 7 FLOPs
- 8 memory references (7 reads, 1 store) per point
- Ideally, cache will filter all but 1 read and 1 write per point



```
#pragma omp parallel for
for(k=1;k<dim+1;k++){
for(j=1;j<dim+1;j++){
  for(i=1;i<dim+1;i++){
    new[k][j][i] = -6.0*old[k][j][i]
                  + old[k][j][i-1]
                  + old[k][j][i+1]
                  + old[k][j-1][i]
                  + old[k][j+1][i]
                  + old[k-1][j][i]
                  + old[k+1][j][i]
  }
}
```


# Roofline Example #2

## ■ Conversely, 7-point constant coefficient stencil...

- 7 FLOPs
- 8 memory references (7 reads, 1 store) per point
- Ideally, cache will filter all but 1 read and 1 write per point
- **$7 / (8+8) = 0.44$  FLOPs per byte (DRAM)**

```
#pragma omp parallel for
for(k=1;k<dim+1;k++){
for(j=1;j<dim+1;j++){
for(i=1;i<dim+1;i++){
    new[k][j][i] = -6.0*old[k ][j ][i ]
                  + old[k ][j ][i-1]
                  + old[k ][j ][i+1]
                  + old[k ][j-1][i ]
                  + old[k ][j+1][i ]
                  + old[k-1][j ][i ]
                  + old[k+1][j ][i ];
}}}

```



# Roofline Example #2


## ■ Conversely, 7-point constant coefficient stencil...

- 7 FLOPs
- 8 memory references (7 reads, 1 store) per point
- Ideally, cache will filter all but 1 read and 1 write per point
- **$7 / (8+8) = 0.44$  FLOPs per byte (DRAM)**

**== memory bound, but 5x the FLOP rate as TRIAD**

```
#pragma omp parallel for
for(k=1;k<dim+1;k++){
for(j=1;j<dim+1;j++){
for(i=1;i<dim+1;i++){
    new[k][j][i] = -6.0*old[k][j][i]
                  + old[k][j][i-1]
                  + old[k][j][i+1]
                  + old[k][j-1][i]
                  + old[k][j+1][i]
                  + old[k-1][j][i]
                  + old[k+1][j][i];
}}}

```



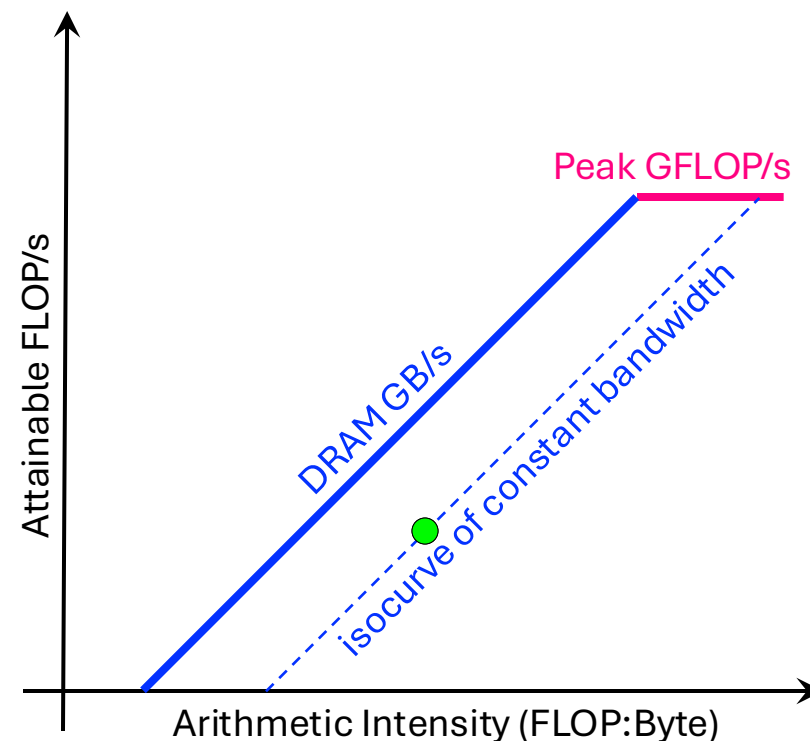


# Roofline Example #3

- Roofline makes it obvious what the bound on FLOP rate is, but let's ask the reverse...
- *Given a low FLOP rate and AI, what DRAM bandwidth are we attaining?*


$$\text{average GB/s} = \frac{\text{GFLOP/s}}{\text{AI}_{\text{DRAM}}}$$

- This is just a slope (y/x)
- Thus we can define an isocurve of constant bandwidth



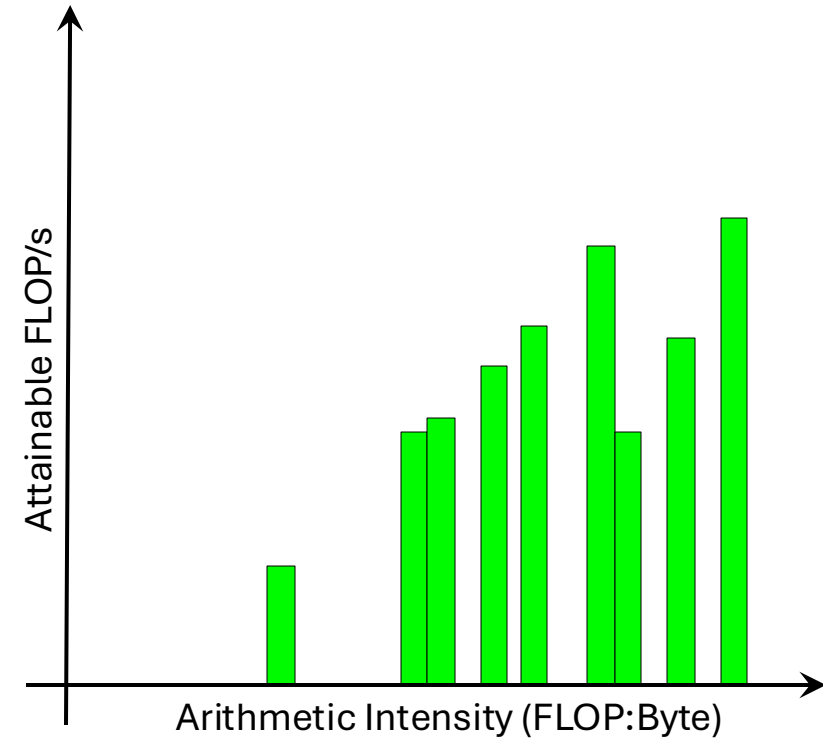
# Are we getting good performance?

- Think back to our mix of benchmarks...




# Are we getting good performance?

- We can sort benchmarks by arithmetic intensity...




# Are we getting good performance?

- We can sort benchmarks by arithmetic intensity...
- ... and compare performance relative to machine capabilities



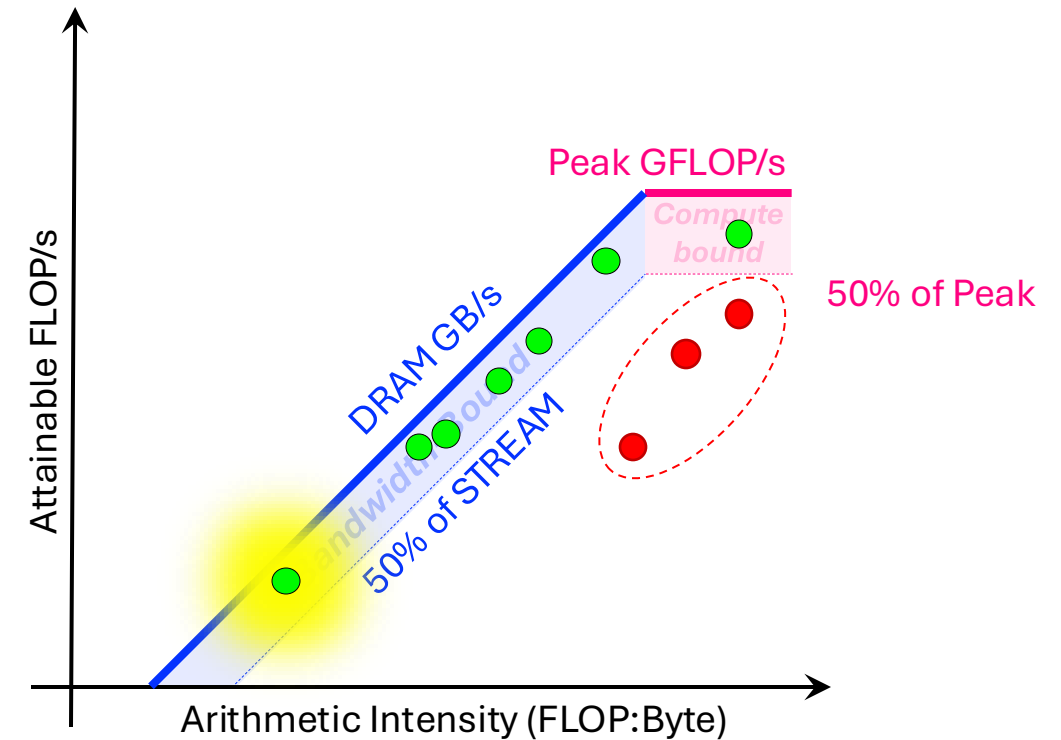
# Are we getting good performance?

- Benchmarks near the roofline are making **good use** of computational resources



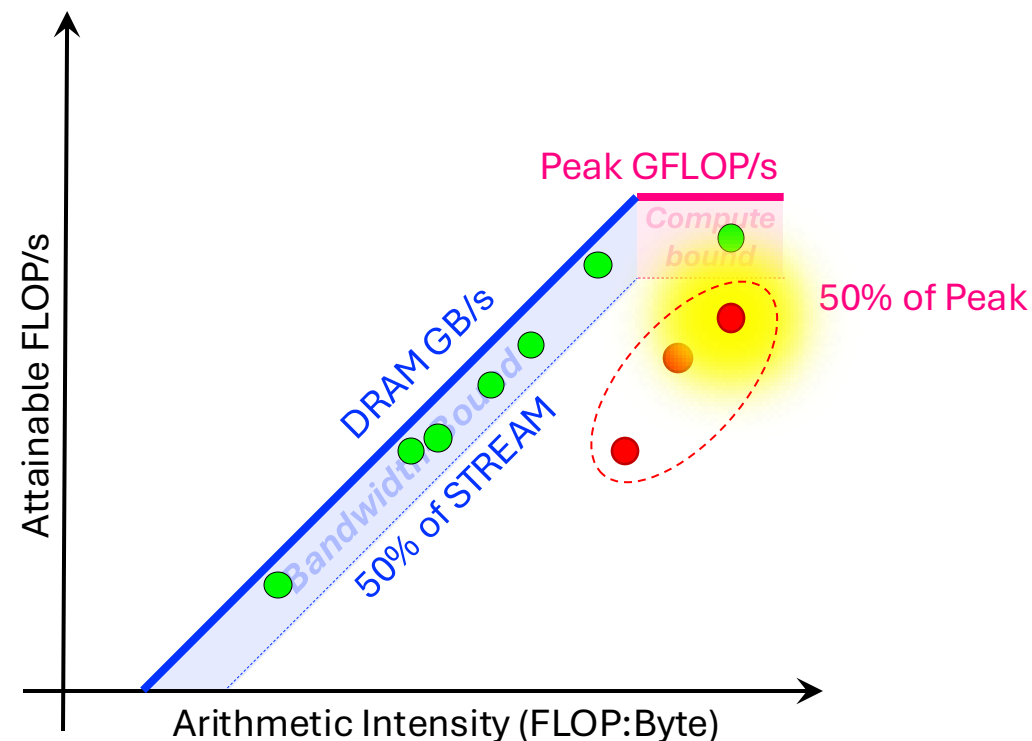
# Are we getting good performance?

