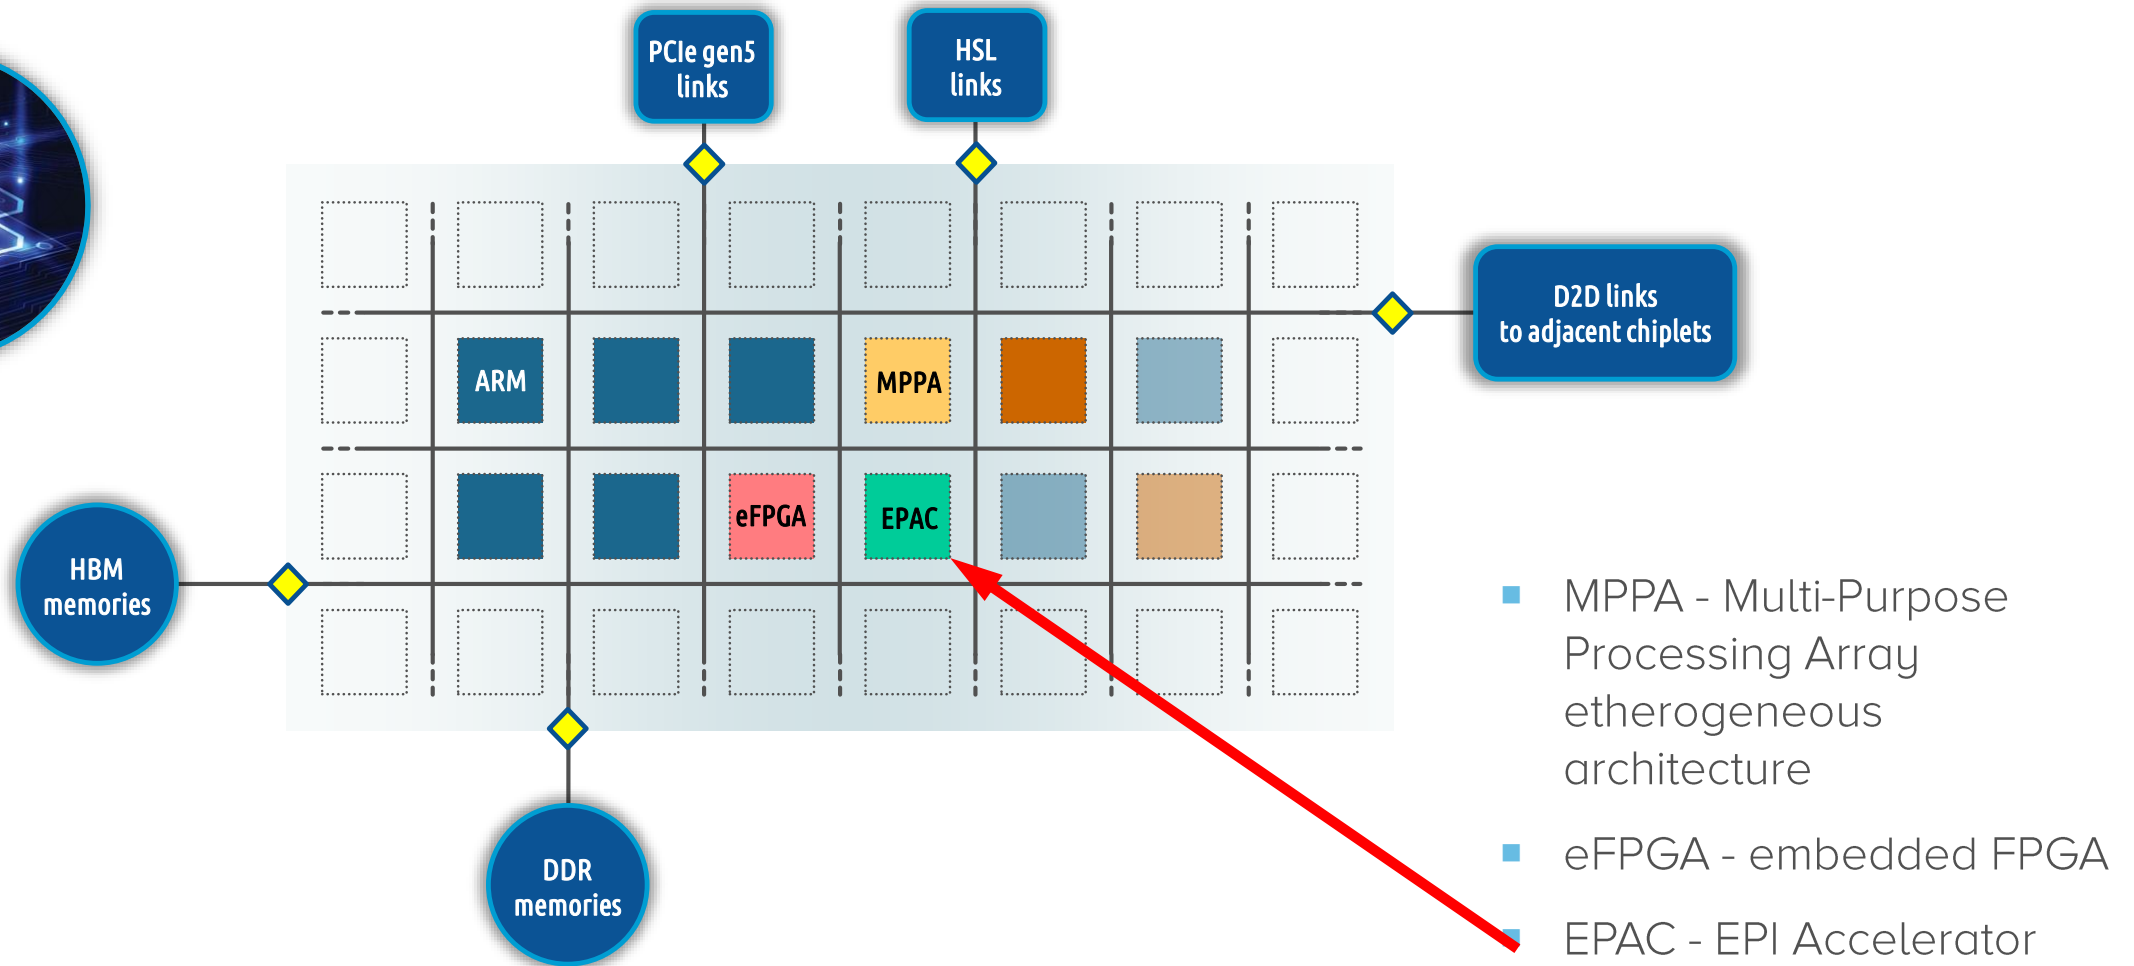




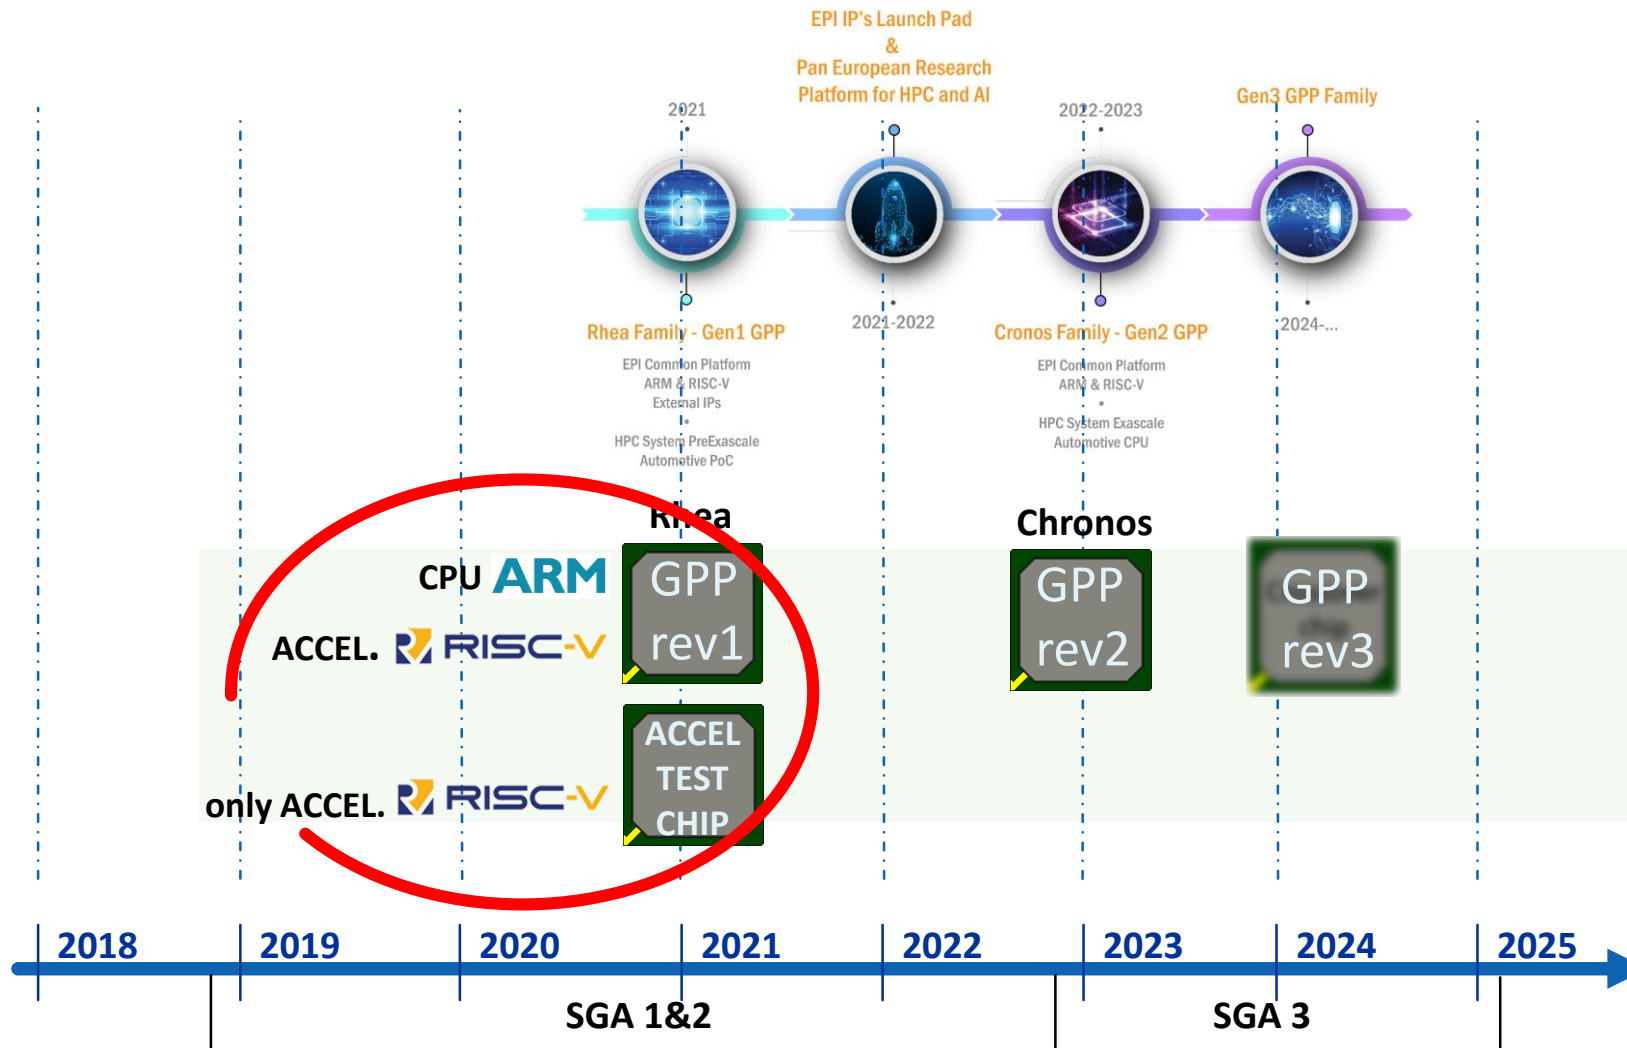
ACCELERATOR TILE

MAURO OLIVIERI

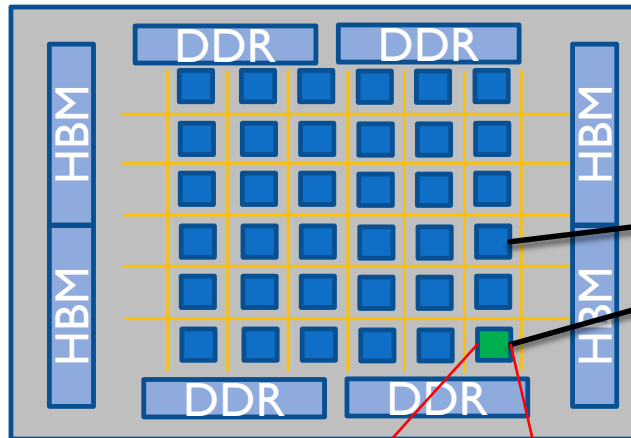
RECALL... THE GPP AND COMMON PLATFORM ARCHITECTURE



ACCELERATOR STREAM IN EPI ROADMAP

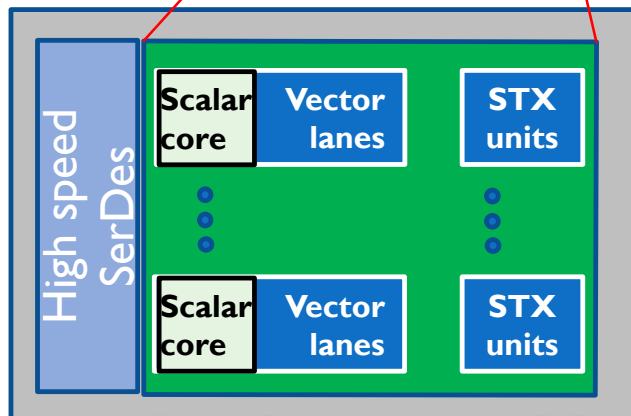


1ST GENERATION EPI CHIPS



General Purpose Processor (GPP) chip

- 6 nm, chip-let technology
- ARM-SVE tiles
- EPAC RISC-V vector+AI accelerator tiles
- L1, L2, L3 cache subsystem + HBM + DDR

