


The CARM Tool

Cache-aware Roofline Model for HPC

José Morgado
Leonel Sousa
Aleksandar Ilic

28th POP Seminar
5 September 2024

- Modern HPC systems and applications are **complex** and **heterogeneous**
 - Hard to model
 - Hard to optimize
- The CARM as a solution
 - Easy to understand
 - Accurate performance overview
 - Good optimization hints
- CARM is only supported by Intel Advisor



Intel Sapphire Rapids CPU

Modern HPC systems and applications

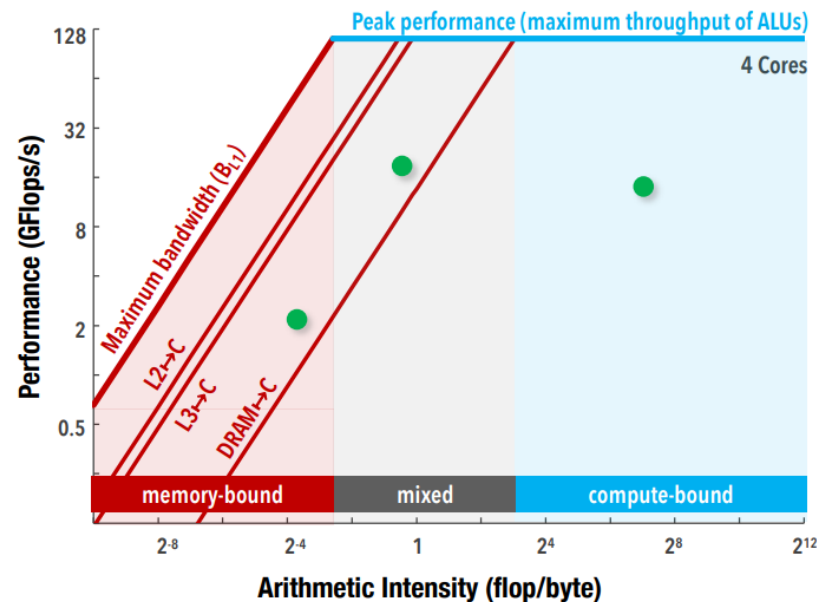
are **complex** and **heterogeneous**

- Hard to model
- Hard to optimize

The CARM as a solution




- Easy to understand
- Accurate performance overview
- Good optimization hints

CARM is only supported by Intel Advisor





The Cache-Aware Roofline Model

Main contributions of the CARM Tool

-  Porting CARM to AMD/ARM/RISC-V
-  Providing application analysis in the scope of CARM
-  Combining all features in a single tool

Porting CARM to other architectures

-  Requires tailored microbenchmarks
-  Understanding of underlying architecture

Experimental results

-  CARM Architecture Analysis
-  CARM Application Analysis



 **State of the Art**

 The CARM Tool

 Results

 Conclusion



State of the art CPUs



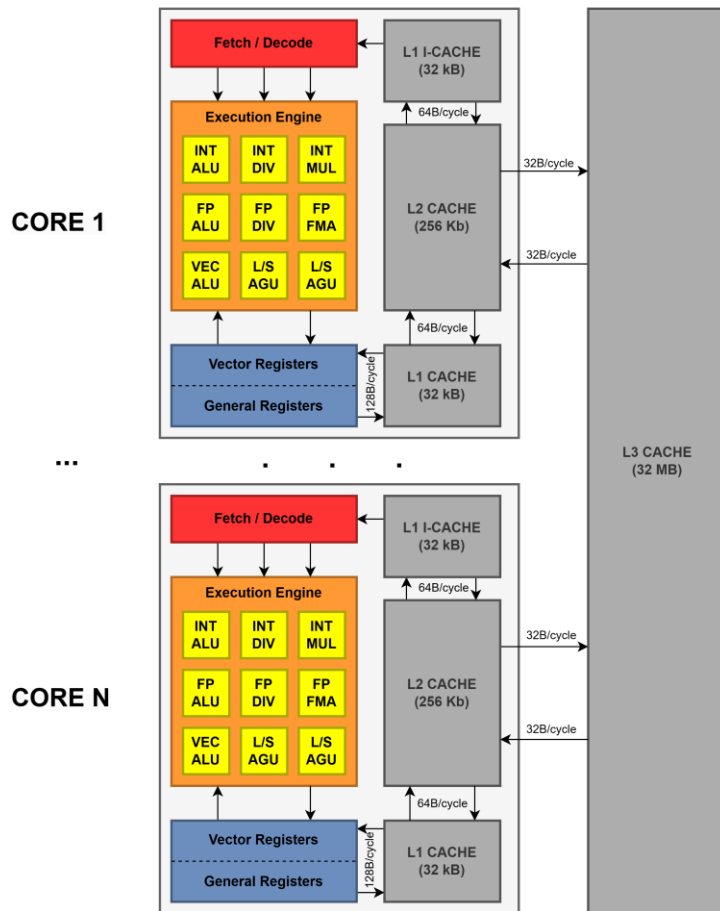
The Cache-Aware Roofline Model



State of the art Roofline Tools



State of the art PMU and DBI Tools



Vendor	Vector ISA Extensions		
Intel	SSE	AVX2	AVX-512
AMD	SSE	AVX2	AVX-512
ARM	Neon	SVE	
RISC-V	RVV		

How does CARM work?

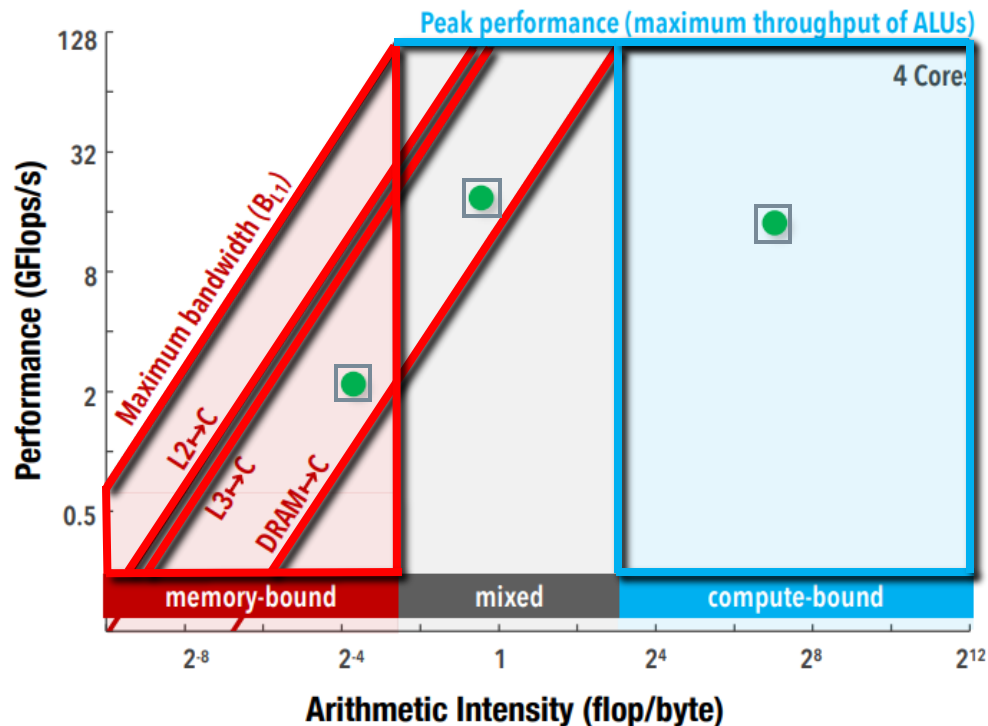
- Sloped Roof
- Flat Roof

How to generate CARM?