- Benchmarks near the roofline are making **good use** of computational resources
  - benchmarks can have low performance (GFLOP/s), but make good use (%STREAM) of a machine



# Are we getting good performance?


- Benchmarks near the roofline are making **good use** of computational resources
  - benchmarks can have low performance (GFLOP/s), but make good use (%STREAM) of a machine
  - benchmarks can have high performance (GFLOP/s), but still make poor use of a machine (%peak)



# Recap: Roofline is made of two components

## ■ Machine Model

- Lines defined by peak GB/s and GF/s (**Benchmarking**)
- Unique to each architecture
- Common to all apps on that architecture






# Recap: Roofline is made of two components

## ■ Machine Model

- Lines defined by peak GB/s and GF/s (**Benchmarking**)
- Unique to each architecture
- Common to all apps on that architecture

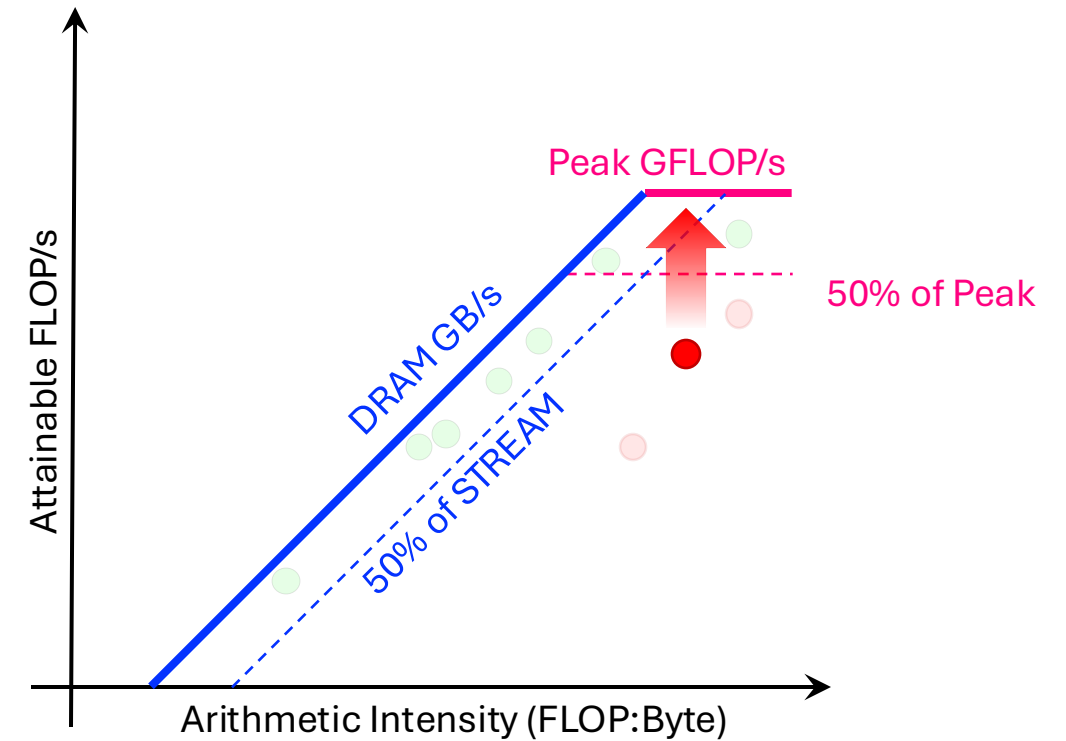
## ■ Application Characteristics

- Dots defined by application GFLOP's and GB's (**Application Instrumentation**)
- Unique to each application
- Unique to each architecture



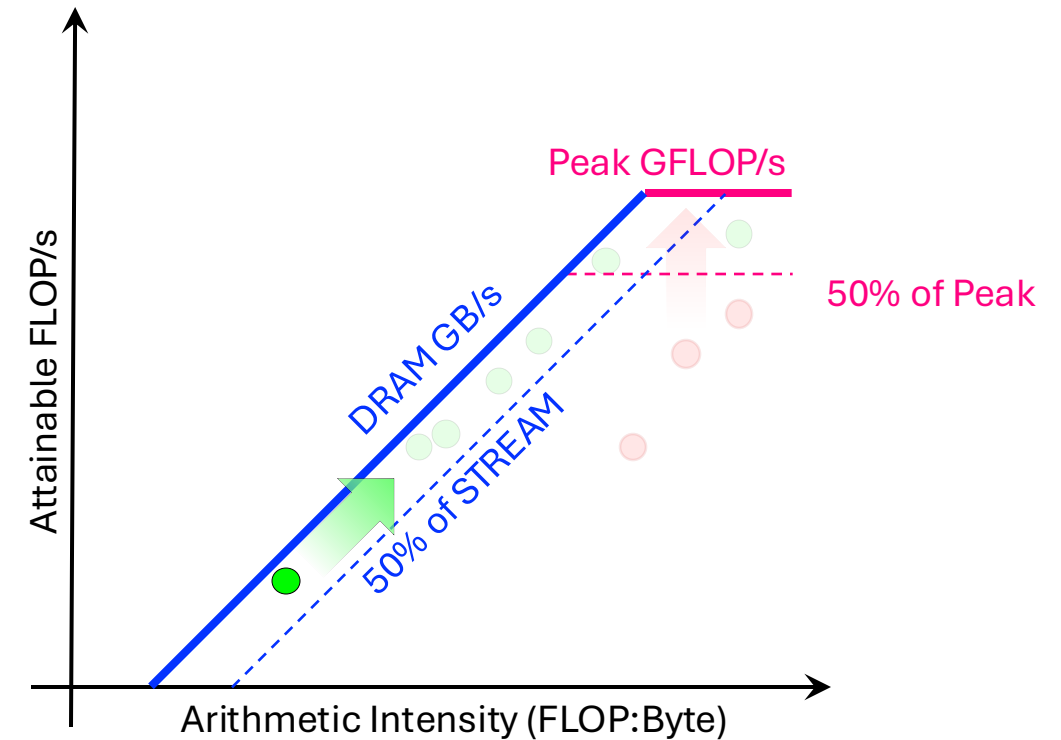
# Recap: Optimization Strategy

## 1. Get to the Roofline



# Recap: Optimization Strategy

1. Get to the Roofline
2. Increase Arithmetic Intensity when bandwidth-limited
  - Reducing data movement increases AI
  - Increasing AI increases performance when bandwidth-bound



**How can performance ever be  
below the Roofline?**

# How can performance be below the Roofline?

*Simple DRAM model can be insufficient for a variety of reasons...*

## DRAM's not the bottleneck...

- Cache bandwidth and cache locality
- PCIe bandwidth

## Not enough of Vector/Tensor instr.


- No FMA
- Mixed Precision
- No Tensor Core OPs

## Integer-heavy Codes... Lack of Parallelism...

- Non-FP inst. impede FLOPs
- No FP instructions
- Idle Cores/SMs
- Insufficient ILP/TLP
- Divergence and Predication


## ...The Hierarchical Roofline Model

T. Koskela, Z. Matveev, C. Yang, A. Adedoyin, R. Belenov, P. Thierry, Z. Zhao, R. Gayatri, H. Shan, L. Oliker, J. Deslippe, R. Green, S. Williams, "A Novel Multi-Level Integrated Roofline Model Approach for Performance Characterization", ISC, 2018.



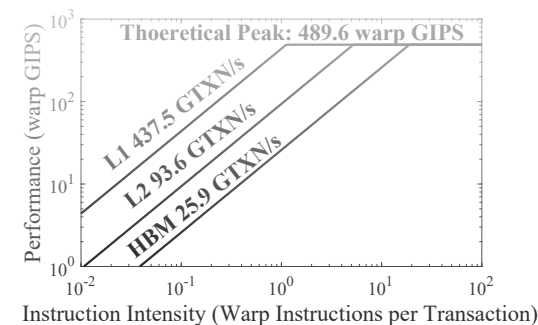
## ... Additional Ceilings

C. Yang, T. Kurth, S. Williams, "Hierarchical Roofline analysis for GPUs: Accelerating performance optimization for the NERSC-9 Perlmutter system", CCPE, 2019.




## ... The Instruction Roofline Model

N. Ding, S. Williams, "An Instruction Roofline Model for GPUs", BEST PAPER, PMBS, 2019.



## ... Roofline Scaling Trajectories

K. Ibrahim, S. Williams, L. Oliker, "Performance Analysis of GPU Programming Models using the Roofline Scaling Trajectories", BEST PAPER, Bench, 2019.




# Below the Roofline?

## Memory Hierarchy and Cache Bottlenecks


# Memory Hierarchy

- CPUs/GPUs have multiple levels of memory/cache
  - Registers
  - L1, L2, L3 cache
  - HBM (KNL/GPU device memory)
  - DDR (main memory)
  - NVRAM (non-volatile memory)



# Memory Hierarchy


- CPUs/GPUs have different bandwidths for each level





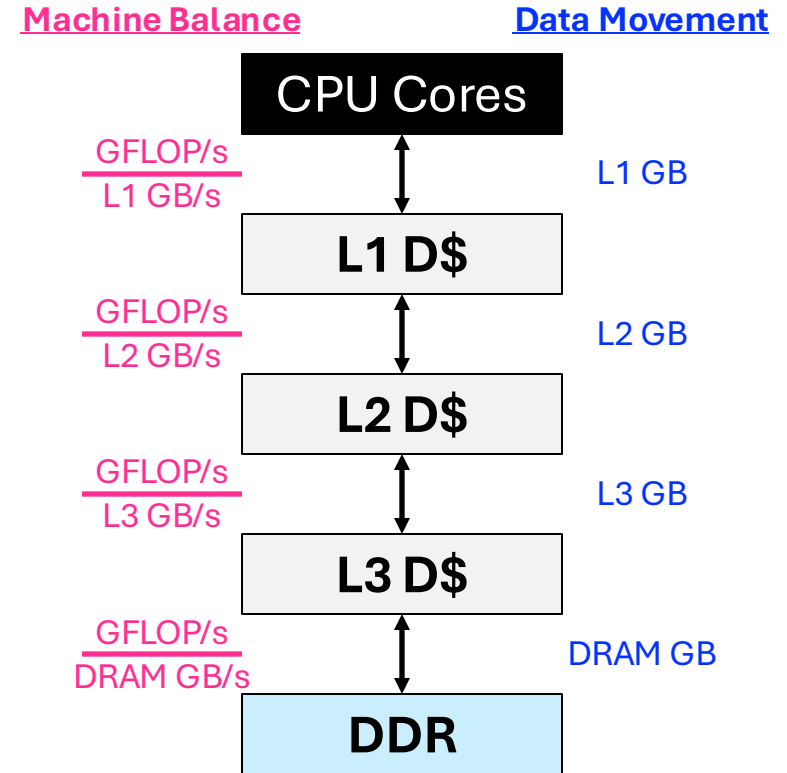
# Memory Hierarchy

- CPUs/GPUs have different bandwidths for each level
  - different machine balances for each level



