semidynamics

Tailor-Made High Performance RISCV64 IP

In Order Core

OOO Core

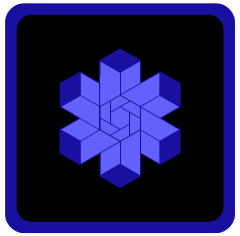

OOO Vector Unit

About Semidynamics

Semidynamics is a **European** supplier of RISC-V IP cores, specializing in **customization** of **high bandwidth high performance cores with vector units** for **tailored projects**

Experts in open core surgery

semidynamics

Our RISC-V Core IP Families

Atrevido

2, 3 or 4-wide **out-of-order**
RISCV64GCV AXI and CHI

Avispado

2-wide **in-order**
RISCV64GCV AXI and CHI

World's first, **fully customizable**, 64-bit RISC-V cores for ultra fast, big memory applications, optimized for a companion RISC-V **vector unit**

Unique tailor-made PPA solutions include customer's secret sauce for product differentiation and IP protection.

semidynamics

What's special about our RISC-V Cores?

We fully customize each core to the customer's precise application needs

We can include unique customer features in a few weeks

Support for OOO RISC-V Vector Unit

Gazzillion™ turbo-charges memory retrieval

Fastest cores on the market for moving big data

Process agnostic — already done 5nm

semidynamics

Avispado 223

- 64b RISC-V Core
- Mid-Range Performance
- **In-order execution**
- User, Supervisor and Machine privilege levels
- Hypervisor available Q4
- Linux-Ready memory subsystem
  - Memory Management Unit (MMU)
  - Supports SV39/48/57
  - Coherent caches with ECC, Parity
  - Hardware support for unaligned accesses
  - Hardware support for Atomics
- PMP Regions (0 to 16)
- AXI4 or AMBA CHI.B compliant interface
- Advanced Debug Capabilities
  - RISC-V debug spec compliant interface over JTAG
  - HW/SW Breakpoint support
- RISC-V Extensions supported:
  - **Vector, Crypto, Bit Manipulation, CMOs, Zifencei**
- Quad-Core Ready

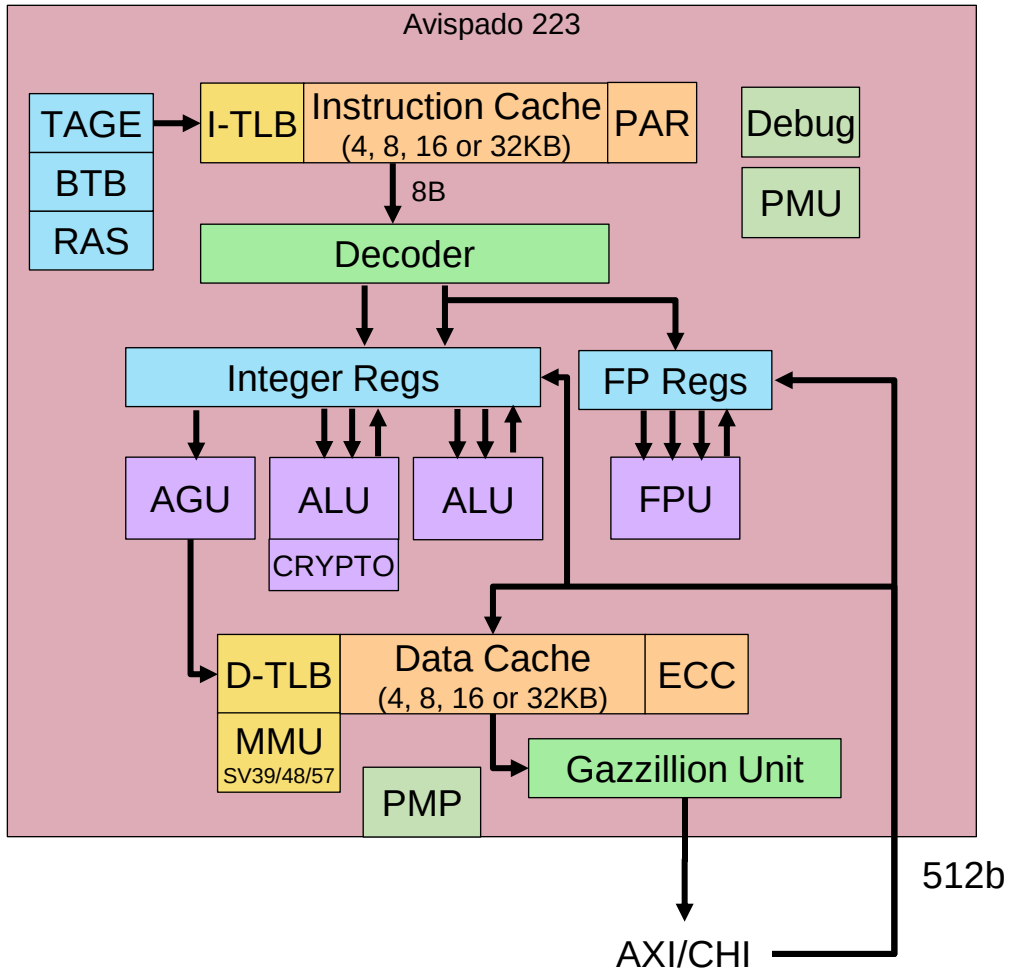

semidynamics

Atrevido 323

- 64b RISC-V Core
- High-End Performance
- **Out-of-order** execution
- **3-wide decode, rename, retire**
- User, Supervisor and Machine privilege levels
- Hypervisor available Q4
- Linux-Ready memory subsystem
  - Memory Management Unit (MMU)
  - Supports SV39/48/57
  - Coherent caches with ECC, Parity
  - Hardware support for unaligned accesses
  - Hardware support for Atomics
- PMP Regions (0 to 16)
- AXI4 or AMBA CHI.B compliant interface
- Advanced Debug Capabilities
  - RISC-V debug spec compliant interface over JTAG
  - HW/SW Breakpoint support
- RISC-V Extensions supported:
  - **Vector, Crypto, Bit Manipulation, CMOs, Zifencei**
- Quad-Core Ready