- Floating-Point Microbenchmarks
- Memory Microbenchmarks

The CARM Tool

- Automatic Benchmarking
- Automatic CARM Generation



How does CARM work?

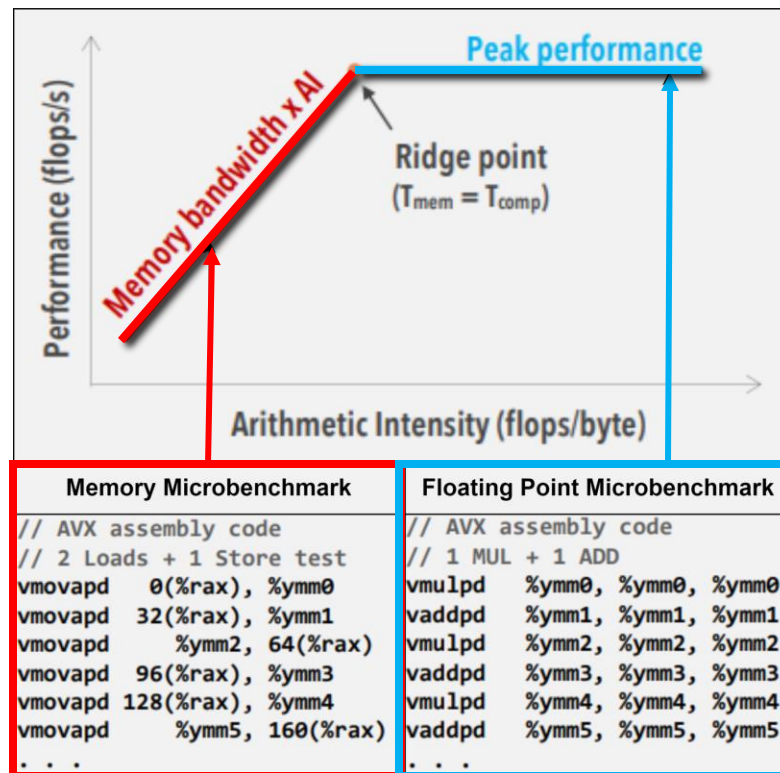
- Sloped Roof
- Flat Roof

How to generate CARM?


- Floating-Point Microbenchmarks
- Memory Microbenchmarks

The CARM Tool

- Automatic Benchmarking
- Automatic CARM Generation



Tools Features	Intel Advisor	AMD uProf	ERT	CARM Tool
Supported Architectures	Intel	AMD	Intel A	Intel AMD ARM RISC-V
Supported Roofline	CARM	ORM	? CARM	CARM
Application Analysis	DBI	PMUs	No sup	DBI PMUs
Open-Source	No	No	Yes	Yes




Tools \ Features	PAPI	Perf
Supports	Intel AMD ARM	Intel AMD ARM RISC-V
Advantage	Portability	Availability
Disadvantage	Only ROI*	Portability

*ROI – Region of Interest



Tools	Intel SDE	DynamoRIO
Features		
Supports	Intel AMD	Intel AMD ARM ? RISC-V ?
Advantage	Detailed Analysis	Customizable Analysis
Disadvantage	Not Customizable	Overhead
Open-Source	No	Yes



 State of the Art

 **The CARM Tool**

 Results

 Conclusion



High-Level Overview



Graphical User Interface



Automatic Benchmarking



Application Analysis



Low-Level Overview



Benchmark Generation



Frequency Measuring



Benchmarking

🔌 Interfacing

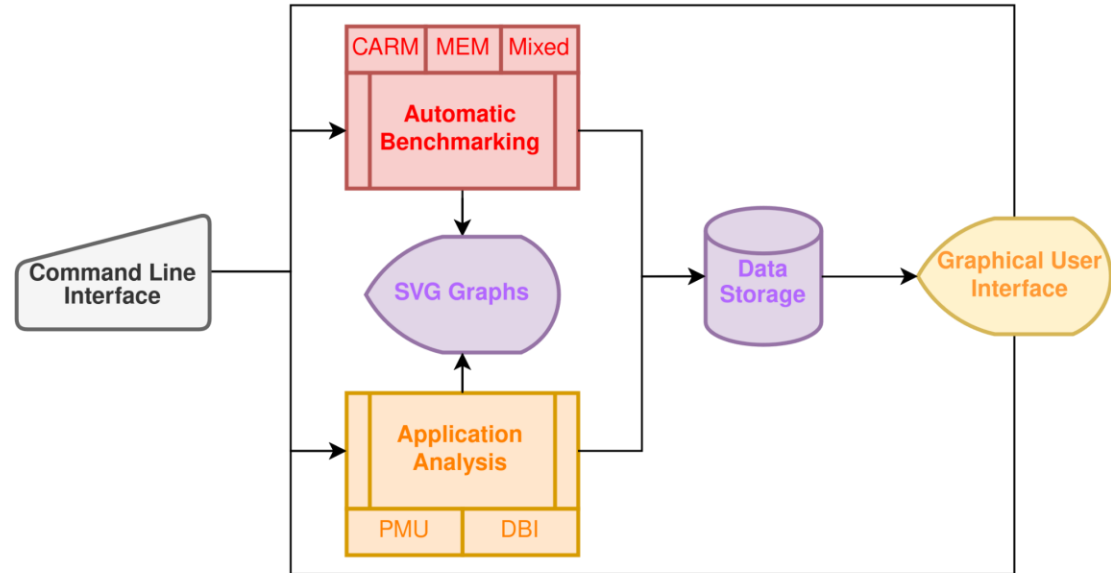
- 🔌 Command Line Interface
- 🔌 Graphical User Interface

