

RISC-V DSP (P) Extension Proposal

2018-5-9



Dr. Chuan-Hua Chang
Senior Director, RD/Architecture

Andes Technology



- ❖ Taiwan-based CPU IP company with over 2.5-billion Andes-Embedded SoCs shipped in diverse applications.



- ❖ Taking RISC-V to those markets with the solutions we developed in the past 13 years.
- ❖ A major contributor to RISC-V tools such as GCC, binutils, newlib, and recently LLVM and LLD.

RISC-V DSP (P) Extension TG



- ❖ P extension task group charter
 - Define and ratify Packed-SIMD DSP extension instructions operating on XLEN-bit integer registers for embedded RISC-V processors.
 - Define compiler intrinsic functions that can be directly used in high-level programming languages.
- ❖ Chair: Chuan-Hua Chang, Andes Technology
- ❖ Co-chair: Richard Herveille, RoaLogic

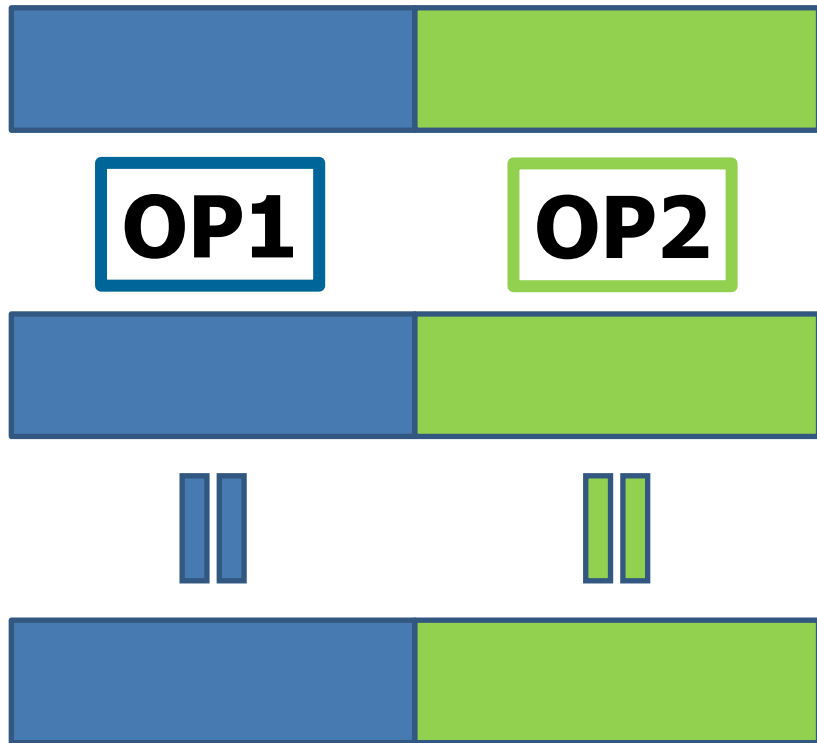
RISC-V DSP (P) Extension Proposal



❖ DSP instruction set proposal based on AndeStar™ V3 DSP ISA.


- User XLEN-bit GPRs.
- Support saturation and rounding.
- Support fixed-point and integer data types.
- **SIMD**-instructions with 8b, 16b, 32b element size.
- **Non-SIMD** DSP instructions operating on 16-bit, 32-bit and 64-bit data types.
- 64-bit signed/unsigned addition & subtraction
- 64-bit signed/unsigned multiplication & addition
 - ◆ E.g., $64 = 64 + 16 \times 16 + 16 \times 16$ or
 - ◆ E.g., $64 = 64 + 32 \times 32$

