

# The CARM Tool

José Morgado Leonel Sousa Aleksandar Ilic

46th VI-HPS Tuning Workshop 4 September 2024

## **Motivation**



- 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

## Challenges



- 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





## The Cache-aware Roofline Model

#### 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



A. Ilic, F. Pratas and L. Sousa, "Cache-aware Roofline model: Upgrading the loft"



## The Cache-aware Roofline Model

#### 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







#### Intro

#### The CARM Tool

#### Results

#### Conclusion



### **The CARM Tool**

#### High-Level Overview

- Graphical User Interface
- Automatic Benchmarking
- Application Analysis

#### Low-Level Overview

- Benchmark Generation
- Frequency Measuring
- Benchmarking



## The CARM Tool – High Level

#### Interfacing

- Command Line Interface
- Graphical User Interface

- Automatic Benchmarking
  - CARM Benchmarks
  - Memory / FP / Mixed Benchmarks
- Application Analysis
  - PMU Analysis
  - DBI Analysis





### 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 – High Level



6



### The CARM Tool – Graphical User Interface

#### Results Visualization





### The CARM Tool – Graphical User Interface

#### Benchmark Execution





### The CARM Tool – Graphical User Interface

#### Application Profiling

CARM Tool Functions				•	
Run CARM Benchmarks	î Sele	Application to Profile	X		ResultsGUI.py
CARM Benchmarks Configuration					
Machine Name Machine Cache Sizes per Core (Kb):		Machine Name			្ន
L1 L2 L3 Total Size		Application Analysis Method			raph
1 2 4 8 16 32 64 Interleave Threads (NUMA)		🗌 DBI 📄 DBI (ROI) 📄 PMU (ROI)			lical
ISA Extensions to Benchmark: AVX512 AVX2 SSE Scalar Precisions to Benchmark: DP SP		Application Specification			User
Load/Store Ratio Configuration: Custom Load/Store Ratio		Enter executable file path			Assets
DRAM Test Size Configuration: Custom Size (Kb)		Enter executable arguments			
Run Application Analysis		Application Source Code must be Injected to Profile Region o	of Interest		CLILImagaa
		Run Application Close			

### The CARM Tool – Automatic Benchmarking

	AVX512 Microbenchmarks Pseudo-Code
1	asmvolatile (
	<pre>"movq %0, %%r8" //Outer Loop Variable Iterations</pre>
	"Loop2_%=:"
	<pre>"movq %1, %%rax" //Pointer to Test Data Array</pre>
	<pre>"movq \$388, %%rdi" //Inner Loop Iterations</pre>
	"Loop1_%=:"
	<pre>"vmovapd 0(%%rax), %%zmm0" //Vector Load</pre>
	<pre>"vmovapd %%zmm1, 64(%%rax)" //Vector Store</pre>
	"vfmadd132pd %%zmm2, %%zmm2, %%zmm2" //Vector FM
LO	
11	"vmovapd 16384(‰rax), ‰zmm29"
12	"vmovapd %%zmm30, 16448(%%rax)"
13	"vfmadd132pd %%zmm31, %%zmm31, %%zmm31"
L4	<pre>"addq \$16512, %%rax" //Pointer Bump</pre>
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)
	:"rax","rdi","r8","zmm0-31"
25	);

TÉCNICO

LISBOA





### The CARM Tool – Application Analysis

Application Analysis							
dbi_carm_roi.h	DBI_AI_Calculator.py Intel_SDE_AiCalculator.py PMU_AI_Calculator.py						
	PAPI DynamoRIO Intel SDE						
	Analysis Type	PMU Analysis	DBI Analysis	DBI Analysis			
CustomClient	Full Application Analysis	No	Yes	Yes			
	ROI Analysis	Yes	Yes	Yes			
opcoder.c	Supported Architectures	Intel   AMD   ARM	Intel   AMD   ARM   ? RISC-V ?	Intel   AMD			





### The CARM Tool – Low Level

#### Benchmark Generation

- Following a general structure
- Adapted to each ISA extension

#### Frequency Measuring

- Assembly based
- Adapted to each ISA

### Benchmarking

- Timing tests
- Actual benchmarking





### The CARM Tool – Low Level





### The CARM Tool – Benchmark Generation



AVX512 Microbenchmarks Pseudo-Code					
1	asmvolatile_ (				
2	<pre>"movq %0, %%r8" //Outer Loop Variable Iterations</pre>				
3	"Loop2_%=:"				
4	<pre>"movq %1, %%rax" //Pointer to Test Data Array</pre>				
5	<pre>"movq \$388, %%rdi" //Inner Loop Iterations</pre>				
6	"Loop1_%=:"				
7	"vmovapd 0(‰rax), ‰zmm0" //Vector Load				
8	<pre>"vmovapd %%zmm1, 64(%%rax)" //Vector Store</pre>				
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	);				



#### Frequency Measuring Approach

- Assembly function that leads to 1 IPC
- Adapted to each ISA

TÉCNICO

LISBOA

Under 1% error on all tested machines

#### Timing methods considered

- Time Stamp Counter (TSC) Intel | AMD
- Clockgettime function ARM | RISC-V

AA	RCH64
1	clktestarm:
0	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



## The CARM Tool – Bechmarking

#### 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









State of the Art

The CARM Tool

Results

Conclusion







#### CARM based Architecture Analysis

- Comparison with Intel Advisor
- CARM based SpMV Application Analysis



### **Results - CARM Architecture 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

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





#### Best Load/Store Ratio

- Follows Ld/St unit ratio
- Loads outperform stores

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





#### Best Load/Store Ratio

- Follows Ld/St unit ratio
- Loads outperform stores

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





- Best Load/Store Ratio
  - Follows Ld/St unit ratio
  - Loads outperform stores

- 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



### **Results - CARM FP Arithmetic Analysis**

#### 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





### **Results - CARM FP Arithmetic Analysis**

#### 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%



### **Results - CARM FP Arithmetic Analysis**

#### 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%



### **Results – CARM Comparison**

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





## **Results – Application Analysis**

- Cross-Architecture SpMV Analysis
  - Using the Eigen library
  - SpMV performance comparison

#### SpMV Performance

RCM Re-Ordering improves performance





## **Results – Application Analysis**

- Cross-Architecture SpMV Analysis
  - Using the Eigen library
  - SpMV performance comparison

- SpMV Performance
  - RCM Re-Ordering improves performance





## **Results – Application Analysis**

- Cross-Architecture SpMV Analysis
  - Using the Eigen library
  - SpMV performance comparison

- SpMV Performance
  - RCM Re-Ordering improves performance









State of the Art

The CARM Tool

Results

Conclusion







#### CARM Tool / CHAMP Hub Github and Paper



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

The CARM Tool's paper is currently in press and will soon be published 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

J. Morgado, L. Sousa, A. Ilic. "CARM Tool: Cache-Aware Roofline Model Automatic Benchmarking and Application Analysis", IISWC, 2024



# Thank You

### Any questions?

José Morgado Leonel Sousa Aleksandar Ilic

46th VI-HPS Tuning Workshop 4 September 2024