Introduction to Arm Telemetry Solution & Performance Analysis Methodology

Linaro Connect
LONDON 2023

Jumana Mundichipparakkal, <jumana.mp@arm.com>
Principal System Architect at Arm
April 2023

Arm Telemetry Solution

"Arm Telemetry solution aims to provide a solution framework to enable runtime monitoring (**performance**/power) and optimization of software and hardware configuration on Aarch64 based based platforms"

Agenda

- Introduction to Arm Telemetry Solution
- Arm Performance Analysis Methodology
- Telemetry Tools
  - Topdown Tool
- Telemetry Collaboration

Arm Telemetry Solution for Performance Analysis

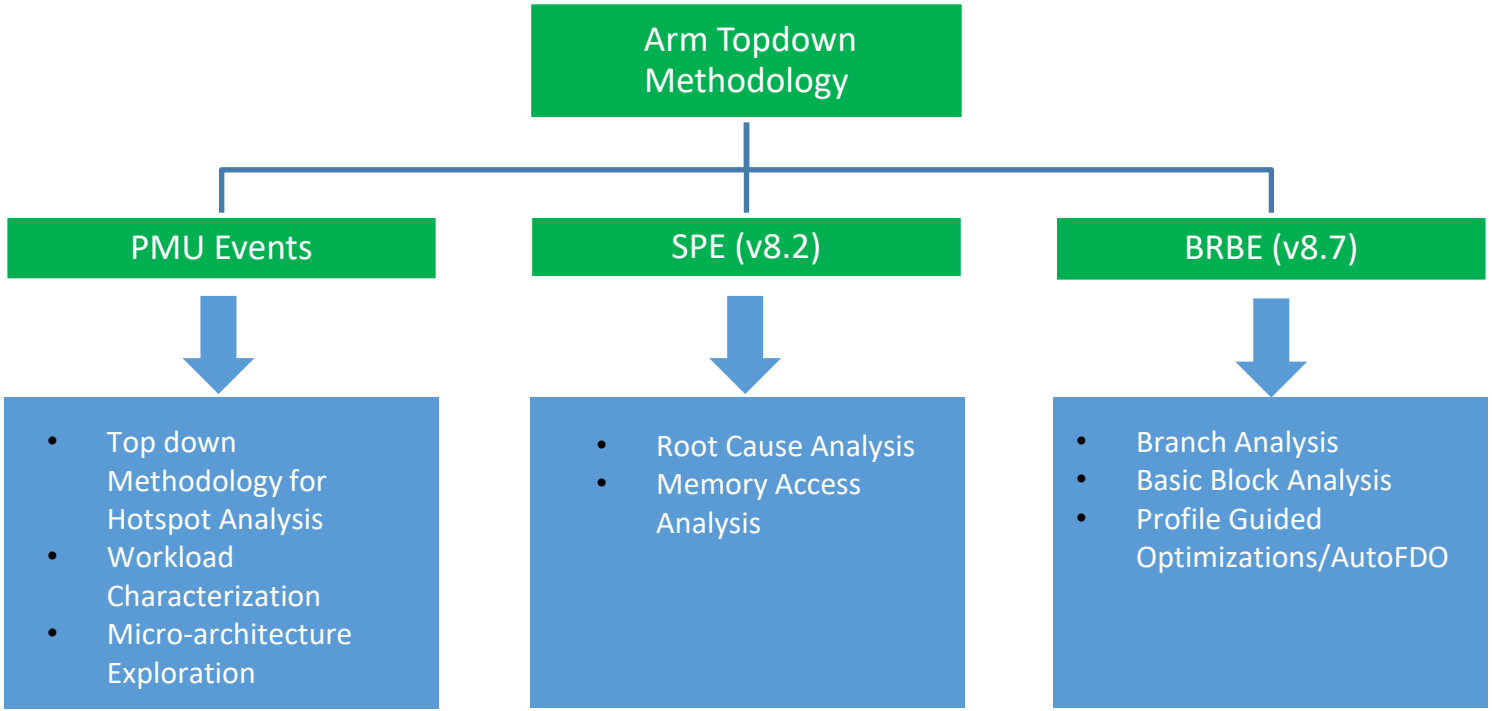

PMU Events -> Metrics -> Methodology

Events Cheat Sheet

Topdown Level 1	Cycle Accounting	Misses Per Kilo Instructions	
• BR_MIS_PRED (r10) • CPU_CYCLES (r11) • OP_RETIRED (r3a) • OP_SPEC (r3b) • STALL_SLOT (r3f) • STALL_SLOT_BACKEND (r3d) • STALL_SLOT_FRONTEND (r3e)	• CPU_CYCLES (r11) • STALL_BACKEND (r24) • STALL_FRONTEND (r23)	• BR_MIS_PRED_RETIRED (r22) • DTLB_WALK (r34) • INST_RETIRED (r08) • ITLB_WALK (r35) • L1D_CACHE_REFILL (r03) • L1D_TLB_REFILL (r05)	• L1I_CACHE_REFILL (r01) • L1I_TLB_REFILL (r02) • L2D_CACHE_REFILL (r17) • L2D_TLB_REFILL (r2d) • LL_CACHE_MISS_RD (r37)
Branch Effectiveness	**Instruction TLB Effectiveness**	**Miss Ratio**	