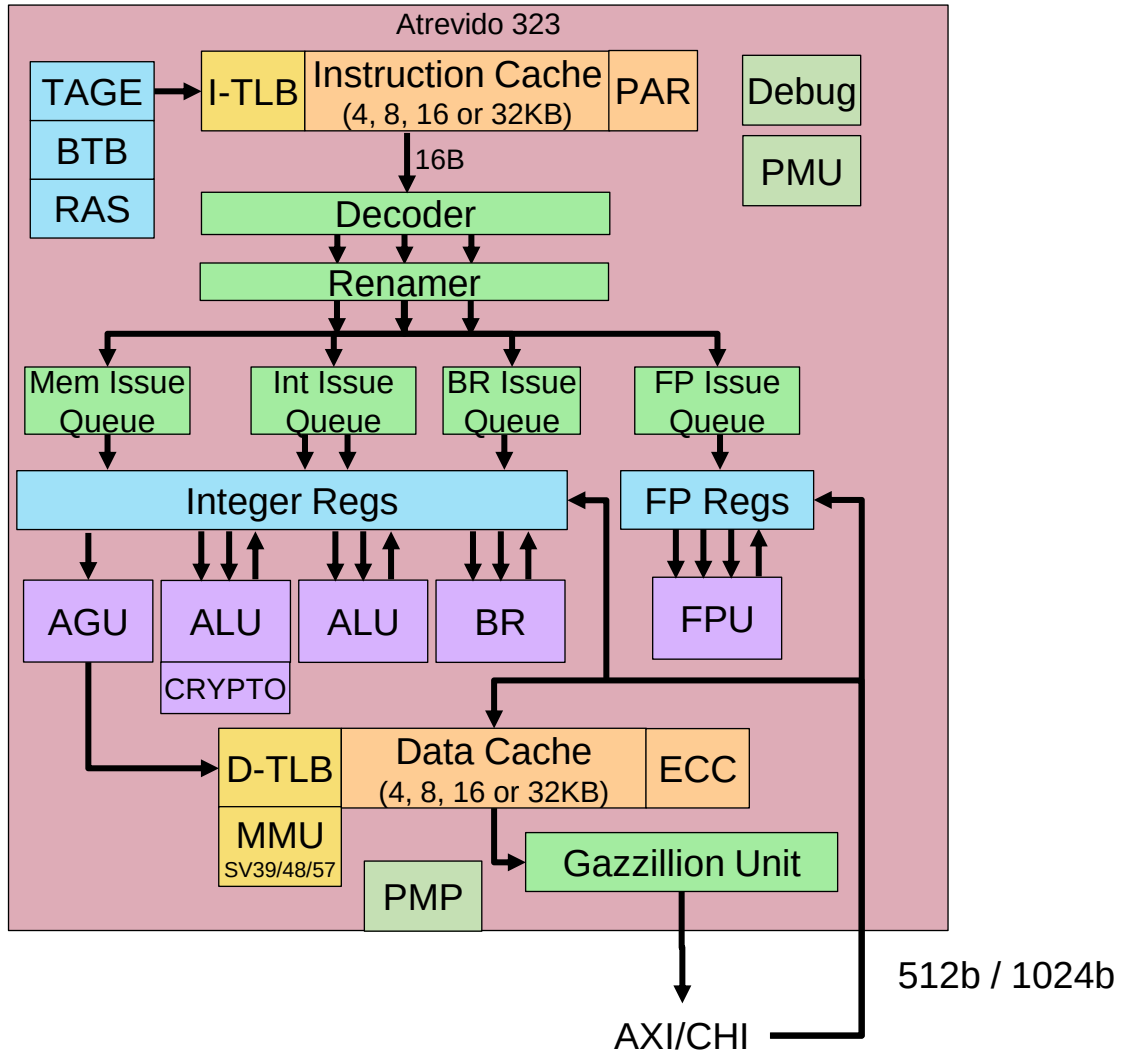

Atrevido 423

- 64b RISC-V Core
- High-End Performance
- **Out-of-order** execution
- **4-wide decode, rename, retire**
- User, Supervisor and Machine privilege levels
- Hypervisor available Q4
- Linux-Ready memory subsystem
  - Memory Management Unit (MMU)
  - Supports SV39/48/57
  - Coherent caches with ECC, Parity
  - Hardware support for unaligned accesses
  - Hardware support for Atomics
- PMP Regions (0 to 16)
- AXI4 or AMBA CHI.B compliant interface
- Advanced Debug Capabilities
  - RISC-V debug spec compliant interface over JTAG
  - HW/SW Breakpoint support
- RISC-V Extensions supported:
  - **Vector, Crypto, Bit Manipulation, CMOs, Zifencei**
- Quad-Core Ready