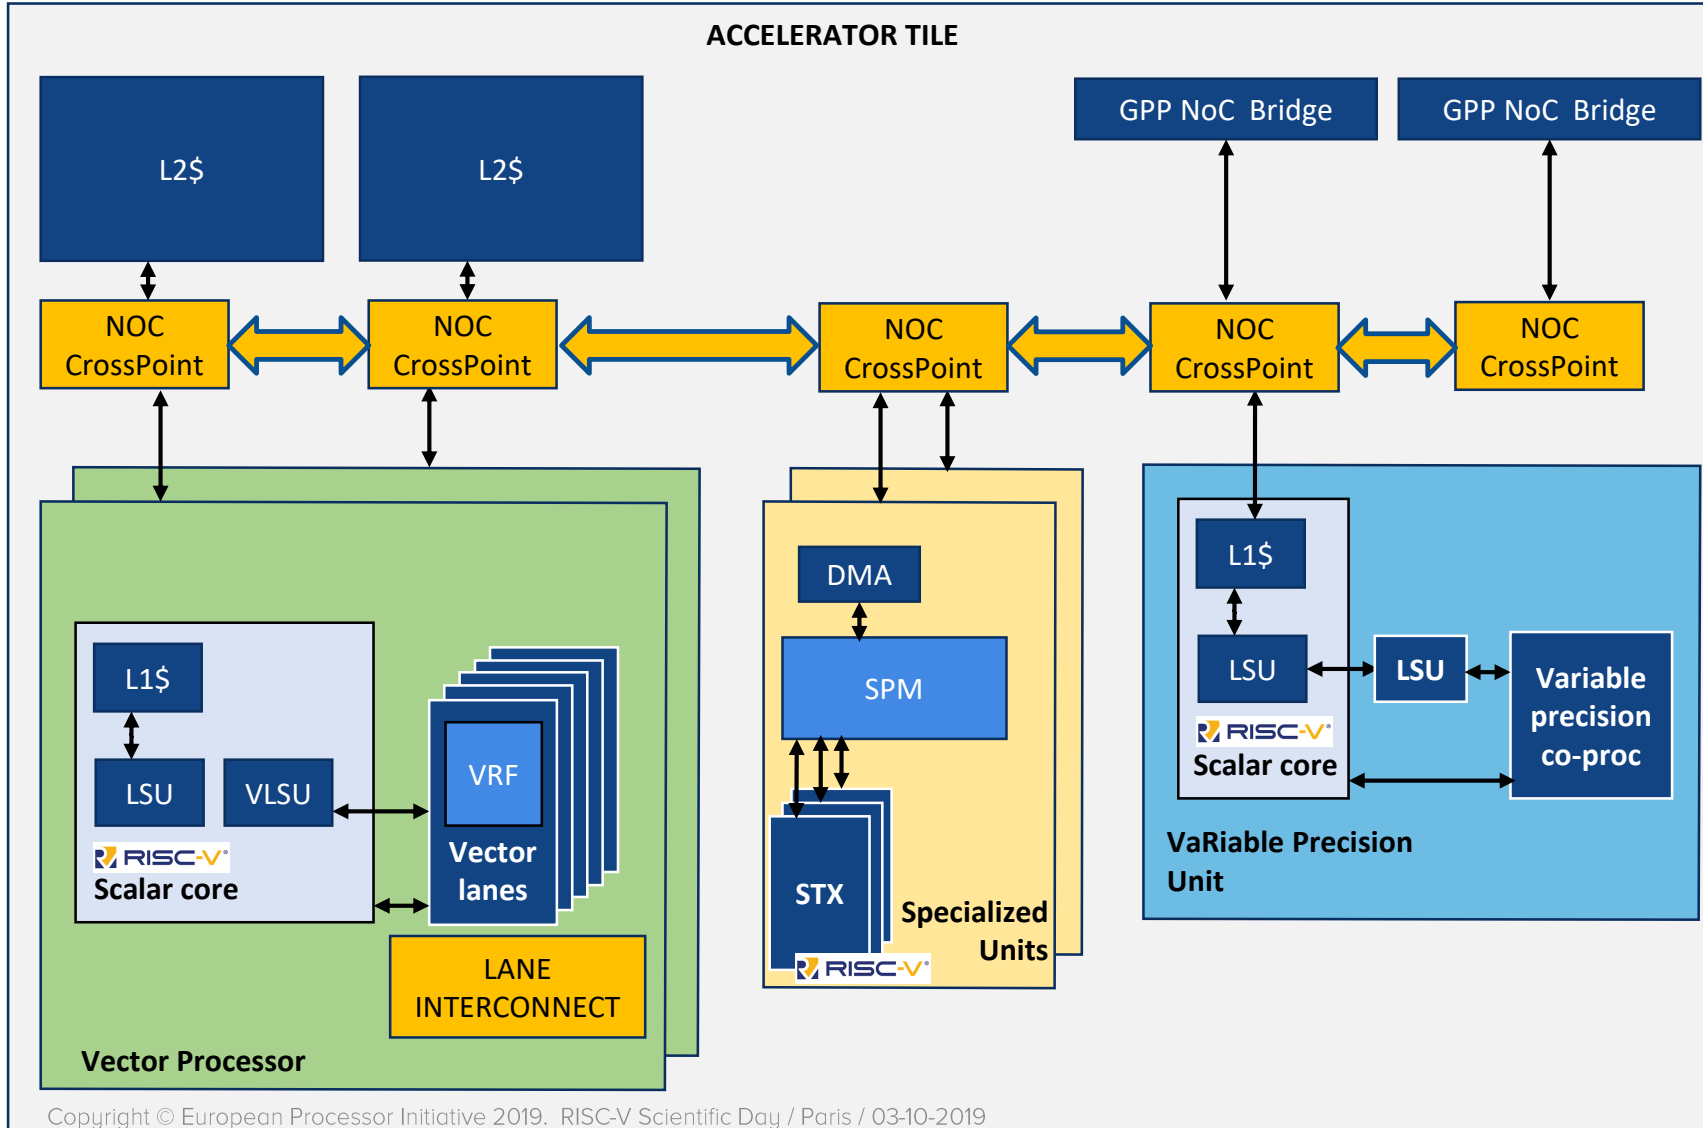


RISC-V Accelerator Demonstrator Test Chip

- 22 nm FDSOI
- Only one RISC-V accelerator tile
- On-chip L1, L2 + off-chip HBM + DDR PHY
- Targets 16 DP GFLOPS per core (vector processor only)

