



**Barcelona
Supercomputing
Center**
Centro Nacional de Supercomputación



The RISC-V vector processor in EPI

Prof. Jesús Labarta
BSC & UPC

Barcelona, September 27rd 2021

Age before beauty

(disclaimer)

- Behavior (insight/models)
- Detail performance analytics
- Work instantiation and order
- Malleability
- Possibilities
- Elegance

before syntax
before aggregated profiles
before overhead
before fitted rigid structure
before how tos
before one day shine

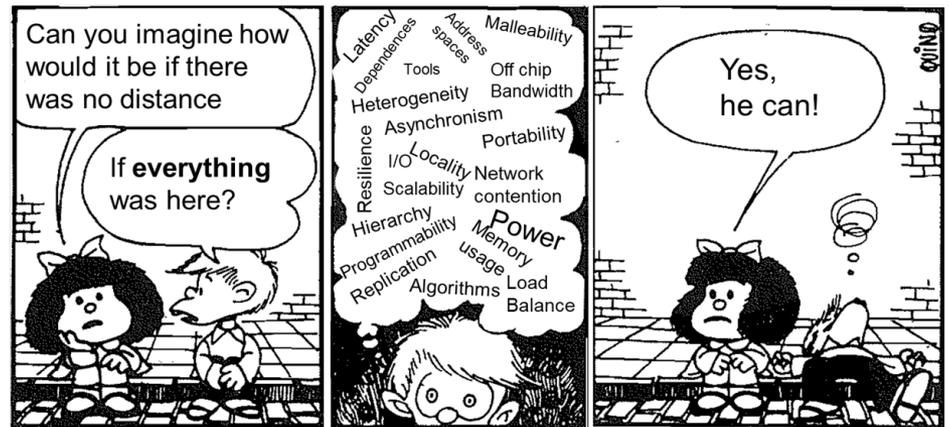


El abuelo cebolleta
ataca de nuevo

EPI Objectives

- Components (low power microprocessor technologies) ...
 - ARM based SoC
 - RISC-V based accelerator

- ... to be combined to target
 - HPC
 - HPDA
 - Emerging
 - Automotive
 - ...

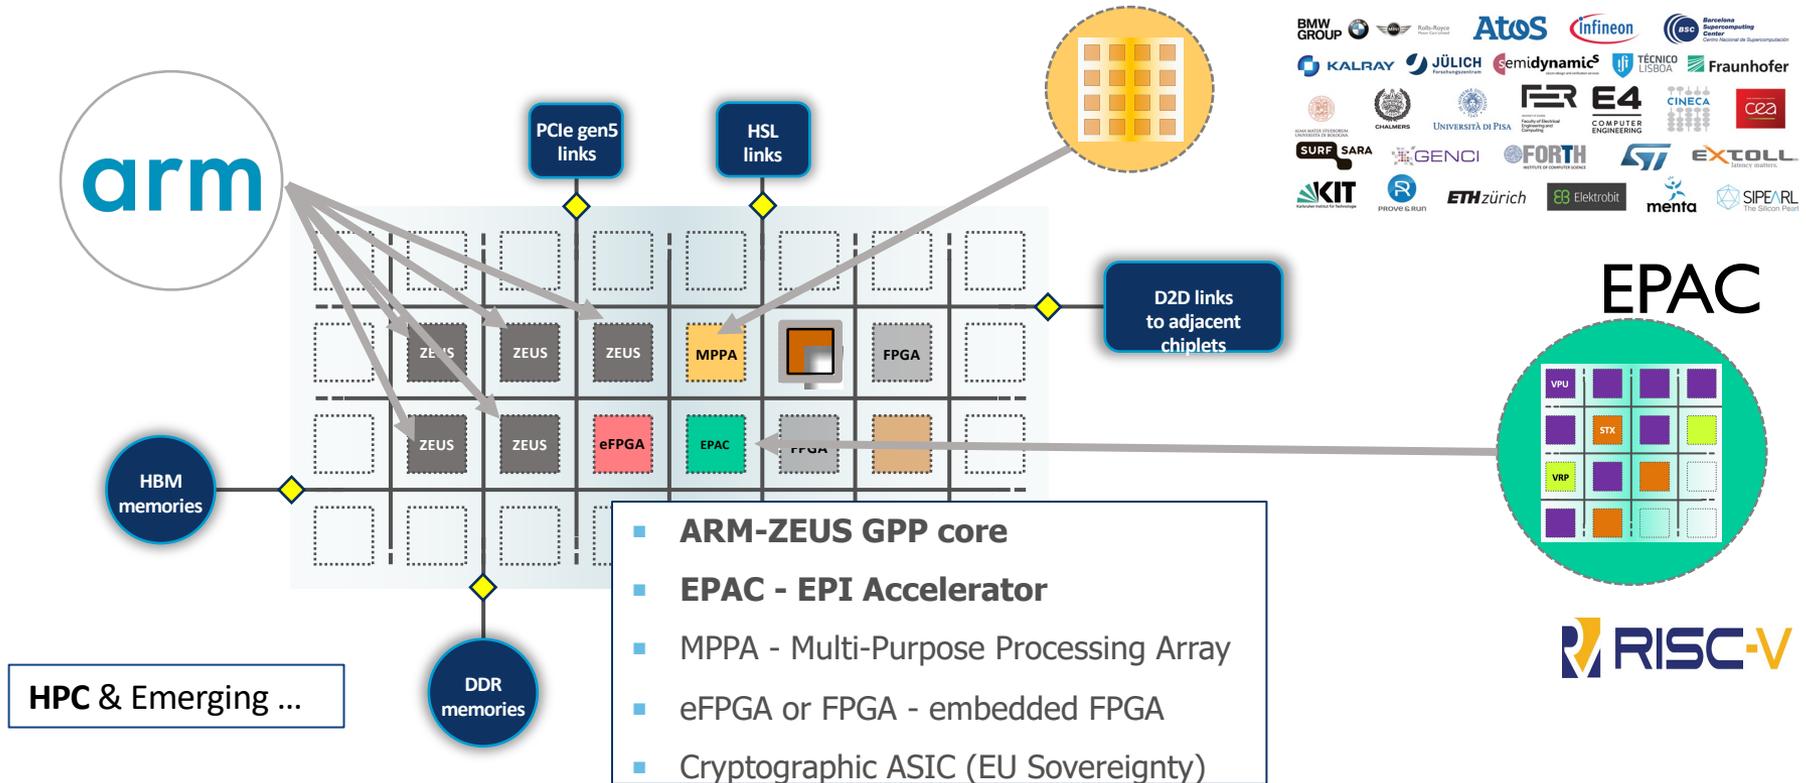


- Co-design role:
 - Application characterization. Identification of fundamentals !!!

Three streams

- General purpose
 - ARM SVE
 - BULL: System integrator → chip integrator
 - Spinoff → SiPearl
- Accelerator
 - RISC-V
 - EU design: BSC, CEA, Chalmers, ETHZ, EXTOLL, E4, IST, FORTH, FhG, Semidynamics, UNIBO, UNIZG
- Automotive
 - Infineon, ...

