Barcelona Supercomputing Center
Centro Nacional de Supercomputación

EXCELENCIA SEVERO OCHOA

Levels of Detail in Performance Analysis: a RISC-V vector codesign case

Prof. Jesús Labarta

BSC & UPC

Scalable Tools Workshop

Granlibakken, June 19th 2023

A bit on Vision

- Scalability: not about size, about dynamic range !

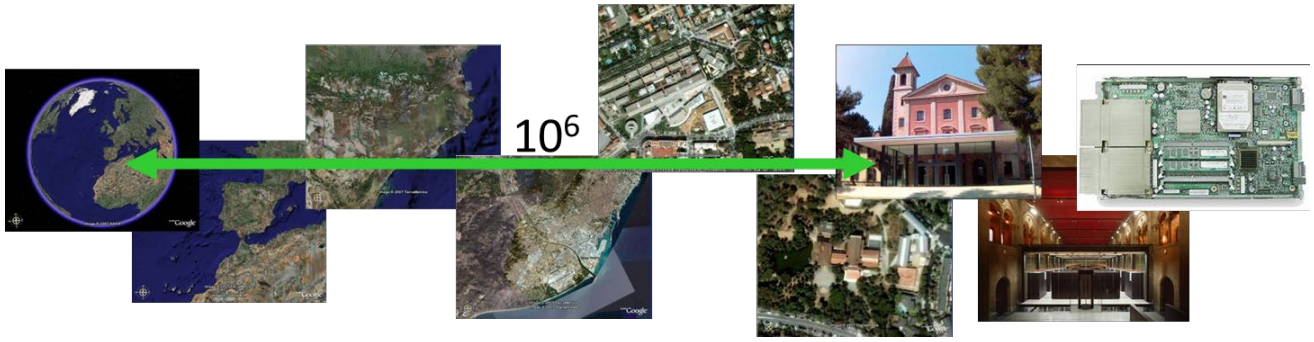

- “Universal” principles: “Similar” at all scales

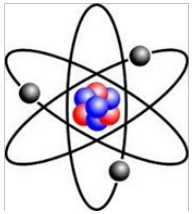

A bit on Vision

- BSC tools focus

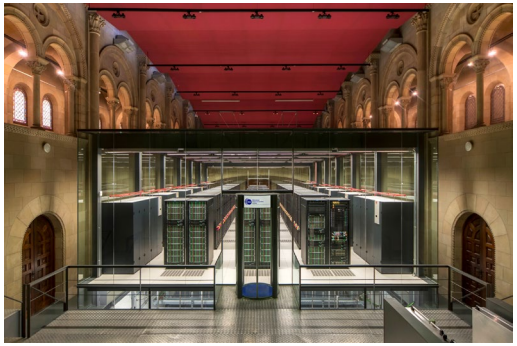

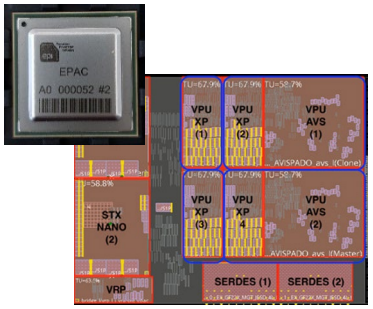

A bit on Vision

- BSC tools focus

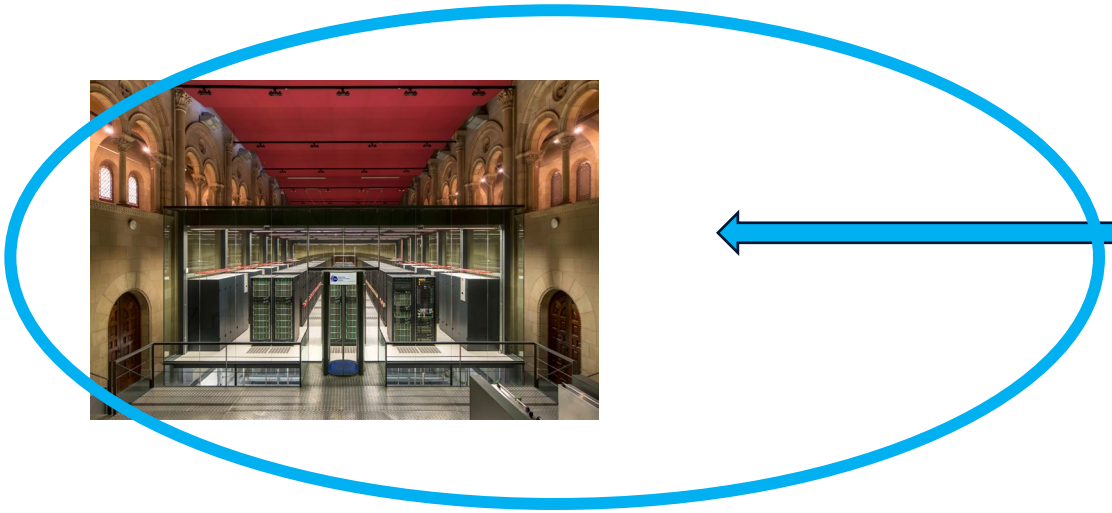

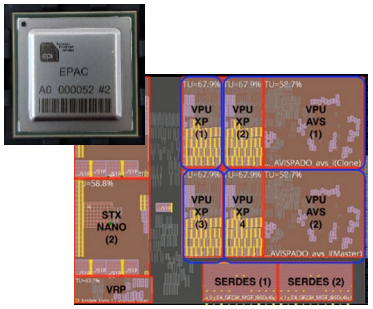

Traditionally

A bit on Vision

- BSC tools focus

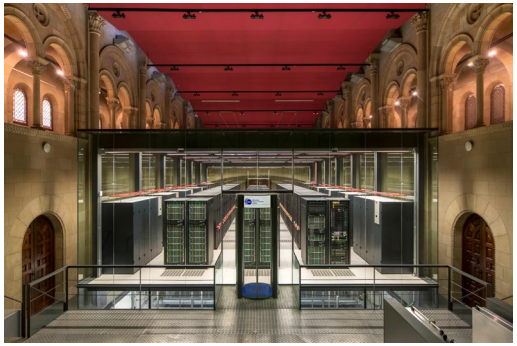

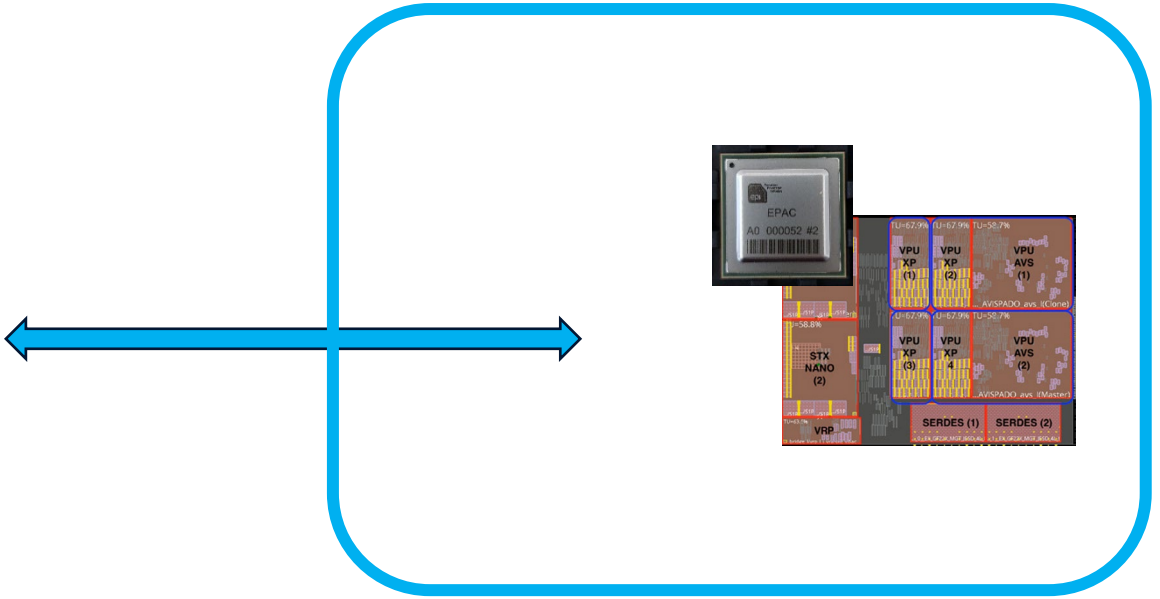

Useful ? Possibilities?

… on RISC-V Vector processor ?