🔌 Automatic Benchmarking

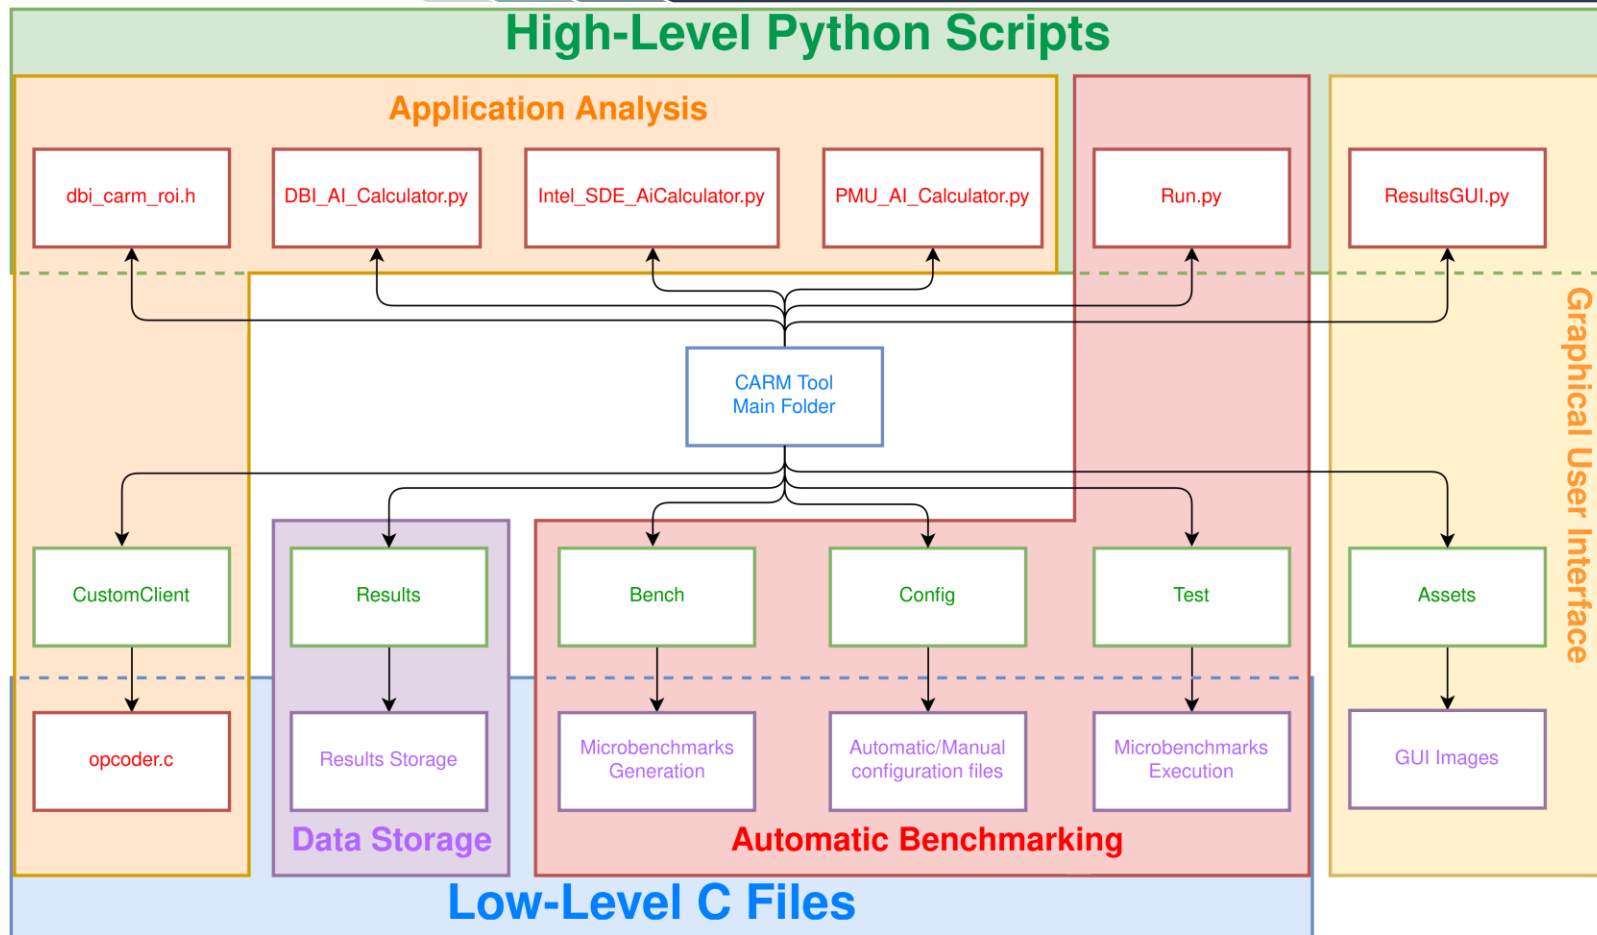
- 🔌 CARM Benchmarks
- 🔌 Memory / FP / Mixed Benchmarks

🔌 Application Analysis

- 🔌 PMU Analysis
- 🔌 DBI Analysis



intel	AMD	arm	RISC-V
Scalar	Scalar	Scalar	Scalar
SSE	SSE	Neon	RVV 0.7.1 1.0
AVX2	AVX2	SVE	
AVX512	AVX512		





Results Visualization



ResultsGUI.py

Assets

GUI Images

Graphical User Interface



Benchmark Execution



The screenshot displays the CARM Tool GUI. On the left, a sidebar titled "CARM Tool Functions" contains a "Run CARM Benchmarks" button and a "CARM Benchmarks Configuration" panel. The configuration panel includes fields for "Machine Name", "Machine Cache Sizes per Core (Kb)" (L1, L2, L3 Total Size), "Thread Counts to Benchmark" (1 2 4 8 16 32 64...), and checkboxes for "Interleave Threads (NUMA)", "ISA Extensions to Benchmark" (AVX512, AVX2, SSE, Scalar), "Precisions to Benchmark" (DP, SP), "Load/Store Ratio Configuration" (Custom Load/Store Ratio, Only Loads, Only Stores), and "DRAM Test Size Configuration" (Custom Size (Kb), Auto_Adjust). Below the configuration panel are buttons for "Run Application Analysis" and "Stop Benchmark/Analysis". The main area shows a "Select a Machine to View CARM Results" prompt and a large "CHAMP" logo with a stylized triangle graphic.