• BR_MIS_PRED_RETIRED (r22) • BR_RETIRED (r21) • INST_RETIRED (r08)	• INST_RETIRED (r08) • ITLB_WALK (r35) • L1I_TLB (r26) • L1I_TLB_REFILL (r02) • L2D_TLB (r2f) • L2D_TLB_REFILL (r2d)	• BR_MIS_PRED_RETIRED (r22) • BR_RETIRED (r21) • DTLB_WALK (r34) • ITLB_WALK (r35) • L1D_CACHE (r04) • L1D_CACHE_REFILL (r03) • L1D_TLB (r25) • L1D_TLB_REFILL (r05) • L1I_CACHE (r14)	• L1I_CACHE_REFILL (r01) • L1I_TLB (r26) • L1I_TLB_REFILL (r02) • L2D_CACHE (r16) • L2D_CACHE_REFILL (r17) • L2D_TLB (r2f) • L2D_TLB_REFILL (r2d) • LL_CACHE_MISS_RD (r37) • LL_CACHE_RD (r36)
Data TLB Effectiveness	**L1 Instruction Cache Effectiveness**	**L1 Data Cache Effectiveness**	
• DTLB_WALK (r34) • INST_RETIRED (r08) • L1D_TLB (r25) • L1D_TLB_REFILL (r05) • L2D_TLB (r2f) • L2D_TLB_REFILL (r2d)	• INST_RETIRED (r08) • L1I_CACHE (r14) • L1I_CACHE_REFILL (r01)	• INST_RETIRED (r08) • L1D_CACHE (r04) • L1D_CACHE_REFILL (r03)	
L2 Unified Cache Effectiveness	**Last Level Cache Effectiveness**	**Speculation Operation Mix**	
• INST_RETIRED (r08) • L2D_CACHE (r16) • L2D_CACHE_REFILL (r17)	• INST_RETIRED (r08) • LL_CACHE_MISS_RD (r37) • LL_CACHE_RD (r36)	• ASE_SPEC (r74) • BR_IMMED_SPEC (r78) • BR_INDIRECT_SPEC (r7a) • CRYPTO_SPEC (r77) • DP_SPEC (r73)	• INST_SPEC (r1b) • LD_SPEC (r70) • ST_SPEC (r71) • VFP_SPEC (r75)