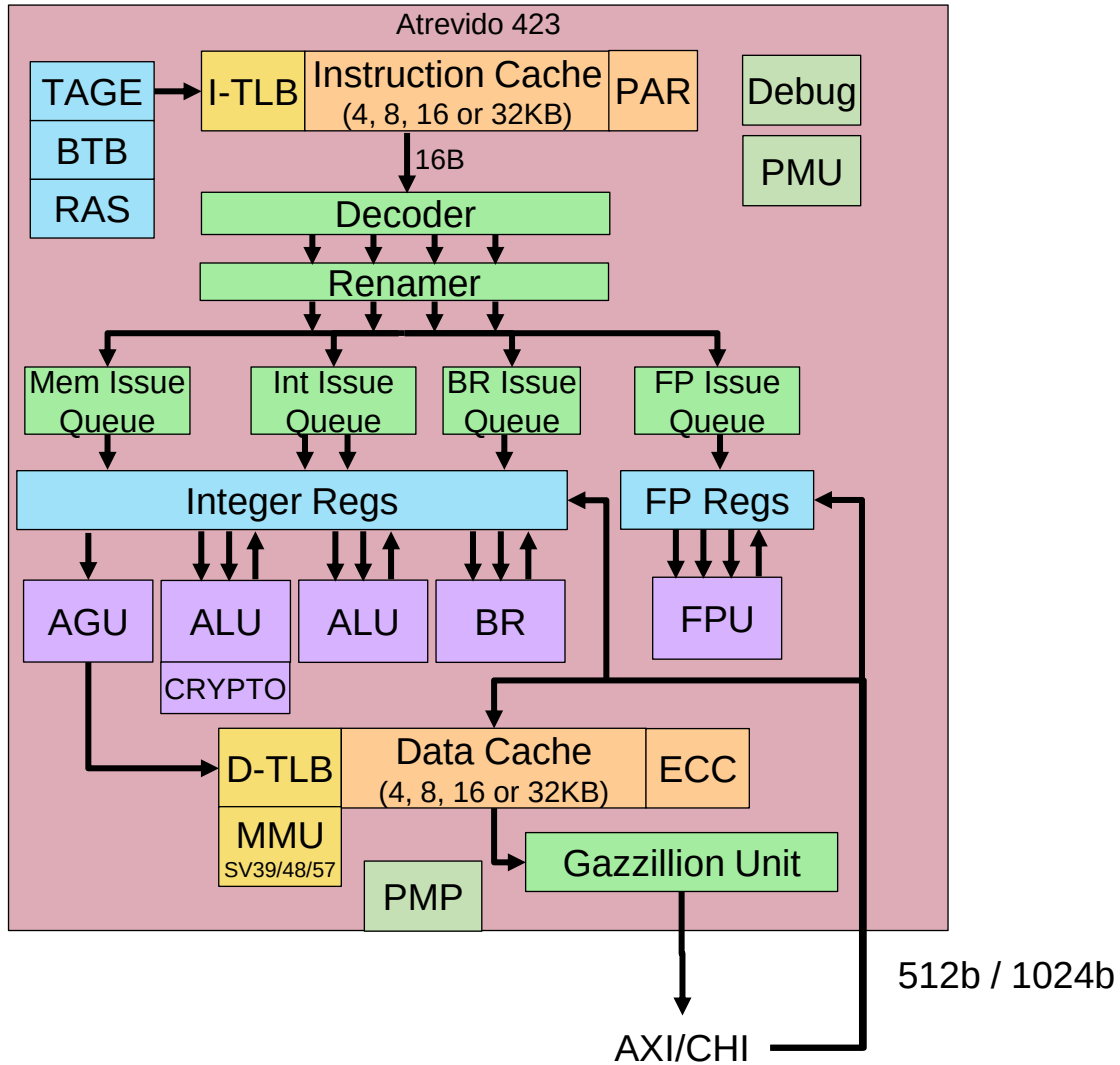

Atrevido 423 + V8 Vector Unit

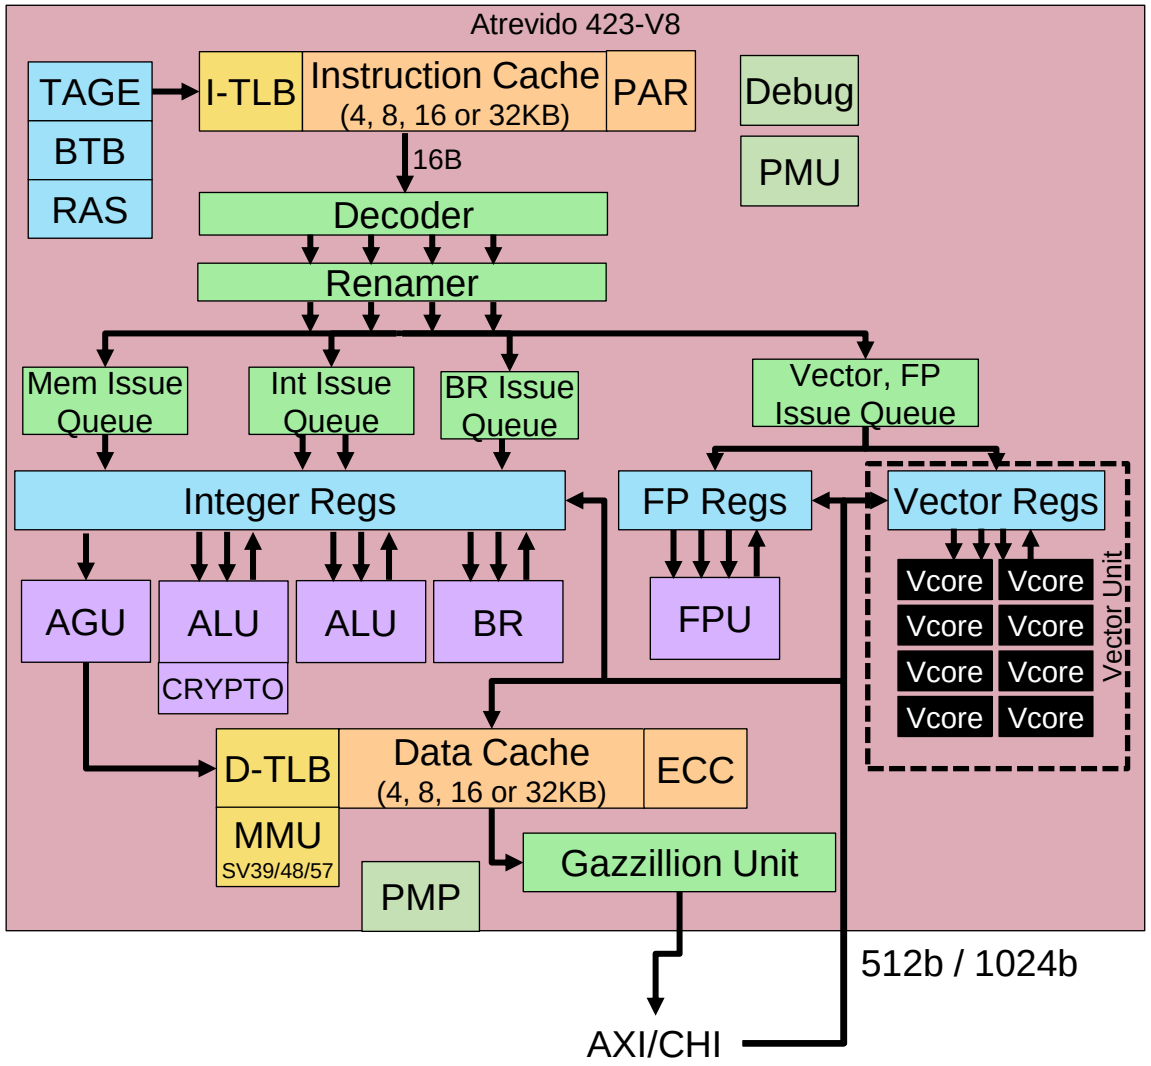

Vle (x5) -> v4
Vle (x6) -> v5
Vfma v4, v5, v1 → v6
Vsqrt v6 → v6
Vse v6 → (x7)