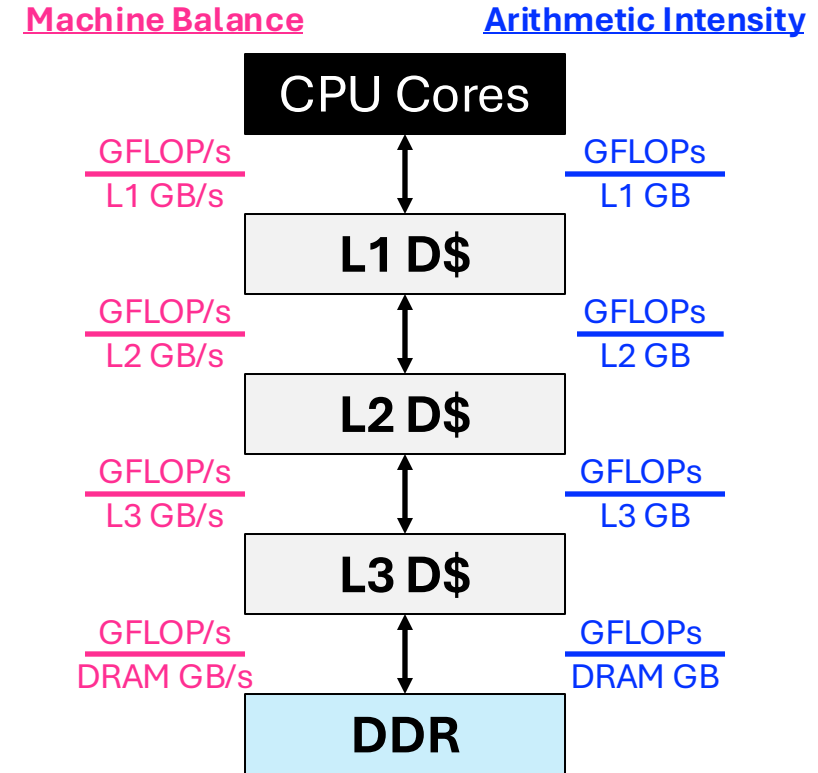
# Memory Hierarchy

- CPUs/GPUs have different bandwidths for each level
  - different machine balances for each level
- Applications have locality in each level
  - different data movements for each level




# Memory Hierarchy

- CPUs/GPUs have different bandwidths for each level
  - different machine balances for each level
- Applications have locality in each level
  - different data movements for each level
  - different arithmetic intensity for each level



# Memory Hierarchy (Discrete GPU)

- CPUs/GPUs have different bandwidths for each level
  - different machine balances for each level
- Applications have locality in each level
  - different data movements for each level
  - different arithmetic intensity for each level
- Same concept applies to GPUs and disaggregated memory
  - DDR is accessed via PCIe, CXL, or NoC



# Cache Bottlenecks

- For each additional level of the memory hierarchy, we can add another term to our model...

$$\text{GFLOP/s} = \min \left\{ \begin{array}{l} \text{Peak GFLOP/s} \\ A_{\text{DRAM}} * \text{DRAM GB/s} \end{array} \right.$$

$A_x$  (Arithmetic Intensity at level “x”) = FLOPs / Bytes (moved to/from level “x”)

# Cache Bottlenecks

- For each additional level of the memory hierarchy, we can add another term to our model...

$$\text{GFLOP/s} = \min \left\{ \begin{array}{l} \text{Peak GFLOP/s} \\ A_{\text{DRAM}} * \text{DRAM GB/s} \\ A_{\text{L2}} * \text{L2 GB/s} \end{array} \right.$$

$A_x$  (Arithmetic Intensity at level “x”) = FLOPs / Bytes (moved to/from level “x”)

# Cache Bottlenecks


- For each additional level of the memory hierarchy, we can add another term to our model...

$$\text{GFLOP/s} = \min \left\{ \begin{array}{l} \text{Peak GFLOP/s} \\ A_{\text{DRAM}} * \text{DRAM GB/s} \\ A_{\text{L2}} * \text{L2 GB/s} \\ A_{\text{L1}} * \text{L1 GB/s} \end{array} \right.$$

$A_x$  (Arithmetic Intensity at level “x”) = FLOPs / Bytes (moved to/from level “x”)

# Cache Bottlenecks


- Plot equation in a single figure...
  - “**Hierarchical Roofline**” Model






# Cache Bottlenecks

- Plot equation in a single figure...
  - “**Hierarchical Roofline**” Model
  - Bandwidth ceiling (diagonal line) for each level of memory



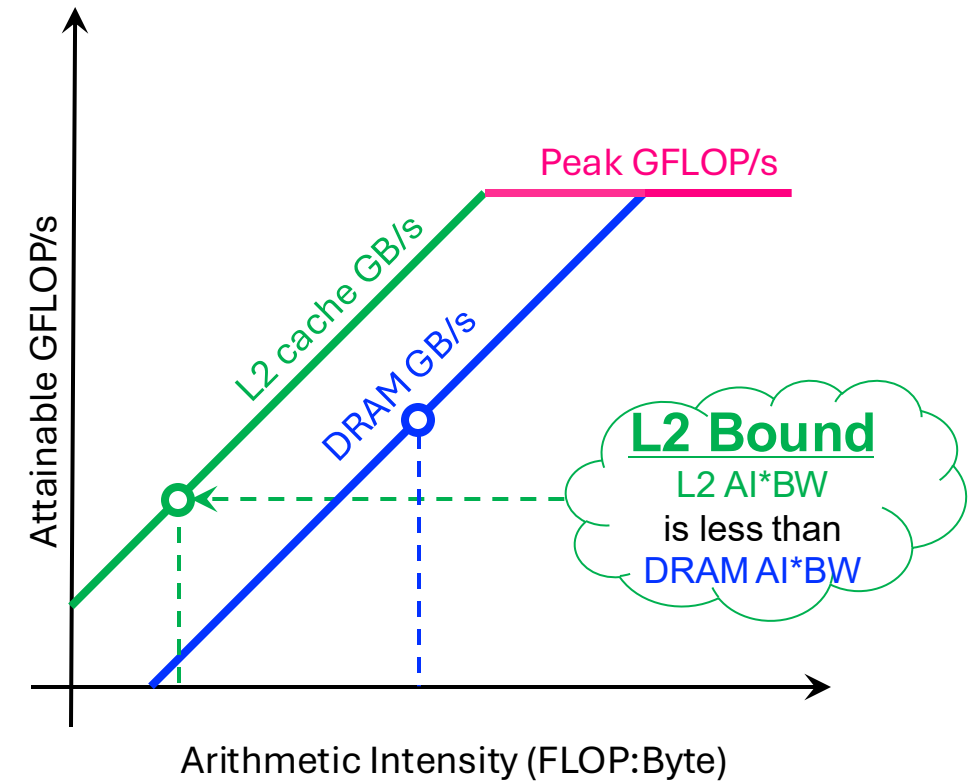
# Cache Bottlenecks

- Plot equation in a single figure...
  - “**Hierarchical Roofline**” Model
  - Bandwidth ceiling (diagonal line) for each level of memory
  - Arithmetic Intensity (dot) for each level of memory



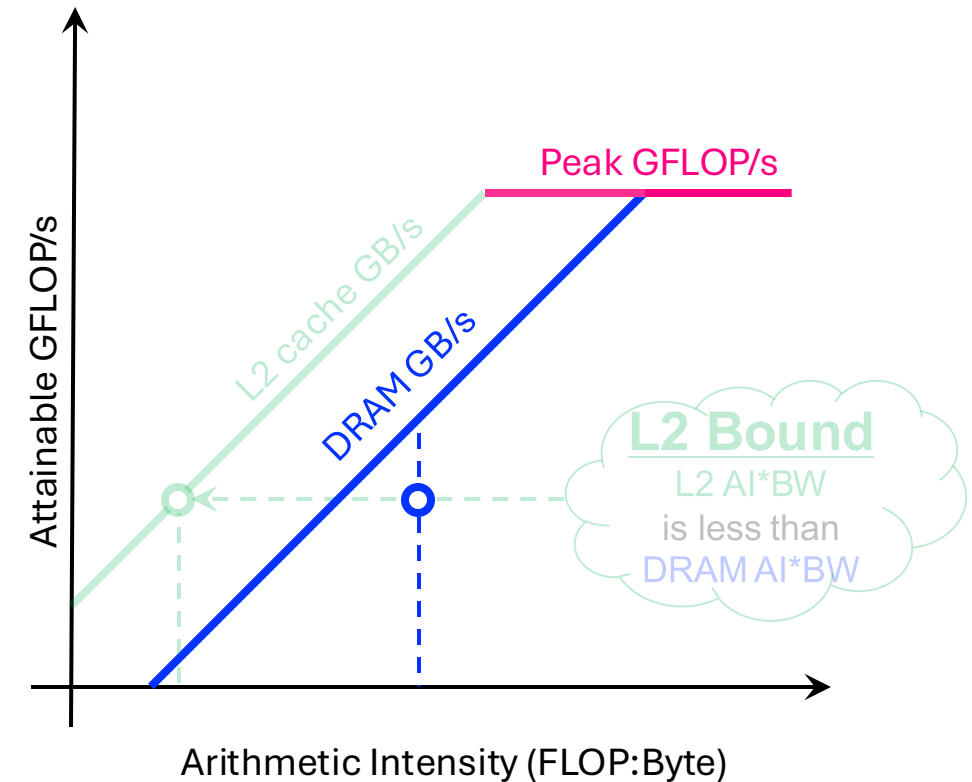
# Cache Bottlenecks

- Plot equation in a single figure...
  - “**Hierarchical Roofline**” Model
  - Bandwidth ceiling (diagonal line) for each level of memory
  - Arithmetic Intensity (dot) for each level of memory
  - **performance is ultimately the minimum of these bounds**




# Cache Bottlenecks

- Plot equation in a single figure...
  - “**Hierarchical Roofline**” Model
  - Bandwidth ceiling (diagonal line) for each level of memory
  - Arithmetic Intensity (dot) for each level of memory
  - **performance is ultimately the minimum of these bounds**
- **If L2 bound, we see DRAM dot well below DRAM ceiling**




# Cache Hit Rates

- Widely separated Arithmetic Intensities indicate high reuse in the (L2) cache



# Cache Hit Rates

- Widely separated Arithmetic Intensities indicate high reuse in the (L2) cache
- Similar Arithmetic Intensities indicate effectively no (L2) cache reuse (**== streaming**)