ResultsGUI.py

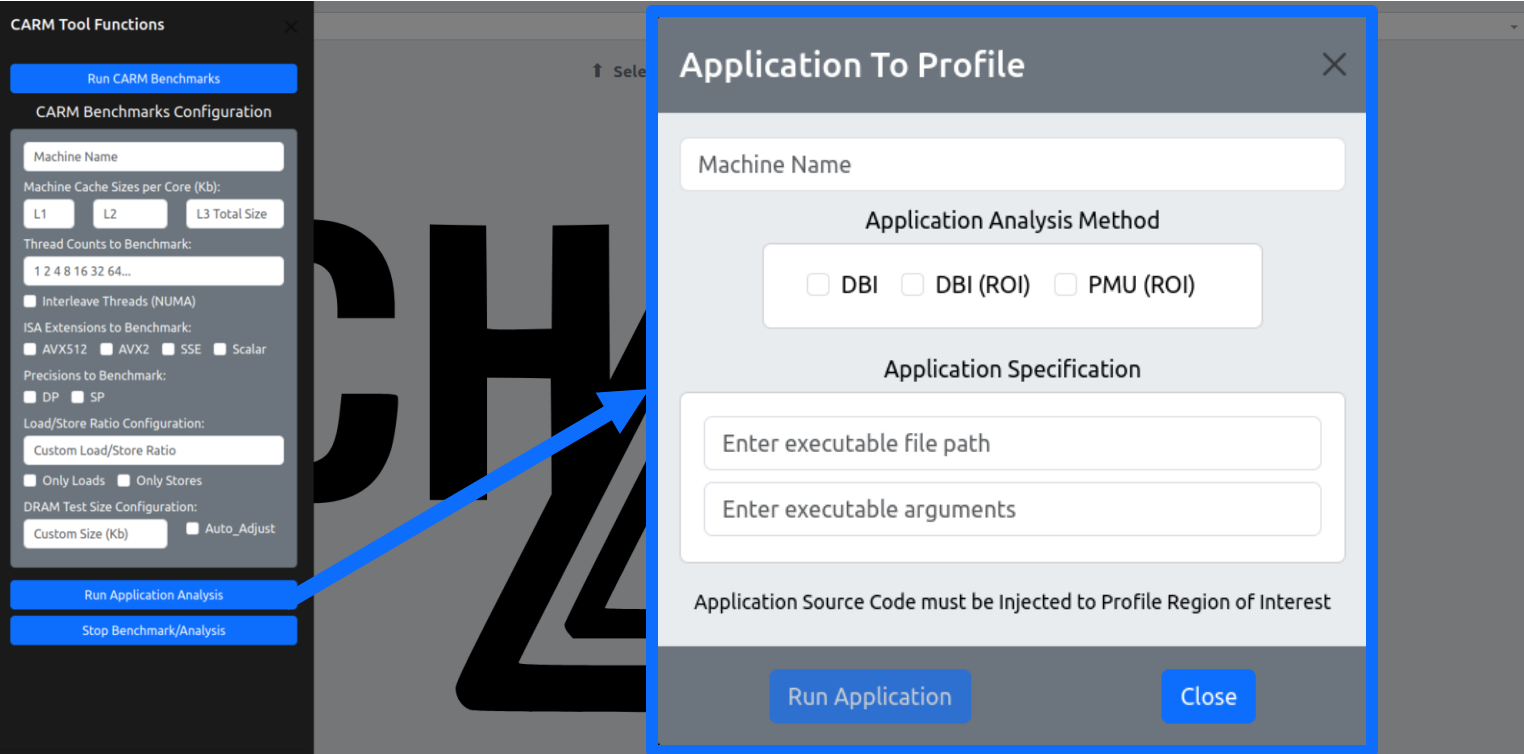
Assets

GUI Images

Graphical User Interface



Application Profiling



CARM Tool Functions

Run CARM Benchmarks

CARM Benchmarks Configuration

Machine Name

Machine Cache Sizes per Core (Kb):
L1 L2 L3 Total Size

Thread Counts to Benchmark:
1 2 4 8 16 32 64...

Interleave Threads (NUMA)

ISA Extensions to Benchmark:
 AVX512 AVX2 SSE Scalar

Precisions to Benchmark:
 DP SP

Load/Store Ratio Configuration:
Custom Load/Store Ratio

Only Loads Only Stores

DRAM Test Size Configuration:
Custom Size (Kb) Auto_Adjust

Run Application Analysis

Stop Benchmark/Analysis

Application To Profile

Machine Name

Application Analysis Method

DBI DBI (ROI) PMU (ROI)

Application Specification

Enter executable file path

Enter executable arguments

Application Source Code must be Injected to Profile Region of Interest

Run Application Close



ResultsGUI.py

Assets

GUI Images

Graphical User Interface

AVX512 Microbenchmarks Pseudo-Code

```

1  asm    volatile    (
2      "movq %0, %%r8" //Outer Loop Variable Iterations
3      "Loop2_%=:"
4      "movq %1, %%rax" //Pointer to Test Data Array
5      "movq $388, %%rdi" //Inner Loop Iterations
6      "Loop1_%=:"
7      "vmovapd 0(%%rax), %%zmm0" //Vector Load
8      "vmovapd %%zmm1, 64(%%rax)" //Vector Store
9      "vmadd132pd %%zmm2, %%zmm2, %%zmm2" //Vector FMA
10     //.....//
11     "vmovapd 16384(%%rax), %%zmm29"
12     "vmovapd %%zmm30, 16448(%%rax)"
13     "vmadd132pd %%zmm31, %%zmm31, %%zmm31"
14     "addq $16512, %%rax" //Pointer Bump
15     "subq $1, %%rdi"
16     "jnz Loop1_%= " //Inner Loop End
17     "vmovapd 0(%%rax), %%zmm0"
18     "vmovapd %%zmm1, 64(%%rax)"
19     "vmadd132pd %%zmm2, %%zmm2, %%zmm2"
20     //.....//
21     "subq $1, %%r8"
22     "jnz Loop2_%= " //Outer Loop End
23     : "r"(num_reps_t), "r" (test_var)
24     : "rax", "rdi", "r8", "zmm0-31"
25 );
  
```

RISC-V RVV Vector Length Detection

```

asm volatile
(
    "li      t0, 8192\n\t"
    "vsetvli t0, t0, e64, m1\n\t"
    "sw      t0, %[vl]\n\t"
    :
    : [vl] "m" (vec_length)
    : "t0", "t1", "t2"
);
  
```

Run.py

Bench

Config

Test

Microbenchmarks
Generation

Automatic/Manual
configuration files

Microbenchmarks
Execution

Automatic Benchmarking

Application Analysis

dbi_carm_roi.h

DBI_AI_Calculator.py

Intel_SDE_AiCalculator.py


PMU_AI_Calculator.py

CustomClient

opcoder.c

	PAPI	DynamoRIO	Intel SDE
Analysis Type	PMU Analysis	DBI Analysis	DBI Analysis
Full Application Analysis	No	Yes	Yes
ROI Analysis	Yes	Yes	Yes
Supported Architectures	Intel AMD ARM	Intel AMD ARM ? RISC-V ?	Intel AMD