semidynamics

Atrevido 423 + V16 Vector Unit

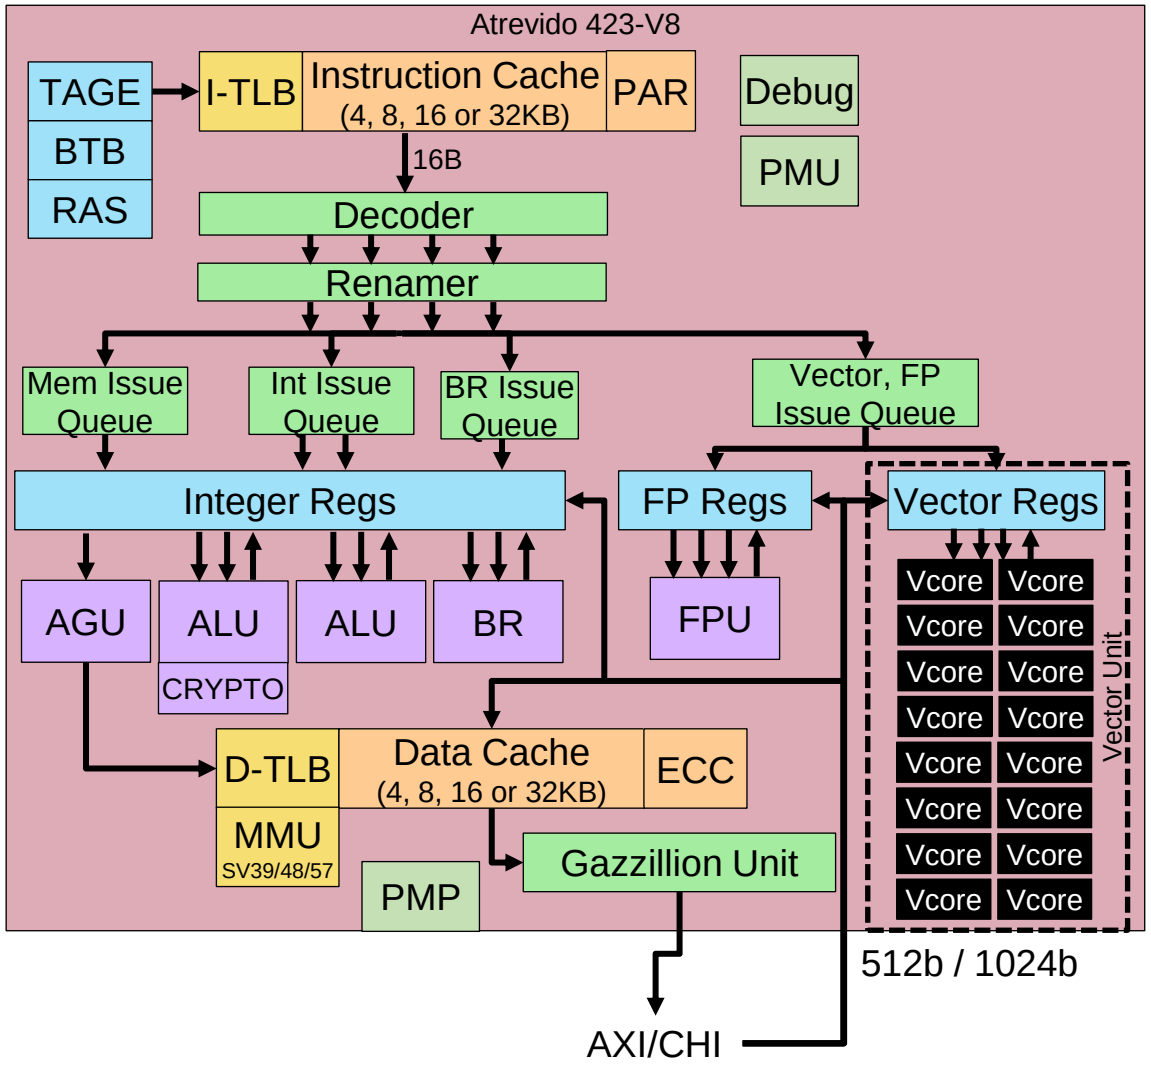

Vle (x5) -> v4
Vle (x6) -> v5
Vfma v4, v5, v1 → v6
Vsqrt v6 → v6
Vse v6 → (x7)

semidynamics

Gazzillion™ Technology

SERIAL

Traditional method of memory retrieval is stop and go. Request some piece of data from memory and wait doing nothing over several hundred clock cycles for it to come back.

PARALLEL

Gazzillion™ sends out up to 128 simultaneous requests for data from memory. Whichever data request comes back first is worked on and then the next one back so that core is always working. 128 of these streams in parallel bringing data back to the core really turbo charges performance.

Turbo charges memory retrieval
Tolerates memory latency like no other

semidynamics

Gazzillion compared to other CPUs...