The EPI RISCV Vector project

European Processor Initiative

The EPI RISCV Vector vision

- Long vectors
  - Raise ISA semantic level
    - "less words, more work"
    - The importance of ISA: Osmotic membrane
      - Convey access pattern semantics to the architecture
      - Potential to optimize throughput tolerating latency
  - Parallelism
    - 1 instruction N cycles
      - Less pressure on frontend
      - Possibilities of ~ OoO execution
        - Decouple Front end – back end
        - Maximize lookahead
      - few cycles "overhead"
      - Opportunity for intelligence
  - Programming productivity
    - Single thread intermixing scalar and vector instructions
      - Consistent
      - Fine granularity
    - Hardware score boarding
      - No need of explicit order control
    - Change in relative importance
      - overheads, parallelism & locality

RISC-V

- technical enabler
- affordable

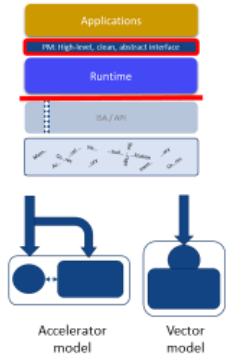

EPAC architecture

RVV
- RV64GCV (→ 8x)
  - 2 way in order core
  - Decoupled VPU
    - 8 lanes
    - Long vectors (256 DP elements)
  - → 128 MSHR
  - L1 - MESI coherency
- CHI interface NoC
  - 1 line / cycle
- L2$: 256KB/module
  - Allocation control mechanisms
- No in tile L3$

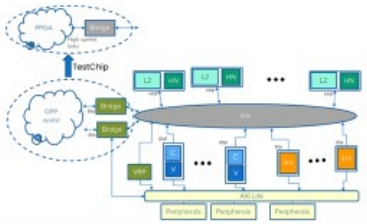

STX
- DL and stencil specific accelerators
- Extensions to planned NTX
  - Programmable address generators →
  - lightweight RISC-v core + fat FPU + (Streaming Semantics & FREP)

VRP
- Variable precision processors

RVV @ FPGA & ecosystem

- HPC software stack @ Commercially available RISC-V platforms
  - SLURM, MPI, OpenMP, BSC tools, RVV software emulator
- EPI SDV platforms
  - Linux
  - Test user codes @ real RTL
  - Early access to EPI technology

Contact : filippo.mantovani@bsc.es

- Holistic CI/CD framework
  - HW & SW
  - Functionality & performance

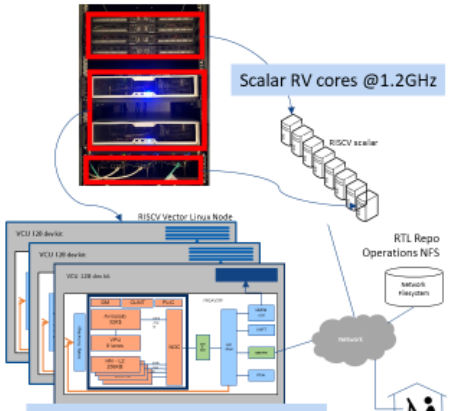

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

LOD for “co-design”