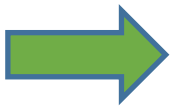

Metric Cheat Sheet

Topdown Level 1	Cycle Accounting	Misses Per Kilo Instructions	
• backend_bound • bad_speculation • frontend_bound • retiring	• backend_stalled_cycles • frontend_stalled_cycles	• branch_mpki • dtlb_mpki • itlb_mpki • l1d_cache_mpki • l1d_tlb_mpki	• l1i_cache_mpki • l1i_tlb_mpki • l2_cache_mpki • l2_tlb_mpki • ll_cache_read_mpki
Branch Effectiveness	**Instruction TLB Effectiveness**	**Miss Ratio**	
• branch_misprediction_ratio • branch_mpki	• itlb_mpki • itlb_walk_ratio • l1i_tlb_miss_ratio • l1i_tlb_mpki • l2_tlb_miss_ratio • l2_tlb_mpki	• branch_misprediction_ratio • dtlb_walk_ratio • itlb_walk_ratio • l1d_cache_miss_ratio • l1d_tlb_miss_ratio	• l1i_cache_miss_ratio • l1i_tlb_miss_ratio • l2_cache_miss_ratio • l2_tlb_miss_ratio • ll_cache_read_miss_ratio
Data TLB Effectiveness	**L1 Instruction Cache Effectiveness**	**L1 Data Cache Effectiveness**	
• dtlb_mpki • dtlb_walk_ratio • l1d_tlb_miss_ratio • l1d_tlb_mpki • l2_tlb_miss_ratio • l2_tlb_mpki	• l1i_cache_miss_ratio • l1i_cache_mpki	• l1d_cache_miss_ratio • l1d_cache_mpki	
L2 Unified Cache Effectiveness	**Last Level Cache Effectiveness**	**Speculation Operation Mix**	
• l2_cache_miss_ratio • l2_cache_mpki	• ll_cache_read_hit_ratio • ll_cache_read_miss_ratio • ll_cache_read_mpki	• branch_percentage • crypto_percentage • integer_dp_percentage • load_percentage	• scalar_fp_percentage • simd_percentage • store_percentage

Hardware Centric – **Micro-architecture oriented**

Software Centric – **HW details abstracted away**

Workload profiling enabled by Topdown Methodology

- Reference: [Neoverse V1 Methodology Whitepaper]