# Below the Roofline?

## Fused Operations and Accelerators


# Fused Operations and Accelerators

- Vectors have their limits (finite DLP, register file energy scales with VL, etc...)
  - Death of Moore's Law is incentivizing operator fusion (e.g. FMA) and compute accelerators (matrix multipliers)
  - Modern CPUs and GPUs are increasingly reliant on special (fused) instructions that perform multiple operations (fuse common instruction sequences)...
    - FMA (Fused Multiply Add):  $z = a * x + y$  ...  $z, x, y$  are vectors or scalars
    - 4FMA (Quad FMA):  $z = A * x + z$  ...  $A$  is a FP32 matrix;  $x, z$  are vectors
    - WMMA (Tensor Core):  $Z = AB + C$  ...  $A, B$  are FP16 matrices;  $Z, C$  are FP32
- **Define a set of “ceilings” based on instruction type (all tensor, all FMA, or all FADD)**




# Floating-Point and Mixed Precision Ceilings

- Consider NVIDIA Volta GPU
- We may define 3 performance ceilings...
  - 15 TFLOPS for FP32 FMA
  - 7.5 TFLOPs for FP32 Add
  - ~100 TFLOPs for FP16 Tensor




# Floating-Point and Mixed Precision Ceilings

- When calculating (AI,GFLOP/s), count the total FLOPs from all types of instructions
- DL performance can often be well below nominal Tensor Core peak



# Floating-Point and Mixed Precision Ceilings

- When calculating (AI,GFLOP/s), count the total FLOPs from all types of instructions
- DL performance can often be well below nominal Tensor Core peak
- DL applications are a mix Tensor, FP16, and FP32 instructions
- Thus, there is a ceiling on performance defined by the mix of instructions



# Below the Roofline?

## Lack of Parallelism

# Roofline and Parallelism

- We've assumed we can always hit either peak GFLOP/s or peak GB/s

$$\text{GFLOP/s} = \min \left\{ \begin{array}{l} \text{GFLOP/s}_{\text{Peak}} \\ \text{AI}_{\text{DRAM}} * \text{GB/s}_{\text{DRAM}} \end{array} \right.$$

$\text{AI}_x$  (Arithmetic Intensity at level “x”) = FLOPs / Bytes (moved to/from level “x”)

# Roofline and Parallelism

- We've assumed we can always hit either peak GFLOP/s or peak GB/s
- But all CPUs and GPUs are highly parallel architectures
- GFLOP/s and GB/s are a function of how much parallelism we utilize...


$$\text{GFLOP/s}(\mathbf{P}) = \min \left\{ \begin{array}{l} \text{GFLOP/s}_{\text{Peak}}(\mathbf{P}) \\ \text{AI}_{\text{DRAM}}(\mathbf{P}) * \text{GB/s}_{\text{DRAM}}(\mathbf{P}) \end{array} \right.$$

$\text{AI}_x$  (Arithmetic Intensity at level “x”) = FLOPs / Bytes (moved to/from level “x”)

*$\text{AI}_{\text{DRAM}}$  is a function of parallelism because cache contention can generate superfluous LLC capacity misses (==DRAM data movement)*

# Roofline and Parallelism

- How do we visualize parallelism in the Roofline?
  - Naively, GFLOP/s(P) and GB/s(P) are proportional to parallelism P
  - SMs are capable of pulling more than their fair share of HBM
  - DVFS implies not true for GFLOP/s




# Roofline and Parallelism

## ■ How do we visualize parallelism in the Roofline?

- Naively, GFLOP/s(P) and GB/s(P) are proportional to parallelism P
- SMs are capable of pulling more than their fair share of HBM
- DVFS implies not true for GFLOP/s


➤ **one must benchmark GFLOP/s and GB/s at each concurrency**






# Roofline and Parallelism

- Consider CUDA kernel optimized for Fermi (16 SMs) running on Volta (80 SMs)
  - Performance looks very poor



# Roofline and Parallelism


- Consider CUDA kernel optimized for Fermi (16 SMs) running on Volta (80 SMs)
  - Performance looks very poor
  - Kernels using only 16 SMs underutilize the V100 architecture.
  - Roofline highlights the fact that performance is constrained by a lack of software parallelism



# Roofline Scaling Trajectories

## ■ Traditional Scalability:

- Plot performance vs. concurrency (#cores or #SMs)
- Observation without much insight
- Why does performance decrease?




# Roofline Scaling Trajectories

- Khaled Ibrahim leveraged Roofline to understand the interplay between concurrency, data locality, and performance

## ➤ Roofline Scaling Trajectories

- Measure (AI,GFLOP/s) for each concurrency
- Plot as a trendline on Roofline




# Roofline Scaling Trajectories

- Khaled Ibrahim leveraged Roofline to understand the interplay between concurrency, data locality, and performance

## ➤ Roofline Scaling Trajectories

- Measure (AI, GFLOP/s) for each concurrency
- Plot as a trendline on Roofline
- **Perfect scaling is a vertical line**




# Roofline Scaling Trajectories

- Khaled Ibrahim leveraged Roofline to understand the interplay between concurrency, data locality, and performance

## ➤ Roofline Scaling Trajectories

- Measure (AI, GFLOP/s) for each concurrency
- Plot as a trendline on Roofline
- Perfect scaling is a vertical line
- **Turnover in AI indicates cache capacity exhaustion (extra L2 misses drives down AI)**



# Recap

# Recap

- Roofline bounds performance as a function of Arithmetic Intensity
  - Horizontal Lines = Compute Ceilings
  - Diagonal Lines = Bandwidth Ceilings
  - Bandwidth ceilings are always parallel on log-log scale
  - **Collectively, define an upper limit on performance (speed-of-light)**
- Loop Arithmetic Intensity (for each level of memory)
  - **Total FLOPs / Total Data Movement** (for that level of memory)
  - Measure of a loop's temporal locality
  - Includes all cache effects
- Plotting loops on the (Hierarchical) Roofline
  - **Each loop has one dot per level of memory**
  - x-coordinate = arithmetic intensity at that level
  - y-coordinate = performance (e.g. GFLOP/s)
  - Proximity to associated ceiling is indicative of a performance bound
  - Proximity of dots to each other is indicative of **streaming** behavior (low cache hit rate)



# What is Roofline used for?

- Understand performance differences between Architectures, Programming Models, implementations, etc...
  - Why do some Architectures/Implementations move more data than others?
  - Why do some compilers outperform others?
- Predict performance on future machines / architectures
  - Set realistic performance expectations
  - Drive for HW/SW Co-Design
- Identify performance bottlenecks & motivate software optimizations
- Determine when we're done optimizing code
  - Assess performance relative to machine capabilities
  - Track progress towards optimality
  - Motivate need for algorithmic changes

Hands-on example on Aurora

# Vendor tools for Roofline analysis

- Intel
  - Intel Advisor
    - <https://www.intel.com/content/www/us/en/developer/tools/oneapi/advisor.html>
- NVIDIA
  - NVIDIA Nsight Compute
    - <https://developer.nvidia.com/nsight-compute>
    - <https://docs.nvidia.com/nsight-compute/NsightCompute/index.html#details-page>
- AMD
  - AMD ROCm Compute Profiler (Omniperf, previously)
    - <https://rocm.docs.amd.com/projects/rocprowler-compute/en/latest/what-is-rocprow-compute.html>
- We won't try every tool. They have different instructions for the same concept. (I know it is annoying, but that is what we have. ☹️)

# Intel Advisor for roofline analysis

# Getting Roofline data in Intel® Advisor: **two-pass approach**

Roofline :  Axis X: $AI = \#FLOP / \#Bytes$  Axis Y: $FLOP/S = \#FLOP \text{ (mask aware)} / \#Seconds$	Overhead
<p>Step 1: Survey (-collect survey)</p> <ul style="list-style-type: none"><li>- Provide <b>#Seconds</b></li><li>- <b>Root access not needed</b></li><li>- User mode sampling, non-intrusive.</li></ul>	<b>1x</b>
<p>Step 2: FLOPS (-collect tripcounts -flops)</p> <ul style="list-style-type: none"><li>- Provide <b>#FLOP</b>, <b>#Bytes</b>, AVX-512 Mask</li><li>- <b>Root access not needed</b></li><li>- Precise, instrumentation based, count number of instructions</li></ul>	<b>3-5x</b>

# Original, Cache-Aware (**CARM**) and **Memory-Level** Roofline

## CARM (cache-aware roofline)

- Single **AI** based on aggregated traffic:  
**CPU core (GPU EUs) <-> memory sub-system**
- Ceilings for compute, cache/memory levels
- AI independent of problem size

Unique features: algorithmic focus and simplicity

## Original Roofline

- AI based on external memory :  
**DDR (GPU GTI)**
- Ceilings for DDR and compute
- AI dependent of problem size



Unique features: DDR bound focus and simplicity

## Memory Level Roofline - **MLR** (see also “Hierarchical Roofline” by LBL)

- **AI** for all memory sub-system levels, combines (1), CARM, (2)Original and (3) Lx-only perspectives
- Harder to interpret for multiple kernels at a time