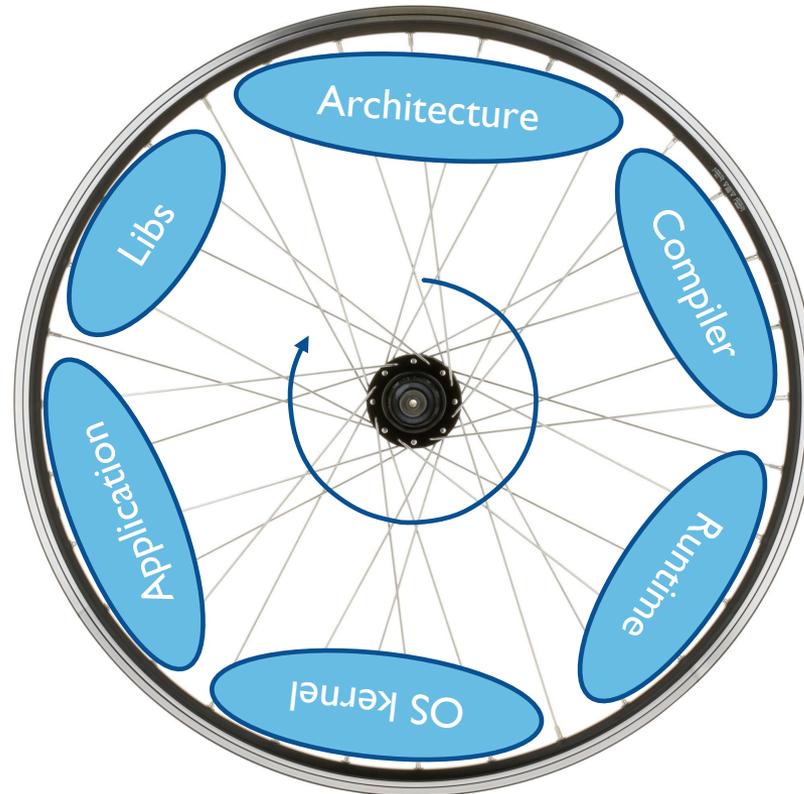
Overall architecture



Towards "Exascale"



Holistic Co-design



Best place to address an issue

Fundamentals

Balance

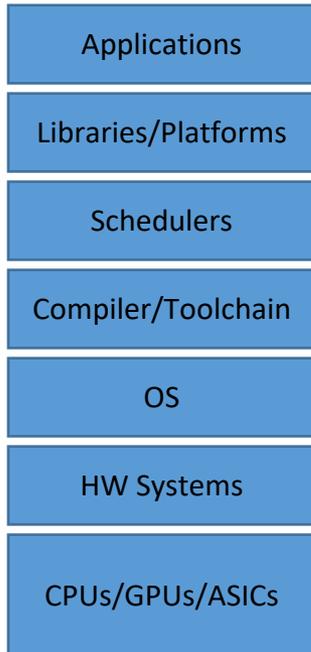
Mindset

Productivity

Efficiency

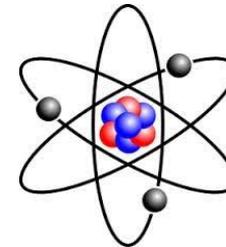
The need/ideal of a vision

- A unified theory
 - Technical
 - Economy / Politics



“As above, so below”
Similar concepts/mechanisms
at all levels

“Of fundamental principles”
“Steered by detailed insight”



Still not there ...
... still not able to communicate



principles

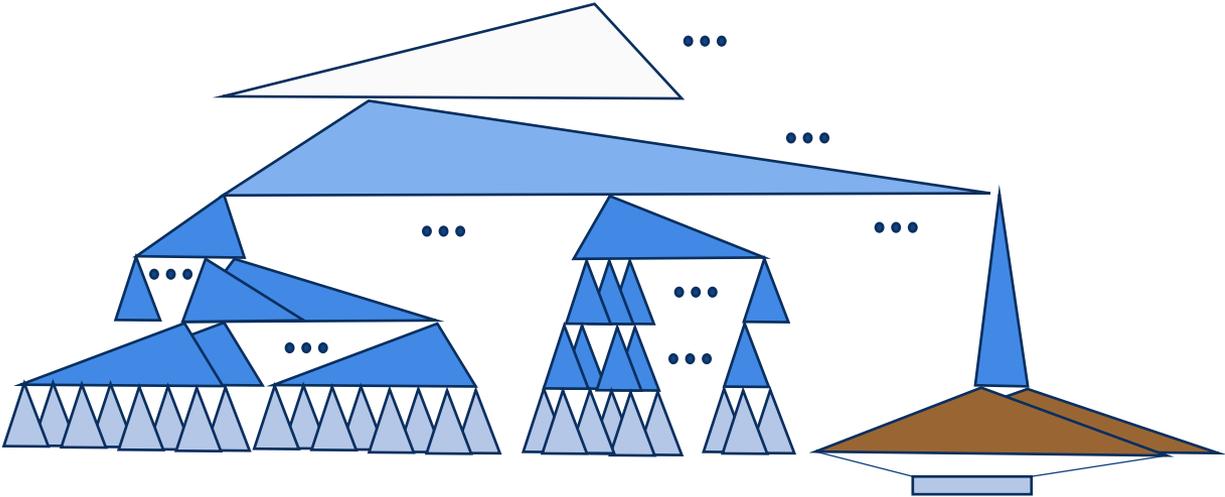
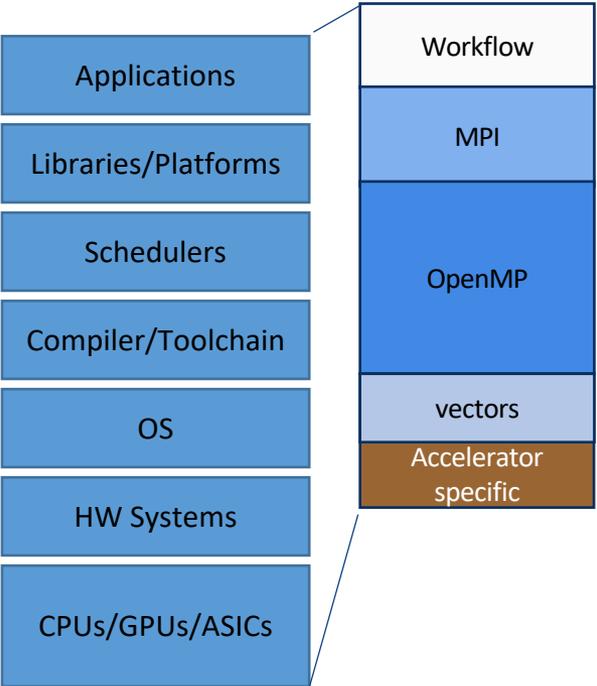


- Balanced hierarchy
- Throughput oriented: asynchrony and overlap
- Malleability and coordinated scheduling
- Homogenizing heterogeneity
- Detailed analysis and insight on behavior

Balanced hierarchy

Expression & exploitation of Parallelism

$$10^6 = 1 \times 10^6 = 10 \times 10^5 = 10^2 \times 10^2 \times 10^2$$



Long vectors. 8 lanes per core
“Limited” number of “general purpose” control flows within tile

Latency → Throughput: asynchrony and overlap

- Applications
- Libraries/Platforms
- Schedulers
- Compiler/Toolchain
- OS
- HW Systems
- CPUs/GPUs/ASICs

Interoperability MPI + OpenMP

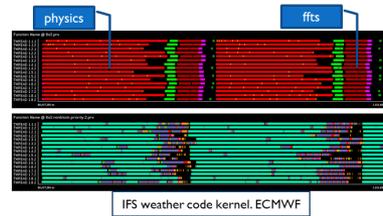
- Taskify MPI calls

Task based models

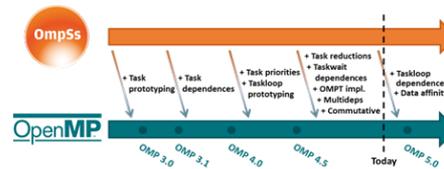
- Single mechanism
 - Concurrency
 - Locality & data management

Long vectors

- decouple Front end – back end
- Convey access pattern semantics to the architecture. Potential to optimize memory throughput