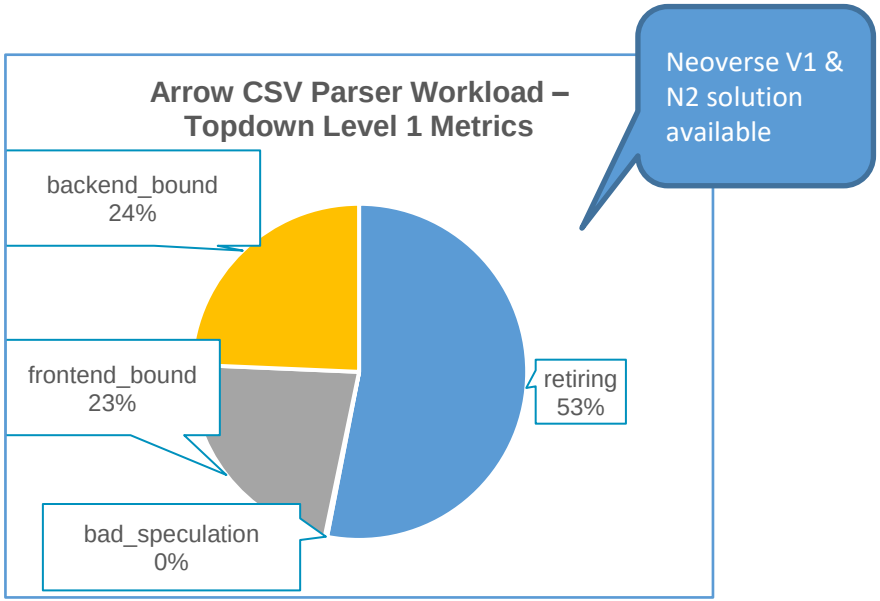

- Data Center Profiling

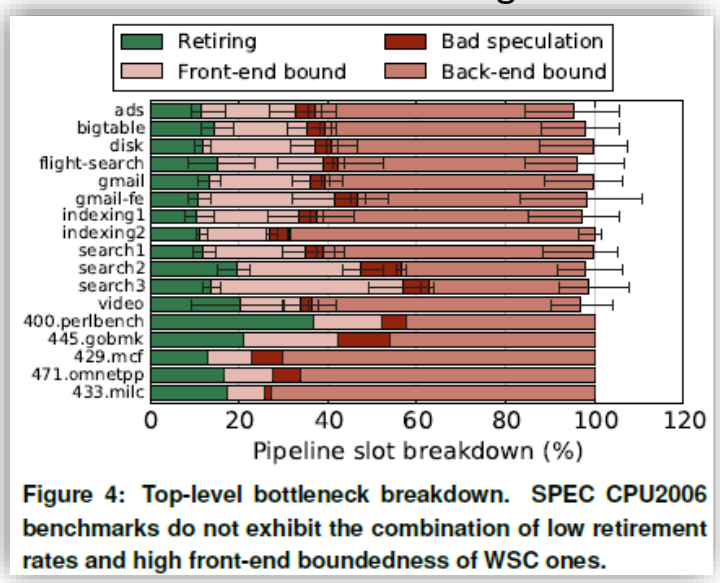

Figure 4: Top-level bottleneck breakdown. SPEC CPU2006 benchmarks do not exhibit the combination of low retirement rates and high front-end boundedness of WSC ones.

Profiling a Warehouse-Scale Computer -S. Kanev, J. P. Darago, K. Hazelwood, P. Ranganathan, T. Moseley, G. Wei and D. Brooks, in International Symposium on Computer Architecture (ISCA), June 2015.

Arm Topdown Methodology Specification

Arm Top-down Analysis Methodology specification for how to conduct performance analysis and processor bottleneck identification with supporting derived metrics and PMU events.