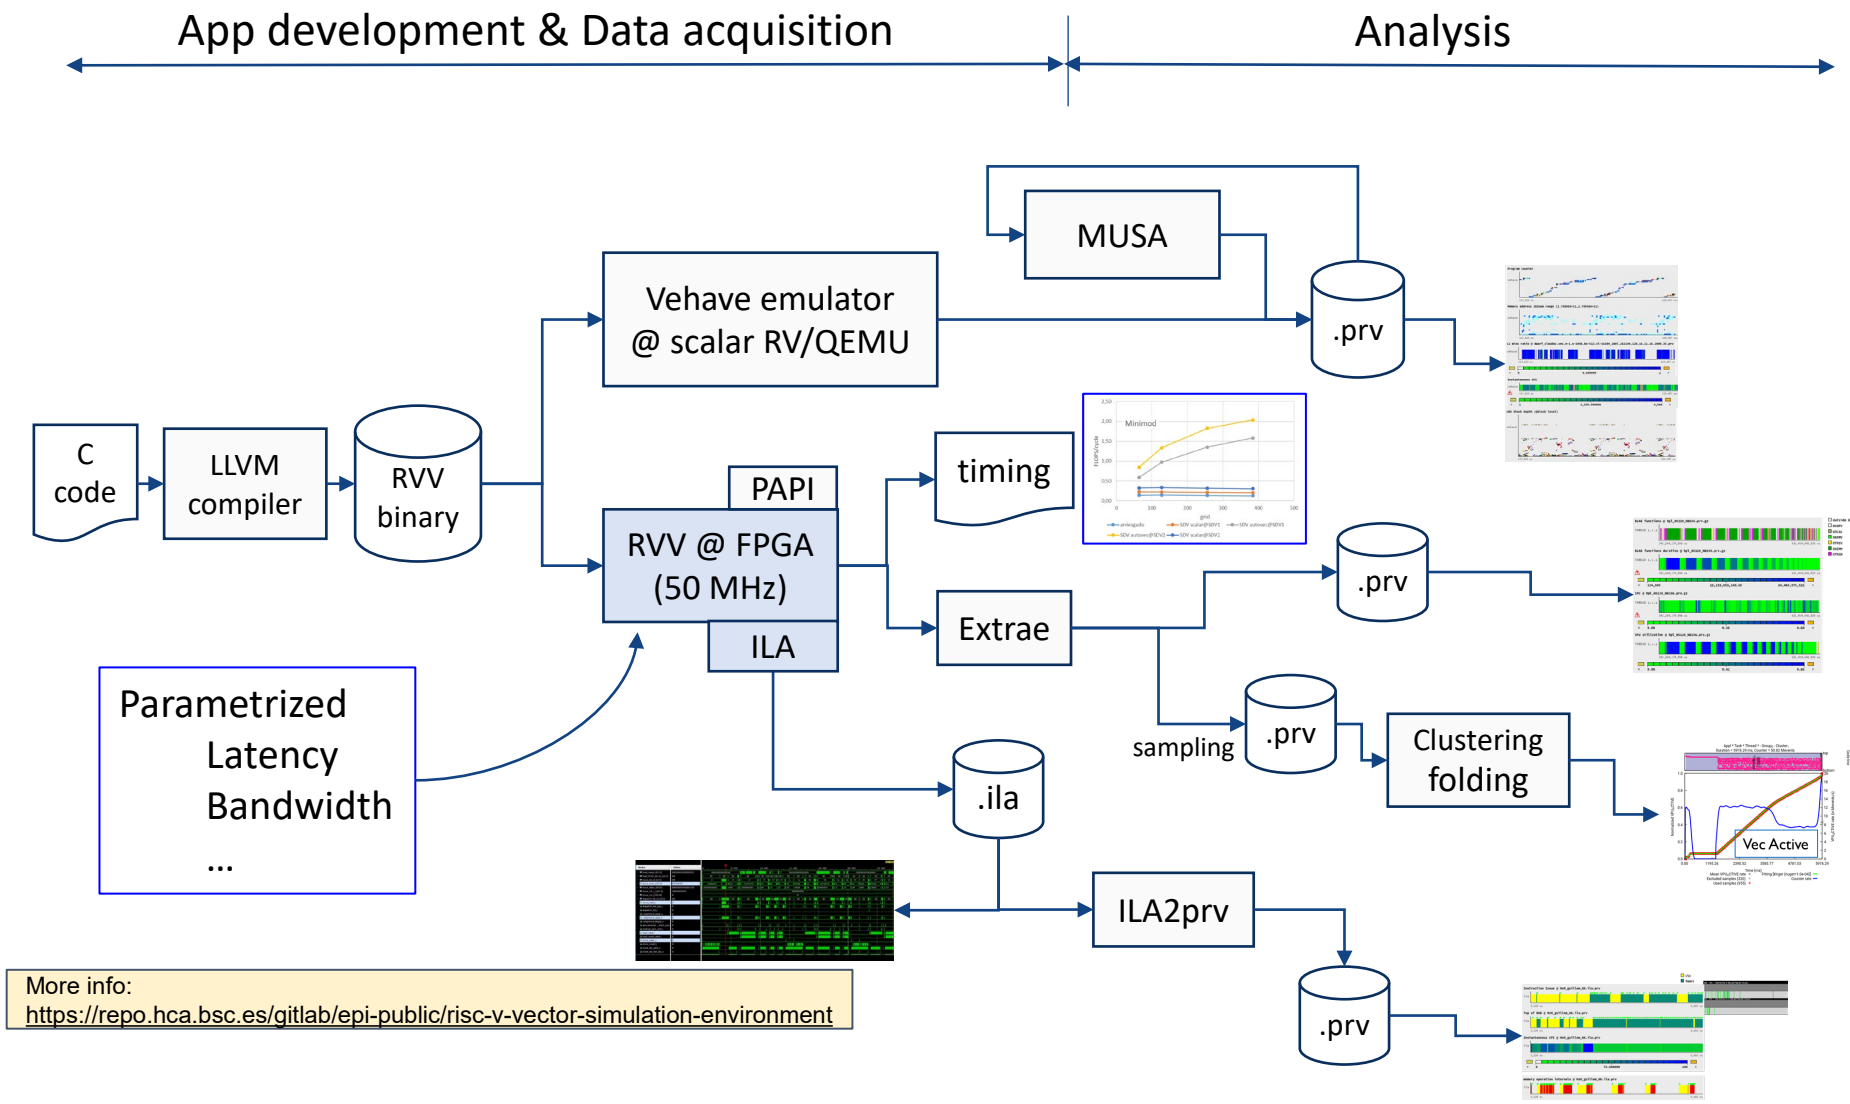

Granularities

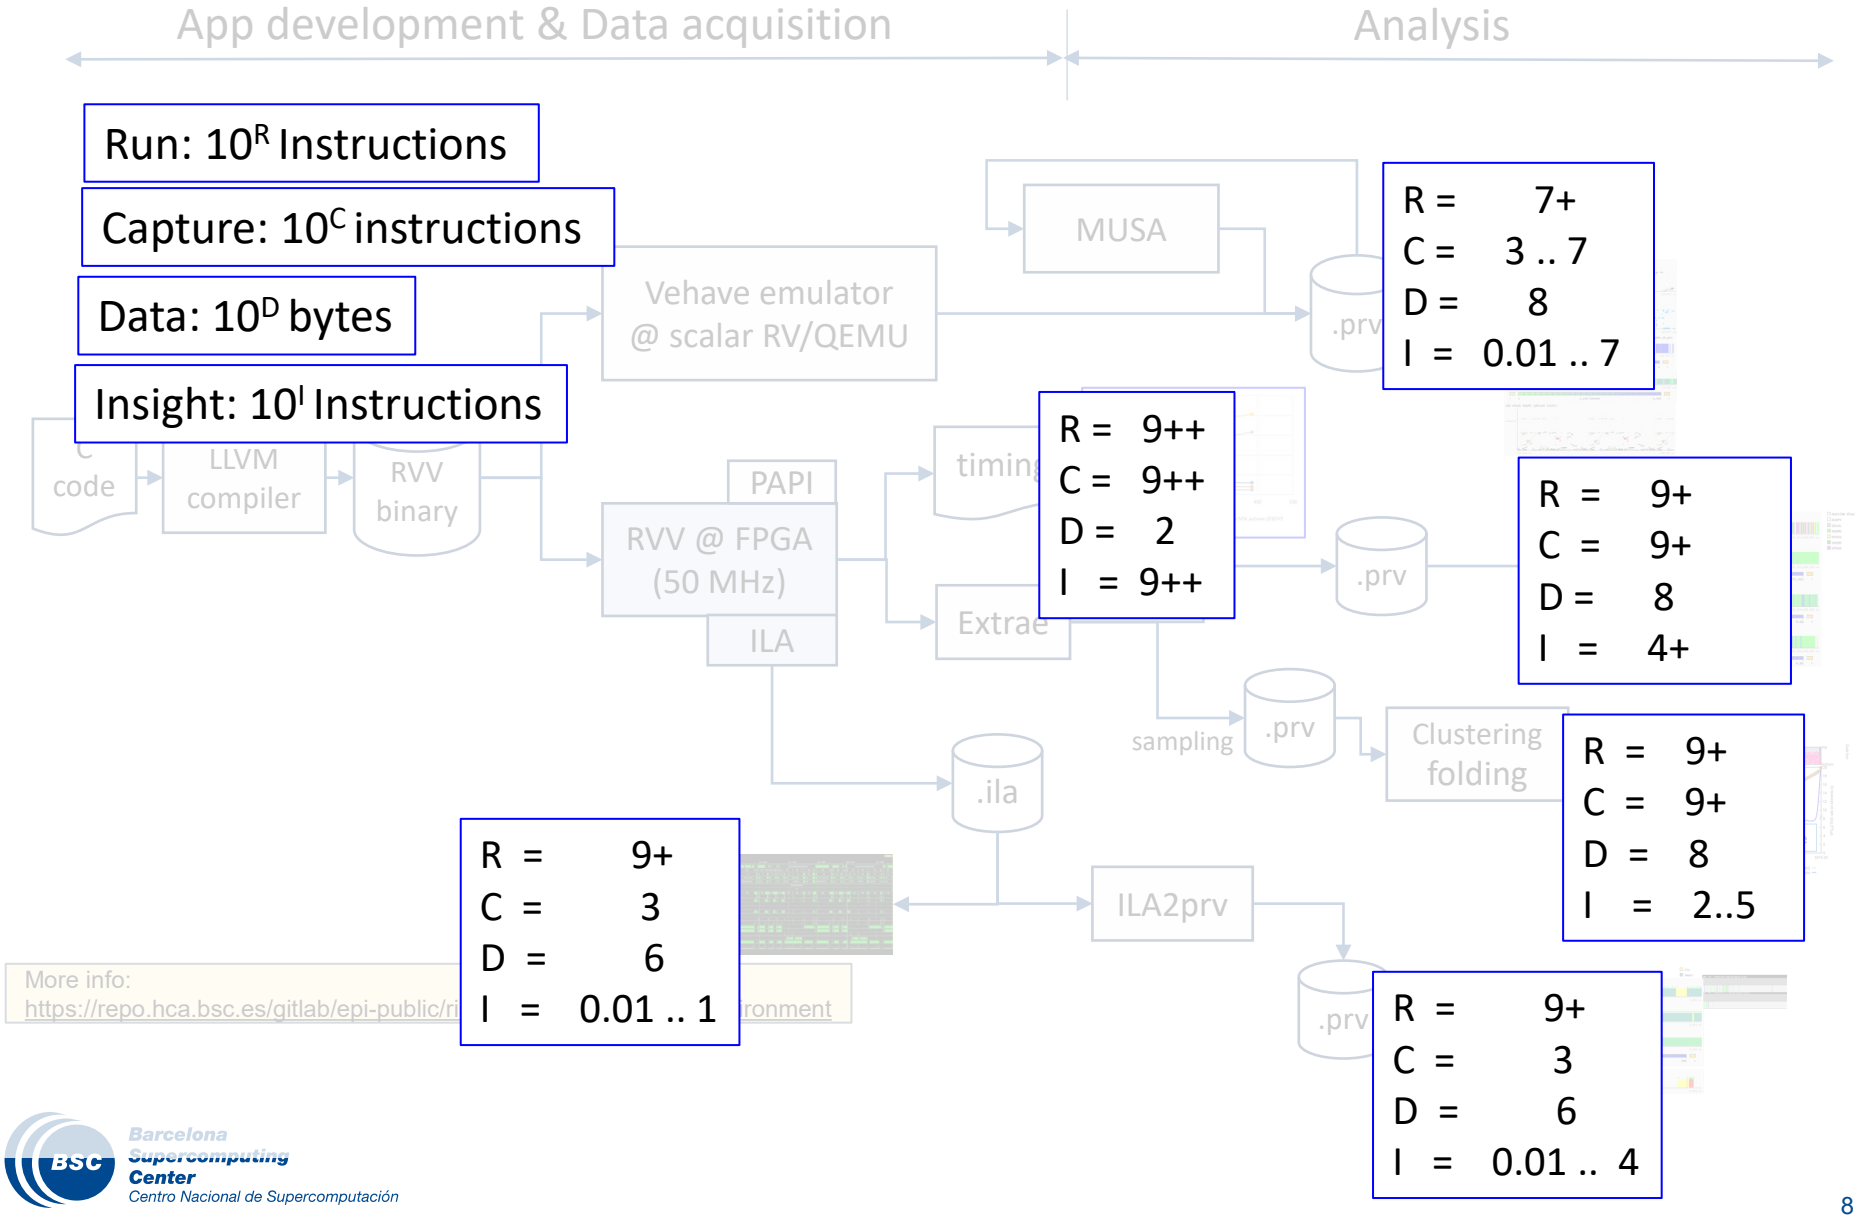

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

EXCELENCIA SEVERO OCHOA

Extrae traces

Prof. Jesús Labarta

BSC & UPC

Scalable Tools Workshop

Granlibakken, June 19th 2023

Extrae traces

- User events
  - Code region events
  - User events M, N, K
    - Derive MFLOPS
  - Sampling or not
  - HW counters
    - Typical: I, cyc, L1, TLB
    - Other: VEC stalls & utilization

HPL: N=4096

% of peak @ dgemm @ Linpack?

