

The Future of Microprocessors: RISC-V

Driving Innovations™



RISC-V Day Vietnam
September 18, 2020

Thang Tran, Ph.D.
Principal Engineer

Agenda

❖ Background

- X86/ARM®
- RISC-V

❖ RISC-V advantages

- In academia
- ACE

❖ Vector processor

- Design
- Advantages

❖ Summary

X86 and the PCs/Laptops

❖ **Microprocessor is always proliferated by the applications that change people lives**

- X86 is the microprocessor for work stations, PCs, laptops
- Intel® became the largest semiconductor in the world

❖ **X86 and Moore's law**

- 1000 of inventions from X86 microprocessors
- Many terminologies became buzz-words in microprocessor: superscalar, OOO, ROB, register renaming, reservation stations, central window, check-point repair, branch prediction, WAW, WAR, RAW, ...
- The race is the fastest, biggest and baddest microprocessor

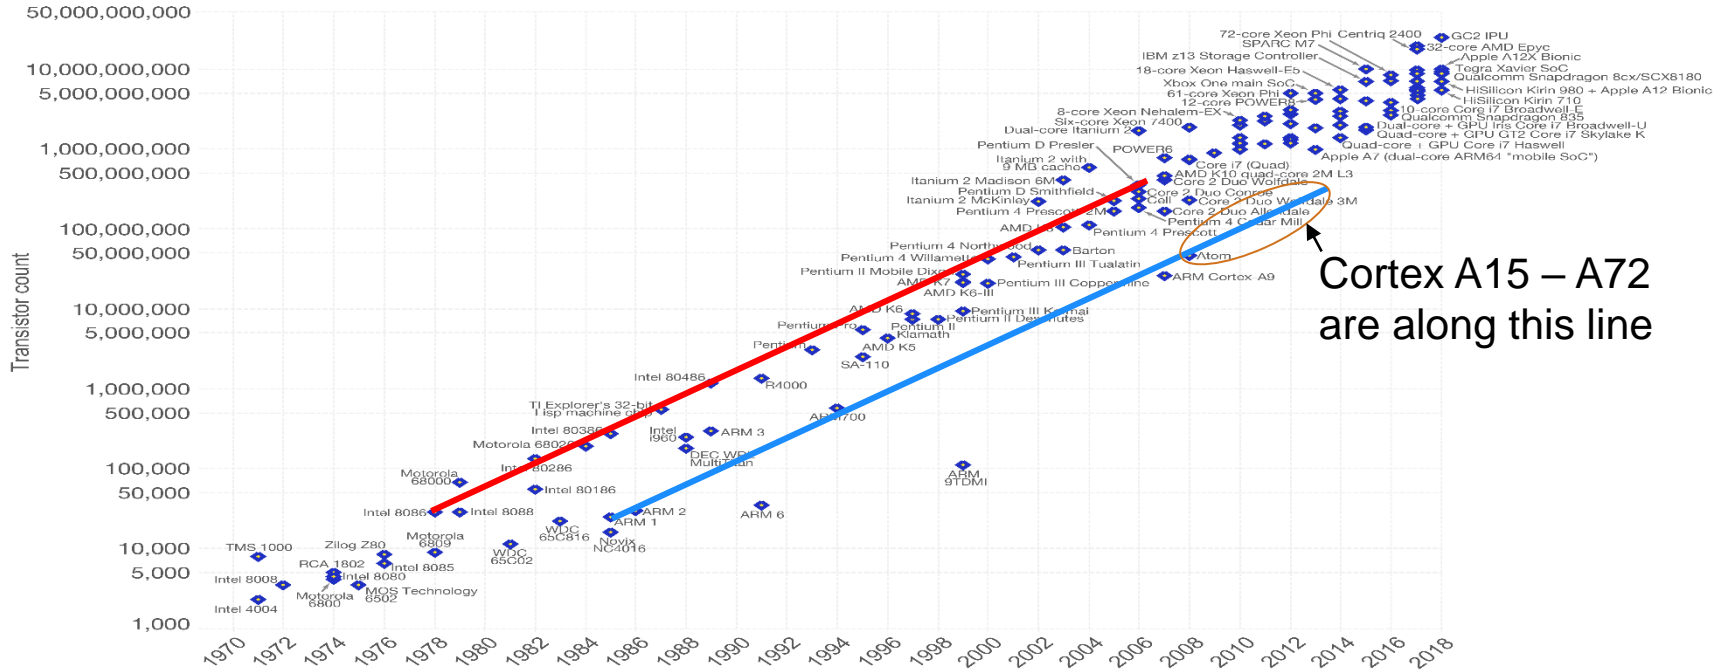
❖ **But the PCs and Laptops became saturated**