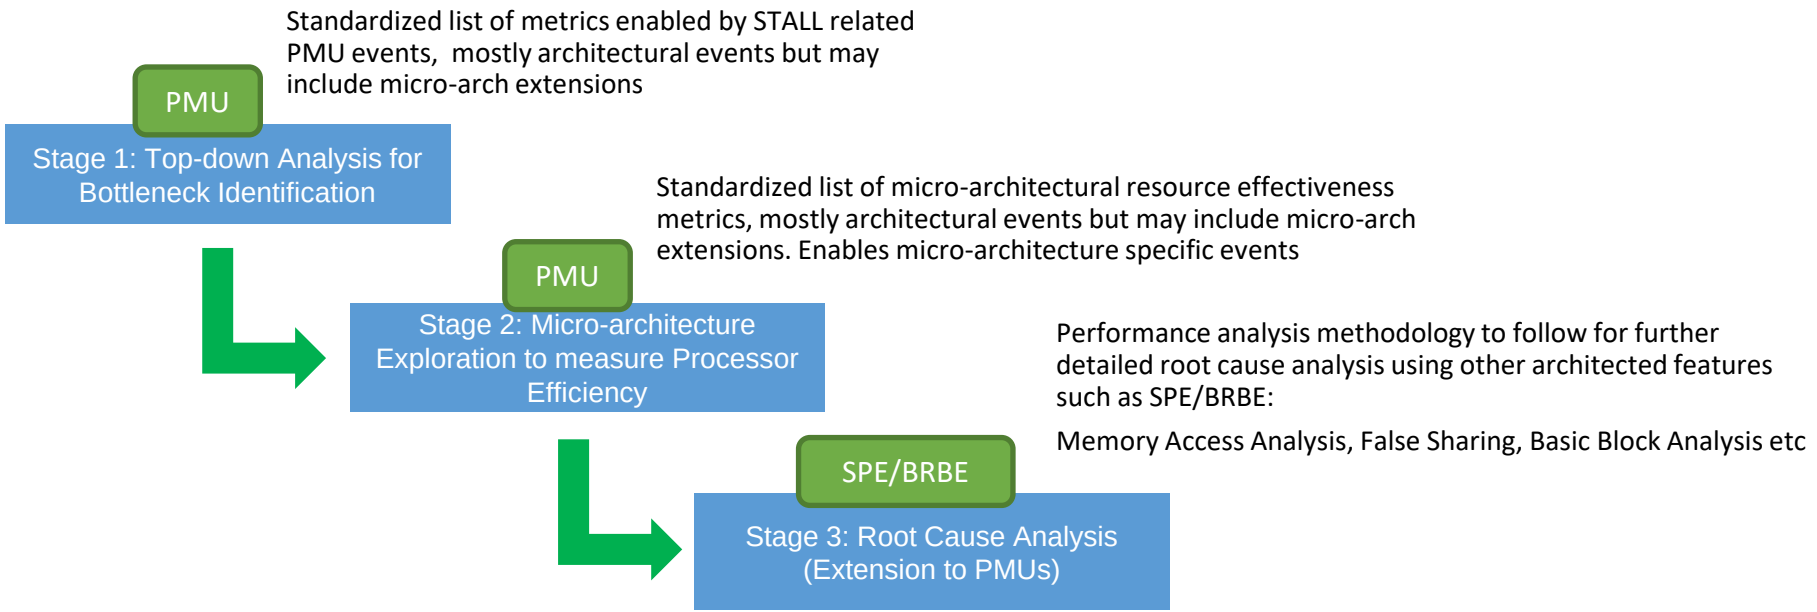

Arm Telemetry Solution: First Release

Arm Neoverse V1 – Top-down Methodology for Performance Analysis & Telemetry Specification

Release Blog with all references

1. Telemetry Specification (Events, Metrics, Methodology)

2. Tooling Collateral (Standardized JSON schema)

3. MVP Tools (Arm topdown tool, SPE analysis tool)

4. Telemetry Methodology & Metrics Validation Suite

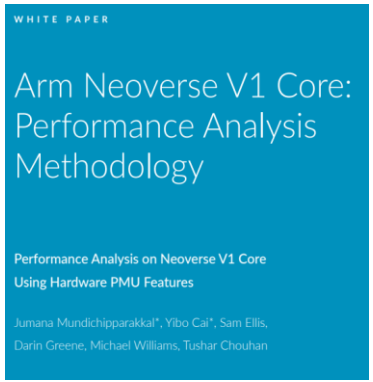

Neoverse V1 Methodology Whitepaper

Telemetry Solution

Project ID: 349

https://gitlab.arm.com/telemetry-solution/telemetry-solution

Topdown Methodology in Action : UStress

4. Telemetry Methodology & Metrics Validation Suite

Reference: [Neoverse V1 Methodology Whitepaper]