PAPI





Benchmark Generation

- Following a general structure
- Adapted to each ISA extension



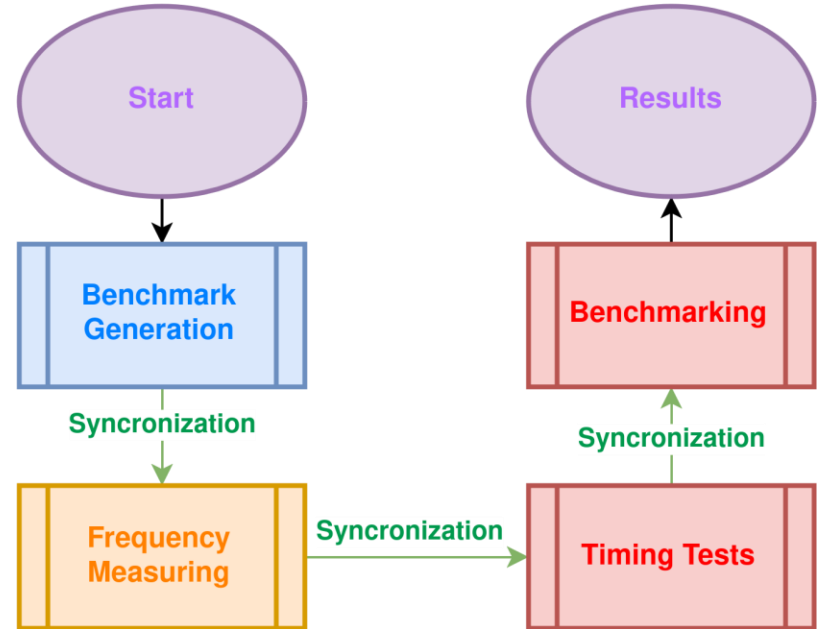
Frequency Measuring

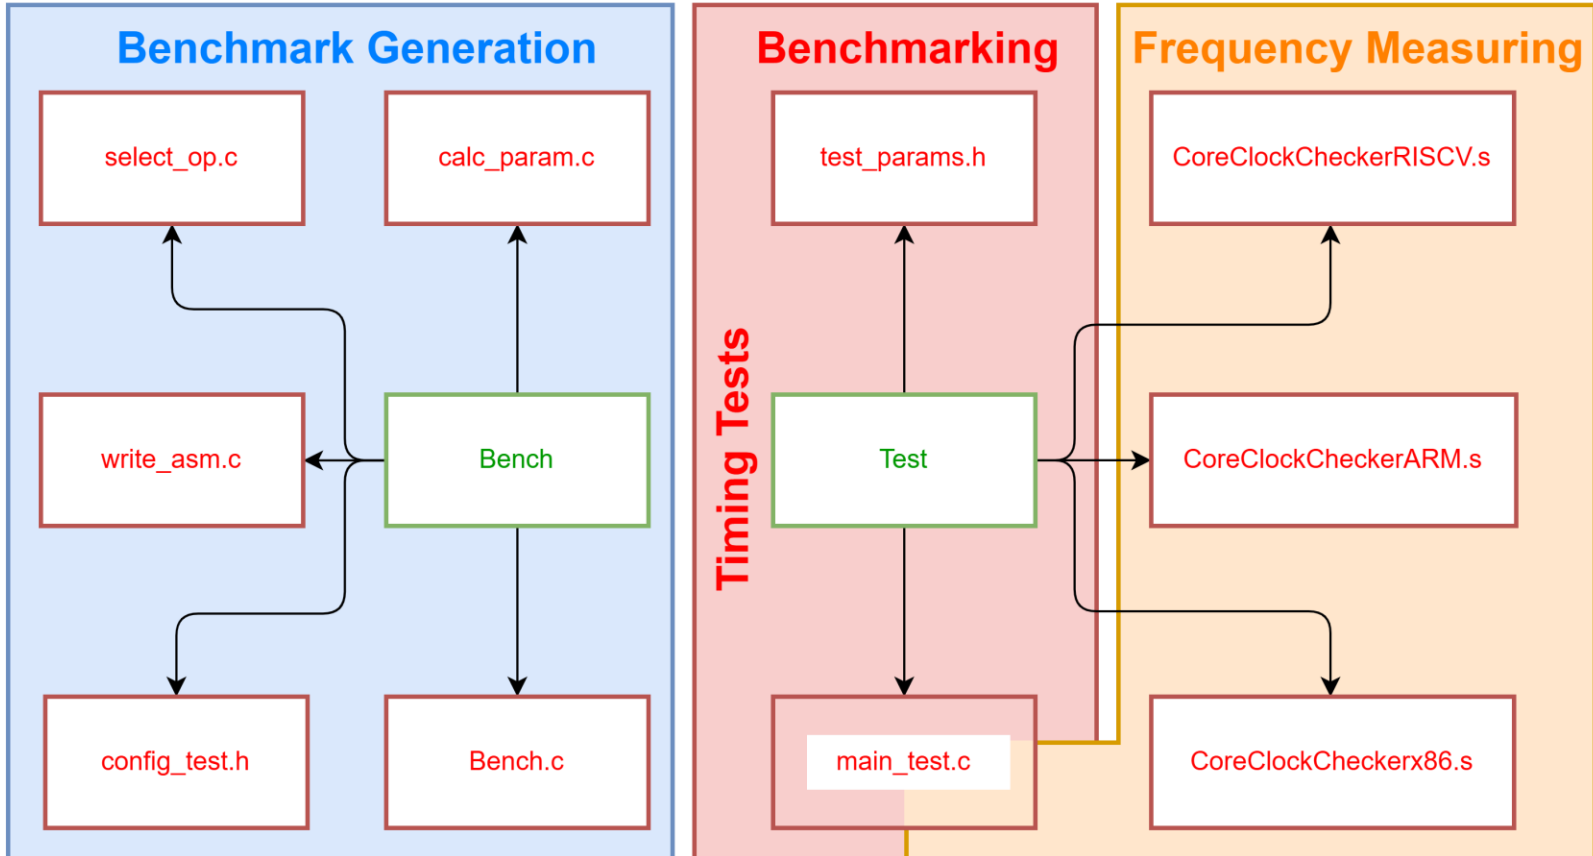
- Assembly based
- Adapted to each ISA



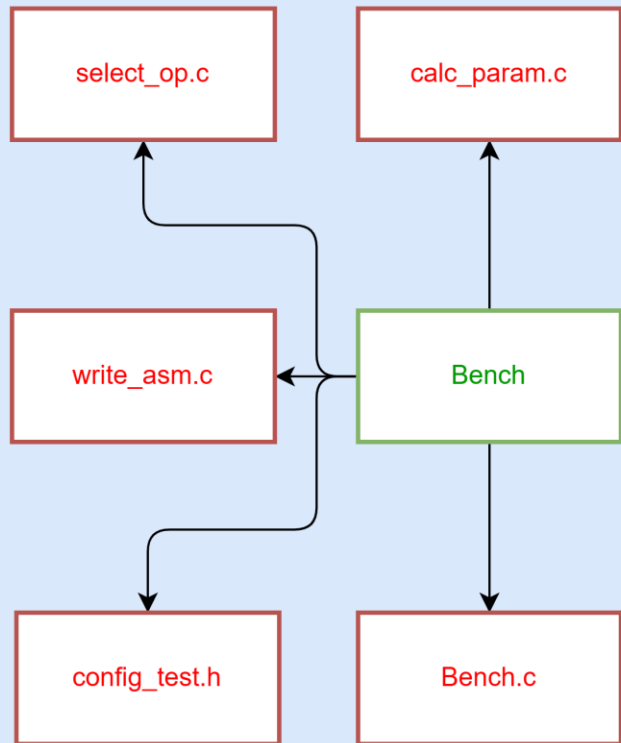
Benchmarking

- Timing tests
- Actual benchmarking





Benchmark Generation



AVX512 Microbenchmarks Pseudo-Code

```

1  asm      volatile (
2      "movq %0, %%r8" //Outer Loop Variable Iterations
3      "Loop2_%=:"
4      "movq %1, %%rax" //Pointer to Test Data Array
5      "movq $388, %%rdi" //Inner Loop Iterations
6      "Loop1_%=:"
7      "vmovapd 0(%%rax), %%zmm0" //Vector Load
8      "vmovapd %%zmm1, 64(%%rax)" //Vector Store
9      "vmadd132pd %%zmm2, %%zmm2, %%zmm2" //Vector FMA
10     //.....//
11     "vmovapd 16384(%%rax), %%zmm29"
12     "vmovapd %%zmm30, 16448(%%rax)"
13     "vmadd132pd %%zmm31, %%zmm31, %%zmm31"
14     "addq $16512, %%rax" //Pointer Bump
15     "subq $1, %%rdi"
16     "jnz Loop1_%" //Inner Loop End
17     "vmovapd 0(%%rax), %%zmm0"
18     "vmovapd %%zmm1, 64(%%rax)"
19     "vmadd132pd %%zmm2, %%zmm2, %%zmm2"
20     //.....//
21     "subq $1, %%r8"
22     "jnz Loop2_%" //Outer Loop End
23     : "r"(num_reps_t), "r"(test_var)
24     : "rax", "rdi", "r8", "zmm0-31"
25 );
  
```

Frequency Measuring Approach

- ⊗ Assembly function that leads to 1 IPC
- ⊗ Adapted to each ISA
- ⊗ Under 1% error on all tested machines

Timing methods considered

- ⊗ Time Stamp Counter (TSC) – Intel | AMD
- ⊗ Clock_gettime function – ARM | RISC-V

```

AARCH64
1  clktestarm:
2  clktest_loop:
3      add x29, x29, x8
4      add x29, x29, x8