16-Bit SIMD Instructions



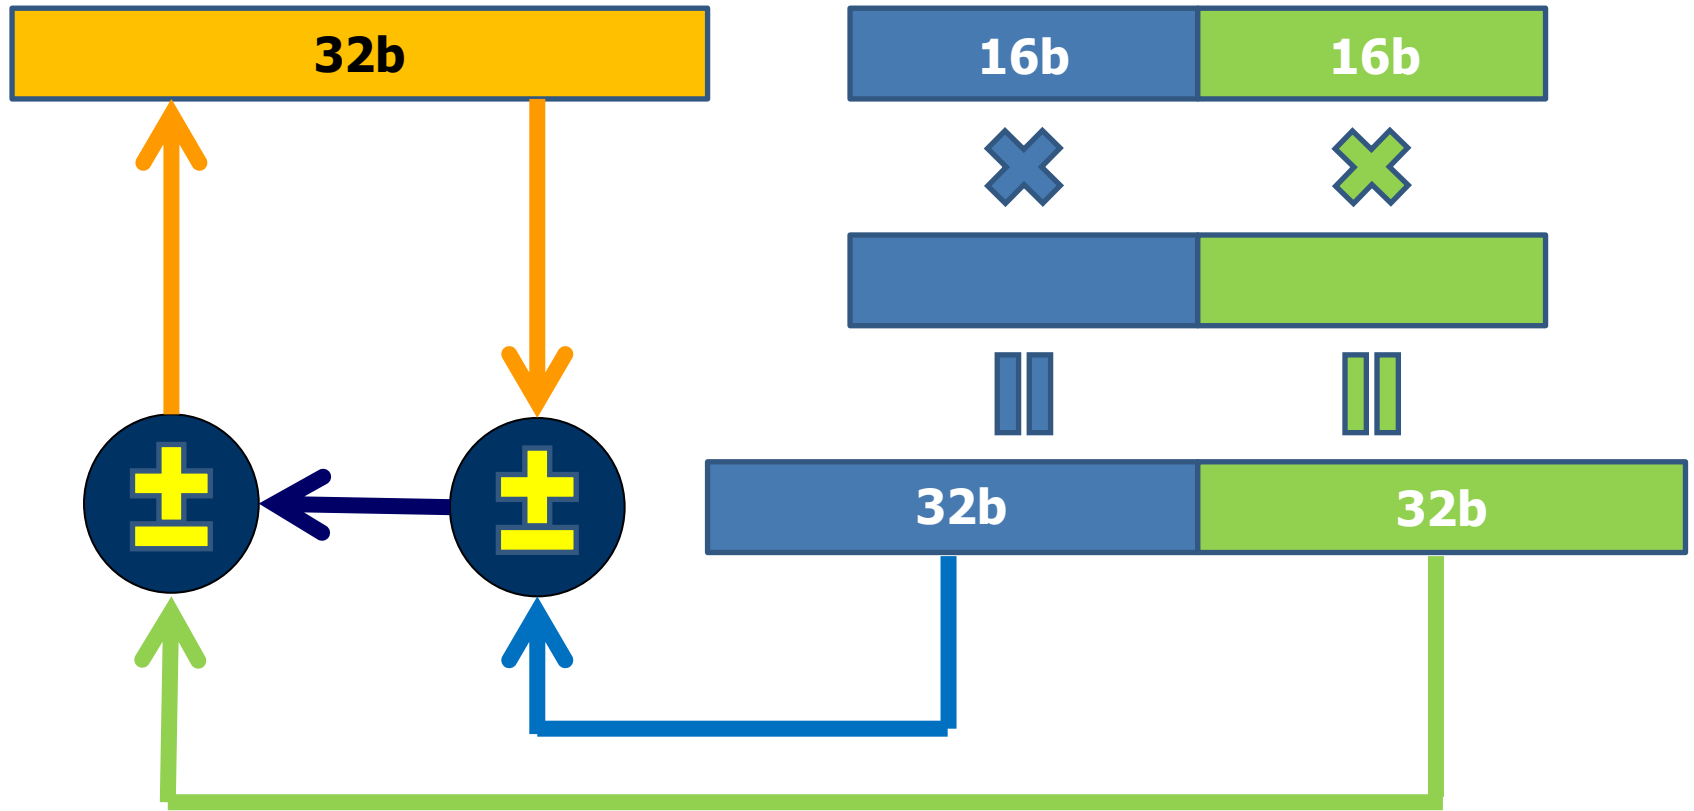
- + ×
<< Min Max
>> Clip ABS
Compare
Signed
Unsigned

