#### TECHNISCHE UNIVERSITÄT DARMSTADT Parallel Programming

#### **Conquering Noise with Hardware Counters on HPC Systems**

Marcus Ritter<sup>1</sup>, **Ahmad Tarraf**<sup>1</sup>, Alexander Geiß<sup>1</sup>, Nour Daoud<sup>2</sup>, Bernd Mohr<sup>2</sup>, Felix Wolf<sup>1</sup> <sup>1</sup>Technical University of Darmstadt

<sup>2</sup>Forschungszentrum Jülich GmbH





# Conquering Noise with Hardware Counters on HPC Systems

Marcus Ritter<sup>1</sup>, Ahmad Tarraf<sup>1</sup>, Alexander Geiß<sup>1</sup>, Nour Daoud<sup>2</sup>, Bernd Mohr<sup>2</sup>, Felix Wolf<sup>1</sup>

<sup>1</sup>Technical University of Darmstadt

<sup>2</sup>Forschungszentrum Jülich GmbH

### Motivation



Programming

Performance and complexity of HPC systems are constantly increasing

- → Important to examine the scaling behavior of an application and identify early performance bottlenecks
- $\rightarrow$  Use empirical performance modeling

#### **Problem:**

In **noisy** environments  $\rightarrow$  difficult to create accurate performance models

- Strong variations in the measurements
- Measurements irreproducible and misleading
- Strong deviations from the actual application behavior

### Motivation



Programming

#### Problem (cont.):

- Application runtime affected by noise
  - Most common performance metric



#### Solution:

Use hardware counters

- Noise has little impact on some hardware counters
  - e.g., double precision operations
- Selecting the right counters requires a thorough analysis

### Contributions



A detailed noise analysis on various hardware counters on different systems:

 $\rightarrow$  Total of **26950** experiments (PAPI preset events only):

Five systems	With and without injected noise
<b>Four</b> hardware architectures	Multiple resource configurations (number of nodes)
Three applications	<b>Five</b> repetitions per setup

**Categorized** the counters across the different systems according to their noise resilience and provided a **user guide** 

## **Analysis Methodology**

Find noise-resilient hardware counters:

- Examine if counters' values change when repeating the measurements
- Expose the counters to different levels of noise
  - Inject different noise patterns using NOIGENA
  - NOIGENA processes were running on the odd processors





PATTERN_1:
Sequence:
- REPEATED_NOISE:
REPEAT: inf
Sequence:
- NETWORK_NOISE: 2
- NO_NOISE: 2
- MEMORY_NOISE: 2
- NO_NOISE: 2

Noise pattern used by NOIGENA to configure the amount and duration of generated noise.

## **Analysis Methodology**



Parallel Programming

Compare counter values of different call paths from three applications for the repeated experiments

 $\rightarrow$  Calculate the **relative deviation** from the arithmetic mean in percent:

$$\frac{|v_i - \bar{v}|}{\bar{v}} * 100\%$$

For the  $a^{th}$  application kernel,  $p^{th}$  MPI rank,  $t^{th}$  OpenMP thread, and  $i^{th}$  repetition,

 $\rightarrow$  Find the arithmetic mean across the **repetitions**:

$$\bar{v}_{a,p,t} = \operatorname{mean}(v_{a,p,t,i}) \rightarrow \bar{v} = \operatorname{mean}(v_i)$$

We do this for all counters on all systems with and without noise

### **Application Benchmarks**



Parallel Programming



For all of them, we use OpenMP and MPI for the measurements

#### **Evaluation Systems**



Alias	Name	Nodes	Processor	RAM	Network	
СМ	DEEP-EST, Cluster Module	50	2x Intel Xeon Skylake Gold 6146 CPUs (12 cores, 24 threads)	192 GB DDR4 RAM (2666 MHz)	InfiniBand EDR (100 GBit/s)	
ESB	DEEP-EST, Extreme Scale Booster	75	1x Intel Xeon Cascade Lake Silver 4215 CPU (8 cores, 16 threads)	48 GB DDR4 RAM (2400 MHz),	InfiniBand EDR (100 GBit/s)	
Jureca	JURECA DC Module, std. compute nodes	480	2x AMD EPYC 7742 CPUs (64 cores, 128 threads)	512 GB DDR4 RAM (3200 MHz)	InfiniBand HDR100 (100 GBit/s)	
Jetson	OACISS, Franken-cluster Jetson ARM64	12x Jetson Tegra TX1	1x Quad-Core ARM Cortex®- A57 MPCore (4 cores, 4 threads)	4 GB 64-bit LPDDR4 RAM	1 GBit/s ethernet	
Cyclops	OACISS, Franken-cluster Cyclops	1	2x 20c IBM Power9 CPUs (20 cores, 80 threads)	384 GB of RAM	BNX2 10G Ethernet NICs, 2x Infiniband EDR (25 GBit/s)	