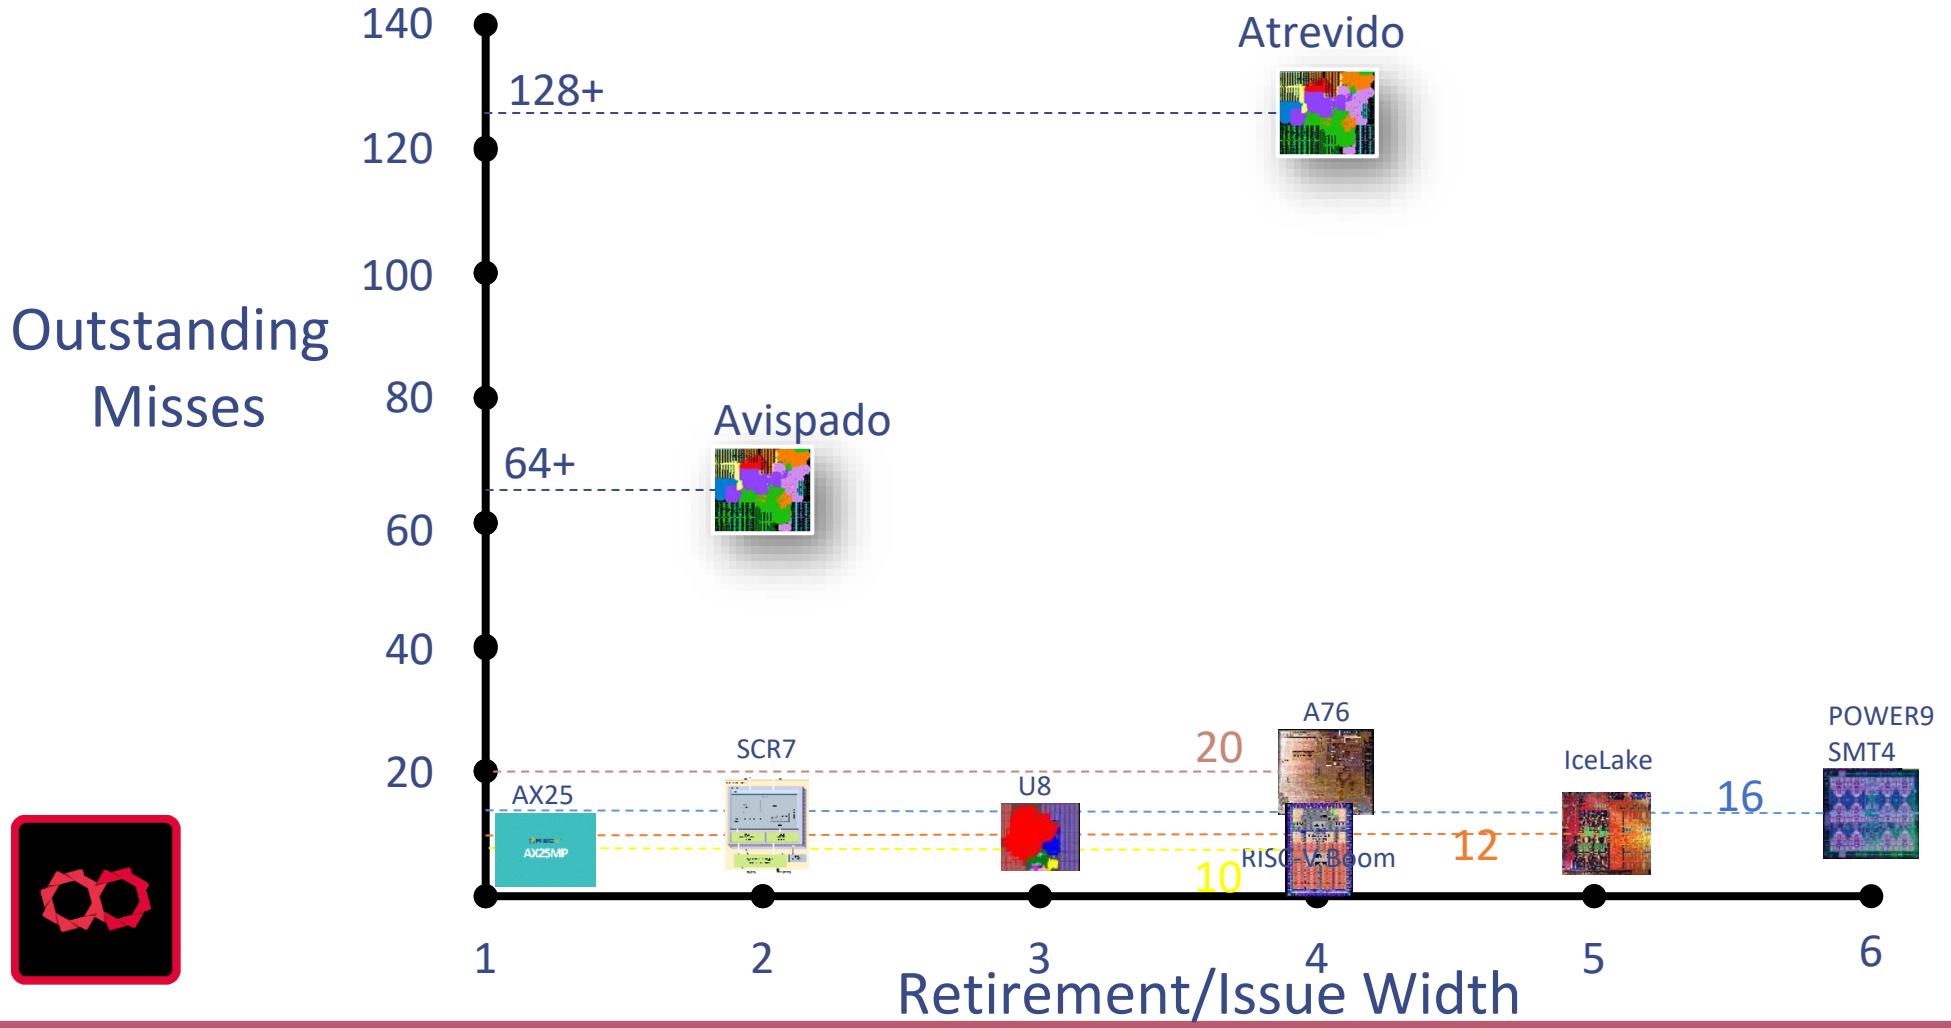

Vector + Gazzillion: A bandwidth rocket!

Can you find a core out there capable of streaming data at over 60 Bytes/cycle?
And from main DDR memory (not from your cache)? We don't think so 😊