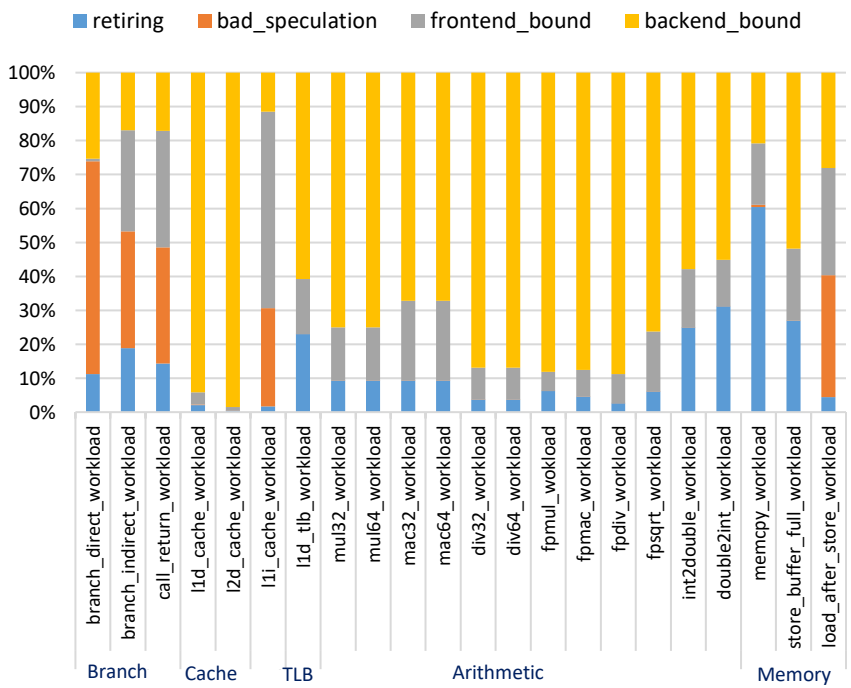

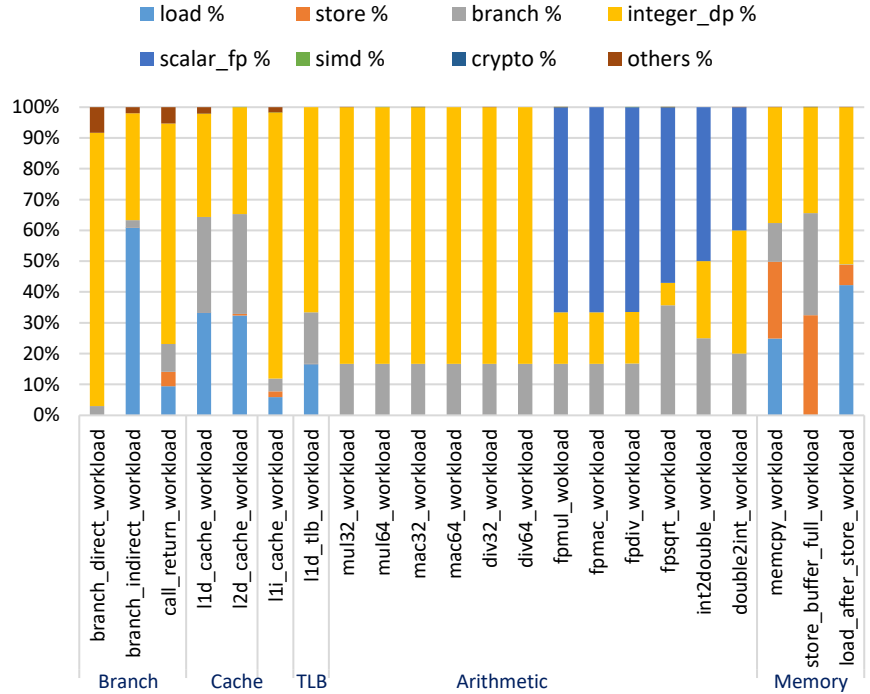

Topdown Methodology & Topdown Tool

3. MVP Tools
(Arm topdown tool)

- **JSON data** describes a set of *metrics* and *event groups* that form a top-down hierarchy (next slide)