8-Bit SIMD Instructions



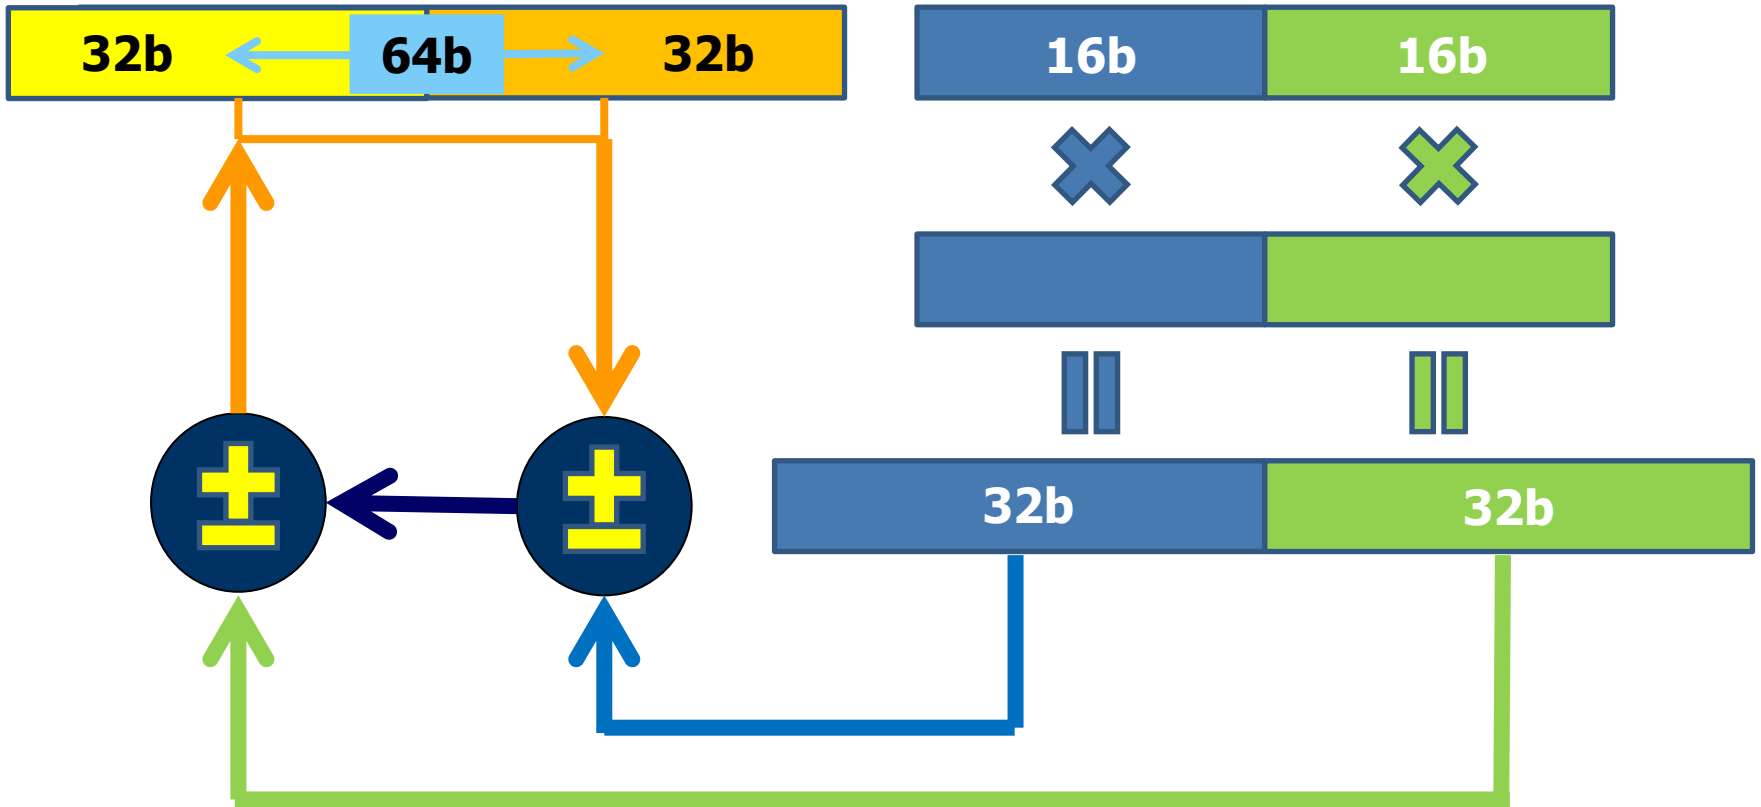
- +
Min Max
Unpack ABS
Compare
Signed
Unsigned