EPAC ARCHITECTURE VIEW

ACCELERATOR TILE

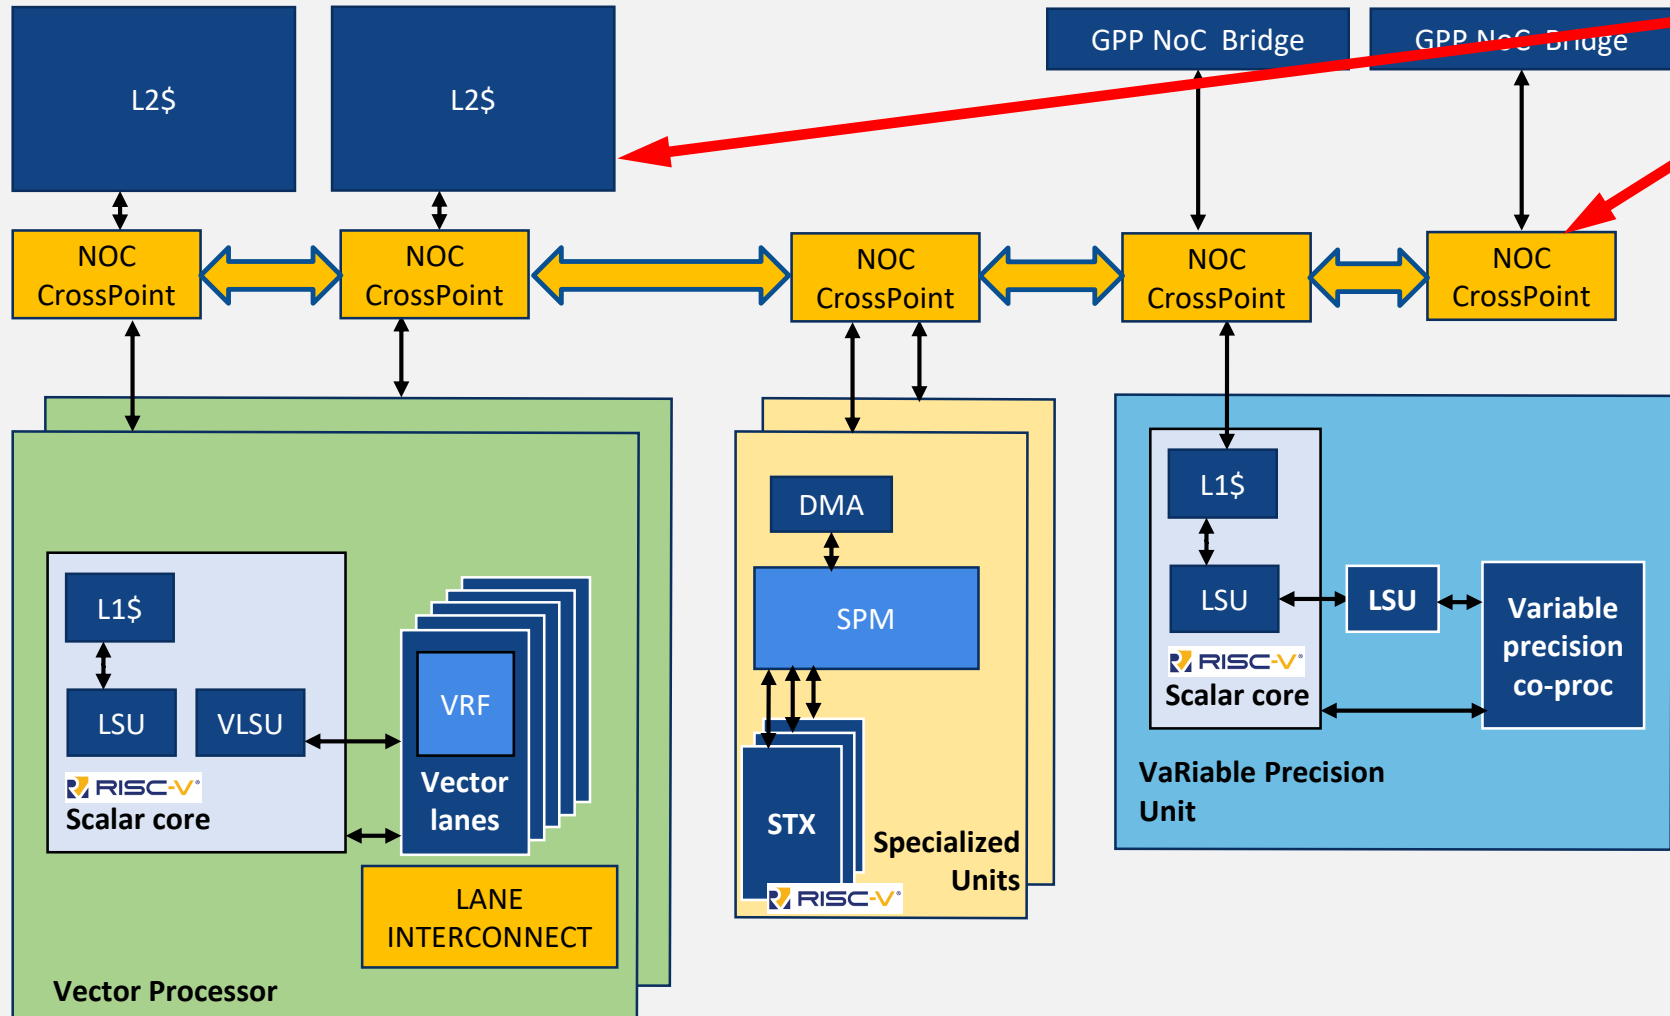


CONTRIBUTING PARTNERS:

- BSC
- CEA
- Chalmers
- E4
- ETH Zurich
- Extoll
- FORTH
- Fraunhofer
- Semidynamics