IFS weather code kernel. ECMWF

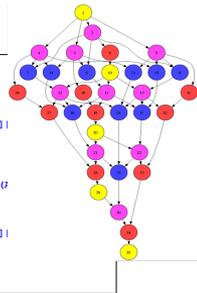


Task based computational workflows



```

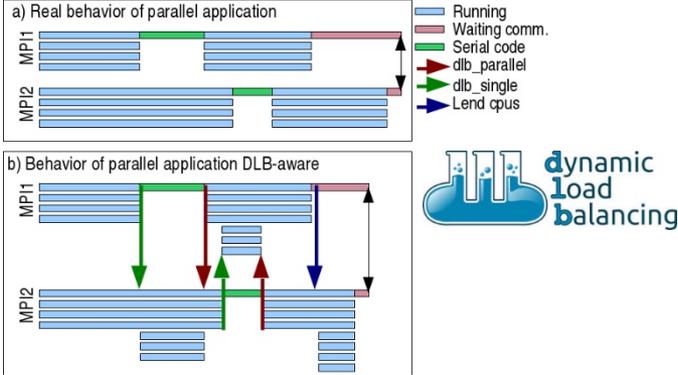
void Cholesky(int NT, float *A[NT][NT]) {
  for (int k=0; k<NT; k++) {
    #pragma omp task inout ([TS][A[k][k]])
    spotrf (A[k][k], TS) ;
    for (int i=k+1; i<NT; i++) {
      #pragma omp task in([TS][A[k][i]]) inout ([TS])
      stram (A[k][k], A[k][i], TS);
    }
    for (int i=k+1; i<NT; i++) {
      for (j=k+1; j<= i) {
        #pragma omp task in([TS][A[k][i]], [TS][TS] (
          inout ([TS][TS] (*A[j][i]))
        )
        sgum( A[k][i], A[k][j], A[j][i], TS);
      }
      #pragma omp task in ([TS][A[k][i]]) inout([TS])
      sayzk (A[k][i], A[i][i], TS);
    }
  }
}
    
```



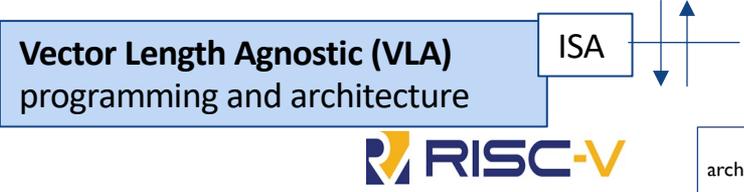
Malleability & Coordinated scheduling



Coordination (policies)
 Composability, interoperability
 (semantic impedance matching)



A wish:
Handoff scheduling



Micro architecture decides

```

saxpy:
vsetvli a4, a0, e32, m8
v1w.v v0, (a1)
sub a0, a0, a4
slli a4, a4, 2
add a1, a1, a4
v1w.v v8, (a2)
vfmacc.vf v8, fa0, v0
vsw.v v8, (a2)
add a2, a2, a4
bnez a0, saxpy
ret
  
```

Homogenizing Heterogeneity

Applications

Libraries/Platforms

Schedulers

Compiler/Toolchain

OS

HW Systems

CPUs/GPUs/ASICs

```
void axpy_omp_nest (double a, double *dx, double *dy, int n) {
    int i, chunk;
    #pragma omp taskloop
    for (i=0; i<n; i+=TS) {
        chunk= n>i+TS? TS : n-i;
        #pragma omp target map(to:dx[i:i+chunk], tofrom:dy[i:i+chunk])
        axpy_omp (a, &dx[i], &dy[i], chunk);
    }
}
```

```
void axpy_omp (double a, double *dx, double *dy, int n) {
    int I, chunk;
    #pragma omp taskloop
    for (i=0; i<n; i+=TS) {
        chunk= n>i+TS? TS : n-i;
        axpy_SIMD (a, &dx[i], &dy[i], chunk);
    }
}
```

```
void axpy_SIMD (double a, double *dx, double *dy, int n) {
    int i;
    #pragma omp simd
    for (i=0; i<n; i++) dy[i] += a*dx[i];
}
```

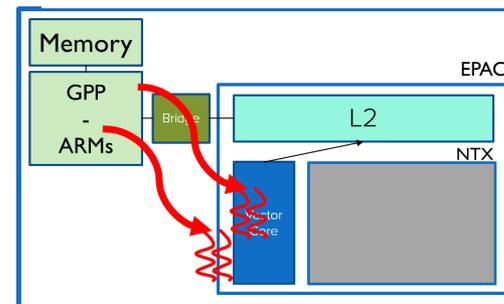
VLA helps homogenize Heterogeneous Performance

- ~ Big – Little cores,

Nested tasked/workshared



- Offload regular OpenMP



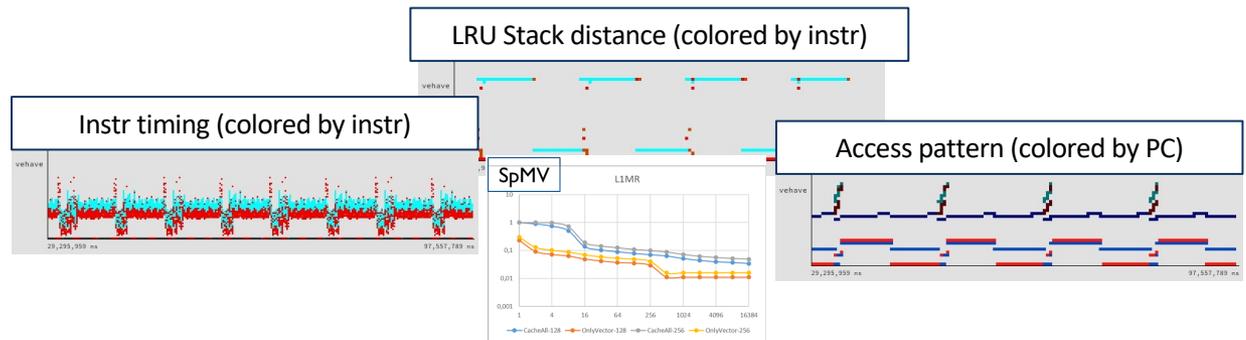
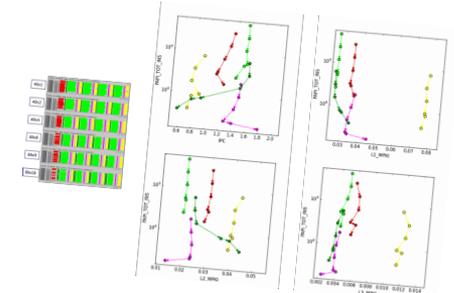
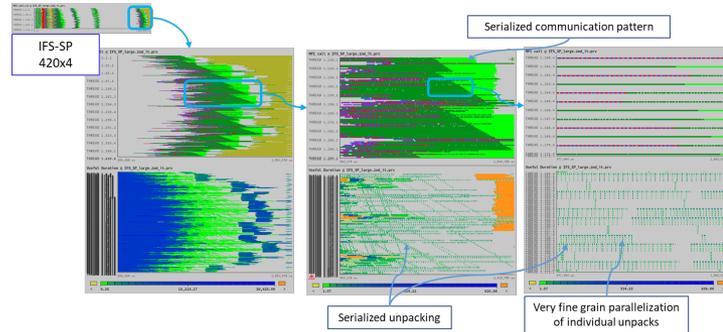
- HW support: IO coherence

Detailed analysis and Insight on behavior



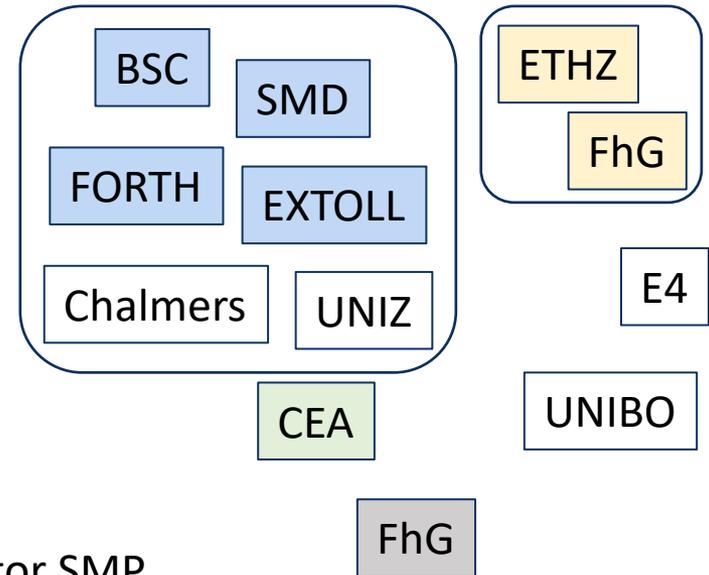
<https://pop-coe.eu/>

- Applications
- Libraries/Platforms
- Schedulers
- Compiler/Toolchain
- OS
- HW Systems
- CPUs/GPUs/ASICs