Unique features: unambiguous bottleneck detection

# How to interpret MLR on CPU ?



Find the minimum of all memory subsystems

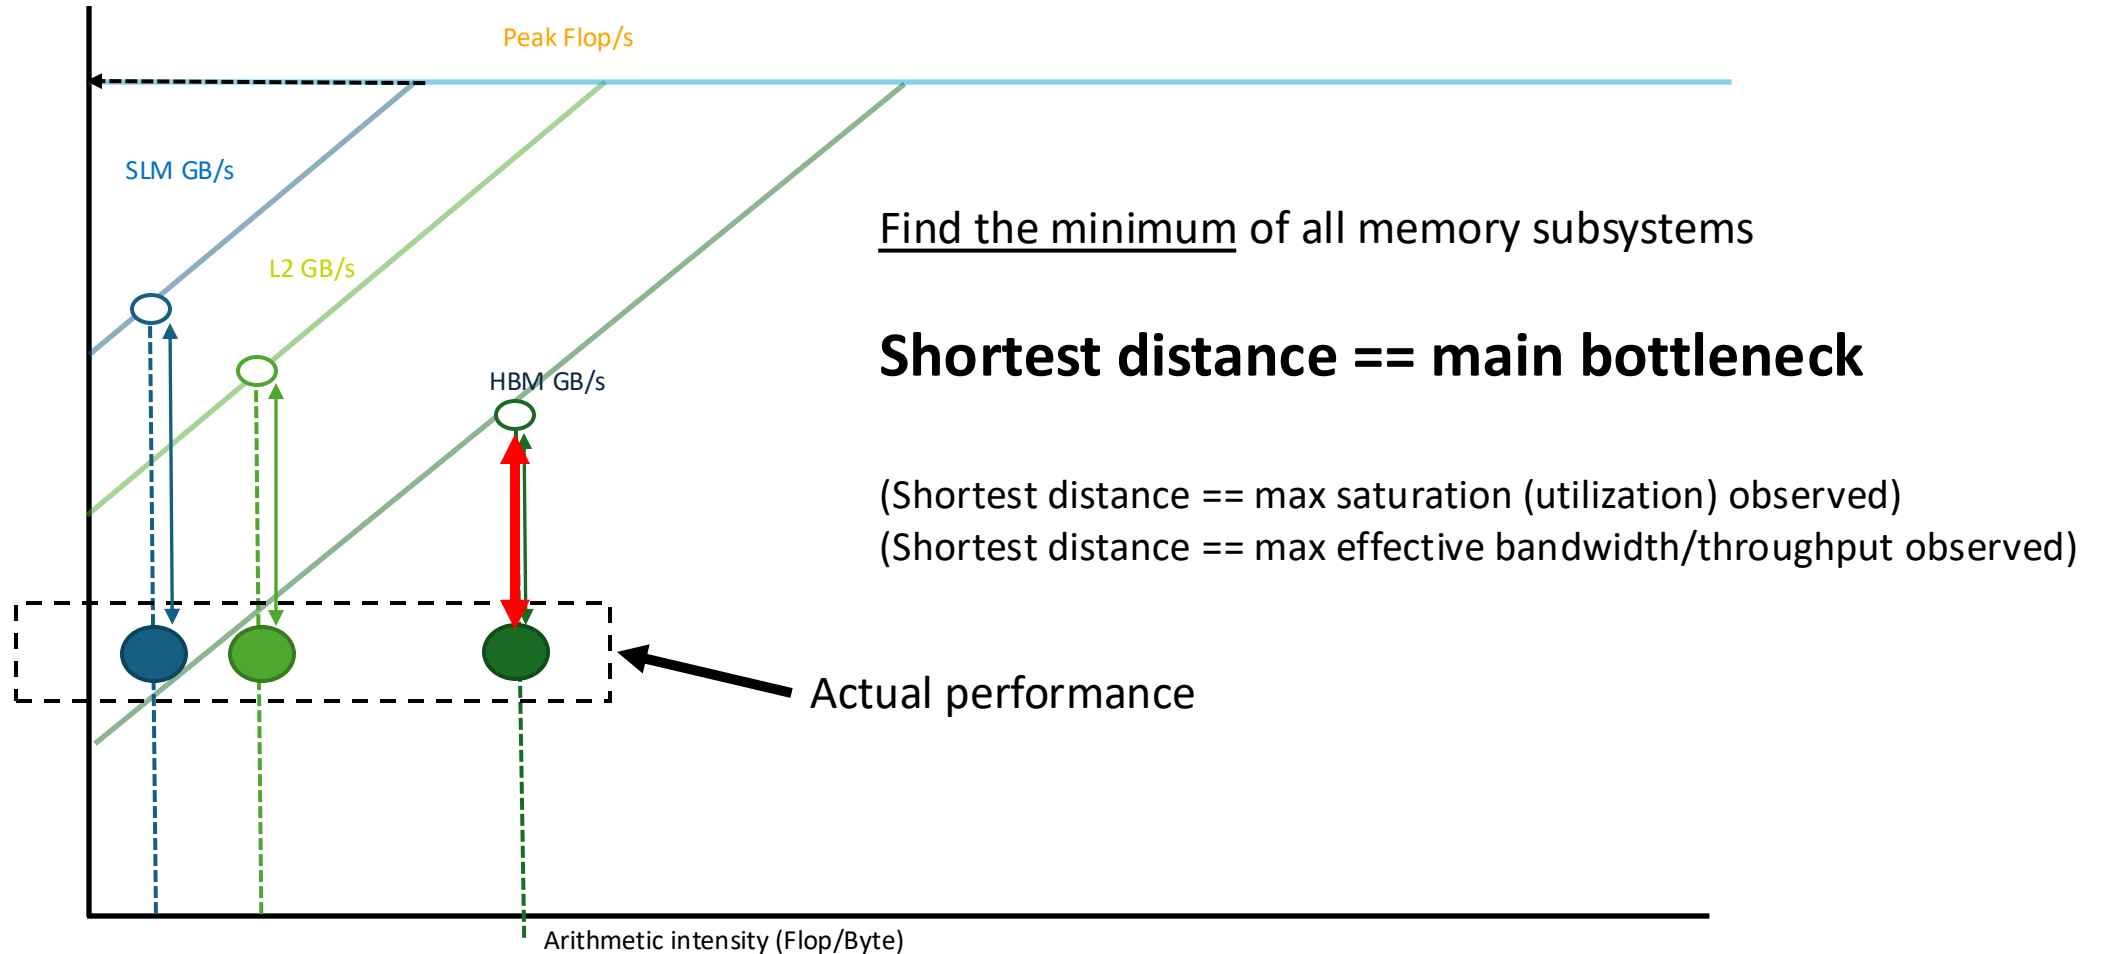
**Shortest distance == main bottleneck**

(Shortest distance == max saturation (utilization) observed)

(Shortest distance == max effective bandwidth/throughput observed)

Actual performance

# How to interpret MLR on GPU ?





# How to generate (profile\*) Roofline for your application

# How to generate **CARM CPU** Roofline profile?

```
As simple as: $ advisor -collect roofline -- <your-executable-with-parameters>
```

## More details / How-To

```
$ source advisor-vars.sh
```

### 1st method. Not compatible with MPI applications:

```
$ advisor -collect roofline --project-dir  
./your_project -- <your-executable-with-  
parameters>
```

*(optional) copy data to your UI desktop system*

```
$ advisor-gui ./your_project
```

```
$ advisor -report roofline --project-dir ./your_project > roofline.html
```

### 2nd method (compatible with MPI, more flexible):

```
$ advisor -collect survey --project-dir ./your_project --  
<your-executable-with-parameters>
```

```
$ advisor -collect tripcounts --flop --project-dir  
./your_project -- <your-executable-with-parameters>
```

# How to generate **MLR+CARM CPU** Roofline profile?

As simple as: `$ advisor -collect roofline -enable-cache-simulation -- <your-executable-with-parameters>`

## More details / How-To

```
$ source advisor-vars.sh
```

1st method. Not compatible with MPI applications :

```
$ advisor -collect roofline -enable-cache-simulation --project-dir ./your_project --<your-executable-with-parameters>
```

*(optional) copy data to your UI desktop system*

```
$ advisor-gui ./your_project
```

```
$ advisor -report roofline --project-dir ./your_project > roofline.html
```

2nd method (compatible with MPI, more flexible):

```
$ advisor -collect survey --project-dir ./your_project --<your-executable-with-parameters>
```

```
$ advisor -collect tripcounts -flop -enable-cache-simulation --project-dir ./your_project -- <your-executable-with-parameters>
```

# How to generate **GPU (MLR & CARM)** Roofline profile?

```
As simple as: $ advisor -collect roofline --profile-gpu --<your-executable-with-parameters>
```

## More details / How-To

```
$ source advisor-vars.sh
```

**1st method. Not compatible with MPI applications :**

```
$ advisor -collect roofline --profile-gpu -  
-project-dir ./your_project -- <your-  
executable-with-parameters>
```

*(optional) copy data to your UI desktop system*

```
$ advisor-gui ./your_project
```

```
$ advisor -report roofline --gpu --project-dir ./your_project > roofline.html
```

**2nd method (compatible with MPI, more flexible):**


```
$ advisor -collect survey --profile-gpu --project-dir  
./your_project -- <your-executable-with-parameters>
```

```
$ advisor -collect tripcounts -flop --profile-gpu --project-  
dir ./your_project -- <your-executable-with-parameters>
```

# GPU Roofline: *Extended* HTML GUI

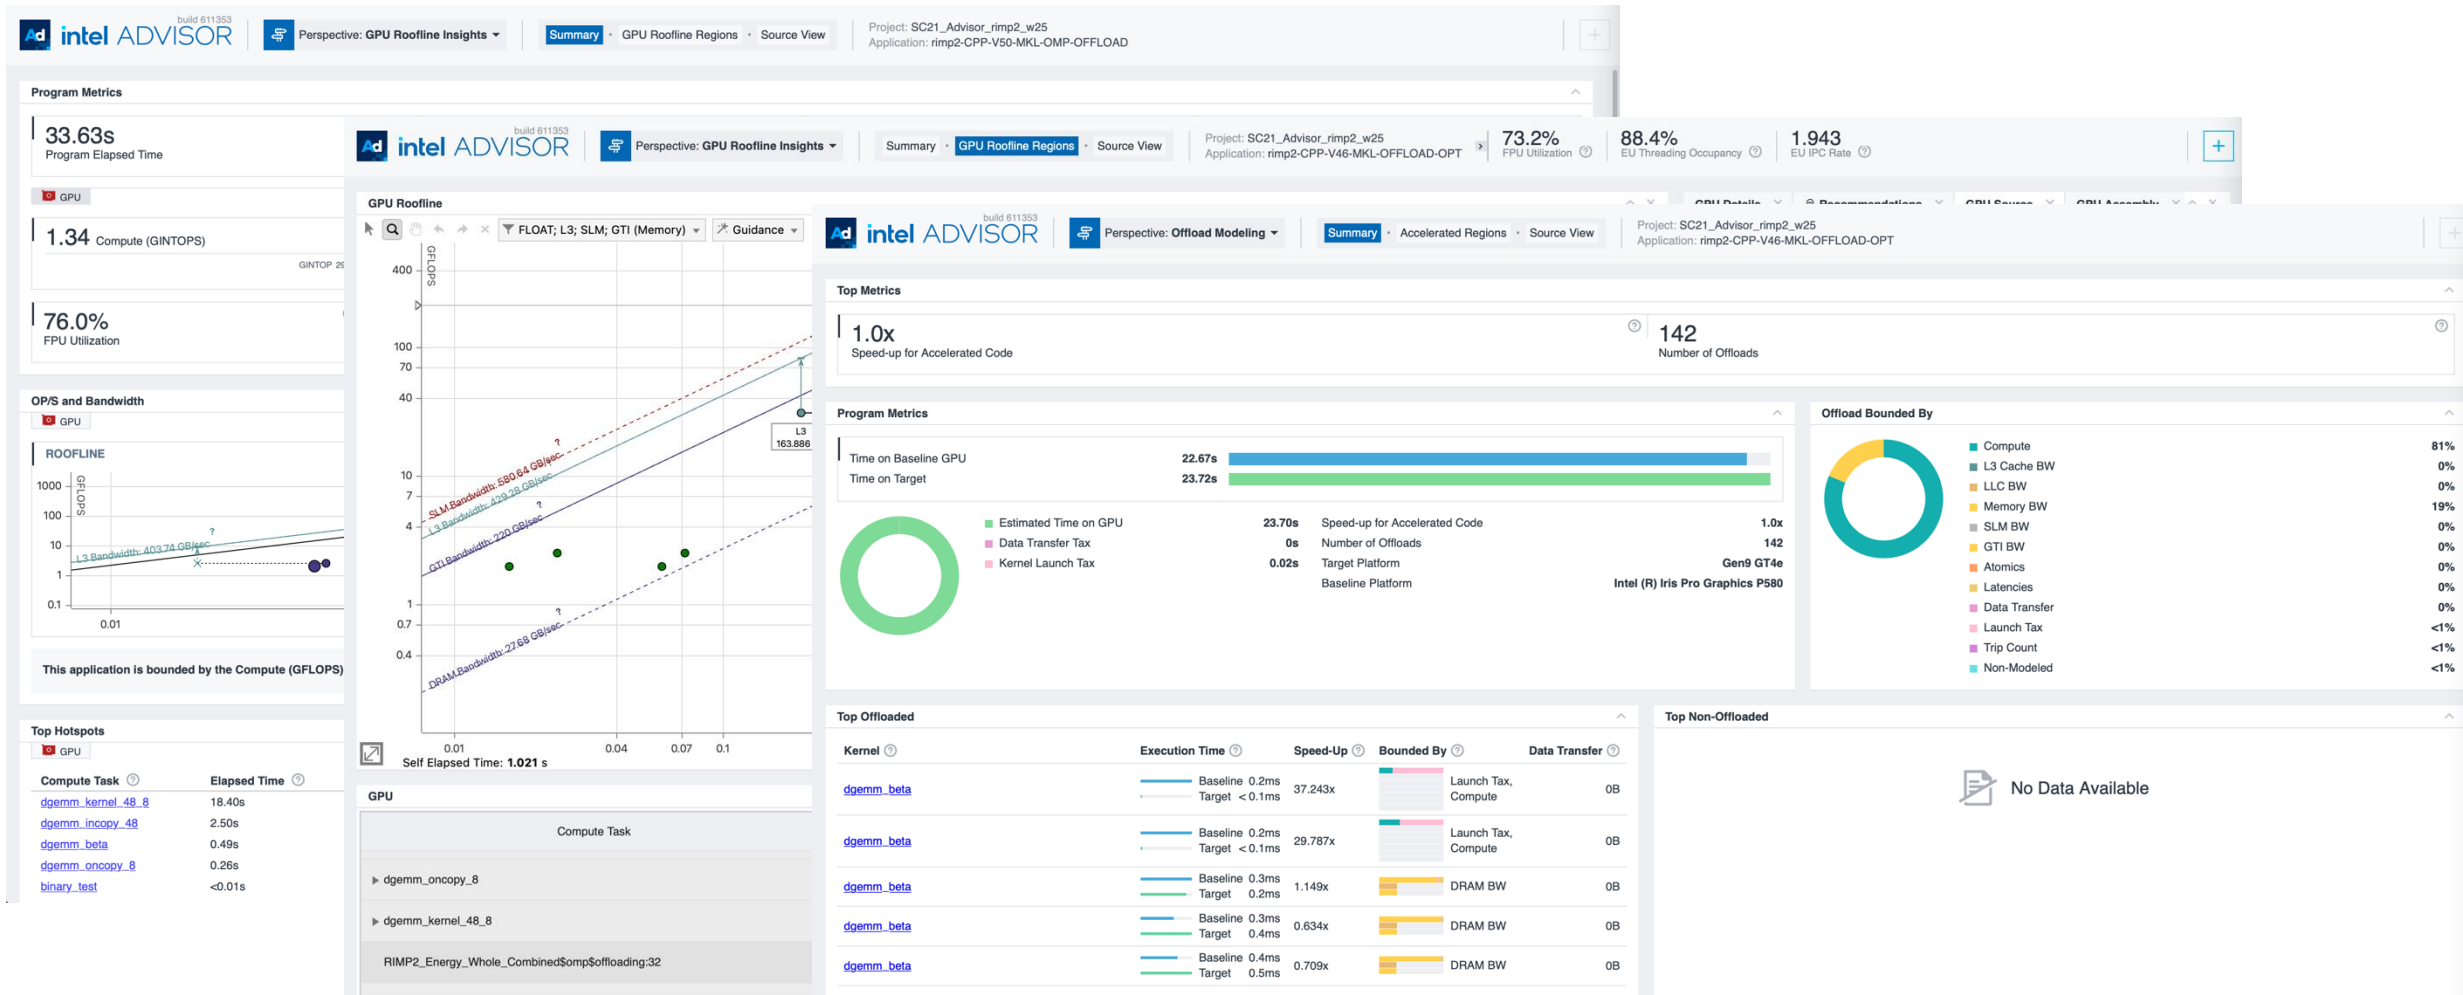
See HTML report in `project-dir/e000|rank.*/report` folder by default

```
source advisor_install_dir/advisor-vars.sh  
  
advisor  
  --report all  
  --project-dir ./your_project  
  --report-output ./roofline.html
```