EPAC ARCHITECTURE VIEW

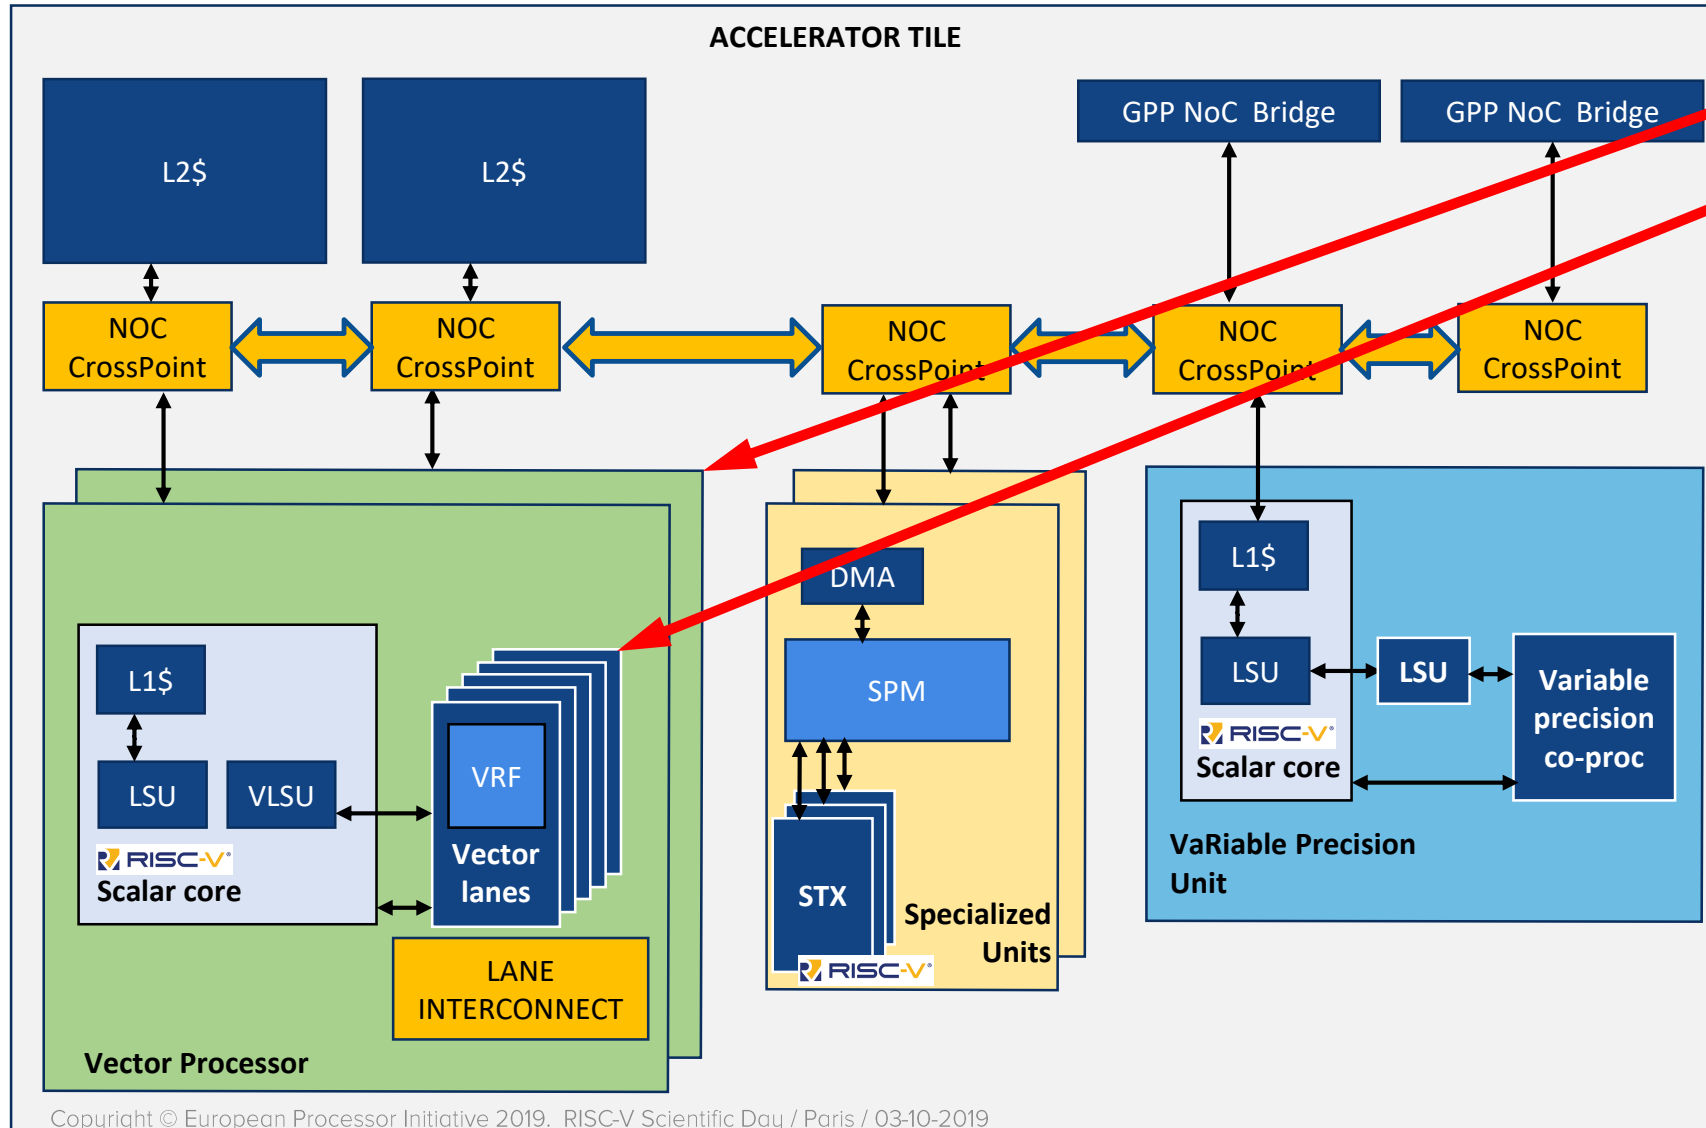
ACCELERATOR TILE



Shared L2 cache banks

Cache coherent NoC

EPAC ARCHITECTURE VIEW



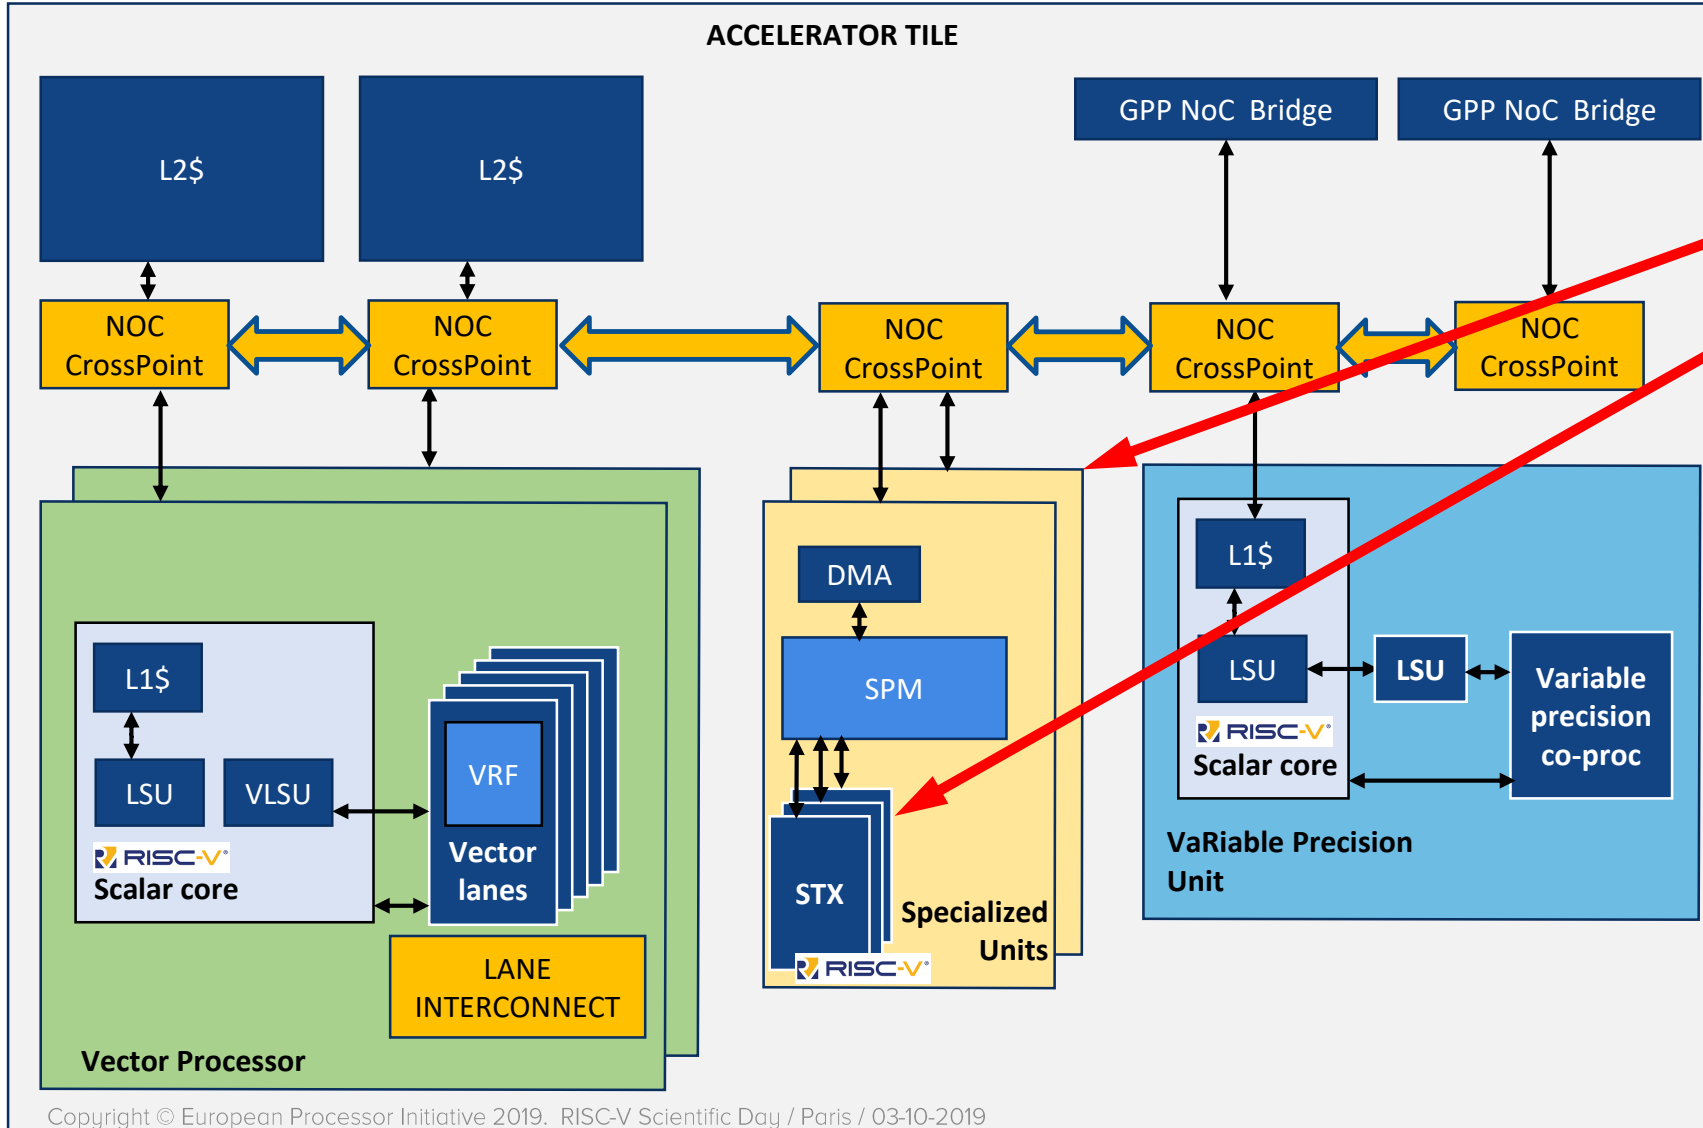
Up to 8 vector processors

Vector Lanes act as tightly coupled (ISA mapped) acceleration units to the scalar core in the vector processor