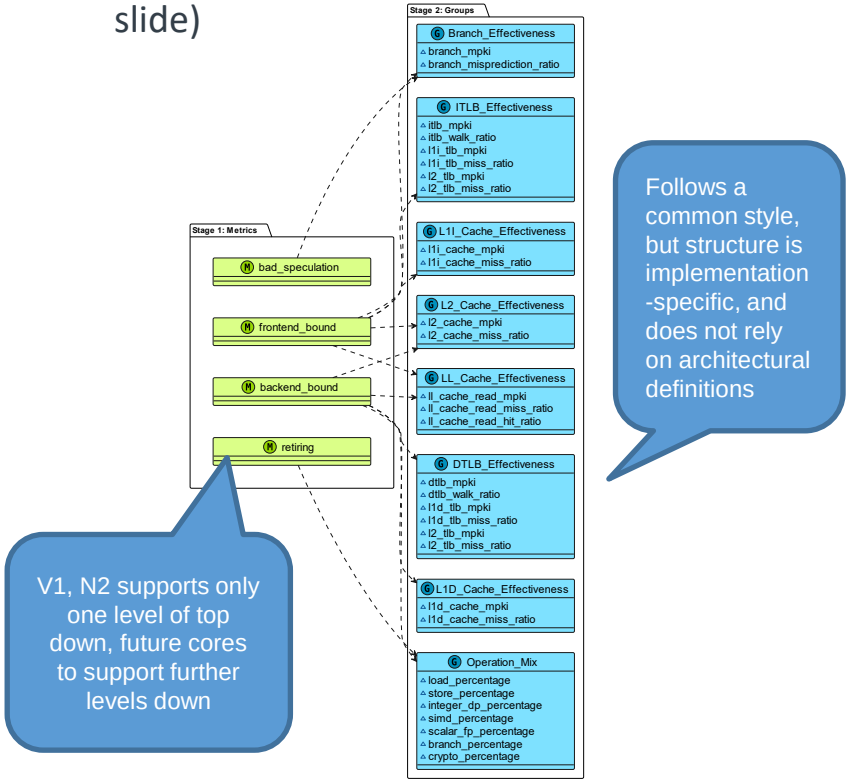

Follows a common style, but structure is implementation -specific, and does not rely on architectural definitions

V1, N2 supports only one level of top down, future cores to support further levels down

- **Software** that uses the machine-readable description to guide the user through the top-down workflow and identify performance issues

```
$ topdown-tool ./a.out
Stage 1 (Topdown metrics)
=========================
[Cycle Accounting]
Frontend Stalled Cycles......... 0.02% cycles
Backend Stalled Cycles.......... 42.59% cycles

Stage 2 (uarch metrics)
=======================
[Branch Effectiveness]
          (follows Frontend Stalled Cycles)
Branch Misprediction Ratio...... 0.001 per branch
Branch MPKI..................... 0.372 misses per 1,000 instructions
Data TLB Effectiveness]
          (follows Backend Stalled Cycles)
DTLB MPKI....................... 0.000 misses per 1,000 instructions
DTLB Walk Ratio................. 0.000 per TLB access
L1 Data TLB Miss Ratio.......... 0.000 per TLB access
L1 Data TLB MPKI................ 0.002 misses per 1,000 instructions
L2 Unified TLB Miss Ratio....... 0.000 per TLB access
L2 Unified TLB MPKI............. 0.006 misses per 1,000 instructions
...
```

- Arm Telemetry Solution GitLab Repository

Arm Telemetry Framework: Building Blocks

2. Tooling Collateral
(Standardized JSON schema)

```
"INST_RETIRED": {
    "code": "0x0008",
    "title": "Instruction architecturally executed",
    "description": "Counts instructions that have been architecturally executed.",
    "common": true,
    "accesses": [
        "PMU",
        "ETE"
    ],
    "architectural": true,
    "impdef": false
},
```

EVENT

```
"L1D_Cache": {
    "title": "L1D CACHE",
    "description": "L1 data cache related events",
    "events": [
        "L1D_CACHE",
        "L1D_CACHE_INVAL",
        "L1D_CACHE_LMISS_RD",
        "L1D_CACHE_RD",
        "L1D_CACHE_REFILL",
        "L1D_CACHE_REFILL_INNER",
        "L1D_CACHE_REFILL_OUTER",
        "L1D_CACHE_REFILL_RD",
        "L1D_CACHE_REFILL_WR",
        "L1D_CACHE_WB",
        "L1D_CACHE_WB_CLEAN",
        "L1D_CACHE_WB_VICTIM",
        "L1D_CACHE_WR"
    ]
},
```

EVENT GROUP

```
"frontend_bound": {
    "title": "Frontend Bound",
    "formula": "(100 * ((STALL_SLOT_FRONTEND / (CPU_CYCLES * 8)) - ((BR_MIS_PRED * 4) / CPU_CYCLES)))",
    "description": "This metric is the percentage of total slots that were stalled due to resource constraints in the frontend of the processor.",
    "units": "percent of slots",
    "events": [
        "BR_MIS_PRED",
        "CPU_CYCLES",
        "STALL_SLOT_FRONTEND"
    ]
},
```