- Power became the big issue

Moore's Law

Moore's Law – The number of transistors on integrated circuit chips (1971-2018)

Moore's law describes the empirical regularity that the number of transistors on integrated circuits doubles approximately every two years. This advancement is important as other aspects of technological progress – such as processing speed or the price of electronic products – are linked to Moore's law.



Data source: Wikipedia (https://en.wikipedia.org/wiki/Transistor_count)
The data visualization is available at [OurWorldinData.org](https://www.ourworldindata.org). There you find more visualizations and research on this topic.

Licensed under CC-BY-SA by the author Max Roser.



ARM and the Smart Phones

- ❖ **The major shift in technology for the mass is the smart phones**
 - Power is important in hand-hold devices
 - The shift from chips to IPs, fabless companies, different business model
 - ARM is much more cheaper than X86
- ❖ **ARM and Moore's law**
 - The same Moore's line is shifted to the right by 10-15 years
 - New terminologies in microprocessor: clock gating, low power technique, dynamic frequency/voltage scaling, power domains, power management
 - The race is the microprocessor with enough performance and long battery life
- ❖ **But the smart phones become saturated**
 - What is next?

Smart Home/Auto, IoT, AI and ML

- ❖ **No longer a single application, the shift is to many applications**
 - Every other aspects of life: home, auto, medical, surveillance, social media, business, ... In general, the basic is IoT, AI and ML
 - Money: something much cheaper than ARM
 - There are so many applications, the critical factor is extensible and configurable
 - RISC-V is enough to learn and open to everyone. RISC-V is the perfect fit
- ❖ **RISC-V and Moore's law**
 - Perhaps the same Moore's line is shifted to the right by 30 years?
 - New terminologies in microprocessor: PPA, configurable, extensible, custom instruction/extension, early/late ALU
 - The race is custom SOC design with the best PPA

RISC-V & RVV Background

- ❖ **An open processor architecture started by UC Berkeley**
 - Compact, modular, extensible
- ❖ **RISC-V International: (formerly RISC-V Foundation):**
 - RISC-V Foundation: formed in 2015 to govern its growth
 - 400+ members including industry and research institute/university
 - 2019 RISC-V Summit has over 1000+ attendees
 - Widely adapted in the world
 - Many universities adopt RISC-V as the basic ISA for computer architecture class;
I am in the process of changing to RISC-V in my next computer architecture class at Santa Clara University

RISC-V Microprocessor

- ❖ **RISC-V ISA has many considerations for implementation**
 - Operands are in fixed bit locations
 - No status bits (carry, negative, zero, ...) which simplifies data dependency
 - No predication instructions, no delayed branch instructions
 - Simple memory/register model and load/store instructions; no load/store double or multiple
- ❖ **Unlike X86 or ARM, RISC-V is developed by many companies**
 - Andes, SiFive, Esparanto, Western Digital, Cudasip, ...
 - Advanced microprocessors, superscalar and/or OOO, have been designed by more companies
 - 64-bit microprocessors are available now, much faster pace in comparison
 - SIMD, DSP, FPU, and Vector Processor are also available

RISC-V Superscalar Microprocessor

- ❖ Common microprocessor design is dual ALU
- ❖ In-order, low-power, simple design
- ❖ Often implemented in dual in-order superscalar
- ❖ This technique is referred as “load latency tolerance” to reduce the “load-to-use” penalty