### **Application Configurations**



App	System	Experiment configuration
MiniFE	СМ	$n = [1, 2, 4, 8], p = n, t = 12p, s = 20^{3}p$
	ESB	$n = [1, 2, 4, 8, 12, 16, 32], p = n, t = 8p, s = 50^3 p$
	Jureca	$n = [4, 8, 12, 16, 20, 24], p = n, t = 128p, s = 20^3 p$
	Jetson	$n = [2, 3, 4, 5, 6],  p = n$ , $t = 2p,  s = 50^3 p$
	Cyclops	$n = 1, p = [4, 8, 12, 16, 20], t = 4p, s \approx 24^3 p$
LULESH	СМ	$n = [1, 8, 27], p = n, t = 12p, s = 10^3 p$
	ESB	$n = [1, 8, 27], p = n, t = 8p, s = 30^3 p$
	Jureca	$n = [1, 8, 27], p = n, t = 128p, s = 10^3 p$
	Jetson	$n = 8, p = n, t = 2p, s = 15^3 p$
	Cyclops	$n = 1, p = [1, 8, 27, 64], t = p, s = 5^3 p$
LAMMPS	СМ	$n = [1, 2, 4, 8], p = n, t = 12p, s = 20^{3}p$
	ESB	$n = [1, 2, 4, 8, 12, 16, 32], p = n, t = 8p, s = 20^3 p$
	Jureca	$n = [1, 2, 4, 8, 16], p = n, t = 128p, s = 20^3 p$
	Jetson	$n = [2, 3, 4, 5, 6],  p = n,  t = 2p,  s = 20^3 p$
	Cyclops	$n = 1, p = [4, 8, 12, 16, 20], t = 4p, s = 20^{3}p$

### **Evaluation Results**



Parallel Programming

Visualizing the results for **each counter** is not an easy task:



To compare distinct counter  $\rightarrow$  scale the plots with the peak occurrence of the relative deviation

### **Evaluation Results**

- The height of each distribution plot describes how often the corresponding relative deviation occurs
- Maximum value of the peak relative deviation was used to scale all callpaths
- We only show callpaths that resulted in more than 1% of the total counter value





### **Evaluation Results**

- The intensity (alpha) shows the importance of the region to the overall behavior
- The importance is the callpath's share of the total counted value across all applications per counter
- Large values indicate the parts of the application that exhibit more of the behavior described by that counter



Programming





#### **Evaluation Results: How to Interpret the Results**



Parallel Programming

#### **Bad Counters**

Counter strongly influenced by noise:

 Large distribution in the presence of noise, while short one in the absence Counter with large deviation:

- Large distribution disregarding the noise
- Not suited for modeling



#### **Evaluation Results: How to Interpret the Results**



Parallel Programming

#### **Good Counters**



Counter robust against noise:

- Similar distribution in the presence and absence of noise
- Short deviation ( < 20%)</p>



Counter to some extent robust against noise:

- Short deviation (< 20%) without noise</p>
- Short deviation (< 20%) in the presence of noise for the significant callpaths

#### **Evaluation Results: Runtime**





#### **Evaluation Results: Floating Point Operations**







#### **Evaluation Results: Branch Instructions, Cycles**

**TECHNISCHE** 

UNIVERSITÄT DARMSTADT

Parallel

#### **Evaluation Results:** L1, L2 Instruction Cache



TECHNISCHE UNIVERSITÄT DARMSTADT



#### **Evaluation Results: Branch Instructions, Cycles**



Programming

18 April 2023

#### **Evaluation Results: Stalls**





### **Evaluation Results:** L2, L3 Cache





#### Impact of the Application: Total Cycles and Load Inst.







LAMMPS







### Best Practice User Guide General



Programming

In a noisy environment, the best counters independent of the system architecture are

- All counters measuring floating point operations, e.g., DP\_OPS, VEC\_SP
- All counters measuring floating point instructions, e.g., FP\_INS

#### Best Practice User Guide Intel CPUs



Parallel Programming

#### Good

- Instructions, Cycles
  - Strong deviation for less important callpaths
  - Small deviation for important callpaths

- Stall and reference cycles
  - Because of large deviations
- L3 counters and TLB\_DM
  - Accurate only in the absence of noise

#### Best Practice User Guide AMD CPUs



#### Parallel Programming

#### Good

- L2 cache accesses and misses, TLB\_DM, TLB\_IM
  - Have a small deviation with a maximum of 20%

- Total and branch instructions have less deviation than time
  - Still too much to be useful for modeling

#### Best Practice User Guide ARM CPUs



Parallel Programming

#### Good

- Instruction and cycle counters
  - Only small deviations in noisy environments
  - Can be used for performance modeling

- Exception: TOT\_CYC
  - Showing a high deviation

#### Best Practice User Guide IBM CPUs



Parallel Programming

#### Good

 Most counters have a small deviation (less than 30%) for significant callpaths

- BR\_TKN, BR\_PRC, LD\_INS, LST\_INS
  - Should be avoided
- L3 related counters
  - High deviations in the presence of noise

#### **Best Practice User Guide**



	Jureca (AMD)	ESB (Intel)	CM (Intel)	Jetson (ARM)	Cyclops (IBM)
Floating point ops./instr.	++	++	++	++	++
Cycles	-	+	+	0	0
Instructions	-	+	+	+	+
L1	+	-	-	0	+
L2	+	0	-	0	+
L3		-	-		-

### Conclusion



- Examined noise resilience of hardware counters on five systems with different architectures (Intel, AMD, ARM, IBM Power9)
- Analyzed all available preset and a selection of native events
- In general, most hardware counters are affected by noise, but still less than the runtime
- Counters measuring floating-point operations or instructions are noise resilient on all systems
- Their reliability significantly depends on the system architecture
- Our best practice guide enables developers to identify the relevant counters for performance analysis for their system

### **Future work**

TECHNISCHE UNIVERSITÄT DARMSTADT Parallel

Programming

- Create a Tool that performs the analysis and evaluation on the system
- Examine more counters (i.e., native counters)
- Examine the correlation between the counters
- Identify the source of noise
- Started working with the PAPI developers





Parallel Programming

# **Questions?**

Thank you for your attention!

18 April 2023

Department of Computer Science | Laboratory for Parallel Programming | Ahmad Tarraf