With sampling

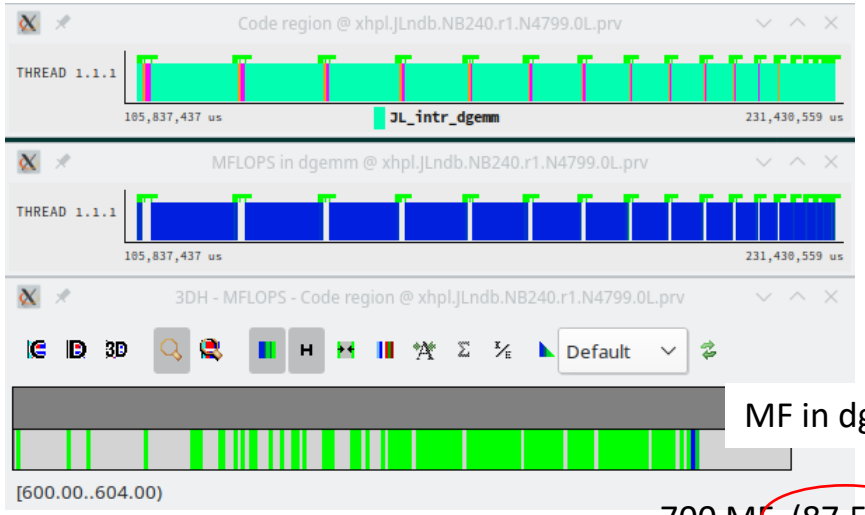

MF in dgemm

700 MF (87.5%)

No sampling

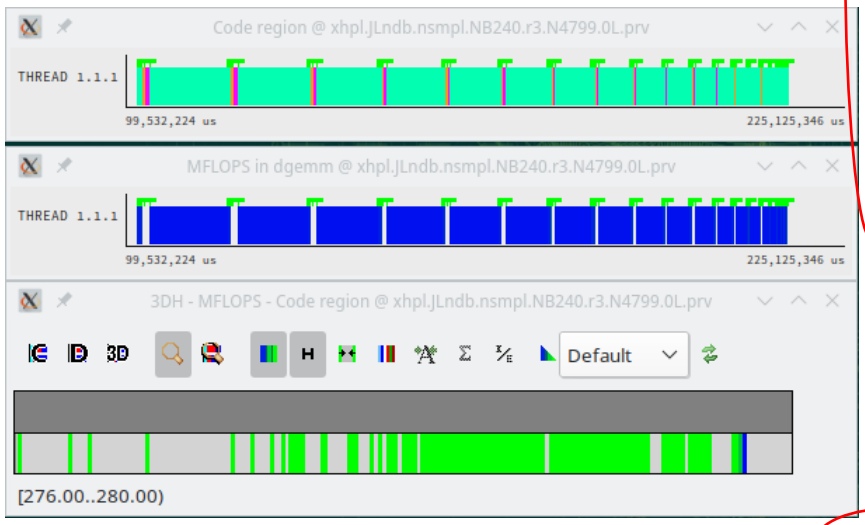

750 MF (93.7%)

Less OS ?
50 MHz emulator
standard Linux interrupts

Extrae HWC metrics

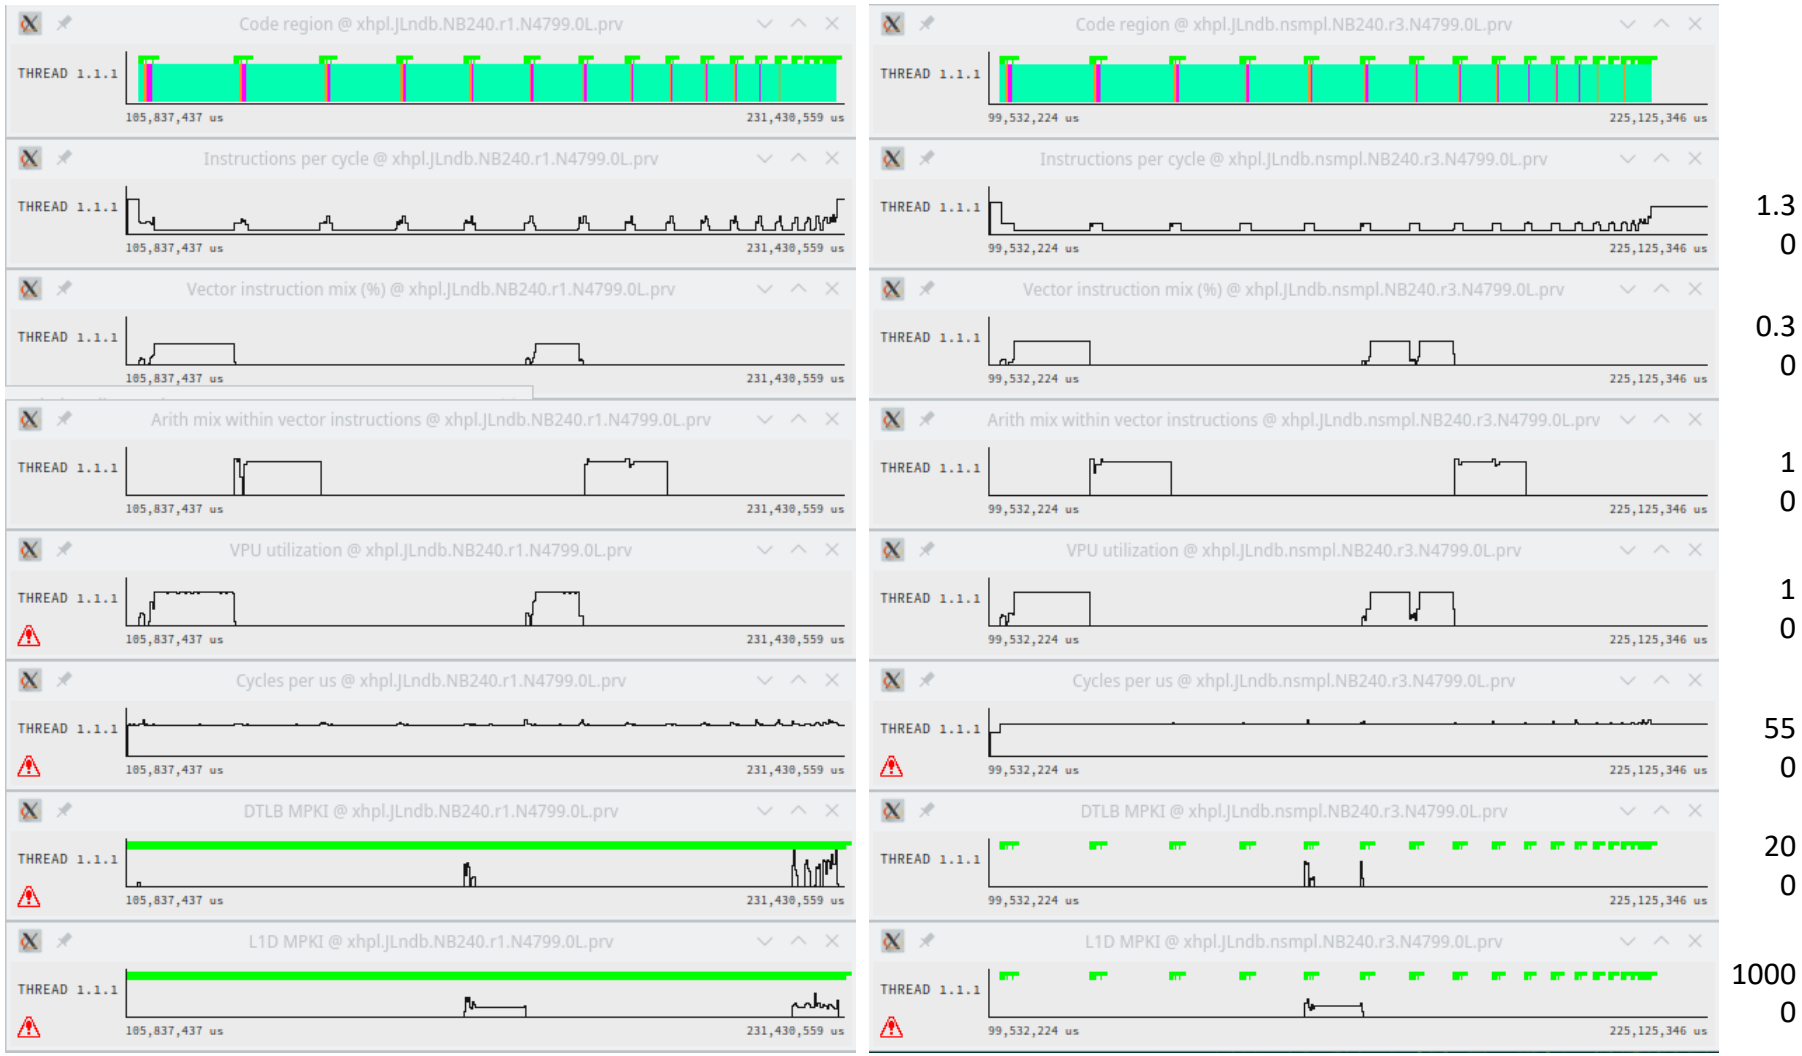

Extrae traces

HPL: N=4096

% of peak @ Linpack?