METRIC

```
"Operation_Mix": {
    "title": "Speculative Operation Mix",
    "description": "This metric group provides the distribution of micro-operation types executed for the program.",
    "metrics": [
        "load_percentage",
        "store_percentage",
        "integer_dp_percentage",
        "simd_percentage",
        "scalar_fp_percentage",
        "branch_percentage",
        "crypto_percentage"
    ]
}
```

METRIC GROUP

```
"topdown_methodology": {
    "title": "Topdown Methodology",
    "description": "Topdown Performance Analysis Methodology",
    "metric_grouping": {
        "stage_1": [
            "Topdown_L1"
        ],
        "stage_2": [
            "Cycle_Accounting",
            "General",
            "MPKI",
            "Miss_Ratio",
            "Branch_Effectiveness",
            "ITLB_Effectiveness",
            "DTLB_Effectiveness",
            "L1I_Cache_Effectiveness",
            "L1D_Cache_Effectiveness",
            "L2_Cache_Effectiveness",
            "LL_Cache_Effectiveness",
            "Operation_Mix"
        ]
    },
    "decision_tree": {
        "root_nodes": [
            "frontend_bound",
            "backend_bound",
            "retiring",
            "bad_speculation"
        ],
        "metrics": [
            {
                "name": "frontend_bound",
                "group": "Topdown_L1",
                "next_items": [
                    "Branch_Effectiveness",
                    "ITLB_Effectiveness",
                    "L1I_Cache_Effectiveness",
```

METHODOLOGY

Summary

- Arm Telemetry Solution currently supports Neoverse cores
  - Neoverse N1, V1 and N2 Telemetry JSON released
  - Neoverse V1 performance analysis methodology whitepaper released
    - V1 Methodology applies to N1, V1 and N2
  - Topdown tool released for developers to conduct performance analysis on Linux/Windows machines
    - Linux perf tool and windowsperf supported
- Future releases planned
  - Neoverse V2 telemetry specification (with JSON)
  - SPE Methodology Whitepaper
    - SPE Analysis Tool released

Handy Resources

- GitLab – Telemetry Solution
  - https://gitlab.arm.com/telemetry-solution/telemetry-solution
    - Telemetry Specification: https://gitlab.arm.com/telemetry-solution/telemetry-solution/-/tree/main/data
    - Telemetry Tools:
  - Contact us: telemetry-solution@arm.com
- Arm Neoverse V1 Performance Analysis Methodology Whitepaper
  - https://www.arm.com/-/media/Files/pdf/white-paper/neoverse-v1-core-performance-analysis
  - V1 Methodology applies to Neoverse N2.
    - N2 Telemetry Specification to be released, JSON available
- Arm Neoverse N1 Performance Analysis Methodology Whitepaper
  - https://armkeil.blob.core.windows.net/developer/Files/pdf/white-paper/neoverse-n1-core-performance-v2.pdf\

Telemetry Solution Collaboration

- Arm has developed a standardized telemetry framework currently supporting CPU PMU events, metrics and topdown methodology. The framework is publicly available at Gitlab telemetry-solution.
- We seek support from the ecosystem to enhance the tooling capabilities and welcome feedback. Contact us at telemetry-solution@arm.com
- For Aarch64 standardization, we invite Arm architectural partners to adopt our telemetry framework to specify their product PMUs and top down methodology. Linux perf tool and WindowsPerf enablement for Arm products are leveraging the telemetry framework (JSON schema – slide 11).

Acknowledgements to the team

Michael Williams (Arm)
Darin Greene (Arm)
Ryan Harkin (Arm)
James Clark (Arm)
Jun He (Arm)
Yibo Cai (Arm)
Rui Chang (Arm)
Xin Yang (Arm)
David Murray (Arm)
Alessandro Di Bella (Arm)
Nathan Hedon (Arm)
Tushar Singh Chouhan (Arm)
Nick Forrington (Arm)
Leo Yan (Linaro)

Telemetry Solution
<telemetry-solution@arm.com>

Linaro Connect
LONDON 2023

Thank you

Jumana Mundichipparakkal <jumana.mp@arm.com>

Presentations will be available post event at
Resource Hub

Telemetry Solution
<telemetry-solution@arm.com>