- Heavily pipelined
- RISC-V vector extension compliant

EPAC ARCHITECTURE VIEW

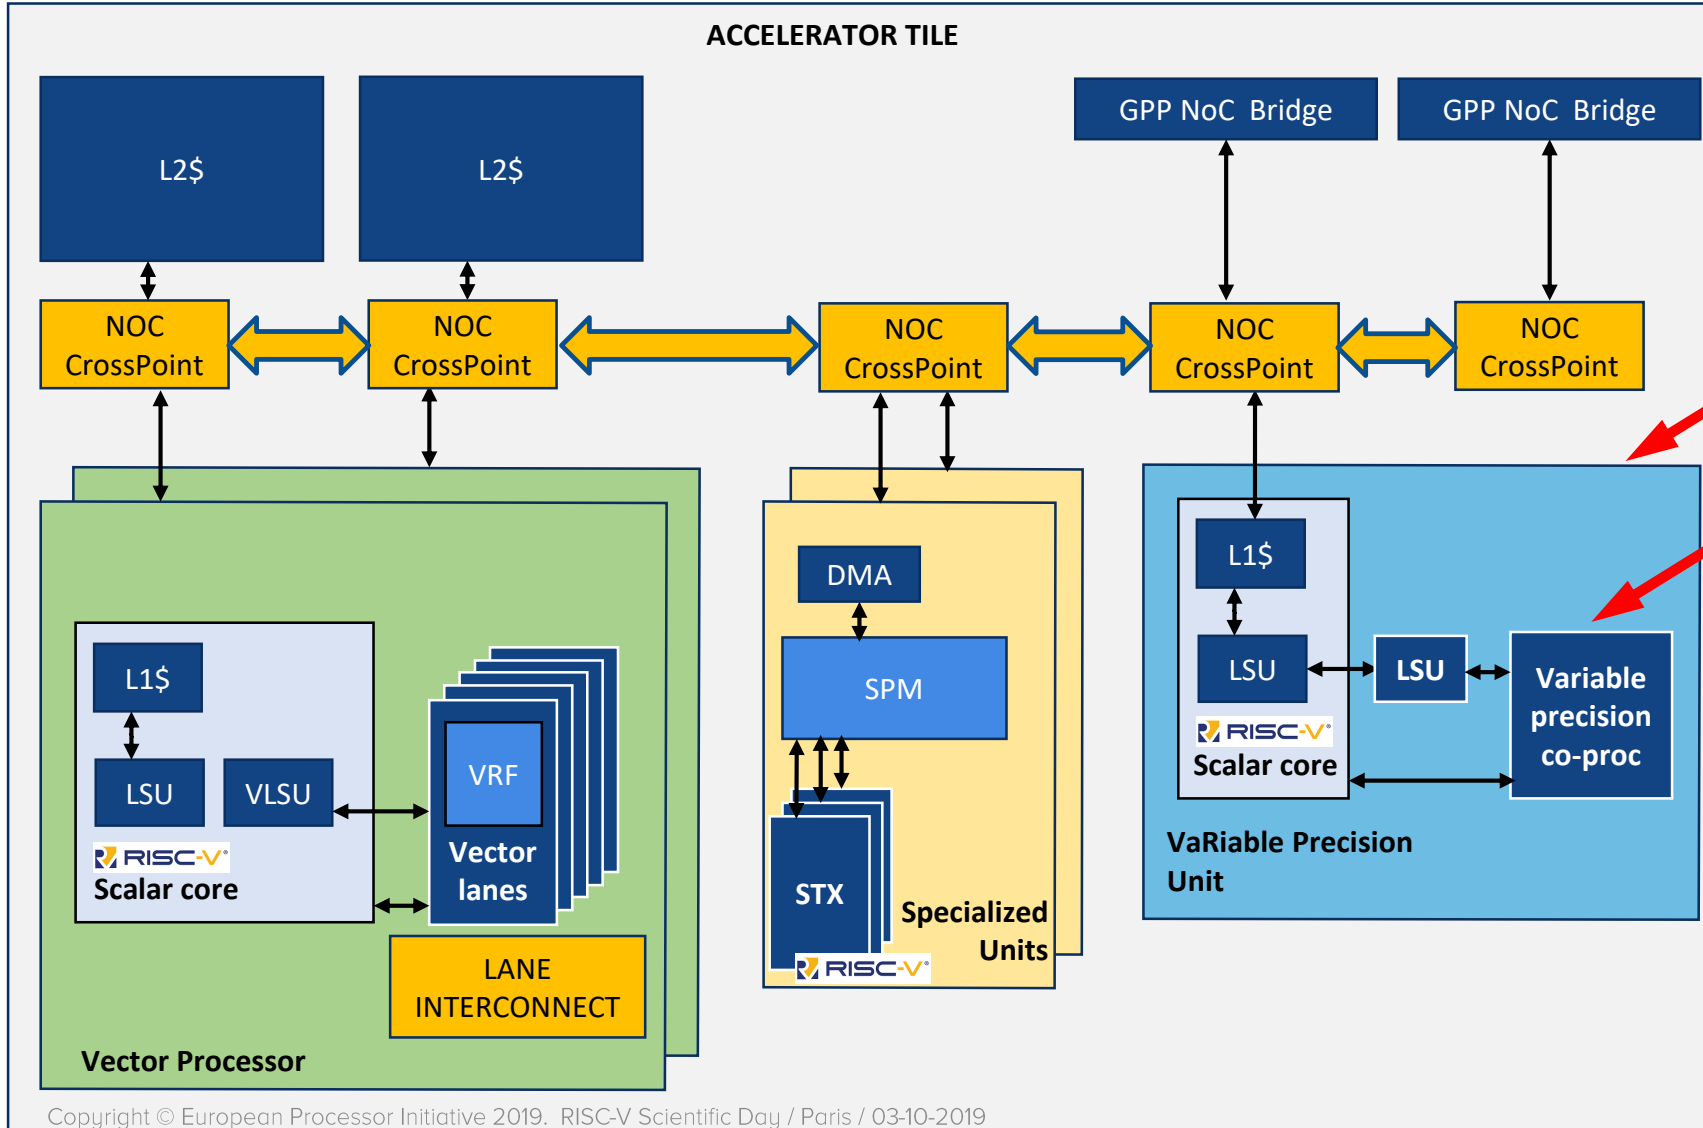
ACCELERATOR TILE



- Up to 8 Specialized Units
- The STX Units act as loosely coupled, memory mapped acceleration units to the scalar cores
- Fast single-cycle MACs in parallel