# Extended HTML GUI

For any system with web browsers



# Roofline on Multi-GPU systems

Add `--target-gpu` option in command line

```
advisor
--collect roofline
--profile-gpu
--project-dir ./your_project
--target-gpu 0:77:0.0
-- <your-executable-with-
parameters>
```

```
Compiled using the dx (Linux OS)/qnx (Windows
OS) option with an Intel compiler.
Tip: Disabling can minimize overhead.
--target-gpu=0:0:2.0 | 0:3:0.0 (0:3:0.0)
The target GPU adapter that will be used to
collect GPU profiling data.
--target-pid=<unsigned integer>
Attach collection to a running process specified
```

# ISO3DFD Code

: A 16th order Finite-Difference Stencil for the 3D Isotropic Wave Equation



# ISO3dfd code

- A Finite Difference stencil kernel for solving the 3D acoustic isotropic wave equation
  - A proxy for propagating a seismic wave
  - 16<sup>th</sup> order in space, with symmetric coefficients
  - 2<sup>nd</sup> order in time scheme without boundary conditions.

- Problem Statement


- Partial Differential Equation (PDE) for wave propagation

$$\frac{d^2 p}{dt^2} = v^2 \left( \frac{d^2 p}{dx^2} + \frac{d^2 p}{dy^2} + \frac{d^2 p}{dz^2} \right)$$

, where  $p$  is pressure,  $v$  is velocity, and  $t$  is time.

- For a 16<sup>th</sup> order finite difference stencil, we use the adjacent 8 values in each direction of the  $x$ ,  $y$ , and  $z$  axis
- The expanded equation looks like this, where the array  $C$  holds the coefficients for the changes in  $x$ ,  $y$  and  $z$ .
 
$$p_{i,j,k}^{n+1} = C[0]p_{i,j,k}^n + C[1](p_{i+1,j,k}^n + p_{i-1,j,k}^n + p_{i,j+1,k}^n + p_{i,j-1,k}^n + p_{i,j,k+1}^n + p_{i,j,k-1}^n) + \dots$$


$$+ C[8](p_{i+8,j,k}^n + p_{i-8,j,k}^n + p_{i,j+8,k}^n + p_{i,j-8,k}^n + p_{i,j,k+8}^n + p_{i,j,k-8}^n)$$



<https://www.intel.com/content/www/us/en/developer/articles/technical/iso3dfd-code-walkthrough.html>

# Code walk through

- Variables
  - $n_1$   $n_2$   $n_3$  : Grid dimensions for the stencil
  - Iterations : No. of timesteps.
  - $n_1$ ,  $n_2$ , and  $n_3$  has the addition of  $2 * kHalfLength$  to represent the entire block including the halo region.



- Dark blue: grid
- Gray: halo of the grid
- Pink: points needed for calculation

```
...
try {
    // Parse command line arguments and increase them by HALO
    n1 = std::stoi(argv[1]) + (2 * kHalfLength);
    n2 = std::stoi(argv[2]) + (2 * kHalfLength);
    n3 = std::stoi(argv[3]) + (2 * kHalfLength);
    num_iterations = std::stoi(argv[4]);
} catch (...) {
    Usage(argv[0]);
    return 1;
}

...

// Compute the total size of grid
size_t nsize = n1 * n2 * n3;

...

// Apply the DX, DY and DZ to coefficients
coeff[0] = (3.0f * coeff[0]) / (dxyz * dxyz);
for (auto i = 1; i <= kHalfLength; i++) {
    coeff[i] = coeff[i] / (dxyz * dxyz);
}
```

# ISO3dFD code versions

- In this hands-on, we use four ISO3DFD variants.
- These variants add progressive/incremental levels of optimization as follow:
  - **1\_CPU\_only.cpp**: an initial CPU version
  - **2\_GPU\_basic.cpp**: basic GPU offloading using SYCL
  - **3\_GPU\_linear.cpp**: reduced index calculation
  - **4\_GPU\_private\_memory\_I.cpp**: addition of private array for coefficients

[https://github.com/oneapi-src/oneAPI-samples/tree/master/DirectProgramming/C%2B%2BSYCL/StructuredGrids/guided\\_iso3dfd\\_GPUOptimization](https://github.com/oneapi-src/oneAPI-samples/tree/master/DirectProgramming/C%2B%2BSYCL/StructuredGrids/guided_iso3dfd_GPUOptimization)

# Build the code

```
$ qsub -l -l select=1 -l walltime=02:00:00 -l filesystems=home:flare -A ATPESC2025 -q ATPESC
```

```
$ cd /flare/ATPESC2025/usr/$USER
```

```
$ git clone https://github.com/oneapi-src/oneAPI-samples.git
```

```
$ cd oneAPI-samples/DirectProgramming/C++SYCL/StructuredGrids/guided_iso3dfd_GPUOptimization/  
or
```

```
$ cp -r /flare/ATPESC2025/EXAMPLES/track6-tools/roofline/oneAPI-  
samples/DirectProgramming/C++SYCL/StructuredGrids/guided_iso3dfd_GPUOptimization/ .
```

```
$ cd guided_iso3dfd_GPUOptimization
```

```
$ mkdir build
```

```
$ cd build
```

```
$ module load cmake
```

```
$ cmake ..
```

```
$ make
```

# 1\_CPU\_only.cpp

```
...
for (auto iz = kHalfLength; iz < n3_end; iz++) {
    for (auto iy = kHalfLength; iy < n2_end; iy++) {
        // Calculate start pointers for the row over X dimension
        float* ptr_next = ptr_next_base + iz * dimn1n2 + iy * n1;
        float* ptr_prev = ptr_prev_base + iz * dimn1n2 + iy * n1;
        float* ptr_vel = ptr_vel_base + iz * dimn1n2 + iy * n1;

        // Iterate over X
        for (auto ix = kHalfLength; ix < n1_end; ix++) {
            // Calculate values for each cell
            float value = ptr_prev[ix] * coeff[0];
            for (int i = 1; i <= kHalfLength; i++) {
                value +=
                    coeff[i] *
                    (ptr_prev[ix + i] + ptr_prev[ix - i] +
                     ptr_prev[ix + i * n1] + ptr_prev[ix - i * n1] +
                     ptr_prev[ix + i * dimn1n2] + ptr_prev[ix - i * dimn1n2]);
            }
            ptr_next[ix] = 2.0f * ptr_prev[ix] - ptr_next[ix] + value * ptr_vel[ix];
        }
    }
}
...
```

**Computing values for each cell**

**Iterate over x**

**Iterate over y and z**

# Run 1\_CPU\_only

**\$ src/1\_CPU\_only 512 512 512 10**

Running on CPU serial version

```
-----  
time           : 11.309 secs  
throughput     : 118.682 Mpts/s  
flops          : 7.23962 GFlops  
bytes          : 1.42419 GBytes/s  
-----
```

## 2\_GPU\_basic: offloading the CPU code to GPU using SYCL

```
// Send a SYCL kernel(lambda) to the device for parallel execution
// Each kernel runs single cell
h.parallel_for(kernel_range, [=](id<3> idx) {
    // Start of device code
    // Add offsets to indices to exclude HALO
    int i = idx[0] + kHalfLength;
    int j = idx[1] + kHalfLength;
    int k = idx[2] + kHalfLength;
```

```
    // Calculate values for each cell
    float value = prev_acc[i][j][k] * coeff_acc[0];
#pragma unroll(8)
    for (int x = 1; x <= kHalfLength; x++) {
        value +=
            coeff_acc[x] * (prev_acc[i][j][k + x] + prev_acc[i][j][k - x] +
                            prev_acc[i][j + x][k] + prev_acc[i][j - x][k] +
                            prev_acc[i + x][j][k] + prev_acc[i - x][j][k]);
    }
    next_acc[i][j][k] = 2.0f * prev_acc[i][j][k] - next_acc[i][j][k] +
        value * vel_acc[i][j][k];
    // End of device code
});
```