With sampling

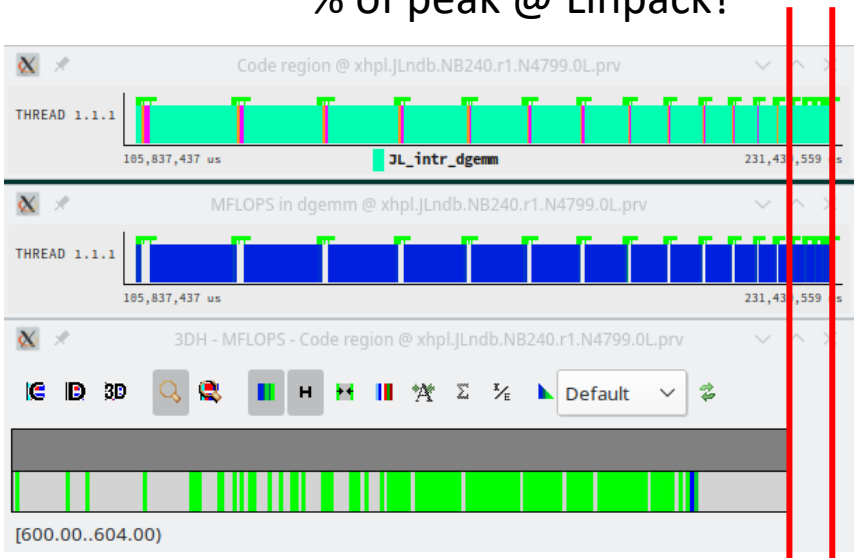

603 → 648

Still only 81%

No sampling

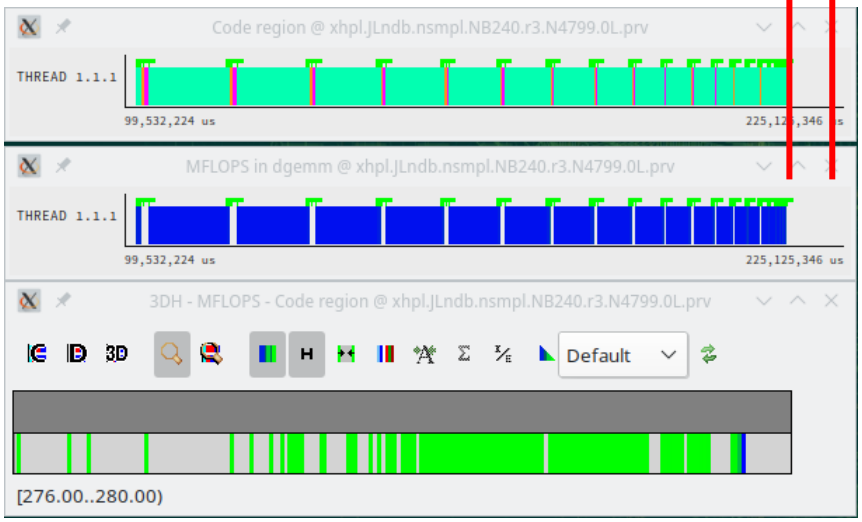

Extrae HWC metrics

With sampling

No sampling

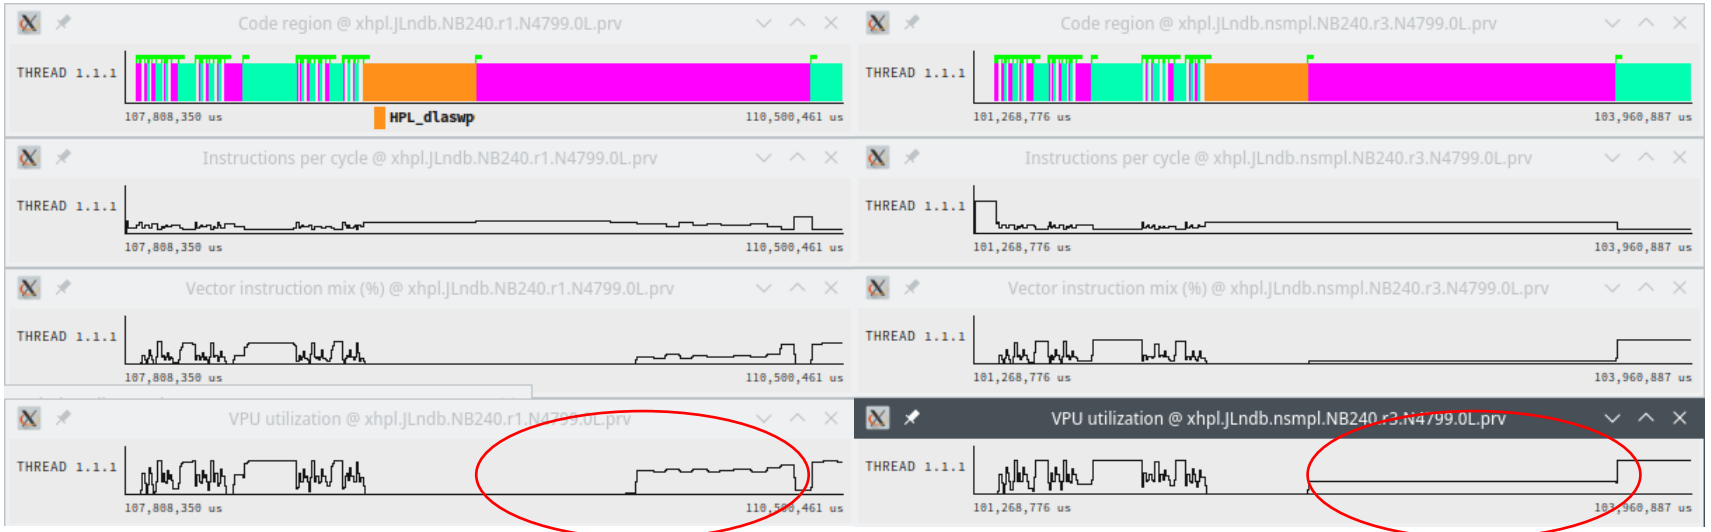

DRTSM: Not all vectorized

intermix scalar and vector instructions

copies to temporary storage

WIP

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

EXCELENCIA SEVERO OCHOA

Sampling & folding

Prof. Jesús Labarta

BSC & UPC

Scalable Tools Workshop

Granlibakken, June 19th 2023

Coarse grain sampling → fine grain detail

- ESiWace Dwarf
  - Climate code miniapp (IFS)
    - Iterative
    - Very flat profile
    - Fine granularity
- Trace with sampling
  - ~70 ms sampled granularity
  - Actual program granularity? Can be identified?