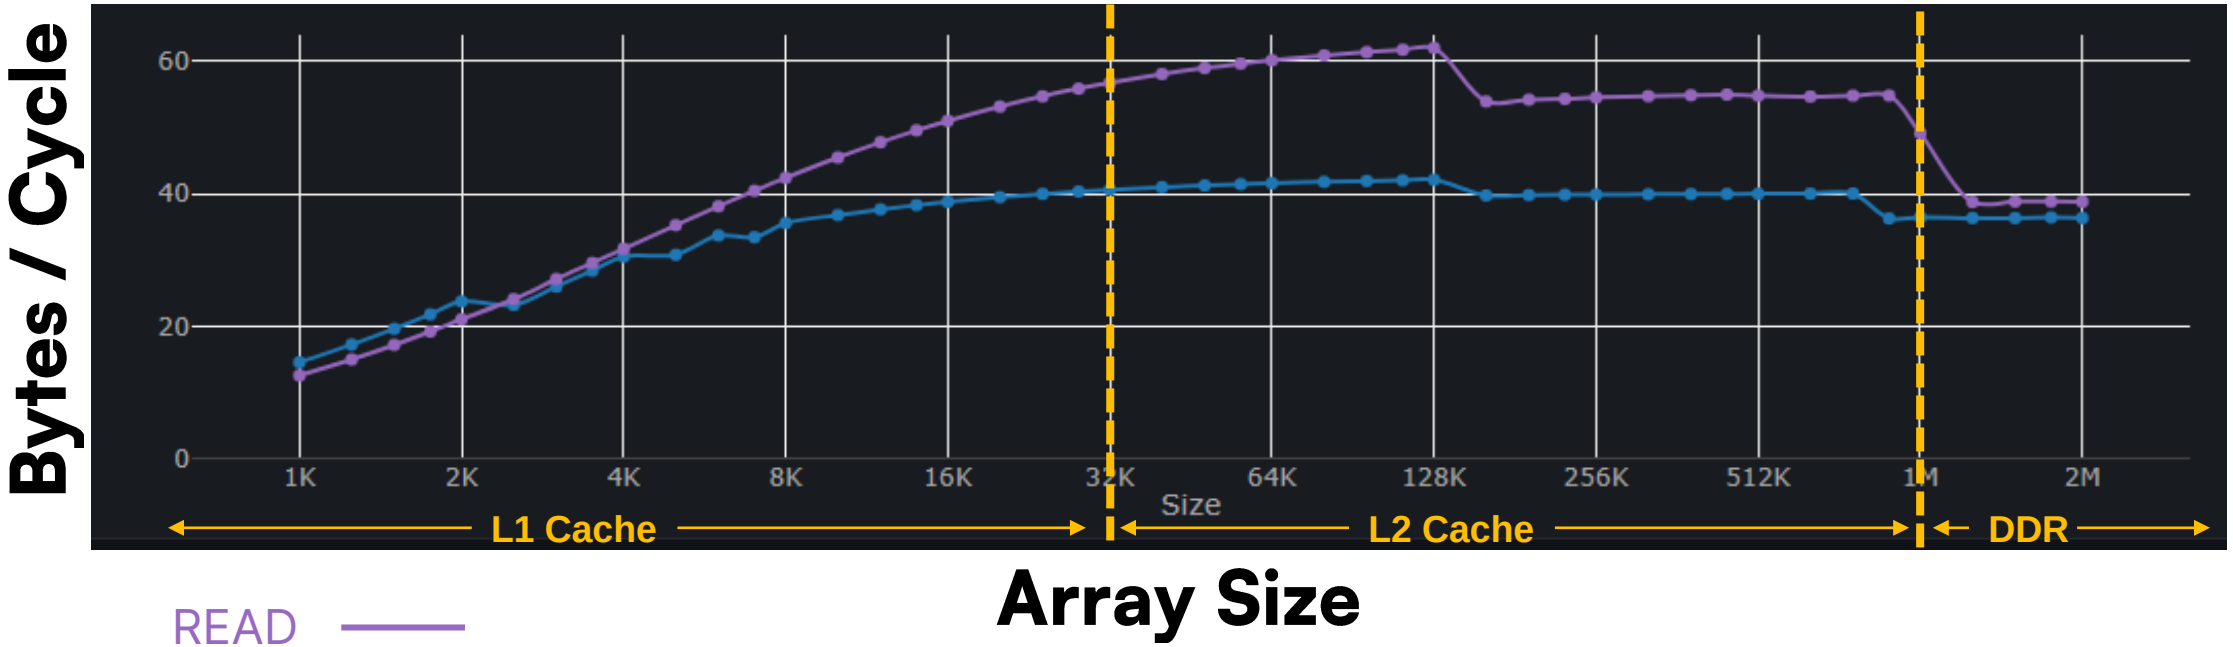

8 vector cores, 32X vector length

semidynamics

DGEMM on Atrevido 423 + V8

(FP64 matrix multiply)