Time	1	2	3	4	5	6	7	8	9	10
Fetch1	I7	I8	I9	I10	I11	I12	I13	I14	I15	I16
Fetch 2	I6	I7	I8	I9	I10	I11	I12	I13	I14	I15
Decode	I5	I6	I7	I8	I9	I10	I11	I12	I13	I14
Reg Read	I4	I5	I6	I7	I8	I9	I10	I11	I12	I13
ALU	Ld			I4	I5	I6	I7	I8	I9	I10
DC1	I2	Ld			I4	I5	I6	I7	I8	I9
DC2	I1	I2	Ld			I4	I5	I6	I7	I8
Write	I0	I1	I2	Ld			I4	I5	I6	I7

2-cycle bubbles

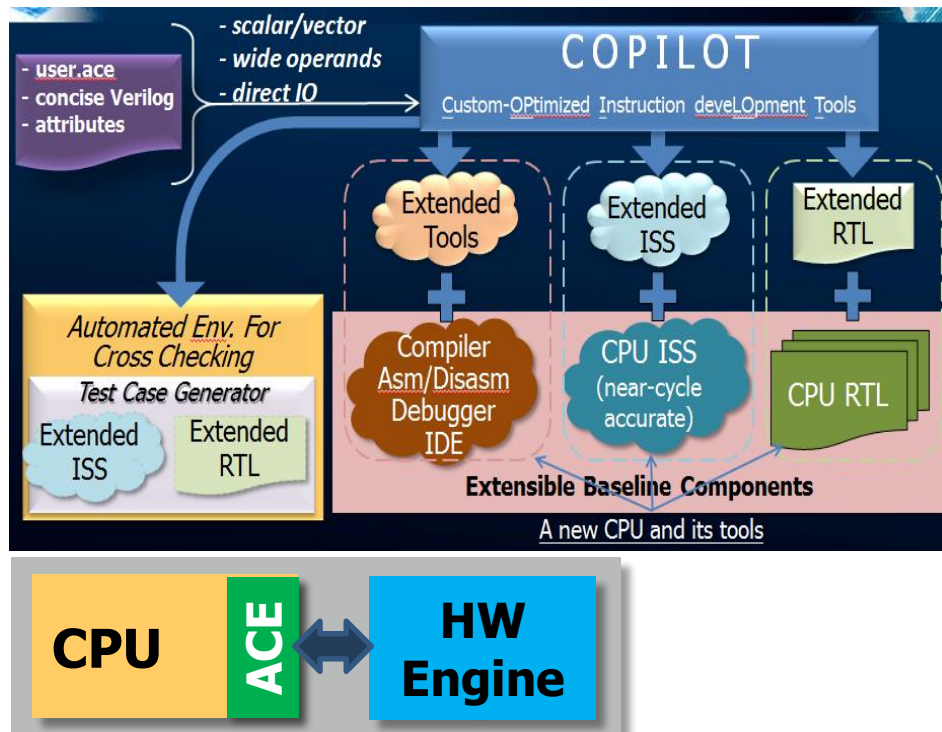
Time	1	2	3	4	5	6	7	8	9	10
Fetch1	I7	I8	I9	I10	I11	I12	I13	I14	I15	I16
Fetch 2	I6	I7	I8	I9	I10	I11	I12	I13	I14	I15
Decode	I5	I6	I7	I8	I9	I10	I11	I12	I13	I14
Reg Read	I4	I5	I6	I7	I8	I9	I10	I11	I12	I13
ALU	Ld	I4	I5	I6	I7	I8	I9	I10	I11	I12
DC1	I2	Ld	I4	I5	I6	I7	I8	I9	I10	I11
DC2/ALU	I1	I2	Ld	I4	I5	I6	I7	I8	I9	I10
Write	I0	I1	I2	Ld	I4	I5	I6	I7	I8	I9