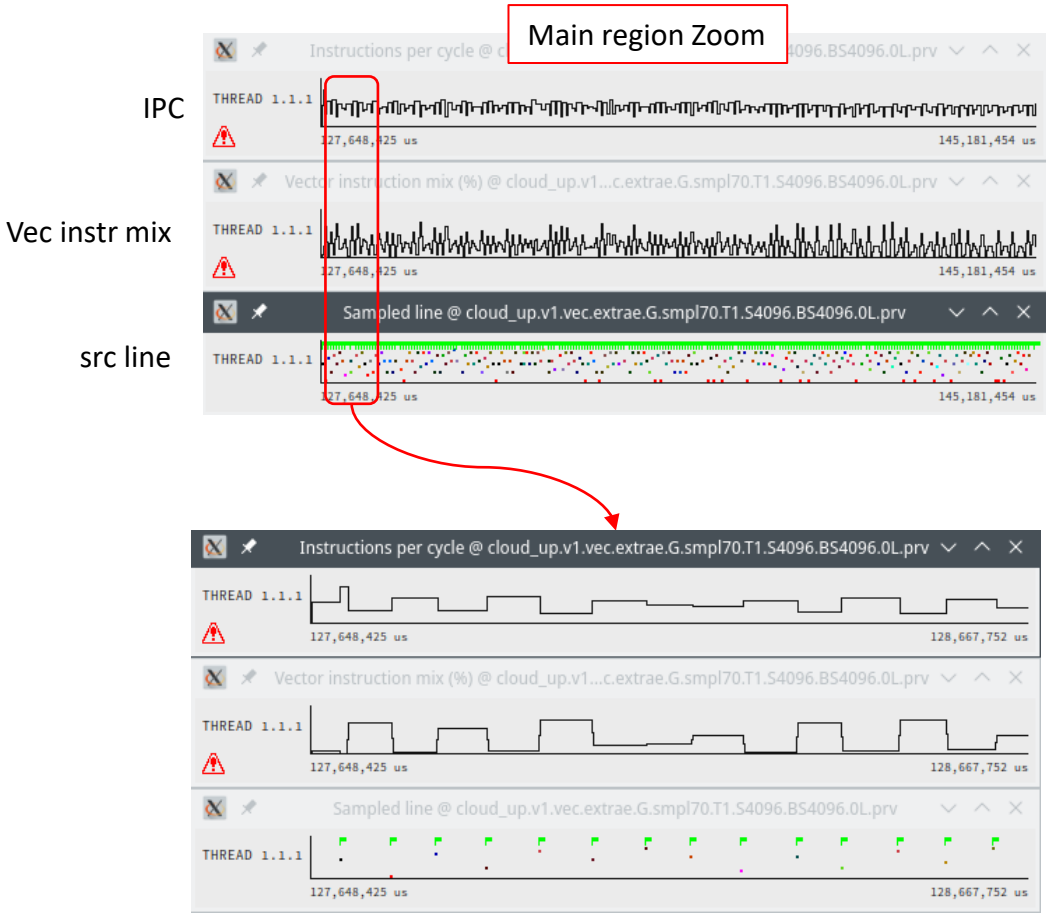

Coarse grain sampling → fine grain detail

- Postprocess trace (folding) to get finer grain details

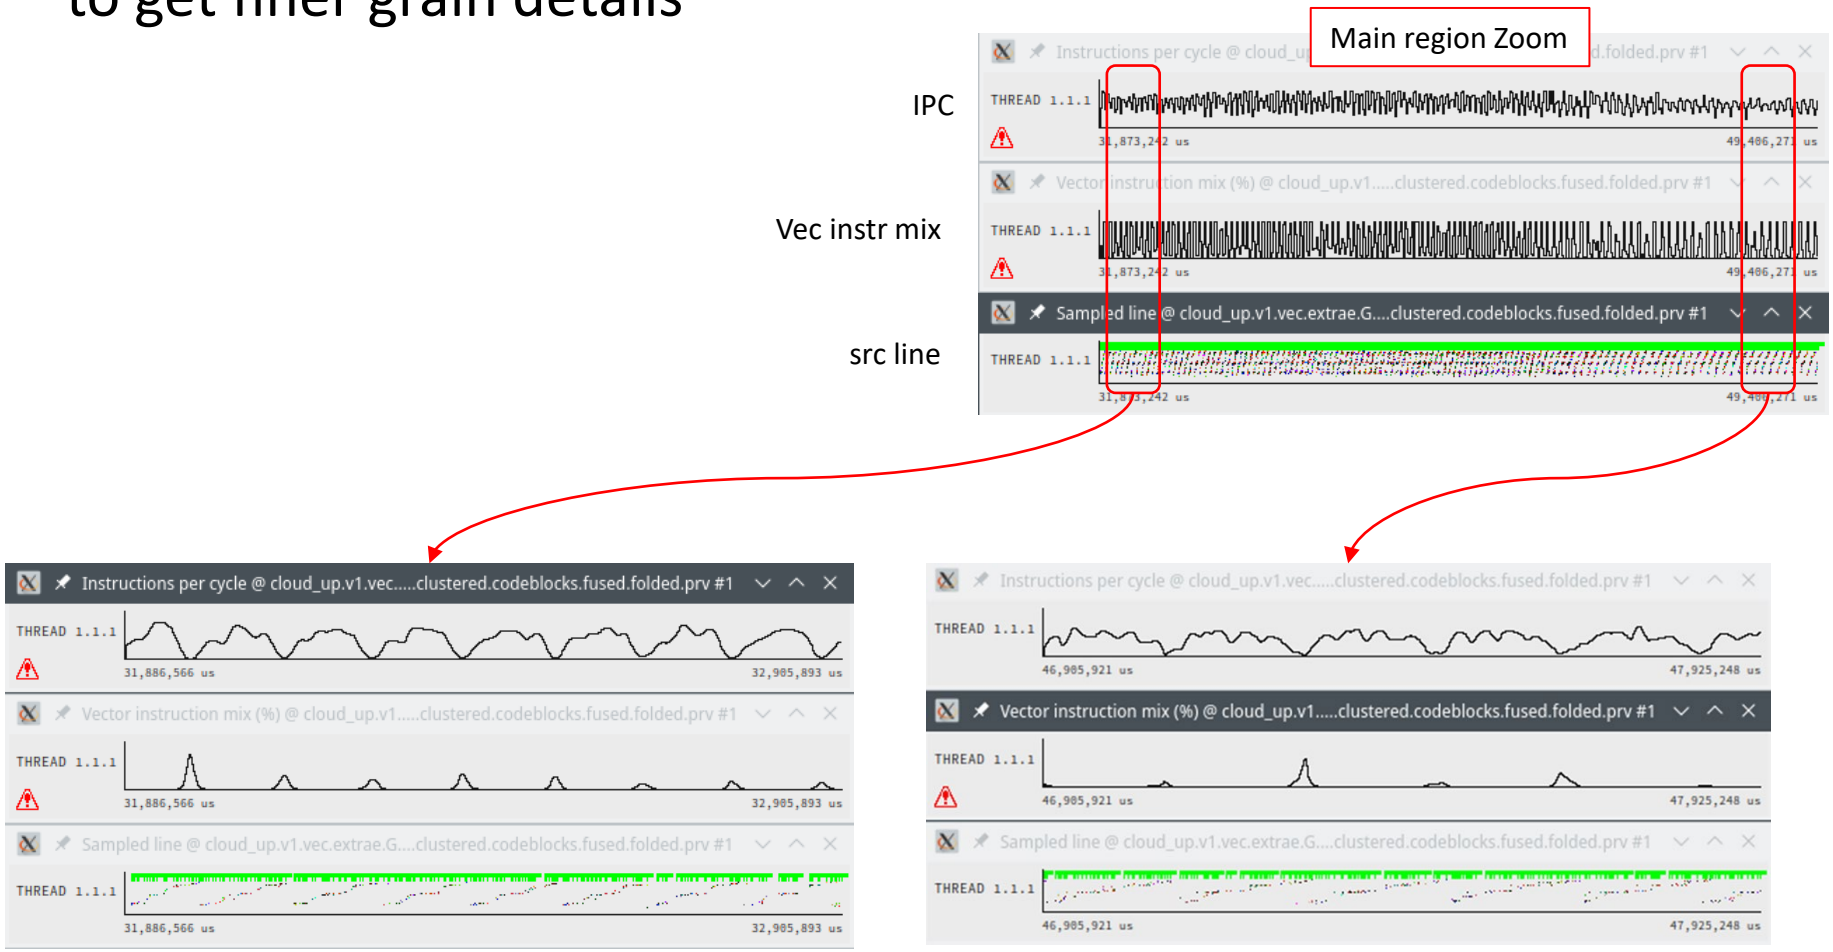

Harald Servat et al. "Detailed performance analysis using coarse grain sampling" PROPER@EUROPAR, 2009

Harald Servat et al. "Unveiling Internal Evolution of Parallel Application Computation Phases" ICPP 2011

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

EXCELENCIA SEVERO OCHOA

Vehave emulation and MUSA simulation

Prof. Jesús Labarta

BSC & UPC

Scalable Tools Workshop

Granlibakken, June 19th 2023

Vehave traces

- EsiWace climate code dwarfs

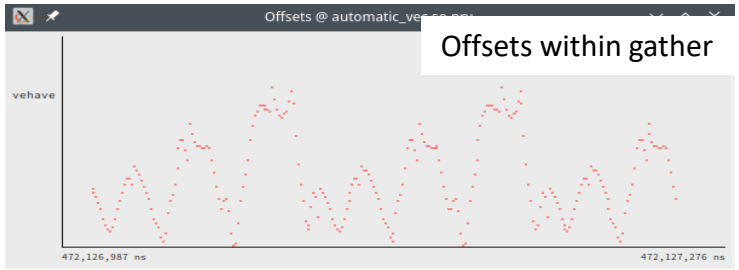

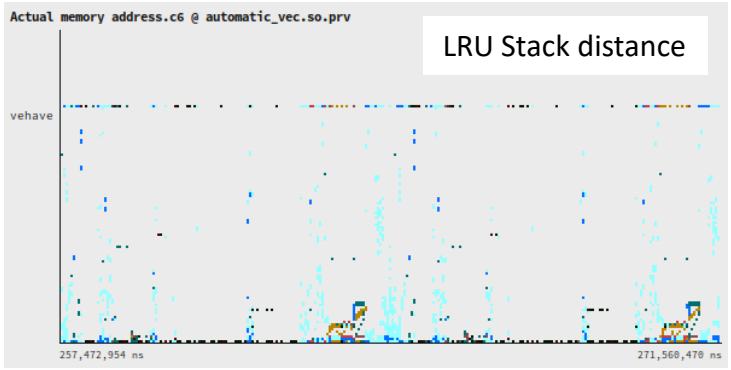

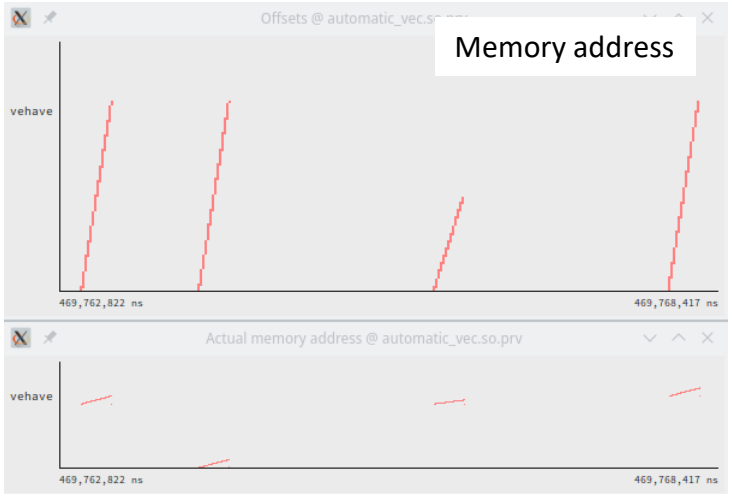

Architectural simulation

- Impact of …
  - Instruction latencies, ROB, Cache sizes, Memory latency, MAXVL, …
- Global metrics & **detail**
  - When are instructions issued, completed, graduated
  - Individual hits & misses