Visions and collaborations

- STX:
 - Specific Accelerator devices
 - AI
 - Stencil
- RVV
 - ISA is important, RISC-V Vector
 - “Accelerator”
 - Easier entry, focus
 - → Standard self hosted, general purpose vector SMP
- VRP:
 - Extended precision arithmetic



EPAC architecture

RVV

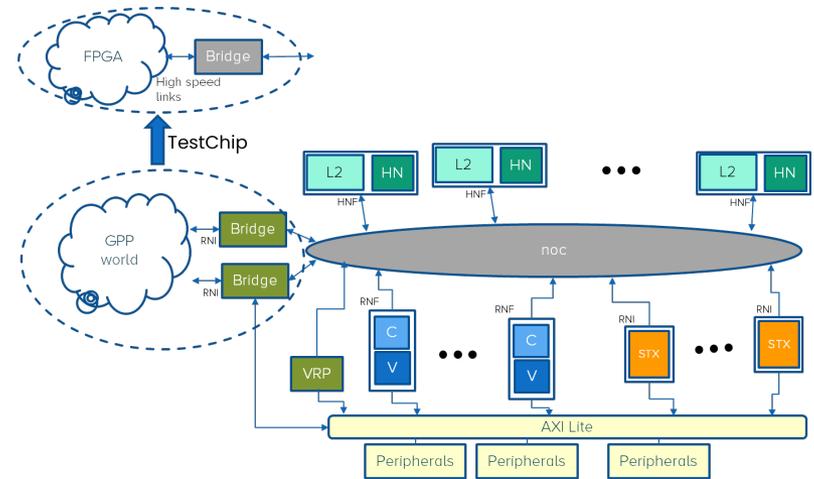
- RV64GCV (→ 8x)
 - 2 way in order core
 - Decoupled VPU
 - 8 lanes
 - Long vectors (256 DP elements)
 - L1 - MESI coherency
 - CHI interface NoC
 - 1 line / cycle
 - L\$2: 256KB/module
 - Allocation control mechanisms
- No in tile L\$3

STX

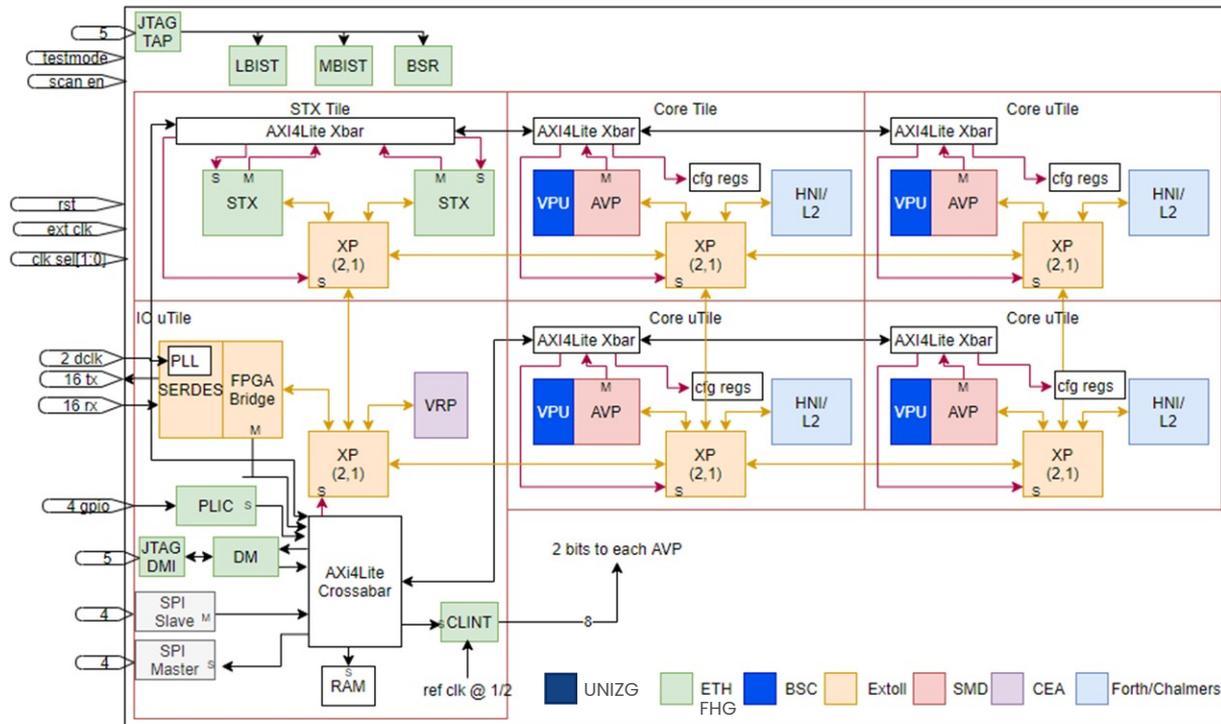
- DL and stencil specific accelerators
- Extensions to planned NTX
 - Programmable address generators →
 - lightweight RISC-v core + fat FPU + (Streaming Semantics & FREP)