5      add x29, x29, x8
6      add x29, x29, x8
7      add x29, x29, x8
8      add x29, x29, x8
9      add x29, x29, x8
10     add x29, x29, x8
11     //.....//
12     add x29, x29, x8
13     add x29, x29, x8
14     add x29, x29, x8
15     add x29, x29, x8
16     add x29, x29, x8
17     add x29, x29, x8
18     sub x0, x0, x9
19     cbnz x0, clktest_loop
20     ret
  
```

Frequency Measuring

CoreClockCheckerRISCV.s

CoreClockCheckerARM.s

CoreClockCheckerx86.s

Timing Tests

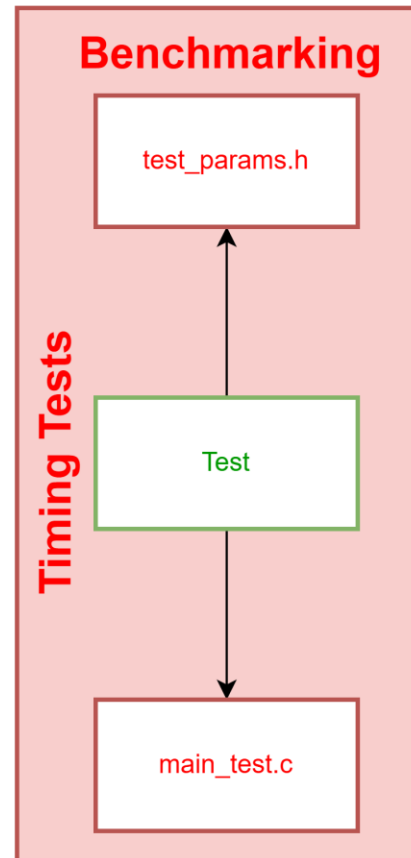
- ⚙️ Preliminary execution of benchmarks
- ⚙️ Ensures benchmarks last enough time
- ⚙️ Avoids unnecessarily long benchmarking

Benchmarking

- ⚙️ Tests are repeated 1024 times
- ⚙️ Thread barriers ensure parallel execution
- ⚙️ Best run per thread is selected

```

AVX512 Microbenchmarks Pseudo-Code
1  asm volatile (
2    "movq %0, %%r8" //Outer Loop Variable Iterations
3    "Loop2 %=: "
4    "movq %1, %%rax" //Pointer to Test Data Array
5    "movq $388, %%rdi" //Inner Loop Iterations
6    "Loop1 %=: "
7    "vmovapd 0(%%rax), %%zmm0" //Vector Load
8    "vmovapd %%zmm1, 64(%%rax)" //Vector Store
9    "vmadd132pd %%zmm2, %%zmm2, %%zmm2" //Vector FMA
10   //.....//
11   "vmovapd 16384(%%rax), %%zmm29"
12   "vmovapd %%zmm30, 16448(%%rax)"
13   "vmadd132pd %%zmm31, %%zmm31, %%zmm31"
14   "addq $16512, %%rax" //Pointer Bump
15   "subq $1, %%rdi"
16   "jnz Loop1 %=" //Inner Loop End
17   "vmovapd 0(%%rax), %%zmm0"
18   "vmovapd %%zmm1, 64(%%rax)"
19   "vmadd132pd %%zmm2, %%zmm2, %%zmm2"
20   //.....//
21   "subq $1, %%r8"
22   "jnz Loop2 %=" //Outer Loop End
23   : "r"(num_reps_t), "r"(test_var)
24   : "rax", "rdi", "r8", "zmm0-31"
25 );
  