RVV Software emulation

- Detailed analysis & insight

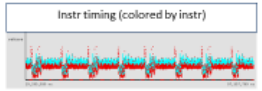

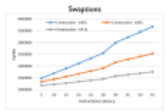

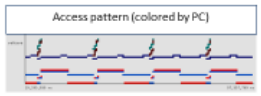

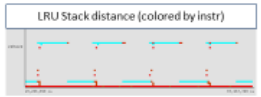

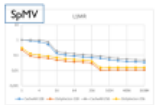

45

RVV Software emulation

- Geophysics

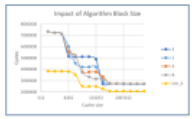

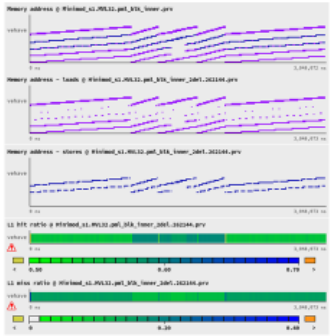

256K cache

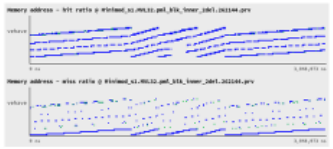

47

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

EXCELENCIA SEVERO OCHOA

ILA

Prof. Jesús Labarta

BSC & UPC

Scalable Tools Workshop

Granlibakken, June 19th 2023

ILA traces

- DGEMM
- Trace: 32K cycles

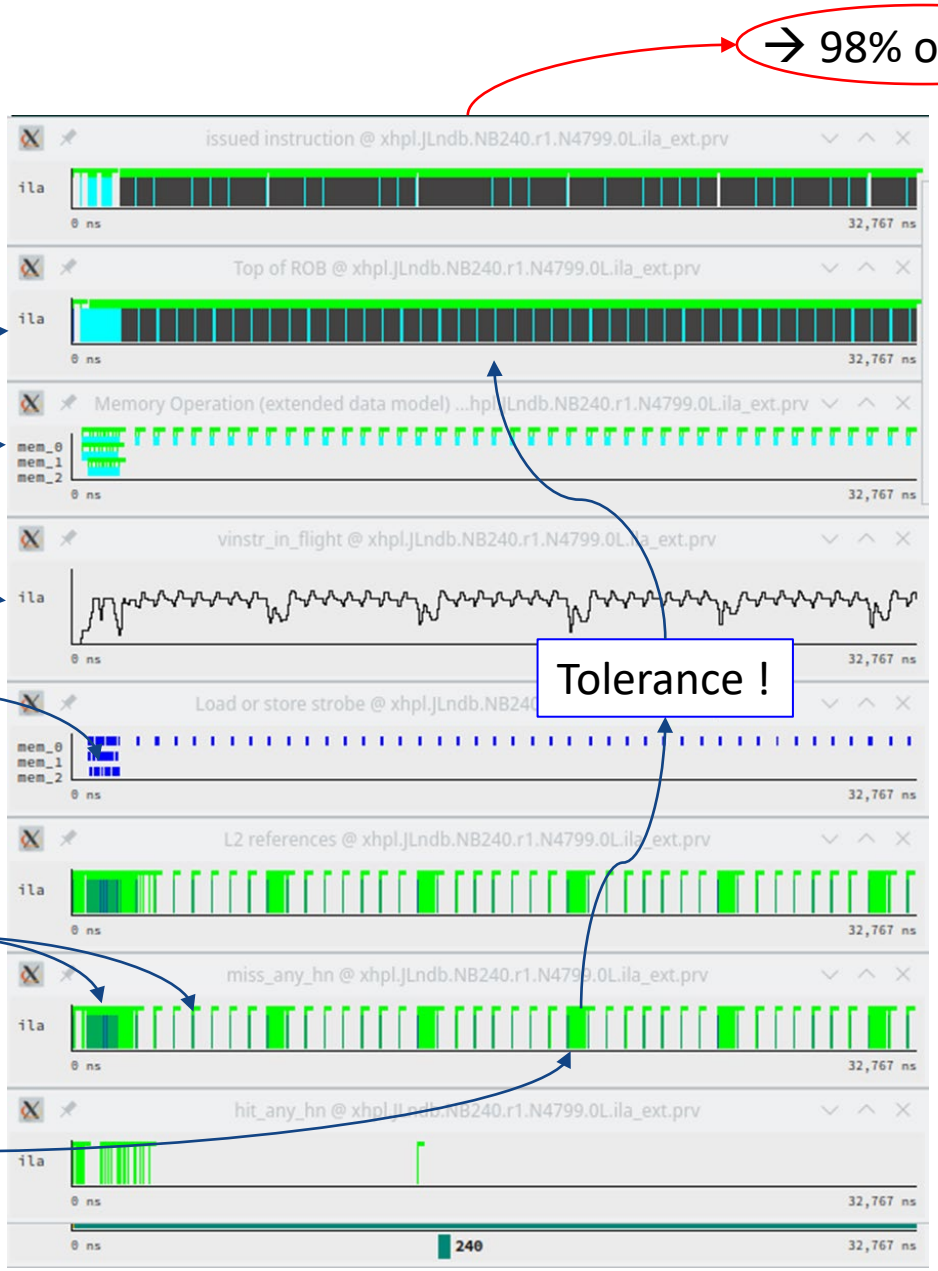

ILA traces

- Zoom

Load not issued till very late …
… still in time to overlap, but ..

Vector loads: always miss

Periodic scalar misses

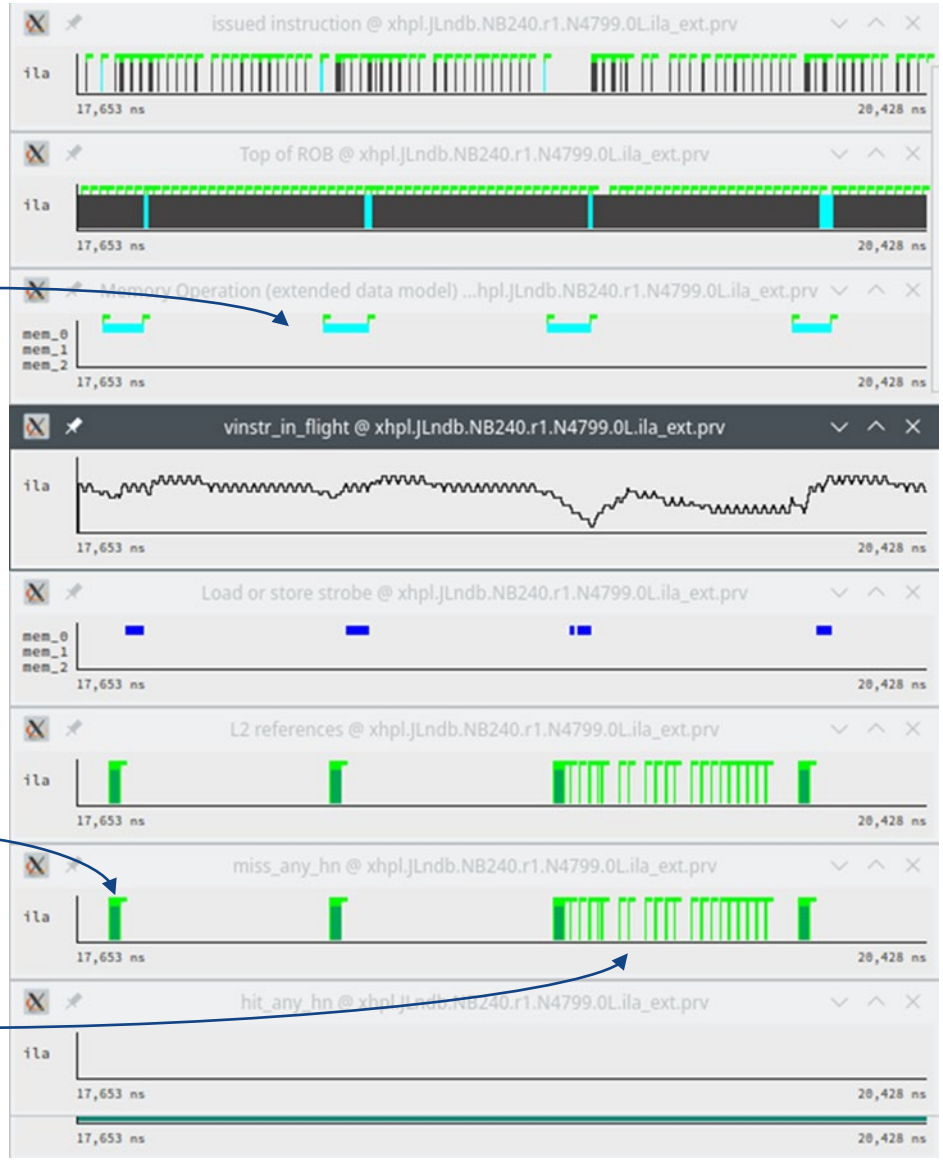

ILA traces

- Impact of memory latency
  - + 0 cycles
  - +100 cycles
  - +200 cycles
  - +300 cycles