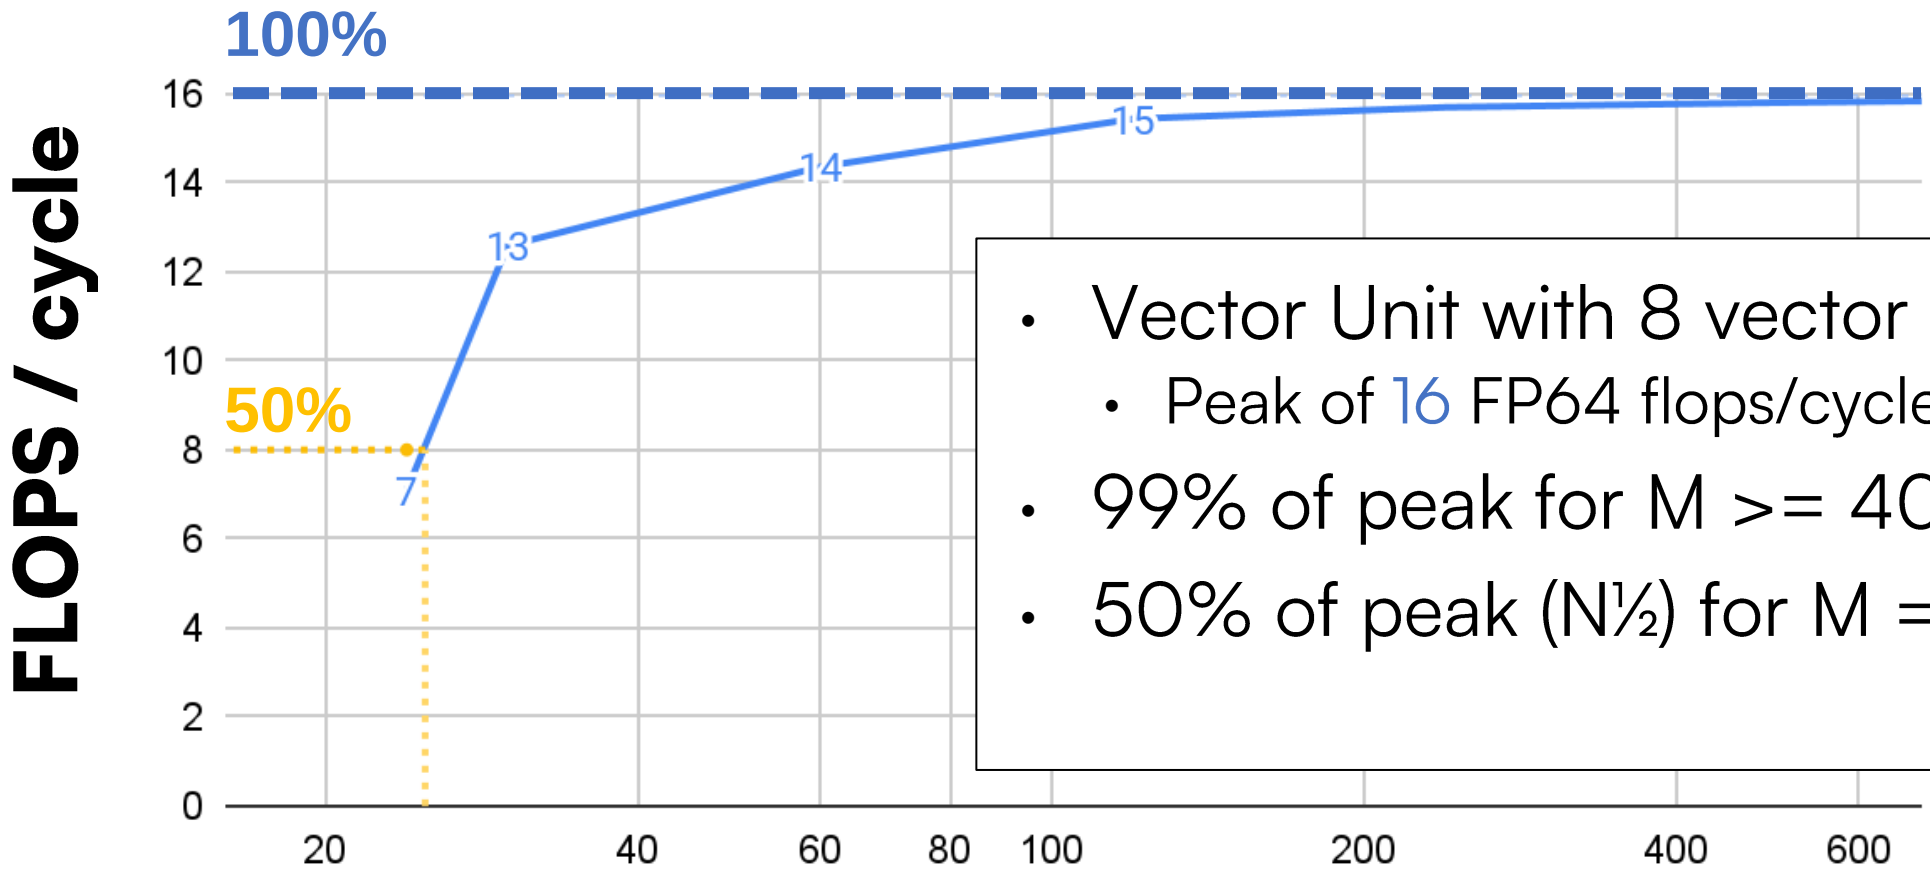

- Vector Unit with 8 vector cores
  - Peak of 16 FP64 flops/cycle
- 99% of peak for M >= 400
- 50% of peak (N½) for M = 24

semidynamics

Yolo on Atrevido 423 + V8

- YOLOv3-tiny:
  - 24 layers, 5.56 Gops/frame, ~9M params
  - Using SGEMM (FP32) for Matrix Multiplication

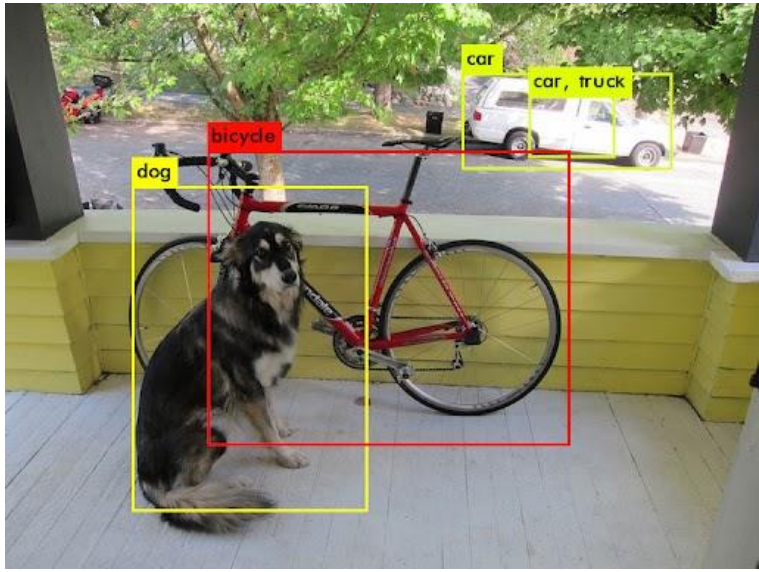

Platform	Vector/Cuda Cores	Frequency (Ghz)	FPS	FPS per 8 vector cores @ 1Ghz
Jetson TX2	256	1.30	19[1]	0.46
Jetson AGX Xavier	512	1.38	32[1]	0.36
GTX Titan X	3072	1.09	220[2]	0.53
Atrevido 423-V8	**8**	**1.00**	**0.84**	**0.84**

58% higher performance per vector core

semidynamics

[1] https://www.researchgate.net/publication/351347699_Real-Time_On_Board_Deep_Learning_Fault_Detection_for_Autonomous_UAV_Inspections
[2] https://pireddie.com/darknet/yolo/

Flexible and customizable Business Model

Customize IP

- AXI, CHI
- Cache Sizes
- Branch predictor
- Custom instructions
- RV32
- Small Core...

Evaluate

- Single Core
- Multi Core
- Vector Unit

License

- License Fee
- Royalties

Maintenance

- Bug Fixes
- Timing fixes
- Area Fixes

semidynamics

THANK YOU!

semidynamics