VRP

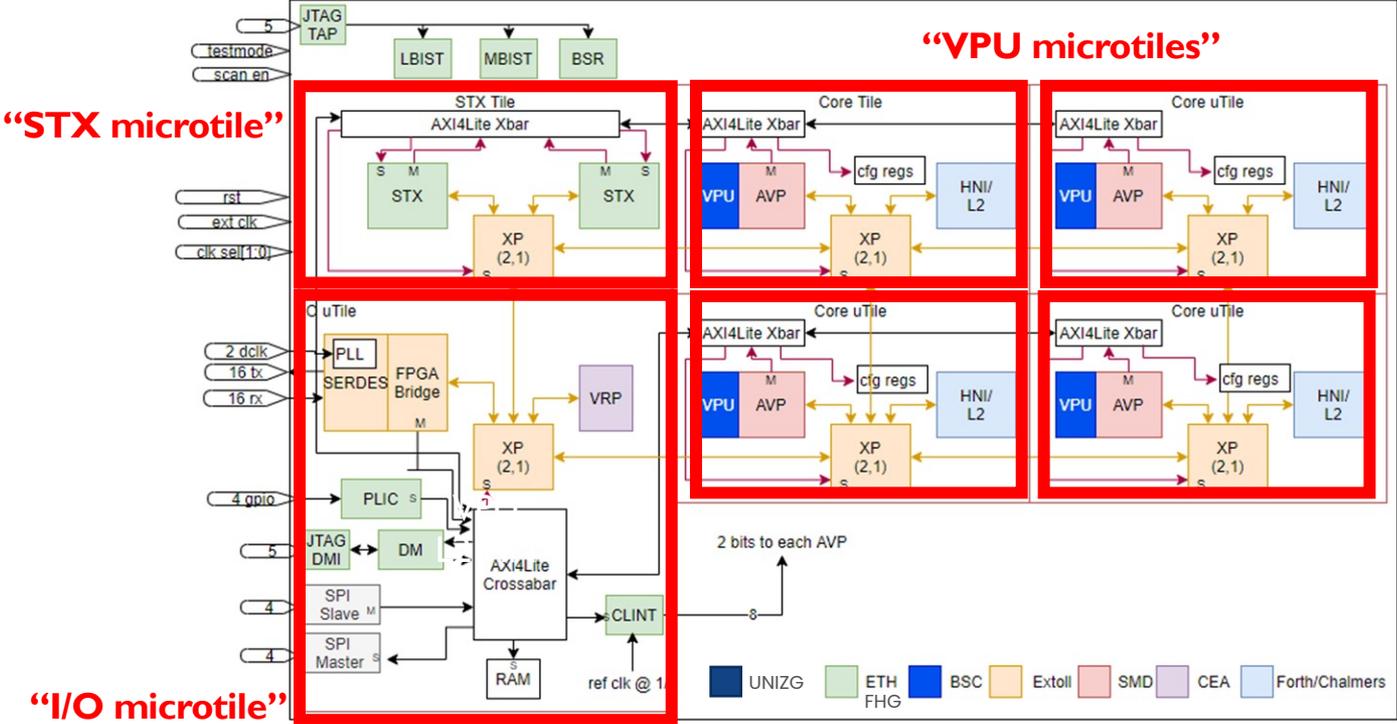
- Variable precision processors



EPAC Test Chip Block Diagram

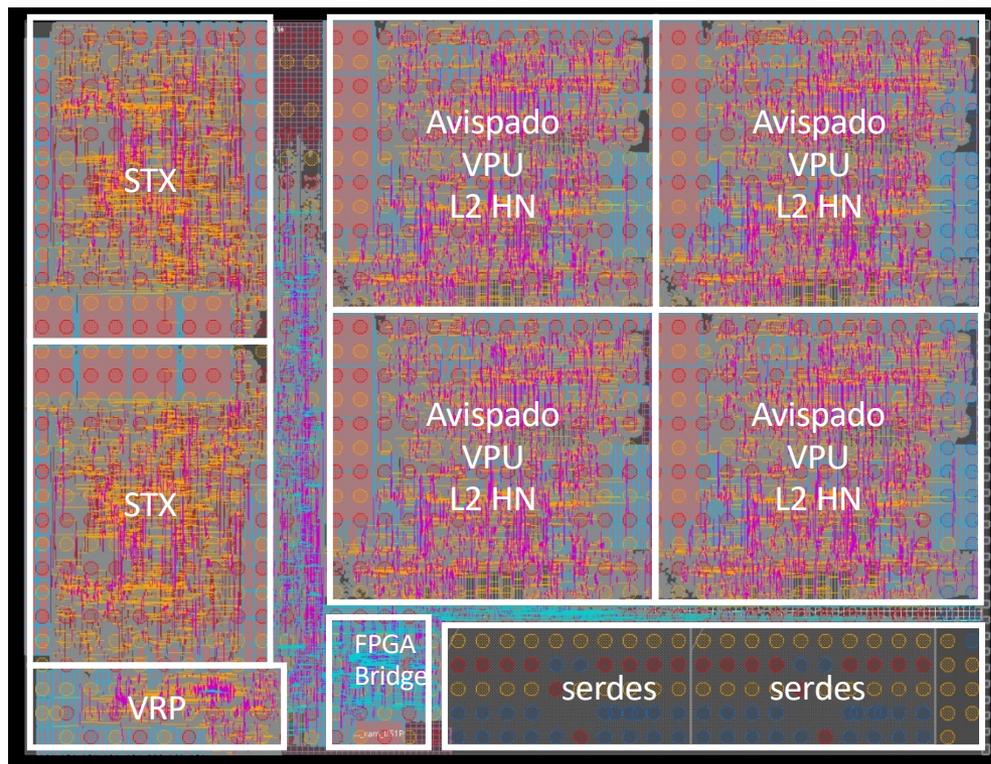


EPAC Test Chip Block Diagram



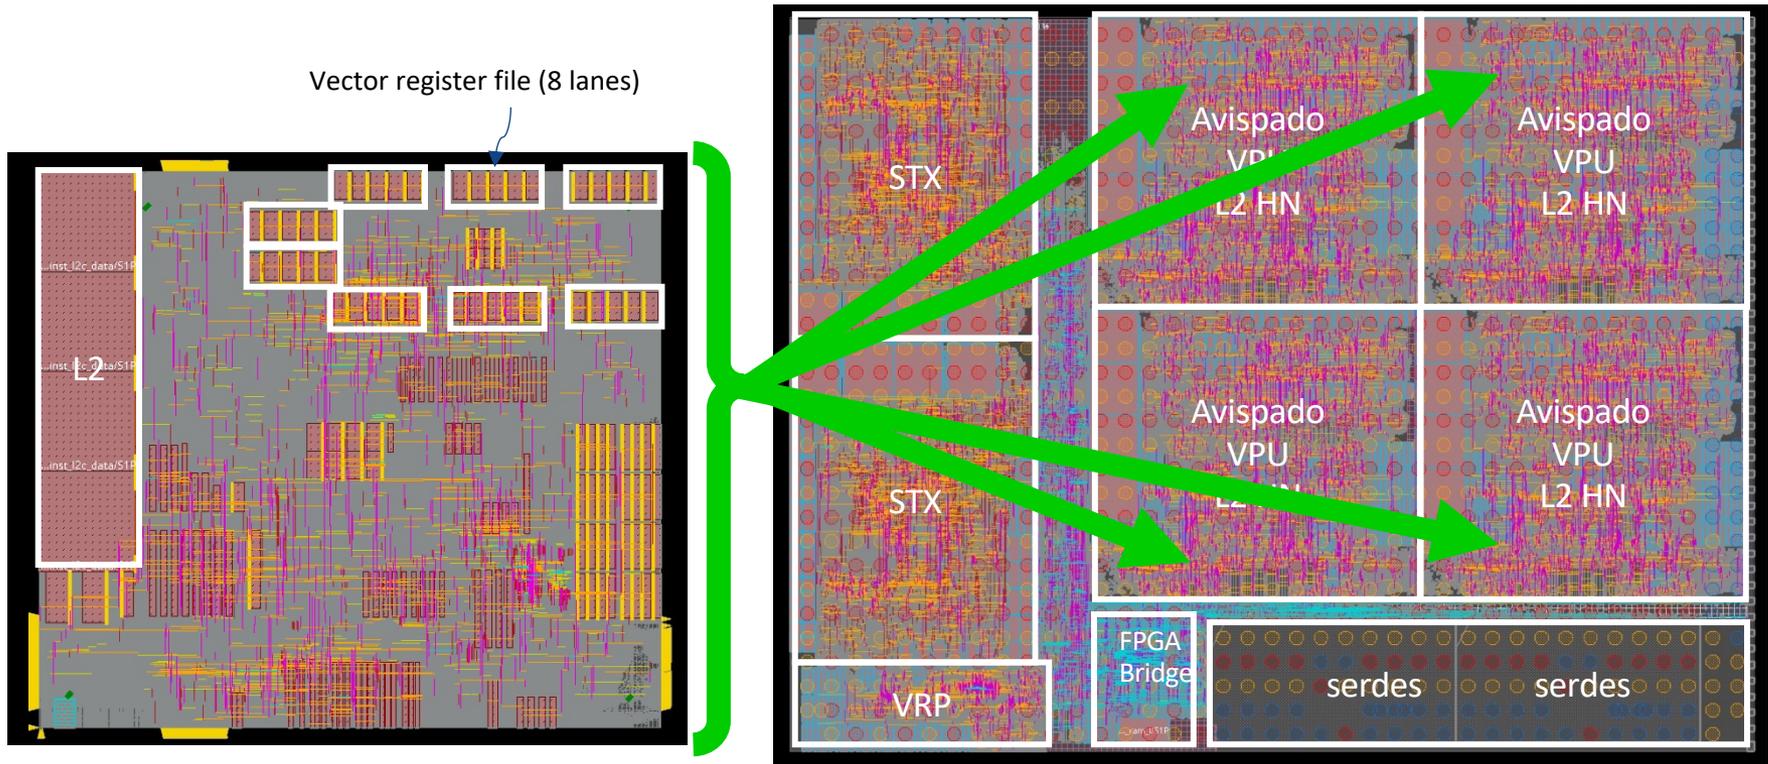
Physical design

- Final Top level chip floorplan
- Total area:
 - 5943 X 4593 um²
 - (27.297 mm²)



