


# The CARM Tool

Cache-aware Roofline Model for  
HPC=CPU+GPU


Aleksandar Ilic  
Leonel Sousa  
José Morgado  
Alexandre Rodrigues  
Diogo Matos

- Modern HPC systems and applications are **complex** and **heterogeneous**
  - Hard to model
  - Hard to optimize
- The CARM as a solution



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 was originally supported only by Intel Advisor (commercial tool)



## 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



# State of the art – Roofline Tools

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

- And about ARM and RISC-V?  
and GPUs?



# The CARM Tool – Supported ISAs

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

- ☐ The CARM Tool
- ☐ Extension to GPU
- ☐ Integration: HPC Tools, Centres and Codesign

## Interfacing


- Command Line Interface
- Graphical User Interface

## Automatic Benchmarking


- CARM Benchmarks
- Memory / FP / Mixed Benchmarks

## Application Analysis

- PMU Analysis
- DBI Analysis




# The CARM Tool – High Level





## Results Visualization



ResultsGUI.py

Graphical User Interface

Assets

GUI Images



## Benchmark Execution

**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

↑ Select a Machine to View CARM Results ↑



ResultsGUI.py


Graphical User Interface

Assets

GUI Images



## Application Profiling




ResultsGUI.py

Graphical User Interface

Assets

GUI Images



## 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      "vfmadd132pd %%zmm2, %%zmm2, %%zmm2" //Vector FMA
10     //.....
11     "vmovapd 16384(%rax), %%zmm29"
12     "vmovapd %%zmm30, 16448(%rax)"
13     "vfmadd132pd %%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     "vfmadd132pd %%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




# The CARM Tool – Application Analysis




# The CARM Tool – Application Analysis





# PAPI







# PAPI




# The CARM Tool – Application Analysis






# PAPI





# PAPI





The screenshot shows a web-based interface for the CARM Tool. At the top, there is a header bar with a profile icon, the text "CARM Tool Demo results...", and a "Share" button. Below the header, a message reads "Select a Machine to View CARM Results". The central feature is a large, bold, black "CHAMP" logo. Inside the letter "A" of the logo is a red YouTube play button icon. At the bottom of the interface, the text "Plot CARM and Application Analysis Results" is displayed in large, bold, black letters. In the bottom left corner, there is a small button with the text "Watch on YouTube".

CHAMP

Watch on YouTube

Plot CARM and Application Analysis Results


# Results – Application Analysis

## Cross-Architecture SpMV Analysis

- Using the Eigen library
- SpMV performance comparison

## SpMV Performance

- RCM Re-Ordering improves performance



- ⌚ The CARM Tool
- ⌚ **Extension to GPU**
- ⌚ Integration: HPC Tools, Centres and Codesign

- Nvidia GPU extension + Tensor CARM
  - Functionality to test both CUDA Cores and Tensor Cores
  - Functionality to test Shared Memory, L2 cache and Global Memory
- Allows for the profiling of GPU applications using Nsight Compute
  - Functionality to profile entire applications or target individual kernels
  - Potential shown in the optimization of CNNs
- AMD GPU extension
  - Complete implementation for AMD MI GPUs using HIP
  - Supporting benchmarking using AMD RocProf

## Maximum measured performance

- Target instruction in PTX assembly
- Thousands of iterations, thousands of threads


## Different data precisions

- Set of different instruction shapes for mat-mul


```
// data preparation in CUDA C++
// ...
mov.u64 %time_start, %%globaltimer;
for i in 0 to Iter do:
    mma.sync.aligned.shape.row.col.precision C, A, B, C;
    // loop unrolling
end
mov.u64 %time_end, %%globaltimer;
```

precision: `s32.s8.s8.s32`  
`f32.f16.f16.f32`  
`f32.tf32.tf32.f32`  
`...`

shape: `int8`  
`m8n8k16`  
`m16n8k16`  
`m16n8k32`



Within 0.5% on the RTX 3080  
 2.5% on the A100 (some exceptions)



- ⌚ Use CARM to evaluate GPU kernels

- ⌚ GEMM kernels provided by cuBLAS and CUTLASS
- ⌚ Deep learning kernels provided by cuDNN






- ⌚ Nsight Compute as profiler tool




- ⌚ Bytes transferred, FLOPs executed and execution time
- ⌚ Cannot measure integer operations

# GEMM and Convolutional Kernels

## GEMM




## Convolve



- The CARM Tool
- Extension to GPU
- **Integration: HPC Tools, Centres and Codesign**

- Tools in HPC Centers are comprehensive, but mostly oriented to performance analysis not architecture modelling
- Close collaboration and extensions with SOTA tools
  - Paraver at BSC
  - GPP Codesign at Codasip
  - Adaptyst at CERN
- Great to integrate computer architecture models such as CARM!



## ⌚ CARM Based Application Profiling


- ⌚ Via the CARM Tool
- ⌚ Via Extrae traces and Paraver integration

## ⌚ Deployment to HPC Centres

- ⌚ Karolina automated deployment complete
- ⌚ Available on EuroHPC centres soon

## ⌚ CARM Dissemination Initiatives

- ⌚ POP3 Webinar
- ⌚ VI-HPS Tools Guide and Blog
  - ⌚ <https://www.vi-hps.org/tools/carm.html>
  - ⌚ <https://pop-coe.eu/blog/28th-pop-webinar-the-carm-tool-cache-aware-roofline-model-for-hpc>



## □ CARM Based Application Profiling in Paraver

- Via analysis of Extrace traces
- CARM Tool provides the CARM results
- Integrated with the Paraver Tool as a separate module


## □ CARM-Paraver Integration Features

- Syncronization with Paraver timelines
- Exporting Paraver timestamp coloring and masks to the CARM GUI
- Exporting CARM related metrics to Paraver



## Synchronization with Paraver Timelines

- Synchronization point defined via Paraver GUI
- Re-sync possible via CARM GUI




## DARE SGA1 Project

- Driving Europe's next-generation HPC systems
- Focus on performance, efficiency, and sovereignty


## Codasip: GPP Pathfinding

- Low-level architecture-aware support for CARM
- Help optimising the apps running on the GPP core
- Insights about characteristics of the architecture

## Collaborate with the DARE SW TA



- The CARM Tool is open source and available on Github
  - <https://github.com/champ-hub/carm-roofline>
- 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?

Aleksandar Ilic  
Leonel Sousa  
José Morgado  
Alexandre Rodrigues  
Diogo Matos