**Computing values for each cell**

**SYCL kernel to the device  
for parallel execution  
over x, y, and z**

# Run 2\_GPU\_basic

```
$ export ZE_AFFINITY_MASK=0.0
```

```
$ src/2_GPU_basic 512 512 512 100
```

```
Running GPU basic offload version
```

```
Running on Intel(R) Data Center GPU Max 1550
```

```
The Device Max Work Group Size is : 1024
```

```
The Device Max EUCount is : 448
```

```
-----  
time           : 7.25 secs  
throughput     : 1851.28 Mpts/s  
flops          : 112.928 GFlops  
bytes          : 22.2153 GBytes/s  
-----
```

**Compared 113.09 sec on CPU for 100 steps**  
**- 15.6X speed-up on GPU from on a CPU core**

```
$ advisor -collect roofline --profile-gpu --project-dir ADV_02_512 -- ./src/2_GPU_basic 512 512 512 100
```

or

```
$ cp /flare/ATPESC2025/EXAMPLES/track6-tools/roofline/ADV_results/ADV_02_512 .
```

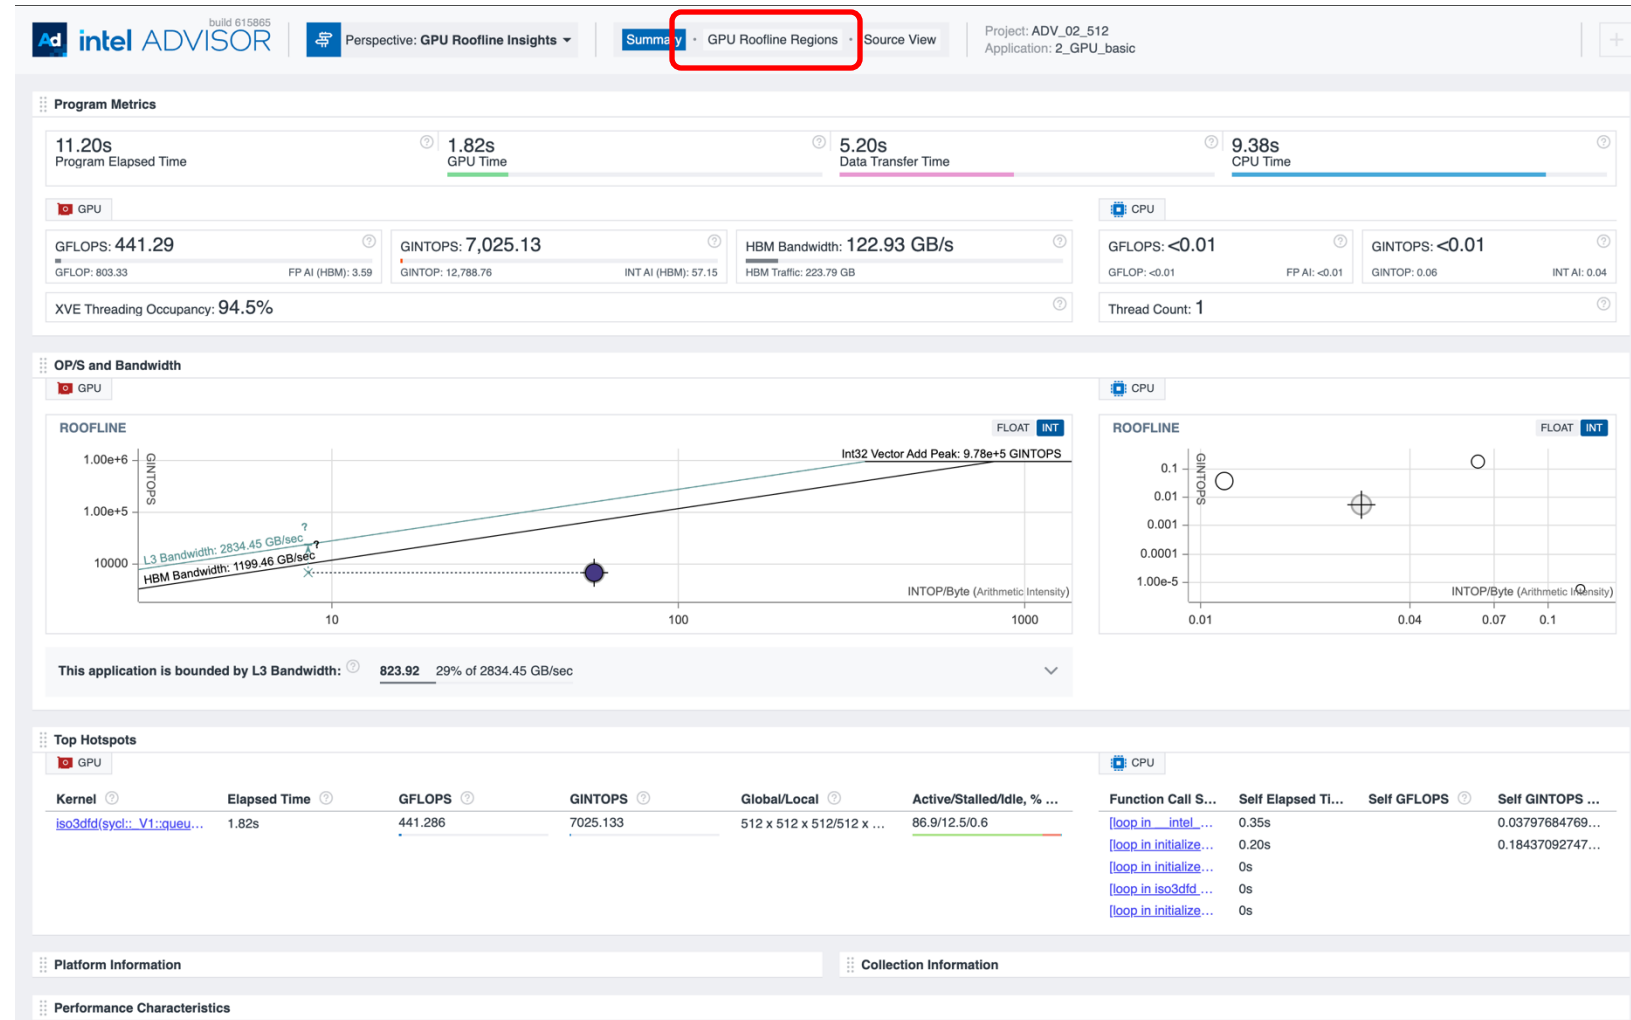
```
# Download advisor-report.html from /ADV_02_512/e000/report
```



# Orient yourself in GPU+CPU Roofline!

- **Quiz:** What are GFLOPS of CPU and GPU?
- **Quiz:** What is a cumulative application bottleneck (Bounded by)?
- **Quiz:** Do you see iso3dfd kernel?

- Now, let's switch to the **main page** ("GPU Roofline Regions")



# GPU MLR Roofline

Note some difference, “HBM” by default

1. Enable **Memory Metrics** and **Point info**
2. Look into **GPU Details** tab and find **INDIVIDUAL ROOFLINE** chart with small Guidance, Hints and BoundBy
  - **Quiz:** what is a main bottleneck (“Bound By”) for the iso3dfd kernel?
  - **Quiz:** what are FP AI and INT AI?
  - **Quiz:** what are Instruction Mix Details?
3. Go back to **Main** Roofline Chart and **double-click** on the circle to get the same guidance on the large chart

Kernel (loop) locality is proportional to the width of the “bound by” line (ratio of DRAM to Lx bytes)




# 2\_GPU\_basic

## SUMMARY

Elapsed Time	GINTOPS	7,025.133
1.820s	GFLOPS	441.286
Global	Local	
512 x 512 x 512	512 x 2 x 1	

## ROOFLINE GUIDANCE

 This kernel is bounded by L3 Bandwidth.




## OP/S AND BANDWIDTH

▶ Compute (GINTOPS):	7025.13	0% of 978000.00 GIN
▶ Compute (GFLOPS):	441.29	1% of 22800.00 GFLC
▶ L3 Bandwidth:	823.92	29% of 2834.45 GB/se
▶ CARM Bandwidth:	1513.01	6% of 23000.00 GB/se
▶ SLM Bandwidth:	0	0% of 23000.00 GB/se
▶ HBM Bandwidth:	122.93	10% of 1199.46 GB/se

## MEMORY METRICS

Impacts	
L3	63%
CARM	14%
SLM	0%
HBM	22%
Shares	
L3	1499.886GB
CARM	2754.324GB
SLM	0GB
HBM	223.786GB


## INSTRUCTION MIX



## INSTRUCTION MIX DETAILS

▶ Compute	584.69	70%
▶ Memory	18.04	2%
▶ Other	233.20	28%

## PERFORMANCE CHARACTERISTICS



# 3\_GPU\_linear: using linearized index to reduce index calculation

```
// Send a SYCL kernel(lambda) to the device for parallel execution
// Each kernel runs single cell
h.parallel_for(kernel_range, [=](id<3> nidx) {
    // Start of device code
    // Add offsets to indices to exclude HALO
    int n2n3 = n2 * n3;
    int i = nidx[0] + kHalfLength;
    int j = nidx[1] + kHalfLength;
    int k = nidx[2] + kHalfLength;
```

```
// Calculate linear index for each cell
```

```
int idx = i * n2n3 + j * n3 + k;
```

**Linearized index**

```
// Calculate values for each cell
```

```
float value = prev_acc[idx] * coeff_acc[0];
```

```
#pragma unroll(8)
```

```
for (int x = 1; x <= kHalfLength; x++) {
```

```
    value +=
```

```
        coeff_acc[x] * (prev_acc[idx + x]          + prev_acc[idx - x] +
                        prev_acc[idx + x * n3]      + prev_acc[idx - x * n3] +
                        prev_acc[idx + x * n2n3] + prev_acc[idx - x * n2n3]);
```

```
}
```

```
next_acc[idx] = 2.0f * prev_acc[idx] - next_acc[idx] +
                value * vel_acc[idx];
```

```
// End of device code
```

```
});
```

**Computing values for each cell**

**SYCL kernel to the device  
for parallel execution  
over x, y, and z**

# Run 3\_GPU\_linear

```
$ export ZE_AFFINITY_MASK=0.0
```

```
$ src/3_GPU_linear 512 512 512 100
```

```
Running linear indexed GPU version
```

```
Running on Intel(R) Data Center GPU Max 1550
```

```
The Device Max Work Group Size is : 1024
```

```
The Device Max EUCount is : 448
```

```
-----  
time           : 0.866 secs  
throughput     : 15498.6 Mpts/s  
flops          : 945.414 GFlops  
bytes          : 185.983 GBytes/s  
-----
```

Compared 7.25 sec from 2\_GPU\_basic  
- 8.4X speed-up

```
$ advisor -collect roofline --profile-gpu --project-dir ADV_03_512 -- ./src/3_GPU_linear 512 512 512 100
```

or

```
$ cp /flare/ATPESC2025/EXAMPLES/track6-tools/roofline/ADV_results/ADV_03_512 .
```

```
# Download advisor-report.html from /ADV_03_512/e000/report
```

# Check 3\_GPU\_linear

- Please open the html for 3\_GPU\_linear
  - **Quiz:** GFLOPS? Did it change?  
(hint: look at Summary , and then go back to GPU Roofline Regions)
  - **Quiz:** what is a main bottleneck for the kernel?
  - **Quiz:** Any changes in INTOP?