Physical design

- 4 “VPU microtiles”
 - 2.517 mm² each



Alive

- Starting bring up process
 - First words, ...
 - ... tiny shell ...
 - ... first vector code @ 1 GHz ...



```
carv@boulker: ~/Desktop/bringup_20210916/tool_new/bringup_20210917
EPAC JTAG Console Client 0.1
Connecting to JTAG Console [0] ...
Press CTRL+AA for exit

-----
| Welcome to EPAC TC Bring-Up Shell |
-----
apac@nikitas$ help
help          Prints the available commands.
echo          Echo the given input.
ping          Pings the core.
banner        Shows a banner with the given input.
cpuinfo       Prints information about the current core.
sleep         Sleeps for a given number of seconds.
uptime        Tell how long the system has been running.
axy           Run the axpy benchmark.
axy_vector    Run the axpy vector version benchmark.
apac@nikitas$ cpuinfo
Chip Frequency   = 1000 Mhz
Core HWID        = 0
Cycle Count      = 216041194340
Instruction Count = 1664256219
apac@nikitas$ uptime
up for 0 hours 03 minutes 46 seconds
apac@nikitas$ axpy 8192
Running AXPY Scalar with 8192 array elements
Init time: 352338 cycles
axy scalar reference time: 281334 cycles
done
Result ok !!!
apac@nikitas$ axpy_vector 8192
Running AXPY Vector with 8192 array elements
Init time: 352341 cycles
axy vector time: 9725 cycles
done
Result ok !!!
apac@nikitas$
```



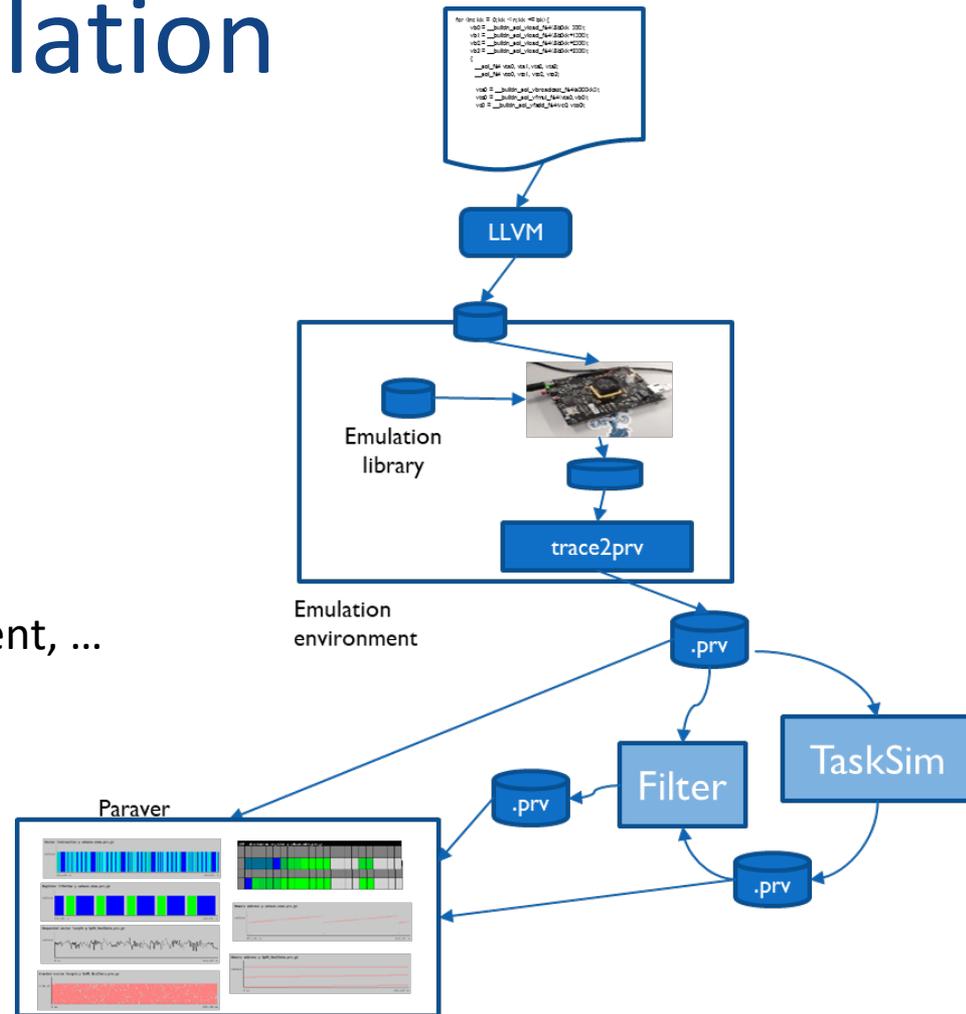
**Barcelona
Supercomputing
Center**
Centro Nacional de Supercomputación



SDVs

RVV Software emulation

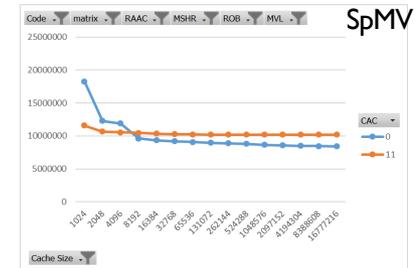
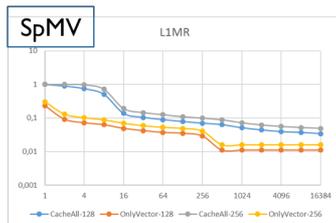
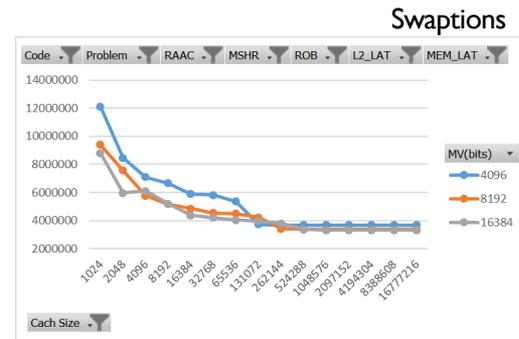
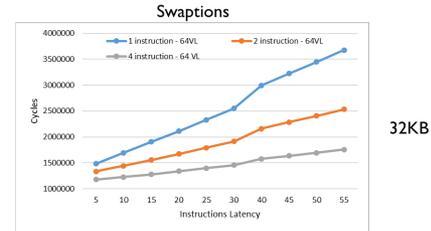
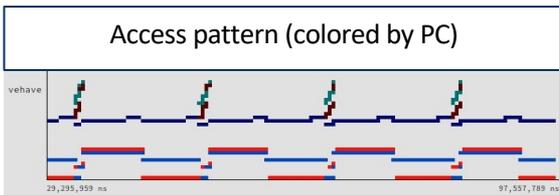
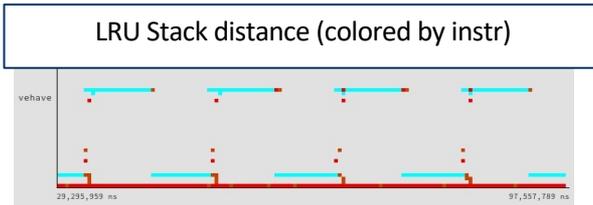
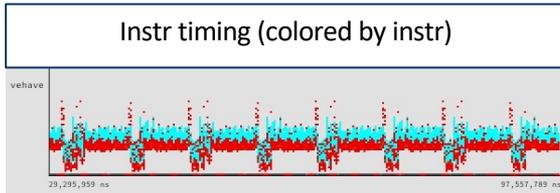
- Functional
 - Full OS,
 - On QEMU & scalar RISC-V cores
- Timing model
 - Microarchitectural parameters
 - MAXVL, OoO, Locality management, ...