Dual 16x16 & 32-Bit Add/Sub



**6x v.s.
Baseline**

Dual 16x16 & 64-Bit Add/Sub



12x v.s. Baseline

GPR vs Separate Register



- ❖ GPR-based SIMD is a more efficient, low power DSP solution for embedded systems running applications in various domains such as audio/speech decoding and processing, IoT sensor data processing, wearable fitness devices, etc.
- ❖ It addresses the need for **high performance generic code processing**, as well as **digital signal processing**.

Ease of Use



- ❖ Provide data types and instructions that can be recognized and used by a compiler.
- ❖ Provide intrinsic functions for software developers to use the DSP instructions directly in C/C++ code.
- ❖ Provide optimized DSP libraries/middlewares covering common DSP functions and algorithms.
 - AndeStar™ V3 core performance: D10 is 79% faster than N10.

64-bit Data Type



- ❖ Use pairs of GPRs on RV32.
- ❖ Use a GPR on RV64.
- ❖ Needed for compiler to generate DSP instructions automatically.
- ❖ The 64-bit operand type is an interface specification. An implementation can still implement 2R1W register file with multi-cycle reads/writes to support the 64-bit type on RV32.

DSP ISA Performance



❖ Helix MP3 decoder

GCC Compiler	Decode (MCPS)
Compile with 32b base ISA	22.64
Compile with 32b base+DSP ISA	10.35
Cycle reduction %	54%
Cycle % of 64b paired-GPR insts	~70%

❖ G.729 codec

	Encode (MCPS)	Decode (MCPS)
Compiler alone	95.45	26.83
Intrinsic + Compiler	26.31	6.02
Cycle reduction %	72%	78%

* MCPS: Millions of Cycles Per Second

AMR-WB Performance



AMR-WB	N10	D10	D15
Encode	282.93	59.64	41.32
Decode	63.76	15.50	10.93

Performance improvements are highlighted with yellow starburst callouts and blue arrows:

- Encode: N10 to D10 is 4.7x; D10 to D15 is 6.9x.
- Decode: N10 to D10 is 4.1x; D10 to D15 is 5.8x.

- N10: Base ISA + 1-issue pipeline
- D10: Base + DSP ISA + 1-issue pipeline
- D15: Base + DSP ISA + 2-issue pipeline


The benchmarking numbers are in MCPS. (Million Cycle Per Second)

For all the test vectors, the average (AVG) numbers are concluded.

ANSI-C Code for the AMR-WB speech codec:

http://www.etsi.org/deliver/etsi_ts/126100_126199/126173/14.00.00_60/ts_126173v140000p0.zip

DSP Library Performance Speedup



P Task Group Progress



- ❖ In the process of discussion with Technical Committee Chair and Co-chair to create the task group.
- ❖ After the task group is created, TG will arrange bi-weekly meetings and publicly invite people to participate.

Thank You!



```
01001001001100100100input [31:0] hwdata;  
01110100100110000110output[31:0] hrdata;  
100100011100011100101output[15:0] hsplit;  
10100if ((retry_en = 1'b1 || split_en = 1'b1)  
010100110110001100begin resp_delay_rx =  
00001100100101010011016'b0 hresp_wait_cnt);  
0100100100110101010010010010010001  
begin resp_delay
```

Andes Core™

www.andestech.com