No stall, I4 uses second ALU

Andes Custom Extension (ACE)

❖ ACE instructions

- ❑ Can be developed by customer
- ❑ Hardware generated automatically including data dependency
- ❑ Added to compiler & debugger
- ❑ Tied into the Andes microprocessor pipeline
- ❑ Support ACR (ACE Register), ACM (ACE Memory), ACP (ACE Port) with arbitrary width and number
- ❑ Including custom load/store instruction
- ❑ Provide proprietary and differentiation for customer



RISC-V Extension

❖ RISC-V must grow up fast, must have many extension to the basic ISA:

- 16/32/64/128-bit instruction sets; C: compress, 16-bit
- M: Multiply-divide, A: Atomic, F: single precision FP, D: double precision FP, Q: quad precision FP
- Exotic extensions – P: SIMD, V: vector, J: Java, B: bit manipulation, ...
- Extensions are defined in a foundation task group with experts from all different companies all over the world

RISC-V Vector Extension

- ❖ **RISC-V Vector Extension (Andes is the first to market)**
 - Scalable vector instruction set, agnostic vector length
 - Scalable data sizes which include 2x and 4x data expansion arithmetic
 - Over 300+ vector instructions, including load/store, integer, fixed-point/ floating-point operations
- ❖ **Vector Applications (important for AI and ML)**
 - Deep Learning
 - Multimedia Processing (compress., graphics, image proc.)
 - Standard benchmark kernels (Matrix Multiply, FFT, Convolution, Sort)
 - Lossy Compression (JPEG, MPEG video)
 - Cryptography (RSA, DES/IDEA, SHA/MD5)
 - Operating systems/Networking (memcpy, memset)

Vector Advantages

- ❖ Register width = 512b
- ❖ 64 elements of 8b
- ❖ An instruction can specify 64-512 elements operation
- ❖ Programmable and configurable

Vector length: $VLEN * LMUL$

- $LMUL=1 \rightarrow 512$ bits
- $LMUL=8 \rightarrow 4096$ bits
- $V0 \sim V7$ group 1
- $V8 \sim V15$ group 2
- $V16 \sim V23$ group 3
- $V24 \sim V31$ group 4

Cycles:	1	2	3	4	5	6	7	8	9	10
vmul v0, v8, v16	X1	<u>X2</u>								
v1, v9, v17		X1	<u>X2</u>							
.....					
v7, v15, v23								X1	<u>X2</u>	
vadd v24, v0, v8			X1							
v25, v1, v9				X1						
.....			

Xn: execution \xrightarrow{LMUL} Chaining $\xrightarrow{\hspace{1cm}}$

Overview of NX27V

❖ AndeStar V5 architecture:

- RV64GCN+ Andes V5 Extensions
- RV Vector extension (RVV)