• <https://repo.hca.bsc.es/gitlab/epi-public/risc-v-vector-simulation-environment>

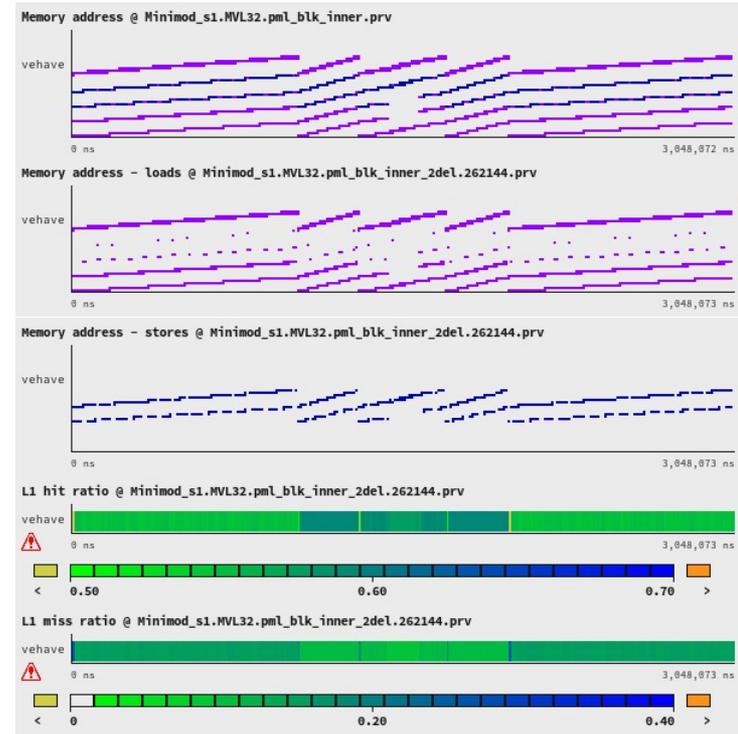
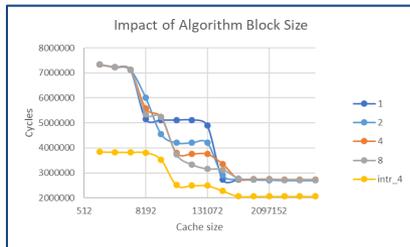
RVV Software emulation

- Detailed analysis & insight

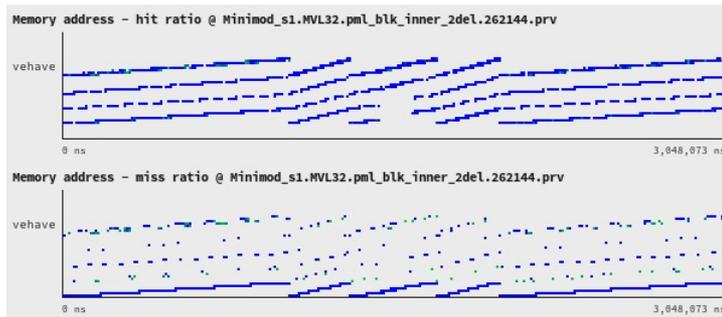


RVV Software emulation

- Detailed analysis & insight



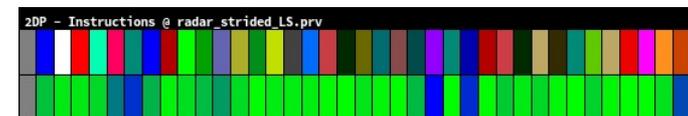
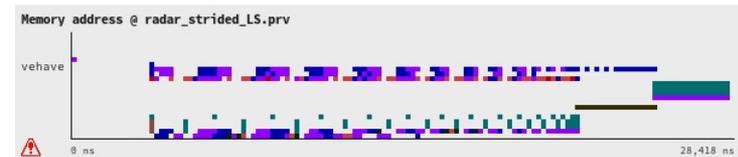
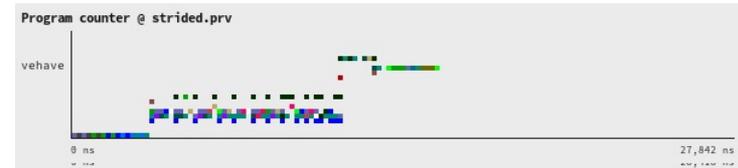
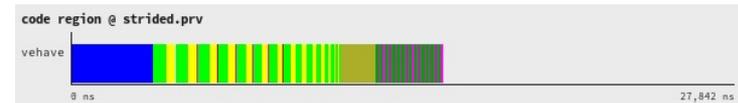
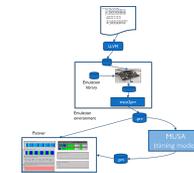
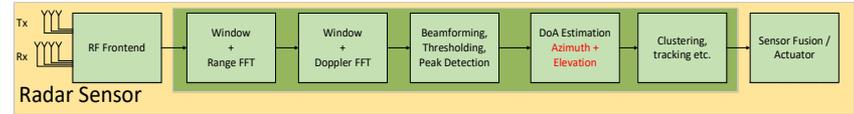
256K cache



RVV Software emulation

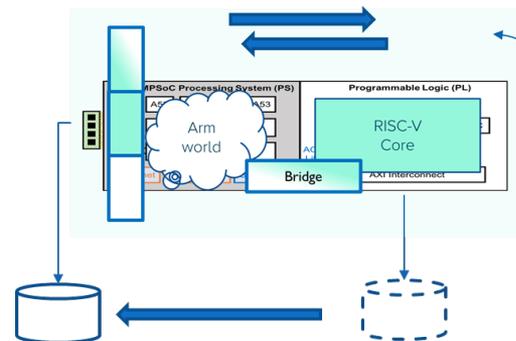


- Radar miniapp
 - Directives + automatic vectorization
 - Incremental way
 - Steering compiler optimizations
 - Complex data types
 - Indexed \rightarrow strided
 - Reuse through registers
 - Avoid optimizations generating extremely short vector lengths
 - Steering code refactoring
 - With productivity in mind !
 - Interchange, collapse, ...



Heterogeneous ARM + RISC-V

- Linux Boot on Arm and RISC-V
 - kernel version updates
 - Tracking mainline, and contributing to ongoing patch testing and review (eg. SV48, huge pages)
- OpenMP offloading
 - Asynchronous calls (via service thread pool)
 - Thread teams
- Reverse offloading
 - access to host-side I/O devices
 - Works for Linux process on RISC-V side
 - WIP for offloaded tasks