# 3\_GPU\_linear

## SUMMARY

Elapsed Time

**0.683s**

GINTOPS

**2,516.945**

GFLOPS

**1,137.126**

Global

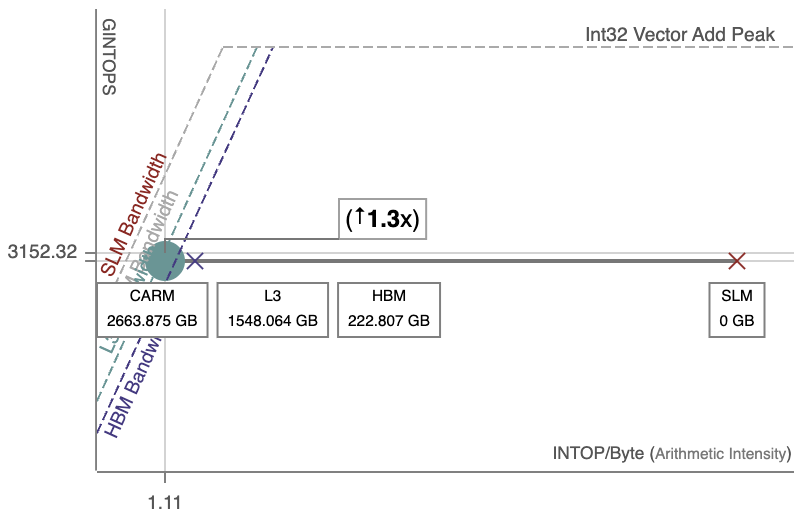
**512 x 512 x 512**

Local

**512 x 2 x 1**

## ROOFLINE GUIDANCE

 This kernel is bounded by **L3 Bandwidth**.



## OP/S AND BANDWIDTH

▶ Compute (GINTOPS):	2516.95	0% of 965000.00 GIN
▶ Compute (GFLOPS):	1137.13	4% of 22800.00 GFLC
▶ L3 Bandwidth:	2266.14	79% of 2838.19 GB/s
▶ CARM Bandwidth:	3899.52	16% of 23000.00 GB/s
▶ SLM Bandwidth:	0	0% of 23000.00 GB/s
▶ HBM Bandwidth:	326.16	27% of 1181.52 GB/s

## MEMORY METRICS


Impacts

L3	64%
CARM	14%
SLM	0%
HBM	22%

Shares

L3	1548.064GB
CARM	2663.875GB
SLM	0GB
HBM	222.807GB

## INSTRUCTION MIX



## INSTRUCTION MIX DETAILS

▶ Compute	119.54	55%
▶ Memory	18.04	8%
▶ Other	80.11	37%

## PERFORMANCE CHARACTERISTICS





Active: 64.0%  
Stalled: 34.4%  
Idle: 1.5%

XVE Threading Occupancy: 84.7%

SIMD Width: 32

# Comparison from 2\_GPU\_basic to 3\_GPU\_linear in FLOAT roofline chart



**ANY OPTIMIZATION OPPORTUNITIES FOR 3\_GPU\_LINEAR?**

Possible to increase AI by re-using data multiple times?



## 4\_GPU\_private\_memory\_I: adding private array for coefficients

```
// Send a SYCL kernel(lambda) to the device for parallel execution
// Each kernel runs single row over first dimension
h.parallel_for(kernel_range, [=](id<2> nidx) {
    // Start of device code
    // Add offsets to indices to exclude HALO
    // Start and end index used in loop
    int n2n3 = n2 * n3;
    int i = kHalfLength;
    int j = nidx[0] + kHalfLength;
    int k = nidx[1] + kHalfLength;
    int end_i = n1 - kHalfLength;

    // Calculate global linear index for each cell
    int idx = i * n2n3 + j * n3 + k;

    // Create arrays to store data used multiple times
    // Local copy of coeff buffer/continuous values over 1st dim which
    // are used to calculate stencil front and back arrays are used to
    // ensure the values over 1st dimension are read once, shifted in`
    // these array and re-used multiple times before being discarded
    // This is an optimization technique to enable data-reuse and
    // improve overall FLOPS to BYTES read ratio
    float coeff[kHalfLength + 1];
    float front[kHalfLength + 1];
    float back[kHalfLength];

    // Fill local arrays, front[0] contains current cell value
    for (int x = 0; x <= kHalfLength; x++) {
        coeff[x] = coeff_acc[x];
        front[x] = prev_acc[idx + n2n3 * x];
    }
    for (int x = 1; x <= kHalfLength; x++) {
        back[x-1] = prev_acc[idx - n2n3 * x];
    }
}
```

**Local array and  
reuse multiple times**

```
// Iterate over first dimension excluding HALO
for (; i < end_i; i++) {
    // Calculate values for each cell
    float value = front[0] * coeff[0];
    #pragma unroll(kHalfLength)
    for (int x = 1; x <= kHalfLength; x++) {
        value += coeff[x] *
            (prev_acc[idx + x] + prev_acc[idx - x] +
             prev_acc[idx + x * n3] + prev_acc[idx - x * n3] +
             front[x] + back[x - 1] );
    }
    next_acc[idx] = 2.0f * front[0] - next_acc[idx] +
        value * vel_acc[idx];

    // Increase linear index, jump to the next cell in first dimension
    idx += n2n3;

    // Shift values in front and back arrays
    for (auto x = kHalfLength - 1; x > 0; x--) {
        back[x] = back[x - 1];
    }
    back[0] = front[0];

    for (auto x = 0; x < kHalfLength; x++) {
        front[x] = front[x + 1];
    }
    front[kHalfLength] = prev_acc[idx + kHalfLength * n2n3];
}

// End of device code
});
```

**Computing values  
for each cell**

**SYCL kernel to the device  
for parallel execution  
over y and z**

# Run 4\_GPU\_private\_memory\_I

```
$ export ZE_AFFINITY_MASK=0.0
```

```
$ src/4_GPU_private_memory_I 512 512 512 100
```

```
Running GPU private memory version with iterations over first dimension
```

```
Running on Intel(R) Data Center GPU Max 1550
```

```
The Device Max Work Group Size is : 1024
```

```
The Device Max EUCount is : 448
```

```
-----  
time           : 0.637 secs  
throughput     : 21070.3 Mpts/s  
flops          : 1285.29 GFlops  
bytes          : 252.843 GBytes/s  
-----
```

Compared 0.866 sec from 3\_GPU\_linear

- 1.35X speed-up

Compared 7.25 sec from 2\_GPU\_basic

- 11.4X speed-up

```
$ advisor -collect roofline --profile-gpu --project-dir ADV_04_512 -- ./src/4_GPU_private_memory_I 512 512 512 100
```

or

```
$ cp /flare/ATPESC2025/EXAMPLES/track6-tools/roofline/ADV_results/ADV_04_512 .
```

```
# Download advisor-report.html from /ADV_04_512/e000/report
```

# Check 4\_gpu\_private\_memory\_l


- Please open the html for 4\_GPU\_Private\_memory\_l
  - **Quiz:** GFLOPS? Did it change?  
(hint: look at Summary , and then go back to GPU Roofline Regions)
  - **Quiz:** what is a main bottleneck for the kernel?

SUMMARY

Elapsed Time	GINTOPS	920.446
0.443s	GFLOPS	1,816.197
Global	Local	
512 x 512	512 x 2	

ROOFLINE GUIDANCE

 This kernel is bounded by L3 Bandwidth.




OP/S AND BANDWIDTH

▶ Compute (GINTOPS):	920.45	0% of 500000.00 GIN
▶ Compute (GFLOPS):	1816.20	7% of 22800.00 GFLC
▶ L3 Bandwidth:	1563.77	55% of 2835.72 GB/s
▶ SLM Bandwidth:	0	0% of 23000.00 GB/s
▶ CARM Bandwidth:	4290.77	18% of 23000.00 GB/s
▶ HBM Bandwidth:	525.13	44% of 1180.92 GB/s

MEMORY METRICS

Impacts	
L3	47%
SLM	0%
CARM	16%
HBM	38%
Shares	
L3	692.461GB
SLM	0GB
CARM	1900.022GB
HBM	232.534GB

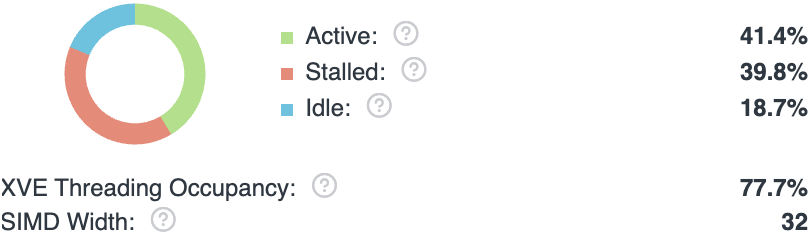
INSTRUCTION MIX





INSTRUCTION MIX DETAILS

▶ Other	27.75	31%
▶ Compute	51.73	58%
▶ Memory	10.08	11%

PERFORMANCE CHARACTERISTICS



# Comparison from 3\_GPU\_linear to 4\_GPU\_private\_memory\_I



AI for L3 increases by re-using data; as a result, GFLOPS improves from 1,137 to 1,816: 1.6X speed-up

# Recap


# Performance results

Grid size: 512x512x512

Number of iteration: 100

Employed Compute Platform

- CPU: Intel Xeon CPU Max
- GPU: Intel Data Center GPU Max



Selected versions	Target platform	GFLOPs	Kernel time (s)	Speed up from 1_CPU_only	Speed-up from 2_GPU_basic
1_CPU_only*	1 core from CPU	7.24	113.09	1 x	-
2_GPU_basic	1 stack from GPU	112.9	7.25	15.7 x	1 x
3_GPU_linear		945.4	0.866	131 x	8.4 x
4_GPU_private_memory_l		1285.3	0.637	179 x	11.4 x

\* 1\_CPU\_only ran 10 iterations instead of 100 iterations since its performance is too low; therefore, the kernel time is projected for 100 iterations.

**Thank you!**





## **ARGONNE TRAINING PROGRAM ON EXTREME-SCALE COMPUTING**

Produced by Argonne National Laboratory, a U.S. Department of Energy Laboratory managed by UChicagoArgonne, LLC under contract DE-AC02-06CH11357.

Special thanks to the National Energy Research Scientific Computing Center (NERSC) and Oak Ridge Leadership Computing Facility (OLCF) for the use of their resources during the training event.

The U.S. Government retains for itself and others acting on its behalf a nonexclusive, royalty-free license in this video, with the rights to reproduce, to prepare derivative works, and to display publicly.