❖ 5-stage pipeline, single-issue

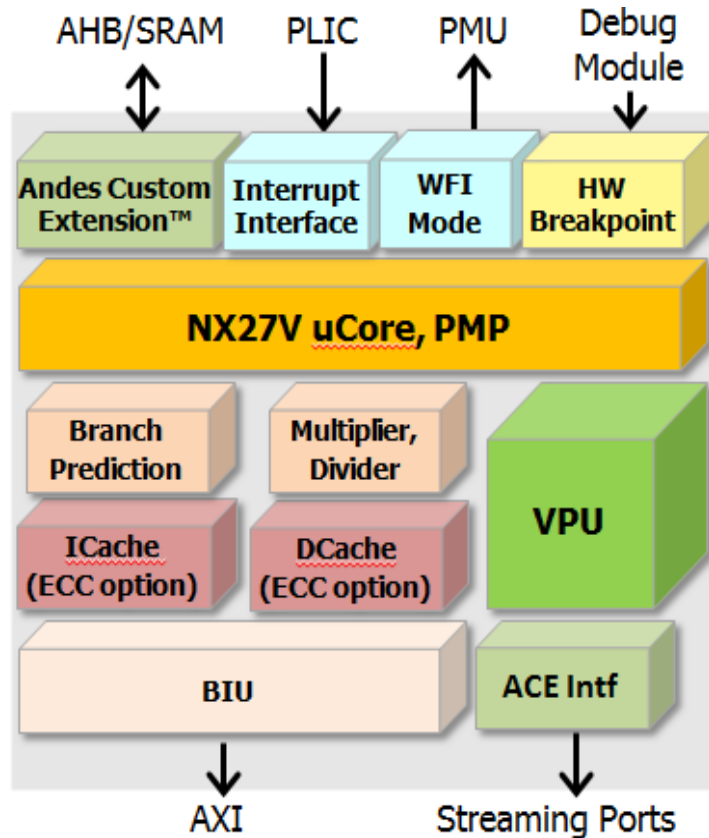
- Optional branch prediction

❖ I/D caches

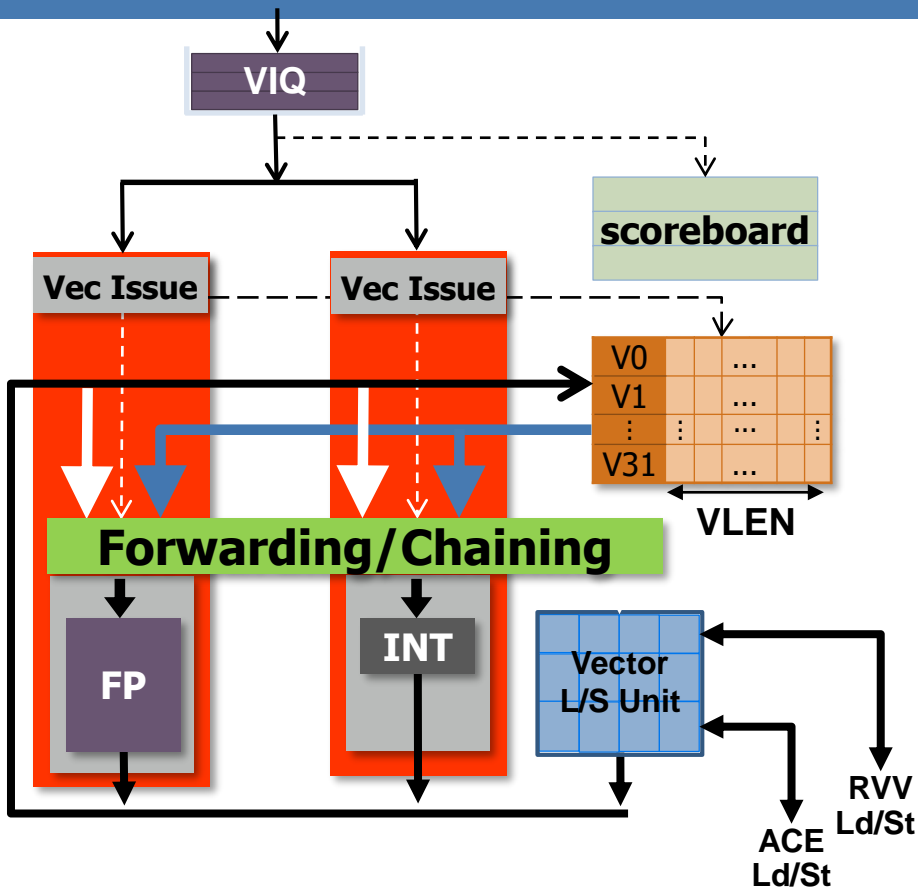
- Caches: 8KB to 64KB
- I\$/D\$ prefetch
- HW unaligned load/store accesses
- 16 non-blocking outstanding data accesses

❖ Wide data paths to feed VPU:

- Cached and uncached RVV load/stores
- Streaming Ports for ACE loads/stores

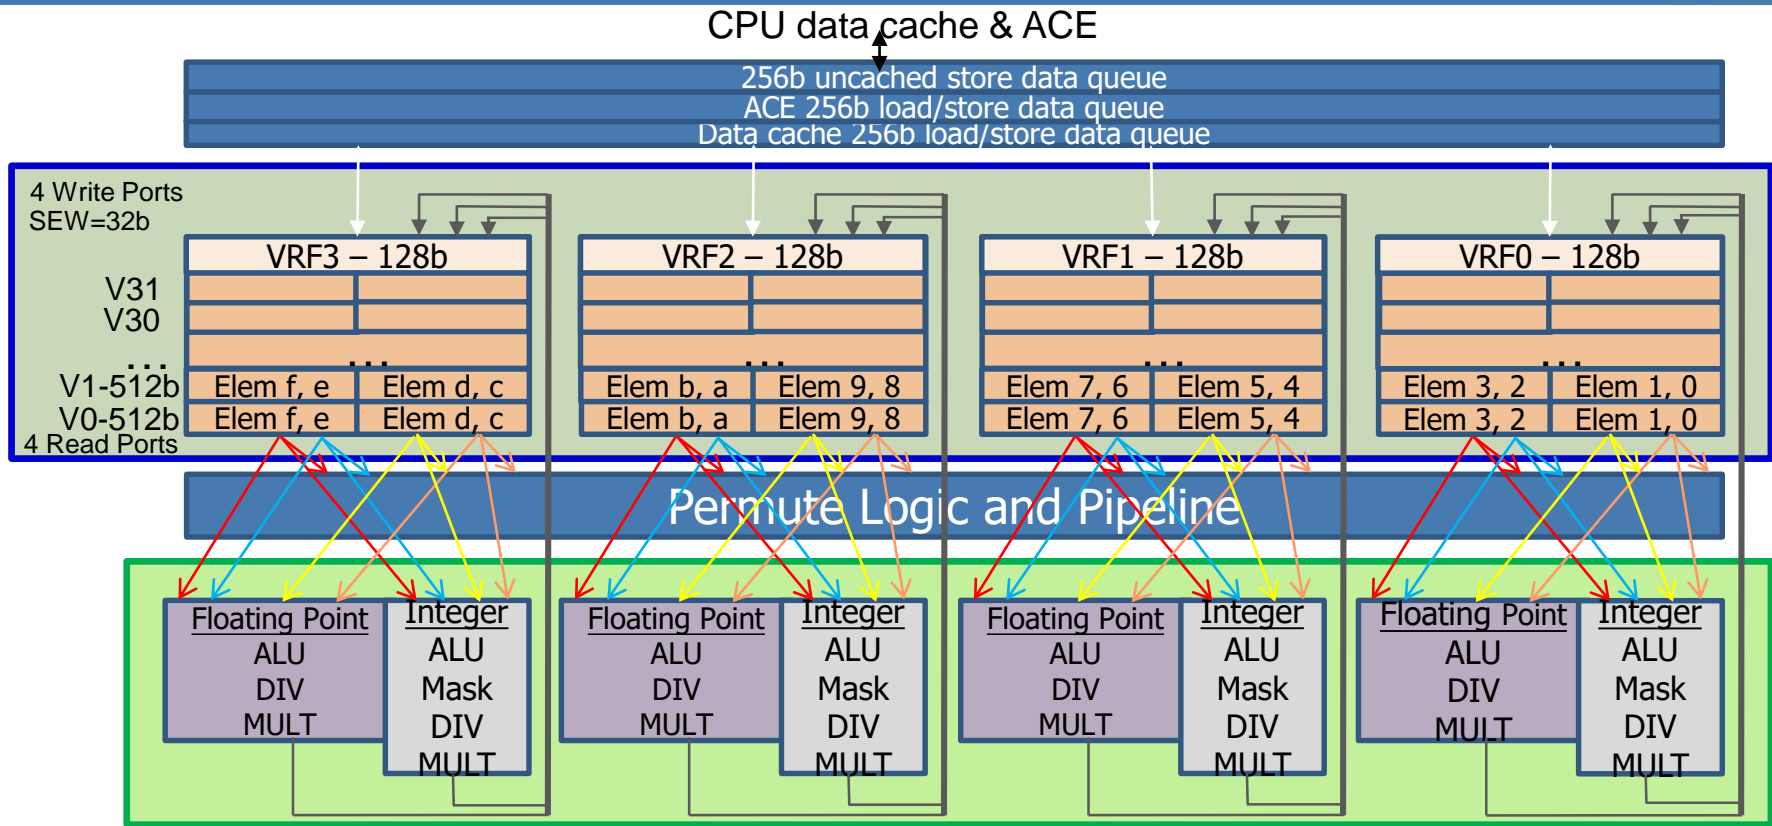


NX27V VPU Overview



- ❖ Supporting the latest RVV spec
- ❖ Data formats:
 - Standard: int8-int64, fp16-fp64
 - Andes-extended: bfloat16 and int4
- ❖ VLEN & SIMD width: 128, 256, 512
- ❖ Vector compute instructions:
 - Multiple Functional Units operating independently (OOO)
 - Chainable, and most fully pipelined
 - 4 SIMD and VRF lanes
- ❖ Independent memory access paths thru RVV load/store and ACE load/store

Vector Processor Data Path



VPU Micro-Architecture

❖ 9 independent vector functional units

- Fixed point (3): ALU, MAC, DIV
- Floating point (3): FMAC, FDIV, FMIS
- Mask
- Permutation
- Load-store: handling 6 independent load/store instructions at one time


❖ In-order issue, out-of-order execution and completion

- Issue up to 8 micro-ops per cycle
- Each execution queue is configured with up to 64 micro-ops per functional unit
- Innovative scoreboard scheme for data dependency
- Innovative yet simple techniques to control reading and writing of vector register file data (configurable)

• Scalar floating points:

- The FP Register File shares the FP functional unit
- Scalar FP instructions have its own 3 execution queues

Vector Processor Pipeline Example









NX27V Performance Gain

Functions	Speedup ¹
F32 basic mathematical functions	19X
RGB CNN functions	18X
Depthwise CNN functions	23X
Pointwise CNN functions	21X
Relu CNN functions	69X
F32 filtering functions	19X
Q7 filtering functions	39X