EPAC ARCHITECTURE VIEW

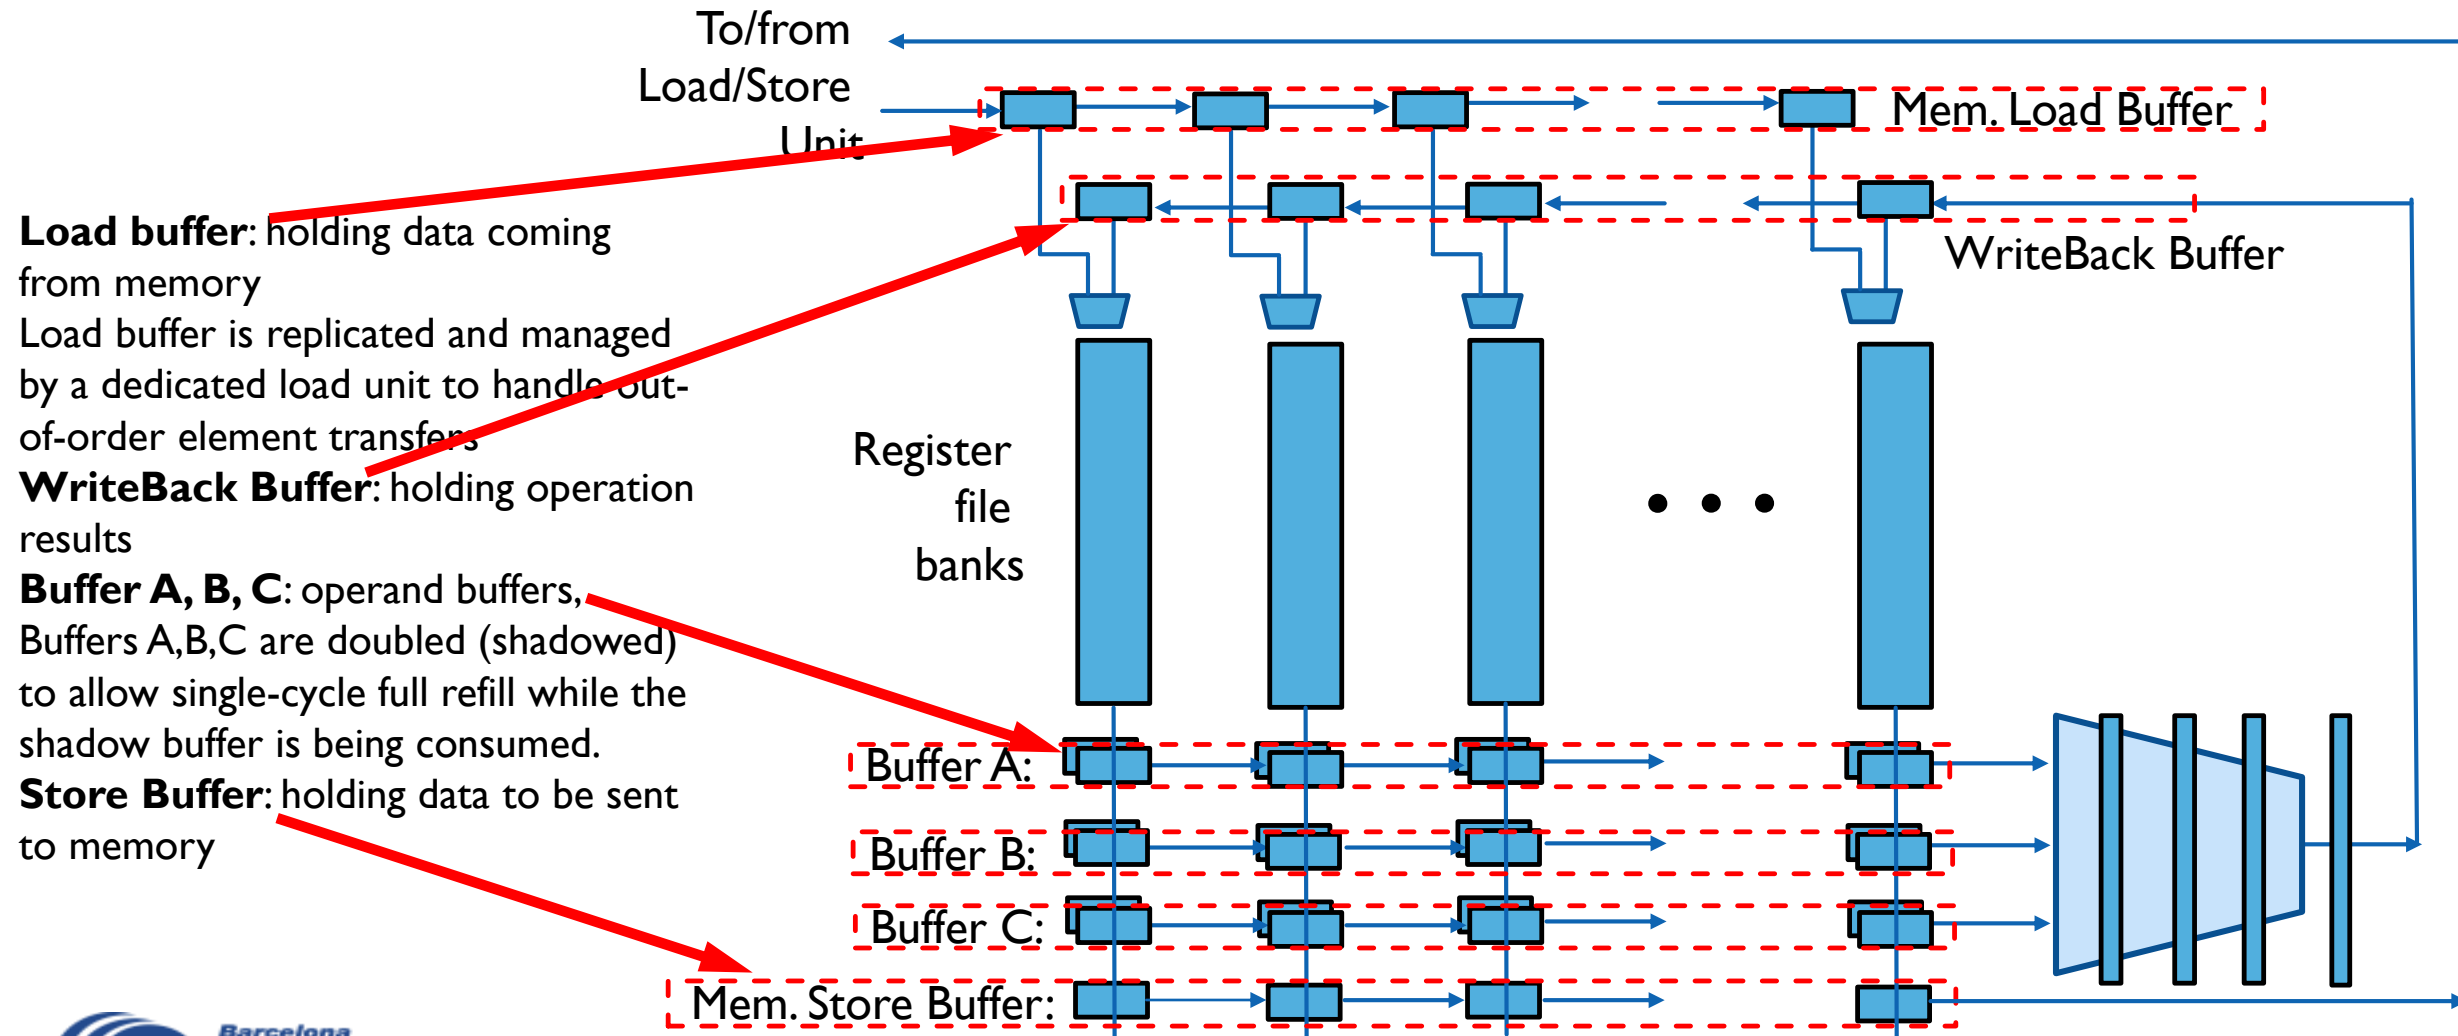
ACCELERATOR TILE



At least one Variable Precision unit

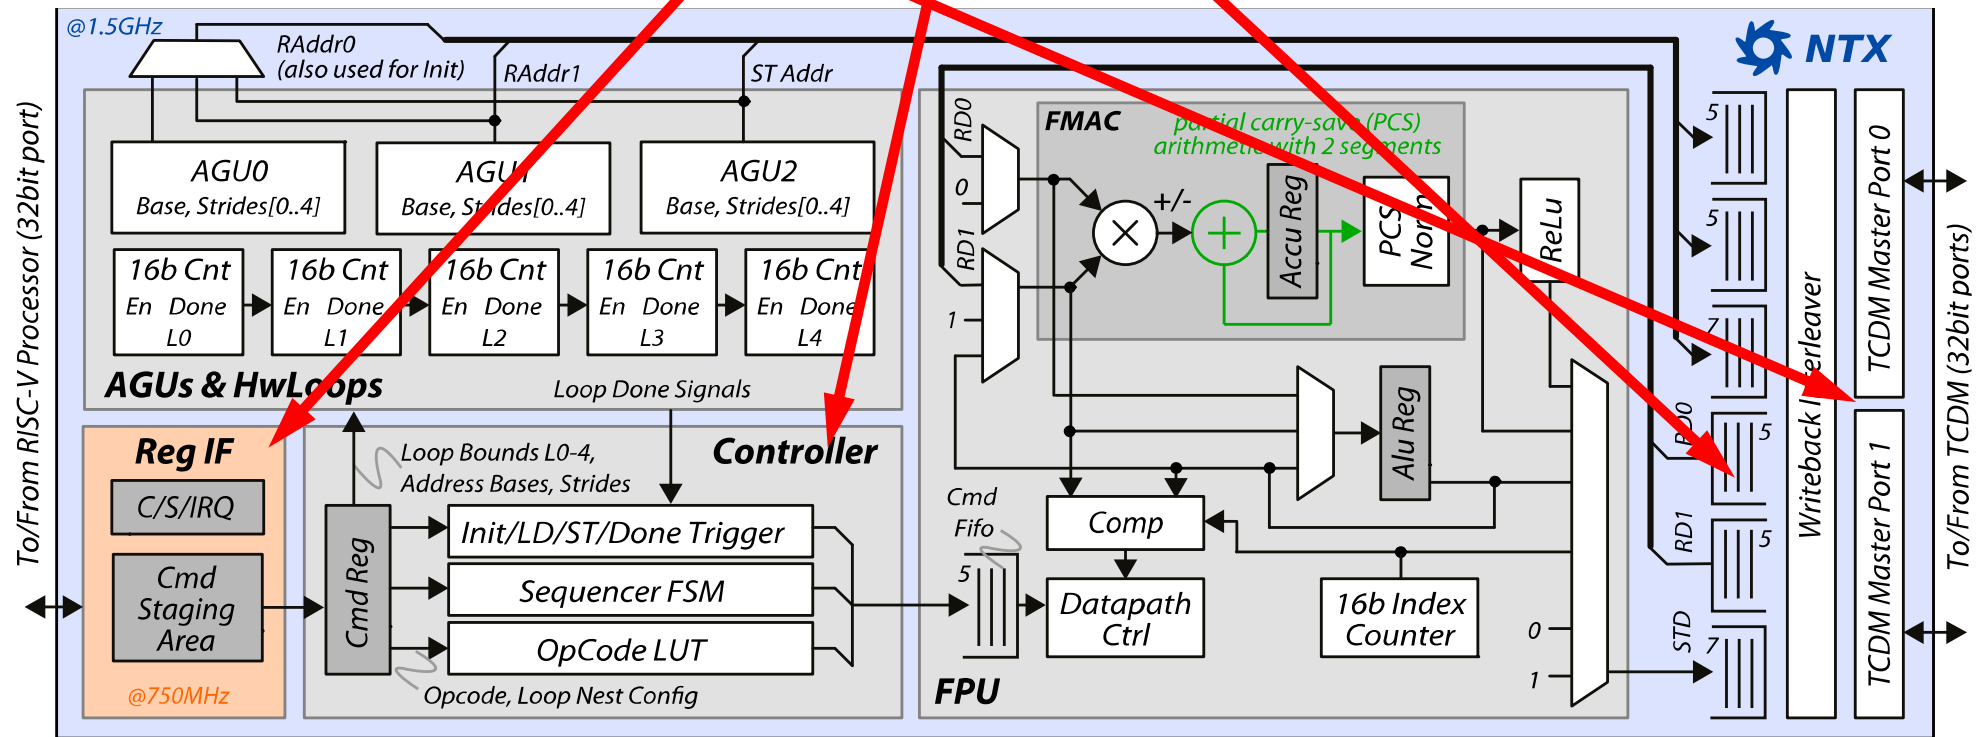
The variable precision co-processor act as execution unit extension to the scalar core

VECTOR LANE MICROARCHITECTURE (SIMPLIFIED VIEW)



STX COPROCESSOR MICROARCHITECTURE

- Processor configures operation via memory-mapped registers
- Controller issues AGU, HWL, and FPU micro-commands based on configuration
- Reads/writes data via 2 memory ports (2 operand and 1 writeback streams)
- FIFOs help buffering data path and memory latencies



- Copyright © European Processor Initiative.





THANK YOU