FORTH

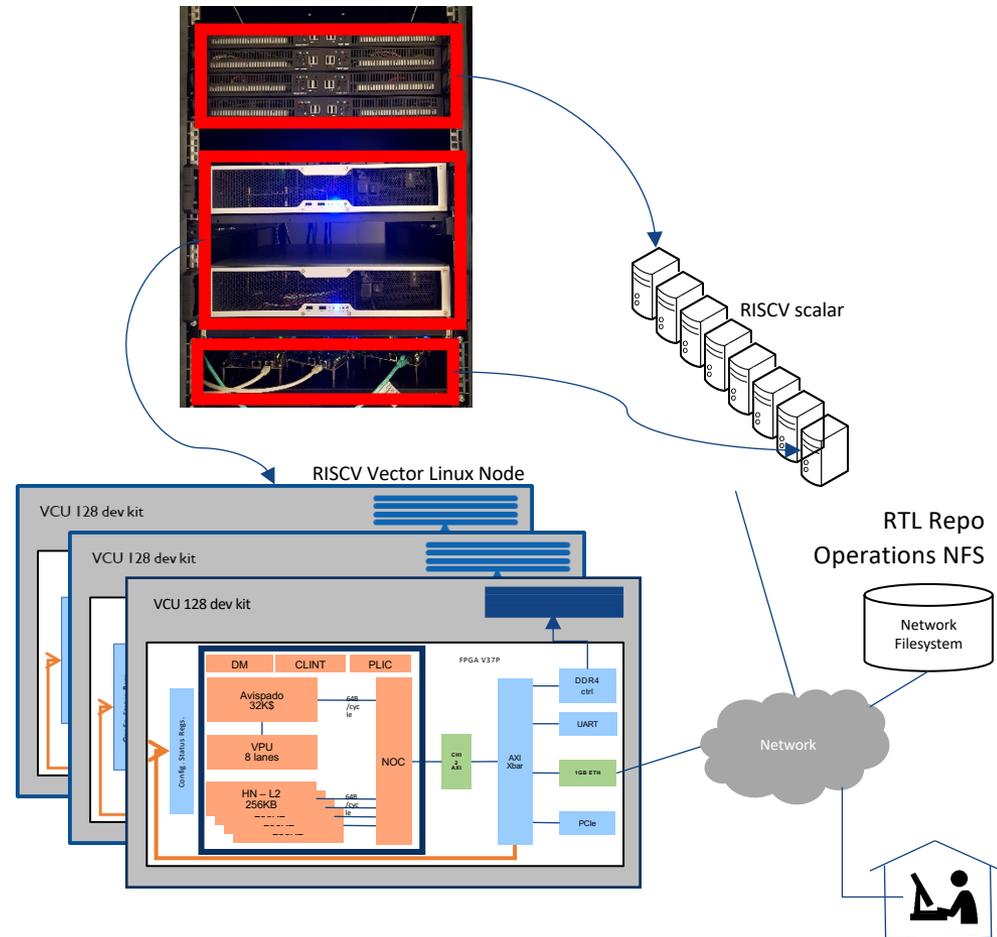
```
#on RISC-V side
$mkfs.ext4 -b 4096 /dev/vda
$mount -t ext4 /dev/vda /mnt/scratchfs
```

```
void main (...)
{ ...
  gemTarget(A, B, C, size);
  ...
}

void gemTarget(double *A, double *B, double *C, long S) {
  #pragma omp target map(to:A[0:S*S],B[0:S*S],S) \
    map(from:C[0:S*S])
  {
    for(int i = 0; i < size; i++)
      for(int j = 0; j < size; j++)
        for(int k = 0; k < size; k++)
          C[i*size + j] += A[i*size + k]*B[k*size + j];
    fp = fopen("/mnt/scratchfs/output_matrix.txt", "w+");
    for (int i=0; i<size*size; i++) fprintf(fp, "%lf ", C[i]);
  }
}
```

RVV @ FPGA & ecosystem

- HPC software stack @ Commercially available RISC-V platforms
 - SLURM, MPI, OpenMP, BSC tools, RVV software emulator
- EPI SDV platforms
 - Test user codes @ real RTL
 - Give to EPI partners an external users early access to EPI technology
 - Two step procedure



RVV @ FPGA & ecosystem

- Very preliminary but promising results
- Providing very useful feedback for EPAC1.5 co-design

```

void SpMV_vec(double *a, long *ia, long *ja, double *x, double *y, int nrows)
{
    for (int row = 0; row < nrows; row++) {
        int nnz_row = ia[row + 1] - ia[row];
        unsigned long gvl; // granted vector lengths
        int idx = ia[row];
        __epi_1xf64 v_a, v_x, v_prod, v_partial_res;
        __epi_1xi64 v_idx_row, v_three;
        y[row]=0.0;
        v_partial_res = __builtin_epi_vbroadcast_1xf64(0.0, gvl);
        for(int colid=0; colid<nnz_row; ) { //blocking on MAXVL
            gvl = __builtin_epi_vsetv1(nnz_row - colid, __epi_e64,
__epi_m1);
            v_a = __builtin_epi_vload_1xf64(&a[idx+colid], gvl);
            v_idx_row = __builtin_epi_vload_1xi64(&ja[idx+colid], gvl);
            v_three = __builtin_epi_vbroadcast_1xi64(3, gvl);
            v_idx_row = __builtin_epi_vadd(v_idx_row, v_three);
            v_x = __builtin_epi_vload_1xf64(v_idx_row, gvl);
            v_prod = __builtin_epi_vmul(v_a, v_x);
            v_partial_res = __builtin_epi_vadd(v_partial_res, v_prod);
            colid += gvl;
        }
        y[row] += __builtin_epi_vload_1xf64(v_partial_res, gvl);
    }
}

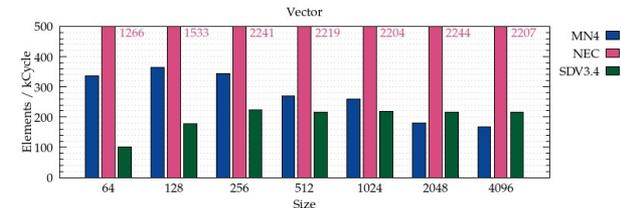
```

```

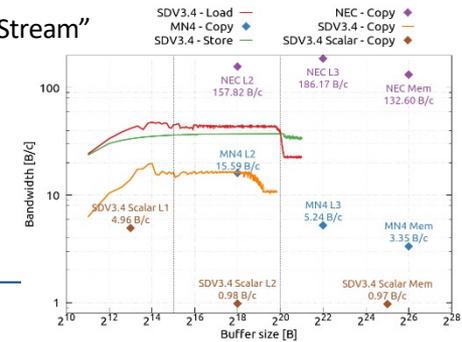
void target_inner_3d (...) {
    for (i) {
        for (j) {
            #pragma clang loop vectorize(enable)
            for (k) {
                lap = LAP;
                v[IDX3_1(i,j,k)] = 2.f*u[IDX3_1(i,j,k)]
                    -v[IDX3_1(i,j,k)]+vp[IDX3(i,j,k)]*lap;
            }
        }
    }
}

```

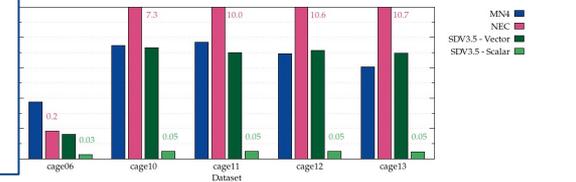
Jacobi 2D



“Stream”



SpMV



The importance of a vision



- Holistic throughput oriented vision based on long vectors and task based models
- Hierarchical concurrency and locality exploitation
- Not massive concurrency at a given level
- High bandwidth per control flow
- Push behaviour exploitation to low levels
- Co-ordination between levels
- Make it all look as classical sequential programming to ensure productivity

EPAC & SAGRADA FAMILIA ?

European Processor Initiative epi

"special" instructions

Program steered cache allocation

LONG vectors

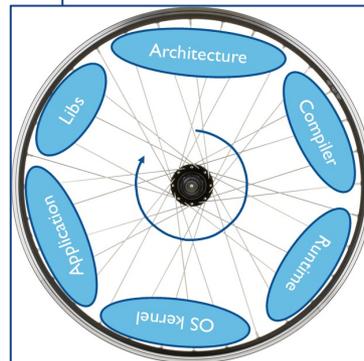
RAA

There is something "special" ...

- ...showing the way ...
- ... sustaining the effort

The throughput oriented facade

Sagrada Familia 1975

A presentation slide titled 'EPAC & SAGRADA FAMILIA ?' featuring a photograph of the Sagrada Família in Barcelona. The slide includes several callout boxes with arrows pointing to the building's spires: 'special' instructions, Program steered cache allocation, LONG vectors, and RAA. A text box at the bottom left of the image says 'The throughput oriented facade'. A small 'European Processor Initiative epi' logo is in the top right. A bulleted list below the title discusses 'special' aspects like 'showing the way' and 'sustaining the effort'. A caption at the bottom of the image reads 'Sagrada Familia 1975'.



**Barcelona
Supercomputing
Center**
Centro Nacional de Supercomputación



**EXCELENCIA
SEVERO
OCHOA**

Thanks