F32 32x32x32 matrix multiplication	57X

¹Compared to pure C scalar code compiled with high optimization; both vector and scalar code ran on the NX27V FPGA with 512-bit VLEN, 256-bit bus.

Andes NX27V vs. Competitions

	Andes NX27V	CAxx	CMxx
Architecture	RVV/Andes VPU	Popular SIMD	Hxxx
Vector registers	32	32	8
Vector Length	Up to 512b 	128b	128b
SIMD width	Up to 512b/cycle 	128b/cycle	64b/cycle
LMUL	Yes 	None	None
Chaining	Yes 	Not applicable	Yes
Custom extension	ACE 	No	No
Streaming Ports	Yes 	No	No

Core/System Performance Comparison

CPU	A64FX*	NX27V
Technology	7	7
Core (GFLOPS)	~56	96
Core Peak Perf 16b (GOPS)	230	320
Core Peak Perf 16b (GFLOPS)	-	128
# of cores	48	48**
System (TFLOPS)	~2.7	~4.6
Memory BW (GB/s)	1024	1536

*Fujitsu presentation at Hot Chip 2018.

**Assume the same number of cores for comparison.

Summary

- ❖ RISC-V is quickly adopted in the world and making impact as the learning tool for computer architecture classes
 - It is most likely the future of microprocessor in many applications
 - Many extensions are developed in the task groups, not proprietary
 - Many companies developed RISC-V microprocessor
 - Extensible and configurable in comparison to other ISAs
- ❖ Andes is one major player in RISC-V microprocessor:
 - The first industry RISC-V vector processor
 - Setting the standard for high performance RISC-V vector processor with innovative design
 - Flexible design configurations to adapt to wide range of applications
- ❖ Andes Technology continues to lead and expand with RISC-V ecosystem

Thank You!