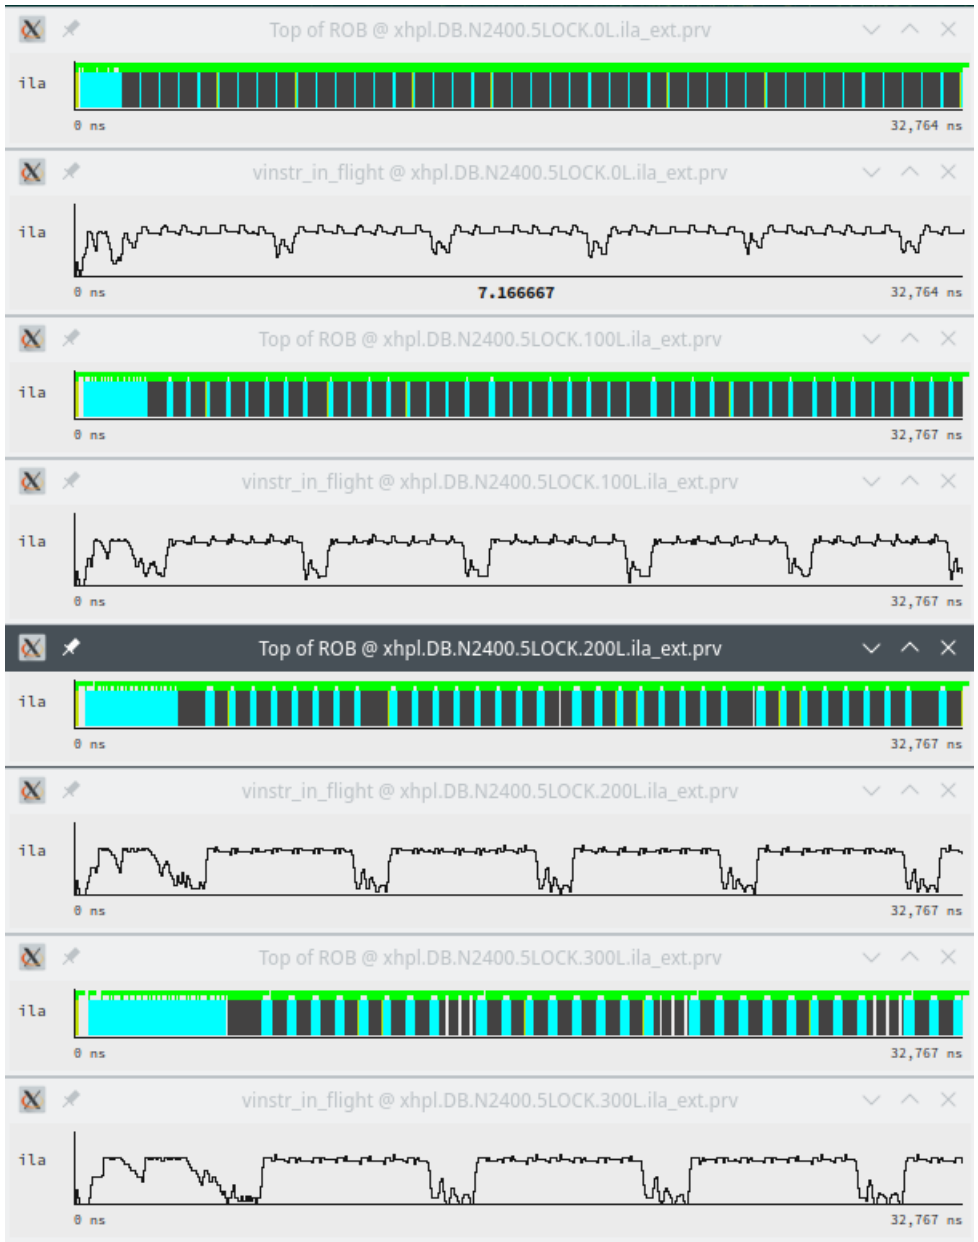

ILA traces

- Impact of memory latency (+300 cycles)

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

EXCELENCIA SEVERO OCHOA

A productivity example

Prof. Jesús Labarta

BSC & UPC

Scalable Tools Workshop

Granlibakken, June 19th 2023

Productivity

- ISO RTM code
- Half an hour with the compiler expert …
  - 4.1x faster than sequential → 15x !!!

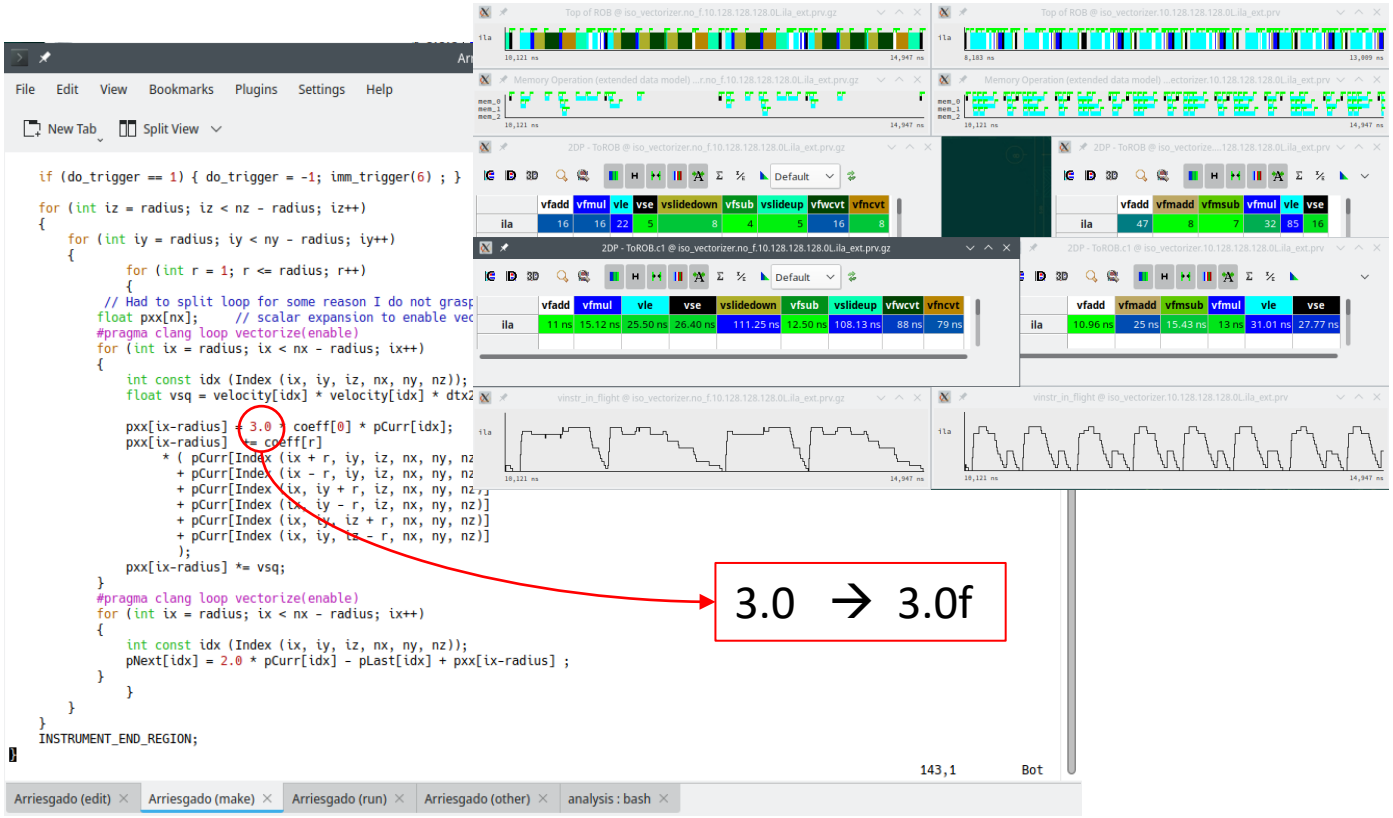

Productivity

- Good ?
  - Still only 15% utilization of core data channel ☹
- Codesign suggestions for
  - Compiler / App
    - Loop interchange/split improvement to reduce memory accesses (store !), reuse register and lookahead (store stalls)
    - Reorder loads to overlap memory and cache access
    - 3x unroll to increase register usage
  - Architecture
    - More concurrent loads
    - Increase cache interleave factor
    - …

Architects often sensitive

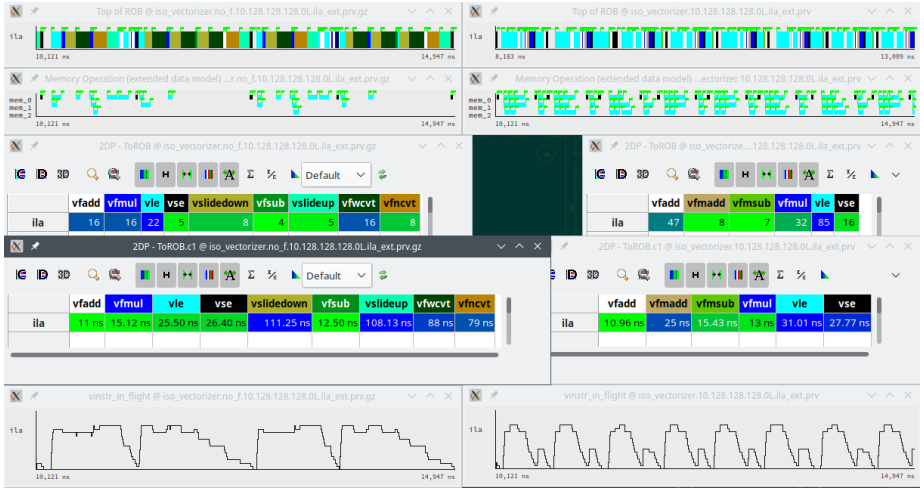

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

EXCELENCIA SEVERO OCHOA

Conclusion

Prof. Jesús Labarta

BSC & UPC

Scalable Tools Workshop

Granlibakken, June 19th 2023

Conclusion

- Hierarchy of tools for Levels of Detail in Performance Analysis
  - Data gathering
    - Extrae, PAPI, …
    - ILA2PRV
  - Visualization → Paraver
  - Analytics

- Some analyses for RVV …
- … to discuss with RTL, Compiler, Application implementors

- An approach that I believe could also be useful to improve productivity in other co-design efforts

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

EXCELENCIA SEVERO OCHOA

END

Prof. Jesús Labarta

BSC & UPC

Granlibakken, June 19th 2023