```



 State of the Art

 The CARM Tool

 **Results**

 Conclusion



CARM based Architecture Analysis



Comparison with state of the art Roofline Tools



Intel Advisor





Empirical Roofline Toolkit




CARM based SpMV Application Analysis

Machines Utilized

-  One from each vendor
-  Covering all ISA extensions supported

Analysis Objective

-  Verify if CARM benchmarks can reach architectural limits


	Venus	Cara	Armq	Milk-V
Vendor	Intel	AMD	ARM	RISC-V
Architecture	Skylake-X	Zen 3	Vulcan	XuanTie C920
ISA Extensions	SSE AVX2 AVX-512	SSE AVX2	Neon	RVV 0.7.1
FP LD/ST Units	2 3	2 3	2 2	2 1

Best Load/Store Ratio

- Follows Ld/St unit ratio
- Loads outperform stores

Memory Architectural Limits

- Accurately achieved for L1
- Progressively harder to reach for lower memory levels




Best Load/Store Ratio

- Follows Ld/St unit ratio
- Loads outperform stores

Memory Architectural Limits

- Accurately achieved for L1
- Progressively harder to reach for lower memory levels




Best Load/Store Ratio

- Follows Ld/St unit ratio
- Loads outperform stores

Memory Architectural Limits

- Accurately achieved for L1
- Progressively harder to reach for lower memory levels




Best Load/Store Ratio



- Follows Ld/St unit ratio
- Loads outperform stores

Memory Architectural Limits



- Accurately achieved for L1
- Progressively harder to reach for lower memory levels



Best Load/Store Ratio

-  Follows Ld/St unit ratio
-  Loads outperform stores

Memory Architectural Limits

-  Accurately achieved for L1
-  Progressively harder to reach for lower memory levels

Deviation / IPC*	Venus	Cara	Armq	Milk-V
L1	-1.54%	-0.093%	-0.01%	-0.01%
L2	+18.57%	-0.01%	0.9 IPC	0.4 IPC
L3	-17%	-16.7%	0.75 IPC	0.1 IPC


*IPC – Instructions per Cycle

Architectural Limits



- Accurately achieved
- Intel slows down with wider ISA extensions

Mixed Benchmark Validation



- FMA leads to more deviation
- Wider ISA extensions lead to more deviation



Architectural Limits



-  Accurately achieved
-  Intel slows down with wider ISA extensions

Mixed Benchmark Validation



-  FMA leads to more deviation
-  Wider ISA extensions lead to more deviation

	Intel	AMD	ARM	RISCV
Scalar	-0.78%	-0.15%	+0.54%	+0.66%
Widest ISA	-5.9%	-0.12%	+0.85%	+0.63%

Architectural Limits

-  Accurately achieved
-  Intel slows down with wider ISA extensions

Mixed Benchmark Validation

-  FMA leads to more deviation
-  Wider ISA extensions lead to more deviation

Deviation	Venus	Cara	Armq	Milk-V
Mixed Scalar Add	-3.44%	-1.11%	-17.72%	-4.55%
Mixed Scalar FMA	-5.34%	-10.27%	-17.25%	-8.29%
Mixed Widest ISA Add	-3.44%	-0.16%	-25.93%	-10.45%
Mixed Widest ISA FMA	-24.83%	-9.14%	-24.88%	-10.57%



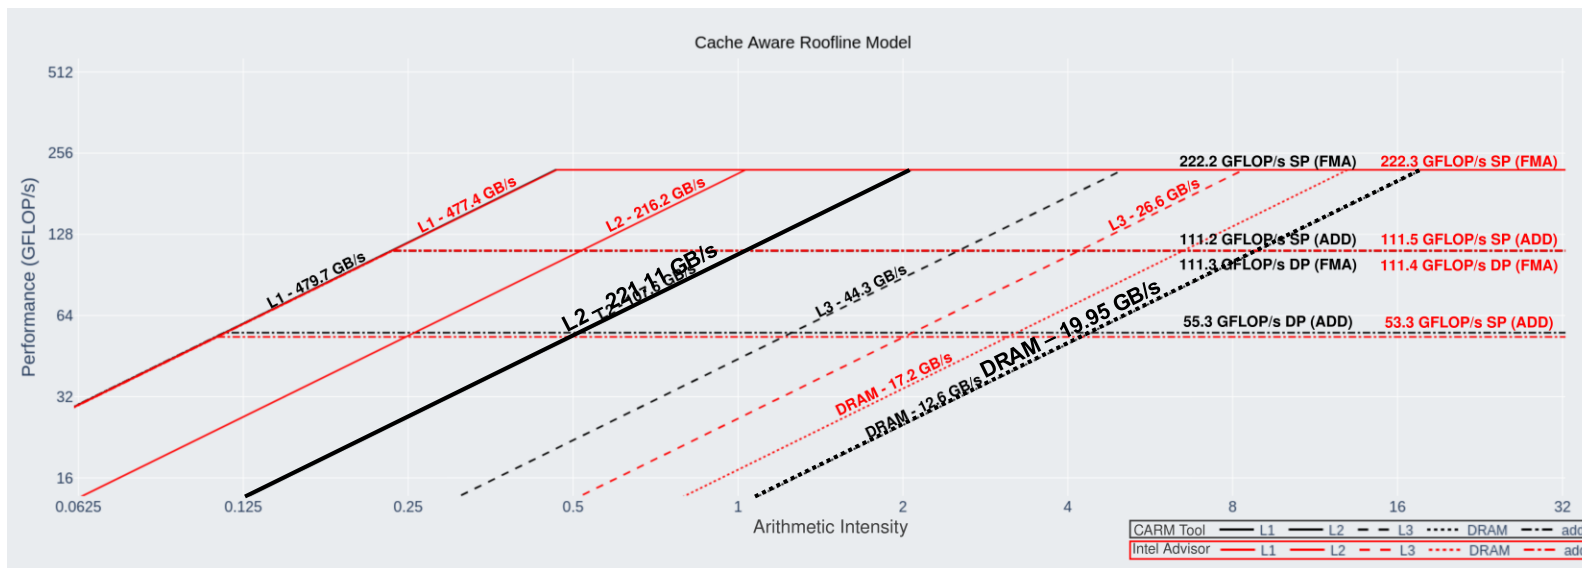
Comparison with ERT

- Higher L1 bandwidth and GFLOPS achieved
- Lower level discrepancies due to imprecise cache size values



Comparison with Intel Advisor

- Similar L1 bandwidth and GFLOPS achieved
- Lower level discrepancies due to Id/st ratio variations

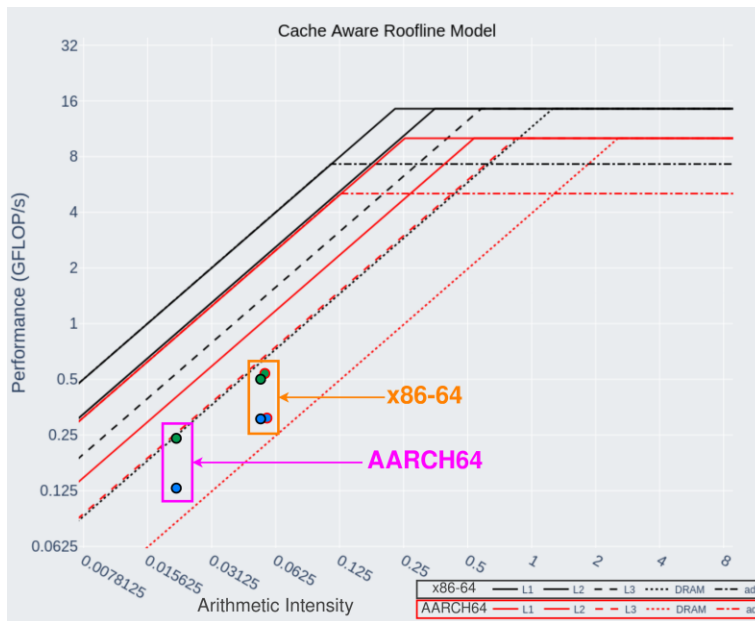


🔧 Cross-Architecture SpMV Analysis

- 🔧 Using the Eigen library
- 🔧 SpMV performance comparison

🔧 SpMV Performance

- 🔧 RCM Re-Ordering improves performance

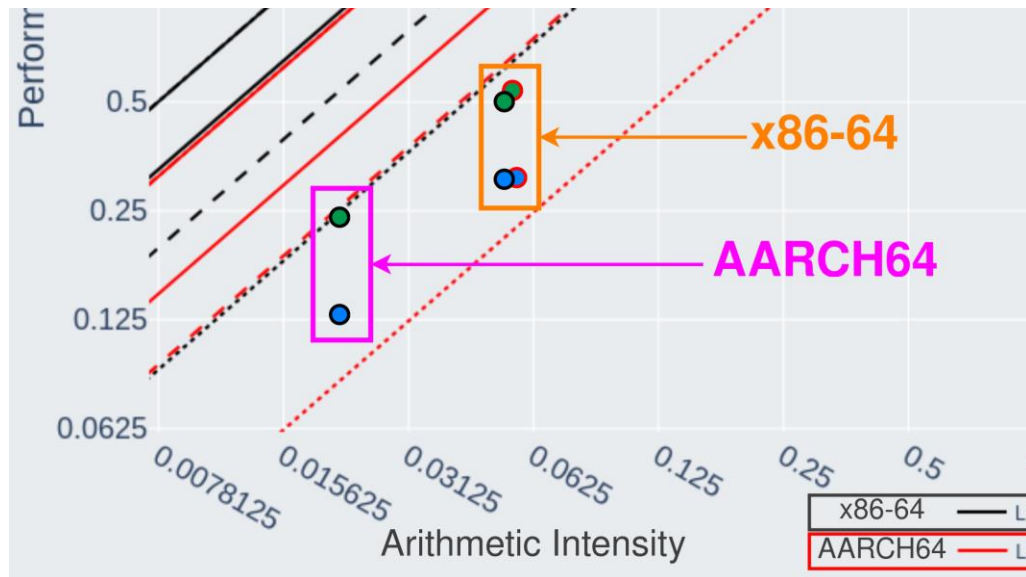


🔧 Cross-Architecture SpMV Analysis

- 🔧 Using the Eigen library
- 🔧 SpMV performance comparison

🔧 SpMV Performance

- 🔧 RCM Re-Ordering improves performance




● DBI - Original Matrix
 ● DBI - RCM Matrix

🔧 Cross-Architecture SpMV Analysis

- 🔧 Using the Eigen library
- 🔧 SpMV performance comparison

🔧 SpMV Performance

- 🔧 RCM Re-Ordering improves performance



● DBI - Original Matrix
 ● DBI - RCM Matrix
 ● PMU - Original Matrix
 ● PMU - RCM Matrix

 State of the Art

 The CARM Tool

 Results

 **Conclusion**

- ❏ Current collaborations
 - ❏ CERN, SYCLOPS, POP3, BSC
- ❏ Complete collaborations
 - ❏ SparCity – SuperTwin Live-CARM
- ❏ CARM Tool / CHAMP Hub Github and Paper

 CERN – Adaptive Perf Tool

 BSC – Paraver Tool

 SYCLOPS – SYCL DB

 POP3 – Application Profiling




🔌 Sparcity – SuperTwin

- 🔌 CARM Tool source code shipped with the tool
- 🔌 SuperTwin interfaces with the CARM Tool to get CARM results

🔌 Live-CARM

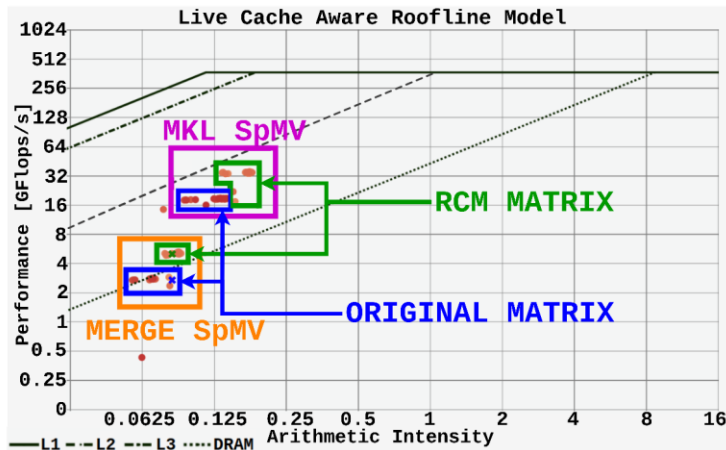
- 🔌 Based on live performance counter data from SuperTwin
- 🔌 Application and benchmark analysis was conducted

Live Cache Aware Roofline Model




🔧 Sparcity – SuperTwin

- 🔧 CARM Tool source code shipped with the tool
- 🔧 SuperTwin interfaces with the CARM Tool to get CARM results



🔧 Live-CARM

- 🔧 Based on live performance counter data from SuperTwin
- 🔧 Application and benchmark analysis was conducted



- 🔌 The CARM Tool is open source and available on Github
 - 🔌 <https://github.com/champ-hub/carm-roofline>
- 🔌 The CARM Tool's paper is accepted for publication in IISWC24
- 🔌 The first of many tools to be developed in the scope of CHAMP hub



Heterogeneous Computing and Performance Modeling Hub

Hub para Computação Heterogénea e Modelação de Performance

Thank You

Any questions?

**José Morgado
Leonel Sousa
Aleksandar Ilic**

**28th POP Seminar
5 September 2024**